DUMPSARENA

Microsoft Azure Fundamentals

Microsoft AZ-900

Version Demo

Total Demo Questions: 151

Total Premium Questions: 1645

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Describe cloud concepts	322
Topic 2, Describe Azure architecture and services	650
Topic 3, Describe Azure management and governance	673
Total	1645

QUESTION NO: 1

A company is planning on deploying a stateless based application based on microservices using the Azure Service Fabric service. You need to design the infrastructure that would be required in the Azure Service Fabric service.

Which of the following should you consider? (Choose two)

A. The network connectivity.

B. The number of node types in the cluster.

C. The properties for each node type.

D. The service tier .

ANSWER: B C

Explanation:

The number of node types in the cluster is a core Service Fabric infrastructure design decision. A node type is a collection of nodes with the same configuration and, in Azure, is typically implemented with a Virtual Machine Scale Set. Planning node types enables workloads with distinct operational or scaling needs to be separated into suitable groups. For example, a cluster can use separate node types for system services and for application services, allowing each group to be scaled and managed independently. This is useful even for stateless microservices because service instances may need different capacity profiles or placement requirements.

The properties for each node type must also be designed because they define the capabilities and resilience of the underlying compute resources. Relevant properties include the virtual machine size, the number of instances, durability level, capacity settings, whether a node type is primary, and placement properties used by services. These choices determine how much compute capacity is available, how the cluster handles infrastructure changes, and where microservices can run. Microsoft's guidance identifies node types and their capacity as fundamental cluster-planning elements. See Service Fabric cluster capacity planning and Service Fabric node types.

QUESTION NO: 2

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:

Deploy the virtual machines to two or more availability zones because availability zones are physically separate locations within an Azure region. Each zone consists of one or more datacenters with independent power, cooling, and networking. When application instances are deployed across zones and the application uses an appropriate load-balancing design, a failure affecting the datacenter infrastructure in one zone does not prevent instances in another zone from continuing to provide the service. This gives the workload datacenter-level fault isolation while remaining in the same Azure region.

Deploy the virtual machines to two or more regions also provides a complete solution. Azure regions are geographically distinct geographic areas containing Azure datacenters. Deploying independent application instances in multiple regions, together with appropriate global traffic routing, failover procedures, and data replication where required, allows the

application to continue serving users from an unaffected region. This architecture provides resilience beyond a single datacenter and can be selected when the application has requirements for wider geographic redundancy. For details, see the Azure availability zones overview and Microsoft's guidance to design reliable applications across regions.

QUESTION NO: 3

This question requires that you evaluate the underlined text to determine if it is correct.

When planning to migrate a public website to Azure, you must plan to pay monthly usage costs .

Instructions: Review the underlined text. If it makes the statement correct, select "No change is needed". If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed

B. Deploy a VPN

C. pay to transfer all the website data to Azure

D. reduce the number of connections to the website

ANSWER: A

Explanation:

No change is needed is correct because an Azure-hosted public website requires planning for the costs of the Azure resources and services that support it. Azure commonly uses a consumption-based pricing model: charges are based on the services selected and the amount of those services consumed over the billing period. For a website, this can include the hosting compute service, storage, database services, networking, monitoring, and outbound data transfer, depending on the architecture. Although some Azure services offer free tiers, included monthly quotas, reservations, or fixed-price plans, these do not remove the need to assess and plan for recurring monthly costs. A public site can generate variable demand, so cost planning should account for expected traffic, scaling behavior, service tier selection, and data egress. Azure Cost Management and Billing provides tools for estimating, monitoring, allocating, and controlling these ongoing costs, while the Azure pricing page provides service-specific pricing information. Therefore, planning to pay monthly usage costs is an appropriate and accurate requirement when preparing a public-website migration to Azure. See Azure pricing and Microsoft Cost Management overview.

QUESTION NO: 4

You plan to deploy an application to Azure virtual machines.

You need to ensure that the application can access an Azure Storage account without storing account keys or passwords in the application code.

Which two Azure features should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a managed identity.

B. an Azure RBAC role assignment.

C. a resource lock.

D. an availability set.

E. a network security group (NSG).

ANSWER: A B

Explanation:

A managed identity provides an identity in Microsoft Entra ID for an Azure resource, such as a virtual machine. The application can use this identity to request tokens without storing credentials in code, configuration files, or scripts.

An Azure RBAC role assignment grants the managed identity the required permissions to the Storage account or to a specific storage scope. For example, a Storage Blob Data Reader or Storage Blob Data Contributor role can authorize access to blob data. Together, the managed identity supplies authentication and the role assignment supplies authorization. See Managed identities for Azure resources and Assign an Azure role for access to blob data.

QUESTION NO: 5

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:

metered pricing and **self-service management** are core characteristics of public-cloud computing. Public cloud services are delivered by a third-party provider over the internet, with customers consuming shared, provider-managed infrastructure and platform services as needed. Metered pricing, also called consumption-based pricing, means that a customer is billed according to the resources and services used. Depending on the Azure service, usage can include compute time, storage consumed, transactions, requests, or outbound network data. This model enables organizations to align costs with actual demand rather than buying and maintaining enough on-premises hardware for peak capacity.

Self-service management enables customers to provision, configure, scale, monitor, and delete their cloud resources independently. In Azure, this can be done through the Azure portal, Azure CLI, PowerShell, APIs, templates, and infrastructure-as-code tools. Resources can therefore be deployed rapidly without the customer needing to own physical datacenter equipment or wait for manual infrastructure provisioning. Although Microsoft operates and secures the underlying cloud datacenters and physical infrastructure, customers retain control over the Azure resources and configurations within their subscriptions. These capabilities support the agility and cost flexibility associated with public cloud. For more information, see Microsoft Learn: What is cloud computing? and Microsoft Learn: Azure resource management.

QUESTION NO: 6

Your company plans to deploy several million IoT sensors that will upload data to Azure. You need to determine which Azure resources need to be created to support this intended deployment.

Which two Azure resources should you choose? Each correct answer forms a 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:

Azure IoT Hub is designed to provide a managed, highly scalable cloud gateway for Internet of Things devices. It supports reliable device-to-cloud telemetry ingestion from a large fleet of devices, including millions of sensors, while providing per-device identities and secure authentication. IoT Hub can also route incoming telemetry to Azure services for downstream storage, processing, and analytics. This makes it the appropriate Azure service for connecting and managing the sensor fleet and receiving its uploaded data.

Azure Data Lake provides the scalable storage foundation for retaining the substantial volume of telemetry produced by millions of sensors. In current Azure implementations, this capability is delivered through Azure Data Lake Storage Gen2, which combines the scale and cost characteristics of Azure Blob Storage with a hierarchical namespace that supports big-data and analytics workloads. It is well suited for storing raw historical sensor data and making that data available for analysis with services such as Azure Synapse Analytics, Azure Databricks, and Azure Data Factory. Together, Azure IoT Hub supplies secure sensor connectivity and ingestion, while Azure Data Lake supplies durable, analytics-ready storage. For details, see Azure IoT Hub overview and Azure Data Lake Storage Gen2 introduction.

QUESTION NO: 7

You have an azure virtual machine named VM1.

You plan to encrypt VM1 by using Azure Disk Encryption. Which Azure resource must you create first?

A. An Azure Storage account.

B. An Azure Information Protection policy.

C. An Encryption Key.

D. An Azure Key Vault.

ANSWER: D

Explanation:

An Azure Key Vault is required before Azure Disk Encryption can be enabled because it provides the protected location where the service stores and manages the encryption secrets used for the virtual machine's disks. Azure Disk Encryption uses BitLocker for supported Windows virtual machines and DM-Crypt for supported Linux virtual machines. During encryption, the Azure Disk Encryption extension relies on the vault to safeguard the disk-encryption secret, keeping that cryptographic material separate from the virtual machine and its managed disks.

The Key Vault must be created and configured before starting the encryption deployment. Its authorization configuration must allow the necessary deployment and disk-encryption operations, and the identity performing the deployment needs appropriate access to the vault's secrets and, if applicable, keys. This arrangement supports centralized control of sensitive cryptographic material, including auditability and access management through Azure Key Vault. Microsoft's Azure Disk Encryption documentation specifically identifies Azure Key Vault as the service used to control and manage disk-encryption keys and secrets. For more information, see the Azure Disk Encryption overview and the Azure Key Vault overview.

QUESTION NO: 8

You are planning to deploy multiple Azure virtual machines and need to manage the ports that can be used by devices on the Internet to access these virtual machines. Which tool should you utilize?

A. a network security group (NSG)

B. an Azure Active Directory (Azure AD) role

C. an Azure Active Directory group

D. an Azure key vault

ANSWER: A

Explanation:

A network security group (NSG) is the appropriate tool because it manages network traffic to and from Azure resources by using stateful inbound and outbound security rules. In this scenario, inbound NSG rules can control exactly which Internet-originated connections may reach the virtual machines. Each rule can specify criteria including the protocol, source and destination IP addresses or service tags, source and destination port ranges, and whether matching traffic is allowed or denied. For example, an organization can permit public HTTPS access on TCP port 443 while limiting remote administration traffic to approved source IP ranges and required ports only.

An NSG can be associated with a subnet, allowing a consistent port-access policy to be applied to multiple virtual machines deployed within that subnet. It can also be associated with an individual virtual machine's network interface when a more specific policy is required. Rules are processed according to their priority values, supporting controlled and predictable Internet-facing access. This makes NSGs the Azure network-layer mechanism for managing ports exposed to devices on the Internet. Microsoft documents NSG rule properties, rule processing, and associations with subnets and network interfaces in the Network security groups overview and the Manage network security groups documentation.

QUESTION NO: 9

A team has an Azure Cosmos DB account. A solution needs to be in place to generate an alert from Azure Log Analytics when a query request charge exceeds 40 units more than 10 times during a 10-minute window.

Which of the following would you recommend? (Choose two)

A. Create a search query to identify when the requestCharge_s exceeds 10.

B. Configure a period of 10 minutes and a frequency of 10 minutes.

C. Create a search query to identify when the requestCharge_s exceeds 40.

D. Create a search query to identify when the duration_s exceeds 10.

ANSWER: B C

Explanation:

Configure a period of 10 minutes and a frequency of 10 minutes because a log search alert rule uses its evaluation period as the lookback interval for the log query. A 10-minute period therefore evaluates the Azure Cosmos DB log events generated during the required 10-minute observation window. Running the evaluation every 10 minutes provides a straightforward nonoverlapping schedule for checking that window. The alert condition can use the count of query results and trigger when the number of matching events is greater than 10.

Create a search query to identify when the `requestCharge_s` exceeds 40 because Azure Cosmos DB measures the cost of database operations in request units. When Cosmos DB resource logs are sent to a Log Analytics workspace, a Kusto Query Language query can filter records where the recorded request charge is greater than 40. Combining that filter with an alert measurement that counts matching records implements the requirement to alert after more than 10 high-charge requests in the selected period. Azure Monitor log search alerts support query-based conditions, aggregation, thresholds, evaluation periods, and evaluation frequencies. For implementation details, see Azure Monitor log search alert rules and Azure Cosmos DB resource logs.

