

AZURE FUNDAMENTAL AZ-900**Principles of cloud computing**

The computing services:

- **Compute power**
When you send an email, book a reservation on the Internet, pay a bill online, or even take this Microsoft Learn module you're interacting with cloud-based servers that are processing each request and returning a response. As a consumer, we're all dependent on the computing services provided by the various cloud providers that make up the Internet.
- **VM** is an emulation of a computer
- **Containers** are similar to VMs except they don't require a guest operating system. 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.
- **Storage**
- **Networking**
- **Analytics**

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. This architecture doesn't work for every app - but when the app logic can be separated to independent units, you can test them separately, update them separately, and launch them in microseconds, making this approach the fastest option for deployment.

Benefits:

- **Cost-efficient (pay-as-you-go)**
- **Scalable**

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.

- **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.
- **Current**
When you use the cloud, you're able to focus on what matters: building and deploying applications.
- **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
- **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.
- **Security**

Compliance Offerings:

- Criminal Justice Information Services (CJIS)
- Cloud Security Alliance (CSA) STAR Certification
- General Data Protection Regulation (GDPR)
- EU Model Clauses
- Health Insurance Portability and Accountability Act (HIPAA)
- International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) 27018.
- Multi-Tier Cloud Security (MTCS) Singapore
- Service Organization Controls (SOC)
- National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF)
- UK Government G-Cloud

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

Two approaches to investment are referred:

- **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.

CapEx computing costs:

- Server cost
- Storage cost
- Network cost
- Backup and archive cost
- Organization continuity and disaster recovery cost
- Datacenter infrastructure costs
- Technical personnel

OpEx computing costs:

- Leasing software and customized features
- Scaling charges based on usage/demand instead of fixed hardware or capacity. Billing categories can include number of users or CPU usage time, allocated RAM, I/O operations per second (IOPS), and storage space.
- Billing at the user or organization level.

Benefits of CapEx

With capital expenditures, you plan your expenses at the start of a project or budget period. Your costs are fixed, meaning you know exactly how much is being spent. This is appealing when you need to predict the expenses before a project starts due to a limited budget.

Benefits of OpEx

Demand and growth can be unpredictable and can outpace expectation

Three deployment methods of cloud computing

- **Public cloud:** (Ex: Microsoft Azure)
- In this case, you have no local hardware to manage or keep up-to-date – everything runs on your cloud provider's hardware. A public cloud is a shared entity where multiple corporations use a portion of the resources in the cloud

- Advantages

- High scalability/agility – you don't have to buy a new server in order to scale
- Pay-as-you-go pricing, no CapEx costs
- You're not responsible for maintenance or updates of the hardware
- Minimal technical knowledge to set up and use

- Disadvantages

- There may be specific security requirements that cannot be met by using public cloud
- There may be government policies, industry standards, or legal requirements which public clouds cannot meet
- You don't own the hardware or services and cannot manage them as you may want to
- Unique business requirements, such as having to maintain a legacy application might be hard to meet

- Private cloud (add additional resources)

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.

- Hybrid cloud

A hybrid cloud combines public and private clouds, allowing you to run your applications in the most appropriate location. For example, you could host a website in the public cloud and link it to a highly secure database hosted in your private cloud (or on-premises datacenter).

Advantages of Hybrid cloud:

- You can keep any systems running and accessible that use out-of-date hardware or an out-of-date operating system
- You have flexibility with what you run locally versus in the cloud
- 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

Disadvantages of hybrid cloud:

- It can be more expensive than selecting one deployment model since it involves some CapEx cost up front
- It can be more complicated to set up and manage

Types of cloud services:

IaaS is commonly used in the following scenarios:

- **Migrating workloads.** Typically, IaaS facilities are managed in a similar way as on-premises infrastructure and provide an easy migration path for moving existing applications to the cloud.
- **Test and development.** Teams can quickly set up and dismantle test and development environments, bringing new applications to market faster. IaaS makes scaling development and testing environments, fast and economical.
- **Storage, backup, and recovery.** Organizations avoid the capital outlay and complexity of storage management, which typically requires skilled staff to manage data and meet legal and compliance requirements. IaaS is useful for managing unpredictable demand and steadily growing storage needs. IaaS can also simplify the planning and management of backup and recovery systems.

