



AZ-900: Microsoft Azure Fundamentals

Describe Cloud Computing

1) What is Cloud Computing?

It's the delivery of computing services over the internet, which is otherwise known as the cloud. These services include servers, storage, databases, networking, software, analytics, and intelligence. Cloud computing offers faster innovation, flexible resources, and economies of scale.

2) Describe the shared responsibility model

The **Shared Responsibility Model** is a security and compliance framework that outlines the responsibilities of **cloud service providers (CSPs)** and **customers** for securing every aspect of the cloud environment, including hardware, infrastructure, endpoints, data, configurations, settings, operating system (OS), network controls and access rights.

In its simplest terms, the Shared Responsibility Model dictates that the cloud provider—such as Amazon Web Service (AWS), Microsoft Azure, or Google Cloud Platform (GCP)—must monitor and respond to security threats related to the cloud itself and its underlying infrastructure. Meanwhile, end users, including individuals and companies, are responsible for protecting data and other assets they store in any cloud

environment.

3) Define cloud models, including public, private, and hybrid

- Public cloud is cloud computing that's delivered via the internet and shared across organizations.
- Private cloud is cloud computing that is dedicated solely to your organization.
- Hybrid cloud is an environment that uses both public and private clouds.

4) Identify appropriate use cases for each cloud model

- Backup as a Service (BaaS).
- Disaster recovery as a service.
- Email.
- Virtual Desktops (VDI) / Desktop as a Service (DaaS).
- Test and Development.
- Infrastructure as a Service (IaaS).
- Private/Public/Hybrid cloud.
- Software-Defined Wide Area Networking (SD-WAN).
- Big Data Analytics.
- Software as a Service (SaaS).

5) Describe the consumption-based mode

A consumption-based pricing model is a service

provision and payment scheme in which the customer pays according to the resources used. The provider tracks how much the customer uses and then bills them for the number of services consumed. It can also be referred to as pay-as-you-go billing, metered billing or usage-based pricing.

Although relatively new in computing, consumption-based pricing is common in many traditional business types. For example, municipal utility companies, such as water and electricity services, charge consumers a fee based on the amount of the service the customer used.

Platform as a service (PaaS) and infrastructure as a service (IaaS) often also use consumption pricing to maintain cost advantage and profitability.

6) Compare cloud pricing models

a) Pay as you go

You can pay for services on Azure according to actual usage, billed per second, with no long-term commitment or upfront payments. This provides complete flexibility to increase or decrease resources as needed. Azure virtual machines (VMs) can be automatically scaled up and down using Azure's autoscaling features.

b) Reserved Instances

Azure provides Reserved Virtual Machine Instances (RVMI)—virtual machines that are pre-purchased for one or three years in a specific region. Committing to reserved instances in advance grants a discount of up to 72% compared to pay-as-you-go prices.

Azure provides the option to replace reserved instances with others during the commitment term. It also allows users to cancel reserved instances before the end of the term, but this incurs an early termination fee.

c) Spot Pricing

Azure lets you buy unused computing power at a discount of up to 90% compared to pay-as-you-go prices. However, spot instances can be interrupted on short notice, so they are considered to be suitable only for workloads that can tolerate disruptions.

Azure provides Virtual Machine Scale Sets (VMSS), an autoscaling mechanism that lets you manage groups of VMs and add spot instances automatically according to predefined policies.

Describe the benefits of using cloud services

1) Describe the benefits of high availability and scalability in the cloud

When you're deploying an application, a service, or any IT resources, it's important the resources are available when needed. High availability focuses on ensuring maximum availability, regardless of disruptions or events that may occur.

When you're architecting your solution, you'll need to account for service availability guarantees. Azure is a highly available cloud environment with uptime guarantees depending on the service. These guarantees are part of the service-level agreements (SLAs).

Some of the benefits are

- Data backup and recovery
- Load Balancing
- Clustering

What is scalability in Azure?

With the Microsoft Azure public cloud and others, scalability is essentially baked into the environment by way of programmatic controls to allow easily extending or shrinking scalability. Let's take a look at some of the specific examples of Microsoft Azure cloud scalability features and functionality to see how scalability is easily accomplished in the Azure cloud.

Microsoft Azure Cloud Scalability Features and Functionality

As opposed to the dilemma of on-premises scaling of resources requiring provisioning more hardware, with the Azure cloud environment, resources can easily be

scaled up and down depending on the customer's needs. Azure features many great capabilities related to scaling, including:

- Scaling up and down
- Scaling in and out
- Autoscaling

Azure's Platform-as-a-Service offering provides services for applications. This managed infrastructure service provided by Azure allows operations and developers to deploy applications on top of the offering without the need to worry about the underlying infrastructure.

2) Describe the benefits of reliability and predictability in the cloud

Reliability

Reliability is the ability of a system to recover from failures and continue to function. It's also one of the pillars of the Microsoft Azure Well-Architected Framework.

The cloud, by virtue of its decentralised design, naturally supports a reliable and resilient infrastructure. With a decentralised design, the cloud enables you to have resources deployed in regions around the world. With this global scale, even if one region has a catastrophic event other regions are still up and running. You can design your applications to

automatically take advantage of this increased reliability. In some cases, your cloud environment itself will automatically shift to a different region for you, with no action needed on your part. You'll learn more about how Azure leverages global scale to provide reliability later in this series.

Predictability

Predictability in the cloud lets you move forward with confidence. Predictability can be focused on performance predictability or cost predictability. Both performance and cost predictability is heavily influenced by the Microsoft Azure Well-Architected Framework. Deploy a solution that's built around this framework and you have a solution that's cost and performance prediction.