QUESTION NO: 10

You need to monitor Azure spending and receive a notification when forecasted costs exceed a defined amount.

Which two Azure Cost Management capabilities should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Cost analysis.

B. Budgets.

C. Azure Policy.

D. Azure Service Health.

ANSWER: A B

Explanation:

Cost analysis. and **Budgets.** provide the required Azure Cost Management functionality. Cost analysis lets you examine cost data at scopes such as a billing account, subscription, resource group, or management group. It provides accumulated cost views, trends, grouping, filtering, and forecast information, enabling you to monitor spending and assess whether projected charges are approaching or passing a target amount. A budget defines that target as a cost or usage threshold for the selected scope and time period. Budget alert rules support notifications based on either actual costs or forecasted costs. In particular, a forecasted alert can be configured to trigger when the projected total cost for the budget period exceeds the specified percentage of the budgeted amount. Notification recipients can be entered directly, and action groups can be used to deliver alerts through supported channels or trigger automation. Using cost analysis to review spending and forecast data together with a budget configured with a forecasted-cost alert therefore satisfies both requirements: monitoring Azure spending and receiving a notification when the forecast exceeds the defined amount. See Cost analysis common uses and Create and manage budgets.

QUESTION NO: 11

Which service provides serverless computing in Azure?

A. Azure Virtual Machines.

B. Azure Container instances.

C. Azure Functions.

D. Azure storage account .

ANSWER: C

Explanation:

Azure Functions is Azure's primary serverless compute service for running small pieces of code in response to events. A function can be triggered by an HTTP request, a scheduled timer, a message placed on a queue, or a change in a data source. This event-driven model lets developers concentrate on application code and configuration rather than provisioning, administering, patching, or maintaining the underlying servers.

Azure automatically manages the infrastructure that hosts the functions and can scale execution capacity according to workload demand. For workloads using the Consumption plan, costs are based on executions and resource consumption, making it well suited to applications with intermittent, variable, or unpredictable activity. The term serverless means that server management is abstracted from the developer; servers still exist and are operated by Azure. Azure Functions is therefore the appropriate service when the requirement is to provide serverless computing in Azure. Microsoft identifies Azure Functions as a serverless solution for building event-driven applications and describes its triggers, bindings, and supported hosting models in the Azure Functions overview. For broader context on the serverless approach, see Microsoft's serverless architecture guidance.

QUESTION NO: 12

An Azure administrator plans to run a PowerShell script that creates Azure resources. You need to recommend which computer configuration to use to run the script.

Which three computers can run the script? Each correct answer presents a complete solution NOTE: Each correct selection is worth one point.

A. a computer that runs macOS and has PowerShell and the Az PowerShell module installed.

B. a computer that runs Windows 10 and has the Azure PowerShell module installed.

C. a computer that runs Chrome OS and uses Azure Cloud Shell.

D. a computer that runs Linux and has the Azure CLI tools installed .

ANSWER: A B C

Explanation:

a computer that runs macOS and has PowerShell and the Az PowerShell module installed. is a complete solution because PowerShell is cross-platform and can run on macOS. The Az PowerShell module provides the Azure-specific cmdlets that a provisioning script uses to sign in to Azure and create or manage resources. The administrator must also authenticate to Azure with an identity authorized to create the required resources.

a computer that runs Windows 10 and has the Azure PowerShell module installed. can also run the script. Windows supports PowerShell, and installing the Az module makes Azure resource-management cmdlets available to the local PowerShell session. Microsoft documents the Az module as the current PowerShell module for interacting with Azure.

a computer that runs Chrome OS and uses Azure Cloud Shell. is valid because Azure Cloud Shell is browser-based and can be accessed from a Chrome OS device. Cloud Shell includes a PowerShell environment and Azure tooling, so the script runs in the Azure-hosted shell rather than requiring PowerShell and modules to be installed on the local Chromebook. For details, see the Azure PowerShell documentation and the Azure Cloud Shell overview.

QUESTION NO: 13

What are two characteristics of the public cloud? Each correct answer presents 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:

metered pricing and **self-service management** are fundamental characteristics of the public cloud. Public-cloud services are commonly billed according to consumption: customers pay for the resources and services they use, such as compute time, storage capacity, and network usage. This consumption-based approach enables organizations to align operational costs with demand rather than purchasing and maintaining all underlying datacenter infrastructure in advance. Azure describes this model as consumption-based pricing, where customers pay only for the resources they use.

Public cloud also supports self-service provisioning and management. Customers can use portals, command-line tools, APIs, infrastructure-as-code templates, and management tools to deploy and configure their own cloud resources, such as virtual machines, web apps, storage accounts, and networks. The cloud provider operates and maintains the physical datacenters and foundational hardware, while customers retain control over the cloud resources and configurations they create. These capabilities provide the agility and scalability associated with public-cloud computing. See Microsoft Learn: Describe cloud computing and Microsoft Learn: Azure subscription and service management.

QUESTION NO: 14

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:

A subscription is the appropriate construct to create for each department because an Azure subscription provides a distinct boundary for consuming Azure services, tracking usage, and managing costs. Every Azure resource is deployed into one subscription, and Azure records the associated service usage and charges at that subscription scope. Creating separate subscriptions therefore enables the organization to associate each department's Azure consumption with its own billing arrangement, supporting separate payment methods where the billing account type and agreement support that configuration.

Subscriptions also provide useful operational boundaries beyond payment. Each department can have separate cost analysis views, budgets, access-control assignments, Azure Policy assignments, and resource quotas. This allows departments to manage their own workloads and spending while the organization retains centralized governance through management groups and the broader Azure billing structure. Microsoft describes subscriptions as logical containers used to provision Azure resources and as a key scope for billing and cost management. For more information, see Azure subscription design considerations and Create additional Azure subscriptions.

QUESTION NO: 15

Your company plans to deploy an Artificial Intelligence (AI) solution in Azure.

What should the company use to build, test, and deploy predictive analytics solutions?

A. Azure Logic Apps

B. Azure Machine Learning designer

C. Azure Batch

D. Azure Cosmos DB

ANSWER: B

Explanation:

Azure Machine Learning designer is the appropriate service because it provides a visual, drag-and-drop environment for building machine learning pipelines without requiring extensive code. A team can use it to ingest and prepare data, select and train machine learning algorithms, evaluate model performance, and compare results. These capabilities directly support the development and testing of predictive analytics solutions, such as models that forecast values or classify future outcomes from historical data.

After a suitable model is trained and evaluated, Azure Machine Learning supports operationalizing it through deployment to managed online endpoints for real-time inference or batch endpoints for large-scale scoring. The designer integrates these steps into an Azure Machine Learning workspace, enabling teams to manage data assets, compute, experiments, model artifacts, and deployment resources in a unified machine learning platform. This makes Azure Machine Learning designer a strong fit for an organization that needs an accessible end-to-end workflow to create, validate, and deploy predictive models

in Azure. Microsoft describes designer as a drag-and-drop interface for creating machine learning pipelines; see Azure Machine Learning designer documentation and Azure Machine Learning endpoints documentation.

QUESTION NO: 16

Your company plans to migrate workloads from an on-premises datacenter to Azure.

You need to estimate the monthly cost of the planned Azure resources and compare the estimated Azure cost to the cost of the existing on-premises environment.

Which two tools should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. the Azure Pricing calculator.

B. the Total Cost of Ownership (TCO) Calculator.

C. Azure Service Health.

D. Azure Resource Graph Explorer.

E. the Azure portal.

ANSWER: A B

Explanation:

The Azure Pricing calculator is used to estimate the expected cost of Azure services before deployment. It lets an organization select services, regions, SKUs, and estimated usage to calculate projected Azure charges.

The Total Cost of Ownership (TCO) Calculator is used to compare the cost of operating on-premises infrastructure with the projected cost of running equivalent workloads in Azure. It considers on-premises cost categories such as servers, storage, networking, electricity, and IT labor. See the Azure Pricing calculator and the Total Cost of Ownership Calculator.

QUESTION NO: 17

An Azure administrator plans to run a PowerShell script that creates Azure resources. You need to recommend which computer configuration to use to run the script.

Which three computers can run the script? Each correct answer presents a complete solution

NOTE: Each correct selection is worth one point.

A. a computer that runs macOS and has PowerShell Core 6.0 installed

B. a computer that runs Windows 10 and has the Azure PowerShell module installed

C. a computer that runs Chrome OS and uses Azure Cloud Shell

D. a computer that runs Linux and has the Azure CLI tools installed

ANSWER: A B C

Explanation:

Azure PowerShell scripts can be run from supported local PowerShell environments or from Azure Cloud Shell's browser-hosted PowerShell environment. A computer that runs macOS and has PowerShell Core 6.0 installed can provide the cross-platform PowerShell host needed to execute PowerShell automation. Azure PowerShell was designed to support PowerShell Core on macOS, allowing administrators to use PowerShell cmdlets for Azure resource management outside Windows.

A computer that runs Windows 10 and has the Azure PowerShell module installed is also a complete local administration configuration. The Azure PowerShell module supplies the Azure-specific cmdlets used by scripts to authenticate, create, configure, and manage Azure resources.

A computer that runs Chrome OS and uses Azure Cloud Shell can run the script because Cloud Shell is accessed through a web browser and includes a maintained PowerShell environment with Azure tooling available. The local device does not need a local Azure PowerShell installation; the script executes in the Microsoft-managed Cloud Shell session after the administrator signs in to Azure. Microsoft documents Azure PowerShell's supported platform availability and Cloud Shell's browser-based PowerShell experience at Install the Azure Az PowerShell module and Azure Cloud Shell overview.

QUESTION NO: 18

In which type of cloud computing model are all the hardware resources owned by a third party and shared among multiple users or tenants?

A. private

B. hybrid

C. public

ANSWER: C

Explanation:

The public cloud model is correct because the cloud service provider owns and operates the physical infrastructure, including datacenters, servers, storage, and networking. Organizations access these services over the internet and do not need to buy, house, or maintain the underlying hardware. Public cloud infrastructure is designed for multitenancy, meaning that multiple independent customers use a shared pool of provider-managed resources. The provider maintains logical separation among customers' identities, subscriptions, data, and workloads, while delivering the economies of scale associated with shared infrastructure. This approach enables customers to provision resources on demand, scale capacity as needs change, and generally pay for the services they consume rather than making large upfront hardware investments. Microsoft Azure is a public cloud platform: Microsoft owns and manages the global infrastructure and provides services to many customers through isolated Azure tenants and subscriptions. The combination of third-party ownership and shared use by multiple tenants is the defining characteristic described for public cloud computing. Microsoft Learn explains that public clouds are owned and operated by a third-party cloud service provider and are shared among organizations or the public. See Microsoft Learn: Describe cloud computing and Microsoft Cloud Adoption Framework: Cloud models overview.

