

DUMPSBOSS.

Microsoft Azure Fundamentals

Microsoft AZ-900

Version Demo

Total Demo Questions: 20

Total Premium Questions: 938

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

In the infrastructure as a service (IaaS) cloud service model, which two components are the responsibility of the cloud service provider? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. the configuration and maintenance of storage.
- B. the installation and configuration of the operating system.
- C. maintaining the hardware.
- D. the network configuration.
- E. physical security of the datacenter infrastructure .

ANSWER: C E

Explanation:

In the IaaS shared responsibility model, the cloud provider is responsible for the underlying physical environment that makes the service possible. That includes the physical datacenter facilities and controls (for example, building access, surveillance, and environmental protections) as well as the physical hardware (servers, storage hardware, and networking equipment) and keeping that hardware operational. Therefore, **maintaining the hardware** and **physical security of the datacenter infrastructure** are provider responsibilities.

By contrast, in IaaS the customer is responsible for what they deploy on top of that infrastructure. The customer installs and configures the operating system, applies OS patches, and manages guest configuration. The customer is also responsible for configuring their virtual network resources (VNets, subnets, NSGs, route tables, etc.) and for configuring and managing their storage at the logical level (for example, choosing disk types, setting up storage accounts, access controls, encryption settings, backups). The provider ensures the underlying storage platform and fabric are available, but not the customer's configuration choices.

For more details, see Microsoft's shared responsibility guidance: [Shared responsibility in the cloud](#) and the IaaS overview: [Azure compute decision tree \(IaaS context\)](#).

QUESTION NO: 2

Your organization intends to migrate various servers to Azure. According to the company's compliance policy, a server called FinServer must be isolated on a distinct network segment. You are considering which Azure services will fulfill the compliance policy requirements. Which Azure solution would you recommend to meet these needs?

- A. a resource group for FinServer and another resource group for all the other servers
- B. a virtual network for FinServer and another virtual network for all the other servers
- C. a VPN for FinServer and a virtual network gateway for each other server
- D. one resource group for all the servers and a resource lock for FinServer

ANSWER: B

Explanation:

To meet a requirement like “FinServer must be isolated on a distinct network segment,” the most appropriate Azure construct is network isolation. In Azure, the fundamental network boundary is the virtual network (VNet). Resources placed in different VNets are separated at the network layer by default, and any connectivity between them must be explicitly configured (for example, via VNet peering, VPN, ExpressRoute, or other routing). Therefore, placing FinServer in its own VNet and placing the other servers in a different VNet best satisfies the idea of a distinct network segment for compliance.

Option A and D use resource groups and locks, which are management/governance features (RBAC, lifecycle, protection from deletion) and do not create network segmentation. Option C describes VPN and virtual network gateways, which are used to connect networks (on-premises-to-Azure or VNet-to-VNet) rather than to isolate a single server; it also doesn't map correctly because gateways are deployed per VNet, not “for each server.”

References: [Microsoft Docs: Virtual network overview](#), [Microsoft Docs: VNet peering overview](#)

QUESTION NO: 3

You have 50 virtual machines hosted on-premises and 50 virtual machines hosted in Azure. The onpremises virtual machines and the Azure virtual machines connect to each other.

Which type of cloud model is this?

- A. hybrid.
- B. public.
- C. private .

ANSWER: A

Explanation:

This scenario describes a **hybrid cloud** model. A hybrid cloud combines resources running in a private environment (in this case, the on-premises virtual machines) with resources running in a public cloud provider (the Azure virtual machines), and crucially, it includes connectivity and integration between the two environments. The question explicitly states that the on-premises VMs and Azure VMs connect to each other, which is a hallmark of hybrid cloud deployments (often implemented using VPN or ExpressRoute, plus shared identity and management where needed).

The **public cloud** option is incorrect because the workload is not entirely hosted in a public cloud; part of it remains on-premises. The **private cloud** option is also incorrect because private cloud refers to cloud resources used exclusively by a single organization, typically hosted on-premises or in a dedicated environment, without relying on a public cloud component. Here, Azure is a public cloud service, so the overall model cannot be purely private.