PaaS is commonly used in the following scenarios:

- **Development framework.** PaaS provides a framework that developers can build upon to develop or customize cloud-based applications. Just like Microsoft Excel macro, PaaS lets developers create applications using built-in software components. Cloud features such as scalability, high-availability, and multi-tenant capability are included, reducing the amount of coding that developers must do.
- **Analytics or business intelligence.** Tools provided as a service with PaaS allow organizations to analyze and mine their data. They can find insights and patterns, and predict outcomes to improve business decisions such as forecasting, product design, and investment returns
- PaaS provides the ability to scale the platform automatically
- PaaS provides professional development services to add features to applications

Serverless computing is the abstraction of servers, infrastructure, and OSs. With *serverless* computing, Azure takes care of managing the server infrastructure and allocation/deallocation of resources based on demand. Serverless computing encompasses three ideas:

1. The abstraction of servers: Serverless computing abstracts the servers you run on. You never explicitly reserve server instances; the platform manages that for you
2. Event-driven scale: Serverless computing is an excellent fit for workloads that respond to incoming events
3. Micro-billing

Additional Azure subscriptions:

- **Environments:** When managing your resources, you can choose to create subscriptions to set up separate environments for development and testing, security, or to isolate data for compliance reasons. This is particularly useful because resource access control occurs at the subscription level.
- **Organizational structures:** You can create subscriptions to reflect different organizational structures. For example, you could limit a team 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.
- **Billing:** You might want to also create additional subscriptions for billing purposes. Because costs are first aggregated at the subscription level, you might want to create subscriptions to manage and track costs based on your needs. For instance, you might want to create a subscription for your production workloads and another subscription for your development and testing workloads.

Developer	Standard	Professional Direct
Non-critical workloads	Production workloads	Business-critical workloads
1 business day response	1-hour response for critical cases	1-hour response + priority tracking of critical cases
Not applicable	Not applicable	Access to a pool of technical experts

Compute

Azure Virtual Machines (IaaS)	<ul style="list-style-type: none"> - When to use VM: <ul style="list-style-type: none"> - During testing and development - When running applications in the cloud - When extending your datacenter to the cloud - During disaster recovery - VM Availability sets: <p>An availability set is a logical grouping of two or more VMs that help keep your application available during planned or unplanned maintenance.</p> <p><i>A planned maintenance event</i> 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</p> <p><i>Unplanned maintenance events</i> 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.</p> <ul style="list-style-type: none"> - With an availability set, you get: <ul style="list-style-type: none"> - Up to three fault domains that each have a server rack with dedicated power and network resources - Five logical update domains which then can be increased to a maximum of 20
Azure Virtual Machine Scale Sets	<ul style="list-style-type: none"> - Azure Virtual Machine Scale Sets let you create and manage a group of identical, load balanced VMs - 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
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	<ul style="list-style-type: none"> - Managed service for parallel and high-performance computing applications - Azure Batch enables large-scale job scheduling and compute management with the ability to scale to tens, hundreds, or thousands of VMs.
Azure Container Instances (PaaS)	<ul style="list-style-type: none"> - Run containerized apps on Azure without provisioning servers or VMs - You can move existing applications to containers and run them within AKS. You can control access via integration with Azure Active Directory (Azure AD) and access Service Level Agreement (SLA)-backed Azure services, such as Azure Database for MySQL for any data needs, via Open Service Broker for Azure (OSBA). AKS is deployed with a virtual network.
Azure Functions	<ul style="list-style-type: none"> - An event-driven, serverless compute service - When you're concerned only about the code running your service, and not the underlying platform or infrastructure. Azure Functions is fast and scale automatically based on demand, so they're a solid choice when demand is variable. Azure Functions can be either stateless (default) or stateful
Azure Logic Apps	<ul style="list-style-type: none"> - Designed in a web-based designer and can execute logic triggered by Azure services without writing any code. - Similar to Azure Function where Functions execute code, Logic Apps execute <i>workflows</i> designed to automate business scenarios and built from predefined logic blocks

-	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.