Benefits of Reliability and Predictability

Platform to design reliable applications and services

- 220 datacenters
- Over 60 regions
- 165,000 miles of fibre optic cable
- Azure services backed by a Service Level Agreement (SLA)
- Performance and availability based on SLA's

- Published pricing data, no surprise costs
- Categorise expenses by subscription, groups of services or individual services with tags

3) Describe the benefits of security and governance in the cloud

Whether you're deploying infrastructure as a service or software as a service, cloud features support governance and compliance. Things like set templates help ensure that all your deployed resources meet corporate standards and government regulatory requirements. Plus, you can update all your deployed resources to new standards as standards change. Cloud-based auditing helps flag any resource that's out of compliance with your corporate standards and provides mitigation strategies.

On the security side, you can find a cloud solution that matches your security needs. If you want maximum control of security, infrastructure as a service provides you with physical resources but lets you manage the operating systems and installed software, including patches and maintenance.

And because the cloud is intended as an over-the-internet delivery of IT resources, cloud providers are typically well suited to handle things like distributed denial of service (DDoS) attacks, making your network more robust and secure.

By establishing a good governance footprint early, you can keep your cloud footprint updated, secure, and well

managed.

4) Describe the benefits of manageability in the cloud

A major benefit of cloud computing is the manageability options. There are two types of manageability for cloud computing.

Management of the cloud

Management of the cloud speaks to managing your cloud resources. In the cloud, you can:

- a) Automatically scale resource deployment based on need.
- b) Deploy resources based on a preconfigured template, removing the need for manual configuration.
- c) Monitor the health of resources and automatically replace failing resources.
- d) Receive automatic alerts based on configured metrics, so you're aware of performance in real-time.

Management in the cloud

Management in the cloud speaks to how you're able to manage your cloud environment and resources. You can manage these:

- a) Through a web portal.
- b) Using a command line interface.
- c) Using APIs.
- d) Using PowerShell.

Describe Cloud service types

1) Describe infrastructure as a service (IaaS)

Infrastructure as a service (IaaS) is the most flexible category of cloud services, as it provides you with the maximum amount of control for your cloud resources. In an IaaS model, the cloud provider is responsible for maintaining the hardware, network connectivity (to the internet), and physical security. You're responsible for everything else: operating system installation, configuration, and maintenance; network configuration; database and storage configuration; and so on. With IaaS, you're essentially renting the hardware in a cloud data centre, but what you do with that hardware is up to you.

The shared responsibility model applies to all the cloud service types. IaaS places the largest share of responsibility with you. The cloud provider is responsible for maintaining the physical infrastructure and its access to the internet. You're responsible for installation and configuration, patching and updates, and security.

2) Describe Platform as a Service (PaaS)

Platform as a service (PaaS) is a middle ground between renting space in a data centre (infrastructure as a service) and paying for a complete and deployed

solution (software as a service). In a PaaS environment, the cloud provider maintains the physical infrastructure, physical security, and connection to the internet. They also maintain the operating systems, middleware, development tools, and business intelligence services that make up a cloud solution. In a PaaS scenario, you don't have to worry about the licensing or patching for operating systems and databases.

PaaS is well suited to provide a complete development environment without the headache of maintaining all the development infrastructure.

The shared responsibility model applies to all the cloud service types. PaaS splits the responsibility between you and the cloud provider. The cloud provider is responsible for maintaining the physical infrastructure and its access to the internet, just like in IaaS.

3) Describe Software as a Service (SaaS)

Software as a service (SaaS) is the most complete cloud service model from a product perspective. With SaaS, you're essentially renting or using a fully developed application. Email, financial software, messaging applications, and connectivity software are all common examples of a SaaS implementation. While the SaaS model may be the least flexible, it's also the easiest to get up and running. It requires the least amount of technical knowledge or expertise to fully employ.

The shared responsibility model applies to all the cloud service types. SaaS is the model that places the most responsibility on the cloud provider and the least responsibility on the user. In a SaaS environment, you're responsible for the data that you put into the system, the devices that you allow to connect to the system, and the users that have access.

4) Identify appropriate use cases for each cloud service (IaaS, PaaS, SaaS)

a) IaaS

- Lift-and-shift migration: You're standing up cloud resources similar to your on-prem datacenter, and then simply moving the things running on-prem to running on the IaaS infrastructure.
- Testing and development: You have established configurations for development and test environments that you need to rapidly replicate. You can stand up or shut down the different environments rapidly with an IaaS structure while maintaining complete control.

b) PaaS

- Development framework: PaaS provides a framework that developers can build upon to develop or customize cloud-based applications. Similar to the way you create an Excel macro, PaaS lets developers create applications using built-in software components.

- Analytics or business intelligence: Tools provided as a service with PaaS allow organizations to analyze and mine their data, finding insights and patterns and predicting outcomes to improve forecasting, product design decisions, investment returns, and other business decisions.

c) SaaS

- Email and messaging.
- Business productivity applications.
- Finance and expense tracking.

Describe Azure architecture and services

Describe the core architectural components of Azure

1) Describe Azure regional, regional pairs, and sovereign regions

Azure is a continually expanding set of cloud services that help you meet current and future business challenges. Azure gives you the freedom to build, manage, and deploy applications on a massive global network using your favourite tools and frameworks.

Regional - A region is a geographical area on the planet that contains at least one, but potentially multiple data centres 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.

When you deploy a resource in Azure, you'll often need to choose the region where you want your resource deployed.

Regional pairs - Most Azure regions are paired with another region within the same geography (such as the US, Europe, or Asia) at least 300 miles away. This approach allows for the replication of resources across geography that helps reduce the likelihood of interruptions because of events such as natural disasters, civil unrest, power outages, or physical network outages that affect an entire region. For example, if a region in a pair was affected by a natural

disaster, services would automatically failover to the other region in its region pair.