References: [Microsoft Learn: Cloud models \(public, private, hybrid\)](#), [Microsoft Learn: Hybrid architecture style](#).

QUESTION NO: 4 - (DRAG DROP)

Match the Azure storage services to the appropriate descriptions.

To answer, drag the appropriate storage service from the column on the left to its description on the right. Each service may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.



ANSWER:



Explanation:

The correct way to think about this question is to map each description to the Azure storage feature it's most closely associated with.

“Used for reliable messaging between applications” maps to **Azure Queue Storage**. Queue Storage is designed for simple, durable message queuing so components can communicate asynchronously (for example, a web front end dropping work items into a queue and a background worker processing them later). This is exactly the “reliable messaging between applications” scenario. See: [Azure Queue Storage overview](#).

“Can be accessed as a network share from a Windows device” maps to **Azure Files**. Azure Files provides fully managed file shares in the cloud that support the SMB protocol, which Windows can mount as a standard network drive/share (just like an on-prem file server). That’s why it’s the right match for “network share from a Windows device.” See: [Azure Files overview](#).

“Can be configured to use the Archive access tier” maps to **Azure Blob storage**. Blob Storage supports access tiers (Hot, Cool, and Archive) to optimize cost based on how frequently data is accessed. Archive is specifically for long-term retention with the lowest storage cost and higher retrieval latency, and it’s a Blob Storage capability. See: [Blob Storage access tiers](#).

Because each of the three descriptions aligns cleanly with these services, the provided drag-and-drop selections are correct.

QUESTION NO: 5 - (DRAG DROP)

Match the cloud computing benefits to the appropriate descriptions.

To answer, drag the appropriate benefit from the column on the left to its description on the right. Each benefit may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Benefits

Disaster recovery
Geo-distribution
High availability
Scalability

Answer Area

	Increase the compute capacity of apps in the cloud.
	Provide a continuous user experience with no apparent downtime.
	Ensure that users always have the best experience by deploying apps to all the regions where there are users.

ANSWER:

Benefits

Disaster recovery
Geo-distribution
High availability
Scalability

Answer Area

Scalability	Increase the compute capacity of apps in the cloud.
High availability	Provide a continuous user experience with no apparent downtime.
Geo-distribution	Ensure that users always have the best experience by deploying apps to all the regions where there are users.

Explanation:

Let's match each description to the cloud benefit it's really talking about. When the prompt says **"Increase the compute capacity of apps in the cloud"**, that's describing the ability to add more CPU/RAM/instances as demand grows (and potentially reduce them when demand drops). In Azure fundamentals terms, that benefit is **scalability** (and often elasticity, but "scalability" is the option provided). Microsoft describes scalability as the ability to increase or decrease resources to meet demand.

The statement **"Provide a continuous user experience with no apparent downtime"** is the classic goal of **high availability**. High availability focuses on keeping an application up and accessible through redundancy, failover, and resilient design so users don't experience outages (or experience minimal interruption). This aligns with Microsoft's definition of high availability as ensuring maximum uptime and availability.

The final description, **"Ensure that users always have the best experience by deploying apps to all the regions where there are users"**, is about placing your application closer to your users geographically to reduce latency and improve responsiveness. That benefit is **geo-distribution** (sometimes discussed as global distribution). Azure supports this by letting you deploy to multiple regions worldwide so users can connect to the nearest region.

Although **disaster recovery** is a real cloud benefit, it's specifically about recovering from major failures (region-wide outages, catastrophic events) using backups, replication, and recovery plans. None of the three descriptions explicitly focuses on recovery after a disaster, so it isn't used here.

References: [Microsoft Cloud Adoption Framework – Resiliency \(availability concepts\)](#), [Microsoft Azure Well-Architected Framework – Scalability overview](#), [Azure Well-Architected Framework – Performance efficiency \(includes latency and regional distribution considerations\)](#).

QUESTION NO: 6

CORRECT TEXT

Azure distributed denial of service (DDoS) protection is an example of protection that is implemented at the _____).