QUESTION NO: 19

Note: This question belongs to a series of questions that present the same scenario. Each question within the series contains a distinct solution that could potentially meet the outlined objectives. Some sets of questions might have more than one correct solution, while others may lack a correct solution entirely. After answering a question in this section, you will **NOT** be able to return to it, and thus these questions will not appear in the review screen.

Your organization intends to migrate all its data and resources to Azure.

According to the company's migration strategy, only Platform as a Service (PaaS) solutions are to be utilized in Azure.

Your task is to deploy an Azure environment that aligns with the company's migration plan.

Proposed Solution: Create an Azure App Service and Azure Storage accounts.

Does this meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes is correct because Azure App Service and Azure Storage accounts are Azure managed platform services, allowing the organization to use application hosting and data-storage capabilities without deploying or administering virtual machines, guest operating systems, or physical storage infrastructure. Azure App Service provides a managed environment for hosting web apps, REST APIs, and mobile back ends. The organization is responsible for its application code, configuration, identity, and data, while Azure operates the underlying infrastructure and platform components needed to run the service.

An Azure Storage account is also a managed Azure service that provides a namespace and configuration boundary for storage services such as Blob Storage, Azure Files, Queue Storage, and Table Storage. Administrators configure such items as redundancy, access control, encryption, networking, and lifecycle policies, but Microsoft manages the storage hardware and service infrastructure. Therefore, creating these services supports a migration strategy restricted to PaaS solutions: the customer consumes managed services rather than infrastructure resources such as Azure virtual machines.

Microsoft describes App Service as a fully managed platform for building, deploying, and scaling web applications. Azure Storage is Microsoft's managed cloud-storage solution for modern data-storage scenarios. See Azure App Service overview and Azure Storage account overview.

QUESTION NO: 20

Which tool should you utilize to assess if your company's Azure environment complies with regulatory standards?

A. the Knowledge Center website

B. the Advisor blade from the Azure portal

C. Compliance Manager from the Service Trust Portal

D. the Solutions blade from the Azure portal

ANSWER: C

Explanation:

Compliance Manager from the Service Trust Portal is the appropriate choice because it is designed to help organizations assess and manage compliance obligations for Microsoft cloud services, including Azure. It supplies built-in assessments that align regulatory, industry, and regional requirements with relevant controls. These assessments give compliance teams a structured view of their compliance posture and identify improvement actions for controls that the customer is responsible for implementing or documenting.

Compliance Manager also provides implementation guidance and supports recording evidence and action status. This enables an organization to track progress against a selected regulatory framework over time, rather than merely viewing Microsoft's own certifications. The tool's scoring and improvement-action model is especially useful for prioritizing compliance work and demonstrating governance activities to internal stakeholders or auditors.

Although Compliance Manager is now provided through the Microsoft Purview portal, the Service Trust Portal remains Microsoft's location for trust resources and compliance documentation. Microsoft documents Compliance Manager's assessments, improvement actions, and regulatory templates at Microsoft Purview Compliance Manager. For information about accessing Microsoft compliance reports and trust materials, see the Service Trust Portal documentation.

QUESTION NO: 21

You have an Azure environment that contains multiple Azure virtual machines.

You plan to implement a solution that enables the client computers on your on-premises network to communicate to the Azure virtual machines.

You need to recommend which Azure resources must be created for the planned solution. Which two Azure resources should you include in the recommendation? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a virtual network gateway.

B. a load balancer.

C. an application gateway.

D. a virtual network.

E. a gateway subnet.

ANSWER: A E

Explanation:

A site-to-site VPN solution needs **a virtual network gateway** to provide the Azure endpoint for the encrypted VPN tunnel between Azure and an on-premises VPN device. The gateway uses IPsec/IKE connectivity and routing to enable clients on the on-premises network to reach private IP addresses assigned to virtual machines in the Azure virtual network. This is the managed Azure resource that performs the VPN gateway role.

Creating a virtual network gateway also requires **a gateway subnet**. This is a dedicated subnet in the virtual network that must be named `GatewaySubnet`. Azure deploys the gateway instances into that subnet, so it must exist before the virtual network gateway can be provisioned. The subnet should be sized appropriately to accommodate the gateway and potential future gateway capabilities. Because the Azure environment already contains virtual machines, the virtual network hosting those workloads is presumed to be available; the gateway-specific components needed for the planned connectivity are the virtual network gateway and its required dedicated subnet.

Microsoft documents the role of VPN gateways and the GatewaySubnet prerequisite in About VPN Gateway and Create a site-to-site VPN connection.

QUESTION NO: 22

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 availability zones. Does this meet the goal?

A. Yes.

B. No.

ANSWER: A

Explanation:

Yes. Deploying redundant virtual machines across two or more Availability Zones meets the requirement to maintain service availability if one datacenter fails. An Availability Zone is a physically separate group of datacenters within an Azure region. Zones have independent power, cooling, and networking infrastructure, so an outage affecting one zone does not inherently affect resources deployed in the other zones. Distributing VM instances across zones therefore provides protection against a datacenter-level failure while keeping the workload in the same Azure region.

For the services themselves to remain available, the deployment should include multiple VM instances hosting the service and a mechanism to route traffic to healthy instances, such as Azure Load Balancer. The workload must also be designed so that its application state, data, and dependencies can tolerate an individual instance or zone becoming unavailable. Those implementation considerations do not change that the stated deployment approach directly provides the required zone-level fault isolation.

Microsoft describes Availability Zones as the resiliency option for protecting workloads from datacenter failures in supported regions. See the Azure Availability Zones overview and Azure virtual machine availability options.

QUESTION NO: 23

Which service provides network traffic filtering capabilities for multiple Azure subscriptions and virtual networks?

A. Azure Firewall

B. an application security group

C. Azure DDoS protection

D. a network security group (NSG)

ANSWER: A

Explanation:

Azure Firewall is the appropriate service because it is a fully managed, stateful firewall designed to protect Azure Virtual Network resources and to apply centralized network-traffic controls at scale. Organizations commonly deploy it in a hub virtual network and connect application virtual networks as spokes. Routing can then direct traffic from those connected networks through the firewall for inspection, enabling consistent filtering of inbound, outbound, and east-west traffic. This architecture can span virtual networks and subscriptions, subject to appropriate connectivity, routing, and Azure role-based access control configuration.

Azure Firewall provides network rules based on source and destination addresses, ports, and protocols, plus application rules for outbound web traffic using fully qualified domain names. It also supports threat-intelligence-based filtering. Azure Firewall Manager and Firewall Policy help organizations centrally manage firewall policies and associate them with multiple Azure Firewall instances, including in environments that use more than one subscription. These capabilities make Azure Firewall the Azure service intended for centrally governed, broad network traffic filtering across an Azure estate. Microsoft describes it as a cloud-native, intelligent network firewall security service for Azure Virtual Network resources. See the Azure Firewall overview and Azure Firewall Manager overview.

QUESTION NO: 24

Fill in the blank

To reduce costs associated with an idle Azure virtual machine running Windows 10, you should (_____).

A. Select Stop for the virtual machine from the Azure portal.

ANSWER: A

Explanation:

Selecting **Stop** for the virtual machine from the Azure portal is the appropriate action because the portal stops and deallocates the VM. In the **Stopped (deallocated)** state, Azure releases the physical compute resources assigned to that VM, including its allocated CPU and memory capacity. As a result, the VM no longer incurs compute charges while it is not in use. The virtual machine's configuration, operating system disk, and attached data disks are preserved, so it can be started again when it is needed without rebuilding it.

This distinction is important for cost management: a VM that is merely shut down from within the Windows operating system can remain allocated in Azure, meaning compute billing can continue. Stopping it through the Azure portal ensures Azure deallocates it. Deallocation does not eliminate every possible charge; managed disks, snapshots, and some networking resources may still be billed because they remain provisioned. Nonetheless, stopping and deallocating an idle VM is the standard Azure action for eliminating its compute cost while retaining the VM for later use. Microsoft documents VM states and the associated billing behavior in Azure virtual machine power states and Azure Virtual Machines pricing and billing.

QUESTION NO: 25

Consider the following scenario: You are preparing to migrate a company to Azure. Each division of the company will have a dedicated administrator to manage Azure resources for their respective division. You aim to ensure that the Azure deployment allows for distinct segmentation for each division, while minimizing administrative overhead.

Proposed Solution: You are considering the use of multiple Azure Active Directory (Azure AD) directories to achieve this segmentation.

Does this solution meet the requirement?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. Multiple Azure AD directories, now known as Microsoft Entra tenants, create separate identity and administrative boundaries. Although this can isolate divisions, it does not minimize administrative overhead because each tenant has separate users, groups, application registrations, policies, privileged roles, and governance configuration. Administrators and users who need to work across divisions may also need cross-tenant access arrangements, such as B2B guest accounts or cross-tenant synchronization.

For divisions within one company, the usual Azure design is to retain a single Microsoft Entra tenant and organize resources through Azure's management hierarchy. Separate subscriptions can provide billing and administrative boundaries for divisions, while management groups can apply governance consistently across related subscriptions. Azure role-based access control can then delegate each division's administrator the required permissions at the relevant management group, subscription, resource group, or resource scope. This approach provides clear segmentation and delegated administration without duplicating identity-management processes.

Microsoft documents that a Microsoft Entra tenant is an identity and access-management boundary, whereas Azure RBAC supports authorization at multiple Azure resource scopes. See What is Microsoft Entra ID? and Azure RBAC overview.

QUESTION NO: 26

Which two types of customers are eligible to use Azure Government to develop a cloud solution? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. a Canadian government contractor

B. a European government contractor

C. a United States government entity

D. a United States government contractor

E. a European government entity

ANSWER: C D

Explanation:

a United States government entity and **a United States government contractor** are eligible to use Azure Government. Azure Government is a separate Azure cloud environment designed for the requirements of US government agencies and organizations that handle sensitive government data. It provides dedicated US datacenters and is operated by screened US personnel, supporting applicable government compliance requirements and standards.

Eligibility includes US federal, state, local, and tribal government entities. It also includes US government contractors when they have an eligible government customer or contractual requirement to use the environment. Contractors can use Azure

Government to develop, deploy, and manage solutions for qualifying US government workloads, including workloads that need to meet requirements such as FedRAMP, DoD Impact Levels, or Criminal Justice Information Services-related controls. Access is subject to an eligibility validation process, helping ensure that the environment remains available for its intended US public-sector customers and their approved partners.

For Microsoft's overview of the Azure Government environment and its intended customers, see What is Azure Government? and Introduction to Azure Government.

QUESTION NO: 27

Your company is planning to migrate its infrastructure to Microsoft Azure. Considering the organization consists of multiple departments, each department will have its own administrator responsible for managing the Azure resources. What are two viable methods to organize and segment Azure resources for each department? Each correct answer provides a complete solution.

NOTE: Each correct answer earns one point.

