



AZURE FUNDAMENTALS

Exam AZ-900

AUGUST 2019

HENDRIK BULENS

CONTENTS

DISCLAIMER I

INTRODUCTION..... III

SKILLS MEASURED III

 UNDERSTAND CLOUD CONCEPTS (15-20%)..... III

 UNDERSTAND CORE AZURE SERVICES (30-35%)..... III

 UNDERSTAND SECURITY, PRIVACY, COMPLIANCE, AND TRUST (25-30%) IV

 UNDERSTAND AZURE PRICING AND SUPPORT (25-30%)..... V

1 CLOUD CONCEPTS - PRINCIPLES OF CLOUD COMPUTING..... 1

 1.1 WHAT IS CLOUD COMPUTING..... 1

 1.2 BENEFITS OF CLOUD COMPUTING 2

 1.2.1 *Cost-effective* 2

 1.2.2 *Scalable*..... 2

 1.2.3 *Elastic* 2

 1.2.4 *Current*..... 3

 1.2.5 *Reliable*..... 3

 1.2.6 *Global* 3

 1.2.7 *Secure* 3

 1.3 COMPLIANCE TERMS AND REQUIREMENTS 3

 1.4 ECONOMIES OF SCALE 3

 1.5 CAPITAL EXPENDITURE (CAPEX) VERSUS OPERATIONAL EXPENDITURE (OPEX) 4

 1.6 CLOUD DEPLOYMENT MODELS 5

 1.6.1 *Public cloud*..... 5

 1.6.2 *Private cloud* 5

 1.6.3 *Hybrid cloud* 6

 1.7 TYPES OF CLOUD SERVICES 6

 1.7.1 *IaaS*..... 6

 1.7.2 *PaaS* 6

 1.7.3 *SaaS* 6

 1.7.4 *Cost and ownership*..... 7

 1.7.5 *Management responsibilities* 7

2 CORE CLOUD SERVICES - INTRODUCTION TO AZURE 8

 2.1 WHAT IS AZURE? 8

 2.2 TOUR OF AZURE SERVICES 8

 2.2.1 *Compute*..... 9

 2.2.2 *Networking* 9

 2.2.3 *Storage* 10

 2.2.4 *Mobile*..... 10

 2.2.5 *Databases* 11

 2.2.6 *Web*..... 11

 2.2.7 *Internet of Things*..... 11

 2.2.8 *Big Data*..... 12

 2.2.9 *Artificial Intelligence* 12

 2.2.10 *DevOps* 13

 2.3 CREATE A VIRTUAL MACHINE 14

2.3.1	<i>Azure Cloud Shell</i>	14
2.3.2	<i>What is a virtual machine</i>	14
2.3.3	<i>Creating resources in Azure</i>	14
2.4	ADD A WEB SERVER.....	16
2.4.1	<i>Configure IIS</i>	16
2.5	SCALE UP.....	18
3	CORE CLOUD SERVICES - AZURE ARCHITECTURE AND SERVICE GUARANTEES	19
3.1	REGIONS	19
3.2	GEOGRAPHIES	19
3.3	AVAILABILITY ZONES	20
3.4	REGION PAIRS	21
3.5	SERVICE LEVEL AGREEMENTS FOR AZURE.....	22
3.5.1	<i>Performance Targets</i>	22
3.5.2	<i>Uptime and Connectivity Guarantees</i>	22
3.5.3	<i>Service Credits</i>	22
3.6	COMPOSING SLAs ACROSS SERVICES.....	23
3.7	IMPROVE YOUR APP RELIABILITY	23
4	CREATE AN AZURE ACCOUNT	24
4.1	AZURE ACCOUNTS AND SUBSCRIPTIONS.....	24
4.1.1	<i>Azure free subscription</i>	24
4.1.2	<i>Azure Pay-As-You-Go subscription</i>	24
4.1.3	<i>Pay-As-You-Go Dev/Test</i>	25
4.1.4	<i>Azure Enterprise Agreement</i>	25
4.1.5	<i>Azure for Students subscription</i>	25
4.2	USING MULTIPLE AZURE SUBSCRIPTIONS.....	25
4.2.1	<i>Access management</i>	25
4.2.2	<i>Billing</i>	26
4.3	AUTHENTICATE ACCESS WITH AZURE ACTIVE DIRECTORY.....	26
4.4	AZURE SUPPORT OPTIONS	27
4.4.1	<i>Paid Azure support plans</i>	27
4.4.2	<i>Support-plan availability and billing</i>	31
4.4.3	<i>Other support options</i>	31
5	CORE CLOUD SERVICES - MANAGE SERVICES WITH THE AZURE PORTAL	32
5.1	AZURE MANAGEMENT OPTIONS	32
5.1.1	<i>Azure portal</i>	32
5.1.2	<i>Azure PowerShell</i>	33
5.1.3	<i>Azure CLI</i>	33
5.1.4	<i>Azure Cloud Shell</i>	33
5.1.5	<i>Azure mobile app</i>	33
5.1.6	<i>Other options</i>	33
5.2	NAVIGATE THE PORTAL.....	34
5.2.1	<i>Azure portal layout</i>	34
5.2.2	<i>What is a blade?</i>	34
5.2.3	<i>What is the Azure Marketplace?</i>	34
5.3	CONFIGURING SETTINGS IN THE AZURE PORTAL	34

6	CORE CLOUD SERVICES - AZURE COMPUTE FUNCTIONS	35
6.1	ESSENTIAL AZURE COMPUTE CONCEPTS	35
6.1.1	What is Azure compute?.....	35
6.1.2	What are virtual machines?	35
6.1.3	What are containers?	35
6.1.4	What is Azure App Service?	35
6.1.5	What is Serverless Computing?.....	35
6.2	EXPLORE AZURE VIRTUAL MACHINES.....	36
6.2.1	Scaling VMs in Azure	36
6.3	EXPLORE CONTAINERS IN AZURE.....	38
6.3.1	Containers in Azure	38
6.4	EXPLORE AZURE APP SERVICE.....	39
6.4.1	Types of web apps	39
6.5	EXPLORE SERVERLESS COMPUTING IN AZURE	40
6.5.1	Azure Functions	40
6.5.2	Azure Logic Apps	40
6.5.3	Functions vs. Logic Apps	40
6.5.4	Azure Event Grid	41
6.6	EXPLORE AZURE IoT	43
6.6.1	IoT Central	45
6.6.2	IoT solution accelerators	45
6.6.3	IoT Hub	46
6.6.4	IoT Hub Device Provisioning Service	47
6.6.5	IoT Edge.....	48
6.6.6	Azure Digital Twins.....	48
6.6.7	Time Series Insights	48
6.6.8	Azure Maps	48
6.7	EXPLORE AZURE AI	48
6.7.1	Machine Learning service	48
6.7.2	Machine Learning Studio	49
6.7.3	How does Azure Machine Learning service differ from Studio?.....	50
6.7.4	Azure Databricks.....	51
7	CORE CLOUD SERVICES - AZURE DATA STORAGE OPTIONS	52
7.1	BENEFITS OF USING AZURE TO STORE DATA	52
7.1.1	Why store your data in the cloud?	52
7.1.2	Types of data.....	52
7.1.3	How Azure data storage can meet your business storage needs	53
7.2	COMPARISON BETWEEN AZURE DATA STORAGE AND ON-PREMISES STORAGE	61
8	CORE CLOUD SERVICES - AZURE NETWORKING OPTIONS	62
8.1	AZURE NETWORKING	62
8.1.1	Connectivity Services.....	63
8.1.2	Application protection services.....	64
8.1.3	Application delivery services	64
8.1.4	Network monitoring services	65
8.2	DEPLOY YOUR SITE TO AZURE	66
8.2.1	Using an N-tier architecture	66
8.3	SCALE WITH AZURE LOAD BALANCER	67

8.3.1	What are availability and high availability?	67
8.3.2	What is resiliency?	68
8.3.3	What is a load balancer?.....	68
8.3.4	What is Azure Load Balancer?.....	68
8.3.5	Azure Application Gateway.....	69
8.3.6	What about DNS?	70
8.4	REDUCE LATENCY WITH AZURE TRAFFIC MANAGER	70
8.4.1	What is network latency?	70
8.4.2	Use Traffic Manager to route users to the closest endpoint	71
8.4.3	Compare Load Balancer to Traffic Manager.....	71
9	SECURITY, RESPONSIBILITY AND TRUST IN AZURE	72
9.1	CLOUD SECURITY IS A SHARED RESPONSIBILITY	72
9.1.1	Share security responsibility with Azure.....	72
9.1.2	A layered approach to security	73
9.2	GET TIPS FROM AZURE SECURITY CENTER	75
9.2.1	Coverage	75
9.2.2	Available tiers.....	75
9.3	IDENTITY AND ACCESS.....	76
9.3.1	Authentication and authorization	76
9.3.2	What is Azure Active Directory?	76
9.3.3	Single sign-on.....	77
9.3.4	Multi-factor authentication.....	77
9.3.5	Providing identities to services	79
9.3.6	Role-based access control	80
9.4	ENCRYPTION	81
9.4.1	What is encryption	81
9.4.2	Encryption on Azure.....	82
9.5	PROTECT YOUR NETWORK	84
9.5.1	Internet protection.....	84
9.5.2	What is firewall?.....	84
9.5.3	Stopping Distributed Denial of Service (DDos) attacks.....	86
9.5.4	Controlling the traffic inside your virtual network.....	87
9.6	PROTECT YOUR SHARED DOCUMENTS.....	88
9.7	AZURE ADVANCED THREAD PROTECTION	88
10	APPLY AND MONITOR INFRASTRUCTURE STANDARDS WITH AZURE POLICY	89
10.1	DEFINE IT COMPLIANCE WITH AZURE POLICY	89
10.1.1	Creating a policy	90
10.1.2	View policy evaluation results.....	91
10.2	ORGANIZE POLICY WITH INITIATIVES	92
10.2.1	Defining initiatives	92
10.2.2	Assigning initiatives	92
10.3	ENTERPRISE GOVERNANCE MANAGEMENT	92
10.4	DEFINE STANDARD RESOURCES WITH AZURE BLUEPRINTS	93
10.4.1	How it's different from Resource Manager templates	94
10.4.2	Blueprint definition	95
10.5	EXPLORE YOUR SERVICE COMPLIANCE WITH COMPLIANCE MANAGER.....	95
10.5.1	Microsoft Privacy Statement.....	95

10.5.2	Microsoft Trust Center.....	95
10.5.3	Service Trust Portal.....	96
10.5.4	Compliance Manager	96
10.6	MONITOR YOUR SERVICE HEALTH	97
10.6.1	Azure Monitor.....	97
10.6.2	Azure Service Health	99
11	CONTROL AND ORGANIZE AZURE RESOURCES WITH AZURE RESOURCE MANAGER	100
11.1	PRINCIPLES OF RESOURCE GROUPS	100
11.1.1	What are resource groups?.....	100
11.1.2	Create a Resource Group.....	101
11.1.3	Use resource groups for organization	101
11.1.4	Organizing principles	101
11.2	USE TAGGING TO ORGANIZE RESOURCES	103
11.2.1	What are tags?.....	103
11.2.2	Use tags for organization	103
11.3	USE POLICIES TO ENFORCE STANDARDS	104
11.3.1	What is Azure Policy?.....	104
11.3.2	Policies to enforce standards	104
11.4	SECURE RESOURCES WITH ROLE-BASED ACCESS CONTROL.....	105
11.4.1	How RBAC works	105
11.4.2	Best Practices for RBAC.....	107
11.5	USE RESOURCE LOCKS TO PROTECT RESOURCES	108
12	PREDICT COST AND OPTIMIZE SPENDING FOR AZURE.....	109
12.1	PURCHASING AZURE PRODUCTS AND SERVICES	109
12.2	FACTORS AFFECTING COSTS.....	110
12.2.1	Resource type	110
12.2.2	Services.....	110
12.2.3	Location.....	110
12.2.4	Azure billing zones.....	110
12.3	ESTIMATE COSTS WITH THE AZURE PRICING CALCULATOR.....	111
12.4	PREDICT AND OPTIMIZE WITH COST MANAGEMENT AND AZURE ADVISOR	112
12.4.1	What is Azure Advisor?.....	112
12.4.2	Azure Cost Management.....	113
12.5	ESTIMATE THE TOTAL COST OF OWNERSHIP WITH THE AZURE TCO CALCULATOR.....	114
12.5.1	Step 1: Open the TCO calculator	114
12.5.2	Step 2: Define your workloads.....	114
12.5.3	Step 3: Adjust assumptions	114
12.5.4	Step 4: View the report.....	114
12.6	SAVE ON INFRASTRUCTURE COSTS.....	115
12.6.1	Use Azure credits	115
12.6.2	Use spending limits	115
12.6.3	Use reserved instances.....	116
12.6.4	Choose low-cost locations and regions	116
12.6.5	Research available cost-saving offers.....	116
12.6.6	Right-size underutilized virtual machines	116
12.6.7	Deallocate virtual machines in off hours	116
12.6.8	Delete unused virtual machines	116

12.6.9	Migrate to PaaS or SaaS services	117
12.7	SAVE ON LICENSING COSTS	117
12.7.1	Linux vs. Windows.....	117
12.7.2	Azure Hybrid Benefit for Windows Server	117
12.7.3	Azure Hybrid Benefit for SQL Server.....	118
12.7.4	Use Dev/Test subscription offers	118
12.7.5	Bring your own SQL Server license.....	118
12.7.6	Use SQL Server Developer Edition.....	118
12.7.7	Use constrained instance sizes for database workloads.....	118
13	SERVICE LIFECYCLE	119
13.1	PRIVATE PREVIEWS.....	119
13.2	PUBLIC PREVIEW	119
13.3	GENERAL AVAILABILITY (GA)	119
13.4	HOW TO ACCESS PREVIEW FEATURES	119
14	READ MORE	120
14.1	STUDY GUIDES.....	120
14.2	USEFUL LINKS	120
14.3	REFERENCES	120
14.3.1	Cloud concepts	120
14.3.2	Understand Core Azure Services	120
14.3.3	Understand Security, Privacy, Compliance and Trust	121
14.3.4	Understand Azure Pricing and support	121
15	INDEX.....	122

DISCLAIMER

All content provided in this document is for informational purposes only. The owner of this document makes no representations as to the accuracy or completeness of any information in this document or found by following any links in this document. The owner will not be liable for any errors or omissions in this information nor for the availability of this information. The owner will not be liable for any losses, injuries, or damages from the display or use of this information.

INTRODUCTION

This exam is designed for candidates looking to demonstrate foundational level knowledge of cloud services and how those services are provided with Microsoft Azure. The exam is intended for candidates with non-technical backgrounds, such as those involved in selling or purchasing cloud based solutions and services or who have some involvement with cloud based solutions and services, as well as those with a technical background who have a need to validate their foundational level knowledge around cloud services. Technical IT experience is not required however some general IT knowledge or experience would be beneficial.

This exam can be taken as an optional first step in learning about cloud services and how those concepts are exemplified by Microsoft Azure. It can be taken as a precursor to Microsoft Azure or Microsoft cloud services exams. While it would be a beneficial first step, validating foundational level knowledge, taking this exam is not a pre-requisite before taking any other Azure-based certifications.

This document is a summary of three main sources:

- The Azure Fundamentals Learning Path, <https://docs.microsoft.com/en-us/learn/paths/azure-fundamentals/>
- Azure Docs, <https://docs.microsoft.com/en-us/azure/>
- Cheshire, J. (2019). *Exam Ref AZ-900, Microsoft Azure Fundamentals*. Pearson Education

I endorse these sources and I suggest you check them out when preparing for the exam.

SKILLS MEASURED

UNDERSTAND CLOUD CONCEPTS (15-20%)

- Describe the benefits and considerations of using cloud services
 - understand terms such as High Availability, Scalability, Elasticity, Agility, Fault Tolerance, and Disaster Recovery
 - understand the principles of economies of scale
 - understand the differences between Capital Expenditure (CapEx) and Operational Expenditure (OpEx)
 - understand the consumption-based model
- Describe the differences between Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) and Software-as-a-Service (SaaS)
 - describe Infrastructure-as-a-Service (IaaS)
 - describe Platform-as-a-Service (PaaS)
 - describe Software-as-a-Service (SaaS)
 - compare and contrast the three different service types
- Describe the differences between Public, Private and Hybrid cloud models
 - describe Public cloud
 - describe Private cloud
 - describe Hybrid cloud
 - compare and contrast the three different cloud models

UNDERSTAND CORE AZURE SERVICES (30-35%)

- Understand the core Azure architectural components
 - describe Regions
 - describe Availability Zones
 - describe Resource Groups
 - describe Azure Resource manager
 - describe the benefits and usage of core Azure architectural components
- Describe some of the core products available in Azure
 - describe products available for Compute such as Virtual Machines, Virtual Machine Scale Sets, App Service and Functions
 - describe products available for Networking such as Virtual Network, Load Balancer, VPN Gateway, Application Gateway and Content Delivery Network
 - describe products available for Storage such as Blob Storage, Disk Storage, File Storage, and Archive Storage
 - describe products available for Databases such as CosmosDB, Azure SQL Database, Azure Database Migration service, and Azure SQL Data Warehouse
 - describe the Azure Marketplace and its usage scenarios

- Describe some of the solutions available on Azure
 - describe Internet of Things (IoT) and products that are available for IoT on Azure such as IoT Fundamentals, IoT Hub and IoT Central
 - describe Big Data and Analytics and products that are available for Big Data and Analytics such as SQL Data Warehouse, HDInsight and Data Lake Analytics
 - describe Artificial Intelligence (AI) and products that are available for AI such as Azure Machine Learning Service and Studio
 - describe Serverless computing and Azure products that are available for serverless computing such as Azure Functions, Logic Apps and App grid
 - describe the benefits and outcomes of using Azure solutions
- Understand Azure management tools
 - understand Azure tools such as Azure CLI, PowerShell, and the Azure Portal
 - understand Azure Advisor

UNDERSTAND SECURITY, PRIVACY, COMPLIANCE, AND TRUST (25-30%)

- Understand securing network connectivity in Azure
 - describe Azure Firewall
 - describe Azure DDoS Protection
 - describe Network Security Group (NSG)
 - choose an appropriate Azure security solution
- Describe core Azure Identity services
 - understand the difference between authentication and authorization
 - describe Azure Active Directory
 - describe Azure Multi-Factor Authentication
- Describe security tools and features of Azure
 - describe Azure Security
 - understand Azure Security center usage scenarios
 - describe Key Vault
 - describe Azure Information Protection (AIP)
 - describe Azure Advanced Threat Protection (ATP)
- Describe Azure governance methodologies
 - describe Azure Policies
 - describe Initiatives
 - describe Role-Based Access Control (RBAC)
 - describe Locks
 - describe Azure Advisor security assistance
- Understand monitoring and reporting options in Azure
 - describe Azure Monitor
 - describe Azure Service Health
 - understand the use cases and benefits of Azure Monitor and Azure Service Health

- Understand privacy, compliance and data protection standards in Azure
 - understand industry compliance terms such as GDPR, ISO and NIST
 - understand the Microsoft Privacy Statement
 - describe the Trust center
 - describe the Service Trust Portal
 - describe Compliance Manager
 - determine if Azure is compliant for a business need
 - understand Azure Government services
 - understand Azure Germany services

UNDERSTAND AZURE PRICING AND SUPPORT (25-30%)

- Understand Azure subscriptions
 - describe an Azure subscription
 - understand the uses and options with Azure subscriptions
- Understand planning and management of costs
 - understand options for purchasing Azure products and services
 - understand options around Azure Free account
 - understand the factors affecting costs such as resource types, services, locations, ingress and egress traffic
 - understand Zones for billing purposes
 - understand the Pricing calculator
 - understand the Total Cost of Ownership (TCO) calculator
 - understand best practices for minimizing Azure costs such as performing cost analysis, creating spending limits and quotas, and using tags to identify cost owners; use Azure reservations; use Azure Advisor recommendations
 - describe Azure Cost Management
- Understand the support options available with Azure
 - understand support plans that are available such as Dev, Standard, Professional Direct and Premier
 - understand how to open a support ticket
 - understand available support channels outside of support plan channels
 - describe the Knowledge Center
- Describe Azure Service Level Agreements (SLAs)
 - describe a Service Level Agreement (SLA)
 - determine SLA for a particular Azure product or service
- Understand service lifecycle in Azure
 - understand Public and Private Preview features
 - understand how to access Preview features
 - understand the term General Availability (GA)
 - monitor feature updates

1 CLOUD CONCEPTS - PRINCIPLES OF CLOUD COMPUTING

In this module, you will:

- Explore common cloud computing services
- Explore the benefits of cloud computing
- Decide which cloud deployment model is best for you

1.1 WHAT IS CLOUD COMPUTING

Cloud computing is renting resources, like storage space or CPU cycles, on another company's computers. You only pay for what you use. The computing services offered tend to vary by cloud provider.

A **virtual machine** (VM) is an emulation of a computer - just like your desktop or laptop you're using now. Each VM includes an operating system and hardware that appears to the user like a physical computer running Windows or Linux. You can then install whatever software you need to do the tasks you want to run in the cloud.

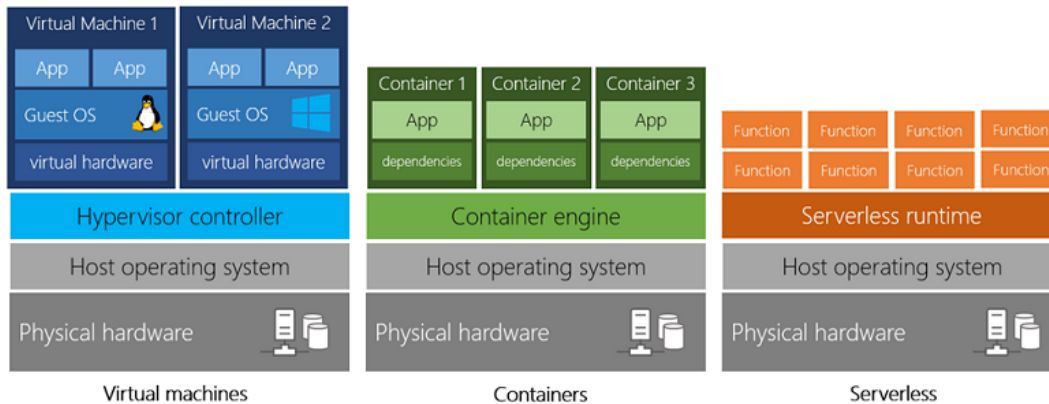
The difference is that you don't have to buy any of the hardware or install the OS. The cloud provider runs your virtual machine on a physical server in one of their datacenters - often sharing that server with other VMs (isolated and secure). With the cloud, you can have a VM ready to go in minutes at less cost than a physical computer.

VMs aren't the only computing choice - there are two other popular options: **containers** and **serverless computing**:

- **Containers** are similar to VMs except they don't require a guest operating system. Instead, the application and all its dependencies is packaged into a "container" and then a standard runtime environment is used to execute the app.
- **Serverless computing** lets you run application code without creating, configuring, or maintaining a server. The core idea is that your application is broken into separate functions that run when triggered by some action. This is ideal for automated tasks - for example, you can build a serverless process that automatically sends an email confirmation after a customer makes an online purchase.

The serverless model differs from VMs and containers in that you only pay for the processing time used by each function as it executes. VMs and containers are charged while they're running - even if the applications on them are idle.

Here's a diagram comparing the three compute approaches we've covered.



VM vs container vs serverless computing

1.2 BENEFITS OF CLOUD COMPUTING

1.2.1 COST-EFFECTIVE

Cloud computing provides a **pay-as-you-go** or **consumption-based** pricing model. Rather than paying upfront for a pre-defined amount of computing resources or hardware, you can rent hardware and pay for the resources that you actually use.

This consumption-based model brings with it many benefits, including:

- No upfront costs
- No need to purchase and manage costly infrastructure that you may not use to its fullest
- The ability to pay for additional resources only when they are needed
- The ability to stop paying for resources that are no longer needed

Renting infrastructure also allows for better cost prediction.

1.2.2 SCALABLE

Both vertical and horizontal scaling are supported, allowing to increase or decrease usage.

Vertical scaling, also known as “scaling up”, is the process of adding resources to increase the power of an existing server. Some examples of vertical scaling are: adding more CPUs, or adding more memory.

Horizontal scaling, also known as “scaling out”, is the process of adding more servers that function together as one unit. For example, you have more than one server processing incoming requests.

Scaling can be done manually or automatically based on specific triggers such as CPU utilization or the number of requests and resources can be allocated or de-allocated in minutes.

1.2.3 ELASTIC

As your workload changes due to a spike or drop in demand, a cloud computing system can compensate by automatically adding or removing resources.

1.2.4 CURRENT

Cloud eliminates the burdens of maintaining software patches, hardware setup, upgrades, and other IT management tasks.

1.2.5 RELIABLE

Cloud computing providers offer data backup, disaster recovery, and data replication services to make sure your data is always safe. In addition, redundancy is often built into cloud services architecture so if one component fails, a backup component takes its place. This is referred to as **fault tolerance** and it ensures that your customers aren't impacted when a disaster occurs.

1.2.6 GLOBAL

Cloud providers have fully redundant datacenters located in various regions all over the globe. This gives you a local presence close to your customers to give them the best response time possible no matter where in the world they are.

You can replicate your services into multiple regions for redundancy and locality, or select a specific region to ensure you meet data-residency and compliance laws for your customers.

1.2.7 SECURE

Cloud providers offer a broad set of policies, technologies, controls, and expert technical skills that can provide better security than most organizations can otherwise achieve. The result is strengthened security, which helps to protect data, apps, and infrastructure from potential threats.

1.3 COMPLIANCE TERMS AND REQUIREMENTS

When selecting a cloud provider to host your solutions, you should understand how that provider can help you comply with regulations and standards.

1.4 ECONOMIES OF SCALE

Economies of scale is the ability to do things more efficiently or at a lower-cost per unit when operating at a larger scale.

1.5 CAPITAL EXPENDITURE (CAPEX) VERSUS OPERATIONAL EXPENDITURE (OPEX)

In the past, companies needed to acquire physical premises and infrastructure to start their business. There was a substantial up-front cost in hardware and infrastructure to start or grow a business. Cloud computing provides services to customers without significant upfront costs or equipment setup time.

These two approaches to investment are referred to as capital expenditure and operational expenditure

- **Capital Expenditure (CapEx):** CapEx is the spending of money on physical infrastructure up front, and then deducting that expense from your tax bill over time. CapEx is an upfront cost, which has a value that reduces over time.
- **Operational Expenditure (OpEx):** OpEx is spending money on services or products now and being billed for them now. You can deduct this expense from your tax bill in the same year. There's no upfront cost. You pay for a service or product as you use it.

Computing costs comparison:

CapEx	OpEx
Server	Leasing cloud-based server
Storage	Leasing software & features
Network	Usage
Backup and archive	
Disaster recovery	
Data center infrastructure	
Technical personnel	

Benefits comparison:

CapEx	OpEx
Expenses are planned at the start of a project	Costs are managed dynamically.
Costs are fixed; you know exactly how much is being spent.	Costs fluctuate along with the demand.
Appealing when you need to predict the expenses before a project starts.	Appealing if the demand fluctuates or is unknown.

1.6 CLOUD DEPLOYMENT MODELS

A cloud deployment model defines where your data is stored and how your customers interact with it – how do they get to it, and where do the applications run? It also depends on how much of your own infrastructure you want or need to manage.

1.6.1 PUBLIC CLOUD

There is no local hardware to manage or keep up-to-date in a public cloud – everything runs on your cloud provider’s hardware.

Advantages	Disadvantages
High scalability/agility	Security requirements that cannot be met by using public cloud
Pay-as-you-go pricing	Government policies, industry standards, or legal requirements which public clouds cannot meet
Not responsible for maintenance or updates of the hardware	You don’t own the hardware or services and cannot manage them as you may want to
Minimal technical knowledge to set up and use	Unique business requirements, such as having to maintain a legacy application might be hard to meet

1.6.2 PRIVATE CLOUD

In a private cloud, you create a cloud environment in your own datacenter and provide self-service access to compute resources to users in your organization. This offers a simulation of a public cloud to your users, but you remain completely responsible for the purchase and maintenance of the hardware and software services you provide.

Advantages	Disadvantages
You can ensure the configuration can support any scenario or legacy application	You have some initial CapEx costs and must purchase the hardware for startup and maintenance
You can control (and responsibility) over security	Owning the equipment limits the agility - to scale you must buy, install, and setup new hardware
Private clouds can meet strict security, compliance, or legal requirements	Private clouds require IT skills and expertise that’s hard to come by
Economies at scale and integration with Azure Security Center	

1.6.3 HYBRID CLOUD

A hybrid cloud combines public and private clouds, allowing you to run your applications in the most appropriate location.

Advantages	Disadvantages
Keep any systems running and accessible that use out-of-date hardware or an out-of-date operating system	Can be more expensive than selecting one deployment model since it involves some CapEx cost up front
You have flexibility with what you run locally versus in the cloud	It can be more complicated to set up and manage
You can take advantage of economies of scale from public cloud providers for services and resources where it's cheaper, and then supplement with your own equipment when it's not	
You can use your own equipment to meet security, compliance, or legacy scenarios where you need to completely control the environment	

1.7 TYPES OF CLOUD SERVICES

1.7.1 IAAS

Infrastructure as a Service is the most flexible category of cloud services. It aims to give you complete control over the hardware that runs your application.

Instead of buying hardware, with IaaS, you rent it. It's an instant computing infrastructure, provisioned and managed over the internet.

1.7.2 PAAS

Platform as a Service provides an environment for building, testing, and deploying software applications. The goal of PaaS is to help you create an application quickly without managing the underlying infrastructure.

1.7.3 SAAS

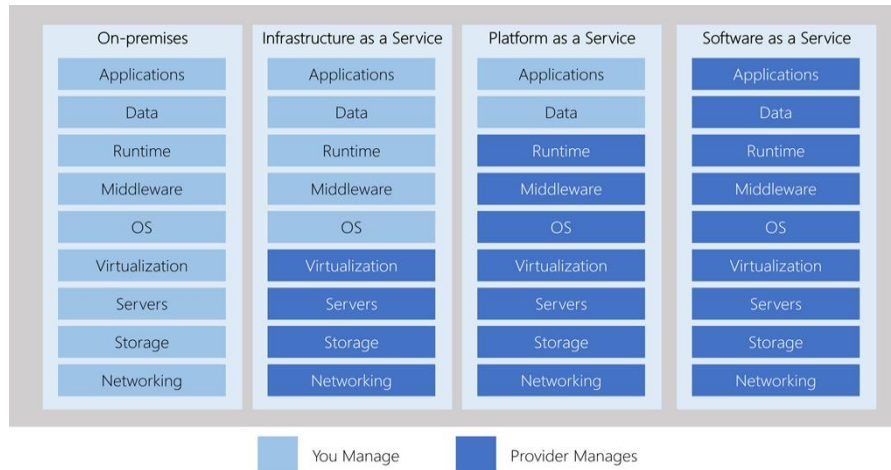
Software as a Service is software that is centrally hosted and managed for the end customer. It is usually based on an architecture where one version of the application is used for all customers, and licensed through a monthly or annual subscription. Office 365, Skype, and Dynamics CRM Online are perfect examples of SaaS software.

1.7.4 COST AND OWNERSHIP

	IaaS	PaaS	SaaS
Upfront costs	No upfront costs. Users pay only for what they consume.	No upfront costs. Users pay only for what they consume.	No upfront costs; they pay a subscription, typically on a monthly or annual basis.
User ownership	The user is responsible for the purchase, installation, configuration, and management of their own software, operating systems, middleware, and applications.	The user is responsible for the development of their own applications. However, they are not responsible for managing the server or infrastructure. This allows the user to focus on the application or workload they want to run.	Users just use the application software; they are not responsible for any maintenance or management of that software.
Cloud provider ownership	The cloud provider is responsible for ensuring that the underlying cloud infrastructure (such as virtual machines, storage, and networking) is available for the user.	The cloud provider is responsible for operating system management, network, and service configuration. Cloud providers are typically responsible for everything apart from the application that a user wants to run. They provide a complete managed platform on which to run the application.	The cloud provider is responsible for the provision, management, and maintenance of the application software.

1.7.5 MANAGEMENT RESPONSIBILITIES

One thing to understand is that these categories are layers on top of each other. For example, PaaS adds a layer on top of IaaS by providing a level of abstraction. The abstraction has the benefit of hiding the details that you may not care about, so that you can get to coding quicker. However, one aspect of the abstraction is that you have less control over the underlying hardware. The following illustration shows a list of resources that you manage and that your service provider manages in each cloud service category.



- **IaaS** requires the most user management of all the cloud services. The user is responsible for managing the operating systems, data, and applications.
- **PaaS** requires less user management. The cloud provider manages the operating systems, and the user is responsible for the applications and data they run and store.
- **SaaS** requires the least amount of management. The cloud provider is responsible for managing everything, and the end user just uses the software.

2 CORE CLOUD SERVICES - INTRODUCTION TO AZURE

In this module, you will:

- Learn what Microsoft Azure is and how it relates to cloud computing
- Use Azure Cloud Shell to launch a Windows or Linux virtual machine
- Configure your virtual machine to run a basic web server
- Scale up your server to give you more compute power

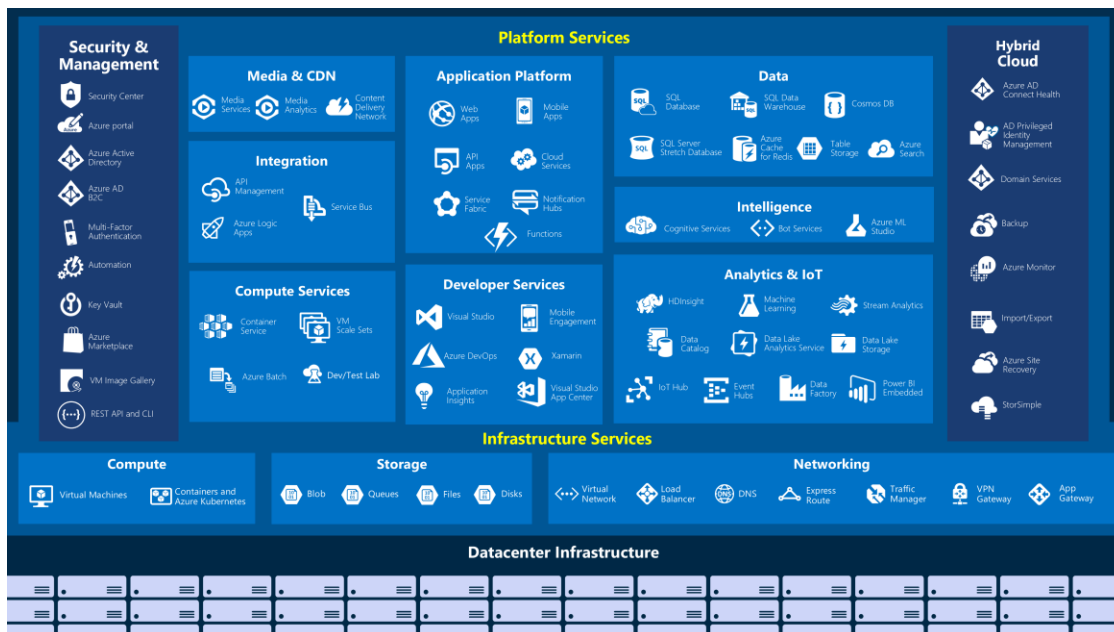
2.1 WHAT IS AZURE?

Azure is Microsoft's cloud computing platform. Azure gives you the freedom to build, manage, and deploy applications on a massive global network using your favorite tools and frameworks. Azure provides over 100 services that enable you to do everything from running your existing applications on virtual machines to exploring new software paradigms such as intelligent bots and mixed reality.