Examples of region pairs in Azure are West US paired with East US and South-East Asia paired with East Asia. Because the pair of regions are directly connected and far enough apart to be isolated from regional disasters, you can use them to provide reliable services and data redundancy.

Sovereign regions - In addition to regular regions, Azure also has sovereign regions. Sovereign regions are instances of Azure that are isolated from the main instance of Azure. You may need to use a sovereign region for compliance or legal purposes.

Azure sovereign regions include:

- a) US DoD Central, US Gov Virginia, US Gov Iowa and more: These regions are physical and logical network-isolated instances of Azure for U.S. government agencies and partners. These data centres are operated by screened U.S. personnel and include additional compliance certifications.
- b) China East, China North, and more: These regions are available through a unique partnership between Microsoft and 21Vianet, whereby Microsoft doesn't directly maintain the data centres.

2) Describe availability zones

Availability zones are physically separate data centres within an Azure region. Each availability zone is made up of one or more data centres equipped with independent power, cooling, and networking. An availability zone 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 fibre-optic networks.

You want to ensure your services and data are redundant so you can protect your information in case of failure. When you host your infrastructure, setting up your redundancy requires that you create duplicate hardware environments. Azure can help make your app highly available through availability zones. Availability zones are primarily for VMs, managed disks, load balancers, and SQL databases.

3) Describe Azure Datacenters

The physical infrastructure for Azure starts with data centres. Conceptually, the data centres are the same as large corporate data centres. They're facilities with resources arranged in racks, with dedicated power, cooling, and networking infrastructure.

As a global cloud provider, Azure has data centres around the world. However, these individual data centres aren't directly accessible. Datacenters are grouped into Azure Regions or Azure Availability Zones that are designed to help you achieve resiliency and reliability for your business-critical workloads.

The Global infrastructure site gives you a chance to interactively explore the underlying Azure infrastructure.

4) Describe Azure resources and the Resource group

Resource - A resource is the basic building block of Azure. Anything you create, provision, deploy, etc. is a resource. Virtual Machines (VMs), virtual networks, databases, cognitive services, etc. are all considered

resources within Azure.

Resource Group - Resource groups are simply groupings of resources. When you create a resource, you're required to place it into a resource group. While a resource group can contain many resources, a single resource can only be in one resource group at a time. Some resources may be moved between resource groups, but when you move a resource to a new group, it will no longer be associated with the former group. Additionally, resource groups can't be nested, meaning you can't put resource group B inside of resource group A.

Resource groups provide a convenient way to group resources. When you apply an action to a resource group, that action will apply to all the resources within the resource group. If you delete a resource group, all the resources will be deleted.

5) Describe Subscriptions

In Azure, subscriptions are management, billing, and scale unit. Similar to how resource groups are a way to logically organize resources, subscriptions allow you to logically organize your resource groups and facilitate billing.

Using Azure requires an Azure subscription. A subscription provides you with authenticated and authorized access to Azure products and services. It also allows you to provision resources. An Azure subscription links to an Azure account, which is an identity in Azure Active Directory (Azure AD) or in a directory that Azure AD trusts.

An account can have multiple subscriptions, but only

one is required. In a multi-subscription account, you can use the subscriptions to configure different billing models and apply different access-management policies.

6) Describe management groups

Resources are gathered into resource groups, and resource groups are gathered into subscriptions. If you're just starting in Azure that might seem like enough hierarchy to keep things organized. But imagine if you're dealing with multiple applications, and multiple development teams, in multiple geographies.

If you have many subscriptions, you might need a way to efficiently manage access, policies, and compliance for those subscriptions. Azure management groups provide a level of scope above subscriptions. You organize subscriptions into containers called management groups and apply governance conditions to the management groups.

7) Describe the hierarchy of Resource groups, management, and subscription groups.

You can build a flexible structure of management groups and subscriptions to organize your resources into a hierarchy for unified policy and access management. The following diagram shows an example of creating a hierarchy for governance by using management groups.

Some examples of how you could use management groups might be:

- a) Create a hierarchy that applies a policy. You could limit VM locations to the US West Region in a

group called Production. This policy will inherit all the subscriptions that are descendants of that management group and will apply to all VMs under those subscriptions. This security policy can't be altered by the resource or subscription owner, which allows for improved governance.

b) Provide user access to multiple subscriptions. By moving multiple subscriptions under a management group, you can create one Azure role-based access control (Azure RBAC) assignment on the management group. Assigning Azure RBAC at the management group level means that all sub-management groups, subscriptions, resource groups, and resources underneath that management group would also inherit those permissions. One assignment on the management group can enable users to have access to everything they need instead of scripting Azure RBAC over different subscriptions.

Describe Azure compute and networking services

1) Compare compute types, including container instances, virtual machines (VMs), and functions

2) Describe VM options, including Azure Virtual Machines, Azure Virtual Machine Scale Sets, availability sets, and Azure Virtual Desktop

With Azure Virtual Machines (VMs), you can create

and use VMs in the cloud. VMs 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 your VM. VMs are an ideal choice when you need:

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

An Azure VM gives you the flexibility of virtualization without having to buy and maintain the physical hardware that runs the VM. However, as an IaaS offering, you still need to configure, update, and maintain the VM's software.

Virtual machine scale sets let you create and manage a group of identical, load-balanced VMs. If you simply created multiple VMs with the same purpose, you'd need to ensure they were all configured identically and then set up network routing parameters to ensure efficiency. You'd also have to monitor the utilization to determine if you need to increase or decrease the number of VMs.