A. multiple subscriptions

B. multiple Azure Active Directory (Azure AD) directories

C. multiple regions

D. multiple resource groups

ANSWER: A D

Explanation:

multiple subscriptions and **multiple resource groups** are viable approaches for separating departmental Azure resources and delegating administration. An Azure subscription is a management, governance, billing, and quota boundary. A company can assign a distinct subscription to each department, then use Azure role-based access control (Azure RBAC) to grant that department's administrator permissions at the subscription scope. This model supports independent cost tracking, spending controls, quotas, and governance requirements for each department.

Resource groups are logical containers within an Azure subscription that hold resources sharing a lifecycle or management scope. A company can create a separate resource group for each department and assign the appropriate Azure RBAC role at that resource-group scope. This gives each departmental administrator authority to manage the resources in their group without granting access to resources held by other groups. Resource groups also provide a useful scope for applying Azure Policy, tags, locks, and deployments consistently. The appropriate design depends on whether the organization requires separate billing and quota boundaries, which subscriptions provide, or delegated administration within a shared subscription, which resource groups provide. For more information, see Organize subscriptions and Manage resource groups.

QUESTION NO: 28

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:

the Azure region., **the volume of outbound data.**, and **the service tier .** are factors that can directly affect Azure resource costs. Azure pricing varies across regions because the costs of operating datacenters, local market conditions, and service availability can differ by location. Consequently, the same resource configuration may have a different rate when deployed in another Azure region.

A service tier, SKU, or performance level determines the capabilities and capacity supplied by a service, such as compute resources, throughput, storage performance, availability features, or included functionality. Selecting the tier that appropriately matches workload requirements is therefore an important part of managing ongoing spend.

Outbound data transfer is also a metered cost in many Azure scenarios. Azure bandwidth pricing commonly applies to data leaving Azure, and charges depend on the amount transferred and the relevant source and destination scenario. Reviewing egress patterns, placing related workloads appropriately, and minimizing unnecessary outbound transfers can help control these charges. Microsoft describes the pricing influences that should be considered when estimating Azure costs and publishes the applicable bandwidth rates by transfer scenario. See Azure cost considerations and Azure Bandwidth Pricing.

QUESTION NO: 29

Suppose you have an Azure environment already established. You want to create a new Azure virtual machine using a tablet that operates on the Android operating system. Identify the three possible methods available to achieve this task. Each correct method offers a full solution.

NOTE: Each correct choice is worth one point.

A. Use Bash in Azure Cloud Shell.

B. Use PowerShell in Azure Cloud Shell.

C. Use the PowerApps portal.

D. Use the Security & Compliance admin center.

E. Use the Azure portal.

ANSWER: A B E

Explanation:

Use Bash in Azure Cloud Shell, Use PowerShell in Azure Cloud Shell, and Use the Azure portal are complete methods for creating an Azure virtual machine from an Android tablet. The relevant factor is access to a supported web browser and an Azure account with sufficient permissions, rather than the tablet's operating system. The Azure portal is a browser-based management interface that provides a guided virtual-machine creation experience, including configuration of the subscription, resource group, image, size, administrator authentication, disks, and networking. Azure Cloud Shell is also browser accessible and integrated with the Azure portal. It offers both Bash and PowerShell sessions that are authenticated for Azure management. A Bash session can provision a VM with Azure CLI commands, such as `az vm create`. A PowerShell session can perform the same kind of deployment through Az PowerShell cmdlets, including `New-AzVM`. These methods support full VM provisioning as long as the account has the required Azure RBAC permissions and the deployment configuration specifies the necessary resources. Microsoft describes Cloud Shell as a browser-based service with both Bash and PowerShell environments, and its virtual-machine documentation describes creating and configuring VMs through the Azure portal. See Azure Cloud Shell overview and Create a Windows VM in the Azure portal.

QUESTION NO: 30



You are planning to store 20 terabytes of data in Microsoft Azure. The data will be accessed infrequently but will need to be visualized using Microsoft Power BI.

What storage solutions should you consider recommending for this scenario? Choose two options, with each correct selection being worth one point.

A. Azure Data Lake

B. Azure Cosmos DB

C. Azure SQL Data Warehouse

D. Azure SQL Database

E. Azure Database for PostgreSQL

ANSWER: A C

Explanation:

Azure Data Lake is appropriate because Azure Data Lake Storage is designed for highly scalable storage of large quantities of structured, semistructured, and unstructured data used in analytics workloads. A 20-terabyte dataset is well suited to this storage model. Because Data Lake Storage is built on Azure Blob Storage, organizations can use blob access tiers and lifecycle management policies to help control costs for data that is accessed infrequently, while retaining the data for downstream analytical use. It also integrates broadly with Azure analytics services that can prepare and serve data for reporting.

Azure SQL Data Warehouse is appropriate for the analytical and business-intelligence layer needed to support Power BI. This service is now known as a dedicated SQL pool in Azure Synapse Analytics. Dedicated SQL pools are designed for large-scale analytical queries and reporting, using massively parallel processing to query and aggregate data efficiently. Data can be retained in a data lake and then transformed into curated warehouse tables that Power BI can query for models, reports, and visualizations. Together, these services support economical large-volume storage and performant BI analysis. See Introduction to Azure Data Lake Storage and Azure Synapse dedicated SQL pool overview.

QUESTION NO: 31

Your company plans to host an application by using Azure App Service.

Which two tasks are managed by Microsoft as part of the platform as a service (PaaS) model? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Maintaining the operating system.

B. Managing application code.

C. Maintaining the web hosting platform.

D. Classifying application data.

E. Managing application user permissions.

ANSWER: A C

Explanation:

When an application is hosted in Azure App Service, **Maintaining the operating system.** and **Maintaining the web hosting platform.** are Microsoft responsibilities under the platform as a service model. Azure App Service supplies a fully managed environment for hosting web applications, REST APIs, and background workloads. Microsoft operates and maintains the underlying datacenter infrastructure, physical servers, networking, host operating systems, and the web-hosting platform that makes application deployment and execution possible. This includes applying platform-level updates, patches, reliability improvements, and security maintenance to the managed service components. As a result, an organization can focus its operational effort on building, configuring, deploying, and monitoring its application rather than administering web servers or performing operating-system maintenance. This managed-platform model is the key distinction between App Service and infrastructure as a service, where the customer would administer the guest operating system and

web-server software. Microsoft's Azure App Service documentation describes App Service as a managed platform for web apps, while its shared-responsibility guidance explains that cloud providers manage the physical infrastructure and, for PaaS, much of the platform layer. See Azure App Service overview and Shared responsibility in the cloud.

QUESTION NO: 32

This question is part of a series of related scenarios. Each question in the series contains a unique solution that may or may not achieve the stated goals. Some questions may have multiple correct solutions, while others may have none. Once you submit an answer, you cannot revisit the question, and they will not appear in review screens.

Your role as an Azure administrator involves running a PowerShell script to create resources in Azure.

You need to recommend a suitable computer configuration to execute the script.

Solution: Execute the script using a computer with Chrome OS, utilizing Azure Cloud Shell.

Does this solution meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes is correct because Azure Cloud Shell provides an Azure-hosted, browser-based command-line environment, including a PowerShell experience. A computer running Chrome OS can access the Azure portal and Cloud Shell through a supported web browser, so the device does not need a locally installed PowerShell environment or Azure PowerShell modules. After starting Cloud Shell and selecting PowerShell, the administrator can sign in to the appropriate Azure subscription and run Azure PowerShell commands or scripts to create and manage resources.

Cloud Shell is designed to make Azure administration available from virtually any internet-connected device. The required command-line tooling is maintained in the managed Cloud Shell environment rather than on the Chrome OS device itself. Scripts can be entered directly, uploaded, or stored persistently through the Cloud Shell storage integration, enabling them to be executed during later sessions. Therefore, using Chrome OS as the client platform together with Azure Cloud Shell satisfies the requirement to execute a PowerShell script for creating Azure resources. Microsoft documents that Cloud Shell supports PowerShell and is accessible through a web browser. See the Azure Cloud Shell overview and the PowerShell in Azure Cloud Shell quickstart.

QUESTION NO: 33

Which two features or services can be integrated with Azure Monitor' Each correct answer presents part of the solution.

NOTE: Each correct answer is worth one point.

A. Application Insights.

B. Azure Advisor.

C. Azure status.

D. Azure Service Health.

E. Log Analytics.

ANSWER: A E

Explanation:

Application Insights. and **Log Analytics.** are integrated Azure Monitor capabilities. Application Insights is Azure Monitor's application performance monitoring capability. It collects and analyzes telemetry from live applications, including requests, dependencies, exceptions, response times, availability-test results, and distributed traces. This enables teams to assess application performance, identify failures, and investigate the behavior of cloud-hosted and hybrid applications using Azure Monitor's monitoring experience.

Log Analytics provides the workspace-based Azure Monitor Logs capability. A Log Analytics workspace stores log and telemetry data collected from Azure resources, virtual machines, applications, agents, and supported external sources. Users can query and correlate this data by using Kusto Query Language, build visualizations and workbooks, and use the resulting queries as the basis for monitoring and alerting workflows. These capabilities complement each other: Application Insights supplies application-specific observability data, while Log Analytics provides centralized collection, retention, and analysis of monitoring logs. Microsoft identifies Azure Monitor Logs and Application Insights as core elements of the Azure Monitor platform. For more information, see the Azure Monitor overview and Application Insights overview.

QUESTION NO: 34

Who is eligible to use the Azure Total Cost of Ownership (TCO) calculator?

A. billing readers for an Azure subscription only

B. owners for an Azure subscription only

C. anyone

D. all users who have an account in Azure Active Directory (Azure AD) that is linked to an Azure subscription only

ANSWER: C

Explanation:

anyone is correct because the Azure Total Cost of Ownership (TCO) calculator is a publicly accessible web-based estimation tool. Its purpose is to help organizations develop a business case for migration by comparing estimated costs for equivalent on-premises infrastructure and Azure resources. A user enters assumptions about existing infrastructure, such as servers, databases, storage, and networking, and the calculator produces an estimated comparison and potential savings report.

The calculator supports planning before an organization has deployed resources in Azure. It does not operate on Azure resources, retrieve subscription billing data, or require permissions to manage a subscription. Therefore, prospective customers, technical professionals, business stakeholders, partners, and organizations evaluating Azure can use it without an Azure subscription, a Microsoft Entra tenant account, or a specific Azure role. This broad availability makes the tool useful during the early strategy and assessment stages of cloud adoption, when an organization may be estimating potential value before establishing its Azure environment.

Microsoft describes the calculator as a way to estimate the cost savings from migrating workloads to Azure. For details, see the Azure TCO calculator overview and Microsoft Learn: Compare costs by using the TCO calculator.

QUESTION NO: 35

This question requires that you evaluate the underlined text to determine if it is correct. You can create an Azure support request from support.microsoft.com.

Instructions: Review the underlined text. If it makes the statement correct, select "œNo change is needed." If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed.