While migrating your existing apps to virtual machines is a good start, the cloud is more than just “a different place to run your virtual machines”. For example, Azure provides AI and machine-learning services that can naturally communicate with your users through vision, hearing, and speech. It also provides storage solutions that dynamically grow to accommodate massive amounts of data. Azure services enable solutions that are simply not feasible without the power of the cloud.

2.2 TOUR OF AZURE SERVICES

Here's a big-picture view of the available services and features in Azure.



Azure Services

Commonly-used categories include:

- **Compute**
- **Networking**
- **Storage**
- **Mobile**
- **Databases**
- **Web**
- **Internet of Things**
- **Big Data**
- **Artificial Intelligence**
- **DevOps**

2.2.1 COMPUTE

Compute services are often one of the primary reasons why companies move to the Azure platform. Azure provides a range of options for hosting applications and services. Here are some examples of compute services in Azure:

Service name	Service function
Azure Virtual Machines	Windows or Linux virtual machines (VMs) hosted in Azure
Azure Virtual Machine Scale Sets	Scaling for Windows or Linux VMs hosted in Azure
Azure Kubernetes Service	Enables management of a cluster of VMs that run containerized services
Azure Service Fabric	Distributed systems platform. Runs in Azure or on-premises
Azure Batch	Managed service for parallel and high-performance computing applications
Azure Container Instances	Run containerized apps on Azure without provisioning servers or VMs
Azure Functions	An event-driven, serverless compute service

2.2.2 NETWORKING

Linking compute resources and providing access to applications is the key function of Azure networking. Networking functionality in Azure includes a range of options to connect the outside world to services and features in the global Microsoft Azure datacenters. Azure networking facilities have the following features:

Service name	Service function
Azure Virtual Network	Connects VMs to incoming Virtual Private Network (VPN) connections
Azure Load Balancer	Balances inbound and outbound connections to applications or service endpoints
Azure Application Gateway	Optimizes app server farm delivery while increasing application security
Azure VPN Gateway	Accesses Azure Virtual Networks through high-performance VPN gateways
Azure DNS	Provides ultra-fast DNS responses and ultra-high domain availability
Azure Content Delivery Network	Delivers high-bandwidth content to customers globally
Azure DDoS Protection	Protects Azure-hosted applications from distributed denial of service (DDoS) attacks
Azure Traffic Manager	Distributes network traffic across Azure regions worldwide
Azure ExpressRoute	Connects to Azure over high-bandwidth dedicated secure connections
Azure Network Watcher	Monitors and diagnoses network issues using scenario-based analysis
Azure Firewall	Implements high-security, high-availability firewall with unlimited scalability
Azure Virtual WAN	Creates a unified wide area network (WAN), connecting local and remote sites

2.2.3 STORAGE

Azure provides four main types of storage services. These services are:

Service name	Service function
Azure Blob storage	Storage service for very large objects, such as video files or bitmaps
Azure File storage	File shares that you can access and manage like a file server
Azure Queue storage	A data store for queuing and reliably delivering messages between applications
Azure Table storage	A NoSQL store that hosts unstructured data independent of any schema

These services all share several common characteristics:

- Durable and highly available with redundancy and replication.
- Secure through automatic encryption and role-based access control.
- Scalable with virtually unlimited storage.
- Managed, handling maintenance and any critical problems for you.
- Accessible from anywhere in the world over HTTP or HTTPS.

2.2.4 MOBILE

Azure enables developers to create mobile backend services for iOS, Android, and Windows apps quickly and easily. Features that used to take time and increase project risks, such as adding corporate sign-in and then connecting to on-premises resources such as SAP, Oracle, SQL Server, and SharePoint, are now simple to include.

Other features of this service include:

- Offline data synchronization.
- Connectivity to on-premises data.
- Broadcasting push notifications.
- Autoscaling to match business needs.

2.2.5 DATABASES

Azure provides multiple database services to store a wide variety of data types and volumes. And with global connectivity, this data is available to users instantly.

Service name	Service function
Azure Cosmos DB	Globally distributed database that supports NoSQL options
Azure SQL Database	Fully managed relational database with auto-scale, integral intelligence, and robust security
Azure Database for MySQL	Fully managed and scalable MySQL relational database with high availability and security
Azure Database for PostgreSQL	Fully managed and scalable PostgreSQL relational database with high availability and security
SQL Server on VMs	Host enterprise SQL Server apps in the cloud
Azure SQL Data Warehouse	Fully managed data warehouse with integral security at every level of scale at no extra cost
Azure Database Migration Service	Migrates your databases to the cloud with no application code changes
Azure Cache for Redis	Caches frequently used and static data to reduce data and application latency
Azure Database for MariaDB	Fully managed and scalable MariaDB relational database with high availability and security

2.2.6 WEB

Having a great web experience is critical in today's business world. Azure includes first-class support to build and host web apps and HTTP-based web services. The Azure services focused on web hosting include:

Service name	Service function
Azure App Service	Quickly create powerful cloud web-based apps
Azure Notification Hubs	Send push notifications to any platform from any back end.
Azure API Management	Publish APIs to developers, partners, and employees securely and at scale.
Azure Search	Fully managed search as a service.
Web Apps feature of Azure App Service	Create and deploy mission-critical web apps at scale.
Azure SignalR Service	Add real-time web functionalities easily.

2.2.7 INTERNET OF THINGS

People are able to access more information than ever before. It began with personal digital assistants (PDAs), then morphed into smartphones. Now there are smart watches, smart thermostats, even smart refrigerators. Personal computers used to be the norm. Now the internet allows any item that's online-capable to access valuable information. This ability for devices to garner and then relay information for data analysis is referred to as the Internet of Things (IoT).

The Azure Internet of Things (IoT) is a collection of Microsoft-managed cloud services that connect, monitor, and control billions of IoT assets. In simpler terms, an IoT solution is made up of one or more IoT devices and one or more back-end services running in the cloud that communicate with each other.

Service name	Service function
IoT Central	Fully-managed global IoT software as a service (SaaS) solution that makes it easy to connect, monitor, and manage your IoT assets at scale
Azure IoT Hub	Messaging hub that provides secure communications and monitoring between millions of IoT devices
IoT Edge	Push your data analysis onto your IoT devices instead of in the cloud allowing them to react more quickly to state changes.

2.2.8 BIG DATA

Data comes in all formats and sizes. When we talk about Big Data, we're referring to large volumes of data. Data from weather systems, communications systems, genomic research, imaging platforms, and many other scenarios generate hundreds of gigabytes of data. This amount of data makes it hard to analyze and make decisions around. It's often so large that traditional forms of processing and analysis are no longer appropriate.

Open source cluster technologies have been developed to deal with these large data sets. Microsoft Azure supports a broad range of technologies and services to provide big data and analytic solutions.

Service name	Service function
Azure SQL Data Warehouse	Run analytics at a massive scale using a cloud-based Enterprise Data Warehouse (EDW) that leverages massive parallel processing (MPP) to run complex queries quickly across petabytes of data
Azure HDInsight	Process massive amounts of data with managed clusters of Hadoop clusters in the cloud
Azure Databricks (preview)	Collaborative Apache Spark-based analytics service that can be integrated with other Big Data services in Azure.

2.2.9 ARTIFICIAL INTELLIGENCE

Artificial Intelligence, in the context of cloud computing, is based around a broad range of services, the core of which is Machine Learning. Machine Learning is a data science technique that allows computers to use existing data to forecast future behaviors, outcomes, and trends. Using machine learning, computers learn without being explicitly programmed.

Forecasts or predictions from machine learning can make apps and devices smarter. For example, when you shop online, machine learning helps recommend other products you might like based on what you've purchased. Or when your credit card is swiped, machine learning compares the transaction to a database of transactions and helps detect fraud. And when your robot vacuum cleaner vacuums a room, machine learning helps it decide whether the job is done.

Some of the most common Artificial Intelligence and Machine Learning service types in Azure are:

Service name	Service function
Azure Machine Learning Service	Cloud-based environment you can use to develop, train, test, deploy, manage, and track machine learning models. It can auto-generate a model and auto-tune it for you. It will let you start training on your local machine, and then scale out to the cloud
Azure Machine Learning Studio	Collaborative, drag-and-drop visual workspace where you can build, test, and deploy machine learning solutions using pre-built machine learning algorithms and data-handling modules

A closely related set of products are the cognitive services. These are pre-built APIs you can leverage in your applications to solve complex problems.

Service name	Service function
Vision	Image-processing algorithms to smartly identify, caption, index, and moderate your pictures and videos.
Speech	Convert spoken audio into text, use voice for verification, or add speaker recognition to your app.
Knowledge mapping	Map complex information and data in order to solve tasks such as intelligent recommendations and semantic search.
Bing Search	Add Bing Search APIs to your apps and harness the ability to comb billions of webpages, images, videos, and news with a single API call
Natural Language processing	Allow your apps to process natural language with pre-built scripts, evaluate sentiment and learn how to recognize what users want.

2.2.10 DEVOPS

DevOps (Development and Operations) brings together people, processes, and technology, automating software delivery to provide continuous value to your users. Azure DevOps Services allows you to create build and release pipelines that provide continuous integration, delivery, and deployment for your applications. You can integrate repositories and application tests, perform application monitoring, and work with build artifacts. You can also work with and backlog items for tracking, automate infrastructure deployment and integrate a range of third-party tools and services such as Jenkins and Chef. All of these functions and many more are closely integrated with Azure to allow for consistent, repeatable deployments for your applications to provide streamlined build and release processes.

Some of the main DevOps services available with Azure are Azure DevOps Services and Azure DevTest Labs.

Service name	Service function
Azure DevOps	Provides development collaboration tools including high-performance pipelines, free private Git repositories, configurable Kanban boards, and extensive automated and cloud-based load testing
Azure DevTest Labs	Quickly create on-demand Windows and Linux environments you can use to test or demo your applications directly from your deployment pipelines

2.3 CREATE A VIRTUAL MACHINE

2.3.1 AZURE CLOUD SHELL

Azure Cloud Shell is a browser-based command-line experience for managing and developing Azure resources. Think of Cloud Shell as an interactive console that you run in the cloud.

Cloud Shell provides two experiences to choose from: *Bash* and *PowerShell*. Both include access to the Azure CLI, the command-line interface for Azure.

You can use any Azure management interface, including the Azure portal, Azure CLI, and Azure PowerShell, to manage any kind of VM.

2.3.2 WHAT IS A VIRTUAL MACHINE

A **virtual machine**, or VM, is a software emulation of a physical computer. A snapshot of a running VM is called an *image*. Azure provides images for Windows and several flavors of Linux. You can also create your own preconfigured images to make deployments go faster.

A virtual machine is defined by a number of factors, including its size and location. Before you bring up your VM, let's briefly cover what's involved.

- **Size:** A VM's size defines its processor speed, amount of memory, initial amount of storage, and expected network bandwidth.
- **Region:** A region is a set of Azure data centers in a named geographic location. Every Azure resource, including virtual machines, is assigned a region. East US and North Europe are examples of regions.
- **Network:** A virtual network is a logically isolated network on Azure. Each virtual machine on Azure is associated with a virtual network. Azure provides cloud-level firewalls for your virtual networks called **network security groups**.
- **Resource groups:** Virtual machines and other cloud resources are grouped into logical containers called resource groups. Groups are typically used to organize sets of resources that are deployed together as part of an application or service. You refer to a resource group by its name.

2.3.3 CREATING RESOURCES IN AZURE

2.3.3.1.1 CREATING RESOURCE GROUP

Normally, the first thing we'd do is to **create a resource group** to hold all the things that we need to create. This allows us to administer all the VMs, disks, network interfaces, and other elements that make up our solution as a unit.

We can use the Azure CLI to create a resource group with the **az group create** command. It takes a **--name** to give it a unique name in our subscription, and a **--location** to tell Azure what area of the world we want the resources to be located by default.

2.3.3.1.2 CREATING VIRTUAL MACHINE

Example Azure CLI command:

```
az vm create \  
  --name myVM \  
  --resource-group Learn-ed49b49c-4d23-418b-89c2-7d6186579226 \  
  --image Win2019Datacenter \  
  --size Standard_DS2_v2 \  
  --location eastus \  
  --admin-username $USERNAME \  
  --admin-password $PASSWORD
```

By default, Azure assigns a public IP address to your VM. You can configure a VM to be accessible from the Internet or only from the internal network.

Let's review the command you just ran.

- The VM is **named** myVM. This name identifies the VM in Azure. It also becomes the VM's internal hostname, or computer name.
- The **resource group**, or the VM's logical container, is named Learn-ed49b49c-4d23-418b-89c2-7d6186579226.
- Win2019Datacenter specifies the Windows Server 2019 **VM image**.
- Standard_DS2_v2 refers to the **size of the VM**. This size has two virtual CPUs and 7 GB of memory.
- The **username** and **password** enable you to connect to your VM later. For example, you can connect over Remote Desktop or WinRM to work with and configure the system.

When the VM is ready, you see information about it. Here's an example.

```
{  
  "fqdns": "",  
  "id": "/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myResourceGroup/providers/Microsoft.Compute/virtualMachines/myVM",  
  "location": "eastus",  
  "macAddress": "00-0D-3A-1E-1B-3B",  
  "powerState": "VM running",  
  "privateIpAddress": "10.0.0.5",  
  "publicIpAddress": "104.211.9.245",  
  "resourceGroup": "myResourceGroup",  
  "zones": ""  
}
```

2.3.3.1.3 VERIFY YOUR VM IS RUNNING

Run the following `az vm get-instance-view` command to verify that the VM was successfully created and is running:

```
az vm get-instance-view \  
  --name myVM \  
  --resource-group Learn-ed49b49c-4d23-418b-89c2-7d6186579226 \  
  --output table
```

The output you see resembles this.

Name	ResourceGroup	Location	ProvisioningState	PowerState
myVM	Learn-ed49b49c-4d23-418b-89c2-7d6186579226	eastus	Succeeded	VM running

You see the VM's name, its resource group, and its location. You also see that the VM was provisioned, or created, successfully and that it's running.

2.4 ADD A WEB SERVER

To configure a VM, you have several choices. You can connect directly and interactively configure your system. Manual configuration is a good start, but as you add systems, you can automate your deployments. Automation involves running repeatable processes such as programs and scripts that take care of the heavy lifting for you.

2.4.1 CONFIGURE IIS

The **Custom Script Extension** is an easy way to download and run scripts on your Azure VMs. It's just one of the many ways you can configure the system once your VM is up and running.

You can store your scripts in Azure storage or in a public location such as GitHub. You can run scripts manually or as part of a more automated deployment. Here, you'll run an Azure CLI command to download a PowerShell script from GitHub and execute it on your VM. The script configures IIS.

Internet Information Services, or IIS, is a web server that runs on Windows. You can use IIS to serve standard web content (HTML, CSS, and JavaScript) or run ASP.NET and other kinds of web applications. IIS comes with Windows Server, but you need to activate it to start serving web pages.

Here you'll use the Custom Script Extension to configure IIS remotely on your VM from Cloud Shell. You'll also configure the firewall to allow inbound network access on port 80 (HTTP).

From Cloud Shell, run this `az vm extension set` command to download and execute a PowerShell script that installs IIS and configures a basic home page.


```
az vm extension set \
  --resource-group Learn-ed49b49c-4d23-418b-89c2-7d6186579226 \
  --vm-name myVM \
  --name CustomScriptExtension \
  --publisher Microsoft.Compute \
  --settings '{"fileUri':['https://raw.githubusercontent.com/MicrosoftDocs/mslearn-welcome-to-azure/master/configure-iis.ps1']}' \
  --protected-settings '{"commandToExecute': 'powershell -ExecutionPolicy Unrestricted -File configure-iis.ps1'}"
```

The PowerShell file that is executed:

```
# Install IIS.
dism /online /enable-feature /featurename:IIS-WebServerRole

# Set the home page.
Set-Content `
  -Path "C:\inetpub\wwwroot\Default.htm" `
  -Value "<html><body><h2>Welcome to Azure! My name is $($env:computername).</h2></body></html>"
```

The script installs IIS and configures the home page to display a welcome message along with the VM's computer name, "myVM". The process to configure IIS, set the contents of the homepage, and start the service takes a couple minutes to complete.

Run this **az vm open-port** command to open port 80 (HTTP) through the firewall.

```
az vm open-port \
  --name myVM \
  --resource-group Learn-ed49b49c-4d23-418b-89c2-7d6186579226 \
  --port 80
```

Now that IIS is set up, let's verify that it's running. Run this **az vm show** command to list your VM's public IP address.

```
az vm show \
  --name myVM \
  --resource-group Learn-ed49b49c-4d23-418b-89c2-7d6186579226 \
  --show-details \
  --query [publicIps] \
  --output tsv
```

2.5 SCALE UP

*Tip: The cloud is elastic. You can scale down or scale in your deployment if you needed to scale up or scale out only temporarily. Scaling down or scaling in can help you save money. **Azure Advisor and Azure Cost Management** are two services that help you optimize cloud spend. You can use these services to identify where you're using more than you need, and then scale back to the capacity you're actually using.*

From Cloud Shell, run `az vm resize` to increase your VM's size to Standard_DS3_v2.

```
az vm resize \  
  --resource-group Learn-ed49b49c-4d23-418b-89c2-7d6186579226 \  
  --name myVM \  
  --size Standard_DS3_v2
```

Run `az vm show` to verify that your VM is running the new size.

```
az vm show \  
  --resource-group Learn-ed49b49c-4d23-418b-89c2-7d6186579226 \  
  --name myVM \  
  --query "hardwareProfile" \  
  --output tsv
```

3 CORE CLOUD SERVICES - AZURE ARCHITECTURE AND SERVICE GUARANTEES

Azure provides a global network of secure datacenters you can deploy your services into. Learn about the physical architecture of Azure, how redundancy is provided, and what sort of service guarantees Microsoft provides.

In this module, you will:

- Explore the physical structure of Azure infrastructure
- Understand the service level agreements provided by Azure
- Learn how to provide your own service level agreements for your apps

3.1 REGIONS

Microsoft Azure is made up of **datacenters** located around the globe. When you leverage a service or create a resource such as a SQL database or virtual machine, you are using physical equipment in one or more of these locations.

The specific datacenters aren't exposed to end users directly; instead, Azure organizes them into regions. A **region** is a geographical area on the planet containing at least one, but potentially multiple datacenters that are nearby and networked together with a low-latency network.

Azure intelligently assigns and controls the resources within each region to ensure workloads are appropriately balanced. It also provides better scalability, redundancy, and preserves data residency for your services. Azure has specialized regions that you might want to use when building out your applications for compliance or legal purposes.

Azure has more global regions than any other cloud provider. This gives you the flexibility to bring applications closer to your users no matter where they are. It also provides better scalability, redundancy, and preserves data residency for your services.

Regions are what you use to identify the location for your resources, but there are two other terms you should also be aware of: **geographies** and **availability zones**.

3.2 GEOGRAPHIES

Azure divides the world into **geographies** that are defined by geopolitical boundaries or country borders. An Azure geography is a discrete market typically containing two or more regions that preserve data residency and compliance boundaries. This division has several benefits.

- Geographies allow customers with specific data residency and compliance needs to keep their data and applications close.
- Geographies ensure that data residency, sovereignty, compliance, and resiliency requirements are honored within geographical boundaries.
- Geographies are fault-tolerant to withstand complete region failure through their connection to dedicated high-capacity networking infrastructure.

Data residency refers to the physical or geographic location of an organization's data or information. It defines the legal or regulatory requirements imposed on data based on the country or region in which it resides and is an important consideration when planning out your application data storage.

Geographies are broken up into the following areas:


- Americas
- Europe
- Asia Pacific
- Middle East and Africa

Each region belongs to a single geography and has specific service availability, compliance, and data residency/sovereignty rules applied to it.

3.3 AVAILABILITY ZONES

You want to ensure your services and data are redundant so you can protect your information in case of failure. When you are hosting your infrastructure, this requires creating duplicate hardware environments. Azure can help make your app highly available through **Availability Zones**.

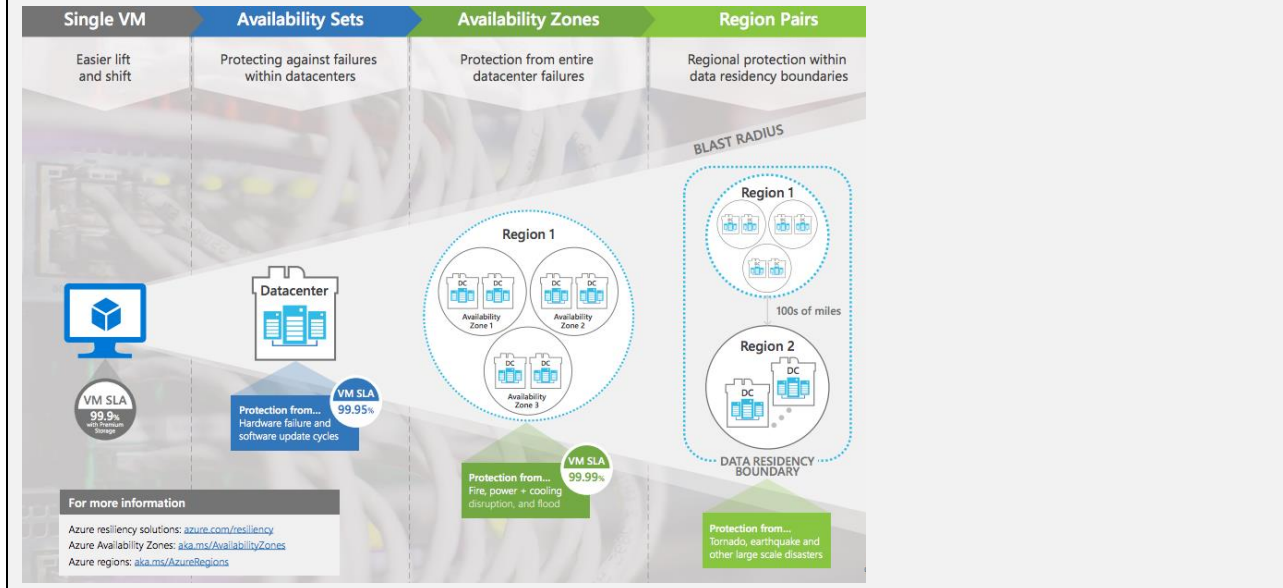
Availability zones are physically separate datacenters within an Azure region.



Availability zones

Each availability zone is made up of one or more datacenters equipped with independent power, cooling, and networking. It is set up to be an *isolation boundary*. If one zone goes down, the other continues working. Availability zones are connected through high-speed, private fiber-optic networks. You can use availability zones to run mission-critical applications and build high-availability into your application architecture by co-locating your compute, storage, networking, and data resources within a zone and replicating in other zones. Keep in mind that there could be a cost to duplicating your services and transferring data between zones.

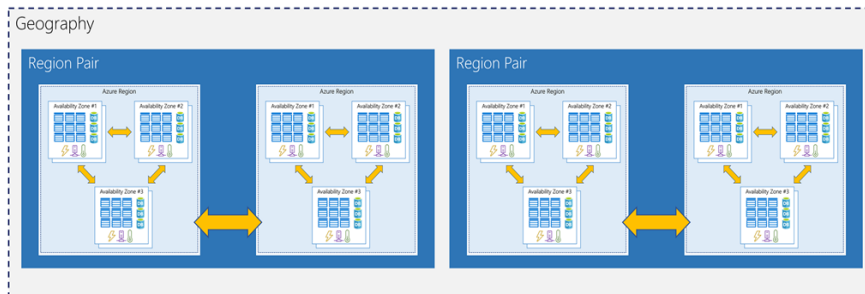
Don't confuse an availability set with an availability zone. An availability set is a group with two or more virtual machines in the same data center. An availability set ensures that at least one of the virtual machines hosted on Azure will be available if something happens. This configuration offers 99.95% SLA. An availability zone is the next level of Azure Virtual Machines high-availability, because virtual machines are in different physical locations within an Azure Region. It can be deployed using one or more Virtual Machines in an Azure Region. Availability zones offer 99.99% SLA where availability sets offer 99.95% SLA.



3.4 REGION PAIRS

Availability zones are created using one or more datacenters, and there are a minimum of three zones within a single region. However, it's possible that a large enough disaster could cause an outage big enough to affect even two datacenters. That's why Azure also creates **region pairs**.

Each Azure region is always paired with another region within the same geography (such as US, Europe, or Asia) at least 300 miles away. This approach allows for the replication of resources (such as virtual machine storage) across a geography that helps reduce the likelihood of interruptions due to events such as natural disasters, civil unrest, power outages, or physical network outages affecting both regions at once.



Region pairs

Since the pair of regions is directly connected and far enough apart to be isolated from regional disasters, you can use them to provide reliable services and data redundancy. Some services offer automatic geo-redundant storage using region pairs.

3.5 SERVICE LEVEL AGREEMENTS FOR AZURE

Microsoft maintains its commitment to providing customers with high-quality products and services by adhering to comprehensive operational policies, standards, and practices. Formal documents called **Service-Level Agreements (SLAs)** capture the specific terms that define the performance standards that apply to Azure.

- SLAs describe Microsoft's commitment to providing Azure customers with specific performance standards.
- There are SLAs for individual Azure products and services.
- SLAs also specify what happens if a service or product fails to perform to a governing SLA's specification.

There are three key characteristics of SLAs for Azure products and services:

- Performance Targets
- Uptime and Connectivity Guarantees
- Service credits

3.5.1 PERFORMANCE TARGETS

An SLA defines performance targets for an Azure product or service. The performance targets that an SLA defines are specific to each Azure product and service. For example, performance targets for some Azure services are expressed as uptime guarantees or connectivity rates.

3.5.2 UPTIME AND CONNECTIVITY GUARANTEES


A typical SLA specifies performance-target commitments that range from 99.9 percent ("three nines") to 99.999 percent ("five nines"), for each corresponding Azure product or service. These targets can apply to such performance criteria as uptime or response times for services.

3.5.3 SERVICE CREDITS

SLAs also describe how Microsoft will respond if an Azure product or service fails to perform to its governing SLA's specification. For example, customers may have a discount applied to their Azure bill, as compensation for an under-performing Azure product or service.


3.6 COMPOSING SLAS ACROSS SERVICES

When combining SLAs across different service offerings, the resultant SLA is called a Composite SLA. The resulting composite SLA can provide higher or lower uptime values, depending on your application architecture. For example: $99.95 \text{ percent} \times 99.99 \text{ percent} = 99.94 \text{ percent}$



This means the **combined probability of failure** is higher than the individual SLA values. This isn't surprising, because an application that relies on multiple services has more potential failure points.

Conversely, you can improve the composite SLA by creating independent fallback paths. For example, if SQL Database is unavailable, you can put transactions into a queue for processing at a later time.



With this design, the application is still available even if it can't connect to the database. However, it fails if both the database and the queue fail simultaneously.

If the expected percentage of time for a simultaneous failure is 0.0001×0.001 , the composite SLA for this combined path of a database or queue would be: $1.0 - (0.0001 \times 0.001) = 99.99999 \text{ percent}$

Therefore, if we add the queue to our web app, the total composite SLA is: $99.95 \text{ percent} \times 99.99999 \text{ percent} = \sim 99.95 \text{ percent}$

Notice we've improved our SLA behavior. However, there are trade-offs to using this approach: the application logic is more complicated, you are paying more to add the queue support, and there may be data-consistency issues you'll have to deal with due to retry behavior.

3.7 IMPROVE YOUR APP RELIABILITY

You can use SLAs to evaluate how your Azure solutions meet business requirements and the needs of your clients and users. By creating your own SLAs, you can set performance targets to suit your specific Azure application. This approach is known as an *Application SLA*. It's important to understand the Azure SLAs that define performance targets for the Azure products and services within your solution. This understanding will help you create achievable Application SLAs.

Resiliency is the ability of a system to recover from failures and continue to function. It's not about avoiding failures, but responding to failures in a way that avoids downtime or data loss. The goal of resiliency is to return the application to a fully functioning state following a failure. High availability and disaster recovery are two crucial components of resiliency.

4 CREATE AN AZURE ACCOUNT

In this module, you will:

- Learn about the different types of Azure accounts and subscriptions
- Understand how billing works in Azure
- Create a free Azure account
- Learn how to get help when you need it with different support options

With a free Azure account and subscription, you can build, test, and deploy enterprise applications, create custom web and mobile experiences, and gain insights from your data through machine learning and powerful analytics.

4.1 AZURE ACCOUNTS AND SUBSCRIPTIONS

An **Azure account** is tied to a specific identity and holds information like:

- Name, email, and contact preferences
- Billing information such as a credit card

An Azure account is what you use to sign in to the Azure website and administer or deploy services. Every Azure account is associated with one or more subscriptions. An **Azure subscription** is a logical container used to provision resources in Microsoft Azure. It holds the details of all your resources like virtual machines, databases, etc.

An Azure subscription has a trust relationship with **Azure Active Directory (Azure AD)**, which means that the subscription trusts Azure AD to authenticate users, services, and devices. Multiple subscriptions can trust the same Azure AD directory, but each subscription can only trust a single directory.

Every Azure Subscription includes:

- Free access to billing and subscription support
- Azure products and services documentation
- Online self-help documentation
- Community support forums

4.1.1 AZURE FREE SUBSCRIPTION

An Azure **free subscription** includes a \$200 credit to spend on any service for the first 30 days, free access to the most popular Azure products for 12 months, and access to more than 25 products that are always free. This is an excellent way for new users to get started. To set up a free subscription, you need a phone number, a credit card, and a Microsoft account.

4.1.2 AZURE PAY-AS-YOU-GO SUBSCRIPTION

A **Pay-As-You-Go (PAYG)** subscription charges you monthly for the services you used in that billing period. This subscription type is appropriate for a wide range of users, from individuals to small businesses, and many large organizations as well.

4.1.3 PAY-AS-YOU-GO DEV/TEST

Designed for the needs of teams of Visual Studio subscribers, the Pay-As-You-Go Dev/Test offer allows you to quickly get your team up and running with dev/test environments in the cloud using pre-configured virtual machines, including Windows 10, and low rates on virtual machines, Cloud Services, SQL Database, HDInsight, App Service and Logic Apps. You have the flexibility to create multiple Azure subscriptions based on this offer, enabling you to maintain isolated environments and a separate bill for different projects or teams.

This offer is exclusively for active Visual Studio subscribers and is limited to development and testing only. This benefit is exclusively for development and testing your applications. Usage within the subscription does not carry a financially-backed SLA, except for use of Visual Studio Team Services and HockeyApp.

4.1.4 AZURE ENTERPRISE AGREEMENT


An **Enterprise Agreement** provides flexibility to buy cloud services and software licenses under one agreement, with discounts for new licenses and Software Assurance. It's targeted at enterprise-scale organizations.

4.1.5 AZURE FOR STUDENTS SUBSCRIPTION

An Azure for Students subscription includes \$100 in Azure credits to be used within the first 12 months plus select free services without requiring a credit card at sign-up. You must verify your student status through your organizational email address.

4.2 USING MULTIPLE AZURE SUBSCRIPTIONS

You can create multiple subscriptions under a single Azure account. This is particularly useful for businesses because access control and billing occur at the subscription level, not the account level.



Accounts and subscriptions

4.2.1 ACCESS MANAGEMENT

You can create separate subscriptions on your Azure account to reflect different organizational structures. For example, you could limit engineering to lower-cost resources, while allowing the IT department a full range. This design allows you to manage and control access to the resources that users provision within each subscription.

4.2.2 BILLING

One bill is generated for every Azure subscription on a monthly basis. The payment is charged automatically to the associated account credit or debit card within 10 days after the billing period ends. On your credit card statement, the line item would say MSFT Azure.

You can set spending limits on each subscription to ensure you aren't surprised at the end of the month. Reports can be generated by subscriptions, if you have multiple internal departments and need to do "chargeback," a possible scenario is to create subscriptions by department or project.

4.3 AUTHENTICATE ACCESS WITH AZURE ACTIVE DIRECTORY

As you've seen, your Azure account is a globally unique entity that gives you access to your Azure subscriptions and services. Authentication for your account is performed using Azure Active Directory (Azure AD). **Azure AD** is a modern identity provider that supports multiple authentication protocols to secure applications and services in the cloud.

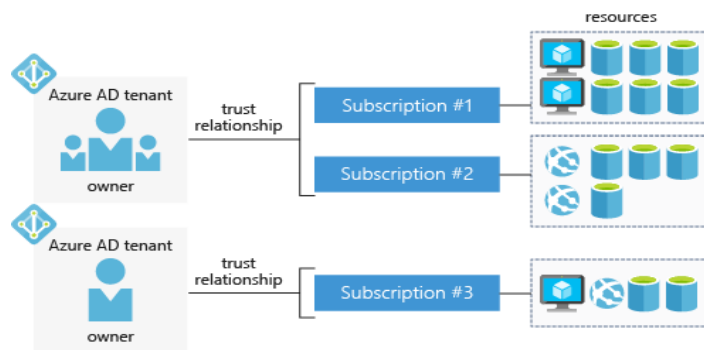
*Azure AD is not the same as **Windows Active Directory**. Windows Active Directory is focused on securing Windows desktops and servers. In contrast, Azure AD is all about web-based authentication standards such as OpenID and OAuth.*

Users, applications, and other entities registered in Azure AD aren't all lumped into a single global service. Instead, Azure AD is partitioned into separate *tenants*. A tenant is a dedicated, isolated instance of the Azure Active Directory service, owned and managed by an organization. When you sign up for a Microsoft cloud service subscription such as Microsoft Azure, Microsoft Intune, or Office 365, a dedicated instance of Azure AD is automatically created for your organization.

When it comes to Azure AD tenants, there is no concrete definition of "organization" — tenants can be owned by individuals, teams, companies, or any other group of people. Tenants are commonly associated with companies. If you sign up for Azure with an email address that's not associated with an existing tenant, the sign-up process will walk you through creating a tenant, owned entirely by you.

Azure AD tenants and subscriptions have a many-to-one trust relationship: A tenant can be associated with multiple Azure subscriptions, but every subscription is associated with only one tenant. This structure allows organizations to manage multiple subscriptions and set security rules across all the resources contained within them.

Here's a simple representation of how accounts, subscriptions, tenants, and resources work together.



Azure AD Tenant

Notice that each Azure AD tenant has an **account owner**. This is the original Azure account that is responsible for billing. You can add additional users to the tenant, and even invite guests from other Azure AD tenants to access resources in subscriptions.

4.4 AZURE SUPPORT OPTIONS

One final thing to know about subscriptions is how to get support when you need it. Every Azure subscription includes free access to the following essential support services:

- Billing and subscription support
- Azure products and services documentation
- Online self-help documentation
- Whitepapers
- Community support forums

4.4.1 PAID AZURE SUPPORT PLANS

Microsoft offers four paid Azure support plans for customers who require technical and operational support:

- Developer
- Standard
- Professional Direct
- Premier

4.4.1.1 DEVELOPER

The Azure Developer offering is appropriate for companies or individuals using Microsoft Azure in a non-production environment or for trial and evaluation.

- Reactive technical support
- Support for non-Microsoft technologies running on Azure
- Lowest priced technical support option

RANGE OF SUPPORT	MICROSOFT AZURE
Unlimited billing & subscription support	Business hours only
Unlimited technical support	Business hours only
Non-Microsoft technologies running on Azure (See FAQ)	Microsoft will assist Azure customers with issues associated with select non-Microsoft technologies.
Incident submission	Online
Initial response time	< 8 hours
Maximum severity	“C” (Minimal business impact)

4.4.1.2 STANDARD

The Azure Standard offering is a good choice for small or mid-size companies with minimal business critical dependence on Microsoft Azure.