A. networkinglayer

ANSWER: A

Explanation:

Azure DDoS Protection is implemented at the networking layer. In Azure, DDoS Protection is a network security service designed to help protect public endpoints (such as public IP addresses associated with Azure Virtual Network resources) from volumetric and protocol-based DDoS attacks. Because it operates by monitoring and mitigating malicious traffic patterns targeting network resources, it's considered a control at the network layer rather than something implemented at the application, identity, or physical datacenter layer. This aligns with the AZ-900 concept of "defense in depth," where different security controls map to different layers, and DDoS mitigation is a classic example of a network-layer protection.

The provided option "networkinglayer" is therefore correct. No additional options are needed. For official details, see the Azure DDoS Protection overview and the defense-in-depth model in Azure security documentation: <https://learn.microsoft.com/en-us/azure/ddos-protection/ddos-protection-overview> and <https://learn.microsoft.com/en-us/azure/security/fundamentals/defense-in-depth>.

QUESTION NO: 7

A support engineer plans to perform several Azure management tasks by using the Azure CLI.

You install the CLI on a computer.

You need to tell the support engineer which tools to use to run the CLI.

Which two tools should you instruct the support engineer to use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Command Prompt
- B. Azure Resource Explorer
- C. Windows PowerShell
- D. Windows Defender Firewall
- E. Network and Sharing Center

ANSWER: A C

Explanation:

Azure CLI is a command-line tool, so it needs to be executed from a shell/terminal environment. On Windows, two common shells that can run the `az` command after installation are **Command Prompt** and **Windows PowerShell**. Both provide an interactive command-line experience where you can authenticate (for example, `az login`) and run management commands against Azure resources.

Azure Resource Explorer is a web-based tool for viewing and interacting with Azure Resource Manager resources via the browser; it is not a shell for running Azure CLI commands. **Windows Defender Firewall** and **Network and Sharing Center** are Windows networking/security configuration utilities and do not provide a command execution environment for Azure CLI.

Microsoft's documentation for installing and using Azure CLI highlights that you run Azure CLI commands from a command-line environment (such as PowerShell or cmd on Windows). For more details, see [Install the Azure CLI](#) and general usage guidance at [Get started with Azure CLI](#).

QUESTION NO: 8

You are planning to deploy a website on Microsoft Azure, which will be accessed by users globally and will host large video files. Which Azure feature should you recommend to ensure the best video playback experience for users worldwide?

- A. an application gateway
- B. an Azure ExpressRoute circuit
- C. a content delivery network (CDN)
- D. an Azure Traffic Manager profile

ANSWER: C

Explanation:

The best choice is **Azure Content Delivery Network (CDN)** because it's designed to deliver high-bandwidth, latency-sensitive content (like large video files) to users around the world by caching content at *edge* locations close to end users. When users request a video, the CDN can serve it from the nearest point of presence (POP), reducing round-trip time, improving throughput, and providing smoother playback compared to always pulling the file from a single origin region.

Why the other options are wrong: **Application Gateway** is a Layer 7 load balancer/WAF for web apps, but it doesn't provide global edge caching for large media. **ExpressRoute** is a private connectivity option between on-premises and Azure; it doesn't help internet users worldwide stream video faster. **Traffic Manager** provides DNS-based global routing/failover to endpoints, but it doesn't cache content at the edge—users could still be routed to a distant origin, and video delivery performance won't match a CDN.

References: [Azure Content Delivery Network \(CDN\) overview](#), [Traffic Manager overview](#).

QUESTION NO: 9

What are two characteristics of using a public cloud? Each correct answer constitutes a complete solution.

Note: Each correct selection is worth one point.

- A. dedicated hardware
- B. unsecured connections
- C. limited storage
- D. metered pricing
- E. self-service management

ANSWER: D E

Explanation:

Two key characteristics of the public cloud are **metered (pay-as-you-go) pricing** and **self-service management**. In a public cloud, you typically pay only for the resources you consume (for example, per second/minute/hour for compute, per GB for storage, per request for some services). This aligns with the cloud's consumption-based model and helps avoid large upfront capital expenditures. Public cloud services are also provisioned and managed through self-service portals, APIs, or automation tools: customers deploy and configure their own resources (VMs, web apps, databases), while the cloud provider operates and maintains the underlying physical infrastructure.