Instead, with virtual machine scale sets, Azure automates most of that work. Scale sets allow you to centrally manage, configure, and update a large number of VMs in minutes. The number of VM instances can automatically increase or decrease in response to demand, or you can set it to scale based on a defined schedule.

Another type of virtual machine is the Azure Virtual Desktop. Azure Virtual Desktop is a desktop and

application virtualization service that runs on the cloud. It enables you to use a cloud-hosted version of Windows from any location. Azure Virtual Desktop works across devices and operating systems and works with apps that you can use to access remote desktops or most modern browsers.

3) Describe resources required for Virtual Machine.

When you provision a VM, you'll also have the chance to pick the resources that are associated with that VM, including:

- Size (purpose, number of processor cores, and amount of RAM)

- Storage disks (hard disk drives, solid state drives, etc.)

- Networking (virtual network, public IP address, and port configuration).

4) Describe application hosting options, including the Web Apps feature of Azure App Service, containers, and virtual machines

Containers are a virtualization environment. Much like running multiple virtual machines on a single physical host, you can run multiple containers on a single physical or virtual host. Unlike virtual machines, you don't manage the operating system for a container. Virtual machines appear to be an instance of an operating system that you can connect to and manage. Containers are lightweight and designed to be created, scaled out, and stopped dynamically. It's possible to create and deploy virtual machines as application

demand increases, but containers are a lighter-weight, more agile method.

Azure Container Instances offer the fastest and simplest way to run a container in Azure; without having to manage any virtual machines or adopt any additional services. Azure Container Instances are a platform as a service (PaaS) offering. Azure Container Instances allow you to upload your containers and then the service will run the containers for you.

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

App Service includes full support for hosting web apps by 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.

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

Azure App Service is a robust hosting option that you can use to host your apps in Azure. Azure App Service

lets you focus on building and maintaining your app, and Azure focuses on keeping the environment up and running.

5) Describe virtual networking, including the purpose of Azure Virtual Networks, Azure virtual subnets, peering, Azure DNS, Azure VPN Gateway, and Azure ExpressRoute

Azure virtual networks and virtual subnets enable Azure resources, such as VMs, web apps, and databases, to communicate with each other, with users on the internet, and with your on-premises client computers. You can think of an Azure network as an extension of your on-premises network with resources that link to other Azure resources.

Azure virtual networks provide the following key networking capabilities:

- a) Isolation and segmentation
- b) Internet communications
- c) Communicate between Azure resources
- d) Communicate with on-premises resources
- e) Route network traffic
- f) Filter network traffic
- g) Connect virtual networks

The azure virtual network allows you to create multiple isolated virtual networks. When you set up a virtual network, you define a private IP address space by using either public or private IP address ranges. The IP range only exists within the virtual network and isn't internet routable. You can divide that IP address space into subnets and allocate part of the defined address

space to each named subnet.

You can link virtual networks together by using virtual network peering. Peering allows two virtual networks to connect directly to each other. Network traffic between peered networks is private and travels on the

Microsoft backbone network, never entering the public internet. Peering enables resources in each virtual network to communicate with each other.

Point-to-site virtual private network connections are from a computer outside your organization back into your corporate network. In this case, the client computer initiates an encrypted VPN connection to connect to azure virtual network.

Site-to-site virtual private networks link your on-premises VPN device or gateway to the Azure VPN gateway in a virtual network. In effect, the devices in Azure can appear as being on the local network.

The connection is encrypted and works over the internet.

Azure ExpressRoute provides dedicated private connectivity to Azure that doesn't travel over the internet. ExpressRoute is useful for environments where you need greater bandwidth and even higher levels of security.

Border Gateway Protocol (BGP) works with Azure VPN gateways, Azure Route Server, or Azure ExpressRoute to propagate on-premises BGP routes to Azure virtual networks.

6) Define public and private endpoints

Private Endpoints grant network access to specific resources behind a given service providing granular segmentation. Traffic can reach the service resource from on-premises without using public endpoints.

Multiple private link resource types support access via Private Endpoints. Resources include Azure PaaS services and your own Private Link Service. It's a one-to-many relationship.

A Service Endpoint remains a publicly routable IP address. A Private Endpoint is a private IP in the address space of the virtual network where the private endpoint is configured.

A Private Link service receives connections from multiple Private Endpoints. A private endpoint connects to one Private Link Service.

Describe Azure Storage services

1) Compare Azure storage services

Azure Blobs is an immensely scalable object store for text and binary data.

Azure Files are organised file shares for cloud or on-premises deployments.

Azure Queue is a messaging store for consistent messaging between application components.

Azure Tables are NoSQL stores for schema-less storage of structured data.

Azure Disks are block-level storage volumes for Azure Virtual Machines.

2) Describe storage Tiers

Data stored in the cloud can grow at an exponential pace. To manage costs for your expanding storage needs, it's helpful to organise your data based on attributes like frequency of access and planned retention period. Data stored in the cloud can be handled differently based on how its generated, processed and accessed over its lifetime. Some data is actively accessed and modified throughout its lifetime. Some data is accessed frequently early in its lifetime, with access dropping drastically as the data ages. Some data remains idle in the cloud and is rarely if ever, accessed after its stored. To accommodate these different access needs, Azure provides several access tiers, which you can use to balance your storage costs with your access needs.

- a) Hot access tier
- b) Cool access tier
- c) Archive access tier

3) Describe redundancy option

Azure Storage always stores multiple copies of your data so that it's protected from planned and unplanned events such as transient hardware failures, network or power outages, and natural disasters. Redundancy ensures that your storage account meets its availability and durability targets even in the face of failures.