- Reactive 24x7 technical support
- Fast initial response for support issues
- Ability to set severity of issues

RANGE OF SUPPORT	MICROSOFT AZURE
Unlimited 24x7 billing & subscription support	V
Unlimited 24x7 technical support	V
Non-Microsoft technologies running on Azure (See FAQ)	Microsoft will assist Azure customers with issues associated with select non-Microsoft technologies.
Incident submission	Online
Initial response time	< 1 hours
Maximum severity	"A" (Critical business impact)

4.4.1.3 PROFESSIONAL DIRECT

The Azure ProDirect offering is most appropriate for mid-size to large companies with substantial business critical utilization of Microsoft Azure.

- Fastest initial response of one hour or less for high-severity support requests
- Escalation management for priority issues
- Pooled team of ProDirect Managers provide account management

RANGE OF SUPPORT	MICROSOFT AZURE
Unlimited 24x7 billing & subscription support	Included
Unlimited 24x7 technical support	Included
Non-Microsoft technologies running on Azure	Microsoft will assist Azure customers with issues associated with select non-Microsoft technologies.
Escalation management	Included
Advisory services	Guidance based on best practices to: plan for deployments and migrations, plan for hybrid cloud solutions, boost performance, improve reliability and recoverability, enhance security
Pooled service account management	A pooled team dedicated to helping you get the most out of your Azure service: Getting started with ProDirect session, service reviews, "Ask the Experts" webinars, enhanced outage communications, Azure Advisor consultations
Incident submission	Online
Initial response time	< 1 hour
Maximum severity 2	"A" (Critical business impact)

4.4.1.4 PREMIER

The Premier offering is well suited for large or global enterprises with strategic and business critical dependence on Microsoft products including Azure.

- Complete coverage for cloud, hybrid, and on premises solutions across all Microsoft products
- Support available onsite in addition to online
- Technical Account Manager is assigned to account

RANGE OF SUPPORT	ALL MICROSOFT PRODUCTS
Unlimited 24x7 billing & subscription support	Available
Unlimited 24x7 technical support	Available
Non-Microsoft technologies running on Azure	Microsoft will assist Azure customers with issues associated with select non-Microsoft technologies.
Escalation management	Available
Advisory services	Assistance based on best practices to resolve how-to scenarios
Service account management	Monthly service review, service delivery planning, remediation planning, reporting and trending advice, operational guidance, onboarding
Proactive services	Risk assessments, operations assessments, Microsoft product support workshops timed with key release cycles, guidance to help maximize current Microsoft platform and support business initiatives, monitor severity "A" support requests
Cloud service dependency mapping	Available
Architecture/code review	Available
Onsite support	Available
Incident submission	Submitted by assigned Technical Account Manager (TAM), online, dedicated phone line
Initial response time	< 15 minutes
Maximum severity	"A" (Critical business impact)

4.4.1.5 COMPARISON CHART

Topic	BASIC	DEVELOPER	STANDARD	PROFESSIONAL DIRECT	PREMIER
Scope	Available to all Microsoft Azure accounts	Trial and non-production environments	Production workload environments	Business-critical dependence	Substantial dependence across multiple products
Customer Service, Self-Help and Communities	24x7 access to billing and subscription support, online self-help, documentation, whitepapers, and support forums	24x7 access to billing and subscription support, online self-help, documentation, whitepapers, and support forums	24x7 access to billing and subscription support, online self-help, documentation, whitepapers, and support forums	24x7 access to billing and subscription support, online self-help, documentation, whitepapers, and support forums	24x7 access to billing and subscription support, online self-help, documentation, whitepapers, and support forums
Best Practices	Access to full set of Azure Advisor recommendations	Access to full set of Azure Advisor recommendations	Access to full set of Azure Advisor recommendations	Access to full set of Azure Advisor recommendations	Access to full set of Azure Advisor recommendations
Health Status and Notifications	Access to personalized Service Health Dashboard & Health API	Access to personalized Service Health Dashboard & Health API	Access to personalized Service Health Dashboard & Health API	Access to personalized Service Health Dashboard & Health API	Access to personalized Service Health Dashboard & Health API
Technical Support	N/A	Business hours access to Support Engineers via email	24x7 access to Support Engineers via email and phone	24x7 access to Support Engineers via email and phone	24x7 access to Support Engineers via email and phone
Who Can Open Cases	N/A	Unlimited contacts / unlimited cases	Unlimited contacts / unlimited cases	Unlimited contacts / unlimited cases	Unlimited contacts / unlimited cases
Third-Party Software Support	N/A	Interoperability & configuration guidance and troubleshooting	Interoperability & configuration guidance and troubleshooting	Interoperability & configuration guidance and troubleshooting	Interoperability & configuration guidance and troubleshooting
Architecture Support	N/A	General guidance	General guidance	Architectural guidance based on best practice delivered by ProDirect Delivery Manager	Customer specific architectural support such as design reviews, performance tuning, configuration and implementation assistance delivered by Microsoft Azure technical specialists.
Operations Support	N/A	N/A	N/A	Onboarding services, service reviews, Azure Advisor consultations	Technical account manager-led service reviews and reporting
Training	N/A	N/A	N/A	Azure Engineering-led web seminars	Azure Engineering-led web seminars, on-demand training
Proactive Guidance	N/A	N/A	N/A	ProDirect Delivery Manager	Designated Technical Account Manager
Launch Support	N/A	N/A	N/A	N/A	Azure Event Management (available for additional fee)
Case Severity/Response Times		Sev C: <8 bh	Sev C: <8 bh Sev B: <4h Sec A: <1h	Sev C: <4 bh Sev B: <2h Sev A: <1h	Sev C: <4bh Sev B: <2h Sev A: <1h or < 15 min with Azure Rapid Response or Azure Event Management

4.4.2 SUPPORT-PLAN AVAILABILITY AND BILLING

The support plans available and how you're charged depends on the type of Azure customer you are, and the type of Azure subscription you have.

For example, Developer support isn't available to Enterprise customers. Enterprise customers can purchase Standard, Professional Direct, and Premier support plans, and be billed for support as part of an Enterprise Agreement (EA). Alternatively, if you purchase a support plan within a pay-as-you-go subscription, your support plan is charged to your monthly Azure subscription bill.

4.4.3 OTHER SUPPORT OPTIONS

Several additional support channels are available outside Azure's official support plans:

- [Azure Knowledge Center](#)
- [Microsoft Developer Network \(MSDN\) Forums](#)
- [Stack Overflow](#)
- [Server Fault](#)
- [Azure Feedback Forums](#)
- [Twitter](#)

5 CORE CLOUD SERVICES - MANAGE SERVICES WITH THE AZURE PORTAL

Azure is a cloud platform that provides the compute, storage, and networking resources needed to build cloud-hosted applications. As a new user, the Azure Portal is likely to be the primary way you will interact with Azure. The Azure Portal lets you create and manage all your Azure resources. For example, you can set up a new database, increase the compute power of your virtual machines, and monitor your monthly costs. It's also a great learning tool since you can survey all available resources and use guided wizards to create the ones you need.

Here you will learn how to sign in to the portal and navigate the portal interface. You will also learn how to customize the dashboard, so it is convenient to locate and monitor your most essential services.

In this module, you will:

- Learn about Azure management options
- Navigate the Azure portal
- Customize the dashboard

5.1 AZURE MANAGEMENT OPTIONS

You can configure and manage Azure using a broad range of tools and platforms. There are tools available for the command line, language-specific Software Development Kits (SDKs), developer tools, tools for migration, and many others.

Tools that are commonly used for day-to-day management and interaction include:

- **Azure portal** for interacting with Azure via a Graphical User Interface (GUI)
- **Azure PowerShell** and **Azure Command-Line Interface (CLI)** for command line and automation-based interactions with Azure
- **Azure Cloud Shell** for a web-based command-line interface
- **Azure mobile app** for monitoring and managing your resources from your mobile device

5.1.1 AZURE PORTAL

The Azure portal is a public website that you can access with any web browser. Once you sign in with your Azure account, you can create, manage and monitor any available Azure services. The dashboard view provides high-level details about your Azure environment. You can customize the dashboard by moving and resizing tiles, and displaying services you're interested in.

The portal doesn't provide any way to automate repetitive tasks. For example, to set up multiple VMs, you would need to create them one at a time by completing the wizard for each VM. This makes the portal approach time-consuming and error-prone for complex tasks.

5.1.2 AZURE POWERSHELL

Azure PowerShell enables you to connect to your Azure subscription and manage resources. Windows PowerShell and PowerShell Core provide services such as the shell window and command parsing. Azure PowerShell then adds the Azure-specific commands.

For example, Azure PowerShell provides the **New-AzureRmVM** command that creates a virtual machine for you inside your Azure subscription. To use it, you would launch PowerShell, install the Azure PowerShell module, sign in to your Azure account using the command **Connect-AzureRMAccount**, and then issue a command such as:

```
New-AzureRmVM `
  -ResourceGroupName "MyResourceGroup" `
  -Name "TestVm" `
  -Image "UbuntuLTS"
  ...
```

Creating administration scripts and using automation tools is a powerful way to optimize your workflow. You can automate repetitive tasks. Once a script is verified, it runs consistently, which can reduce errors. Another scripting environment is the **Azure CLI**.

5.1.3 AZURE CLI

Azure CLI is a cross-platform command-line program that connects to Azure and executes administrative commands on Azure resources.

```
az vm create `
  --resource-group MyResourceGroup `
  --name TestVm `
  --image UbuntuLTS
  --generate-ssh-keys
  ...
```

5.1.4 AZURE CLOUD SHELL

Azure Cloud Shell is a browser-based scripting environment for command-line administration of Azure resources. It provides support for two shell environments. Linux users can opt for a Bash experience, while Windows users can use PowerShell. In addition to these administrative tools, the Cloud Shell has a suite of developer tools, text editors, and other tools.

5.1.5 AZURE MOBILE APP

The Microsoft Azure mobile app allows you to access, manage, and monitor all your Azure accounts and resources from your iOS or Android phone or tablet.

5.1.6 OTHER OPTIONS

There are also Azure SDKs for a range of languages and frameworks, and REST APIs that you can use to manage and control Azure resources programmatically.

5.2 NAVIGATE THE PORTAL

The portal is a web-based administration site that lets you interact with all of your subscriptions and resources you have created. Almost everything you do with Azure can be done through this web interface.

5.2.1 AZURE PORTAL LAYOUT

The Azure portal is the primary graphical user interface (GUI) for controlling Microsoft Azure. You can carry out the majority of management actions in the portal, and it is typically the best interface for carrying out single tasks or where you want to look at the configuration options in detail.

5.2.2 WHAT IS A BLADE?

The Azure portal uses a **blades model** for navigation. A blade is a slide-out panel containing the UI for a single level in a navigation sequence. For example, each of these elements in this sequence would be represented by a blade: Virtual machines > Compute > Ubuntu Server. Each blade contains some information and configurable options. Some of these options generate another blade, which reveals itself to the right of any existing blade. On the new blade, any further configurable options will spawn another blade, and so on. Soon, you can end up with several blades open at the same time. You can maximize blades as well so that they fill the entire screen.

5.2.3 WHAT IS THE AZURE MARKETPLACE?

The Marketplace allows customers to find, try, purchase, and provision applications and services from hundreds of leading service providers, all certified to run on Azure. With Azure Marketplace, customers can discover technical applications built for or built on Azure. It combines Microsoft Azure's market of solutions and services into a single, unified platform to discover, try, buy, or deploy solutions in just a few clicks.

The solution catalog spans several industry categories, including but not limited to open-source container platforms, virtual machine images, databases, application build and deployment software, developer tools, threat detection, and blockchain. Using Azure Marketplace, you can provision end-to-end solutions quickly and reliably, hosted in your own Azure environment. At the time of writing, this includes over 8,000 listings.

The Azure Marketplace offers technical solutions and services from Microsoft and partners built to extend Azure products and services. The solution catalog spans several categories, including but not limited to:

- base operating systems
- databases
- security
- identity
- networking
- blockchain
- developer tools
- ...and more

Azure Marketplace offers SaaS applications, Virtual Machines, Solution Templates, Azure-Managed applications, and consulting services.

5.3 CONFIGURING SETTINGS IN THE AZURE PORTAL

If you click the Cloud Shell icon (>_), you will create a new Azure Cloud Shell session. Recall that Azure Cloud Shell is an interactive, browser-accessible shell for managing Azure resources. It provides the flexibility of choosing the shell experience that best suits the way you work. Linux users can opt for a Bash experience, while Windows users can opt for PowerShell. This browser-based terminal lets you control and administer all of your Azure resources in the current subscription through a command-line interface built right into the portal.

6 CORE CLOUD SERVICES - AZURE COMPUTE FUNCTIONS

In this module, you will:

- Identify compute options in Azure
- Select compute options that are appropriate for your business

6.1 ESSENTIAL AZURE COMPUTE CONCEPTS

6.1.1 WHAT IS AZURE COMPUTE?

Azure compute is an on-demand computing service for running cloud-based applications. It provides computing resources like multi-core processors and supercomputers via virtual machines and containers. It also provides serverless computing to run apps without requiring infrastructure setup or configuration.

There are four common techniques for performing compute in Azure:

- Virtual machines
- Containers
- Azure App Service
- Serverless computing

6.1.2 WHAT ARE VIRTUAL MACHINES?

Virtual machines, or VMs, are software emulations of physical computers. They include a virtual processor, memory, storage, and networking resources. They host an operating system (OS), and you're able to install and run software just like a physical computer. And by using a remote desktop client, you can use and control the virtual machine as if you were sitting in front of it.

6.1.3 WHAT ARE CONTAINERS?

Containers are a virtualization environment for running applications. Just like virtual machines, containers are run on top of a host operating system but unlike VMs, they don't include an operating system for the apps running inside the container. Instead, containers bundle the libraries and components needed to run the application and use the existing host OS running the container. For example, if five containers are running on a server with a specific Linux kernel, all five containers and the apps within them share that same Linux kernel.

6.1.4 WHAT IS AZURE APP SERVICE?

Azure App Service is a platform-as-a-service (PaaS) offering in Azure that is designed to host enterprise-grade web-oriented applications. You can meet rigorous performance, scalability, security, and compliance requirements while using a fully managed platform to perform infrastructure maintenance.

6.1.5 WHAT IS SERVERLESS COMPUTING?

Serverless computing is a cloud-hosted execution environment that runs your code but completely abstracts the underlying hosting environment. You create an instance of the service, and you add your code; no infrastructure configuration or maintenance is required, or even allowed.

6.2 EXPLORE AZURE VIRTUAL MACHINES

Azure Virtual Machines (VMs) let you create and use virtual machines in the cloud. They provide infrastructure as a service (IaaS) in the form of a virtualized server and can be used in many ways. Just like a physical computer, you can customize all of the software running on the VM. VMs are an ideal choice when you need:

- Total control over the operating system (OS)
- The ability to run custom software
- To use custom hosting configurations

You can create and provision a VM in minutes when you select a pre-configured VM image. Selecting an image is one of the most important decisions you'll make when creating a VM. An **image** is a template used to create a VM. These templates already include an OS and often other software, like development tools or web hosting environments.

6.2.1 SCALING VMS IN AZURE

You can run single VMs for testing, development, or minor tasks, or group VMs together to provide high availability, scalability, and redundancy. Azure has several features so that no matter what your uptime requirements are, Azure can meet them. These features include:


- Availability sets
- Virtual Machine Scale Sets
- Azure Batch

6.2.1.1 WHAT ARE AVAILABILITY SETS?

An **availability set** is a logical grouping of two or more VMs that help keep your application available during planned or unplanned maintenance.

A *planned maintenance event* is when the underlying **Azure fabric** that hosts VMs is updated by Microsoft. A planned maintenance event is done to patch security vulnerabilities, improve performance, and add or update features. Most of the time these updates are done without any impact to the guest VMs. But sometimes VMs require a reboot to complete an update. When the VM is part of an availability set, the Azure fabric updates are sequenced so not all of the associated VMs are rebooted at the same time. VMs are put into different **update domains**. Update domains indicate groups of VMs and underlying physical hardware that can be rebooted at the same time. Update domains are a logical part of each data center and are implemented with software and logic.

Unplanned maintenance events involve a hardware failure in the data center, such as a power outage or disk failure. VMs that are part of an availability set automatically switch to a working physical server so the VM continues to run. The group of virtual machines that share common hardware are in the same fault domain. A **fault domain** is essentially a rack of servers. It provides the physical separation of your workload across different power, cooling, and network hardware that support the physical servers in the data center server racks. In the event the hardware that supports a server rack becomes unavailable, only that rack of servers is affected by the outage.



With an availability set, you get:

- Up to three fault domains that each have a server rack with dedicated power and network resources
- Five logical update domains

6.2.1.2 WHAT ARE VIRTUAL MACHINE SCALE SETS?

Azure Virtual Machine Scale Sets let you create and manage a group of identical, load balanced VMs. Imagine you're running a website that enables scientists to upload astronomy images that need to be processed. If you duplicated the VM, you'd normally need to configure an additional service to route requests between multiple instances of the website. VM Scale Sets could do that work for you.

Scale sets allow you to centrally manage, configure, and update a large number of VMs in minutes to provide highly available applications. The number of VM instances can automatically increase or decrease in response to demand or a defined schedule. With VM Scale Sets, you can build large-scale services for areas such as compute, big data, and container workloads.

Scale sets are built from virtual machines. With scale sets, the management and automation layers are provided to run and scale your applications. You could instead manually create and manage individual VMs, or integrate existing tools to build a similar level of automation. The following table outlines the benefits of scale sets compared to manually managing multiple VM instances.

Scenario	Manual group of VMs	Virtual machine scale set
Add additional VM instances	Manual process to create, configure, and ensure compliance	Automatically create from central configuration
Traffic balancing and distribution	Manual process to create and configure Azure load balancer or Application Gateway	Can automatically create and integrate with Azure load balancer or Application Gateway
High availability and redundancy	Manually create Availability Set or distribute and track VMs across Availability Zones	Automatic distribution of VM instances across Availability Zones or Availability Sets
Scaling of VMs	Manual monitoring and Azure Automation	Autoscale based on host metrics, in-guest metrics, Application Insights, or schedule

There is no additional cost to scale sets. You only pay for the underlying compute resources such as the VM instances, load balancer, or Managed Disk storage. The management and automation features, such as autoscale and redundancy, incur no additional charges over the use of VMs.

6.2.1.3 WHAT IS AZURE BATCH?

Azure Batch enables large-scale job scheduling and compute management with the ability to scale to tens, hundreds, or thousands of VMs.

When you're ready to run a job, Batch:

- Starts a pool of compute VMs for you
- Installs applications and staging data
- Runs jobs with as many tasks as you have
- Identifies failures
- Requeues work
- Scales down the pool as work completes

6.3 EXPLORE CONTAINERS IN AZURE

If you wish to run multiple instances of an application on a single virtual machine, containers are an excellent choice. The container orchestrator can start, stop, and scale out application instances as needed.

Containers are meant to be lightweight, created, scaled out, and stopped dynamically. This design allows you to respond quickly to changes in demand or failure.

Another benefit of containers is you can run multiple isolated applications on a single VM host. Since containers are secured and isolated, you don't need separate VMs for each app.

6.3.1 CONTAINERS IN AZURE

Azure supports Docker containers, and there are several ways to manage containers in Azure.

- Azure Container Instances (ACI)
- Azure Kubernetes Service (AKS)

6.3.1.1 AZURE CONTAINER INSTANCES

Azure Container Instances (ACI) offers the fastest and simplest way to run a container in Azure. You don't have to manage any virtual machines or configure any additional services. It is a PaaS offering that allows you to upload your containers and execute them directly.

6.3.1.2 AZURE KUBERNETES SERVICE

The task of automating and managing and interacting with a large number of containers is known as orchestration. **Azure Kubernetes Service (AKS)** is a complete orchestration service for containers with distributed architectures with multiple containers.

6.3.1.3 USING CONTAINERS IN YOUR SOLUTIONS

Containers are often used to create solutions using a *microservice architecture*. This is where you break solutions into smaller, independent pieces. For example, you may split a website into a container hosting your front end, another hosting your back end, and a third for storage. This allows you to separate portions of your app into logical sections that can be maintained, scaled, or updated independently.

6.4 EXPLORE AZURE APP SERVICE

Azure App Service enables you to build and host web apps, background jobs, mobile backends, and RESTful APIs in the programming language of your choice without managing infrastructure. It offers auto-scaling and high availability, supports both Windows and Linux, and enables automated deployments from GitHub, Azure DevOps, or any Git repo to support a continuous deployment model.

This platform as a service (PaaS) allows you to focus on the website and API logic while Azure takes care of the infrastructure to run and scale your web applications.

6.4.1 TYPES OF WEB APPS

With Azure App Service, you can host most common web app styles including:

- **Web Apps**
- **API Apps**
- **WebJobs**
- **Mobile Apps**

Azure App Service handles most of the infrastructure decisions you deal with in hosting web apps: deployment and management are integrated into the platform, endpoints can be secured, sites can be scaled quickly to handle high traffic loads, and the built-in load balancing and traffic manager provide high availability. All of these app styles are hosted in the same infrastructure and share these benefits. This makes App Service the ideal choice to host web-oriented applications.

6.4.1.1 WEB APPS

App Service includes full support for hosting web apps using ASP.NET, ASP.NET Core, Java, Ruby, Node.js, PHP, or Python. You can choose either Windows or Linux as the host operating system.

6.4.1.2 API APPS

Much like hosting a website, you can build REST-based Web APIs using your choice of language and framework. You get full Swagger support, and the ability to package and publish your API in the Azure Marketplace. The produced apps can be consumed from any HTTP(s) based client.

6.4.1.3 WEB JOBS

WebJobs allows you to run a program (.exe, Java, PHP, Python or Node.js) or script (.cmd, .bat, PowerShell, or Bash) in the same context as a web app, API app, or mobile app. They can be scheduled, or run by a trigger. This is often used to run background tasks as part of your application logic.

6.4.1.4 MOBILE APPS

Use the Mobile Apps feature of Azure App Service to quickly build a back-end for iOS and Android apps. With just a few clicks in the Azure portal you can:

- Store mobile app data in a cloud-based SQL database
- Authenticate customers against common social providers such as MSA, Google, Twitter and Facebook
- Send push notifications
- Execute custom back-end logic in C# or Node.js

6.5 EXPLORE SERVERLESS COMPUTING IN AZURE

With *serverless computing*, Azure takes care of managing the server infrastructure and allocation/deallocation of resources based on demand. Infrastructure isn't your responsibility. Scaling and performance are handled automatically, and you are billed only for the exact resources you use. There's no need to even reserve capacity.

You focus solely on the logic you need to execute and the trigger that is used to run your code. You configure your serverless apps to respond to events. This could be a REST endpoint, a periodic timer, or even a message received from another Azure service. The serverless app runs only when it's triggered by an event.

Azure has two implementations of serverless compute:

- **Azure Functions** which can execute code in almost any modern language.
- **Azure Logic Apps** which are designed in a web-based designer and can execute logic triggered by Azure services without writing any code.

6.5.1 AZURE FUNCTIONS

When you're concerned only about the code running your service, and not the underlying platform or infrastructure, Azure Functions are ideal. Azure Functions scale automatically based on demand, so they're a solid choice when demand is variable.

Furthermore, Azure Functions can be either stateless (the default) where they behave as if they're restarted every time they respond to an event), or stateful (called "**Durable Functions**") where a context is passed through the function to track prior activity.

6.5.2 AZURE LOGIC APPS

Azure Logic Apps are similar to Functions - both enable you to trigger logic based on an event. Where Functions execute code, Logic Apps execute workflows built from predefined logic blocks. They are specifically designed to automate your business processes.

You create Logic App workflows using a visual designer on the Azure Portal or in Visual Studio. The workflows are persisted as a JSON file with a known workflow schema.

Azure provides over 200 different connectors and processing blocks to interact with different services - including most popular enterprise apps. You can also build custom connectors and workflow steps if the service you need to interact with isn't covered. You then use the visual designer to link connectors and blocks together, passing data through the workflow to do custom processing - often all without writing any code.

6.5.3 FUNCTIONS VS. LOGIC APPS

Functions and Logic Apps can both create complex orchestrations. An orchestration is a collection of functions or steps, that are executed to accomplish a complex task. With Azure Functions, you write code to complete each step, with Logic Apps, you use a GUI to define the actions and how they relate to one another.

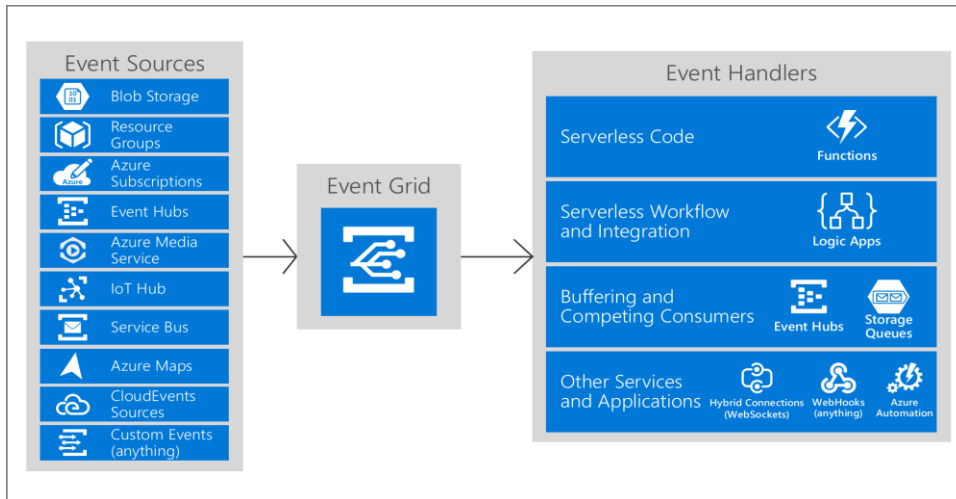
You can mix and match services when you build an orchestration, calling functions from logic apps and calling logic apps from functions. Here are some common differences between the two.

/	Functions	Logic Apps
State	Normally stateless, but Durable Functions provide state	Stateful
Development	Code-first (imperative)	Designer-first (declarative)
Connectivity	About a dozen built-in binding types, write code for custom bindings	Large collection of connectors, Enterprise Integration Pack for B2B scenarios, build custom connectors
Actions	Each activity is an Azure function; write code for activity functions	Large collection of ready-made actions
Monitoring	Azure Application Insights	Azure portal, Log Analytics
Management	REST API, Visual Studio	Azure portal, REST API, PowerShell, Visual Studio
Execution context	Can run locally or in the cloud	Runs only in the cloud.

6.5.4 AZURE EVENT GRID

Azure Event Grid allows you to easily build applications with event-based architectures. First, select the Azure resource you would like to subscribe to, and then give the event handler or WebHook endpoint to send the event to. Event Grid has built-in support for events coming from Azure services, like storage blobs and resource groups. Event Grid also has support for your own events, using custom topics.

You can use filters to route specific events to different endpoints, multicast to multiple endpoints, and make sure your events are reliably delivered.



Azure Event Grid

There are five concepts in Azure Event Grid that let you get going:

- **Events** - What happened.
- **Event sources** - Where the event took place.
- **Topics** - The endpoint where publishers send events.
- **Event subscriptions** - The endpoint or built-in mechanism to route events, sometimes to more than one handler.
- **Subscriptions** are also used by handlers to intelligently filter incoming events.
- **Event handlers** - The app or service reacting to the event.

6.5.4.1 CHOOSE BETWEEN AZURE MESSAGING SERVICES

Azure offers three services that assist with delivering event messages throughout a solution:

- **Event Grid**
- **Event Hubs**
- **Service Bus**

Although they have some similarities, each service is designed for particular scenarios. This article describes the differences between these services, and helps you understand which one to choose for your application. In many cases, the messaging services are complementary and can be used together.

Service	Purpose	Type	When to use
Event Grid	Reactive programming	Event distribution (discrete)	React to status changes
Event Hubs	Big data pipeline	Event streaming (series)	Telemetry and distributed data streaming
Service Bus	High-value enterprise messaging	Message	Order processing and financial transactions

Event Grid is an eventing backplane that enables event-driven, reactive programming. It uses a publish-subscribe model. Publishers emit events, but have no expectation about which events are handled. Subscribers decide which events they want to handle. Event Grid is deeply integrated with Azure services and can be integrated with third-party services. It simplifies event consumption and lowers costs by eliminating the need for constant polling. Event Grid efficiently and reliably routes events from Azure and non-Azure resources.

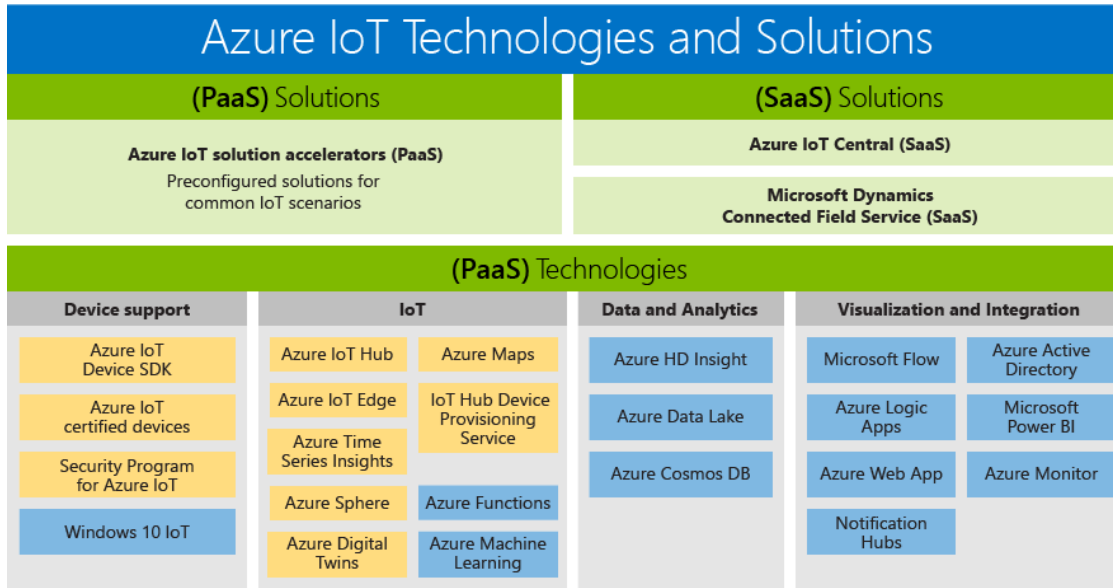
Azure Event Hubs is a big data pipeline. It facilitates the capture, retention, and replay of telemetry and event stream data. The data can come from many concurrent sources. Event Hubs allows telemetry and event data to be made available to a variety of stream-processing infrastructures and analytics services. It is available either as data streams or bundled event batches. This service provides a single solution that enables rapid data retrieval for real-time processing as well as repeated replay of stored raw data. It can capture the streaming data into a file for processing and analysis.

Service Bus is intended for traditional enterprise applications. These enterprise applications require transactions, ordering, duplicate detection, and instantaneous consistency. Service Bus enables cloud-native applications to provide reliable state transition management for business processes. When handling high-value messages that cannot be lost or duplicated, use Azure Service Bus. Service Bus also facilitates highly secure communication across hybrid cloud solutions and can connect existing on-premises systems to cloud solutions.

Service Bus is a brokered messaging system. It stores messages in a “broker” (for example, a queue) until the consuming party is ready to receive the messages.

6.6 EXPLORE AZURE IOT

There are several IoT-related services in Azure and it can be confusing to figure out which one you want to use. Some, such as IoT Central and the IoT solution accelerators, provide templates to help you create your own solution and get started quickly. You can also fully develop your own solutions using other services available – it all depends on how much help you want, and how much control. Here is a list of the services available, as well as what you may use them for.



Azure IoT solutions

The **Azure Internet of Things (IoT)** is a collection of Microsoft-managed cloud services that connect, monitor, and control billions of IoT assets. In simpler terms, an IoT solution is made up of one or more IoT devices and one or more back-end services running in the cloud that communicate with each other.

The main parts of an IoT solution are as follows: devices, back-end services, and the communications between the two. Devices are generally made up of a circuit board with sensors attached that connect to the internet. Many devices communicate via a Wi-Fi chip. Here are some examples of IoT devices:

- pressure sensors on a remote oil pump
- temperature and humidity sensors in an air-conditioning unit
- accelerometers in an elevator
- presence sensors in a room

Your device can communicate with back-end services in both directions. Here are some examples of ways that the device can communicate with the back-end solution:

- Your device may send temperature from a mobile refrigeration truck every 5 minutes to an IoT Hub.
- The back-end service can ask the device to send telemetry more frequently to help diagnose a problem.
- Your device can send alerts based on the values read from its sensors. For example, if monitoring a batch reactor in a chemical plant, you may want to send an alert when the temperatures exceeds a certain value.

- Your device can send information to a dashboard for viewing by human operators. For example, a control room in a refinery may show the temperature and pressure of each pipe, as well as the volume flowing through that pipe, allowing the operators to watch it.

Connecting devices securely and reliably is often the biggest challenge in IoT solutions. This is because IoT devices have different characteristics when compared to other clients such as browsers and mobile apps. Specifically, IoT devices:

- Are often embedded systems with no human operator (unlike a phone).
- Can be deployed in remote locations, where physical access is expensive.
- May only be reachable through the solution back end. There is no other way to interact with the device.
- May have limited power and processing resources.
- May have intermittent, slow, or expensive network connectivity.
- May need to use proprietary, custom, or industry-specific application protocols.

Here are some of the functions a back-end service can provide.

- Receiving telemetry at scale from your devices, and determining how to process and store that data.
- Analyzing the telemetry to provide insights, either in real time or after the fact.
- Sending commands from the cloud to a specific device.
- Provisioning devices and control which devices can connect to your infrastructure.
- Control the state of your devices and monitor their activities.

For example, in a predictive maintenance scenario, the cloud back end stores historical telemetry. The solution uses this data to identify potential anomalous behavior on specific pumps before they cause a real problem. Using data analytics, it can identify that the preventative solution is to send a command back to the device to take a corrective action. This process generates an automated feedback loop between the device and the cloud that greatly increases the solution efficiency.

There are several IoT-related services in Azure and it can be confusing to figure out which one you want to use. Some, such as IoT Central and the IoT solution accelerators, provide templates to help you create your own solution and get started quickly. You can also fully develop your own solutions using other services available – it all depends on how much help you want, and how much control.

Here is a list of the services available, as well as what you may use them for.

- **IoT Central:** This is a SaaS solution that helps you connect, monitor, and manage your IoT devices. To start, you select a template for your device type and create and test a basic IoT Central application that the operators of the device will use. The IoT Central application will also enable you to monitor the devices and provision new devices. This service is for straightforward solutions that don't require deep service customization.
- **IoT solution accelerators:** This is a collection of PaaS solutions you can use to accelerate your development of an IoT solution. You start with a provided IoT solution and then fully customize that solution to your requirements. You need Java or .NET skills to customize the back-end, and JavaScript skills to customize the visualization.
- **IoT Hub:** This service allows you to connect from your devices to an IoT hub, and monitor and control billions of IoT devices. This is especially useful if you need bi-directional communication between your IoT devices and your back end. This is the underlying service for IoT Central and IoT solution accelerators.

- **IoT Hub Device Provisioning Service:** This is a helper service for IoT Hub that you can use to provision devices to your IoT hub securely. With this service, you can easily provision millions of devices rapidly, rather than provisioning them one by one.
- **IoT Edge:** This service builds on top of IoT Hub. It can be used to analyze data on the IoT devices rather than in the cloud. By moving parts of your workload to the edge, fewer messages need to be sent to the cloud.
- **Azure Digital Twins:** This service enables you to create comprehensive models of the physical environment. You can model the relationships and interactions between people, spaces, and devices. For example, you can predict maintenance needs for a factory, analyze real-time energy requirements for an electrical grid, or optimize the use of available space for an office.
- **Time Series Insights:** This service enables you to store, visualize, and query large amounts of time series data generated by IoT devices. You can use this service with IoT Hub.
- **Azure Maps:** This service provides geographic information to web and mobile applications. There is a full set of REST APIs as well as a web-based JavaScript control that can be used to create flexible applications that work on desktop or mobile applications for both Apple and Windows devices.