Networking

Azure Virtual Network	<ul style="list-style-type: none"> - Connects VMs to incoming Virtual Private Network (VPN) connections - 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 <i>subnets</i>. Subnets help you organize and secure your resources in discrete sections. The web, application, and data tiers each have a single VM. All three VMs are in the same virtual network but are in separate subnets. - You configure virtual networks and gateways through software, which enables you to treat a virtual network just like your own network. - For communication between virtual machines, <i>Network Security Groups</i> (NSGs) are a critical piece to restrict unnecessary communication.
Azure Load Balancer	<ul style="list-style-type: none"> - Balances inbound and outbound connections to applications or service endpoints 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 - Availability refers to how long your service is up and running without interruption - Resiliency refers to a system's ability to stay operational during abnormal conditions. - A load balancer distributes traffic evenly among each system in a pool. A load balancer can help you achieve both high availability and resiliency. - We need to define the forwarding rules based on the source IP and port to a set of destination IP/ports.
Azure Application Gateway	<ul style="list-style-type: none"> - Optimizes app server farm delivery while increasing application security - 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 is designed to protect HTTP traffic.

	<p>Benefits of using Azure Application Gateway over a simple load balancer:</p> <ul style="list-style-type: none"> • 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. • 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 <i>content delivery network</i>. • 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 VPN Gateway	<ul style="list-style-type: none"> - Accesses Azure Virtual Networks through high-performance VPN gateways - VPN provide a secure connection between an Azure Virtual Network and an on-premises location over the internet.
Azure DNS	Provides ultra-fast DNS responses and ultra-high domain availability
Azure Content Delivery Network	<ul style="list-style-type: none"> - 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 - A content delivery network (CDN) is used to deploy a website in Azure that will be accessed by users worldwide and that host large video files for best video playback experiences.
Azure DDoS Protection	<ul style="list-style-type: none"> - Protects Azure-hosted applications from distributed denial of service (DDoS) attacks <p>Azure DDoS Protection provides the following service tiers:</p> <ul style="list-style-type: none"> • Basic - The Basic service tier is automatically enabled as part of the Azure platform. Always-on traffic monitoring and real-time mitigation of common network-level attacks provide the same defenses that Microsoft's online services use. Azure's global network is used to distribute and mitigate attack traffic across regions. • Standard - The Standard service tier provides additional mitigation capabilities that are tuned specifically to Microsoft Azure Virtual Network resources. DDoS Protection Standard is simple to enable and requires no application changes. Protection policies are tuned through dedicated traffic monitoring and machine learning algorithms. Policies are applied to public IP addresses associated with resources deployed in virtual networks, such as Azure Load Balancer and Application Gateway. DDoS standard protection can mitigate the following types of attacks: <ul style="list-style-type: none"> ○ Volumetric attacks. The attackers goal is to flood the network layer with a substantial amount of seemingly legitimate traffic. ○ Protocol attacks. These attacks render a target inaccessible, by exploiting a weakness in the layer 3 and layer 4 protocol stack. ○ Resource (application) layer attacks. These attacks target web application packets to disrupt the transmission of data between hosts
Azure Traffic Manager (Reduce latency)	<ul style="list-style-type: none"> - Distributes network traffic across Azure regions worldwide - Traffic Manager uses the DNS server that's closest to the user to direct user traffic to a globally distributed endpoint.