When deciding which redundancy option is best for your scenario, consider the tradeoffs between lower costs and higher availability. The factors that help determine which redundancy option you should choose to include:

- a) How your data is replicated in the primary region.
- b) Whether your data is replicated to a second region that is geographically distant to the primary region, to protect against regional disasters.
- c) Whether your application requires read access to the replicated data in the secondary region if the primary region becomes unavailable.

Data in an Azure Storage account is always replicated three times in the primary region. Azure Storage offers two options for how your data is replicated in the primary region, locally redundant storage (LRS) and zone-redundant storage (ZRS).

4) Describe storage account options and storage types.

A storage account provides a unique namespace for your Azure Storage data that's accessible from anywhere in the world over HTTP or HTTPS. Data in this account is secure, highly available, durable, and massively scalable.

When you create your storage account, you'll start by picking the storage account type. The type of account determines the storage services and redundancy options and has an impact on the use cases. Below is a list of redundancy options that will be covered later in this module:

- a) Locally redundant storage (LRS)
- b) Geo-redundant storage (GRS)
- c) Read-access geo-redundant storage (RA-GRS)
- d) Zone-redundant storage (ZRS)
- e) Geo-zone-redundant storage (GZRS)

5) Identify options for moving files, including **AzCopy**, **Azure Storage Explorer**, and **Azure File Sync**.

In addition to large-scale migration using services like Azure Migrate and Azure Data Box, Azure also has tools designed to help you move or interact with individual files or small file groups.

AzCopy is a command-line utility that you can use to copy blobs or files to or from your storage account. With AzCopy, you can upload files, download files, copy files between storage accounts, and even synchronise files.

Azure Storage Explorer is a standalone app that provides a graphical interface to manage files and blobs in your Azure Storage Account. It works on Windows, macOS, and Linux operating systems and uses AzCopy on the backend to perform all of the file and blob management tasks.

Azure File Sync is a tool that lets you centralize your file shares in Azure Files and keep the flexibility, performance, and compatibility of a Windows file server. It's almost like turning your Windows file server into a miniature content delivery network. Once you install Azure File Sync on your local Windows server, it will automatically stay bi-directionally synced with your files in Azure.

AzCopy can even be configured to work with other cloud providers to help move files back and forth between clouds.

With Storage Explorer, you can upload to Azure, download from Azure, or move between storage accounts.

With **Azure File Sync**, you can:

- a) Use any protocol that's available on Windows Server to access your data locally, including SMB, NFS, and FTPS.
- b) Have as many caches as you need across the world.
- c) Replace a failed local server by installing Azure File Sync on a new server in the same data centre.

6) Describe migration options, including Azure Migrate and Azure Data Box.

Now that you understand the different storage options within Azure, it's important to also understand how to get your data and information into Azure. Azure supports both real-time migration of infrastructure, applications, and data using Azure Migrate as well as asynchronous migration of data using Azure Data Box.

Azure Migrate is a service that helps you migrate from an on-premises environment to the cloud. Azure Migrate functions as a hub to help you manage the assessment and migration of your on-premises data centre to Azure. It provides the following:

- a) Unified migration platform
- b) Range of tools
- c) Assessment and migration

Azure Data Box is a physical migration service that helps transfer large amounts of data in a quick, inexpensive, and reliable way. The secure data transfer is accelerated by shipping you a proprietary Data Box storage device that has a maximum usable storage capacity of 80 terabytes. The Data Box is transported to

and from your data centre via a regional carrier. A rugged case protects and secures the Data Box from damage during transit.

You can order the Data Box device via the Azure portal to import or export data from Azure.

Describe Azure identity, access and security

1) Describe directory services in Azure, including Azure Active Directory (Azure AD) and Azure Active Directory Domain Services (Azure AD DS).

Azure Active Directory (Azure AD) is a directory service that enables you to sign in and access both Microsoft cloud applications and the cloud applications that you develop. Azure AD can also help you maintain your on-premises Active Directory deployment. For on-premises environments, Active Directory running on Windows Server provides an identity and access management service that's managed by your organization. Azure AD is Microsoft's cloud-based identity and access management service.

Azure Active Directory Domain Services (Azure AD DS) is a service that provides managed domain services such as domain join, group policy, lightweight directory access protocol (LDAP), and Kerberos/NTLM authentication. Just like Azure AD lets you use directory services without having to maintain the infrastructure supporting it, with Azure AD DS, you get the benefit of domain services without the need to deploy, manage, and patch domain controllers (DCS) in the cloud. When you create an Azure AD DS managed domain,

you define a unique namespace. This namespace is the domain name. Two Windows Server domain controllers are then deployed into your selected Azure region. This deployment of DCs is known as a replica set.

2) Describe authentication methods in Azure, including single sign-on (SSO), multi-factor authentication, and password-less.

Authentication is the process of establishing the identity of a person, service, or device. It requires the person, service, or device to provide some type of credential to prove who they are. Authentication is like presenting an ID when you're travelling. It doesn't confirm that you're ticketed, it just proves that you're who you say you are.

Single sign-on (SSO) enables a user to sign in one time and use that credential to access multiple resources and applications from different providers. For SSO to work, the different applications and providers must trust the initial authenticator.

More identities mean more passwords to remember and change. Password policies can vary among applications. As complexity requirements increase, it becomes increasingly difficult for users to remember them.

Multi-factor authentication is the process of prompting a user for an extra form (or factor) of identification during the sign-in process. MFA helps protect against a password compromise in situations where the password was compromised but the second factor wasn't.

Think about how you sign into websites, email, or online services. After entering your username and password,

have you ever needed to enter a code that was sent to your phone? If so, you've used multifactor authentication to sign in.

Features like MFA are a great way to secure your organization, but users often get frustrated with the additional security layer on top of having to remember their passwords. People are more likely to comply when it's easy and convenient to do so. **Password-less** authentication methods are more convenient because the password is removed and replaced with something you have, plus something you are, or something you **know**. **Password-less** authentication needs to be set up on a device before it can work.