6.6.1 IOT CENTRAL

Azure IoT Central is a fully managed IoT software-as-a-service solution that makes it easy to create products that connect the physical and digital worlds. You can bring your connected product vision to life by:

- Deriving new insights from connected devices to enable better products and experiences for your customers.
- Creating new business opportunities for your organization.

Azure IoT Central, as compared to a typical IoT project:

- Reduces the management burden.
- Reduces operational costs and overheads.
- Makes it easy to customize your application, while working with:
- Industry-leading technologies such as Azure IoT Hub and Azure Time Series Insights.
- Enterprise-grade security features such as end-to-end encryption.

6.6.2 IOT SOLUTION ACCELERATORS

Azure IoT solution accelerators are customizable PaaS solutions that provide a high level of control over your IoT solution. If your business is implementing IoT for connected operations or has specific customization requirements for connected products, Azure IoT solution accelerators provide the control you need.

Organizations with a large number of devices or device models, and manufacturers seeking connected factory solutions, are examples of companies that can benefit from IoT solution accelerators. Creating highly customizable solutions tailored to complex needs, IoT solution accelerators provide:

- Prebuilt solutions
- Remote monitoring
- Connected factory
- Predictive maintenance
- Device simulation
- Ability to deploy in minutes
- Accelerated time to value
- Solutions that give ultimate control

6.6.3 IOT HUB

IoT Hub is a managed service, hosted in the cloud, that acts as a central message hub for *bi-directional communication* between your IoT application and the devices it manages. You can use Azure IoT Hub to build IoT solutions with reliable and secure communications between millions of IoT devices and a cloud-hosted solution backend. You can connect virtually any device to IoT Hub.

IoT Hub supports communications both from the device to the cloud and from the cloud to the device. IoT Hub supports multiple messaging patterns such as device-to-cloud telemetry, file upload from devices, and request-reply methods to control your devices from the cloud. IoT Hub monitoring helps you maintain the health of your solution by tracking events such as device creation, device failures, and device connections.

IoT Hub's capabilities help you build scalable, full-featured IoT solutions such as managing industrial equipment used in manufacturing, tracking valuable assets in healthcare, and monitoring office building usage. This is the underlying service for IoT Central and IoT solution accelerators.

You can integrate IoT Hub with other Azure services to build complete, end-to-end solutions. For example, use:

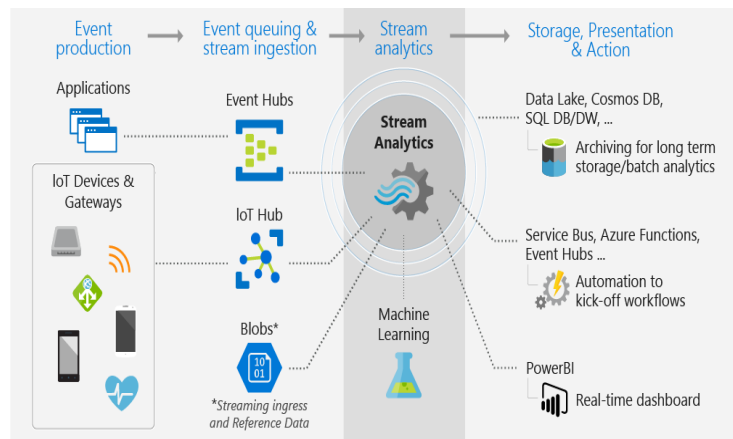
- **Azure Event Grid** to enable your business to react quickly to critical events in a reliable, scalable, and secure manner.
- **Azure Logic Apps** to automate business processes.
- **Azure Machine Learning** to add machine learning and AI models to your solution.
- **Azure Stream Analytics** to run real-time analytic computations on the data streaming from your devices.

6.6.3.1 AZURE STREAM ANALYTICS

Azure Stream Analytics is a real-time analytics and complex event-processing engine that is designed to analyze and process high volumes of fast streaming data from multiple sources simultaneously. Patterns and relationships can be identified in information extracted from a number of input sources including devices, sensors, clickstreams, social media feeds, and applications. These patterns can be used to trigger actions and initiate workflows such as creating alerts, feeding information to a reporting tool, or storing transformed data for later use. Also, Stream Analytics is available on Azure IoT Edge runtime, and supports the same exact language or syntax as cloud.

The following scenarios are examples of when you can use Azure Stream Analytics:

- Analyze real-time telemetry streams from IoT devices
- Web logs/clickstream analytics
- Geospatial analytics for fleet management and driverless vehicles
- Remote monitoring and predictive maintenance of high value assets
- Real-time analytics on Point of Sale data for inventory control and anomaly detection



An Azure Stream Analytics job consists of an input, query, and an output. Stream Analytics ingests data from Azure Event Hubs, Azure IoT Hub, or Azure Blob Storage. The query, which is based on SQL query language, can be used to easily filter, sort, aggregate, and join streaming data over a period of time. You can also extend this SQL language with JavaScript and C# user defined functions (UDFs). You can easily adjust the event ordering options and duration of time windows when performing aggregation operations through simple language constructs and/or configurations.

Each job has an output for the transformed data, and you can control what happens in response to the information you've analyzed. For example, you can:

- Send data to services such as Azure Functions, Service Bus Topics or Queues to trigger communications or custom workflows downstream.
- Send data to a Power BI dashboard for real-time dashboarding.
- Store data in other Azure storage services to train a machine learning model based on historical data or perform batch analytics.

6.6.3.2 IOT HUB TIERS

The standard tier of IoT Hub enables all features, and is required for any IoT solutions that want to make use of the bi-directional communication capabilities. The basic tier enables a subset of the features and is intended for IoT solutions that only need uni-directional communication from devices to the cloud. Both tiers offer the same security and authentication features. IoT Hub also offers a free tier that is meant for testing and evaluation. It has all the capabilities of the standard tier, but limited messaging allowances. You cannot upgrade from the free tier to either basic or standard.

FEATURE	BASIC	STANDARD / FREE
Device-to-cloud telemetry	V	V
Per-device identity	V	V
Message Routing, Event Grid Integration	V	V
HTTP, AMQP, MQTT Protocols	V	V
DPS Support	V	V
Monitoring and diagnostics	V	V
Device Streams (PREVIEW)	X	V
Cloud-to-device messaging	X	V
Device Management, Device Twin, Module Twin	X	V
IoT Edge	X	V

6.6.4 IOT HUB DEVICE PROVISIONING SERVICE

The **IoT Hub Device Provisioning Service** is a helper service for IoT Hub that enables zero-touch, just-in-time provisioning to the right IoT hub without requiring human intervention, allowing customers to provision millions of devices in a secure and scalable manner.

6.6.5 IOT EDGE

Azure IoT Edge is an Internet of Things (IoT) service that builds on top of IoT Hub. This service is meant for customers who want to analyze data on devices, or “at the edge,” instead of in the cloud. By moving parts of your workload to the edge, your devices can spend less time sending messages to the cloud and react more quickly to events.

6.6.6 AZURE DIGITAL TWINS

This service enables you to create comprehensive models of the physical environment. You can model the relationships and interactions between people, spaces, and devices. For example, you can predict maintenance needs for a factory, analyze real-time energy requirements for an electrical grid, or optimize the use of available space for an office.

6.6.7 TIME SERIES INSIGHTS

This service enables you to store, visualize, and query large amounts of time series data generated by IoT devices. You can use this service with IoT Hub.

6.6.8 AZURE MAPS

This service provides geographic information to web and mobile applications. There is a full set of REST APIs as well as a web-based JavaScript control that can be used to create flexible applications that work on desktop or mobile applications for both Apple and Windows devices.

6.7 EXPLORE AZURE AI

AI is the capability of a machine to imitate intelligent human behavior. Through AI, machines can analyze images, comprehend speech, interact in natural ways and make predictions using data.

Make AI real for your business today across the cloud and the edge:

- **Machine Learning:** Quickly and easily build, train, deploy and manage your models.
- **Knowledge Mining:** Uncover latent insights from all your content.
- **AI apps and agents:** Deliver breakthrough experiences in your apps.

6.7.1 MACHINE LEARNING SERVICE

Azure Machine Learning service is a cloud service that you use to train, deploy, automate, and manage machine learning models, all at the broad scale that the cloud provides.

Machine learning is a data science technique that allows computers to use existing data to forecast future behaviors, outcomes, and trends. By using machine learning, computers learn without being explicitly programmed. Forecasts or predictions from machine learning can make apps and devices smarter. For example, when you shop online, machine learning helps recommend other products you might want based on what you’ve bought. Or when your credit card is swiped, machine learning compares the transaction to a database of transactions and helps detect fraud. And when your robot vacuum cleaner vacuums a room, machine learning helps it decide whether the job is done.

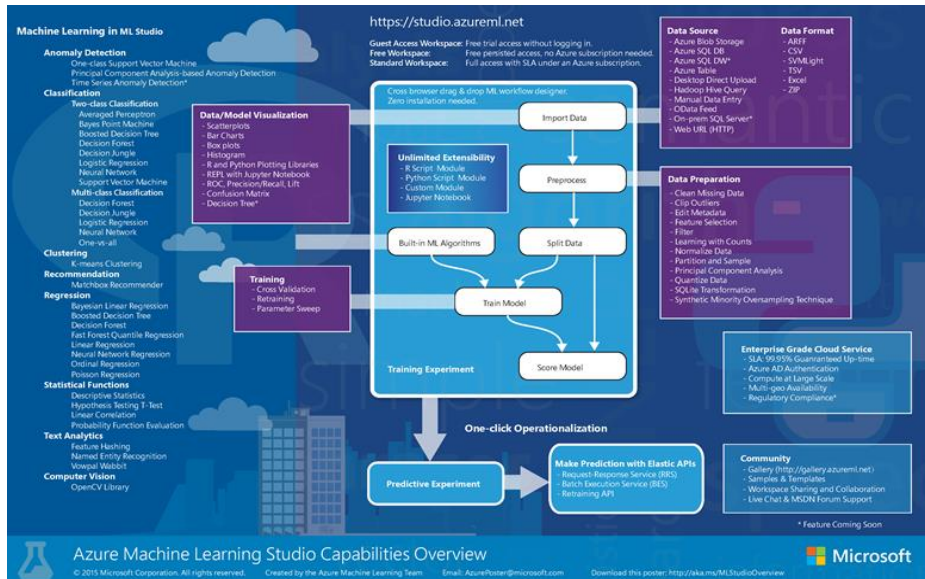
Azure Machine Learning service provides a cloud-based environment you can use to prep data, train, test, deploy, manage, and track machine learning models. Start training on your local machine and then scale out to the cloud. The service fully supports open-source technologies such as PyTorch, TensorFlow, and scikit-learn and can be used for any kind of machine learning, from classical ml to deep learning, supervised and unsupervised learning.

Explore and prepare data, train and test models, and deploy them using rich tools such as:

- A visual interface in which you can drag and drop modules to build your experiments and then deploy models
- Jupyter notebooks in which you use the SDKs to write your own code, such as these sample notebooks
- Visual Studio Code extension.

6.7.2 MACHINE LEARNING STUDIO

Azure Machine Learning Studio is a fully-managed cloud service that enables you to easily build, deploy, and share predictive analytics solutions.

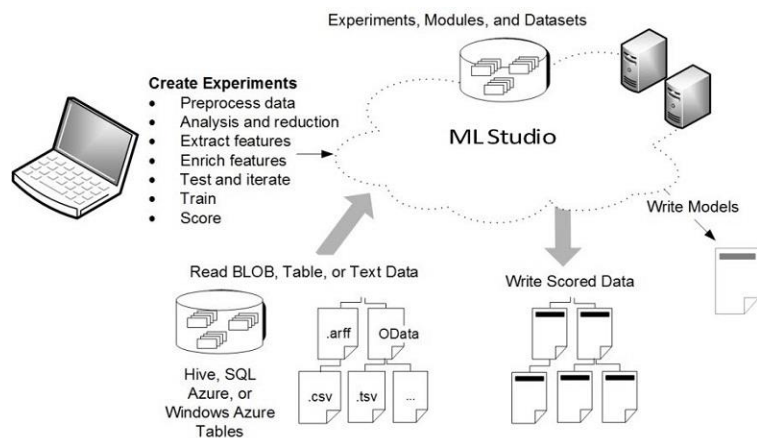


Azure Machine Learning Studio

To develop a predictive analysis model, you typically use data from one or more sources, transform, and analyze that data through various data manipulation and statistical functions, and generate a set of results. Developing a model like this is an iterative process. As you modify the various functions and their parameters, your results converge until you are satisfied that you have a trained, effective model.

Azure Machine Learning Studio gives you an interactive, visual workspace to easily build, test, and iterate on a predictive analysis model. You drag-and-drop datasets and analysis modules onto an interactive canvas, connecting them together to form an experiment, which you run in Machine Learning Studio. To iterate on your model design, you edit the experiment, save a copy if desired, and run it again. When you're ready, you can convert your training experiment to a predictive experiment, and then publish it as a web service so that your model can be accessed by others.

There is no programming required, just visually connecting datasets and modules to construct your predictive analysis model.



Azure Machine Learning Studio

6.7.3 HOW DOES AZURE MACHINE LEARNING SERVICE DIFFER FROM STUDIO?

Machine Learning Studio is a collaborative, drag-and-drop visual workspace where you can build, test, and deploy machine learning solutions without needing to write code. It uses prebuilt and preconfigured machine learning algorithms and data-handling modules as well as a proprietary compute platform.

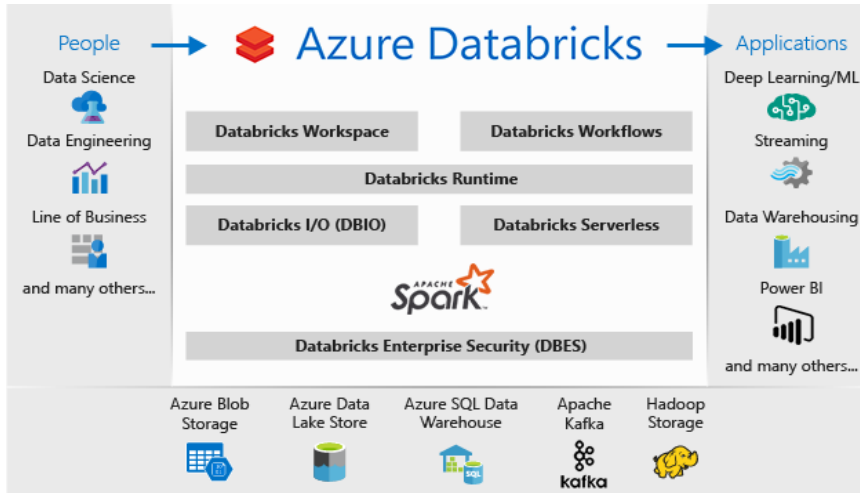
Azure Machine Learning service provides both SDKs -and- a visual interface (preview), to quickly prep data, train and deploy machine learning models. This visual interface (preview) provides a similar drag-and-drop experience to Studio. However, unlike the proprietary compute platform of Studio, the visual interface uses your own compute resources and is fully integrated into Azure Machine Learning service.

Here is a quick comparison.

Machine Learning Studio	Azure Machine Learning service: Visual interface
Generally available (GA)	In preview
Modules for interface	Many
Training compute targets	Proprietary compute target, CPU support only
Deployment compute targets	Proprietary web service format, not customizable
Automated model training and hyperparameter tuning	No

6.7.4 AZURE DATABRICKS

Azure Databricks is an Apache Spark-based analytics platform optimized for the Microsoft Azure cloud services platform. Designed with the founders of Apache Spark, Databricks is integrated with Azure to provide one-click setup, streamlined workflows, and an interactive workspace that enables collaboration between data scientists, data engineers, and business analysts.



Azure Databricks

Azure Databricks is a fast, easy, and collaborative Apache Spark-based analytics service. For a big data pipeline, the data (raw or structured) is ingested into Azure through Azure Data Factory in batches, or streamed near real-time using Kafka, Event Hub, or IoT Hub. This data lands in a data lake for long term persisted storage, in Azure Blob Storage or Azure Data Lake Storage. As part of your analytics workflow, use Azure Databricks to read data from multiple data sources such as Azure Blob Storage, Azure Data Lake Storage, Azure Cosmos DB, or Azure SQL Data Warehouse and turn it into breakthrough insights using Spark.

7 CORE CLOUD SERVICES - AZURE DATA STORAGE OPTIONS

In this module, you will:

- Survey the data storage options in Azure
- Discover how Azure data storage can meet your business demands
- Compare Azure data storage with on-premises storage

7.1 BENEFITS OF USING AZURE TO STORE DATA

7.1.1 WHY STORE YOUR DATA IN THE CLOUD?

Here are some of the important benefits of Azure data storage:

- **Automated backup and recovery:** mitigates the risk of losing your data if there is any unforeseen failure or interruption.
- **Replication across the globe:** copies your data to protect it against any planned or unplanned events, such as scheduled maintenance or hardware failures. You can choose to replicate your data at multiple locations across the globe.
- **Support for data analytics:** supports performing analytics on your data consumption.
- **Encryption capabilities:** data is encrypted to make it highly secure; you also have tight control over who can access the data.
- **Multiple data types:** Azure can store almost any type of data you need. It can handle video files, text files, and even large binary files like virtual hard disks. It also has many options for your relational and NoSQL data.
- **Data storage in virtual disks:** Azure also has the capability of storing up to 8 TB of data in its virtual disks. This is a significant capability when you're storing heavy data such as videos and simulations.
- **Storage tiers:** storage tiers to prioritize access to data based on frequently used versus rarely used information.

7.1.2 TYPES OF DATA

There are three primary types of data that Azure Storage is designed to hold.

- **Structured data.** Structured data is data that adheres to a schema, so all of the data has the same fields or properties. Structured data can be stored in a database table with rows and columns. Structured data relies on keys to indicate how one row in a table relates to data in another row of another table. Structured data is also referred to as *relational data*, as the data's schema defines the table of data, the fields in the table, and the clear relationship between the two. Structured data is straightforward in that it's easy to enter, query, and analyze. All of the data follows the same format. Examples of structured data include sensor data or financial data.
- **Semi-structured data.** Semi-structured data doesn't fit neatly into tables, rows, and columns. Instead, semi-structured data uses tags or keys that organize and provide a hierarchy for the data. Semi-structured data is also referred to as non-relational or NoSQL data.

- Unstructured data.** Unstructured data encompasses data that has no designated structure to it. This also means that there are no restrictions on the kinds of data it can hold. For example, a blob can hold a PDF document, a JPG image, a JSON file, video content, etc. As such, unstructured data is becoming more prominent as businesses try to tap into new data sources.

7.1.3 HOW AZURE DATA STORAGE CAN MEET YOUR BUSINESS STORAGE NEEDS

Azure provides several storage options that accommodate specific types of data storage needs.

7.1.3.1 AZURE SQL DATABASE


Azure SQL Database is a relational database as a service (DaaS) based on the latest stable version of the Microsoft SQL Server database engine. SQL Database is a high-performance, reliable, fully managed and secure database.

Azure Database Migration Service integrates some of the functionality of our existing tools and services. It provides customers with a comprehensive, highly available solution. The service uses the Data Migration Assistant to generate assessment reports that provide recommendations to guide you through the changes required prior to performing a migration. It's up to you to perform any remediation required. When you're ready to begin the migration process, Azure Database Migration Service performs all of the required steps. You can fire and forget your migration projects with peace of mind, knowing that the process takes advantage of best practices as determined by Microsoft.

7.1.3.2 AZURE SQL DATA WAREHOUSE

SQL Data Warehouse is a cloud-based Enterprise Data Warehouse (EDW) that uses Massively Parallel Processing (MPP) to quickly run complex queries across petabytes of data. Use SQL Data Warehouse as a key component of a big data solution. Import big data into SQL Data Warehouse with simple PolyBase T-SQL queries, and then use the power of MPP to run high-performance analytics. As you integrate and analyze, the data warehouse will become the single version of truth your business can count on for insights.

SQL Data Warehouse is a key component of an end-to-end big data solution in the Cloud.



SQL Data Warehouse

In a cloud data solution, data is ingested into big data stores from a variety of sources. Once in a big data store, Hadoop, Spark, and machine learning algorithms prepare and train the data. When the data is ready for complex analysis, SQL Data Warehouse uses PolyBase to query the big data stores. PolyBase uses standard T-SQL queries to bring the data into SQL Data Warehouse.

SQL Data Warehouse stores data into relational tables with columnar storage. This format significantly reduces the data storage costs, and improves query performance. Once data is stored in SQL Data Warehouse, you can run analytics at massive scale. Compared to traditional database systems, analysis queries finish in seconds instead of minutes, or hours instead of days.

The analysis results can go to worldwide reporting databases or applications. Business analysts can then gain insights to make well-informed business decisions.

7.1.3.3 AZURE COSMOS DB

Today's applications are required to be highly responsive and always online. To achieve low latency and high availability, instances of these applications need to be deployed in datacenters that are close to their users. Applications need to respond in real time to large changes in usage at peak hours, store ever increasing volumes of data, and make this data available to users in milliseconds.

Azure Cosmos DB is Microsoft's globally distributed, multi-model database service. With a click of a button, Cosmos DB enables you to elastically and independently scale throughput and storage across any number of Azure regions worldwide. You can elastically scale throughput and storage, and take advantage of fast, single-digit-millisecond data access using your favorite API including SQL, MongoDB, Cassandra, Tables, or Gremlin. Azure Cosmos DB is a globally distributed database service. It supports schema-less data that lets you build highly responsive and Always On applications to support constantly changing data.

7.1.3.4 AZURE STORAGE

Azure Storage is Microsoft's cloud storage solution for modern data storage scenarios. Azure Storage offers a massively scalable object store for data objects, a file system service for the cloud, a messaging store for reliable messaging, and a NoSQL store.

Azure Storage includes these data services:

- **Azure Blobs:** A massively scalable object store for text and binary data.
- **Azure Files:** Managed file shares for cloud or on-premises deployments.
- **Azure Queues:** A messaging store for reliable messaging between application components.
- **Azure Tables:** A NoSQL store for schemaless storage of structured data.

Each service is accessed through a storage account.

7.1.3.4.1 AZURE BLOB STORAGE

Azure Blob storage is Microsoft's object storage solution for the cloud. Blob storage is optimized for storing massive amounts of unstructured data. Unstructured data is data that does not adhere to a particular data model or definition, such as text or binary data.

Blob storage is designed for:

- Serving images or documents directly to a browser.
- Storing files for distributed access.
- Streaming video and audio.
- Writing to log files.
- Storing data for backup and restore, disaster recovery, and archiving.
- Storing data for analysis by an on-premises or Azure-hosted service.

Users or client applications can access objects in Blob storage via HTTP/HTTPS, from anywhere in the world. Objects in Blob storage are accessible via the Azure Storage REST API, Azure PowerShell, Azure CLI, or an Azure Storage client library. Client libraries are available for a variety of languages, including .NET, Java, Node.js, Python, Go, PHP, and Ruby.

7.1.3.4.2 AZURE DATA LAKE STORAGE GEN2

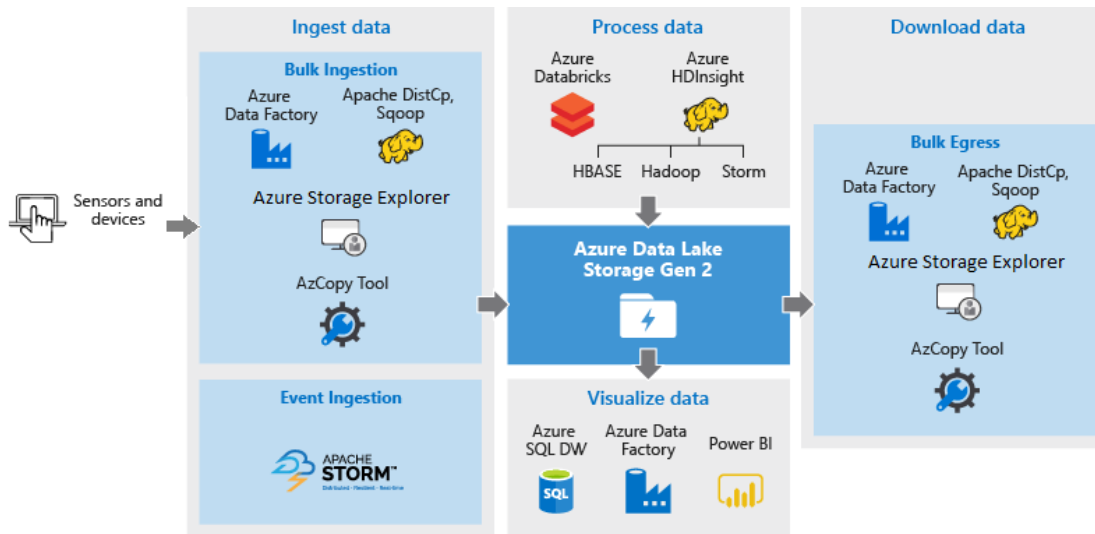
Blob storage supports **Azure Data Lake Storage Gen2**, Microsoft's enterprise *big data analytics solution* for the cloud. Data Lake Storage Gen2 makes Azure Storage the foundation for building enterprise data lakes on Azure. Designed from the start to service multiple petabytes of information while sustaining hundreds of gigabits of throughput, Data Lake Storage Gen2 allows you to easily manage massive amounts of data.

Data lakes and data warehouses are both widely used for storing big data, but they are not interchangeable terms:

- A data lake is a vast pool of raw data, the purpose for which is not yet defined.
- A data warehouse is a repository for structured, filtered data that has already been processed for a specific purpose.

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. A fundamental part of Data Lake Storage Gen2 is the addition of a *hierarchical namespace* to Blob storage. The hierarchical namespace organizes objects/files into a hierarchy of directories for efficient data access. A common object store naming convention uses slashes in the name to mimic a hierarchical directory structure. This structure becomes real with Data Lake Storage Gen2. Operations such as renaming or deleting a directory become single atomic metadata operations on the directory rather than enumerating and processing all objects that share the name prefix of the directory.

The **Data Lake feature** allows you to perform analytics on your data usage and prepare reports. Data Lake is a large repository that stores both structured and unstructured data. Azure Data Lake Storage Gen2 combines the scalability and cost benefits of object storage with the reliability and performance of the Big Data file system capabilities. It is ideal for performing analysis against large amounts of data that aren't stored in a relational way.



Big data analytics

Here's a list of tools that you can use to run data analysis jobs on data that is stored in Data Lake Storage Gen2.

Tool	Guidance
Azure HDInsight	Use Azure Data Lake Storage Gen2 with Azure HDInsight clusters
Azure Databricks	Azure Data Lake Storage Gen2

In contrast, **Azure Data Lake Analytics** is an on-demand analytics job service that simplifies big data. Instead of deploying, configuring, and tuning hardware, you write queries to transform your data and extract valuable insights. The analytics service can handle jobs of any scale instantly by setting the dial for how much power you need. You only pay for your job when it is running, making it cost-effective.

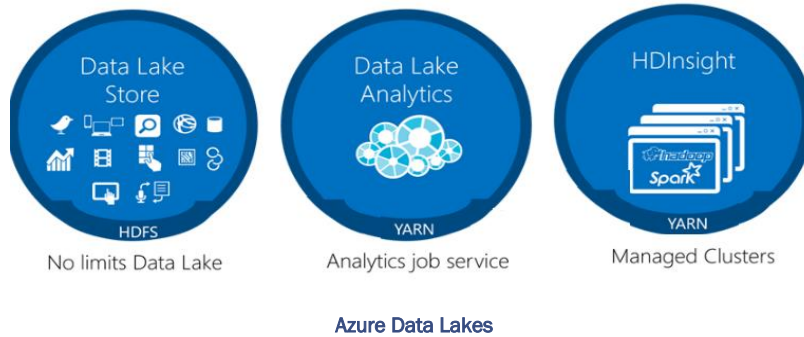
Azure HDInsight is a managed, full-spectrum, open-source analytics service for enterprises. HDInsight is a cloud service that makes it easy, fast, and cost-effective to process massive amounts of data. HDInsight also supports a broad range of scenarios, like extract, transform, and load (ETL); data warehousing; machine learning; and IoT.

The easiest way to think of Data Lake is to think of this large container that has like a real lake with rivers coming into the river you never know where the rivers are coming from (or what "type" of river). Azure Data Lake was introduced to make big data easy for developers, data scientists, and analysts to store data of any size. It removes the complexities of ingesting and storing all your data while making it faster to get up and running with big data. Data Lake is able to store the mass different types of data (Structured data, unstructured data, log files, real-time, images, etc.) and to blend that together, to correlate many different data types. The key thing here is as we are moving from traditional way to the modern tools (like Hadoop, Cassandra, NoSQL DB, etc.).

Azure Data Lake includes three services:

- **Azure Data Lake Store**, a no limits data lake that powers big data analytics
- **Azure Data Lake Analytics**, a massively parallel on-demand job service
- **Azure HDInsight**, a full managed Cloud Hadoop and Spark offering

Azure Data Lake



Azure Data Lake Store is like a cloud-based file service or file system that is pretty much unlimited in size. We can run services on top of the data that's in that store. So you could use Hadoop or Spark in an HDInsight cluster, or you could use the Azure Data Lake analytic service, which is a complement to the Azure Data Lake Store. And what that service will let you do is to run jobs that effectively query the data you have stored in the Azure Data Lake store and generate output results.

7.1.3.4.3 AZURE FILES

Azure Files enables you to set up highly available network file shares that can be accessed by using the standard Server Message Block (SMB) protocol. That means that multiple VMs can share the same files with both read and write access. You can also read the files using the REST interface or the storage client libraries.

One thing that distinguishes Azure Files from files on a corporate file share is that you can access the files from anywhere in the world using a URL that points to the file and includes a shared access signature (SAS) token. You can generate SAS tokens; they allow specific access to a private asset for a specific amount of time.

File shares can be used for many common scenarios:

- Many on-premises applications use file shares. This feature makes it easier to migrate those applications that share data to Azure. If you mount the file share to the same drive letter that the on-premises application uses, the part of your application that accesses the file share should work with minimal, if any, changes.
- Configuration files can be stored on a file share and accessed from multiple VMs. Tools and utilities used by multiple developers in a group can be stored on a file share, ensuring that everybody can find them, and that they use the same version.
- Diagnostic logs, metrics, and crash dumps are just three examples of data that can be written to a file share and processed or analyzed later.

7.1.3.4.4 AZURE QUEUE

Azure Queue storage is a service for storing large numbers of messages that can be accessed from anywhere in the world.

Azure Queue Storage can be used to help build flexible applications and separate functions for better durability across large workloads. When application components are decoupled, they can scale independently. Queue storage provides asynchronous message queuing for communication between application components, whether they are running in the cloud, on the desktop, on-premises, or on mobile devices.

You can use queue storage to:

- Create a backlog of work and to pass messages between different Azure web servers.
- Distribute load among different web servers/infrastructure and to manage bursts of traffic.
- Build resilience against component failure when multiple users access your data at the same time.

7.1.3.4.5 DISK STORAGE

An **Azure managed disk** is a **virtual hard disk (VHD)**. You can think of it like a physical disk in an on-premises server but, virtualized. Azure managed disks are stored as page blobs, which are a random IO storage object in Azure. We call a managed disk 'managed' because it is an abstraction over page blobs, blob containers, and Azure storage accounts. With managed disks, all you have to do is provision the disk, and Azure takes care of the rest.

Disk storage provides disks for virtual machines, applications, and other services to access and use as they need, similar to how they would in on-premises scenarios. Disk storage allows data to be persistently stored and accessed from an attached virtual hard disk. The disks can be managed or unmanaged by Azure, and therefore managed and configured by the user. Typical scenarios for using disk storage are if you want to lift and shift applications that read and write data to persistent disks, or if you are storing data that is not required to be accessed from outside the virtual machine to which the disk is attached.

7.1.3.4.6 AZURE TABLE

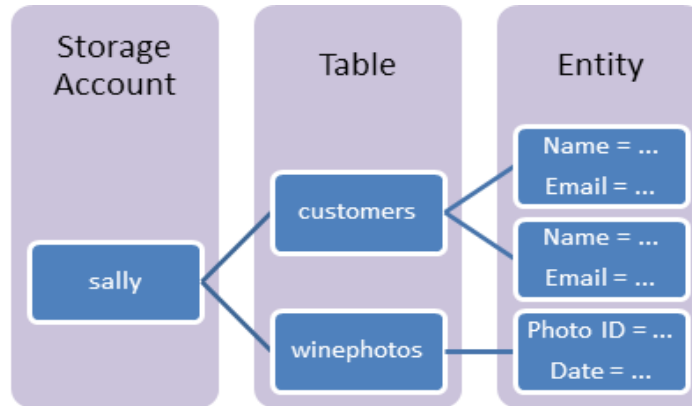
Azure Table storage is a service that stores structured NoSQL data in the cloud, providing a key/attribute store with a schemaless design. Because Table storage is schemaless, it's easy to adapt your data as the needs of your application evolve. Access to Table storage data is fast and cost-effective for many types of applications, and is typically lower in cost than traditional SQL for similar volumes of data.

You can use Table storage to store flexible datasets like user data for web applications, address books, device information, or other types of metadata your service requires. You can store any number of entities in a table, and a storage account may contain any number of tables, up to the capacity limit of the storage account.

Azure Table storage stores large amounts of structured data. The service is a NoSQL datastore which accepts authenticated calls from inside and outside the Azure cloud. Azure tables are ideal for storing structured, non-relational data. Common uses of Table storage include:

- Storing TBs of structured data capable of serving web scale applications
- Storing datasets that don't require complex joins, foreign keys, or stored procedures and can be denormalized for fast access
- Quickly querying data using a clustered index
- Accessing data using the OData protocol and LINQ queries with WCF Data Service .NET Libraries

You can use Table storage to store and query huge sets of structured, non-relational data, and your tables will scale as demand increases.



Azure Table Storage

7.1.3.4.7 DECIDING WHEN TO USE AZURE BLOBS, AZURE FILES, OR AZURE DISKS

Microsoft Azure provides several features in Azure Storage for storing and accessing your data in the cloud. This article covers Azure Files, Blobs, and Disks, and is designed to help you choose between these features.

The following table compares Files, Blobs, and Disks, and shows example scenarios appropriate for each.

Feature	Description	When to use
Azure Files	Provides an SMB interface, client libraries, and a REST interface that allows access from anywhere to stored files	<p>You want to “lift and shift” an application to the cloud which already uses the native file system APIs to share data between it and other applications running in Azure.</p> <p>You want to store development and debugging tools that need to be accessed from many virtual machines.</p>
Azure Blobs	<p>Provides client libraries and a REST interface that allows unstructured data to be stored and accessed at a massive scale in block blobs.</p> <p>Also supports Azure Data Lake Storage Gen2 for enterprise big data analytics solutions.</p>	<p>You want your application to support streaming and random access scenarios; You want to be able to access application data from anywhere.</p> <p>You want to build an enterprise data lake on Azure and perform big data analytics.</p>
Azure Disks	Provides client libraries and a REST interface that allows data to be persistently stored and accessed from an attached virtual hard disk	<p>You want to lift and shift applications that use native file system APIs to read and write data to persistent disks.</p> <p>You want to store data that is not required to be accessed from outside the virtual machine to which the disk is attached</p>

7.1.3.5 STORAGE TIERS

Azure storage offers different access tiers, which allow you to store blob object data in the most cost-effective manner. The available access tiers include:

- **Hot storage tier:** optimized for storing data that is accessed frequently.
- **Cool storage tier:** optimized for data that is infrequently accessed and stored for at least 30 days.
- **Archive storage tier:** for data that is rarely accessed and stored for at least 180 days with flexible latency requirements.