Azure ExpressRoute	<ul style="list-style-type: none"> - To provide a dedicated, private connection between your network and Azure - Extend your on-premises networks into the Microsoft cloud over a private connection facilitated by a connectivity provider - Improve the security of your on-premises communication by sending this traffic over the private circuit instead of over the public internet.
Azure Network Watcher	Monitors and diagnoses network issues using scenario-based analysis
Azure Firewall	<ul style="list-style-type: none"> - Implements high-security, high-availability firewall with unlimited scalability - It 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 Virtual WAN	Creates a unified wide area network (WAN), connecting local and remote sites
Azure Network Security Group	<ul style="list-style-type: none"> - A <i>network security group</i>, 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. - 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.
Network Virtual Appliances	<ul style="list-style-type: none"> - Are ideal options for non-HTTP services or advanced configurations, and are similar to hardware firewall appliances.

Storage (IaaS)

Azure Blob storage	<p>Storage service for very large objects, such as video files or bitmaps</p> <p>Three storage tiers for blob object storage:</p> <ol style="list-style-type: none"> 1. Hot storage tier: optimized for storing data that is accessed frequently. 2. Cool storage tier: optimized for data that are infrequently accessed and stored for at least 30 days. 3. Archive storage tier: for data that are rarely accessed and stored for at least 180 days with flexible latency requirements.
Azure File storage	File shares that you can access and manage like a file server. Azure Files offers fully managed file shares in the cloud that are accessible via the industry standard Server Message Block (SMB) protocol that ensures the data is encrypted at rest and in transit.
Azure Queue storage	<p>A data store for queuing and reliably delivering messages between applications</p> <p>You can use queue storage to:</p> <ul style="list-style-type: none"> • 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.
Azure Table storage	A NoSQL store that hosts unstructured data independent of any schema
Azure DataLake storage	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.

Databases

Azure Cosmos DB (PaaS)	Globally distributed database that supports NoSQL options. You can use this feature to store data that is updated and maintained by users around the world
Azure SQL Database (PaaS)	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 Synapse Analytics	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
Azure SQL Data Warehouse	High availability SQL, for data analytics

Web

Azure App Service (PaaS)	<ul style="list-style-type: none"> - Quickly create powerful cloud web-based apps - HTTP-based service that enables you to build and host many types of web-based solutions without managing infrastructure - For example, you can host web apps, mobile back ends, and RESTful APIs in several supported programming languages - With Azure App Service, you can host most common web app styles including: <ul style="list-style-type: none"> • Web Apps: full support for hosting web apps ex: ASP, PHP • API Apps: build REST based Web API • WebJobs: Allows run a program (app logic) • Mobile Apps: Store app data in SQL database, notification, authenticate customers, execute backend logic in C#
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 Cognitive 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.

Internet of Things

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 between and monitoring of millions of IoT devices
IoT Edge	Push your data analysis models directly onto your IoT devices, allowing them to react quickly to state changes without needing to consult cloud-based AI models.

Big Data

Azure Synapse Analytics (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	Collaborative Apache Spark-based analytics service that can be integrated with other Big Data services in Azure.

Artificial Intelligence

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

DevOps

Azure DevOps	Azure DevOps Services (formerly known as Visual Studio Team Services, or VSTS), 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

Additional notes:

- Create a *resource group* to hold all the things that we need to create. The *resource group* allows us to administer all the services, disks, network interfaces, and other elements that potentially make up our solution as a unit.
- *Scale* refers to adding network bandwidth, memory, storage, or compute power to achieve better performance.
- Scaling up, or vertical scaling means to increase the memory, storage, or compute power on an existing virtual machine. For example, you can add additional memory to a web or database server to make it run faster.
- Scaling out, or horizontal scaling means to add extra virtual machines to power your application. For example, you might create many virtual machines configured in exactly the same way and use a load balancer to distribute work across them.
- Scaling down or scaling in can help you save money
- Workload categories: Dev/Test < Production < Isolated

Review:

- Azure provide flexibility between CapEx and OpEx
- Even if an azure VM is stopped, you continue to pay storage costs of the VM
- An organization that hosts its infrastructure in a public cloud can decommission its data center.
- A company can extend the capacity of its internal network using the public cloud
- The azure portal is: portal.azure.com
- Azure database that can add concurrently from multiple regions and can store JSON documents is Azure Cosmos DB

- 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
- Examples of regions are *West US*, *Canada Central*, *West Europe*, *Australia East*, and *Japan West*.