3) Describe external identities and guest access in Azure.

An **external identity** is a person, device, service, etc. that is outside your organisation. Azure AD External Identities refers to all the ways you can securely interact with users outside of your organisation. If you want to collaborate with partners, distributors, suppliers, or vendors, you can share your resources and define how your internal users can access external organisations. External identities may sound similar to single sign-on. With External Identities, external users can "bring their own identities."

With Azure Active Directory (Azure AD), you can easily enable collaboration across organisational boundaries by using the Azure AD B2B feature. Guest users from other tenants can be invited by administrators or by other users.

You also can easily ensure that guest users have the

appropriate access. You can ask the guests themselves or a decision maker to participate in an access review and rectify (or attest) the guests' access.

4) Describe Azure AD Conditional Access.

Conditional Access is a tool that Azure Active Directory uses to allow (or deny) access to resources based on identity signals. These signals include who the user is, where the user is, and what device the user is requesting access from.

Conditional Access helps IT, administrators:

- Empower users to be productive wherever and whenever.

- Protect the organisation's assets.

Conditional Access also provides a more granular multi-factor authentication experience for users. For example, a user might not be challenged for a second authentication factor if they're at a known location. During sign-in, Conditional Access collects signals from the user, makes decisions based on those signals, and then enforces that decision by allowing or denying the access request or challenging for a multi-factor authentication response.

5) Describe Azure role-based access control (RBAC).

When you have multiple IT and engineering teams, how can you control what access they have to the resources in your cloud environment? The principle of least privilege says you should only grant access up to the level needed to complete a task. If you only need

read access to a storage blob, then you should only be granted read access to that storage blob. However, managing that level of permissions for an entire team would become tedious. Instead of defining the detailed access requirements for each individual, and then updating access requirements when new resources are created or new people join the team, Azure enables you to control access through Azure role-based access control (Azure RBAC). Azure provides built-in roles that describe common access rules for cloud resources. You can also define your roles. Each role has an associated set of access permissions that relate to that role.

6) Describe the concept of Zero Trust.

Zero Trust is a security model that assumes the worst-case scenario and protects resources with that expectation. Zero Trust assumes breach at the outset and then verifies each request as though it originated from an uncontrolled network.

To address this new world of computing, Microsoft highly recommends the Zero Trust security model, which is based on these guiding principles:

- a) **Verify explicitly** - Always authenticate and authorize based on all available data points.
- b) **Use least privilege access** - Limit user access with Just-In-Time and Just-Enough-Access (JIT/JEA), risk-based adaptive policies, and data protection.

c) **Assume breach** - Minimise blast radius and segment access. Verify end-to-end encryption. Use analytics to get visibility, drive threat detection, and improve defences.

7) Describe the purpose of the defence-in-depth model

The objective of defence-in-depth is to protect information and prevent it from being stolen by those who aren't authorized to access it.

A defence-in-depth strategy uses a series of mechanisms to slow the advance of an attack that aims at acquiring unauthorized access to data.

Layers of defence-in-depth :

- a. Physical Security.
- b. Identity and Access.
- c. Perimeter.
- d. Network.
- e. Compute.
- f. Application.
- g. Data.

8) Describe the purpose of Microsoft Defender for Cloud.

Defender for Cloud is a monitoring tool for security posture management and threat protection. It monitors your cloud, on-premises, hybrid, and multi-cloud environments to provide guidance and notifications aimed at strengthening your security posture.

Defender for Cloud provides the tools needed to harden your resources, track your security posture, protect

against cyber attacks, and streamline security management. Deployment of Defender for Cloud is easy, it's already natively integrated into Azure.

- a) Protection everywhere you're deployed.
- b) Assess, Secure, and Defend.
- c) Secure.
- d) Defend.

Describe Azure management and governance

Describe cost management in Azure

1) Describe factors that can affect costs in Azure.

Azure shifts development costs from the capital expense (CapEx) of building out and maintaining infrastructure and facilities to an operational expense (OpEx) of renting infrastructure as you need it, whether it's computed, storage, networking, and so on.

That OpEx cost can be impacted by many factors. Some of the impacting factors are:

- a) Resource type.
- b) Consumption.
- c) Maintenance.
- d) Geography.
- e) Subscription type.
- f) Azure Marketplace.

2) Compare the Pricing calculator and the Total Cost of Ownership (TCO) calculator.

The pricing calculator and the total cost of ownership (TCO) calculator are two calculators that help you understand potential Azure expenses. Both calculators are accessible from the internet, and both calculators allow you to build out a configuration. However, the two calculators have very different purposes.

The **pricing calculator** is designed to give you an estimated cost for provisioning resources in Azure. You

can get an estimate for individual resources, build out a solution, or use an example scenario to see an estimate of the Azure spend. The pricing calculator's focus is on the cost of provisioned resources in Azure.

The **TCO calculator** is designed to help you compare the costs for running an on-premises infrastructure compared to an Azure Cloud infrastructure. With the TCO calculator, you enter your current infrastructure configuration, including servers, databases, storage, and outbound network traffic. The TCO calculator then compares the anticipated costs for your current environment with an Azure environment supporting the same infrastructure requirements. With the TCO calculator, you enter your configuration, add in assumptions like power and IT labour costs, and are presented with an estimation of the cost difference to run the same environment in your current data centre or Azure.

3) Describe the Azure Cost management and Billing tool.

Cost Management provides the ability to quickly check Azure resource costs, create alerts based on resource spending, and create budgets that can be used to automate the management of resources.