The other options are not characteristics of public cloud. **Dedicated hardware** is generally associated with private cloud or specialized offerings (some public cloud services can offer dedicated hosts, but it's not a defining characteristic of public cloud). **Unsecured connections** is incorrect because public cloud connectivity can be secured using encryption, VPN, private endpoints, and strong identity controls. **Limited storage** is also incorrect; public cloud is designed to be highly scalable/elastic, with practical limits managed via quotas rather than being inherently "limited."

References: [Microsoft Learn: Cloud deployment models](#), [Microsoft Learn: Financial considerations \(consumption-based\)](#)

QUESTION NO: 10

Your company has an Azure subscription that contains several resources.

You need to identify which department is responsible for the cost of each resource. What should you use?

- A. tags.
- B. alerts.
- C. budgets .

ANSWER: A

Explanation:

Use **tags**. Azure tags are name/value pairs you apply to resources (or resource groups) to logically organize them and, importantly for this scenario, to support **cost allocation and chargeback**. A common pattern is to tag resources with values like `Department=Finance` or `CostCenter=1234`, then use Azure Cost Management to filter/group costs by those tags to see which department owns which spend.

Budgets are used to set spending thresholds and track/notify when costs approach or exceed a defined amount, but they don't inherently identify departmental ownership per resource unless you already have a grouping mechanism (like tags) in place. **Alerts** can notify you about conditions (including cost alerts), but they also don't provide a systematic way to attribute each resource's cost to a department.

So, to identify which department is responsible for the cost of each resource, tagging is the correct tool.

References: [Tag resources, resource groups, and subscriptions for logical organization](#), [Use tags for cost allocation \(Cost Management\)](#).

QUESTION NO: 11

You need to manage containers.

Which two services can you use? Each correct answer presents a complete solution,

NOTE: Each correct selection is worth one point.

- A. Azure Virtual Desktop
- B. Azure virtual machines
- C. Azure Functions
- D. Azure Kubernetes Service (AKS)
- E. Azure Container Instances

ANSWER: D E

Explanation:

To manage containers in Azure, the two most directly relevant services here are **Azure Kubernetes Service (AKS)** and **Azure Container Instances (ACI)**. AKS is Azure's managed Kubernetes offering, designed for orchestrating and managing containerized applications at scale (cluster management, scheduling, scaling, rolling updates, etc.). ACI is a serverless container runtime that lets you run containers on-demand without managing underlying VMs; it's commonly used for simple container workloads, burst scenarios, or as a building block alongside orchestration solutions.

Azure Virtual Desktop is for delivering virtualized Windows desktops and apps, not container management. **Azure virtual machines** can host containers (e.g., you can install Docker), but VMs themselves are not a container management service; you'd be managing the infrastructure and container runtime manually. **Azure Functions** is serverless compute for event-driven code; while it can be packaged with containers, it's not primarily a container management/orchestration service in the AZ-900 sense.

References: [What is Azure Kubernetes Service \(AKS\)?](#), [Azure Container Instances overview](#).

QUESTION NO: 12 - (HOTSPOT)

Select the answer that correctly completes the sentence.

Azure China

- is operated by Microsoft.
- has feature parity with Azye global.
- services can be Accessed from China only.
- is a distinct separate instance of Microsoft Azure.

Answer selections

ANSWER:

Azure China	
<input type="checkbox"/>	is operated by Microsoft.
<input type="checkbox"/>	has feature parity with Azure global.
<input type="checkbox"/>	services can be accessed from China only.
<input checked="" type="checkbox"/>	is a distinct separate instance of Microsoft Azure.
Answer selections	

Explanation:

For AZ-900, the important concept is that Azure has a few special sovereign/isolated cloud environments that are *separate* from the public Azure cloud (often called “Azure global”). Azure China is one of these environments. It runs in datacenters located in China and is operated by a local partner (21Vianet), which is why it’s treated as its own cloud instance with its own portal endpoints and service availability timeline.