7.1.3.5.1 HOT ACCESS TIER

The hot access tier has higher storage costs than cool and archive tiers, but the lowest access costs. Example usage scenarios for the hot access tier include:

- Data that's in active use or expected to be accessed (read from and written to) frequently.
- Data that's staged for processing and eventual migration to the cool access tier.

7.1.3.5.2 COOL ACCESS TIER

The cool access tier has lower storage costs and higher access costs compared to hot storage. This tier is intended for data that will remain in the cool tier for at least 30 days. Example usage scenarios for the cool access tier include:

- Short-term backup and disaster recovery datasets.
- Older media content not viewed frequently anymore but is expected to be available immediately when accessed.
- Large data sets that need to be stored cost effectively while more data is being gathered for future processing. (For example, long-term storage of scientific data, raw telemetry data from a manufacturing facility)

7.1.3.5.3 ARCHIVE ACCESS TIER

The archive access tier has the lowest storage cost and higher data retrieval costs compared to hot and cool tiers. This tier is intended for data that can tolerate several hours of retrieval latency and will remain in the archive tier for at least 180 days.

While a blob is in archive storage, the blob data is offline and cannot be read, copied, overwritten, or modified. You can't take snapshots of a blob in archive storage. However, the blob metadata remains online and available, allowing you to list the blob and its properties. For blobs in archive, the only valid operations are GetBlobProperties, GetBlobMetadata, ListBlobs, SetBlobTier, and DeleteBlob.

Example usage scenarios for the archive access tier include:

- Long-term backup, secondary backup, and archival datasets
- Original (raw) data that must be preserved, even after it has been processed into final usable form. (For example, Raw media files after transcoding into other formats)
- Compliance and archival data that needs to be stored for a long time and is hardly ever accessed. (For example, security camera footage, old X-Rays/MRIs for healthcare organizations, audio recordings, and transcripts of customer calls for financial services)

7.1.3.6 ENCRYPTION AND REPLICATION









Azure provides security and high availability to your data through encryption and replication features. The following encryption types are available for your resources:

- **Azure Storage Service Encryption (SSE)** for data at rest helps you secure your data to meet the organization’s security and regulatory compliance. It encrypts the data before storing it and decrypts the data before retrieving it. The encryption and decryption are transparent to the user.
- **Client-side encryption** is where the data is already encrypted by the client libraries. Azure stores the data in the encrypted state at rest, which is then decrypted during retrieval.

7.1.3.7 REPLICATION FOR STORAGE AVAILABILITY

A replication type is set up when you create a storage account. The replication feature ensures that your data is durable and always available. Azure provides regional and geographic replications to protect your data against natural disasters and other local disasters like fire or flooding.

7.2 COMPARISON BETWEEN AZURE DATA STORAGE AND ON-PREMISES STORAGE

Needs	On-premise	Azure Data Storage
 Compliance and Security	1 Dedicated servers required for privacy and security	Client side encryption and encryption at rest
 Store structured and unstructured data	2 Additional IT resources with dedicated servers required	Azure Data Lake and portal analyzes and manages all types of data
 Replication and High Availability	3 More resources, licensing, and servers required	Built-in replication and redundancy features available
 Application sharing and access to shared resources	4 File sharing requires additional administration resources	File sharing options available without additional license
 Relational Data storage	5 Needs a database server with database admin role	Offers Database-as-a-Service option
 Distributed storage and data access	6 Expensive storage, networking, and compute resources needed	Azure Cosmos DB provides price-winning distributed access
 Messaging and load balancing	7 Hardware redundancy impacts budget and resources	Azure Queue provides effective load balancing
 Tiered storage	8 Management of tiered storage needs technology and labor skillset	Azure offers automated tiered storage of data

8 CORE CLOUD SERVICES - AZURE NETWORKING OPTIONS

In this module, you will:

- Learn how virtual networking helps you isolate network and compute resources
- Learn how Azure Load Balancer helps improve resiliency, or the ability to recover when your service goes down
- Learn how Traffic Manager can route traffic to different endpoints, including the endpoint with the lowest latency to the user

8.1 AZURE NETWORKING

The networking services in Azure provide a variety of networking capabilities that can be used together or separately:

- **Connectivity services:** Connect Azure resources and on-premises resources using any or a combination of these networking services in Azure:
 - Virtual Network (VNet)
 - Virtual WAN
 - ExpressRoute
 - VPN Gateway
 - Azure DNS
 - Azure Bastion
- **Application protection services:** Protect your applications using any or a combination of these networking services in Azure:
 - DDoS protection
 - Firewall
 - Network Security Groups
 - Web Application Firewall
 - Virtual Network Endpoints
- **Application delivery services:** Deliver applications in the Azure network using any or a combination of these networking services in Azure:
 - Content Delivery Network (CDN)
 - Azure Front Door Service
 - Traffic Manager
 - Application Gateway
 - Load Balancer
- **Network monitoring** – Monitor your network resources using any or a combination of these networking services in Azure:
 - Network Watcher
 - ExpressRoute Monitor
 - Azure Monitor
 - VNet Terminal Access Point (TAP)

8.1.1 CONNECTIVITY SERVICES


This section describes services that provide connectivity between Azure resources, connectivity from an on-premises network to Azure resources, and branch to branch connectivity in Azure.

Service	Why use?
Virtual network	Enables Azure resources to securely communicate with each other, the internet, and on-premises networks.
ExpressRoute	Extends your on-premises networks into the Microsoft cloud over a private connection facilitated by a connectivity provider.
VPN Gateway	Sends encrypted traffic between an Azure virtual network and an on-premises location over the public Internet
Virtual WAN	Optimizes and automates branch connectivity to, and through, Azure. Azure regions serve as hubs that you can choose to connect your branches to.
Azure DNS	Hosts DNS domains that provide name resolution by using Microsoft Azure infrastructure.
Azure Bastion (Preview)	Configure secure and seamless RDP/SSH connectivity to your virtual machines directly in the Azure portal over SSL. When you connect via Azure Bastion, your virtual machines do not need a public IP address

Azure Virtual Network (VNet) is the fundamental building block for your private network in Azure. You can use a VNets to:

- **Communicate between Azure resources:** You can deploy VMs, and several other types of Azure resources to a virtual network
- **Communicate between each other:** You can connect virtual networks to each other, enabling resources in either virtual network to communicate with each other, using virtual network peering. The virtual networks you connect can be in the same, or different, Azure regions.
- **Communicate to the internet:** All resources in a VNet can communicate outbound to the internet, by default. You can communicate inbound to a resource by assigning a public IP address or a public Load Balancer. You can also use Public IP addresses or public Load Balancer to manage your outbound connections.
- **Communicate with on-premises networks:** You can connect your on-premises computers and networks to a virtual network using VPN Gateway or ExpressRoute.

Azure Virtual Wide Area Network (WAN) is a networking service that provides optimized and automated branch connectivity to, and through, Azure. Azure regions serve as hubs that you can choose to connect your branches to. You can leverage the Azure backbone to also connect branches and enjoy branch-to-VNet connectivity.



8.1.2 APPLICATION PROTECTION SERVICES

This section describes networking services in Azure that help protect your network resources.

Service	Why use?
DDoS protection	High availability for your applications with protection from excess IP traffic charges
Web Application Firewall	Azure WAF with Application Gateway provides regional protection to entities in public and private address space Azure WAF with Front Door provides protection at the network edge to public endpoints.
Azure Firewall	Azure Firewall is a managed, cloud-based network security service that protects your Azure Virtual Network resources. It's a fully stateful firewall as a service with built-in high availability and unrestricted cloud scalability.
Network security groups	Full granular distributed end node control at VM/subnet for all network traffic flows
Virtual network service endpoints	Enables you to limit network access to some Azure service resources to a virtual network subnet

8.1.3 APPLICATION DELIVERY SERVICES

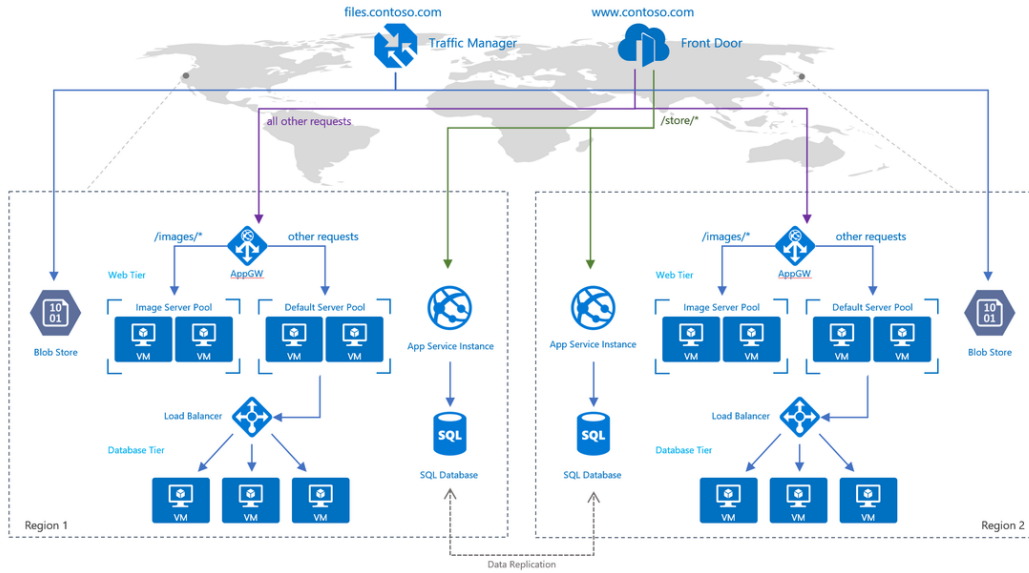
This section describes networking services in Azure that help deliver applications.

Service	Why use?
Content Delivery Network	Delivers high-bandwidth content to users. CDNs store cached content on edge servers in point-of-presence (POP) locations that are close to end users, to minimize latency
Azure Front Door Service	Enables you to define, manage, and monitor the global routing for your web traffic by optimizing for best performance and instant global failover for high availability.
Traffic Manager	Distributes traffic based on DNS to services across global Azure regions, while providing high availability and responsiveness
Load Balancer	Provides regional load-balancing by routing traffic across availability zones and into your VNets. Provides internal load-balancing by routing traffic across and between your resources to build your regional application.
Application Gateway	Azure Application Gateway is a web traffic load balancer that enables you to manage traffic to your web applications.

The services are broken into two categories.

Global load balancing services such as **Traffic Manager** and **Front Door** distribute traffic from your end users across your regional backends, across clouds or even your hybrid on-premise services. Global load balancing routes your traffic to your closest service backend and reacts to changes in service reliability or performance to maintain always-on, maximal performance for your users

Regional load balancing services such as **Standard Load Balancer** or **Application Gateway** provide the ability to distribute traffic within virtual networks (VNets) across your virtual machines (VMs) or zonal service endpoints within a region.



In summary:

- **Azure Front Door:** a service that offers a single global entry point for customers accessing web apps, APIs, content and cloud services.
- **Application Gateway:** uses Azure Load Balancer at the transport level and then applies the routing rules to support layer-7 (HTTP) load balancing.
- **Azure Traffic Manager:** load balancer for geographically distributed datacenters. Azure Traffic Manager uses DNS to redirect requests to an appropriate geographical location endpoint. Traffic Manager does not see the traffic passing between the client and the service. It simply redirects the request based on most appropriate endpoints

8.1.4 NETWORK MONITORING SERVICES

This section describes networking services in Azure that help monitor your network resources.

Service	Why use?
Network Watcher	Helps monitor and troubleshoot connectivity issues, helps diagnose VPN, NSG, and routing issues, capture packets on your VM, automates triggering diagnostics tools using Azure Functions and Logic Apps
ExpressRoute Monitor	Provides real-time monitoring of network performance, availability, and utilization, helps with auto-discovery of network topology, provides faster fault isolation, detects transient network issues, helps analyze historical network performance characteristics, supports multi-subscription
Azure Monitor	Helps you understand how your applications are performing and proactively identifies issues affecting them and the resources they depend on.

8.2 DEPLOY YOUR SITE TO AZURE

8.2.1 USING AN N-TIER ARCHITECTURE


An architectural pattern that can be used to build loosely coupled systems is N-tier. An N-tier architecture divides an application into two or more logical tiers. Architecturally, a higher tier can access services from a lower tier, but a lower tier should never access a higher tier.

Tiers help separate concerns and are ideally designed to be reusable. Using a tiered architecture also simplifies maintenance. Tiers can be updated or replaced independently, and new tiers can be inserted if needed.

Three-tier refers to an n-tier application that has three tiers. Your e-commerce web application follows this three-tier architecture:

- The web tier provides the web interface to your users through a browser.
- The application tier runs business logic.
- The data tier includes databases and other storage that hold product information and customer orders.

The following illustration shows the flow of a request from the user to the data tier.



8.2.1.1 WHAT'S AN AZURE REGION?

A region is one or more Azure data centers within a specific geographic location. East US, West US, and North Europe are examples of regions. In this instance, you see that the application is running in the East US region.

8.2.1.2 WHAT'S A VIRTUAL NETWORK?

A **virtual network** is a logically isolated network on Azure. A virtual network allows Azure resources to securely communicate with each other, the internet, and on-premises networks. A virtual network is scoped to a single region; however, multiple virtual networks from different regions can be connected together using **virtual network peering**.

Virtual networks can be segmented into one or more **subnets**. Subnets enable you to segment the virtual network into one or more sub-networks and allocate a portion of the virtual network's address space to each subnet. You can then deploy Azure resources in a specific subnet. Just like in a traditional network, subnets allow you to segment your VNet address space into segments that are appropriate for the organization's internal network. This also improves address allocation efficiency. You can secure resources within subnets using Network Security Groups.

Users interact with the web tier directly, so that VM has a public IP address along with a **private IP address**. Users don't interact with the application or data tiers, so these VMs each have a private IP address only.

You can also keep your service or data tiers in your on-premises network, placing your web tier into the cloud, but keeping tight control over other aspects of your application. A **VPN gateway** (or virtual network gateway), enables this scenario. It can provide a secure connection between an Azure Virtual Network and an on-premises location over the internet. A VPN gateway is a specific type of virtual network gateway that is used to send encrypted traffic between an Azure virtual network and an on-premises location over the public Internet. You can also use a VPN gateway to send encrypted traffic between Azure virtual networks over the Microsoft

network. Each virtual network can have only one VPN gateway. However, you can create multiple connections to the same VPN gateway. When you create multiple connections to the same VPN gateway, all VPN tunnels share the available gateway bandwidth.

Azure manages the physical hardware for you. You configure virtual networks and gateways through software, which enables you to treat a virtual network just like your own network. You choose which networks your virtual network can reach, whether that's the public internet or other networks in the private IP address space.

8.2.1.3 WHAT'S A NETWORK SECURITY GROUP?

A **network security group**, or NSG, allows or denies inbound network traffic to your Azure resources. Think of a network security group as a *cloud-level firewall* for your network.

For example, notice that the VM in the web tier allows inbound traffic on ports 22 (SSH) and 80 (HTTP). This VM's network security group allows inbound traffic over these ports from all sources. You can configure a network security group to accept traffic only from known sources, such as IP addresses that you trust.

8.3 SCALE WITH AZURE LOAD BALANCER

8.3.1 WHAT ARE AVAILABILITY AND HIGH AVAILABILITY?

Availability refers to how long your service is up and running without interruption. *High availability*, or highly available, refers to a service that's up and running for a long period of time. High Availability is the concept or goal of ensuring your critical systems are always functioning. In practice, this means creating and managing the ability to automatically "failover" to a secondary system if the primary system goes down for any reason as well as eliminating all single points of failure from your infrastructure.

Fault Tolerance describes a computer system or technology infrastructure that is designed in such a way that when one component fails (be it hardware or software), a backup component takes over operations immediately so that there is no loss of service. The concept of having backup components in place is called redundancy and the more backup components you have in place, the more tolerant your network is hardware and software failure.

The main and most important difference between high availability and fault tolerance, is actually that if an error occurs during an active action, a highly available system does not ensure the correct end state of that action, whilst a fault tolerant one, does. In other words, if, for instance, a web request is being processed by your highly available platform, and one of the nodes crashes, that user will probably get a 500 error back from the API, but the system will still be responsive for following requests. In the case of a fault-tolerant platform, the failure will somehow (more on this in a minute) be worked-around and the request will finish correctly, so the user can get a valid response. The second case will most likely take longer, due to the extra steps.

Disaster Recovery refers to the set of policies and procedures in place to ensure the continuity and recovery of mission critical systems in the event of a disruptive event such as a power outage, flood, or cyberattack. In other words, how quickly can you get your computers and systems up and running after a disastrous event? It might seem as though you don't need a disaster recovery infrastructure if your systems are configured with HA or FT. After all, if your servers can survive downtime with 99.999% or better availability, why set up a separate DR site? DR goes beyond FT or HA and consists of a complete plan to recover critical business systems and normal operations in the event of a catastrophic disaster like a major weather event (hurricane, flood, tornado, etc), a cyberattack, or any other cause of significant downtime. HA is often a major component of DR, which can also consist of an entirely separate physical infrastructure site with a 1:1 replacement for every critical infrastructure component, or at least as many as required to restore the most essential business functions.


8.3.2 WHAT IS RESILIENCY?

Resiliency refers to a system's ability to stay operational during abnormal conditions. These conditions include:

- Natural disasters
- System maintenance, both planned and unplanned, including software updates and security patches.
- Spikes in traffic to your site
- Threats made by malicious parties, such as distributed denial of service, or DDoS, attacks

8.3.3 WHAT IS A LOAD BALANCER?

A **load balancer** distributes traffic evenly among each system in a pool. A load balancer can help you achieve both high availability and resiliency. The load balancer becomes the entry point to the user. The user doesn't know (or need to know) which system the load balancer chooses to receive the request.



Load balancing enables you to run maintenance tasks without interrupting service. For example, you can **stagger** the maintenance window for each VM. During the maintenance window, the load balancer detects that the VM is unresponsive, and directs traffic to other VMs in the pool.

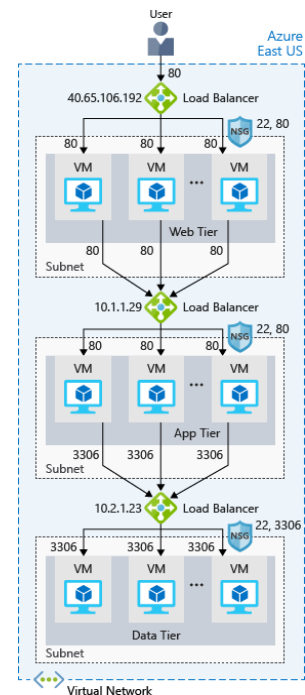
The load balancer receives the user's request and directs the request to one of the VMs in the web tier. If a VM is unavailable or stops responding, the load balancer stops sending traffic to it. The load balancer then directs traffic to one of the responsive servers.

Load balancing enables you to run maintenance tasks without interrupting service. For example, you can stagger the maintenance window for each VM. During the maintenance window, the load balancer detects that the VM is unresponsive, and directs traffic to other VMs in the pool.

8.3.4 WHAT IS AZURE LOAD BALANCER?

Azure Load Balancer is a load balancer service that Microsoft provides that helps take care of the maintenance for you. Load Balancer supports inbound and outbound scenarios, provides low latency and high throughput, and scales up to millions of flows for all Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) applications.

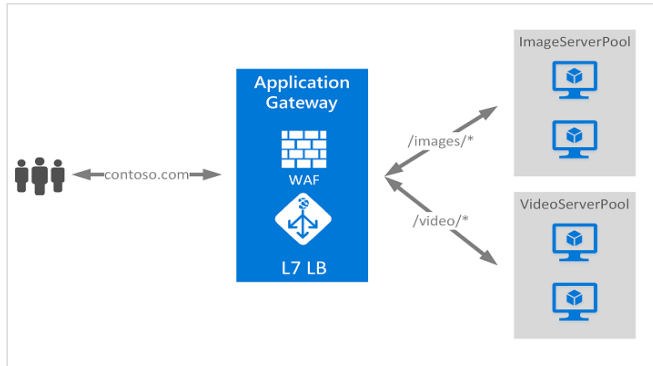
You can use Load Balancer with incoming internet traffic, internal traffic across Azure services, port forwarding for specific traffic, or outbound connectivity for VMs in your virtual network.



8.3.5 AZURE APPLICATION GATEWAY

If all your traffic is HTTP, a potentially better option is to use **Azure Application Gateway**. Application Gateway is a load balancer designed for web applications. It uses Azure Load Balancer at the transport level (TCP) and applies sophisticated URL-based routing rules to support several advanced scenarios.

With Application Gateway, you can make routing decisions based on additional attributes of an HTTP request, such as URI path or host headers. For example, you can route traffic based on the incoming URL. So if `/images` is in the incoming URL, you can route traffic to a specific set of servers (known as a pool) configured for images. If `/video` is in the URL, that traffic is routed to another pool that's optimized for videos.



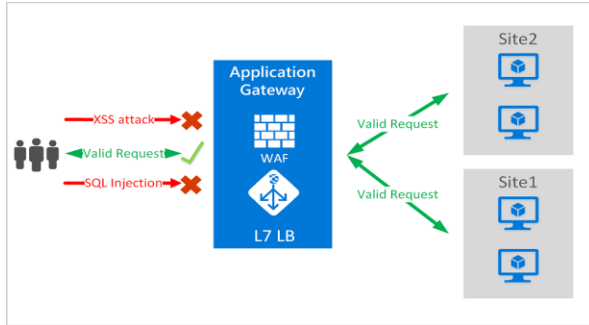
Azure application gateway

This type of routing is known as application layer (OSI layer 7) load balancing since it understands the structure of the HTTP message.

Here are some of the benefits of using Azure Application Gateway over a simple load balancer:

- **Cookie affinity.** Useful when you want to keep a user session on the same backend server.
- **SSL termination.** Application Gateway can manage your SSL certificates and pass unencrypted traffic to the backend servers to avoid encryption/decryption overhead. It also supports full end-to-end encryption for applications that require that.
- **Web application firewall.** Application gateway supports a sophisticated firewall (WAF) with detailed monitoring and logging to detect malicious attacks against your network infrastructure.
- **URL rule-based routes.** Application Gateway allows you to route traffic based on URL patterns, source IP address and port to destination IP address and port. This is helpful when setting up a content delivery network.
- **Rewrite HTTP headers.** You can add or remove information from the inbound and outbound HTTP headers of each request to enable important security scenarios, or scrub sensitive information such as server names.

Azure Application Gateway offers a **web application firewall (WAF)** that provides centralized protection of your web applications from common exploits and vulnerabilities. Web applications are increasingly targeted by malicious attacks that exploit commonly known vulnerabilities. SQL injection and cross-site scripting are among the most common attacks. A WAF solution can react to a security threat faster by centrally patching a known vulnerability, instead of securing each individual web application. Existing application gateways can easily be converted into firewall-enabled application gateways. The Application Gateway WAF is based on Core Rule Set (CRS) 3.0 or 2.2.9 from the Open Web Application Security Project (OWASP). The WAF automatically updates to include protection against new vulnerabilities, with no additional configuration needed.




Web application firewall

A **content delivery network (CDN)** is a distributed network of servers that can efficiently deliver web content to users. It is a way to get content to users in their local region to minimize latency. CDN can be hosted in Azure or any other location. You can cache content at strategically placed physical nodes across the world and provide better performance to end users. Typical usage scenarios include web applications containing multimedia content, a product launch event in a particular region, or any event where you expect a high-bandwidth requirement in a region.

8.3.6 WHAT ABOUT DNS?

DNS, or **Domain Name System**, is a way to map user-friendly names to their IP addresses. You can think of DNS as the phonebook of the internet. For example, your domain name, contoso.com, might map to the IP address of the load balancer at the web tier, 40.65.106.192.

You can bring your own DNS server or use **Azure DNS**, a hosting service for DNS domains that runs on Azure infrastructure.



8.4 REDUCE LATENCY WITH AZURE TRAFFIC MANAGER

Previously, you saw how Azure Load Balancer helps you achieve high availability and minimize downtime. Although your e-commerce site is more highly available, it doesn't solve the issue of latency or create resiliency across geographic regions. How can you make your site, which is located in the United States, load faster for users located in Europe or Asia?

8.4.1 WHAT IS NETWORK LATENCY?


Latency refers to the time it takes for data to travel over the network. Latency is typically measured in milliseconds. Compare latency to bandwidth. Bandwidth refers to the amount of data that can fit on the connection. Latency refers to the time it takes for that data to reach its destination.

Factors such as the type of connection you use and how your application is designed can affect latency. But perhaps the biggest factor is distance.

Think about your e-commerce site on Azure, which is in the East US region. It would typically take less time to transfer data to Atlanta (a distance of around 400 miles) than to transfer data to London (a distance of around 4,000 miles). Your e-commerce site delivers standard HTML, CSS, JavaScript, and images. The network latency for many files can add up. How can you reduce latency for users located far away geographically?

8.4.2 USE TRAFFIC MANAGER TO ROUTE USERS TO THE CLOSEST ENDPOINT

Recall that Azure provides data centers in regions across the globe. **Azure Traffic Manager** uses the DNS server that's closest to the user to direct user traffic to a globally distributed endpoint.



Azure Traffic Manager

Traffic Manager doesn't see the traffic that's passed between the client and server. Rather, it directs the client web browser to a preferred endpoint. Traffic Manager can route traffic in a few different ways, such as to the endpoint with the lowest latency.

8.4.3 COMPARE LOAD BALANCER TO TRAFFIC MANAGER

Azure Load Balancer distributes traffic within the same region to make your services more highly available and resilient. **Traffic Manager** works at the DNS level, and directs the client to a preferred endpoint. This endpoint can be to the region that's closest to your user.

Load Balancer and Traffic Manager both help make your services more resilient, but in slightly different ways. When Load Balancer detects an unresponsive VM, it directs traffic to other VMs in the pool. Traffic Manager monitors the health of your endpoints. In contrast, when Traffic Manager finds an unresponsive endpoint, it directs traffic to the next closest endpoint that is responsive.

9 SECURITY, RESPONSIBILITY AND TRUST IN AZURE

In this module, you will:

- Learn how security responsibility is shared with Azure
- Learn how identity management provides protection, even outside your network
- Learn how encryption capabilities built into Azure can protect your data
- Learn how to protect your network and virtual networks
- Learn about advanced services and features Azure provides to keep your services and data secure and safe

9.1 CLOUD SECURITY IS A SHARED RESPONSIBILITY

9.1.1 SHARE SECURITY RESPONSIBILITY WITH AZURE

The first shift you'll make is from on-premises data centers to **infrastructure as a service** (IaaS). With IaaS, you are leveraging the lowest-level service and asking Azure to create virtual machines (VMs) and virtual networks. At this level, it's still your responsibility to patch and secure your operating systems and software, as well as configure your network to be secure.

Moving to **platform as a service** (PaaS) outsources a lot of security concerns. At this level, Azure is taking care of the operating system and of most foundational software like database management systems. Everything is updated with the latest security patches and can be integrated with Azure Active Directory for access controls. PaaS also comes with a lot of operational advantages. Rather than building whole infrastructures and subnets for your environments by hand, you can “point and click” within the Azure portal or run automated scripts to bring complex, secured systems up and down, and scale them as needed. Contoso Shipping uses Azure Event Hubs for ingesting telemetry data from drones and trucks — as well as a web app with an Azure Cosmos DB back end with its mobile apps — which are all examples of PaaS.

With **software as a service** (SaaS), you outsource almost everything. SaaS is software that runs with an internet infrastructure. The code is controlled by the vendor but configured to be used by the customer.


Responsibility	On-prem	IaaS	PaaS	SaaS
Data governance & rights management	Customer	Customer	Customer	Customer
Client endpoints	Customer	Customer	Customer	Customer
Account & access management	Customer	Customer	Customer	Customer
Identity & directory infrastructure	Customer	Customer	Shared	Shared
Application	Customer	Customer	Shared	Microsoft
Network controls	Customer	Customer	Shared	Microsoft
Operating system	Customer	Customer	Microsoft	Microsoft
Physical hosts	Customer	Microsoft	Microsoft	Microsoft
Physical network	Customer	Microsoft	Microsoft	Microsoft
Physical datacenter	Customer	Microsoft	Microsoft	Microsoft

■ Microsoft ■ Customer

9.1.2 A LAYERED APPROACH TO SECURITY

Defense in depth is a strategy that employs a series of mechanisms to slow the advance of an attack aimed at acquiring unauthorized access to information. Each layer provides protection so that if one layer is breached, a subsequent layer is already in place to prevent further exposure. Microsoft applies a layered approach to security, both in physical data centers and across Azure services. The objective of defense in depth is to protect and prevent information from being stolen by individuals who are not authorized to access it.

Defense in depth can be visualized as a set of concentric rings, with the data to be secured at the center. Each ring adds an additional layer of security around the data. This approach removes reliance on any single layer of protection and acts to slow down an attack and provide alert telemetry that can be acted upon, either automatically or manually. Let's take a look at each of the layers.



9.1.2.1 DATA

In almost all cases, attackers are after data:

- Stored in a database
- Stored on disk inside virtual machines
- Stored on a SaaS application such as Office 365
- Stored in cloud storage

It's the responsibility of those storing and controlling access to data to ensure that it's properly secured. Often, there are regulatory requirements that dictate the controls and processes that must be in place to ensure the confidentiality, integrity, and availability of the data.

9.1.2.2 APPLICATION

- Ensure applications are secure and free of vulnerabilities.
- Store sensitive application secrets in a secure storage medium.
- Make security a design requirement for all application development.

Integrating security into the application development life cycle will help reduce the number of vulnerabilities introduced in code. We encourage all development teams to ensure their applications are secure by default, and that they're making security requirements non-negotiable.

9.1.2.3 COMPUTE

- Secure access to virtual machines.
- Implement endpoint protection and keep systems patched and current.

Malware, unpatched systems, and improperly secured systems open your environment to attacks. The focus in this layer is on making sure your compute resources are secure, and that you have the proper controls in place to minimize security issues.

9.1.2.4 NETWORKING

- Limit communication between resources.
- Deny by default.
- Restrict inbound internet access and limit outbound, where appropriate.
- Implement secure connectivity to on-premises networks.

At this layer, the focus is on limiting the network connectivity across all your resources to allow only what is required. By limiting this communication, you reduce the risk of lateral movement throughout your network.

9.1.2.5 PERIMETER

- Use distributed denial of service (DDoS) protection to filter large-scale attacks before they can cause a denial of service for end users.
- Use perimeter firewalls to identify and alert on malicious attacks against your network.

At the network perimeter, it's about protecting from network-based attacks against your resources. Identifying these attacks, eliminating their impact, and alerting you when they happen are important ways to keep your network secure.

9.1.2.6 IDENTITY AND ACCESS

- Control access to infrastructure and change control.
- Use single sign-on and multi-factor authentication.
- Audit events and changes.

The identity and access layer is all about ensuring identities are secure, access granted is only what is needed, and changes are logged.

9.1.2.7 PHYSICAL SECURITY

- Physical building security and controlling access to computing hardware within the data center is the first line of defense.

With physical security, the intent is to provide physical safeguards against access to assets. This ensures that other layers can't be bypassed, and loss or theft is handled appropriately.

9.2 GET TIPS FROM AZURE SECURITY CENTER

A great place to start when examining the security of your Azure-based solutions is **Azure Security Center**. Security Center is a monitoring service that provides threat protection across all of your services both in Azure, and on-premises. Security Center can:

- Provide security recommendations based on your configurations, resources, and networks.
- Monitor security settings across on-premises and cloud workloads, and automatically apply required security to new services as they come online.
- Continuously monitor all your services, and perform automatic security assessments to identify potential vulnerabilities before they can be exploited.
- Use machine learning to detect and block malware from being installed on your virtual machines and services. You can also define a list of allowed applications to ensure that only the apps you validate are allowed to execute.
- Analyze and identify potential inbound attacks, and help to investigate threats and any post-breach activity that might have occurred.
- Provide just-in-time access control for ports, reducing your attack surface by ensuring the network only allows traffic that you require.

9.2.1 COVERAGE

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.

Microsoft Threat Intelligence is used to identify security threats. It does this by using historical data and machine learning. This service is obtained by using the **Azure Advanced Threat Protection** service.

9.2.2 AVAILABLE TIERS

Azure Security Center is available in two tiers:

- **Free.** Available as part of your Azure subscription, this tier is limited to assessments and recommendations of Azure resources only.
- **Standard.** This tier provides a full suite of security-related services including continuous monitoring, threat detection, just-in-time access control for ports, and more.

9.3 IDENTITY AND ACCESS

Network perimeters, firewalls, and physical access controls used to be the primary protection for corporate data. But network perimeters have become increasingly porous with the explosion of bring your own device (BYOD), mobile apps, and cloud applications.

Identity has become the new primary security boundary. Therefore, proper authentication and assignment of privileges is critical to maintaining control of your data.

9.3.1 AUTHENTICATION AND AUTHORIZATION

Two fundamental concepts that need to be understood when talking about identity and access control are authentication and authorization. They underpin everything else that happens and occur sequentially in any identity and access process:

- **Authentication** is the process of establishing the identity of a person or service looking to access a resource. It involves the act of challenging a party for legitimate credentials, and provides the basis for creating a security principal for identity and access control use. It establishes if they are who they say they are.
- **Authorization** is the process of establishing what level of access an authenticated person or service has. It specifies what data they're allowed to access and what they can do with it.

Azure provides services to manage both authentication and authorization through **Azure Active Directory** (Azure AD).

9.3.2 WHAT IS AZURE ACTIVE DIRECTORY?

Azure AD is a cloud-based identity service. It has built in support for synchronizing with your existing on-premises Active Directory or can be used stand-alone. This means that all your applications, whether on-premises, in the cloud (including Office 365), or even mobile can share the same credentials. Administrators and developers can control access to internal and external data and applications using centralized rules and policies configured in Azure AD.

Azure AD provides services such as:

- **Authentication.** This includes verifying identity to access applications and resources, and providing functionality such as self-service password reset, multi-factor authentication (MFA), a custom banned password list, and smart lockout services.
- **Single-Sign-On (SSO).** SSO enables users to remember only one ID and one password to access multiple applications. A single identity is tied to a user, simplifying the security model. As users change roles or leave an organization, access modifications are tied to that identity, greatly reducing the effort needed to change or disable accounts.
- **Application management.** You can manage your cloud and on-premises apps using Azure AD Application Proxy, SSO, the My apps portal (also referred to as Access panel), and SaaS apps.
- **Business to business (B2B) identity services.** Manage your guest users and external partners while maintaining control over your own corporate data Business-to-Customer (B2C) identity services.

Customize and control how users sign up, sign in, and manage their profiles when using your apps with services.

- **Business to consumer (B2C) identity services.** Azure Active Directory (Azure AD) B2C is a business-to-consumer identity management service. This service enables you to customize and control how users securely interact with your web, desktop, mobile, or single-page applications. Using Azure AD B2C, users can sign up, sign in, reset passwords, and edit profiles. Azure AD B2C implements a form of the OpenID Connect and OAuth 2.0 protocols. The important key in the implementation of these protocols is the security tokens and their claims that enable you to provide secure access to resources.
- **Device Management.** Manage how your cloud or on-premises devices access your corporate data.

9.3.3 SINGLE SIGN-ON

The more identities a user has to manage, the greater the risk of a credential-related security incident. More identities mean more passwords to remember and change. Password policies can vary between applications and, as complexity requirements increase, it becomes increasingly difficult for users to remember them.