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.

Geographies are broken up into the following areas:

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

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.

Availability Zones are primarily for VMs, managed disks, load balancers, and SQL databases. Azure services that support Availability Zones fall into two categories:

- **Zonal services** – you pin the resource to a specific zone (for example, virtual machines, managed disks, IP addresses)
- **Zone-redundant services** – platform replicates automatically across zones (for example, zone-redundant storage, SQL Database).

Each Azure region is always paired with another region within the same geography (such as US, Europe, or Asia) at least **300 miles away**.

Advantages of region pairs include:

- If there's an extensive Azure outage, one region out of every pair is prioritized to make sure at least one is restored as quick as possible for applications hosted in that region pair.
 - Reliable services and data redundancy
 - Planned Azure updates are rolled out to paired regions one region at a time to minimize downtime and risk of application outage.
 - Data continues to reside within the same geography as its pair (except for Brazil South) for tax and law enforcement jurisdiction purposes.
- 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:

1. Performance Targets
2. Uptime and Connectivity Guarantees

3. Service credits
 - 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.
 - **Availability** refers to the time that a system is functional and working
-

Feature preview categories

There are two types of previews available:

- **Private Preview.** An Azure feature marked "private preview" is available to *specific* Azure customers for evaluation purposes. This is typically by invite only and issued directly by the product team responsible for the feature or service.
- **Public Preview.** An Azure feature marked "public preview" is available to *all* Azure customers for evaluation purposes. These previews can be turned on through the preview features page as detailed below.

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 32 TB of data in its virtual disks. This capability is significant 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.

Three primary types of data that Azure Storage is designed to hold.

1. **Structured data.** Structured data is data that adheres to a schema, so all of the data has the same fields or properties. Examples of structured data include sensor data or financial data.
2. **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.
3. **Unstructured data.** Unstructured data encompasses data that has no designated structure to it. 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.

Encryption types are available for your resources:

1. **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.
2. **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.

The following table describes the differences between on-premises storage and Azure data storage.

Needs	On-premises	Azure data storage
Compliance and security	Dedicated servers required for privacy and security	Client-side encryption and encryption at rest
Store structured and unstructured data	Additional IT resources with dedicated servers required	Azure Data Lake and portal analyzes and manages all types of data
Replication and high availability	More resources, licensing, and servers required	Built-in replication and redundancy features available
Application sharing and access to shared resources	File sharing requires additional administration resources	File sharing options available without additional license
Relational data storage	Needs a database server with database admin role	Offers database-as-a-service options
Distributed storage and data access	Expensive storage, networking, and compute resources needed	Azure Cosmos DB provides distributed access
Messaging and load balancing	Hardware redundancy impacts budget and resources	Azure Queue provides effective load balancing
Tiered storage	Management of tiered storage needs technology and labor skill set	Azure offers automated tiered storage of data

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.

One way to reduce latency is to provide exact copies of your service in more than one region

Application and network control are security need to be focus by the customer and microsoft in PaaS but not in SaaS.

Regardless of the deployment type, you always retain responsibility for the following items:

- Data
- Endpoints
- Accounts
- Access management

Security layer by layer:

(most depth) ----- (least depth)

Data < Application < Compute < Network < Perimeter < Identity and Access < Physical Security

Application: Ensure applications are secure and free of vulnerabilities.