Because of that, the statement “Azure China **is operated by Microsoft**” is not correct—Microsoft does not directly operate Azure China. Also, “Azure China **has feature parity with Azure global**” is not correct; services and features can lag or differ compared to Azure global. While there are connectivity and regulatory considerations for accessing services in China, the exam-friendly, always-correct way to describe Azure China is that it is a **distinct, separate instance** of Microsoft Azure.

So the dropdown should be set to: “**is a distinct separate instance of Microsoft Azure.**”

References: [Azure China 21Vianet - Overview and operations](#), [Azure sovereign clouds overview \(concept of separate instances\)](#).

QUESTION NO: 13

What are two benefits of cloud computing? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. enables the rapid provisioning of resources
- B. has increased administrative complexity

- C. has the same configuration options as on-premises
- D. shifts capital expenditures (CAPEX) to operating expenditures (OPEX)

ANSWER: A D

Explanation:

The correct benefits are **rapid provisioning of resources** and the ability to **shift CAPEX to OPEX**. Cloud platforms like Azure let you provision compute, storage, and other services quickly (often in minutes) using portals, templates, or automation, which supports agility and scaling without waiting for hardware procurement and setup. Cloud computing also changes the financial model: instead of large upfront purchases of servers and datacenter equipment (CAPEX), you typically pay for what you use as an ongoing operational expense (OPEX), improving cost flexibility and aligning spend with demand.

Option B is incorrect because increased administrative complexity is not a benefit; while cloud introduces new governance and management considerations, a core value proposition is reducing the burden of owning and maintaining physical infrastructure. Option C is incorrect because cloud services do not necessarily provide the same configuration options as on-premises; many services are abstracted/managed, trading some low-level control for simplicity, resiliency, and speed.

References: [Microsoft Learn: Cloud financial models \(CAPEX vs OPEX\)](#), [Microsoft Learn: Agility and speed in the cloud](#).

QUESTION NO: 14

Your company plans to deploy several million sensors that will upload data to Azure. You need to identify which Azure resources must be created to support the planned solution. Which two Azure resources should you identify? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point

- A. Azure Data Lake.
- B. Azure Queue storage.
- C. Azure File Storage.
- D. Azure IoT Hub.
- E. Azure Notification Hubs .

ANSWER: A D

Explanation:

For a solution where millions of sensors upload telemetry to Azure, you typically need (1) a cloud gateway/service to securely ingest and manage device connectivity at scale, and (2) a durable storage/analytics landing zone for the incoming data. **Azure IoT Hub** is the core Azure service designed for large-scale device-to-cloud ingestion, device identity, authentication, and bi-directional communication, making it the right choice for the sensor connectivity layer. Once telemetry is ingested, it commonly needs to be persisted for downstream processing and analytics. **Azure Data Lake (via Azure Data Lake Storage Gen2)** is a common target for IoT telemetry at scale and is supported as an IoT Hub message routing endpoint, enabling long-term storage and analytics-friendly organization of data.

Why the others are wrong: Azure Queue Storage is a general-purpose queueing service but isn't the primary device ingestion and management service for IoT fleets. Azure File Storage is for SMB/NFS file shares, not a typical high-throughput telemetry sink. Azure Notification Hubs is for pushing notifications to mobile devices, not ingesting sensor telemetry.

References: [Azure IoT Hub concepts](#), [IoT Hub device-to-cloud messages and routing endpoints](#).

QUESTION NO: 15

Which statement correctly describes the Modern Lifecycle Policy for services offered by Azure?

- A. Microsoft provides mainstream support for a service for five years.
- B. Microsoft provides a minimum of 12 months' notice before ending support for a service.
- C. After a service is made generally available, Microsoft provides support for the service for a minimum of four years.
- D. When a service is retired, you can purchase extended support for the service for up to five years.