With **single sign-on (SSO)**, users need to remember only one ID and one password. Access across applications is granted to a single identity tied to a user, simplifying the security model. As users change roles or leave an organization, access modifications are tied to the single identity, greatly reducing the effort needed to change or disable accounts. Using single sign-on for accounts will make it easier for users to manage their identities and will increase the security capabilities in your environment.

9.3.4 MULTI-FACTOR AUTHENTICATION

Multi-factor authentication (MFA) provides additional security for your identities by requiring two or more elements for full authentication. These elements fall into three categories:

- Something you know
- Something you possess
- Something you are

Something you know would be a password or the answer to a security question. **Something you possess** could be a mobile app that receives a notification or a token-generating device. **Something you are** is typically some sort of biometric property, such as a fingerprint or face scan used on many mobile devices.

Using MFA increases security of your identity by limiting the impact of credential exposure. An attacker who has a user's password would also need to have possession of their phone or their face in order to fully authenticate. Authentication with only a single factor verified is insufficient, and the attacker would be unable to use those credentials to authenticate. The benefits this brings to security are huge, and we can't emphasize enough the importance of enabling MFA wherever possible.

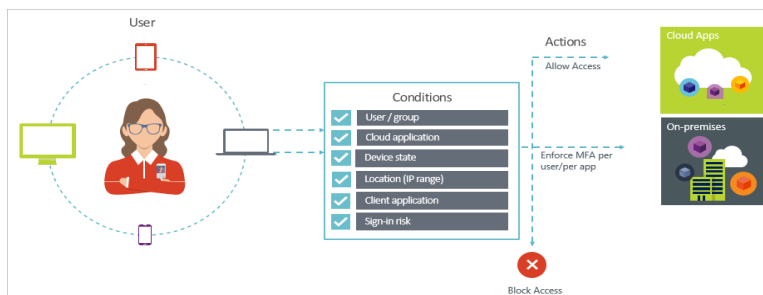
Azure AD has MFA capabilities built in and will integrate with other third-party MFA providers. It's provided free of charge to any user who has the Global Administrator role in Azure AD, because these are highly sensitive accounts. All other accounts can have MFA enabled by purchasing licenses with this capability – as well as assigning a license to the account.

Multi-Factor Authentication comes as part of the following offerings:

- Azure Active Directory Premium or Microsoft 365 Business - Full featured use of Azure Multi-Factor Authentication using Conditional Access policies to require multi-factor authentication.
- Azure AD Free or standalone Office 365 licenses - Use pre-created Conditional Access baseline protection policies to require multi-factor authentication for your users and administrators.
- Azure Active Directory Global Administrators - A subset of Azure Multi-Factor Authentication capabilities are available as a means to protect global administrator accounts.

9.3.4.1 CONDITIONAL ACCESS

Conditional Access is a capability of Azure Active Directory. With Conditional Access, you can implement automated access control decisions for accessing your cloud apps that are based on conditions.



Conditional Access

In a mobile-first, cloud-first world, Azure Active Directory enables single sign-on to devices, apps, and services from anywhere. With the proliferation of devices (including BYOD), work off corporate networks, and third-party SaaS apps, you are faced with two opposing goals:

- Empower users to be productive wherever and whenever
- Protect the corporate assets at any time

By using Conditional Access policies, you can apply the right access controls under the required conditions. Azure AD Conditional Access provides you with added security when needed and stays out of your user's way when it isn't.

Following are some common access concerns that Conditional Access can help you with:

- **Sign-in risk: Azure AD Identity Protection** detects sign-in risks. How do you restrict access if a detected sign-in risk indicates a bad actor? What if you would like to get stronger evidence that a sign-in was performed by the legitimate user? What if your doubts are strong enough to even block specific users from accessing an app?
- **Network location:** Azure AD is accessible from anywhere. What if an access attempt is performed from a network location that is not under the control of your IT department? A username and password combination might be good enough as proof of identity for access attempts from your corporate network. What if you demand a stronger proof of identity for access attempts that are initiated from other unexpected countries or regions of the world? What if you even want to block access attempts from certain locations?

- **Device management:** In Azure AD, users can access cloud apps from a broad range of devices including mobile and also personal devices. What if you demand that access attempts should only be performed with devices that are managed by your IT department? What if you even want to block certain device types from accessing cloud apps in your environment?
- **Client application:** Today, you can access many cloud apps using different app types such as web-based apps, mobile apps, or desktop apps. What if an access attempt is performed using a client app type that causes known issues? What if you require a device that is managed by your IT department for certain app types?

Note 1: Using this feature requires an Azure AD Premium P1 license.

Note 2: Azure MFA is enabled via the Conditional Access tab

9.3.4.2 AZURE ACTIVE DIRECTORY IDENTITY PROTECTION

Azure Active Directory Identity Protection enables organizations to configure automated responses to detected suspicious actions related to user identities. Azure Active Directory Identity Protection is more than a monitoring and reporting tool. To protect your organization's identities, you can configure risk-based policies that automatically respond to detected issues when a specified risk level has been reached. These policies, in addition to other Conditional Access controls provided by Azure Active Directory and Enterprise Mobility + Security (EMS), can either automatically block or initiate adaptive remediation actions including password resets and multi-factor authentication enforcement.

9.3.5 PROVIDING IDENTITIES TO SERVICES

It's usually valuable for services to have identities. Often, and against best practices, credential information is embedded in configuration files. With no security around these configuration files, anyone with access to the systems or repositories can access these credentials and risk exposure.

Azure AD addresses this problem through two methods: **service principals** and **managed identities** for Azure services.

9.3.5.1 SERVICE PRINCIPALS

To understand service principals, it's useful to first understand the words **identity** and **principal**, because of how they are used in the identity management world.

An **identity** is just a thing that can be authenticated. Obviously, this includes users with a user name and password, but it can also include applications or other servers, which might authenticate with secret keys or certificates.

A **principal** is an identity acting with certain roles or claims. Usually, it is not useful to consider identity and principal separately, but think of using `sudo` on a Bash prompt in Linux or on Windows using "run as Administrator." In both those cases, you are still logged in as the same identity as before, but you've changed the role under which you are executing. Groups are often also considered principals because they can have rights assigned.

A **service principal** is an identity that is used by a service or application. And like other identities, it can be assigned roles.

9.3.5.2 MANAGED IDENTITIES FOR AZURE SERVICES

The creation of service principals can be a tedious process, and there are a lot of touch points that can make maintaining them difficult. Managed identities for Azure services are much easier and will do most of the work for you.

A **managed identity** can be instantly created for any Azure service that supports it—and the list is constantly growing. When you create a managed identity for a service, you are creating an account on the Azure AD tenant. The Azure infrastructure will automatically take care of authenticating the service and managing the account. You can then use that account like any other Azure AD account, including securely letting the authenticated service access other Azure resources.


9.3.6 ROLE-BASED ACCESS CONTROL

Roles are sets of permissions, like “Read-only” or “Contributor”, that users can be granted to access an Azure service instance.

Identities are mapped to roles directly or through group membership. Separating security principals, access permissions, and resources provides simple access management and fine-grained control. Administrators are able to ensure the minimum necessary permissions are granted.

Roles can be granted at the individual service instance level, but they also flow down the Azure Resource Manager hierarchy.

Here's a diagram that shows this relationship. Roles assigned at a higher scope, like an entire subscription, are inherited by child scopes, like service instances.



Role Assignment Scope

9.3.6.1 PRIVILEGED IDENTITY MANAGEMENT

In addition to managing Azure resource access with role-based access control (RBAC), a comprehensive approach to infrastructure protection should consider including the ongoing auditing of role members as their organization changes and evolves. **Azure Active Directory (Azure AD) Privileged Identity Management (PIM)** is a service that enables you to manage, control, and monitor access to important resources in your organization. This includes access to resources in Azure AD, Azure resources, and other Microsoft Online Services like Office 365 or Microsoft Intune.

PIM essentially helps you manage the who, what, when, where, and why for resources that you care about. Here are some of the key features of PIM:

- Provide **just-in-time** privileged access to Azure AD and Azure resources
- Assign **time-bound access** to resources using start and end dates
- Require **approval** to activate privileged roles
- Enforce **multi-factor authentication** to activate any role
- Use **justification** to understand why users activate
- Get **notifications** when privileged roles are activated
- Conduct **access reviews** to ensure users still need roles
- Download **audit history** for internal or external audit

Using this feature requires an Azure AD Premium P2 license.

9.4 ENCRYPTION

9.4.1 WHAT IS ENCRYPTION

Encryption is the process of making data unreadable and unusable to unauthorized viewers. To use or read the encrypted data, it must be decrypted, which requires the use of a secret key. There are two top-level types of encryption: symmetric and asymmetric.

Symmetric encryption uses the same key to encrypt and decrypt the data. Consider a desktop password manager application. You enter your passwords and they are encrypted with your own personal key (your key is often derived from your master password). When the data needs to be retrieved, the same key is used, and the data is decrypted.

Asymmetric encryption uses a public key and private key pair. Either key can encrypt but a single key can't decrypt its own encrypted data. To decrypt, you need the paired key. Asymmetric encryption is used for things like Transport Layer Security (TLS) (used in HTTPS) and data signing.

Both symmetric and asymmetric encryption play a role in properly securing your data. Encryption is typically approached in two ways:

- Encryption at rest
- Encryption in transit

9.4.1.1 ENCRYPTION AT REST

Data at rest is the data that has been stored on a physical medium. This could be data stored on the disk of a server, data stored in a database, or data stored in a storage account. Regardless of the storage mechanism, encryption of data at rest ensures that the stored data is unreadable without the keys and secrets needed to decrypt it. If an attacker was to obtain a hard drive with encrypted data and did not have access to the encryption keys, the attacker would not compromise the data without great difficulty.

9.4.1.2 ENCRYPTION IN TRANSIT

Data in transit is the data actively moving from one location to another, such as across the internet or through a private network. Secure transfer can be handled by several different layers. It could be done by encrypting the data at the application layer prior to sending it over a network. HTTPS is an example of application layer in transit encryption.

You can also set up a secure channel, like a virtual private network (VPN), at a network layer, to transmit data between two systems.

Encrypting data in transit protects the data from outside observers and provides a mechanism to transmit data while limiting risk of exposure.

9.4.2 ENCRYPTION ON AZURE

9.4.2.1 ENCRYPT RAW STORAGE

Azure Storage Service Encryption (SSE) for data at rest helps you protect your data to meet your organizational security and compliance commitments. With this feature, the Azure storage platform automatically encrypts your data before persisting it to Azure Managed Disks, Azure Blob storage, Azure Files, or Azure Queue storage, and decrypts the data before retrieval. The handling of encryption, encryption at rest, decryption, and key management in Storage Service Encryption is transparent to applications using the services.

9.4.2.2 ENCRYPT VIRTUAL MACHINE DISKS

Azure Disk Encryption (ADE) is a capability that helps you encrypt your Windows and Linux IaaS virtual machine disks. Azure Disk Encryption leverages the industry-standard BitLocker feature of Windows and the dm-crypt feature of Linux to provide volume encryption for the OS and data disks. The solution is integrated with Azure Key Vault to help you control and manage the disk encryption keys and secrets (and you can use managed service identities for accessing Key Vault).

9.4.2.3 ENCRYPT DATABASES

Transparent Data Encryption (TDE) helps protect Azure SQL Database and Azure Data Warehouse against the threat of malicious activity. It performs real-time encryption and decryption of the database, associated backups, and transaction log files at rest without requiring changes to the application. By default, TDE is enabled for all newly deployed Azure SQL Database instances.

TDE encrypts the storage of an entire database by using a symmetric key called the database encryption key. By default, Azure provides a unique encryption key per logical SQL Server instance and handles all the details. Bring your own key (BYOK) is also supported with keys stored in **Azure Key Vault**.

9.4.2.4 ENCRYPT SECRETS

We've seen that the encryption services all use keys to encrypt and decrypt data, so how do we ensure that the keys themselves are secure? Corporations may also have passwords, connection strings, or other sensitive pieces of information that they need to securely store. In Azure, we can use **Azure Key Vault** to protect our secrets.

Azure Key Vault is a centralized cloud service for storing your application secrets. Key Vault helps you control your applications' secrets by keeping them in a single, central location and by providing secure access, permissions control, and access logging capabilities. It is useful for a variety of scenarios:

- *Secrets management.* You can use Key Vault to securely store and tightly control access to tokens, passwords, certificates, Application Programming Interface (API) keys, and other secrets.
- *Key management.* You also can use Key Vault as a key management solution. Key Vault makes it easier to create and control the encryption keys used to encrypt your data.
- *Certificate management.* Key Vault lets you provision, manage, and deploy your public and private Secure Sockets Layer/ Transport Layer Security (SSL/ TLS) certificates for your Azure, and internally connected, resources more easily.
- *Store secrets backed by hardware security modules (HSMs).* The secrets and keys can be protected either by software, or by FIPS 140-2 Level 2 validated HSMs.

The benefits of using Key Vault include:

- *Centralized application secrets.* Centralizing storage for application secrets allows you to control their distribution, and reduces the chances that secrets may be accidentally leaked.
- *Securely stored secrets and keys.* Azure uses industry-standard algorithms, key lengths, and HSMs, and access requires proper authentication and authorization.
- *Monitor access and use.* Using Key Vault, you can monitor and control access to company secrets.
- *Simplified administration of application secrets.* Key Vault makes it easier to enroll and renew certificates from public Certificate Authorities (CAs). You can also scale up and replicate content within regions, and use standard certificate management tools.
- *Integrate with other Azure services.* You can integrate Key Vault with storage accounts, container registries, event hubs and many more Azure services.

9.5 PROTECT YOUR NETWORK

Securing your network from attacks and unauthorized access is an important part of any architecture. Here, we'll take a look at what network security looks like, how to integrate a layered approach into your architecture, and how Azure can help you provide network security for your environment.

9.5.1 INTERNET PROTECTION


If we start on the perimeter of the network, we're focused on limiting and eliminating attacks from the internet. We suggest first assessing the resources that are internet-facing, and to only allow inbound and outbound communication where necessary. Make sure you identify all resources that are allowing inbound network traffic of any type, and then ensure they are restricted to only the ports and protocols required. Azure Security Center is a great place to look for this information, because it will identify internet-facing resources that don't have network security groups associated with them, as well as resources that are not secured behind a **firewall**.

9.5.2 WHAT IS FIREWALL?

A firewall is a service that grants server access based on the originating IP address of each request. You create firewall rules that specify ranges of IP addresses. Only clients from these granted IP addresses will be allowed to access the server. Firewall rules, generally speaking, also include specific network protocol and port information.

To provide inbound protection at the perimeter, you have several choices.

- **Azure Firewall** is a managed, cloud-based, network security service that protects your Azure Virtual Network resources. It is a fully stateful firewall as a service with built-in high availability and unrestricted cloud scalability. Azure Firewall provides inbound protection for non-HTTP/S protocols. Examples of non-HTTP/S protocols include: Remote Desktop Protocol (RDP), Secure Shell (SSH), and File Transfer Protocol (FTP). It also provides outbound, network-level protection for all ports and protocols, and application-level protection for outbound HTTP/S.



Azure Firewall

- **Azure Application Gateway** is a load balancer that includes a Web Application Firewall (WAF) that provides protection from common, known vulnerabilities in websites. It is specifically designed to protect HTTP traffic.
- **Network virtual appliances (NVAs)** are ideal options for non-HTTP services or advanced configurations, and are similar to hardware firewall appliances.

9.5.2.1 AZURE FIREWALL

You can deploy Azure Firewall on any virtual network, but customers typically deploy it on a central virtual network and peer other virtual networks to it in a **hub-and-spoke** model. You can then set the default route from the peered virtual networks to point to this central firewall virtual network. Global VNet peering is supported, but it isn't recommended because of potential performance and latency issues across regions. For best performance, deploy one firewall per region.

The advantage of this model is the ability to centrally exert control on multiple spoke VNets across different subscriptions. There are also cost savings as you don't need to deploy a firewall in each VNet separately. The cost savings should be measured versus the associate peering cost based on the customer traffic patterns.

There are three types of rule collections available in Azure Firewall:

- **Network Address Translation (NAT)** rules are used to forward traffic from the firewall to another device on the network
- **Network rules** are rules that allow traffic on specific IP address ranges and ports that you specify
- **Application rules** are used to allow applications such as Windows Update to communicate across you network. They can also be used to allow particular domain names such as azure.com and microsoft.com.

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, an all further rule processing stops.

9.5.2.2 WHAT IS THE DIFFERENCE BETWEEN APPLICATION GATEWAY WAF AND AZURE FIREWALL?

The **Web Application Firewall (WAF)** is a feature of **Application Gateway** that provides centralized inbound protection of your web applications from common exploits and vulnerabilities.

Azure Firewall provides inbound protection for non-HTTP/S protocols (for example, RDP, SSH, FTP), outbound network-level protection for all ports and protocols, and application-level protection for outbound HTTP/S.

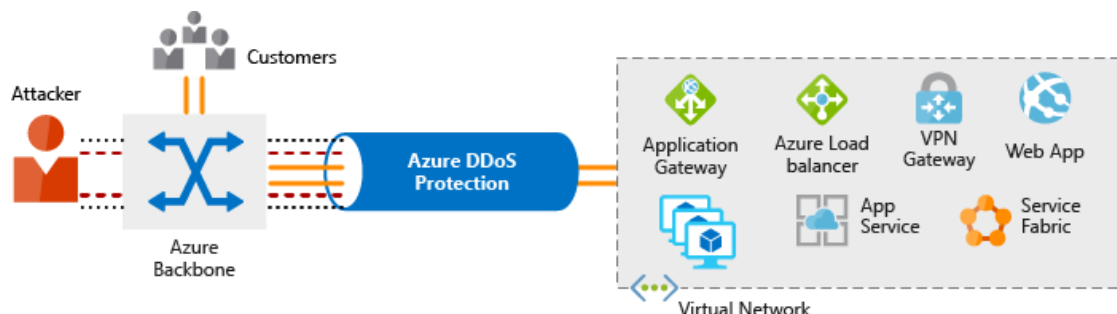
9.5.2.3 WHAT IS THE DIFFERENCE BETWEEN NSGS AND AZURE FIREWALL?

The Azure Firewall service complements network security group functionality. Together, they provide better "defense-in-depth" network security:

- Network security groups provide distributed network layer traffic filtering to limit traffic to resources within virtual networks in each subscription.
- Azure Firewall is a fully stateful, centralized network firewall as-a-service, which provides network- and application-level protection across different subscriptions and virtual networks.

9.5.3 STOPPING DISTRIBUTED DENIAL OF SERVICE (DDOS) ATTACKS

Any resource exposed on the internet is at risk of being attacked by a denial of service attack. These types of attacks attempt to overwhelm a network resource by sending so many requests that the resource becomes slow or unresponsive.



Azure DDoS Protection

When you combine **Azure DDoS Protection** with application design best practices, you help provide defense against DDoS attacks. DDoS Protection leverages the scale and elasticity of Microsoft’s global network to bring DDoS mitigation capacity to every Azure region. The Azure DDoS Protection service protects your Azure applications by scrubbing traffic at the Azure network edge before it can impact your service’s availability.

Azure DDoS protection provides the following service tiers:

Feature	DDoS Protection Basic	DDoS Protection Standard
Active traffic monitoring & always on detection	Yes	Yes
Automatic attack mitigations	Yes	Yes
Availability guarantee	Azure region	Application
Mitigation policies	Tuned for Azure region traffic volume	Tuned for application traffic volume
Metrics & alerts	No	Real time attack metrics & diagnostic logs via Azure monitor
Mitigation reports	No	Post attack mitigation reports
Mitigation flow logs	No	NRT log stream for SIEM integration
Mitigation policy customizations	No	Engage DDoS experts
Support	Best effort	Access to DDoS Experts during an active attack
SLA	Azure region	Application SLA guarantee & cost protection
Pricing	Free	Monthly & usage based

Azure DDoS Protection

DDoS Protection Standard can mitigate the following types of attacks:

- Volumetric attacks:** The attack’s goal is to flood the network layer with a substantial amount of seemingly legitimate traffic. It includes UDP floods, amplification floods, and other spoofed-packet floods. DDoS Protection Standard mitigates these potential multi-gigabyte attacks by absorbing and scrubbing them, with Azure’s global network scale, automatically.
- Protocol attacks:** These attacks render a target inaccessible, by exploiting a weakness in the layer 3 and layer 4 protocol stack. It includes, SYN flood attacks, reflection attacks, and other protocol attacks. DDoS

Protection Standard mitigates these attacks, differentiating between malicious and legitimate traffic, by interacting with the client, and blocking malicious traffic.

- **Resource (application) layer attacks:** These attacks target web application packets, to disrupt the transmission of data between hosts. The attacks include HTTP protocol violations, SQL injection, cross-site scripting, and other layer 7 attacks. Use the Azure Application Gateway web application firewall, with DDoS Protection Standard, to provide defense against these attacks.

There are also third-party web application firewall offerings available in the Azure Marketplace. DDoS Protection Standard protects resources in a virtual network including public IP addresses associated with virtual machines, load balancers, and application gateways. When coupled with the Application Gateway web application firewall, DDoS Protection Standard can provide full layer 3 to layer 7 mitigation capability.

9.5.4 CONTROLLING THE TRAFFIC INSIDE YOUR VIRTUAL NETWORK

9.5.4.1 VIRTUAL NETWORK SECURITY

Once inside a **virtual network (VNet)**, it's crucial that you limit communication between resources to only what is required. For communication between virtual machines, **Network Security Groups (NSGs)** are a critical piece to restrict unnecessary communication.

Network Security Groups allow you to filter network traffic to and from Azure resources in an Azure virtual network. An NSG can contain multiple inbound and outbound security rules that enable you to filter traffic to and from resources by source and destination IP address, port, and protocol. They provide a list of allowed and denied communication to and from network interfaces and subnets, and are fully customizable.

You can completely remove public internet access to your services by restricting access to service endpoints. With service endpoints, Azure service access can be limited to your virtual network.

9.5.4.2 NETWORK INTEGRATION

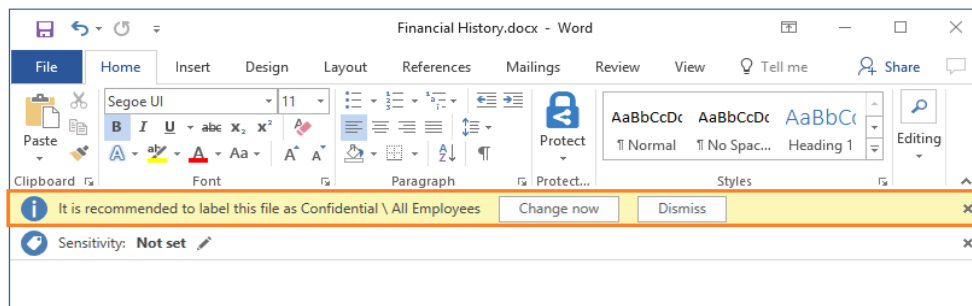
Virtual private network (VPN) connections are a common way of establishing secure communication channels between networks. Connection between Azure Virtual Network and an on-premises VPN device is a great way to provide secure communication between your network and your VNet on Azure.

To provide a dedicated, private connection between your network and Azure, you can use **Azure ExpressRoute**. ExpressRoute lets you extend your on-premises networks into the Microsoft cloud over a private connection facilitated by a connectivity provider. With ExpressRoute, you can establish connections to Microsoft cloud services, such as Microsoft Azure, Office 365, and Dynamics 365. This improves the security of your on-premises communication by sending this traffic over the private circuit instead of over the public internet. You don't need to allow access to these services for your end users over the public internet, and you can send this traffic through appliances for further traffic inspection.

9.6 PROTECT YOUR SHARED DOCUMENTS

Microsoft Azure Information Protection (MSIP or sometimes referred to as AIP) is a cloud-based solution that helps organizations classify and optionally protect documents and emails by applying labels.

Labels can be applied automatically based on rules and conditions, manually, or a combination of both where users are guided by recommendations.



Azure Information Protection

After your content is classified, you can track and control how the content is used. For example, you can:

- Analyze data flows to gain insight into your business
- Detect risky behaviors and take corrective measures
- Track access to documents
- Prevent data leakage or misuse of confidential information

The protection technology uses **Azure Rights Management** (often abbreviated to Azure RMS). This technology is integrated with other Microsoft cloud services and applications, such as Office 365 and Azure Active Directory. It can also be used with your own line-of-business applications and information protection solutions from software vendors, whether these applications and solutions are on-premises, or in the cloud.

This protection technology uses encryption, identity, and authorization policies. Similarly to the labels that are applied, protection that is applied by using Rights Management stays with the documents and emails, independently of the location—inside or outside your organization, networks, file servers, and applications. This information protection solution keeps you in control of your data, even when it is shared with other people.

9.7 AZURE ADVANCED THREAD PROTECTION

Azure Advanced Threat Protection (Azure ATP) is a cloud-based security solution that identifies, detects, and helps you investigate advanced threats, compromised identities, and malicious insider actions directed at your organization.

Azure ATP is capable of detecting known malicious attacks and techniques, security issues, and risks against your network. It enables SecOp analysts and security professionals struggling to detect advanced attacks in hybrid environments to:

- Monitor users, entity behavior, and activities with learning-based analytics
- Protect user identities and credentials stored in Active Directory
- Identify and investigate suspicious user activities and advanced attacks throughout the kill chain
- Provide clear incident information on a simple timeline for fast triage

10 APPLY AND MONITOR INFRASTRUCTURE STANDARDS WITH AZURE POLICY

Good IT governance involves planning your initiatives and setting priorities on a strategic level to help manage and prevent issues.

You need good governance when:

- You have multiple engineering teams working in Azure
- You have multiple subscriptions in your tenant
- You have regulatory requirements which must be enforced
- You want to ensure standards are followed for all IT allocated resources

You could enforce standards by not allowing teams to directly create Azure resources - and instead have the IT team define and deploy all cloud-based assets. This is often the solution in on-premises solutions, but this reduces the team agility and ability to innovate. Instead, Azure provides several tools you can use to enforce and validate your standards, while still allowing your engineering teams to create and own their own resources in the cloud.

In addition to providing IT standards, you need to be able to monitor your resources to make sure they are responsive and performing properly. Azure provides several built-in features to track and analyze your resource utilization and performance.

In this module, you will:

- Apply policies to control and audit resource creation
- Learn how role-based security can fine-tune access to your resources
- Understand Microsoft's policies and privacy guarantees
- Learn how to monitor your resources

10.1 DEFINE IT COMPLIANCE WITH AZURE POLICY

Planning out a consistent cloud infrastructure starts with setting up policy. Your policies will enforce your rules for created resources, so your infrastructure stays compliant with your corporate standards, cost requirements, and service-level agreements (SLAs) you have with your customers.

Azure Policy is a service in Azure that you use to define, assign, and, manage standards for resources in your environment. It can prevent the creation of disallowed resources, ensure new resources have specific settings applied, and run evaluations of your existing resources to scan for non-compliance.

Azure Policy comes with many built-in policy and initiative definitions that you can use, under categories such as Storage, Networking, Compute, Security Center, and Monitoring.

Imagine we allow anyone in our organization to create virtual machines (VMs). We want to control costs, so the administrator of our Azure tenant defines a policy that prohibits the creation of any VM with more than 4 CPUs. Once the policy is implemented, Azure Policy will stop anyone from creating a new VM outside the list of allowed stock keeping units (SKUs). Also, if you try to update an existing VM, it will be checked against policy. Finally, Azure Policy will audit all the existing VMs in our organization to ensure our policy is enforced. It can audit non-compliant resources, alter the resource properties, or stop the resource from being created.

10.1.1 CREATING A POLICY

The process of creating and implementing an Azure Policy begins with creating a policy definition. Every policy definition has conditions under which it is enforced. And, it has an accompanying effect that takes place if the conditions are met. To apply a policy, you will:

- Create a policy definition
- Assign a definition to a scope of resources
- View policy evaluation results

10.1.1.1 CREATE A POLICY DEFINITION

A **policy definition** expresses what to evaluate and what action to take. For example, you could ensure all public websites are secured with HTTPS, prevent a particular storage type from being created, or force a specific version of SQL Server to be used.

The policy definition itself is represented as a JSON file - you can use one of the pre-defined definitions in the portal or create your own.

Here is an example of a Compute policy that only allows specific virtual machine sizes:

```
{
  "if": {
    "allOf": [
      {
        "field": "type",
        "equals": "Microsoft.Compute/virtualMachines"
      },
      {
        "not": {
          "field": "Microsoft.Compute/virtualMachines/sku.name",
          "in": "[parameters('listOfAllowedSKUs')]"
        }
      }
    ]
  },
  "then": {
    "effect": "Deny"
  }
}
```

Notice the [parameters('listOfAllowedSKUs')] value; this is a replacement token that will be filled in when the policy definition is applied to a scope. When a parameter is defined, it's given a name and optionally given a value.

10.1.1.2 ASSIGN A DEFINITION TO A SCOPE OF RESOURCES

Once you've defined one or more policy definitions, you'll need to assign them. A **policy assignment** is a policy definition that has been assigned to take place within a specific scope.

This scope could range from a full subscription down to a resource group. Policy assignments are inherited by all child resources. This means that if a policy is applied to a resource group, it is applied to all the resources within that resource group. However, you can exclude a subscope from the policy assignment. For example, we could enforce a policy for an entire subscription and then exclude a few select resource groups.

You can assign any of these policies through the Azure portal, PowerShell, or Azure CLI. When you assign a policy definition, you will need to supply any parameters which are defined.

10.1.1.3 POLICY EFFECTS

Requests to create or update a resource through Azure Resource Manager are evaluated by Azure Policy first. Policy creates a list of all assignments that apply to the resource and then evaluates the resource against each definition. Policy processes several of the effects before handing the request to the appropriate Resource Provider to avoid any unnecessary processing if the resource violates policy.

Each policy definition in Azure Policy has a single effect. That effect determines what happens when the associated policy rule is matched. When that happens, Azure Policy will take a specific action based on the assigned effect.

Policy Effect	What happens?
Deny	The resource creation/update fails due to policy.
Disabled	The policy rule is ignored (disabled). Often used for testing.
Append	Adds additional parameters/fields to the requested resource during creation or update. A common example is adding tags on resources such as Cost Center or specifying allowed IPs for a storage resource.
Audit, AuditIfNotExists	Creates a warning event in the activity log when evaluating a non-compliant resource, but it doesn't stop the request.
DeployIfNotExists	Executes a template deployment when a specific condition is met. For example, if SQL encryption is enabled on a database, then it can run a template after the DB is created to set it up a specific way.

10.1.2 VIEW POLICY EVALUATION RESULTS

Azure Policy can allow a resource to be created even if it doesn't pass validation. In these cases, you can have it trigger an audit event which can be viewed in the Azure Policy portal, or through command-line tools. The easiest approach is in the portal as it provides a nice graphical overview which you can explore. You can find the **Azure Policy** section through the search field or All Services.

10.2 ORGANIZE POLICY WITH INITIATIVES

Managing a few policy definitions is easy, but once you have more than a few, you will want to organize them. That's where **initiatives** come in.

Initiatives work alongside policies in Azure Policy. An initiative definition is a set or group of policy definitions to help track your compliance state for a larger goal. Even if you have a single policy, we recommend using initiatives if you anticipate increasing the number of policies over time.

Like a policy assignment, an initiative assignment is an initiative definition assigned to a specific scope. Initiative assignments reduce the need to make several initiative definitions for each scope. This scope could also range from a management group to a resource group.

Once defined, initiatives can be assigned just as policies can - and they apply all the associated policy definitions.

10.2.1 DEFINING INITIATIVES

Initiative definitions simplify the process of managing and assigning policy definitions by grouping a set of policies into a single item. For example, you could create an initiative named Enable Monitoring in Azure Security Center, with a goal to monitor all the available security recommendations in your Azure Security Center.

You can define initiatives using the Azure portal, or command-line tools. In the portal, you use the “*Authoring*” section.

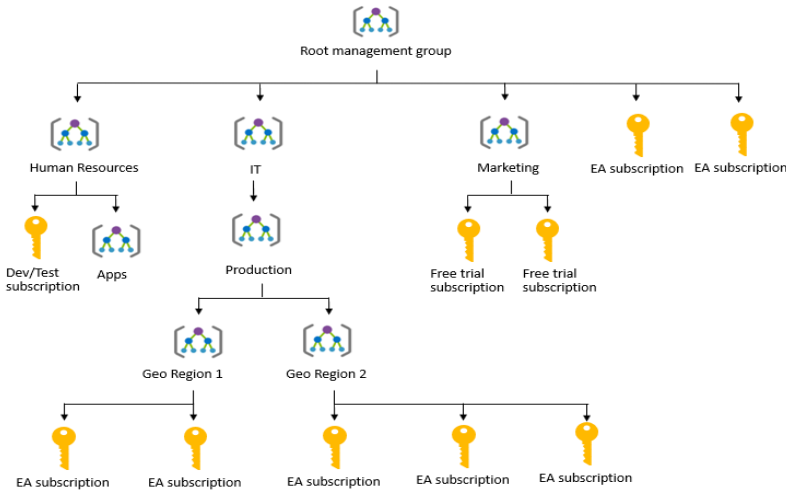
10.2.2 ASSIGNING INITIATIVES

Like a policy assignment, an initiative assignment is an initiative definition assigned to a specific scope. Initiative assignments reduce the need to make several initiative definitions for each scope. This scope could also range from a management group to a resource group.

10.3 ENTERPRISE GOVERNANCE MANAGEMENT

Access management occurs at the Azure subscription level. This allows an organization to configure each division of the company in a specific fashion based on their responsibilities and requirements. Planning and keeping rules consistent across subscriptions can be challenging without a little help.

Azure Management Groups are containers for managing access, policies, and compliance across multiple Azure subscriptions. Management groups allow you to order your Azure resources hierarchically into collections, which provide a further level of classification that is above the level of subscriptions. All subscriptions within a management group automatically inherit the conditions applied to the management group. Management groups give you enterprise-grade management at a large scale no matter what type of subscriptions you might have.



Create a hierarchy so you can apply a policy, for example, limit VM locations to US West Region on the group “Infrastructure Team management group”. This policy will inherit onto both EA subscriptions under that management group and will apply to all VMs under those subscriptions. This security policy cannot be altered by the resource or subscription owner allowing for improved governance.

Another scenario where you would use management groups is to provide user access to multi subscriptions. By moving many subscriptions under that management group, you can create one role-based access control (RBAC) assignment on the management group, which will inherit that access to all the subscriptions. One assignment on the management group can enable users to have access to everything they need instead of scripting RBAC rules over different subscriptions.

You can manage your Azure subscriptions more effectively by using Azure Policy and **Azure role-based access controls (RBACs)**. These provide distinct governance conditions that you can apply to each management group. The resources and subscriptions you assign to a management group automatically inherit the conditions that you apply to that management group.

10.4 DEFINE STANDARD RESOURCES WITH AZURE BLUEPRINTS

Adhering to security or compliance requirements, whether government or industry requirements, can be difficult and time-consuming. To help you with auditing, traceability, and compliance with your deployments, use **Azure Blueprint** artifacts and tools.

Azure Blueprint allows you to define a repeatable set of Azure resources that implement and adhere to your organization’s standards, patterns, and requirements. Blueprint enables development teams to rapidly build and deploy new environments with the knowledge that they’re building within organizational compliance with a set of built-in components that speed up development and delivery.