B. the Azure portal.

C. the Knowledge Center.

D. the Security & Compliance admin center.

ANSWER: B

Explanation:

the Azure portal. is correct because Azure support requests are created and managed in the Azure portal within the context of the relevant Azure subscription. After signing in, a user opens **Help + support** and selects **Create a support request**. The portal guides the user through choosing an issue type, specifying the affected Azure service or resource, supplying technical and contact details, and selecting an available severity level based on the subscription's support plan. Associating the request with the subscription is important because Azure support eligibility, support-plan benefits, and resource details are evaluated for that subscription. The same Help + support experience enables users to view request status, exchange updates with Microsoft support, and manage existing cases. Microsoft's official guidance identifies the Azure portal as the place to create a new Azure support request and provides the specific navigation path. For subscription and billing support, users may also use this support-request workflow even without a paid technical support plan. See Create an Azure support request and Azure supportability and Help + support.

QUESTION NO: 36

What types of data can be encrypted by Azure Information Protection?

A. network traffic

B. documents and email messages

C. an Azure Storage account

D. an Azure SQL database

ANSWER: B

Explanation:

Azure Information Protection is designed to protect unstructured business information in documents and email messages. In its current Microsoft Purview integration, administrators configure sensitivity labels that can classify content and apply encryption by using Azure Rights Management capabilities. When users apply a label configured for encryption to a supported file or email message, the content is encrypted and receives persistent usage rights. Those rights can define which users or groups may access the protected item and what actions they can take, such as viewing, editing, printing, copying, or forwarding it.

This protection is content-centric: it stays with the document or email message even after it is copied, sent outside the organization, or stored in a different supported location. Access is evaluated when a recipient attempts to open the protected content, helping organizations safeguard sensitive information wherever it travels. This scenario is distinct from infrastructure-level encryption services; Azure Information Protection focuses on applying classification, labeling, and rights management protections to supported files and email. Microsoft documents the encryption settings available for sensitivity labels, including permission assignment and content-access controls, at Encryption for sensitivity labels. For an overview of Azure Information Protection and its role in protecting documents and emails, see What is Azure Information Protection?.

QUESTION NO: 37

Which two types of customers qualify to utilize Azure Government for developing a cloud solution? Each correct answer represents a complete solution.

Note: Each correct selection is worth one point.

A. a Canadian government contractor

B. a European government contractor

C. a United States government entity

D. a United States government contractor

E. a European government entity

ANSWER: C D

Explanation:

Azure Government is a separate Microsoft cloud environment intended for eligible United States public-sector customers that need to meet specific security, compliance, and data-residency requirements. A United States government entity qualifies because Azure Government is available to US federal, state, local, and tribal government organizations. These organizations can build and operate workloads in an environment designed to support public-sector requirements, including compliance offerings relevant to government workloads and operations performed by screened US personnel.

A United States government contractor also qualifies when it supports an eligible US government customer. Contractors may require the same dedicated cloud boundary to develop, host, or manage solutions involving government data and workloads. Microsoft validates customer eligibility before granting access to Azure Government, helping ensure that the environment is used by authorized government entities and their qualifying partners. This model enables an agency and its approved contractor to use compatible Azure Government services while addressing the agency's applicable compliance and operational requirements. Microsoft's eligibility guidance identifies US government entities and their partners as the customer groups for this cloud environment. For details, see Azure Government documentation overview and Compare Azure Government and global Azure.

QUESTION NO: 38

Your company has an on-premises network that contains multiple servers.

The company plans to reduce the following administrative responsibilities of network administrators: Backing up application data

Replacing failed server hardware Managing physical server security Updating server operating systems

Managing permissions to shared documents

The company plans to migrate several servers to Azure virtual machines.

You need to identify which administrative responsibilities will be reduced after the planned migration.

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

A. Replacing failed server hardware.

B. Backing up application data.

C. Managing physical server security.

D. Updating server operating systems.

E. Managing permissions to shared documents.

ANSWER: A C

Explanation:

Azure virtual machines are an infrastructure as a service offering. With this model, Microsoft is responsible for the underlying physical cloud infrastructure, while the customer retains responsibility for the guest operating system, applications, data, identity, and access configurations within each virtual machine. Therefore, **Replacing failed server hardware.** is reduced after migration because Microsoft operates the physical host servers in its datacenters. If a host component fails, Microsoft performs the maintenance, repair, and replacement necessary to maintain the Azure platform; the customer does not manage or replace that physical hardware.

Managing physical server security. is likewise reduced because Microsoft secures its datacenters and the physical infrastructure that supports Azure services. This includes the facilities, physical host hardware, and associated environmental and access protections. The organization continues to secure its workloads through measures such as

operating system configuration, identity controls, and network security settings, but it no longer has responsibility for physically protecting the servers hosting its virtual machines. Microsoft documents these boundaries in the Azure shared responsibility model and describes Azure virtual machines as customer-managed guest operating systems running on Microsoft-managed infrastructure in the Azure Virtual Machines overview.

QUESTION NO: 39

You manage an on-premises network infrastructure comprising 100 servers. Your task is to propose a solution that offers additional resources to your users while minimizing both capital and operational expenditure costs. Which solution should you recommend?

A. a complete migration to the public cloud

B. an additional data center

C. a private cloud

D. a hybrid cloud

ANSWER: D

Explanation:

A hybrid cloud is the appropriate recommendation because it lets the organization retain and use its existing on-premises server estate while extending capacity into Azure as demand requires. This model avoids a large immediate capital investment in additional physical servers, storage, networking equipment, and data-center space. Instead, the organization can obtain cloud resources through consumption-based pricing, paying for capacity when it is used and scaling it back when it is no longer needed.

It can also lower operational costs by reducing the need to deploy, maintain, power, cool, patch, and refresh enough on-premises hardware to accommodate occasional peak demand. Workloads can remain on-premises where that is appropriate, while eligible applications, backup capacity, disaster recovery resources, or burst workloads can use Azure services. This enables a gradual, practical expansion of the existing environment rather than requiring a disruptive replacement of all infrastructure at once. Microsoft defines hybrid cloud as computing that combines public-cloud and on-premises resources, enabling organizations to use the most suitable location for each workload. Azure also provides flexible scaling and pay-as-you-go consumption for cloud resources. See Microsoft Learn: Hybrid and multicloud scenarios and Microsoft Learn: Describe cloud models.

QUESTION NO: 40

You need to grant a user permission to manage resources in a resource group named RG1. The user must not be able to manage access to the resources.

Which two Azure roles can you assign to the user? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Owner.

B. Contributor.

C. User Access Administrator.

D. Virtual Machine Contributor.

E. Reader.

ANSWER: B D

Explanation:

The **Contributor** role permits a user to create and manage Azure resources, but it does not allow the user to assign Azure roles or manage access to resources. Assigning Contributor at the RG1 scope grants the required permissions only for resources in that resource group.

The **Virtual Machine Contributor** role is also a complete solution when the resources to be managed are virtual machines. This role permits management of virtual machines and related VM operations without granting permission to manage role assignments. Azure RBAC roles can be assigned at management group, subscription, resource group, or resource scope.

The Owner and User Access Administrator roles include permissions to manage access and therefore do not meet the requirement. For more information, see Azure built-in roles and Azure RBAC overview.

QUESTION NO: 41

You are planning to deploy multiple Azure virtual machines, and you need to ensure that the services running on these virtual machines remain available even if a single data center experiences a failure. What are two possible solutions that would meet this requirement? Each correct answer represents a complete solution.

A. Deploy the virtual machines to two or more availability zones.

B. Deploy the virtual machines to two or more resource groups.

C. Deploy the virtual machines to a scale set.

D. Deploy the virtual machines to two or more regions.

ANSWER: A D

Explanation:

Deploy the virtual machines to two or more availability zones. Availability Zones are physically separate locations within an Azure region, and each zone consists of one or more datacenters with independent power, cooling, and networking. Distributing workload instances across zones protects the service from a failure that affects a single datacenter or zone. To deliver application availability, the independently deployed instances should also be reachable through an appropriate load-balancing design that sends traffic to healthy virtual machines. Deploy the virtual machines to two or more regions. Azure regions are geographically separate areas containing one or more datacenters, so deploying equivalent workload instances in multiple regions also ensures that a failure of one datacenter does not eliminate all service instances. A production multi-region architecture normally includes global traffic distribution, such as Azure Front Door or Azure Traffic Manager, plus a data replication and failover approach appropriate to the application. These deployment approaches provide distinct fault domains for the workload and therefore meet the stated availability requirement. Microsoft documents that Availability Zones protect applications and data from datacenter failures, while cross-region architectures provide resiliency across geographically separate Azure regions. See the Availability Zones overview and Azure cross-region replication guidance.

QUESTION NO: 42

This question requires you to assess the accuracy of the underlined text.

The statement is: "Azure Key Vault is used to store secrets for Azure Active Directory (Azure AD) user accounts."

Directions: Evaluate the underlined text. If it is correct, select “No change is needed.” If it is incorrect, choose the option that accurately completes the statement.

A. No change is needed

B. Azure Active Directory (Azure AD) administrative accounts

C. Personally Identifiable Information (PII)

D. server applications

ANSWER: D

Explanation:

server applications accurately completes the statement. Azure Key Vault is an Azure service designed to securely store and manage sensitive application values, including secrets, cryptographic keys, and certificates. In this context, a secret is typically a value that a server-side workload needs to access while running, such as a database password, connection string, API key, shared access signature, or token. A web application, API, background worker, virtual-machine-hosted service, or automation process can retrieve these values from Key Vault at runtime rather than placing them in source code, deployment artifacts, or application configuration files.

Key Vault supports identity-based access through Microsoft Entra ID, including managed identities, so an application can authenticate without developers having to embed credentials in code. It also provides centralized management capabilities for secret creation, expiration, rotation, monitoring, and auditing. This supports a security model in which sensitive values are separated from the application and access is limited to authorized workloads. Microsoft documents Key Vault as a service for storing secrets such as passwords and connection strings used by applications. See Azure Key Vault basic concepts and About Azure Key Vault secrets.

QUESTION NO: 43

You have developed an Azure web application and you need to manage its settings from an iPhone device. Which two Azure management tools can you utilize? Select each correct option as each represents a complete solution. NOTE: Each correct selection is worth one point.

A. Azure CLI

B. the Azure portal

C. Azure Cloud Shell

D. Windows PowerShell

E. Azure Storage Explorer

ANSWER: B C

Explanation:

The Azure portal is correct because it is Azure's web-based management interface and can be accessed from a browser on an iPhone. After authenticating, an administrator can open the Azure App Service resource for the web application and manage supported settings, including application settings, connection strings, deployment configuration, networking configuration, and monitoring-related settings. Since the portal is hosted by Microsoft and delivered through the browser, the iPhone does not need a locally installed Azure administration client.