Compute:

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

Networking:

- Implement secure connectivity to on-premises networks.
- Limit communication between resources.

Perimeter:

- DDoS protection
- Use perimeter firewalls to identify and alert on malicious attacks against your network.

Identity and Access:

- Use single sign-on and multi-factor authentication.

Azure Security Center is available in two tiers:

1. *Free*. Available as part of your Azure subscription, this tier is limited to assessments and recommendations of Azure resources only.
 2. *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.
 - Use Security Center for incident response. (Detect, Access, Diagnose)
 - Use Security Center recommendations to enhance security.
- *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 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.
- **Device Management.** Manage how your cloud or on-premises devices access your corporate data.

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 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.

Azure AD:

Service principals:

- An **identity** is just a thing that can be authenticated.
- 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.

Managed identities for Azure services

- 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 your organization's Active Directory (a specific organization's Active Directory instance is known as an "Active Directory 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 allowing the authenticated service secure access of other Azure resources.

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

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:

1. Encryption at rest:

- Data at rest is the data that has been stored on a physical medium and encrypted. This data could be stored on the disk of a server, data stored in a database, or data stored in a storage account.

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.

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 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*
- *Monitor access and use.*
- *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 **create certificates** in Key Vault, or import existing certificates
- You can securely **store and manage certificates** without interaction with private key material.
- You can **create a policy** that directs Key Vault to manage the life cycle of a certificate.
- You can **provide contact information** for notification about life-cycle events of expiration and renewal of certificate.
- You can automatically **renew certificates** with selected issuers - Key Vault partner x509 certificate providers / certificate authorities.

Transport Layer Security (TLS) is the basis for encryption of website data in transit. TLS uses *certificates* to encrypt and decrypt data. Certificates used in Azure are **x.509 v3** and can be signed by a trusted certificate authority, or they can be self-signed. A self-signed certificate is signed by its own creator; therefore, it is not trusted by default. However, you should only use self-signed certificates when developing and testing your cloud services.

Certificates are used in Azure for two primary purposes and are given a specific designation based on their intended use.

- **Service certificates** are used for cloud services and enable secure communication to and from the service. Service certificates, which are defined in your service definition, are automatically deployed to the VM that is running an instance of your role. To update the certificate, it's only necessary to upload a new certificate and change the thumbprint value in the service configuration file.
- **Management certificates** are used for authenticating with the management API. Many programs and tools (such as Visual Studio or the Azure SDK) use these certificates to automate configuration and deployment of various Azure services. However, these types of certificates are not related to cloud services.

Review:

- Each Azure subscription can contain multiple account administrators
- Each Azure subscription can be managed by using a Microsoft account only
- If you plan to create an Azure VM, you need Azure Blobs to store the data disks of the VM
- If you plan to map a network drive from several computers that run Windows to Azure Storage, use Azure Files
- AZ are not used to replicate data and application to multiple regions

Microsoft Security Development Lifecycle (SDL) introduces security and privacy considerations throughout all phases of the development process:

- **Provide training**
Security is everyone's job.
- **Define security requirements**
The optimal time to define the security requirements is during the initial design and planning stages.

Factors that influence security requirements include, but are not limited to:

- Legal and industry requirements
- Internal standards and coding practices
- Review of previous incidents
- Known threats
- **Define metrics and compliance reporting**
It's essential for an organization to define the minimum acceptable levels of security quality, and to hold engineering teams accountable to meeting that criteria
- **Perform threat modeling**
it allows development teams to consider, document, and discuss the security implications of designs in the context of their planned operational environment, and in a structured fashion
- **Establish design requirements**
- **Define and use cryptography standards**
- **Manage security risks from using third-party components**
- **Use approved tools**
- **Perform Static Analysis Security Testing**
To identify vulnerabilities each time the software is built or packaged
- **Perform Dynamic Analysis Security Testing**
Performing run-time verification of your fully compiled or packaged software checks functionality that is only apparent when all components are integrated and running.
- **Perform penetration testing**
- **Establish a standard incident response process**

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.