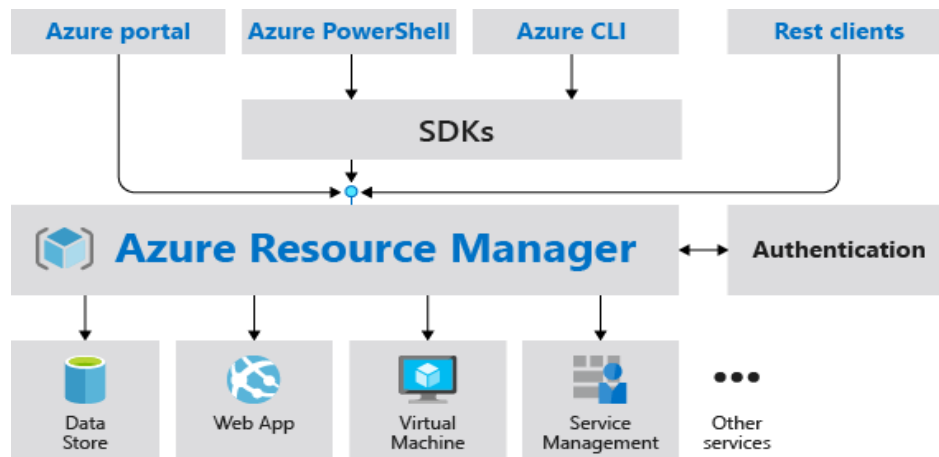
Azure Blueprint is a declarative way to orchestrate the deployment of various resource templates and other artifacts, such as:

- Role assignments
- Policy assignments
- Azure Resource Manager templates
- Resource groups

With Azure Blueprint, the relationship between the blueprint definition (what should be deployed) and the blueprint assignment (what was deployed) is preserved. This connection supports improved deployment tracking and auditing.

10.4.1 HOW IT'S DIFFERENT FROM RESOURCE MANAGER TEMPLATES

Azure Blueprints are different from **Azure Resource Manager Templates**. When Azure Resource Manager Templates deploy resources, they have no active relationship with the deployed resources (they exist in a local environment or source control). By contrast, with Azure Blueprint, each deployment is tied to an Azure Blueprint package. This means that the relationship with resources will be maintained, even after deployment. Managing relationships, in this way, improves auditing and tracking capabilities.



The service is designed to help with environment setup. This setup often consists of a set of resource groups, policies, role assignments, and Resource Manager template deployments. A blueprint is a package to bring each of these artifact types together and allow you to compose and version that package – including through a CI/CD pipeline. Ultimately, each is assigned to a subscription in a single operation that can be audited and tracked.

Nearly everything that you want to include for deployment in Blueprints can be accomplished with a Resource Manager template. However, a Resource Manager template is a document that doesn't exist natively in Azure – each is stored either locally or in source control. The template gets used for deployments of one or more Azure resources, but once those resources deploy there's no active connection or relationship to the template.

With Blueprints, the relationship between the blueprint definition (what should be deployed) and the blueprint assignment (what was deployed) is preserved. This connection supports improved tracking and auditing of deployments. Blueprints can also upgrade several subscriptions at once that are governed by the same blueprint.

There's no need to choose between a Resource Manager template and a blueprint. Each blueprint can consist of zero or more Resource Manager template artifacts. This support means that previous efforts to develop and maintain a library of Resource Manager templates are reusable in Blueprints.

10.4.2 BLUEPRINT DEFINITION

A blueprint is made up of artifacts. Blueprints currently support the following resources as artifacts:

Resource	Hierarchy options	Description
Resource Groups	Subscription	Create a new resource group for use by other artifacts within the blueprint. These placeholder resource groups enable you to organize resources exactly the way you want them structured and provides a scope limiter for included policy and role assignment artifacts and Azure Resource Manager templates.
Azure Resource Manager template	Subscription, Resource Group	Templates are used to compose complex environments. Example environments: a SharePoint farm, Azure Automation State Configuration, or a Log Analytics workspace.
Policy Assignment	Subscription, Resource Group	Allows assignment of a policy or initiative to the subscription the blueprint is assigned to. The policy or initiative must be within the scope of the blueprint definition location. If the policy or initiative has parameters, these parameters are assigned at creation of the blueprint or during blueprint assignment.
Role Assignment	Subscription, Resource Group	Add an existing user or group to a built-in role to make sure the right people always have the right access to your resources. Role assignments can be defined for the entire subscription or nested to a specific resource group included in the blueprint.

10.5 EXPLORE YOUR SERVICE COMPLIANCE WITH COMPLIANCE MANAGER

Governing your own resources and how they are used is only part of the solution when using a cloud provider. You also have to understand how the provider manages the underlying resources you are building on.

Microsoft takes this management very seriously and provides full transparency with four sources:

- Microsoft Privacy Statement
- Microsoft Trust Center
- Service Trust Portal
- Compliance Manager

10.5.1 MICROSOFT PRIVACY STATEMENT

The Microsoft privacy statement explains what personal data Microsoft processes, how Microsoft processes it, and for what purposes.

The statement applies to the interactions Microsoft has with you and Microsoft products such as Microsoft services, websites, apps, software, servers, and devices. It is intended to provide openness and honesty about how Microsoft deals with personal data in its products and services.

10.5.2 MICROSOFT TRUST CENTER

Trust Center is a website resource containing information and details about how Microsoft implements and supports security, privacy, compliance, and transparency in all Microsoft cloud products and services. The Trust Center is an important part of the Microsoft Trusted Cloud Initiative, and provides support and resources for the legal and compliance community.

10.5.3 SERVICE TRUST PORTAL

The **Service Trust Portal** (STP) hosts the Compliance Manager service, and is the Microsoft public site for publishing audit reports and other compliance-related information relevant to Microsoft's cloud services.

STP also includes information about how Microsoft online services can help your organization maintain and track compliance with standards, laws, and regulations, such as ISO, GDPR or NIST.

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.

10.5.4 COMPLIANCE MANAGER

Compliance Manager is a workflow-based risk assessment dashboard within the Trust Portal that enables you to track, assign, and verify your organization's regulatory compliance activities related to Microsoft professional services and Microsoft cloud services such as Office 365, Dynamics 365, and Azure.

Compliance Manager provides the following features:

- Combines the following three items:
 - Detailed information provided by Microsoft to auditors and regulators, as part of various third-party audits of Microsoft's cloud services against various standards (for example, ISO 27001, ISO 27018, and NIST).
 - Information that Microsoft compiles internally for its compliance with regulations (such as HIPAA and the EU GDPR).
 - An organization's self-assessment of their own compliance with these standards and regulations.
- Enables you to assign, track, and record compliance and assessment-related activities, which can help your organization cross team barriers to achieve your organization's compliance goals.
- Provides a Compliance Score to help you track your progress and prioritize auditing controls that will help reduce your organization's exposure to risk.
- Provides a secure repository in which to upload and manage evidence and other artifacts related to compliance activities.
- Produces richly detailed reports in Microsoft Excel that document the compliance activities performed by Microsoft and your organization, which can be provided to auditors, regulators, and other compliance stakeholders.

Compliance Manager is a dashboard that provides a summary of your data protection and compliance stature and recommendations for improvement. The Customer Actions provided in Compliance Manager are recommendations only; it is up to each organization to evaluate the effectiveness of these recommendations in their respective regulatory environment prior to implementation. Recommendations found in Compliance Manager should not be interpreted as a guarantee of compliance.

10.6 MONITOR YOUR SERVICE HEALTH


Defining policy and access provides fine-grained control over resources in your cloud IT infrastructure. Once those resources are deployed, you will want to know about any issues or performance problems they might encounter.

Azure provides two primary services to monitor the health of your apps and resources.

- Azure Monitor
- Azure Service Health

10.6.1 AZURE MONITOR

Azure Monitor maximizes the availability and performance of your applications by delivering a comprehensive solution for collecting, analyzing, and acting on telemetry from your cloud and on-premises environments. It helps you understand how your applications are performing and proactively identifies issues affecting them and the resources they depend on.



Azure Monitor

10.6.1.1 DATA SOURCES

Azure Monitor can collect data from a variety of sources. You can think of monitoring data for your applications in tiers ranging from your application, any operating system and services it relies on, down to the platform itself. Azure Monitor collects data from each of the following tiers:

- **Application monitoring data:** Data about the performance and functionality of the code you have written, regardless of its platform.
- **Guest OS monitoring data:** Data about the operating system on which your application is running. This could be running in Azure, another cloud, or on-premises.
- **Azure resource monitoring data:** Data about the operation of an Azure resource.
- **Azure subscription monitoring data:** Data about the operation and management of an Azure subscription, as well as data about the health and operation of Azure itself.
- **Azure tenant monitoring data:** Data about the operation of tenant-level Azure services, such as Azure Active Directory.

All data collected by Azure Monitor fits into one of two fundamental types, **metrics** and **logs**:

- **Metrics** are numerical values that describe some aspect of a system at a particular point in time. They are lightweight and capable of supporting near real-time scenarios.
- **Logs** contain different kinds of data organized into records with different sets of properties for each type. Telemetry such as events and traces are stored as logs in addition to performance data so that it can all be combined for analysis.

As soon as you create an Azure subscription and start adding resources such as virtual machines and web apps, Azure Monitor starts collecting data. **Activity Logs** record when resources are created or modified and **Metrics** tell you how the resource is performing and the resources that it's consuming.

10.6.1.2 INSIGHTS

Data monitoring is only useful if it improves your visibility of the operations in your computing environment. Azure Monitor includes several features and tools that provide valuable insights into your applications, and the other resources they may depend on.

Application Insights is a service that monitors the availability, performance, and usage of your web applications, whether they're hosted in the cloud or on-premises. It leverages the powerful data analysis platform in **Log Analytics** to provide you with deeper insights into your application's operations. Application Insights can diagnose errors, without waiting for a user to report them. Application Insights includes connection points to a variety of development tools, and integrates with Microsoft Visual Studio to support your DevOps processes.

Azure Monitor for containers is a service that is designed to monitor the performance of container workloads, which are deployed to managed Kubernetes clusters hosted on Azure Kubernetes Service (AKS). It gives you performance visibility by collecting memory and processor metrics from controllers, nodes, and containers, which are available in Kubernetes through the metrics API. Container logs are also collected.

Azure Monitor for VMs is a service that monitors your Azure VMs at scale, by analyzing the performance and health of your Windows and Linux VMs (including their different processes and interconnected dependencies on other resources, and external processes). Azure Monitor for VMs includes support for monitoring performance and application dependencies for VMs hosted on-premises, and for VMs hosted with other cloud providers.

Integrating any, or all, of these monitoring services with Azure Service Health has additional benefits. Staying informed of the health status of Azure services will help you understand if, and when, an issue affecting an Azure service is impacting your environment. What may seem like a localized problem could be the result of a more widespread issue, and **Azure Service Health** provides this kind of insight. Azure Service Health identifies any issues with Azure services that might affect your application. Azure Service Health also helps you to plan for scheduled maintenance.

10.6.1.3 RESPONDING TO ALERT CONDITIONS

In addition to allowing you to analyze your monitoring data interactively, an effective monitoring solution must respond proactively to any critical conditions that are identified within the data it collects. This might involve, for example, sending a text or email to an administrator who is responsible for investigating an issue, or launching an automated process that attempts to correct an error condition.

Alerts. Azure Monitor proactively notifies you of critical conditions using alerts, and can potentially attempt to take corrective actions. Alert rules based on metrics can provide alerts in almost real-time, based on numeric values. Alert rules based on logs allow for complex logic across data, from multiple sources.

Autoscale. Azure Monitor uses Autoscale to ensure that you have the right amount of resources running to manage the load on your application effectively. Autoscale enables you to create rules that use metrics, collected by Azure Monitor, to determine when to automatically add resources to handle increases in load. Autoscale can also help reduce your Azure costs by removing resources that are not being used. You can specify a minimum and maximum number of instances, and provide the logic that determines when Autoscale should increase or decrease resources.

10.6.2 AZURE SERVICE HEALTH

Azure Service Health is a suite of experiences that provide personalized guidance and support when issues with Azure services affect you. It can notify you, help you understand the impact of issues, and keep you updated as the issue is resolved. Azure Service Health can also help you prepare for planned maintenance and changes that could affect the availability of your resources.

Azure Service Health is composed of the following views:

- **Azure Status** provides a global view of the health state of Azure services. With Azure Status, you can get up-to-the-minute information on service availability. Everyone has access to Azure Status and can view all services that report their health state.
- **Service Health** provides you with a customizable dashboard that tracks the state of your Azure services in the regions where you use them. In this dashboard, you can track active events such as ongoing service issues, upcoming planned maintenance, or relevant Health advisories. When events become inactive, they are placed in your Health history for up to 90 days. Finally, you can use the Service Health dashboard to create and manage service Health alerts, which notify you whenever there are service issues that affect you.
- **Resource Health** helps you diagnose and obtain support when an Azure service issue affects your resources. It provides you details with about the current and past state of your resources. It also provides technical support to help you mitigate problems. In contrast to Azure Status, which informs you about service problems that affect a broad set of Azure customers, Resource Health gives you a personalized dashboard of your resources' health. Resource Health shows you times, in the past, when your resources were unavailable because of Azure service problems. It's then easier for you to understand if an SLA was violated.

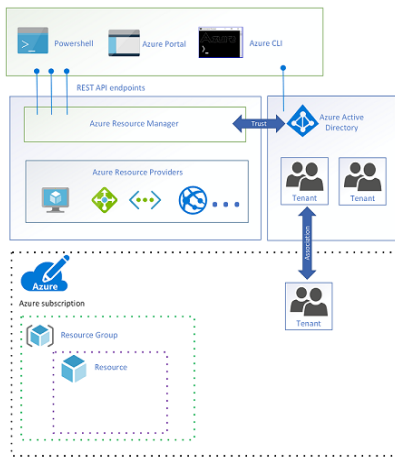
Together, the Azure Service Health components provide you with a comprehensive view of the health status of Azure, at the level of granularity that is most relevant to you.

11 CONTROL AND ORGANIZE AZURE RESOURCES WITH AZURE RESOURCE MANAGER

Azure Resource Manager has a number of features that you can use to organize resources, enforce standards, and protect critical Azure resources from accidental deletion.

In this module, you will:

- Use resource groups to organize Azure resources
- Use tags to organize resources
- Apply policies to enforce standards in your Azure environments
- Use resource locks to protect critical Azure resources from accidental deletion



Azure Resource Manager


11.1 PRINCIPLES OF RESOURCE GROUPS

11.1.1 WHAT ARE RESOURCE GROUPS?

Resource groups are a fundamental element of the Azure platform. A resource group is a logical container for resources deployed on Azure. These resources are *anything* you create in an Azure subscription like virtual machines, Application Gateways, and CosmosDB instances. All resources must be in a resource group and a resource can only be a member of a single resource group. Resources can be moved between resource groups at any time. Resource groups can't be nested. Before any resource can be provisioned, you need a resource group for it to be placed in.

11.1.1.1 LOGICAL GROUPING

Resource groups exist to help manage and organize your Azure resources. By placing resources of similar usage, type, or location, you can provide some order and organization to resources you create in Azure. Logical grouping is the aspect that we're most interested in here, since there's a lot of disorder among our resources.



Resource group

11.1.1.2 LIFE CYCLE

If you delete a resource group, all resources contained within are also deleted. Organizing resources by life cycle can be useful in non-production environments, where you might try an experiment, but then dispose of it when done. Resource groups make it easy to remove a set of resources at once.

11.1.1.3 AUTHORIZATION

Resource groups are also a scope for applying **role-based access control (RBAC) permissions**. By applying RBAC permissions to a resource group, you can ease administration and limit access to allow only what is needed.

11.1.2 CREATE A RESOURCE GROUP

Resource groups can be created by using the following methods:

- Azure portal
- Azure PowerShell
- Azure CLI
- Templates
- Azure SDKs (like .NET, Java)

11.1.3 USE RESOURCE GROUPS FOR ORGANIZATION

So how can you use resource groups to your advantage in your new organization? There are some guidelines and best practices that can help with the organization.

11.1.3.1 CONSISTENT NAMING CONVENTION

You can start with using an understandable naming convention. We named our resource group `msftlearn-core-infrastructure-rg`. We've given some indication of what it's used for (`msftlearn`), the types of resources contained within (`core-infrastructure`), and the type of resource it is itself (`rg`).

11.1.4 ORGANIZING PRINCIPLES

Resource groups can be organized in a number of ways, let's take a look at a few examples. We might put all resources that are core infrastructure into this resource group. But we could also organize them strictly by resource type. For example, put all VNets in one resource group, all virtual machines in another resource group, and all Azure Cosmos DB instances in yet another resource group.



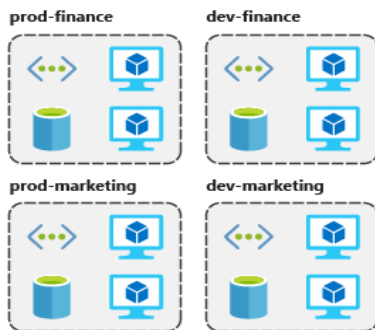
We could organize them by environment (`prod`, `qa`, `dev`). In this case, all production resources are in one resource group, all test resources are in another resource group, and so on.



We could organize them by department (marketing, finance, human resources). Marketing resources go in one resource group, finance in another resource group, and HR in a third resource group.



We could even use a combination of these strategies and organize by environment and department. Put production finance resources in one resource group, dev finance resources in another, and the same for the marketing resources.



There are a few factors that can play into the strategy you use to organize resources: authorization, resource life cycle, and billing.

11.1.4.1 ORGANIZING FOR AUTHORIZATION

Since resource groups are a scope of RBAC, you can organize resources by who needs to administer them. If your database administration team is responsible for managing all of your Azure SQL Database instances, putting them in the same resource group would simplify administration. You could give them the proper permissions at the resource group level to administer the databases within the resource group. Similarly, the database administration team could be denied access to the resource group with virtual networks, so they don't inadvertently make changes to resources outside the scope of their responsibility.

11.1.4.2 ORGANIZING FOR LIFE CYCLE

We mentioned earlier that resource groups serve as the life cycle for the resources within it. If you delete a resource group, you delete all the resources in it. Use this to your advantage, especially in areas where resources are more disposable, like non-production environments. If you deploy 10 servers for a project that you know will only last a couple of months, you might put them all in a single resource group. One resource group is easier to clean up than 10 or more resource groups.

11.1.4.3 ORGANIZING FOR BILLING

Lastly, placing resources in the same resource group is a way to group them for usage in billing reports. If you're trying to understand how your costs are distributed in your Azure environment, grouping them by resource group is one way to filter and sort the data to better understand where costs are allocated.

11.2 USE TAGGING TO ORGANIZE RESOURCES

You've gone through your resources and moved them into resource groups that are more organized than before. But what if resources have multiple uses? How do you better search, filter, and organize these resources? Tags can be helpful as you look to improve organization of your Azure resources.

11.2.1 WHAT ARE TAGS?

Tags are name/value pairs of text data that you can apply to resources and resource groups. Tags allow you to associate custom details about your resource, in addition to the standard Azure properties a resource has:

- department (like finance, marketing, and more)
- environment (prod, test, dev),
- cost center
- life cycle and automation (like shutdown and startup of virtual machines).

Tags can be added and manipulated through the Azure portal, Azure CLI, Azure PowerShell, Resource Manager templates, and through the REST API. For example, to add a resource tag to a virtual network using the Azure CLI, you could use the following command:

```
az resource tag --tags Department=Finance \
  --resource-group msftlearn-core-infrastructure-rg \
  --name msftlearn-vnet1 \
  --resource-type "Microsoft.Network/virtualNetworks"
```

You can use **Azure Policy** to automatically add or enforce tags for resources your organization creates based on policy conditions that you define. For example, you could require that a value for the Department tag is entered when someone in your organization creates a virtual network in a specific resource group.

11.2.2 USE TAGS FOR ORGANIZATION

The above example is just one example of where you can use tags to organize your resources. With their flexibility, there are several ways you can use tags to your advantage.

You can use tags to group your billing data. For example, if you're running multiple VMs for different organizations, use the tags to group usage by cost center. You can also use tags to categorize costs by runtime environment, such as the billing usage for VMs running in the production environment. When exporting billing data or accessing it through billing APIs, tags are included in that data and can be used to further slice your data from a cost perspective.

You can retrieve all the resources in your subscription with a specific tag name or value. Tags enable you to retrieve related resources from different resource groups. This approach is helpful when you need to organize resources for **billing or management**.

Tagging resources can also help in **monitoring** to track down impacted resources. Monitoring systems could include tag data with alerts, giving you the ability to know exactly who is impacted. In our example above, we applied the Department:Finance tag to the msftlearn-vnet1 resource. If an alarm was thrown on msftlearn-vnet1 and the alarm included the tag, we'd know that the finance department may be impacted by the condition that triggered the alarm. This contextual information can be valuable if an issue occurs.

It's also common for tags to be used in **automation**. If you want to automate the shutdown and startup of virtual machines in development environments during off-hours to save costs, you can use tags to assist in this. Add a shutdown:6PM and startup:7AM tag to the virtual machines, then create an automation job that looks for these tags, and shuts them down or starts them up based on the tag value. There are several solutions in the Azure Automation Runbooks Gallery that use tags in a similar manner to accomplish this.

11.3 USE POLICIES TO ENFORCE STANDARDS

You're organizing your resources better in resource groups, and you've applied tags to your resources to use them in billing reports and in your monitoring solution. Resource grouping and tagging have made a difference in the existing resources, but how do you ensure that new resources follow the rules? Let's take a look at how policies can help you enforce standards in your Azure environment.

11.3.1 WHAT IS AZURE POLICY?

Governance validates that your organization can achieve its goals through effective and efficient use of IT. It meets this need by creating clarity between business goals and IT projects.

Does your company experience a significant number of IT issues that never seem to get resolved? Good IT governance involves planning your initiatives and setting priorities on a strategic level to help manage and prevent issues. This strategic need is where Azure Policy comes in.

Azure Policy is a service you can use to create, assign, and manage policies. These policies apply and enforce rules that your resources need to follow. These policies can enforce these rules when resources are created, and can be evaluated against existing resources to give visibility into compliance.

Policies can enforce things such as only allowing specific types of resources to be created, or only allowing resources in specific Azure regions. You can enforce naming conventions across your Azure environment. You can also enforce that specific tags are applied to resources. Let's take a look at how policies work.

11.3.2 POLICIES TO ENFORCE STANDARDS

We've seen how we could use policies to ensure that our resources have the tags that organize our resources. There are other ways policies can be used to our benefit.

We could use policy to restrict which Azure regions we can deploy resources to. For organizations that are heavily regulated or have legal or regulatory restrictions on where data can reside, policies help to ensure that resources aren't provisioned in geographic areas that would go against these requirements.

We could use policy to restrict which types of virtual machine sizes can be deployed. You may want to allow large VM sizes in your production subscriptions, but maybe you'd like to ensure that you keep costs minimized in your dev subscriptions. By denying the large VM sizes through policy in your dev subscriptions, you can ensure they don't get deployed in these environments.

We could also use policy to enforce naming conventions. If our organization has standardized on specific naming conventions, using policy to enforce the conventions helps us to keep a consistent naming standard across our Azure resources.

11.4 SECURE RESOURCES WITH ROLE-BASED ACCESS CONTROL

Implementing Azure Policy ensured that all our employees with Azure access are following our internal standards for creating resources, but we have a second issue we need to solve: how do we protect those resources once they are deployed? We have IT personnel that need to manage settings, developers that need to have read-only access, and administrators that need to be able to control them completely. Enter **Role-Based Access Control (RBAC)**.

RBAC provides fine-grained access management for Azure resources, enabling you to grant users the specific rights they need to perform their jobs. RBAC is considered a core service and is included with all subscription levels at no cost.

Using RBAC, you can:

- Allow one user to manage VMs in a subscription, and another user to manage virtual networks.
- Allow a database administrator (DBA) group to manage SQL databases in a subscription.
- Allow a user to manage all resources in a resource group, such as VMs, websites, and virtual subnets.
- Allow an application to access all resources in a resource group.

RBAC uses an *allow model* for access. When you are assigned to a role, RBAC allows you to perform specific actions, such as read, write, or delete. Therefore, if one role assignment grants you read permissions to a resource group, and a different role assignment grants you write permissions to the same resource group, you will have write permissions on that resource group.

11.4.1 HOW RBAC WORKS

The way you control access to resources using RBAC is to create role assignments. This is a key concept to understand – it's how permissions are enforced. A role assignment consists of three elements: security principal, role definition, and scope.

11.4.1.1 SECURITY PRINCIPAL

A security principal is an object that represents a user, group, service principal, or managed identity that is requesting access to Azure resources.

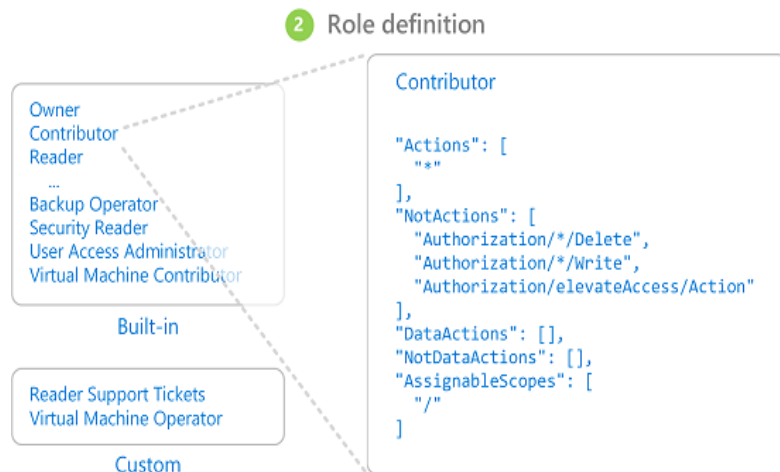


Security principal

- **User** - An individual who has a profile in Azure Active Directory. You can also assign roles to users in other tenants. For information about users in other organizations, see Azure Active Directory B2B.
- **Group** - A set of users created in Azure Active Directory. When you assign a role to a group, all users within that group have that role.
- **Service principal** - A security identity used by applications or services to access specific Azure resources. You can think of it as a user identity (username and password - or certificate) for an application.
- **Managed identity** - An identity in Azure Active Directory that is automatically managed by Azure. You typically use managed identities when developing cloud applications to manage the credentials for authenticating to Azure services.

11.4.1.2 ROLE DEFINITION

A role definition is a collection of permissions. It's sometimes just called a role. A role definition lists the operations that can be performed, such as read, write, and delete. Roles can be high-level, like owner, or specific, like virtual machine reader.

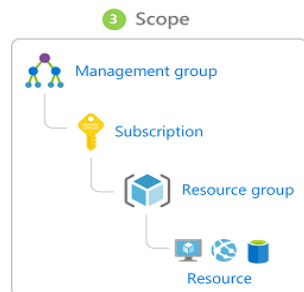


Role definition

11.4.1.3 SCOPE

Scope is the set of resources that the access applies to. When you assign a role, you can further limit the actions allowed by defining a scope. This is helpful if you want to make someone a Website Contributor, but only for one resource group.

In Azure, you can specify a scope at multiple levels: management group, subscription, resource group, or resource. Scopes are structured in a parent-child relationship.



RBAC Scope

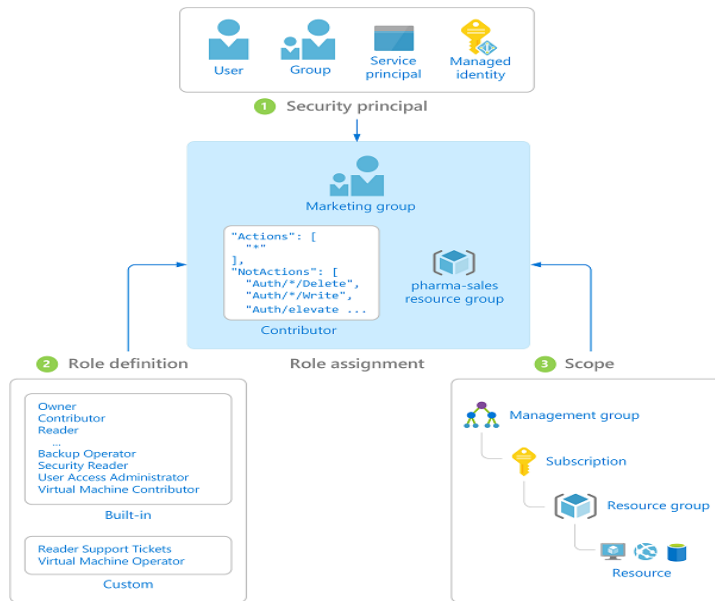
When you grant access at a parent scope, those permissions are inherited to the child scopes. For example:

- If you assign the Owner role to a user at the management group scope, that user can manage everything in all subscriptions in the management group.
- If you assign the Reader role to a group at the subscription scope, the members of that group can view every resource group and resource in the subscription.
- If you assign the Contributor role to an application at the resource group scope, it can manage resources of all types in that resource group, but not other resource groups in the subscription.

11.4.1.4 ROLE ASSIGNMENTS

A role assignment is the process of attaching a role definition to a user, group, service principal, or managed identity at a particular scope for the purpose of granting access. Access is granted by creating a role assignment, and access is revoked by removing a role assignment.

The following diagram shows an example of a role assignment. In this example, the Marketing group has been assigned the Contributor role for the pharma-sales resource group. This means that users in the Marketing group can create or manage any Azure resource in the pharma-sales resource group. Marketing users do not have access to resources outside the pharma-sales resource group, unless they are part of another role assignment.



Role assignment

11.4.2 BEST PRACTICES FOR RBAC

Here are some best practices you should use when setting up resources.

- Segregate duties within your team and grant only the amount of access to users that they need to perform their jobs. Instead of giving everybody unrestricted permissions in your Azure subscription or resources, allow only specific actions at a particular scope.
- When planning your access control strategy, grant users the lowest privilege level that they need to do their work.
- Use **Resource Locks** to ensure critical resources aren't modified or deleted (more on that next!)

Azure includes several built-in roles that you can use. The following lists four fundamental built-in roles. The first three apply to all resource types.

- **Owner** - Has full access to all resources including the right to delegate access to others.
- **Contributor** - Can create and manage all types of Azure resources but can't grant access to others.
- **Reader** - Can view existing Azure resources.
- **User Access Administrator** - Lets you manage user access to Azure resources.

The rest of the built-in roles allow management of specific Azure resources. For example, the Virtual Machine Contributor role allows a user to create and manage virtual machines. If the built-in roles don't meet the specific needs of your organization, you can create your own custom roles for Azure resources.

Azure has data operations that enable you to grant access to data within an object. For example, if a user has read data access to a storage account, then they can read the blobs or messages within that storage account. For more information, see [Understand role definitions for Azure resources](#).

11.5 USE RESOURCE LOCKS TO PROTECT RESOURCES

Resource locks are a setting that can be applied to any resource to block modification or deletion. Resource locks can be set to either Delete or Read-only:

- **Delete** will allow all operations against the resource but block the ability to delete it.
- **Read-only** will only allow read activities to be performed against it, blocking any modification or deletion of the resource.

Resource locks can be applied to subscriptions, resource groups, and to individual resources, and are inherited when applied at higher levels.

When you apply a lock at a parent scope, all resources within that scope inherit the same lock. Even resources you add later inherit the lock from the parent. The most restrictive lock in the inheritance takes precedence.

Unlike role-based access control, you use management locks to apply a restriction across all users and roles. In order to create a lock, you must either be in the 'Owner' or 'User Access Administrator' role in RBAC. Alternatively, an administrator can create a custom role that grants the right to create a lock.

Applying ReadOnly can lead to unexpected results because some operations that don't seem to modify the resource actually require actions that are blocked by the lock. The ReadOnly lock can be applied to the resource or to the resource group containing the resource. Some common examples of the operations that are blocked by a ReadOnly lock are:

- A ReadOnly lock on a storage account prevents all users from listing the keys. The list keys operation is handled through a POST request because the returned keys are available for write operations.
- A ReadOnly lock on an App Service resource prevents Visual Studio Server Explorer from displaying files for the resource because that interaction requires write access.
- A ReadOnly lock on a resource group that contains a virtual machine prevents all users from starting or restarting the virtual machine. These operations require a POST request.

12 PREDICT COST AND OPTIMIZE SPENDING FOR AZURE

When planning a solution in the cloud, there's always the challenge of balancing cost against performance. It can feel like a guessing game whether your selected options will stay within budget or if you'll have a surprise on your next bill.

You need to be able to confidently answer several questions:

- What will this solution cost this fiscal year?
- Is there an alternate configuration you could use to save money?
- Can you estimate how a change would impact your cost and performance without putting it into a production system?

In this module, we'll explore the tools you can use to answer these questions and more. In this module, you will:

- Learn the different options you have to purchase Azure services
- Estimate costs with the Azure pricing calculator
- Predict and optimize costs with Azure Cost Management and Azure Advisor
- Apply best practices for saving on infrastructure costs
- Apply best practices for saving on licensing costs

12.1 PURCHASING AZURE PRODUCTS AND SERVICES

Let's start by examining the **purchasing options** you have with Azure. There are three main customer types on which the available purchasing options for Azure products and services are contingent, including:

- **Enterprise** - Enterprise customers sign an Enterprise Agreement with Azure that commits them to spend a negotiated amount on Azure services, which they typically pay annually. Enterprise customers also have access to customized Azure pricing.
- **Web direct** - Direct Web customers pay general public prices for Azure resources, and their monthly billing and payments occur through the Azure website.
- **Cloud Solution Provider** - Cloud Solution Provider (CSP) typically are Microsoft partner companies that a customer hires to build solutions on top of Azure. Payment and billing for Azure usage occur through the customer's CSP.

When you provision an Azure resource, Azure creates one or more meter instances for that resource. The **meters** track the resources' usage, and generate a usage record that is used to calculate your bill.

For example, a single virtual machine that you provision in Azure might have the following meters tracking its usage:

- Compute Hours
- IP Address Hours
- Data Transfer In
- Data Transfer Out
- Standard Managed Disk
- Standard Managed Disk Operations
- Standard IO-Disk
- Standard IO-Block Blob Read
- Standard IO-Block Blob Write
- Standard IO-Block Blob Delete

The meters and pricing vary per product and often have different pricing tiers based on the size or capacity of the resource. At the end of each monthly billing cycle, the usage values will be charged to your payment method and the meters are reset.

The key takeaway is that resources are always charged based on usage. For example, if you de-allocate a VM then you will not be billed for compute hours, I/O reads or writes or the private IP address since the VM is not running and has no allocated compute resources. However you will incur storage costs for the disks.

12.2 FACTORS AFFECTING COSTS

Just like your on-premises equipment costs, there are several elements that will affect your monthly costs when using Azure services. Let's look at a few of the primary factors including resource type, services, the user's location, and the billing zone.

12.2.1 RESOURCE TYPE

Costs are resource-specific, so the usage that a meter tracks and the number of meters associated with a resource depend on the resource type.

Each meter tracks a particular kind of usage. For example, a meter might track bandwidth usage (ingress or egress network traffic in bits-per-second), the number of operations, size (storage capacity in bytes), or similar items.

The usage that a meter tracks correlates to a number of **billable units**. Those are charged to your account for each billing period, and the rate per billable unit depends on the resource type you are using.

12.2.2 SERVICES

Azure usage rates and billing periods can differ between Enterprise, Web Direct, and Cloud Solution Provider (CSP) customers. Some subscription types also include usage allowances, which affect costs.

The Azure team develops and offers first-party products and services, while products and services from third-party vendors are available in the Azure Marketplace. Different billing structures apply to each of these categories.

12.2.3 LOCATION

Azure has datacenters all over the world. Usage costs vary between locations that offer particular Azure products, services, and resources based on popularity, demand, and local infrastructure costs.

12.2.4 AZURE BILLING ZONES

Bandwidth refers to data moving in and out of Azure datacenters. Most of the time inbound data transfers (data going into Azure datacenters) are free. For outbound data transfers (data going out of Azure datacenters), the data transfer pricing is based on **Billing Zones**.

A **Zone** is a geographical grouping of Azure Regions for billing purposes. The following zones exist and include the listed countries (regions) listed.