Azure Cloud Shell is also correct because it provides an authenticated Azure command-line environment in a browser. Cloud Shell supports both Bash and PowerShell experiences and includes Azure management tooling, enabling an administrator to run commands for Azure resources from an iPhone browser without installing the Azure CLI or PowerShell locally on iOS. For example, supported Azure commands can retrieve or update App Service configuration after the user signs in to the appropriate Azure subscription. Cloud Shell can be launched from the Azure portal or accessed directly through a browser, making it a complete mobile-accessible management solution. Microsoft describes the portal as a unified web-based management experience and Cloud Shell as a browser-accessible shell for managing Azure resources. See Azure portal overview and Azure Cloud Shell overview.

QUESTION NO: 44

You are planning to migrate a web application to Microsoft Azure. The application is accessed by external users.

You need to recommend a cloud deployment solution that minimizes the administrative effort required to manage the web application.

Which deployment solution should you recommend?

A. Software as a Service (SaaS)

B. Platform as a Service (PaaS)

C. Infrastructure as a Service (IaaS)

D. Database as a Service (DaaS)

ANSWER: B

Explanation:

Platform as a Service (PaaS) is the appropriate recommendation because it provides a managed environment for deploying and operating a custom web application while minimizing the administrative work required from the application team. In Azure, a PaaS service such as Azure App Service enables an organization to deploy web application code and configure the application without managing the underlying virtual machines, operating systems, web-server software, or much of the supporting platform infrastructure. This is particularly suitable for an internet-facing application because Azure App Service is designed to host web apps, APIs, and back-end services with built-in platform capabilities for scaling, availability, monitoring, authentication integrations, and secure deployment practices.

With PaaS, the organization retains responsibility for the application itself: its code, configuration, data, identity design, and application-level security. Azure manages the platform components and performs infrastructure-related operational tasks such as maintaining the host environment and applying platform updates. This shared-responsibility balance allows the team to focus on delivering features and maintaining the web application rather than performing routine server administration. Microsoft describes PaaS as a complete development and deployment environment in the cloud, and Azure App Service as a fully managed service for hosting web applications. For more information, see the Azure App Service overview and Microsoft's SaaS, PaaS, and IaaS comparison.

QUESTION NO: 45

You need to ensure that all existing and new resources in a resource group have a tag named CostCenter.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Assign an Azure Policy that adds the CostCenter tag when it is missing.

B. Create a remediation task.

C. Create a delete lock on the resource group.

D. Assign the Reader role to all users.

E. Create an availability set.

ANSWER: A B

Explanation:

Assign an Azure Policy that adds the CostCenter tag when it is missing, scoped to the target resource group. Azure Policy's *modify* effect can add or update tags as resources are created or updated, enabling ongoing governance for new resources and for existing resources whenever they are subsequently modified. The policy assignment requires a managed identity with appropriate permissions to make tag changes, such as the permissions provided through the Tag Contributor role at the applicable scope.

Create a remediation task after assigning the policy to address resources that already exist and have not undergone a create or update operation since the assignment. A remediation task evaluates resources identified as noncompliant and uses the policy assignment's managed identity to apply the modify operation, adding the CostCenter tag where it is absent. Together, the policy assignment provides continuous enforcement for future resource operations, while remediation brings the current set of resources into the required tagged state. Microsoft documents tag modification behavior and required identity configuration in its Azure Policy modify effect documentation, and describes how to use remediation tasks for existing noncompliant resources in Remediate non-compliant resources with Azure Policy.

QUESTION NO: 46 - (SIMULATION)

(_____) in Azure Firewall enables users on the internet to access a server on a virtual network.

ANSWER: See the explanation for the answer

Explanation:

Answer: Destination network address translation (DNAT)

In Azure Firewall, a **DNAT rule** allows inbound connections from the internet to reach a server hosted on an Azure virtual network. The rule translates the firewall's public IP address and specified port to the private IP address and port of the internal server.

QUESTION NO: 47

Note: This question is part of a series that presents a scenario. Each question in the series contains a unique solution that might or might not meet the stated goals. Once you answer a question, you will NOT be able to return to it. Therefore, these questions will not appear in the review screen.

Your company is planning to purchase an Azure subscription. According to the company's support policy, the Azure environment must offer the option to access support engineers through phone or email.

You need to identify which support plan satisfies the support policy requirement.

Solution: Recommend a Basic support plan.

Does this solution meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because the Basic support plan does not provide access to Microsoft technical support engineers for Azure technical issues. Basic support is included with an Azure subscription and supplies self-help capabilities, including access to Azure documentation, Microsoft Learn, Azure Advisor, service-health information, and community resources. It also allows support requests for billing and subscription-management matters, but those services do not satisfy a requirement for technical engagement with support engineers.

The policy requires the organization to have the option to contact Azure support engineers by phone or email. Microsoft provides technical support through paid Azure support plans, where customers can create technical support requests and work with support professionals based on the plan's stated availability and response-time targets. Therefore, the organization must select an eligible paid support plan rather than Basic if it needs technical support access. The specific paid plan should be chosen according to operational requirements such as business-hours or around-the-clock coverage, severity level, and required initial response times.

Microsoft's comparison of available offerings is available at Azure support plans. For details on opening and managing technical support requests, see Create an Azure support request.

QUESTION NO: 48

Your Azure environment contains multiple Azure virtual machines.

You need to ensure that a virtual machine named VM1 is accessible from the Internet over HTT

P. What are two possible solutions? Each correct answer presents a complete solution.

A. Modify a DDoS protection plan.

B. Modify an Azure firewall.

C. Modify an Azure Traffic Manager profile.

D. Modify a network security group (NSG).

E. What are two possible solutions? Each correct answer presents a complete solution.

ANSWER: B D

Explanation:

Modify an Azure firewall. is a complete solution because Azure Firewall can provide a controlled public ingress point through its public IP address. A destination network address translation (DNAT) rule can translate inbound TCP requests on port 80 received by the firewall to VM1's private IP address and port 80. With appropriate firewall policy rules and routing, the firewall forwards the HTTP traffic to the virtual machine while retaining centralized inspection and control of Internet-facing access.

Modify a network security group (NSG). is also a complete solution when VM1 is assigned a public IP address or is otherwise exposed through a public endpoint. An inbound NSG security rule can allow TCP port 80 from Internet sources to the VM's network interface or subnet. NSGs are stateful, so response traffic for an allowed inbound HTTP session is automatically permitted. The rule can be scoped to all Internet addresses or, preferably, limited to known source ranges where appropriate.

Microsoft describes NSG inbound rules, their association with subnets and network interfaces, and their stateful behavior in the Network security groups overview. For publishing inbound traffic through Azure Firewall using DNAT, see Tutorial: Filter inbound traffic with Azure Firewall DNAT.

QUESTION NO: 49

Which Azure service provides a set of version control tools to manage code?

A. Azure Repos.

B. Azure DevTest Labs.

C. Azure Storage.

D. Azure Cosmos DB.

ANSWER: A

Explanation:

Azure Repos provides the set of version-control tools used to manage source code during the development lifecycle. It is a service within Azure DevOps and supports both Git repositories and Team Foundation Version Control (TFVC), enabling a team to choose the version-control approach that fits its requirements. Teams can centrally store code, maintain a complete history of changes, create branches, merge contributions, and use pull requests for collaborative review. These capabilities allow multiple developers to work on the same codebase while retaining traceability and control over changes. Azure Repos also supports branch policies, which can enforce practices such as required reviewers, linked work items, and successful build validation before code is merged into protected branches. This helps teams implement consistent quality and governance practices around their source code. Microsoft explicitly defines Azure Repos as a set of version-control tools for managing code and identifies Git and TFVC as the supported repository types. Learn more in What is Azure Repos? and What is Azure DevOps?.

QUESTION NO: 50

You plan to migrate a web application to Azure. The web application is accessed by external users. You need to recommend a cloud deployment solution to minimize the amount of administrative effort used to manage the web application.

What should you include in the recommendation?

A. software as a service (SaaS).

B. platform as a service (PaaS).

C. infrastructure as a service (IaaS).

D. database as a service (DaaS) .

ANSWER: B

Explanation:

Platform as a service (PaaS) is the appropriate recommendation because it provides a managed application hosting environment while reducing the operational work required from the customer. For an externally accessible web application, Azure App Service is a typical PaaS solution: the organization deploys its web application code and manages application configuration and data, while Azure manages the underlying servers, operating systems, runtime infrastructure, platform patching, and much of the availability and scaling infrastructure. This lets the team focus on the application rather than administering virtual machines and web-server software.

Azure App Service is designed to host HTTP-based web applications, REST APIs, and mobile back ends. It supports capabilities commonly needed for public web apps, such as custom domains, TLS/SSL certificates, authentication, deployment slots, autoscaling, monitoring integration, and built-in load balancing. These platform-managed features minimize administrative overhead while retaining control over the application itself. Microsoft's shared-responsibility model identifies PaaS as a model in which the cloud provider assumes management of more of the technology stack than the customer, which directly satisfies the stated goal. See the Azure App Service overview and Microsoft Learn: describe cloud service types.

QUESTION NO: 51

You manage an Azure environment encompassing 10 virtual networks and 100 virtual machines. Your objective is to restrict all inbound traffic to these Azure virtual networks. **Which configuration should you employ?**

A. one Azure firewall

B. 10 Azure ExpressRoute circuits

C. 10 virtual network gateways

D. one application security group (ASG)

ANSWER: A

Explanation:

One Azure firewall is the appropriate configuration because Azure Firewall is a fully managed, stateful firewall service designed to centrally inspect, filter, and log network traffic for Azure workloads. In an environment with multiple virtual networks, it can be deployed as a shared security service in a hub virtual network. The other virtual networks can then operate as spokes, with routing configured so that relevant traffic traverses the firewall before reaching protected resources. This design provides a single, consistent enforcement point for inbound-access policies across the environment rather than requiring an independent perimeter security deployment for each virtual network or virtual machine.

Azure Firewall supports network rules for controlling traffic by protocol, source, destination, and port, as well as application rules for approved outbound web traffic. For published inbound workloads, it can apply destination NAT rules and filtering policies to allow only explicitly authorized connections. Its centralized logging and integration with Azure monitoring services also improve visibility into traffic and security events. Microsoft describes Azure Firewall as a cloud-native network firewall service that protects Azure virtual network resources, while its hub-and-spoke architecture guidance shows how centralized security services can serve multiple peered virtual networks. See Azure Firewall overview and Hub-spoke network topology in Azure.

QUESTION NO: 52

Your company uses tags on Azure resources.

Which two tasks can be performed by using tags? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Organize resources by business metadata.

B. Group costs in cost analysis.

C. Grant users access to Azure resources.

D. Prevent resources from being deleted.

E. Encrypt data stored in a resource.

ANSWER: A B

Explanation:

Organizing resources by business metadata is a primary use of tags. A tag consists of a name and value pair, such as `Department:Finance` or `Environment:Production`. Tags help identify resource ownership, workload purpose, cost center, and other business information.

Grouping costs in cost analysis is also supported by tags. When tags are applied consistently, Azure Cost Management can group and filter usage and costs by tag values, which helps allocate cloud costs to departments, projects, or applications. Tags do not grant permissions or prevent resource deletion. Microsoft documents tag usage in the Use tags to organize your Azure resources and management hierarchy.