Here are some of the most common policy definitions you can apply.

Policy definition	Description
Allowed Storage Account SKUs	This policy definition has a set of conditions/rules that determine whether a storage account that is being deployed is within a set of SKU sizes. Its effect is to deny all storage accounts that do not adhere to the set of defined SKU sizes.
Allowed Resource Type	This policy definition has a set of conditions/rules to specify the resource types that your organization can deploy. Its effect is to deny all resources that are not part of this defined list.
Allowed Locations	This policy enables you to restrict the locations that your organization can specify when deploying resources. Its effect is used to enforce your geographic compliance requirements.
Allowed Virtual Machine SKUs	This policy enables you to specify a set of VM SKUs that your organization can deploy.
Not allowed resource types	Prevents a list of resource types from being deployed.

Policy assignments are inherited by all child resources. This inheritance 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.

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.

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 that can be viewed in the Azure Policy portal, or through command-line tools.

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. Initiative definitions simplify the process of managing and assigning policy definitions by grouping a set of policies into a single item.

Azure Blueprints 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. Blueprint objects are replicated to multiple Azure regions. This replication provides low latency, high availability, and consistent access to your blueprint objects, regardless of which region Blueprint deploys your resources to.

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. Resource Manager templates are 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.

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.

A policy can be included as one of many artifacts in a blueprint definition. Blueprints also support using parameters with policies and initiatives.

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

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.
- Recommended resources
- Information specific to key organizational roles
- Cross-company document search
- Direct guidance and support for when you can't find what you're looking for.

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

Compliance Manager is a workflow-based risk assessment dashboard within the Service 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. Provides a Compliance Score and provides a secure repository.

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

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.

Data tier	Description
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.

Diagnostics: *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.

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 Status** provides a global view of the health state of Azure services.
- **Service Health** provides you with a customizable dashboard that tracks the state of your Azure services in the regions where you use them
- **Resource Health** helps you diagnose and obtain support when an Azure service issue affects your resources

Review:

- Data that is copied to an Azure Storage account is maintained automatically in at least three copies
 - Even if you have Azure resources deployed to every region, you still cannot implement AZ in all the regions
 - If your VPN is using global IP address, use Local network gateways. If it is private, use virtual network gateways
 - If you plan to migrate all your network resources to Azure, you need to create a subscription first
 - Azure AD connect application to retrieve security tokens
 - In Azure AD Basic and Premium, at least 99.9% availability is guaranteed
-

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. Many resources can be moved between resource groups with some services having specific limitations or requirements to move. Resource groups can't be nested

If you delete a resource group, all resources contained within are also deleted

Since resource groups are a scope of RBAC, you can organize resources by *who* needs to administer them.

Tags are name/value pairs of text data that you can apply to resources and resource groups

A resource can have up to 50 tags. The name is limited to 512 characters for all types of resources except storage accounts, which have a limit of 128 characters. The tag value is limited to 256 characters for all types of resources. Tags aren't inherited from parent resources. Not all resource types support tags, and tags can't be applied to classic resources. Tags can be added and manipulated through the Azure portal, Azure CLI, Azure PowerShell, Resource Manager templates, and through the REST API.

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 automation. 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

Resource locks are a setting that can be applied to any resource to block modification or deletion. Resource locks can 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.

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

Factor affecting costs:

- Resource type
- Services
- Location
- 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**

Azure pricing calculator:

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 customers 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.

On the pricing calculator page, you'll see several tabs:

1. **Products.** This tab is where you'll do most of your activity. This tab has all the Azure services listed and is where you'll add or remove services to put together your estimate.
2. **Example Scenarios** This tab has several examples of infrastructure involved in common cloud-based solutions. You can add all the components of the entire scenario to estimate the cost.
3. **Saved Estimates.** This tab has all of your previously saved estimates. We'll go through this process in a moment.
4. **FAQ.** Just as it says, this tab has answers to some frequently asked questions.