ANSWER: B

Explanation:

Azure services generally fall under Microsoft's Modern Lifecycle Policy. A key promise of this policy is that customers will be given advance notice before a service is ended/retired when Microsoft is discontinuing support. Specifically, Microsoft states it will provide at least 12 months' notice prior to ending support for products governed by the Modern Lifecycle Policy (with some exclusions such as preview offerings and certain free services). This is why option B is the best description: it captures the policy's customer-protection commitment around retirement/end of support timelines.

The other options describe fixed-duration support terms (for example "five years of mainstream support" or "minimum of four years after GA") or the ability to buy extended support. Those concepts align more with older "Fixed Lifecycle" style policies for traditional on-premises products, not Azure's Modern Lifecycle approach, which is not defined by a guaranteed multi-year mainstream/extended support schedule. Likewise, Azure service retirement does not generally offer a purchasable extended support period "up to five years" as a standard policy.

References: [Microsoft Modern Lifecycle Policy](#), [Modern Lifecycle Policy \(Support article\)](#)

QUESTION NO: 16

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You plan to deploy several Azure virtual machines.

You need to ensure that the services running on the virtual machines are available if a single data center fails.

Solution: You deploy the virtual machines to two or more resource groups. Does this meet the goal?

- A. Yes.
- B. No.

ANSWER: B

Explanation:

No. Deploying VMs into two or more *resource groups* does not provide datacenter-level resiliency. A resource group is primarily a logical management boundary used to organize resources, apply RBAC, tags, policies, and manage lifecycle operations (deploy, update, delete) as a unit. It is not a high-availability construct and it does not control or guarantee that resources are placed in different datacenters.

To stay available if a single datacenter fails, you need to place VMs across fault domains at minimum, and ideally across physically separate datacenters within a region by using **Availability Zones** (zonal deployment) or across regions using paired regions plus traffic management. In Azure, the standard way to achieve this for VMs is to use **Availability Sets** (protects against rack/power/network failures within a datacenter) or **Availability Zones** (protects against a full datacenter/zone outage). Simply splitting VMs into multiple resource groups could still leave all VMs in the same zone/datacenter, so a single datacenter failure could take them all down.

References: [Azure Resource Manager resource groups](#), [Availability Zones overview](#).

QUESTION NO: 17

Your company has 10 departments.

The company plans to implement an Azure environment.

You need to ensure that each department can use a different payment option for the Azure services it consumes.

What should you create for each department?

- A. a reservation.
- B. a subscription.
- C. a resource group.
- D. a container instance .

ANSWER: B

Explanation:

The correct choice is to create a **subscription** for each department. In Azure, billing and payment are fundamentally scoped to the subscription: usage is metered and charges are accumulated at the subscription level, and cost management, invoices, and chargeback are typically organized per subscription. By giving each department its own subscription, you can separate costs cleanly and align each subscription with the appropriate billing arrangement/payment method available under your organization's billing account (for example, different invoicing profiles or cost allocation structures depending on the agreement type). This is the standard approach in Azure for departmental separation of spend and governance.

Option A (reservation) is incorrect because reservations are a pricing discount commitment for specific resource types (e.g., VMs) and don't create separate billing entities or payment options. Option C (resource group) is incorrect because resource groups are only logical containers for resources; they can help with organizing and reporting costs, but they don't provide separate payment methods or billing boundaries. Option D (container instance) is an Azure compute service and unrelated to billing segregation.

References: [Create an Azure subscription](#), [Understand cost management data \(costs are scoped and reported by subscription\)](#).

QUESTION NO: 18

You plan to reduce ongoing Azure expenditures.

You need to identify which factors affect the costs of a resource.

Which three factors should you identify? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. the Azure region.
- B. the type of processed data.
- C. the volume of inbound data.
- D. the volume of outbound data.
- E. the service tier .