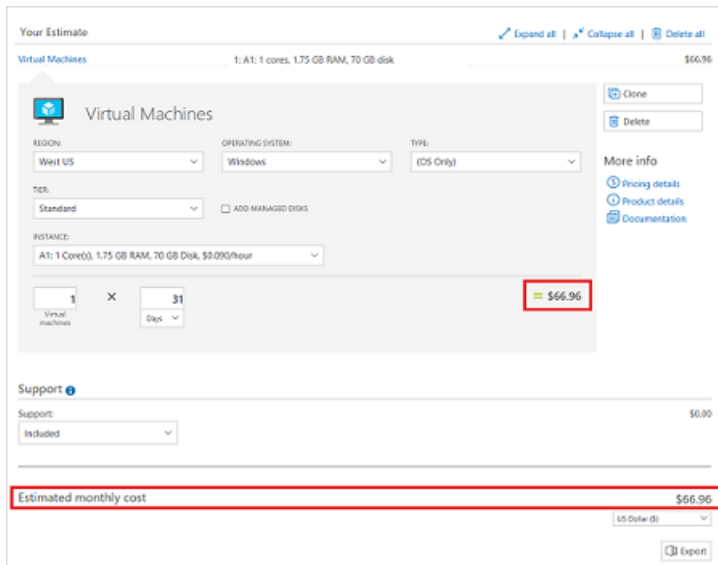
Zone	Areas
Zone 1	United States, Europe, Canada, UK, France
Zone 2	Asia Pacific, Japan, Australia, India, Korea
Zone 3	Brazil
DE Zone 1	Germany

The cheapest outbound networking costs are in zone 1, followed by DE Zone 1, Zone 2 and Zone 3.

Note: Billing zones aren't the same as an Availability Zone. In Azure, the term zone is for *billing purposes* only, and the full term Availability Zone refers to the failure protection that Azure provides for datacenters.

12.3 ESTIMATE COSTS WITH THE AZURE PRICING CALCULATOR

To make estimates easy for customers to create, Microsoft developed the **Azure pricing calculator**. The Azure pricing calculator is a free web-based tool that allows you to input Azure services and modify properties and options of the services. It outputs the costs per service and total cost for the full estimate.



Pricing calculator

The options that you can configure in the pricing calculator vary between products, but basic configuration options include:

Option	Description
Region	Lists the regions from which you can provision a product. Southeast Asia, central Canada, the western United States, and Northern Europe are among the possible regions available for some resources.
Tier	Sets the type of tier you wish to allocate to a selected resource, such as Free Tier, Basic Tier, etc.
Billing Options	Highlights the billing options available to different types of customer and subscriptions for a chosen product.
Support Options	Allows you to pick from included or paid support pricing options for a selected product.
Programs and Offers	Allows you to choose from available price offerings according to your customer or subscription type.
Azure Dev/Test Pricing	Lists the available development and test prices for a product. Dev/Test pricing applies only when you run resources within an Azure subscription that is based on a Dev/Test offer.

12.4 PREDICT AND OPTIMIZE WITH COST MANAGEMENT AND AZURE ADVISOR

We learned how to estimate your costs before you deploy services on Azure, but what if you already have resources deployed? How do you get visibility into the costs you're already accruing? If we had deployed our previous solution to Azure and now want to make sure that we've sized the virtual machines properly and predict how much our bill will be, how can we do this? Let's look at a few tools on Azure that you can use to help you solve this problem.

12.4.1 WHAT IS AZURE ADVISOR?

Azure Advisor is a free service built into Azure that provides recommendations on:

- high availability
- security,
- performance
- cost

Advisor analyzes your deployed services and looks for ways to improve your environment across those four areas. We'll focus on the cost recommendations, but you'll want to take some time to review the other recommendations as well.

Advisor makes **high availability recommendations** in the following areas:

- Ensure virtual machine fault tolerance
- Ensure availability set fault tolerance
- Use Managed Disks to improve data reliability
- Ensure application gateway fault tolerance
- Protect your virtual machine data from accidental deletion
- Configure Traffic Manager endpoints for resiliency
- Use production VPN gateways to run your production workloads

Advisor makes **performance** in the following areas:

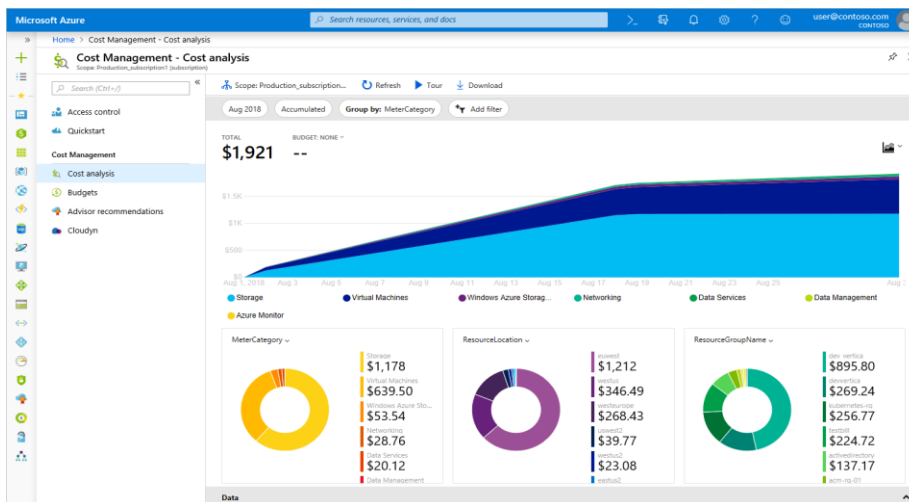
- Improve database performance with SQL DB Advisor
- Improve App Service performance and reliability
- Use Managed Disks to prevent disk I/O throttling
- Improve the performance and reliability of virtual machine disks by using Premium Storage
- Design your storage accounts to prevent hitting the maximum subscription limit
- Add regions with traffic to your Azure Cosmos DB account

Advisor makes **cost recommendations** in the following areas:

- **Reduce costs by eliminating unprovisioned Azure ExpressRoute circuits.** This identifies ExpressRoute circuits that have been in the provider status of Not Provisioned for more than one month and recommends deleting the circuit if you aren't planning to provision the circuit with your connectivity provider.
- **Buy reserved instances to save money over pay-as-you-go.** This will review your virtual machine usage over the last 30 days and determine if you could save money in the future by purchasing reserved instances. Advisor will show you the regions and sizes where you potentially have the most savings and will show you the estimated savings you might achieve from purchasing reserved instances.
- **Right-size or shutdown underutilized virtual machines.** This monitors your virtual machine usage for 14 days and then identifies underutilized virtual machines. Virtual machines whose average CPU utilization is 5 percent or less and network usage is 7 MB or less for four or more days are considered underutilized virtual machines. The average CPU utilization threshold is adjustable up to 20 percent. By identifying these virtual machines, you can decide to resize them to a smaller instance type, reducing your costs.
- Delete unassociated public IP addresses to save money

12.4.2 AZURE COST MANAGEMENT

Azure Cost Management is another free, built-in Azure tool that can be used to gain greater insights into where your cloud money is going. You can see historical breakdowns of what services you are spending your money on and how it is tracking against budgets that you have set. You can set budgets, schedule reports, and analyze your cost areas.



Cost management

12.5 ESTIMATE THE TOTAL COST OF OWNERSHIP WITH THE AZURE TCO CALCULATOR

The pricing calculator and cost management advisor can help you predict and analyze your spend for new or existing services.

If you are starting to migrate to the cloud, a useful tool you can use to predict your cost savings is the **Total Cost of Ownership (TCO)** calculator. The TCO calculator uses a comprehensive list of on-premises 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.

To use the TCO calculator, you need to complete four steps.

12.5.1 STEP 1: OPEN THE TCO CALCULATOR

Start by opening the Total Cost of Ownership calculator website.

12.5.2 STEP 2: DEFINE YOUR WORKLOADS

Start by entering details about your on-premises infrastructure into the TCO calculator according to four groups:

Group	Description
Servers	Enter details of your current on-premises server infrastructure.
Databases	Enter details of your on-premises database infrastructure in the Source section. In the Destination section, select the corresponding Azure service you would like to use.
Storage	Enter the details of your on-premises storage infrastructure.
Networking	Enter the amount of network bandwidth you currently consume in your on-premises environment.


12.5.3 STEP 3: ADJUST ASSUMPTIONS

Adjust the values of assumptions that the TCO calculator makes, which might vary between customers. To improve the accuracy of the TCO calculator, you should adjust the values, so they match the costs of your current on-premises infrastructure. The assumptions you can customize include:

- Storage costs
- IT labor costs
- Hardware costs
- Software costs
- Electricity costs
- Virtualization costs
- Datacenter costs
- Networking costs
- Database costs

12.5.4 STEP 4: VIEW THE REPORT

The TCO calculator generates a detailed report based on the details you enter and the adjustments you make. The report allows you to compare the costs of your on-premises infrastructure with the costs of using Azure products and services to host your infrastructure in the cloud.



12.6 SAVE ON INFRASTRUCTURE COSTS

We have seen how to create cost estimates for environments you'd like to build, walked through some tools to get details on where we're spending money, and projected future expenses. Our next challenge is to look at how to reduce those infrastructure costs.

12.6.1 USE AZURE CREDITS

Visual Studio subscribers can activate a monthly credit benefit which allows you to experiment with, develop, and test new solutions on Azure. Use Azure credits to try out new services such as App Service, Windows 10 VMs, Azure SQL Server databases, Containers, Cognitive Services, Functions, Data Lake, and more without incurring any monetary costs.

When you activate this benefit, you will own a separate Azure subscription under your account with a monthly credit balance that renews each month while you remain an active Visual Studio subscriber.

12.6.2 USE SPENDING LIMITS

By default, Azure subscriptions which have associated monthly credits (which includes trial accounts) have a spending limit to ensure you aren't charged once you have used up your credits. This feature is useful for development teams exploring new solution architectures as it ensures you won't have an unexpectedly large bill at the end of the month.

Azure provides the **Spending Limits** feature to help prevent you from exhausting the credit on your account within each billing period. When your Azure usage results in charges that use all the included monthly credit, the services that you deployed are disabled and turned off for the rest of that billing period. Once a new billing period starts, assuming there are credits available, the resources are re-activated and deployed.

You are notified by email when you hit the spending limit for your subscription. In addition, the Azure portal includes notifications about your credit spend. You can adjust the spending limit as desired or even turn it off.

12.6.3 USE RESERVED INSTANCES

If you have VM workloads that are static and predictable, particularly ones that run 24x7x365, using reserved instances is a fantastic way to potentially save up to 70-80%, depending on the VM size.

Reserved instances are purchased in one-year or three-year terms, with payment required for the full term up front. After it's purchased, Microsoft matches up the reservation to running instances and decrements the hours from your reservation.

12.6.4 CHOOSE LOW-COST LOCATIONS AND REGIONS

The cost of Azure products, services, and resources can vary across locations and regions, and if possible, you should use them in those locations and regions where they cost less.

12.6.5 RESEARCH AVAILABLE COST-SAVING OFFERS

Keep up-to-date with the latest Azure customer and subscription offers, and switch to offers that provide the most significant cost-saving benefit.

12.6.6 RIGHT-SIZE UNDERUTILIZED VIRTUAL MACHINES

Recall from our previous discussion that Azure Cost Management and Azure Advisor might recommend right-sizing or shutting down VMs. Right-sizing a virtual machine is the process of resizing it to a proper size. Over-sized virtual machines are a common unnecessary expense on Azure and one that can be easily fixed. You can change the size of a VM through the Azure portal, Azure PowerShell, or the Azure CLI.

12.6.7 DEALLOCATE VIRTUAL MACHINES IN OFF HOURS

If you have virtual machine workloads that are only used during certain periods, but you're running them every hour of every day, you're wasting money. These VMs are great candidates to shut down when not in use and start back up on a schedule, saving you compute costs while the VM is deallocated.

This approach is an excellent strategy for development environments. It's often the case that development may happen only during business hours, giving you the flexibility to deallocate these systems in the off hours and stopping your compute costs from accruing. Azure now has an *automation solution* fully available for you to leverage in your environment.

You can also use the **auto-shutdown** feature on a virtual machine to schedule automated shutdowns.

12.6.8 DELETE UNUSED VIRTUAL MACHINES

This advice may sound obvious, but if you aren't using a service, you should shut it down. It's not uncommon to find non-production or proof-of-concept systems left around following a project that is no longer needed. Regularly review your environment and work to identify these systems. Shutting down these systems can have a multifaceted benefit by saving you not only on infrastructure costs but also potential savings on licensing and operations.

12.6.9 MIGRATE TO PAAS OR SAAS SERVICES

Lastly, as you move workloads to the cloud, a natural evolution is to start with infrastructure-as-a-service (IaaS) services and then move them to platform-as-a-service (PaaS) as appropriate, in an iterative process.

PaaS services typically provide substantial savings in both resource and operational costs. The challenge is that depending on the type of service, varying levels of effort will be required to move to these services from both a time and resource perspective. You might be able to move a SQL Server database to Azure SQL Database easily, but it might take substantially more effort to transfer your multi-tier application to a container or serverless based architecture. It's a good practice to continuously evaluate the architecture of your applications to determine if there are efficiencies to be gained through PaaS services.

Azure makes it easy to test these services with little risk, giving you the ability to try out new architecture patterns relatively easily. That said, it's typically a longer journey and might not be of immediate help if you're looking for quick wins from a cost-savings perspective. The Azure Architecture Center is a great place to get ideas for transforming your application, as well as best practices across a wide array of architectures and Azure services.

12.7 SAVE ON LICENSING COSTS

Licensing is another area that can dramatically impact your cloud spending. Let's look at some ways you can reduce your licensing costs.

12.7.1 LINUX VS. WINDOWS

Many of the Azure services you deploy have the choice of running on Windows or Linux. In some cases, the cost of the product can be different based on the OS you choose. Where you have a choice, and your application doesn't depend on the underlying OS, it's useful to compare pricing to determine whether you can save money.

12.7.2 AZURE HYBRID BENEFIT FOR WINDOWS SERVER

Many customers have invested in Windows Server licenses and would like to repurpose this investment on Azure. The **Azure Hybrid Benefit** gives customers the right to use these licenses for virtual machines on Azure. That means you won't be charged for the Windows Server license and will instead be billed at the Linux rate.

To be eligible for this benefit, your Windows licenses must be covered by Software Assurance.

12.7.3 AZURE HYBRID BENEFIT FOR SQL SERVER

The Azure Hybrid Benefit for SQL Server helps you maximize the value from your current licensing investments and accelerate your migration to the cloud. Azure Hybrid Benefit for SQL Server is an Azure-based benefit that enables you to use your SQL Server licenses with active Software Assurance to pay a reduced rate.

12.7.4 USE DEV/TEST SUBSCRIPTION OFFERS

The **Enterprise Dev/Test** and **Pay-As-You-Go Dev/Test** offers are a benefit you can take advantage of to save costs on your non-production environments. This benefit gives you several discounts, most notably for Windows workloads, eliminating license charges and only billing you at the Linux rate for virtual machines. This also applies to SQL Server and any other Microsoft software that is covered under a Visual Studio subscription (formerly known as MSDN).

There are a few requirements for this benefit, one being that it's only for non-production workloads, and another being that any users of these environments (excluding testers) must be covered under a Visual Studio subscription. In short, for non-production workloads, this allows you to save money on your Windows, SQL Server, and other Microsoft virtual machine workloads.

Below are the full details of each offer. If you are a customer on an Enterprise Agreement, you'd want to leverage the Enterprise Dev/Test offer, and if you are a customer without an Enterprise Agreement and are instead using PAYG accounts, you'd leverage the Pay-As-You-Go Dev/Test offer.

12.7.5 BRING YOUR OWN SQL SERVER LICENSE

If you are a customer on an Enterprise Agreement and already have an investment in SQL Server licenses, and they have freed up as part of moving resources to Azure, you can provision bring your own license (BYOL) images off the Azure Marketplace, giving you the ability to take advantage of these unused licenses and reduce your Azure VM cost. You've always been able to do this by provisioning a Windows VM and manually installing SQL Server, but this simplifies the creation process by leveraging Microsoft certified images. Search for BYOL in the Marketplace to find these images.

An Enterprise Agreement subscription is required to use these certified BYOL images.

12.7.6 USE SQL SERVER DEVELOPER EDITION

A lot of people are unaware that SQL Server Developer Edition is a free product for nonproduction use. Developer Edition has all the same features that Enterprise Edition has, but for nonproduction workloads, you can save dramatically on your licensing costs.

12.7.7 USE CONSTRAINED INSTANCE SIZES FOR DATABASE WORKLOADS

Many customers have high requirements for memory, storage, or I/O bandwidth but low CPU core counts. Based on this popular request, Microsoft has made available the most popular VM sizes (DS, ES, GS, and MS) in new sizes that constrain the vCPU count to one half or one-quarter of the original VM size, while maintaining the same memory, storage, and I/O bandwidth.

13 SERVICE LIFECYCLE

13.1 PRIVATE PREVIEWS

Private previews are generally the first release of a product and customers are invited to try the service to help provide feedback, to understand if the service is something that is needed as well as help to shape the future of the product.

Generally there is a limited number of customers that are invited into a private preview and these customers are expected to spend time with the Product Group* providing their feedback of the product.

All preview services are provided “as is” and aren’t covered by any Service Level Agreements (SLA). They are generally not supported by customer support either, however some support is provided via the Product Group. There is also no guarantee that a preview feature will go into General Availability.

Previews are generally free of charge.

13.2 PUBLIC PREVIEW

After a product has went through the private preview stage, the next stage is a larger public preview. Again this is used to understand customer needs and help shape the future of the product. Private Preview offerings are often advertised on the Azure blogs or places such as Twitter.

Again, as with the private preview, all preview services are provided “as is” and aren’t covered by any Service Level Agreements (SLA). Again though, previews are generally free of charge.

Getting involved with the private or public previews is a great way of testing services and help drive products in the direction that would be useful for you and your organization.

13.3 GENERAL AVAILABILITY (GA)

When a service goes GA it is in full production mode. It is fully supported by SLAs, customer support and is viable for production workloads. Now that the service is live it is also chargeable.

13.4 HOW TO ACCESS PREVIEW FEATURES

You can activate specific preview features through the preview features page. This page lists the preview features that are available for evaluation. To preview a feature, select the Try it button for the relevant feature. Another preview area you can try is the next version of the Azure portal. Use the URL <https://preview.portal.azure.com>.

14 READ MORE

14.1 STUDY GUIDES

- <https://github.com/AzureMentor/Azure-AZ-900-Study-Guide>
- <https://absolute-sharepoint.com/az-900-study-guide-microsoft-azure-fundamentals>
- <https://www.taygan.co/blog/2019/02/07/az-900-azure-fundamentals-exam-preparation>

14.2 USEFUL LINKS

- The Azure Fundamentals Learning Path, <https://docs.microsoft.com/en-us/learn/paths/azure-fundamentals/>
- Azure Docs, <https://docs.microsoft.com/en-us/azure/>
- Cheshire, J. (2019). *Exam Ref AZ-900, Microsoft Azure Fundamentals*. Pearson Education

14.3 REFERENCES

14.3.1 CLOUD CONCEPTS

- [Pros of cloud computing](#)
- [Cloud computing](#)
- [High Availability vs. Fault Tolerance vs. Disaster Recovery](#)
- [Elastic computing](#)
- [Scalability](#)
- [Reliability](#)
- [Disaster recovery](#)
- [Elasticity vs scalability](#)
- [Economies of scale](#)
- [Capital Expenditure](#)
- [Consumption based pricing model](#)
- [Public cloud](#)
- [Private cloud](#)
- [Hybrid cloud](#)
- [PaaS](#)
- [SaaS](#)
- [IaaS](#)

14.3.2 UNDERSTAND CORE AZURE SERVICES

- [Overview of cloud services](#)
- [Azure Regions](#)
- [Availability Zones](#)
- [Resource Groups](#)
- [Resource Managers](#)
- [Azure Networking](#)
- [Azure Storage](#)
- [Azure IoT](#)
- [Azure IoT Hub](#)
- [Azure IoT Central](#)
- [Azure Machine Learning Service](#)
- [Azure Databricks](#)
- [Azure Virtual WAN](#)
- [Azure Virtual Machine Scale Sets](#)
- [Azure App Service](#)
- [Azure Functions](#)
- [Azure Virtual Network](#)
- [Azure Load Balancer](#)
- [Azure VPN Gateway](#)
- [Azure Application Gateway](#)
- [Azure CDN](#)
- [Azure Blob Storage](#)
- [Azure Disk Storage](#)
- [Azure File Storage](#)
- [Azure Blob Storage Tiers](#)
- [Azure Cosmos DB](#)
- [Azure SQL](#)
- [Azure DB Migration Service](#)
- [Azure SQL Data Warehouse](#)
- [Azure Marketplace](#)
- [Azure HDInsight](#)
- [Azure Data Lake Analytics](#)
- [Azure Machine Learning Service](#)
- [Azure Machine Learning Studio](#)
- [Azure Logic Apps](#)
- [Azure Functions](#)
- [Front Door](#)

14.3.3 UNDERSTAND SECURITY, PRIVACY, COMPLIANCE AND TRUST

- [Azure Firewall](#)
- [Azure DDoS](#)
- [NSG](#)
- [Azure AD](#)
- [Azure MFA](#)
- [Azure Security](#)
- [Azure Security Center](#)
- [Azure Key Vault](#)
- [Azure Information Protection](#)
- [Azure Advanced Threat Protection](#)
- [Azure Policy](#)
- [Azure RBAC](#)
- [Locks](#)

14.3.4 UNDERSTAND AZURE PRICING AND SUPPORT

- [Billing](#)
- [SLAs](#)
- [Cost Management](#)
- [Azure Subscriptions](#)
- [Azure Free Account](#)
- [Azure Support Plans](#)
- [Azure Previews](#)
- [Azure Knowledge Center](#)
- [Azure Support](#)
- [Azure support ticket](#)
- [Azure Advisor](#)
- [Azure Reservations](#)
- [Tags](#)
- [Spending limits](#)
- [Service limits](#)
- [Cost Analysis](#)
- [Prevent unexpected charges](#)
- [Azure Monitor](#)
- [Azure Service Health](#)
- [Privacy Statement](#)
- [Trust Center](#)
- [Service Trust Portal](#)
- [Azure Germany](#)
- [Azure Government](#)
- [Autoscaling](#)

A

- access control · 10, 25, 75, 76, 78, 80, 93, 101, 105, 107, 108
- Activity Logs** · 98
- ADE** · 82
- AI · iv, 8, 46, 48
- analytics · 12, 24, 42, 44, 46, 47, 49, 51, 52, 53, 54, 55, 56, 59, 88
- Apache Spark · 12, 51
- Application Gateway · iii, 37, 62, 64, 69, 85, 87
- Application Insights · 37, 41, 98
- Application SLAs · 23
- Archive storage tier · 60
- Artificial Intelligence · iv, 9, 12
- Authentication · iv, 26, 76, 77, 78
- Authorization** · 76, 101
- Autoscale · 37, 99
- auto-shutdown · 116
- availability set** · 36, 37, 112
- availability zone · 20
- Azure account · 24, 25, 26, 27, 32, 33
- Azure Active Directory · iv, 24, 26, 72, 76, 77, 78, 79, 80, 88, 97, 105
- Azure AD · 24, 26, 27, 76, 77, 78, 79, 80, 81
- Azure AD Identity Protection · 78
- Azure Advanced Threat Protection · iv, 75, 88
- Azure Advisor** · iv, v, 18, 28, 30, 109, 112, 116
- Azure API Management · 11
- Azure App Service · 11, 35, 39
- Azure Application Gateway · 9, 64, 69, 85, 87
- Azure Bastion · 62, 63
- Azure Batch · 9, 36, 37
- Azure Blob storage · 10, 55, 82
- Azure Blobs · 54, 59
- Azure Blueprint · 93, 94
- Azure Cache for Redis · 11
- Azure CLI · iv, 14, 15, 16, 33, 55, 91, 101, 103, 116
- Azure Cloud Shell** · 8, 14, 32, 33, 34
- Azure Container Instances · 9, 38
- Azure Content Delivery Network · 9
- Azure Cosmos DB · 11, 51, 54, 72, 101, 113
- Azure Cost Management** · v, 18, 109, 113, 116
- Azure Data Lake Analytics · 56
- Azure Data Lake Storage Gen2 · 55, 56, 59
- Azure Database for MariaDB · 11
- Azure Database for MySQL · 11
- Azure Database for PostgreSQL · 11
- Azure Database Migration Service · 11, 53
- Azure Databricks · 12, 51, 56, 120
- Azure DDoS Protection · iv, 9, 86
- Azure DevOps · 13, 39
- Azure DevTest Labs · 13
- Azure Disk Encryption · 82
- Azure DNS · 9, 62, 63, 70
- Azure Event Grid · 41, 46
- Azure Event Management · 30
- Azure ExpressRoute · 9, 87, 113
- Azure fabric · 36
- Azure File storage · 10
- Azure Files · 54, 57, 59, 82
- Azure Firewall · iv, 9, 64, 84, 85
- Azure Front Door Service · 62, 64
- Azure Functions · iv, 9, 40, 47, 65
- Azure HDInsight · 12, 56
- Azure Hybrid Benefit · 117, 118
- Azure IoT Hub · 12, 45, 46, 47
- Azure Key Vault · 82, 83
- Azure Knowledge Center · 31
- Azure Kubernetes Service · 9, 38, 98
- Azure Load Balancer · 9, 62, 67, 68, 69, 70, 71
- Azure Logic Apps · 40, 46
- Azure Machine Learning Service · iv, 12
- Azure Machine Learning Studio · 12, 49, 50
- Azure Management Groups · 92
- Azure Monitor · iv, 62, 65, 97, 98, 99
- Azure Network Watcher · 9
- Azure Notification Hubs · 11
- Azure Policy · 89, 90, 91, 92, 93, 103, 104, 105
- Azure portal · 14, 32, 34, 39, 41, 63, 72, 91, 92, 101, 103, 115, 116, 119
- Azure PowerShell · 14, 33
- Azure pricing calculator · 109, 111
- Azure ProDirect · 28
- Azure Queue storage · 10, 58, 82
- Azure Queues · 54
- Azure Rapid Response · 30
- Azure Resource Manager · 80, 91, 93, 94, 95, 100
- Azure Rights Management · 88
- Azure Search · 11
- Azure Security Center · 5, 75, 84, 92
- Azure Service Fabric · 9
- Azure Service Health · iv, 97, 98, 99
- Azure SignalR Service · 11
- Azure SQL Data Warehouse · iii, 11, 12, 51, 53
- Azure SQL Database · iii, 11, 53, 82, 102, 117
- Azure Standard · 28
- Azure Storage Service Encryption · 61, 82
- Azure Stream Analytics · 46, 47
- Azure subscription · v, 24, 26, 27, 31, 33, 75, 92, 97, 98, 100, 107, 112, 115

Azure Table storage · 10, 58
 Azure Tables · 54
 Azure Traffic Manager · 9, 70, 71
 Azure Virtual Machine Scale Sets · 9, 37
 Azure Virtual Machines · 9, 36
 Azure Virtual Network · 9, 63, 64, 66, 84, 87
 Azure Virtual WAN · 9
 Azure VPN Gateway · 9

B

Bash · 14, 33, 34, 39, 79
 Big Data · iv, 9, 12, 55
 billing · v, 24, 25, 26, 27, 28, 29, 30, 31, 102, 103, 104,
 109, 110, 111, 112, 115, 118
 Billing Zones · 110
 BitLocker · 82
 Bring your own key · 82
 BYOK · 82

C

Capital Expenditure · iii, 4
 CDN · 62, 70
 Cloud Solution Provider · 109, 110
 cognitive services · 13
 compliance · iv, v, 3, 5, 6, 19, 20, 35, 37, 61, 75, 82, 89,
 92, 93, 95, 96, 104
 Compliance Manager · v, 95, 96
 composite SLA · 23
 Conditional Access · 78, 79
 container · 1, 2, 15, 24, 34, 35, 37, 38, 56, 83, 98, 100,
 117
 containers · 1, 35, 38, 115
 Content Delivery Network · iii, 62, 64
 Cool storage tier · 60
 CSP · 109, 110
Custom Script Extension · 16

D

data lake · 51, 55, 56, 59
 data residency · 19, 20
 data warehouse · 11, 53, 55
 DDoS · iv, 9, 62, 64, 68, 74, 86, 87
 Developer · 27, 31, 118
 DevOps · 9, 13, 98
 Direct Web customers · 109
 disaster recovery · 3, 23, 55, 60, 67
 Disaster recovery · 4

distributed denial of service · 9, 68, 74
 DNS · 9, 63, 64, 70, 71
 Durable Functions · 40, 41

E

economies of scale · iii, 6
 Economies of scale · 3
 encryption · 10, 45, 61, 69, 72, 81, 82, 83, 88, 91
 Enterprise Agreement · 25, 31, 109, 118
 Enterprise Dev/Test · 118
 ETL · 56
 Event Grid · 41, 42, 47
 Event Hubs · 42, 47, 72
 ExpressRoute · 62, 63, 65, 87, 113

F

fault domain · 36
fault tolerance · 3, 67, 112
 File Transfer Protocol · 84
 firewall · 9, 16, 17, 64, 67, 69, 70, 84, 85, 87
 Firewall · 62, 84, 85
 FTP · 84, 85

G

GA · v, 50, 119
 GDPR · v, 96
 geography · 19, 20, 21
 GUI · 32, 34, 40

H

Hadoop · 12, 54, 56, 57
 high availability · 11, 36, 39, 54, 61, 64, 67, 68, 70, 84,
 112
 High availability · 23, 37, 64, 67
 Hot storage tier · 60
 hybrid cloud · 6, 28, 42

I

IaaS · iii, 6, 7, 36, 72, 82, 117
 IIS · 16, 17
 image · 14, 15, 33, 36, 53
Infrastructure as a Service · 6
 initiative · 89, 92, 95

Initiatives · iv, 92
Internet of Things · iv, 9, 11, 43, 48
IoT · iv, 11, 12, 43, 44, 45, 46, 47, 48, 51, 56, 120
IoT Central · iv, 12, 43, 44, 45, 46
IoT Edge · 12, 45, 46, 47, 48
ISO · v, 96

J

JSON · 40, 53, 90

L

latency · 11, 19, 54, 60, 62, 64, 68, 70, 71, 85
Linux · 1, 8, 9, 13, 14, 33, 34, 35, 39, 79, 82, 98, 117, 118
Load Balancer · iii, 62, 63, 64, 68, 71
Log Analytics · 41, 95, 98
logs · 46, 57, 98, 99

M

Machine Learning · iv, 12, 46, 48, 49, 50
managed identities · 79, 105
managed identity · 80, 105, 107
Marketplace · iii, 34, 39, 87, 110, 118
massive parallel processing · 12
meters · 109, 110
metrics · 37, 57, 98, 99
MFA · 76, 77, 79
microservice · 38
Microsoft Azure Information Protection · 88
Microsoft Privacy Statement · v, 95
Microsoft Threat Intelligence · 75
MPP · 12, 53
MSDN · 31, 118
MSIP · 88
multi-factor authentication · 74, 76, 78, 79, 81
MySQL · 11

N

NAT · 85
network security group · 67, 85
Network security groups · 64, 85
Network Security Groups · 62, 66, 87
Network Watcher · 62, 65
NIST · v, 96
NoSQL · 10, 11, 52, 54, 56, 58
notebooks · 49

NSG · iv, 65, 67, 87
N-tier · 66

O

Operational Expenditure · iii, 4
orchestrator · 38
OS · 1, 35, 36, 82, 97, 117

P

PaaS · iii, 6, 7, 35, 38, 39, 44, 45, 72, 117
Pay-As-You-Go Dev/Test · 25, 118
PIM · 80, 81
Platform as a Service · 6
policy assignment · 91, 92
policy definition · 90, 91
PolyBase · 53, 54
PostgreSQL · 11
Power BI · 47
PowerShell · iv, 14, 16, 32, 33, 34, 39, 41, 55, 91, 101, 103, 116
PowerShell Core · 33
Premier · v, 27, 29, 31
private cloud · 5
Private previews · 119
Privileged Identity Management · 80
ProDirect Managers · 28
Professional Direct · v, 27, 31
public cloud · 5, 6
public IP address · 15, 17, 63, 66
public preview · 119

R

RBAC · iv, 80, 93, 101, 102, 105, 106, 107, 108
RDP · 63, 84, 85
redundancy · 3, 10, 19, 21, 36, 37, 67
region · 3, 14, 19, 20, 21, 66, 70, 71, 85, 86
region pairs · 21
relational database · 11, 53
Remote Desktop Protocol · 84
Reserved instances · 116
resiliency · 19, 23, 62, 68, 70, 112
resource group · 14, 15, 16, 91, 92, 95, 100, 101, 102, 103, 105, 106, 107, 108
Resource locks · 108
REST · 33, 39, 40, 41, 45, 48, 55, 57, 59, 103
Roles · 80, 106

S

SaaS · iii, 6, 7, 12, 34, 44, 72, 73, 76, 78, 117
 Scale up · 8, 18
 Secure Shell · 84
 Server Message Block · 57
 serverless computing · iv, 1, 2, 35, 40
 Service Bus · 42, 47
 service principals · 79, 80
 Service Trust Portal · v, 95, 96
 service-level agreements · 89
 Single-Sign-On · 76
 SLA · v, 22, 23, 25, 99, 119
 SLAs · v, 22, 23, 89, 119, 121
 SMB · 57, 59
Software as a Service · 6
 Spark · 51, 54, 56, 57
 spending limits · v, 26, 115
Spending Limits · 115
 SQL · iii, iv, 10, 11, 12, 19, 23, 25, 39, 47, 51, 53, 54, 58, 69, 82, 87, 90, 91, 102, 105, 113, 115, 117, 118
 SQL Data Warehouse · iv, 53, 54
 SQL Server · 10, 11, 53, 82, 90, 115, 117, 118
 SQL Server on VMs · 11
 SSE · 61
 SSH · 63, 67, 84, 85
 SSO · 76, 77
 Standard · v, 15, 18, 27, 28, 31, 75, 86, 87, 109
 storage account · 54, 58, 61, 81, 108
 subnet · 66, 72, 87, 105
 support · v, 5, 11, 23, 24, 27, 28
 support plans · v, 27, 31

T

tags · v, 52, 91, 100, 103, 104
 TCO calculator · 114
 TCP · 68, 69
 Technical Account Manager · 29, 30
 telemetry · 42, 43, 44, 46, 47, 60, 72, 73, 97
 tenant · 26, 27, 80, 89, 97
 Three-tier · 66
 Traffic Manager · 62, 64, 71, 112

Transmission Control Protocol · 68
 Transparent data encryption · 82
 Trust Center · 95

U

UDP · 68, 86
 update domain · 36, 37
 User Datagram Protocol · 68

V

VHD · 58
 virtual hard disk · 58, 59
 virtual machine · 1, 8, 14, 15, 19, 21, 33, 34, 35, 37, 38, 58, 59, 82, 90, 104, 106, 108, 109, 112, 113, 116, 118
 Virtual Machine Scale Sets · iii, 36
 virtual machines · 14, 34, 35, 113
 virtual network · 14, 63, 64, 66, 67, 68, 85, 87, 103
 Virtual Network · iii, 62
 virtual network peering · 63, 66
 Virtual private network · 87
 Virtual Private Network · 9
 Visual Studio · 25, 40, 41, 49, 98, 108, 115, 118
 Visual Studio Team Services · 25
 VNet · 62, 63, 66, 85, 87
 VPN · iii, 9, 62, 63, 65, 66, 82, 87, 112
 VPN gateway · 66

W

WAF · 64, 69, 85
 WAN · 9, 62, 63
 Web Application Firewall · 62, 64, 85
 WebHook · 41

Z

Zone · 110, 111