Cost analysis is a subset of Cost Management that provides a quick visual for your Azure costs. Using cost analysis, you can quickly view the total cost in a variety of different ways, including by billing cycle, region, resource, and so on.

You use cost analysis to explore and analyze your organizational costs. You can view aggregated costs by

an organization to understand where costs are accrued and to identify spending trends. And you can see accumulated costs over time to estimate monthly, quarterly, or even yearly cost trends against a budget. Cost alerts provide a single location to quickly check on all of the different alert types that may show up in the Cost Management service. The three types of alerts that may show up are:

- a) Budget alerts.
- b) Credit alerts.
- c) Department spending quota alerts.

A budget is where you set a spending limit for Azure.

You can set budgets based on a subscription, resource

group, service type, or other criteria. When you set a budget, you will also set a budget alert. When the budget hits the budget alert level, it will trigger a budget alert that shows up in the cost alerts area.

4) Describe the purpose of Tags.

As your cloud usage grows, it's increasingly important to stay organised. A good organisation strategy helps you understand your cloud usage and can help you manage costs.

One way to organise related resources is to place them in their subscriptions. You can also use resource groups to manage related resources. Resource tags are another way to organise resources. Tags provide extra information, or metadata, about your resources. This metadata is useful for:

- a) Resource management.
- b) Cost management and optimisation.
- c) Operations management.
- d) Security.
- e) Governance and regulatory compliance.
- f) Workload optimisation and automation.

Describe features and tools in Azure for governance and compliance

1) Describe the purpose of Azure blueprints.

What happens when your cloud starts to grow beyond just one subscription or environment? How can you scale the configuration of features? How can you enforce settings and policies in new subscriptions? Azure Blueprints lets you standardise cloud subscription or environment deployments. Instead of having to configure features like Azure Policy for each new subscription, with Azure Blueprints, you can define repeatable settings and policies that are applied as new subscriptions are created. Need a new test/dev environment? Azure Blueprints lets you deploy a new Test/Dev environment with security and compliance settings already configured. In this way, development teams can rapidly build and deploy new environments with the knowledge that they're building within organisational requirements.

Azure Blueprints are version-able, allowing you to create an initial configuration and then make updates later on and assign a new version to the update. With versioning, you can make small updates and keep track of which deployments used which configuration set.

2) Describe the purpose of Azure policy.

How do you ensure that your resources stay compliant? Can you be alerted if a resource's configuration has changed?

Azure Policy is a service in Azure that enables you to create, assign, and manage policies that control or audit your resources. These policies enforce different rules across your resource configurations so that those configurations stay compliant with corporate standards. Azure Policy enables you to define both individual policies and groups of related policies, known as initiatives. Azure Policy evaluates your resources and highlights resources that aren't compliant with the policies you've created. Azure Policy can also prevent non-compliant resources from being created. An Azure Policy initiative is a way of grouping related policies together. The initiative definition contains all of the policy definitions to help track your compliance state for a larger goal.

3) Describe the purpose of resource locks.

A resource lock prevents resources from being accidentally deleted or changed.

Even with Azure role-based access control (Azure RBAC) policies in place, there's still a risk that people with the right level of access could delete critical cloud resources. Resource locks prevent resources from being deleted or updated, depending on the type of lock. Resource locks can be applied to individual resources, resource groups, or even an entire subscription. Resource locks are inherited, meaning that if you place

a resource lock on a resource group, all of the resources within the resource group will also have the resource lock applied.

There are two types of resource locks, one that prevents users from deleting and one that prevents users from changing or deleting a resource.

a) Delete means authorized users can still read and modify a resource, but they can't delete the resource.

b) ReadOnly means authorized users can read a resource, but they can't delete or update the resource. Applying this lock is similar to restricting all authorised users to the permissions granted by the Reader role.

4) Describe the purpose of the Service Trust Portal.

The Microsoft Service Trust Portal is a portal that provides access to various content, tools, and other resources about Microsoft security, privacy, and compliance practices.

The Service Trust Portal contains details about Microsoft's implementation of controls and processes that protect our cloud services and the customer data therein. To access some of the resources on the Service Trust Portal, you must sign in as an authenticated user with your Microsoft cloud services account (Azure Active Directory organisation account). You'll need to review and accept the Microsoft non-disclosure agreement for compliance materials.

The Service Trust Portal features and content are accessible from the main menu. The categories on the main menu are:

- a) Service Trust Portal.
- b) Trust Documents.
- c) Industries and Regions.
- d) Trust Center.
- e) Resources.
- f) My Library.

Describe features and tools for managing and deploying Azure resources

1) Describe Azure Portal.

The Azure portal is a web-based, unified console that provides an alternative to command-line tools. With the Azure portal, you can manage your Azure subscription by using a graphical user interface. You can:

- a) Build, manage, and monitor everything from simple web apps to complex cloud deployments.
- b) Create custom dashboards for an organized view of resources.
- c) Configure accessibility options for an optimal experience.

The Azure portal is designed for resiliency and continuous availability. It maintains a presence in every Azure data centre. This configuration makes the Azure portal resilient to individual data centre failures and avoids network slowdowns by being close to users.

2) Describe Azure Cloud Shell, including Azure CLI and Azure PowerShell.

Azure Cloud Shell is a browser-based shell tool that

allows you to create, configure, and manage Azure resources using a shell. Azure Cloud Shell supports both Azure PowerShell and the Azure Command Line Interface (CLI), which is a Bash shell.

You can access Azure Cloud Shell via the Azure portal by selecting the Cloud Shell icon:

Azure PowerShell is a shell with which developers, DevOps, and IT professionals can run commands called command-lets (cmdlets). These commands call the Azure REST API to perform management tasks in Azure. Cmdlets can be run independently to handle one-off changes, or they may be combined to help orchestrate complex actions such as:

- a) The routine setup, teardown, and maintenance of a single resource or multiple connected resources.
- b) The deployment of an entire infrastructure, which might contain dozens or hundreds of resources, from imperative code.