ANSWER: A D E

Explanation:

Azure resource costs are primarily driven by where the resource is deployed, what pricing/tier you choose, and how much billable usage you generate (including data transfer). **Region** matters because Azure pricing varies by geography due to differences in operating costs and demand; the same service can have different rates in different regions. The **service tier** (for example, Standard vs. Premium, or different performance/feature tiers) directly changes the unit price and included capabilities, so it's a key lever for ongoing cost optimization. Finally, **outbound data volume** (egress) is commonly billed across many services; reducing egress or keeping traffic within the same region/VNet can lower costs.

The other options are less generally applicable as "factors" across Azure resources. **Inbound data volume** is typically free (ingress is not usually charged), so it's not a primary cost driver in most scenarios. **Type of processed data** is not a standard Azure pricing dimension; costs depend on the service's meter (compute time, transactions, storage, throughput, etc.), not the semantic "type" of data.

References: [Understand cost management data](#), [Azure services and pricing vary by region \(service availability/pricing context\)](#).

QUESTION NO: 19

You plan to deploy several Azure virtual machines.

You need to ensure that the services running on the virtual machines remain available if a single data center fails.

What are two possible solutions? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Deploy the virtual machines to a scale set.
- B. Deploy the virtual machines to two or more resource groups.
- C. Deploy the virtual machines to two or more availability zones.
- D. Deploy the virtual machines to two or more regions.

ANSWER: C D

Explanation:

To stay available when a **single datacenter** fails, you need redundancy across physically separate datacenters. In Azure, that is exactly what **Availability Zones** provide: each zone is a unique physical location within an Azure region, with independent power, cooling, and networking. By deploying VMs across **two or more availability zones**, your service can continue running even if one datacenter (zone) goes down.

Deploying to **two or more regions** can also keep services available if a datacenter fails, because regions are geographically separated and contain multiple datacenters. While this is broader than required (it protects against regional outages as well), it still satisfies the requirement of surviving a single datacenter failure, assuming the workload is deployed in an active/active or active/passive architecture across regions.

A **scale set** by itself doesn't guarantee protection from a datacenter failure unless it is explicitly configured to span availability zones; scale sets can be deployed within a single zone or without zones. **Resource groups** are only a management boundary and provide no physical resiliency.

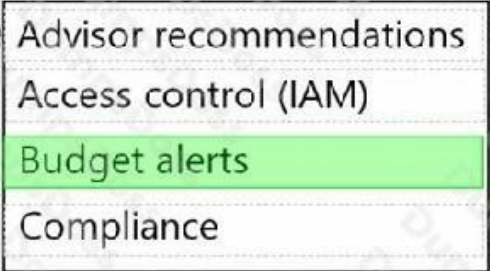
References: [Availability Zones overview](#), [Designing for resiliency and disaster recovery](#).

QUESTION NO: 20 - (HOTSPOT)

To complete the sentence, select the appropriate option in the answer area.

You can use Advisor recommendations
Access control (IAM)
Budget alerts
Compliance in Azure to send email alerts when the cost of the current billing period exceeds a specified limit.

ANSWER:

You can use  in Azure to send email alerts when the cost of the current billing period exceeds a specified limit.

Explanation:

The sentence is asking which Azure feature can send email alerts when your subscription's costs for the current billing period go over a limit you define. In Azure, that capability is provided by **budgets** in **Azure Cost Management**. When you create a budget, you set an amount (for example, \$500/month) and then configure alert thresholds (for example, 80%, 100%). Those budget thresholds can trigger notifications, including sending emails to specified recipients (and optionally triggering action groups/automation).

Looking at the available dropdown choices:

Advisor recommendations focuses on optimization guidance (cost, security, reliability, operational excellence, performance), but it doesn't function as a billing-period cost threshold alerting tool.

Access control (IAM) is for managing permissions and role assignments, not cost alerting.

Compliance relates to regulatory/compliance reporting and policies, not budget notifications.

Budget alerts is the only option that directly matches "send email alerts when the cost exceeds a specified limit," because budgets are designed specifically for cost tracking and alerting.

References: Azure budgets and alerts in Cost Management are documented here: [Create and manage budgets in Azure Cost Management](#) and budget alert behavior is described here: [Monitor cost and usage with alerts](#).