QUESTION NO: 53

You need to control Azure spending and receive a notification when the forecasted cost for a subscription exceeds a defined amount.

Which two features should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a budget.

B. an alert action group.

C. a management lock.

D. an availability set.

E. an Azure Advisor recommendation.

ANSWER: A B

Explanation:

A **budget** in Azure Cost Management defines a spending threshold for a scope such as a subscription or resource group. A budget can be evaluated against actual or forecasted costs and can track spending over a monthly, quarterly, or annual period.

An **alert action group** provides the notification mechanism for a budget alert. Budget alert rules can send notifications to specified email recipients and can integrate with action groups for additional notification or automation actions. Together, a budget and an alert action group meet the requirement to define a threshold and notify stakeholders when forecasted spending reaches it.

Azure Advisor can make cost optimization recommendations, but it does not define spending thresholds. A management lock prevents changes or deletion and is unrelated to cost notification. See Create and manage Azure budgets and Create and manage action groups.

QUESTION NO: 54

You need to identify security recommendations for Azure resources and view a secure score for an Azure subscription.

What should you use?

A. Azure Advisor.

B. Microsoft Defender for Cloud.

C. Azure Service Health.

D. Microsoft Sentinel.

ANSWER: B

Explanation:

Microsoft Defender for Cloud provides security posture management for Azure resources. It continuously assesses resources against security recommendations and calculates a secure score that helps an organization understand and prioritize improvements to its security posture.

Defender for Cloud can also provide workload protection capabilities when applicable Defender plans are enabled. Azure Advisor includes security recommendations, but Defender for Cloud is the service that provides the secure score and cloud security posture management experience. See the Microsoft Defender for Cloud overview.

QUESTION NO: 55

What tool can you use in Azure to identify and optimize underutilized or unused virtual machines?

A. Azure Advisor

B. Azure Cost Management + Billing

C. Azure reservations

D. Azure Policy

ANSWER: A

Explanation:

Azure Advisor is the appropriate tool because it analyzes Azure resource configurations and usage telemetry and provides personalized recommendations aligned with Microsoft best practices. For virtual machines, its cost recommendations can identify instances that have sustained low CPU utilization and low network usage, indicating that they may be idle or oversized for their workloads. Advisor then provides actionable optimization guidance, such as shutting down or deleting an unused virtual machine, resizing it to a more suitable SKU, or otherwise reducing avoidable cost. This makes Advisor especially useful when the goal is not merely to view spending, but to discover specific resources whose observed usage suggests an opportunity for improvement. Recommendations are displayed in the Azure portal and can also be retrieved through Azure APIs, enabling teams to incorporate the guidance into cost-governance and operational processes across subscriptions. Azure Advisor evaluates several optimization dimensions, including Cost, Reliability, Security, Operational Excellence, and Performance; the underutilized-VM scenario is addressed by its Cost recommendations. Microsoft documents Advisor as the service for optimizing Azure deployments and specifically describes recommendations for idle and underutilized virtual machines in its cost guidance. See the Azure Advisor overview and Azure Advisor cost recommendations.

QUESTION NO: 56

You have 50 virtual machines hosted on-premises and 50 virtual machines hosted in Azure. The on-premises 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:

hybrid is correct because the environment combines on-premises infrastructure with resources running in the public cloud, Azure, and those environments are connected. A hybrid cloud model is specifically designed to allow an organization to use both its private, on-premises environment and public cloud services together. Connectivity between the virtual machines enables workloads, applications, identities, and data to operate across the two environments as part of a broader integrated architecture. For example, an organization may retain some systems on-premises because of latency, compliance, or dependency requirements while placing other workloads in Azure to gain scalability, availability, and access to cloud services. The number of virtual machines in either location does not determine the model; the defining characteristic is the combined use of on-premises and cloud resources. Microsoft describes hybrid cloud computing as using both public cloud and private cloud/on-premises infrastructure, with the ability to share data and applications between them. Azure provides networking and hybrid-management services that support this approach. See Azure hybrid connectivity guidance and Microsoft's hybrid cloud overview.

QUESTION NO: 57

You have set up an Azure environment with various virtual machines, and you intend to enable communication between client computers in your on-premises network and your Azure virtual machines.

Which two Azure resources should you include in your implementation plan to facilitate this communication? Each correct choice is essential to the solution.

Note: Each correct selection is worth one point.

A. a load balancer

B. a virtual network

C. an application gateway

D. a gateway subnet

E. a virtual network gateway

ANSWER: D E

Explanation:

A gateway subnet and a virtual network gateway are required on the Azure side when implementing site-to-site VPN connectivity between an on-premises network and virtual machines hosted in an Azure virtual network. The gateway subnet is a dedicated subnet within the virtual network that must be named `GatewaySubnet`. Azure deploys the instances used by the virtual network gateway into this subnet. It must therefore exist before a gateway can be created, and Microsoft recommends sizing it to accommodate the gateway and potential future changes.

The virtual network gateway is the Azure-managed VPN endpoint. For a site-to-site connection, it establishes an encrypted IPsec/IKE tunnel to a compatible on-premises VPN device. Once the connection and routing are configured, clients on the on-premises network can communicate privately with resources, including virtual machines, in the Azure virtual network.

These two resources provide the required Azure VPN Gateway infrastructure for this hybrid connectivity scenario. For further details, see Microsoft's guidance on the gateway subnet requirements and its overview of Azure VPN Gateway.

QUESTION NO: 58

You are managing an Azure application that utilizes the services listed in the table below:



Each service has its own Service Level Agreement (SLA):

- Azure Virtual Machines: 99.9%
- Azure SQL Database: 99.99%

How would you determine the composite SLA for the application?

A. Max(0.999, 0.9999) = 0.9999 = 99.99%

B. 0.999 / 0.9999 = 4.9991 = 99.91%

C. 0.999 * 0.9999 = 0.9989001 = 99.89001%

D. Min(0.999, 0.9999) = 0.999 = 99.9%

ANSWER: C

Explanation:

For an application that requires both Azure Virtual Machines and Azure SQL Database to be available, calculate the composite SLA by multiplying the individual availability values expressed as decimals. The calculation is **0.999 × 0.9999 = 0.9989001**, which is **99.89001%**. This represents the availability of a dependency chain: the application remains available only when both required services are available at the same time. Because either service becoming unavailable affects the workload, adding required dependencies decreases the overall composite availability compared with the SLA of either individual component.

Composite SLA calculations are a key part of Azure workload reliability planning. They help teams set realistic availability targets and identify where redundancy, alternative architectures, or resiliency measures may be needed. The calculation assumes that the service availability values can be multiplied for the workload's dependency model, as is standard for a series of required components. Microsoft documents this approach for dependent services in its reliability and architecture guidance. See Establish reliability targets in the Well-Architected Framework and Business metrics and composite SLAs.

QUESTION NO: 59

What should you use to evaluate whether your company's Azure environment meets regulatory requirements?

A. Azure Security Center

B. Azure Advisor

C. Azure Service Health

D. Azure Knowledge Center

ANSWER: A

Explanation:

Azure Security Center is the correct answer; Microsoft has since renamed and expanded this service as Microsoft Defender for Cloud. Its cloud security posture management capabilities continuously assess Azure resources and configurations, then present the organization's security and compliance posture. In particular, the regulatory compliance experience enables

teams to evaluate resources against supported standards, regulations, and industry benchmarks. This is appropriate for determining whether an Azure environment aligns with applicable regulatory requirements because it provides a consolidated compliance view, identifies controls that need attention, and supplies actionable recommendations for remediation.

The compliance experience uses Azure Policy initiatives and assessments mapped to compliance standards. This mapping helps translate technical resource configurations into a view that compliance, security, and operations teams can use to monitor progress and prioritize work. Evaluation is continuous, rather than a point-in-time manual exercise, so the reported posture can reflect changes to deployed resources and policies. The current Microsoft documentation describes this functionality in the Regulatory compliance dashboard in Microsoft Defender for Cloud and explains the service's overall security-posture capabilities in What is Microsoft Defender for Cloud?.

QUESTION NO: 60

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:

Command Prompt and Windows PowerShell are both appropriate tools for running the Azure CLI on a Windows computer after the CLI has been installed. Azure CLI is a command-line tool whose commands begin with `az`, such as `az login` to authenticate to Azure and `az group list` to view resource groups. The installer makes the Azure CLI executable available through the system path, allowing these commands to be entered directly in either of these command-shell environments.

Command Prompt supports interactive execution of Azure CLI commands and is suitable for straightforward administrative tasks. Windows PowerShell can execute the same `az` commands and additionally provides shell capabilities useful for automation, including variables, scripting, and output processing. Azure CLI command syntax remains consistent regardless of which of these Windows shells the engineer uses. Microsoft's Windows installation guidance specifically shows Azure CLI use from PowerShell and Command Prompt, while Azure CLI documentation describes signing in and managing Azure resources through `az` commands. See Install the Azure CLI on Windows and Get started with Azure CLI.

QUESTION NO: 61

Which two features of services can be integrated with Azure Monitor? Each correct answer presents part of the solution.

NOTE Each correct answer is worth one point:

A. Azure Service Healt h

B. Azure Advisor

C. Azure status

D. Application Insights

E. Log Analytics

ANSWER: D E

Explanation:

Application Insights and **Log Analytics** are integrated Azure Monitor capabilities that provide complementary monitoring functions. Application Insights is Azure Monitor's application performance monitoring feature. It collects telemetry from live applications, including requests, dependencies, exceptions, availability results, and performance data. This enables teams to observe application behavior, investigate failures, and analyze end-user experience through Azure Monitor tools such as metrics, logs, dashboards, and alerts.

Log Analytics provides the workspace-based log data platform for Azure Monitor. A Log Analytics workspace stores and organizes collected log data from Azure resources, virtual machines, applications, and other connected sources. Administrators query this data by using Kusto Query Language, correlate records across sources, create visualizations, and build alert rules based on log-query results. Application Insights telemetry can also be workspace-based, allowing application data to be queried and correlated with broader operational logs in the same Azure Monitor environment. Together, these services support monitoring of both application-level telemetry and centralized log data. See Application Insights overview and Log Analytics workspace overview.

QUESTION NO: 62

You are planning to deploy several virtual machines on Azure.

Your goal is to ensure that the services running on these virtual machines remain operational even if there is a failure at a single data center.

Which two solutions can achieve this requirement? Each correct answer provides a complete solution.

NOTE: Each correct answer 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:

"Deploy the virtual machines to two or more availability zones" meets the requirement because Availability Zones are physically separate datacenter locations within an Azure region. Each zone has independent power, cooling, and networking infrastructure. Placing workload instances in more than one zone prevents a failure affecting one datacenter from taking down every instance of the service. The workload can then continue through instances in the unaffected zone, provided the application uses suitable traffic distribution and health monitoring.