We can save this estimate, so we can come back to it later and adjust it if necessary. We can also export it to Excel for further analysis or share the estimate via a URL.

Azure Advisor makes cost recommendations in the following areas:

1. **Reduce costs by eliminating unprovisioned Azure ExpressRoute circuits.**
2. **Buy reserved instances to save money over pay-as-you-go.** Advisor 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.
3. **Right-size or shutdown underutilized virtual machines.** This analysis 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.

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

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.

Azure monthly credit 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.

If you have virtual machine workloads that are static and predictable, using reserved instances is a fantastic way to potentially save up to 70 to 80 percent off the pay-as-you-go cost. The savings can be significant, depending on the VM size and duration the machine runs. The following illustration shows that using Azure reserved instances saves you up to 72 percent and using reserved instance plus Azure Hybrid Benefit saves up to 80 percent in costs.

PaaS is typically less expensive than IaaS.

The Azure Enterprise Dev/Test and Azure Pay-As-You-Go Dev/Test benefits give you several discounts, most notably for Windows workloads, eliminating license charges and billing you only at the Linux rate for virtual machines. This benefit also applies to SQL Server and any other Microsoft software that is covered under a Visual Studio subscription.

Review:

- If you delete a resource group, all the resources in the resource group will be deleted
- A resource group can contain resources from multiple Azure regions
- Advisor provides recommendations on how to reduce the cost of running Azure VM but doesn't provide recommendations on how to improve the security of an Azure AD environment
- From Azure Service Health, an administrator can create a rule to be alerted if an Azure service fails, but it cannot prevent a service from failing.
- Azure has built-in AuthN and AuthZ that provide secure access to Azure resources
- Identities stored in Azure AD, third-party cloud services and on-prem AD can be used to access Azure resources
- An Azure resource group cannot contain multiple Azure subscriptions

Important Link:

- <https://docs.microsoft.com/en-us/azure/availability-zones/az-overview>
- <https://docs.microsoft.com/en-us/azure/active-directory/identity-protection/howto-sign-in-risk-policy>

List of Azure terminology overview (not exhaustive):

- **Azure Cloud Shell** is a browser-based command-line experience for managing and developing Azure resources. Cloud Shell provides two experiences to choose from, Bash and PowerShell
- **Azure Advisor** is a free service built into Azure that provides recommendations on high availability, security, performance, operational excellence, and cost.
- **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.
- **Azure Security Center** is a monitoring service that provides threat protection across all of your services both in Azure, and on-premises.
- **Azure Active Directory** 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.
- **Azure AD Privileged Identity Management (PIM)** is an additional, paid-for offering that provides oversight of role assignments, self-service, and just-in-time role activation and Azure AD and Azure resource access reviews.
- **Azure Disk Encryption** is a capability that helps you encrypt your Windows and Linux IaaS virtual machine disks together with **Azure Key Vault**
- **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
- **Azure Key Vault** to protect our secrets.
- **Azure Information Protection** (sometimes referred to as AIP) is a cloud-based solution that helps organizations classify and optionally protect documents and emails by applying labels.
- **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 Policy** is an Azure service you use to create, assign and, manage policies. These policies enforce different rules and effects over your resources so that those resources stay compliant with your corporate standards and service level agreements. Azure Policy meets this need by evaluating your resources for noncompliance with assigned policies. Unlike RBAC, Azure Policy is a **default-allow-and-explicit-deny-system**
- **Azure Management Groups** are containers for managing access, policies, and compliance across *multiple* Azure subscriptions.
- **Azure Blueprints** enables cloud architects and central information technology groups to define a repeatable set of Azure resources that implements and adheres to an organization's standards, patterns, and requirements. It help you with auditing, traceability, and compliance of your deployments
- **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.