The Azure CLI is functionally equivalent to Azure PowerShell, with the primary difference being the syntax of commands. While Azure PowerShell uses PowerShell commands, the Azure CLI uses Bash commands. The Azure CLI provides the same benefits of handling discrete tasks or orchestrating complex operations through code. It's also installable on Windows, Linux, and Mac platforms, as well as through Azure Cloud Shell.

3) Describe the purpose of Azure Arc.

Managing hybrid and multi-cloud environments can rapidly get complicated. Azure provides a host of tools to provision, configure, and monitor Azure resources.

What about the on-premises resources in a hybrid configuration or the cloud resources in a multi-cloud configuration?

In utilizing Azure Resource Manager (ARM), Arc lets you extend your Azure compliance and monitoring to your hybrid and multi-cloud configurations. Azure Arc simplifies governance and management by delivering a consistent multi-cloud and on-premises management platform.

Azure Arc provides a centralized, unified way to:

- a) Manage your entire environment together by projecting your existing non-Azure resources into an ARM.
- b) Manage multi-cloud and hybrid virtual machines, Kubernetes clusters, and databases as if they are running in Azure.
- c) Use familiar Azure services and management capabilities, regardless of where they live.

Currently, Azure Arc allows you to manage the following resource types hosted outside of Azure:

- a) Servers
- b) Kubernetes clusters
- c) Azure data services
- d) SQL Server
- e) Virtual machines (preview)

4) Describe Azure Resource Manager and Azure Resource Manager templates (ARM templates).

Azure Resource Manager (ARM) is the deployment and management service for Azure. It provides a management layer that enables you to create, update, and delete resources in your Azure account. Anytime

you do anything with your Azure resources, ARM is involved.

When a user sends a request from any of the Azure tools, APIs, or SDKs, ARM receives the request. ARM authenticates and authorizes the request. Then, ARM sends the request to the Azure service, which takes the requested action. You see consistent results and capabilities in all the different tools because all requests are handled through the same API.

Infrastructure as code is a concept where you manage your infrastructure as lines of code. Leveraging Azure Cloud Shell, Azure PowerShell, or the Azure CLI are some examples of using code to deploy cloud infrastructure. ARM templates are another example of infrastructure as code at work.

By using ARM templates, you can describe the resources you want to use in a declarative JSON format. With an ARM template, the deployment code is verified before any code is run. This ensures that the resources will be created and connected correctly. The template then orchestrates the creation of those resources in parallel.

Describe monitoring tools in Azure

1) Describe the purpose of Azure Advisor.

Azure Advisor evaluates your Azure resources and makes recommendations to help improve reliability, security, and performance, achieve operational excellence, and reduce costs. Azure Advisor is designed to help you save time on cloud optimization. The

recommendation service includes suggested actions you can take right away, postpone, or dismiss.

When you're in the Azure portal, the Advisor dashboard displays personalized recommendations for all your subscriptions. You can use filters to select recommendations for specific subscriptions, resource groups, or services. The recommendations are divided into five categories:

- a) Reliability.
- b) Security.
- c) Performance.
- d) Operational Excellence.
- e) Cost.

2) Describe Azure Service Health.

Microsoft Azure provides a global cloud solution to help you manage your infrastructure needs, reach your customers, innovate, and adapt rapidly. Knowing the status of the global Azure infrastructure and your resources could seem like a daunting task. Azure Service Health helps you keep track of Azure resources, both your specifically deployed resources and the overall status of Azure.

Service Health provides a narrower view of Azure services and regions. It focuses on the Azure services and regions you're using. This is the best place to look for service-impacting communications about outages, planned maintenance activities, and other health advisories because the authenticated Service Health experience knows which services and resources you currently use. You can even set up Service Health alerts to notify you when service issues, planned maintenance,

or other changes may affect the Azure services and regions you use.

3) Describe Azure Monitor, including Log Analytics, Azure Monitor alerts, and Application Insights

Azure Monitor is a platform for collecting data on your resources, analyzing that data, visualizing the information, and even acting on the results. Azure Monitor can monitor Azure resources, your on-premises resources, and even multi-cloud resources like virtual machines hosted with a different cloud provider.

Azure Log Analytics is the tool in the Azure portal where you'll write and run log queries on the data gathered by Azure Monitor. Log Analytics is a robust tool that supports both simple, and complex queries and data analysis. You can write a simple query that returns a set of records and then use features of Log Analytics to sort, filter, and analyze the records. You can write an advanced query to perform statistical analysis and visualize the results in a chart to identify a particular trend.

Azure Monitor Alerts are an automated way to stay informed when Azure Monitor detects a threshold being crossed. You set the alert conditions, and the notification actions, and then Azure Monitor Alerts notify when an alert is triggered. Depending on your configuration, Azure Monitor Alerts can also attempt corrective action. Azure Monitor Alerts use action groups to configure who to notify and what action to take. An action group is simply a collection of notifications and action preferences that you associate with one or multiple alerts.

Application Insights, an Azure Monitor feature, monitors your web applications. Application Insights is capable of monitoring applications that are running in Azure, on-premises, or a different cloud environment. There are two ways to configure Application Insights to help monitor your application. You can either install an SDK in your application, or you can use the Application Insights agent. The Application Insights agent is supported in C#.NET, VB.NET, Java, JavaScript, Node.js, and Python. Not only does Application Insights help you monitor the performance of your application, but you can also configure it to periodically send synthetic requests to your application.