"Deploy the virtual machines to two or more regions" also provides the required resilience. Separate Azure regions are geographically distinct locations, so a workload deployed redundantly across regions has isolation beyond an individual datacenter. A multi-region architecture can direct traffic or fail over to the healthy regional deployment when the other deployment is unavailable. This design also provides broader protection than the stated requirement, including resilience against outages that affect an entire region. Azure guidance recommends designing applications to use redundancy and appropriate failover routing when availability requirements extend across locations.

See the Azure Availability Zones overview and Azure cross-region resiliency guidance for details.

QUESTION NO: 63 - (DRAG DROP)

Match the term to the appropriate description.

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

NOTE: Each correct match is worth one point.



ANSWER:



Explanation:

The correct mappings describe three core identity concepts used throughout Azure and Microsoft Entra ID. Single sign-on (SSO) is the capability that allows a user to authenticate once with a set of credentials and then access multiple connected resources or applications without having to repeatedly sign in. This can apply across applications provided by different providers when they trust the same identity system. SSO improves the sign-in experience while allowing the organization to maintain centralized identity management. Microsoft provides an overview at What is single sign-on in Microsoft Entra ID.

Authorization concerns the access level available to an already identified user, group, managed identity, or service principal. It determines the actions that identity can perform and the resources it can reach. In Azure, this is commonly controlled through Azure role-based access control, which uses role assignments at scopes such as management groups, subscriptions, resource groups, and individual resources. The applicable permissions therefore reflect the assigned role and scope. More detail is available in Azure role-based access control overview.

Multi-factor authentication (MFA) uses multiple identity-verification elements during sign-in. These elements are generally based on something the person knows, has, or is. For example, a password together with a verification prompt in an authenticator application combines separate proofs of identity. The reference to one of several required elements reflects the factor-based nature of MFA, where the overall sign-in verification relies on more than one factor rather than a password alone. Microsoft explains this process at How Microsoft Entra multifactor authentication works.

QUESTION NO: 64

This question requires that you evaluate the underlined text to determine if it is correct.

A support plan solution that gives you best practice information, health status and notifications, and 24 access to billing information at the lowest possible cost is a Standard support plan.

Instructions: Review the underlined text. If it makes the statement correct, select "œNo change is needed" . If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed.

B. Developer.

C. Basic.

D. Premier.

ANSWER: C

Explanation:

Basic. is the correct replacement because Basic support is included with every Azure subscription at no additional charge, making it the lowest-cost support offering. It provides access to billing and subscription support around the clock for matters such as account, subscription, invoice, and payment issues. This is distinct from paid technical support, which is intended for assistance with Azure technical incidents and carries defined response-time targets.

The described best-practice information is available through Azure Advisor. Advisor analyzes Azure resource configurations and provides recommendations across areas such as reliability, security, performance, operational excellence, and cost optimization. Azure customers can also use Azure Service Health to view service issues, planned maintenance, health advisories, and relevant notifications affecting their resources. These self-help, health, and billing capabilities are available without purchasing a paid technical-support plan. Therefore, Basic meets the stated requirements while retaining the lowest possible cost. Microsoft's current support-plan comparison is available at Azure Support Plans. For details about the recommendations supplied by Advisor, see What is Azure Advisor?.

QUESTION NO: 65

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.

Your company plans to migrate all its data and resources to Azure.

The company's migration plan states that only platform as a service (PaaS) solutions must be used in Azure.

You need to deploy an Azure environment that supports the planned migration. Solution: You create an Azure App Service and Azure SQL databases.

Does this meet the goal?

A. Yes.

B. No.

ANSWER: A

Explanation:

Yes. Azure App Service and Azure SQL Database are both platform as a service offerings, so creating these resources satisfies a requirement to use only PaaS solutions. Azure App Service is a fully managed platform for hosting web applications, APIs, and background workloads. The customer deploys and configures the application, while Microsoft operates the underlying infrastructure and platform components required to run it.

Azure SQL Database is a fully managed relational database PaaS service. It provides a SQL Server database engine without requiring the organization to deploy or maintain database servers, virtual machines, operating systems, or physical infrastructure. The service includes managed capabilities such as built-in high availability, automated backups, patching, and maintenance. Using Azure App Service for an application tier and Azure SQL Database for its relational data tier is therefore a typical managed Azure architecture that remains within the PaaS service model.

Microsoft describes App Service as a service for hosting web applications and identifies Azure SQL Database as a fully managed PaaS database engine. See the Azure App Service overview and Azure SQL Database PaaS overview.

QUESTION NO: 66

What do you need to access Azure Cost Management?

A. a Dev/Test subscription

B. Software Assurance

C. an Enterprise Agreement (EA)

D. a pay-as-you-go subscription

E. appropriate Cost Management permissions at the relevant scope

ANSWER: E

Explanation:

You need appropriate permissions at the relevant Azure Cost Management scope to access cost information and Cost Management features. In practice, this commonly means being assigned the Cost Management Reader role, or a role with equivalent or greater permissions, on a scope such as a management group, subscription, resource group, billing account, billing profile, or enrollment account. The selected scope determines the costs the user can view, while the assigned role determines the actions they can perform. For example, Cost Management Reader enables a user to view cost data and configurations, whereas Cost Management Contributor also allows management of Cost Management settings such as budgets and exports. Access can also be granted through applicable billing roles for billing-account scopes. An Enterprise Agreement and a pay-as-you-go subscription are billing arrangements that can support Cost Management, but neither arrangement by itself gives an individual user authorization to see the organization's cost data. Microsoft emphasizes that Cost Management access is governed by scope and role-based permissions. See Understand and work with scopes and Assign access to Cost Management data.

QUESTION NO: 67

You have a resource group named RG1.

You plan to create virtual networks and app services in RG1.

You need to prevent the creation of virtual machines only in RG1. What should you use?

A. a lock.

B. an Azure role.

C. a tag.

D. an Azure policy.

ANSWER: D

Explanation:

An Azure policy is the appropriate control because Azure Policy can enforce rules governing which resource types may be created at a specified scope. Assign a policy to RG1 with a rule that targets the `Microsoft.Compute/virtualMachines` resource type and applies the `Deny` effect. Azure Resource Manager evaluates this policy during deployment, so a request to create a virtual machine in RG1 is rejected before the resource is provisioned. Because the assignment scope is the RG1 resource group, the restriction is limited to that resource group and does not affect virtual machine deployments elsewhere in the subscription.

This meets the requirement precisely: virtual networks and App Service resources can still be deployed in RG1 as long as they do not match the virtual-machine resource type condition. Azure Policy is designed to enforce organizational standards, control allowed resource configurations and types, and report compliance. Policy assignments can be scoped to management groups, subscriptions, resource groups, or individual resources, making a resource-group-level assignment suitable for this targeted restriction. Microsoft documents Azure Policy governance and assignment scopes in the Azure Policy overview, and describes how the `Deny` effect blocks noncompliant resource requests in the Azure Policy deny effect documentation.

QUESTION NO: 68

This question requires that you evaluate the underlined text to determine if it is correct.

If a resource group named RG1 has a delete lock, only a member of the global administrators group can delete RG1.

Instructions: Review the underlined text If it makes the statement correct, select " No change is needed. " If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed.

B. the delete lock must be removed before an administrator

C. an Azure policy must be modified before an administrator

D. an Azure tag must be added before an administrator

ANSWER: B

Explanation:

the delete lock must be removed before an administrator is correct because an Azure resource lock set to *CanNotDelete* protects the resource group and its contained resources from deletion. The lock is evaluated by Azure Resource Manager before a delete operation is completed, so normal Azure role assignments, including highly privileged administrative permissions, do not by themselves bypass the lock. An identity must first have sufficient authorization to remove the lock, such as the relevant Microsoft.Authorization/locks/delete permission. After the delete lock has been removed, an administrator with the required permissions can delete RG1. Locks are intended to prevent accidental or unintended deletion of critical cloud resources and are inherited by resources within a locked resource group. This behavior is distinct from identity-directory administration; membership in a directory administrative role is not the condition that permits deletion of a locked resource group. Microsoft documents that a CanNotDelete lock allows authorized users to read or modify a resource but prevents deletion until the lock is removed. For details, see Lock your Azure resources to protect your infrastructure and Azure resource provider operations.

QUESTION NO: 69

This question requires that you evaluate the underlined BOLD text to determine if it is correct. You can use Advisor recommendations in Azure to send email alerts when the cost of the current billing period for an Azure subscription exceeds a specified limit.

Instructions: Review the underlined text. If it makes the statement correct, select "œNo change is needed." If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed.

B. Access control (IAM).

C. Budget alerts.

D. Compliance .

ANSWER: C

Explanation:

Budget alerts. is correct because Azure Cost Management budgets are specifically designed to track spending against a defined monetary amount over a chosen period and scope, including an Azure subscription. When creating a budget, an organization selects the budget amount, recurrence period, and alert thresholds. Thresholds can be based on actual accumulated costs or forecasted costs, allowing stakeholders to receive notification as spending reaches the configured percentage or amount.

Budget alert notifications can be sent to designated email recipients, and they can also use Azure Monitor action groups for broader notification and automation scenarios. This supports proactive cost governance by giving subscription owners and financial stakeholders timely visibility into current-period spending relative to a specified limit. A budget alert is a monitoring and notification mechanism; it does not itself stop resources or prevent additional charges. Azure documents the configuration of budget amounts, time periods, alert thresholds, and email recipients in Create and manage budgets in Azure Cost Management. For further details on cost and budget alert behavior, see Use cost alerts to monitor usage and spending.

QUESTION NO: 70

Consider the following scenario: Your company's developers plan to deploy a substantial number of custom virtual machines each week. These virtual machines will also be decommissioned within the same week of deployment. Sixty percent of these virtual machines will have Windows Server 2016 installed, while the remaining forty percent will run on Ubuntu Linux. Your task is to ensure that the administrative effort involved in this process is minimized by utilizing an appropriate Azure service.

Solution Proposal: You propose using Azure Reserved Virtual Machine (VM) Instances.

Does this solution fulfill the objective?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. Azure Reserved VM Instances are a pricing benefit, not a deployment or lifecycle-management service. A reservation provides a discounted rate for eligible virtual-machine compute usage when an organization commits to a specific VM family, region, and term, typically one or three years. It does not create virtual machines, apply operating-system configurations, distribute Windows Server and Ubuntu deployments, or automatically remove machines after a short-lived workload completes.

The requirement focuses on reducing administrative work for a large and continually changing fleet of custom VMs that is created and decommissioned each week. This calls for automated provisioning and management, such as Azure Virtual Machine Scale Sets, together with reusable VM images and infrastructure-as-code tools where appropriate. Virtual Machine Scale Sets are intended to create and manage groups of load-balanced VMs and can automatically scale the number of instances in response to demand or a defined schedule. Reservations can potentially complement such an environment for predictable, sustained compute consumption, but they do not fulfill the stated operational automation objective.

Microsoft explains that VM reservations provide a billing discount in Save costs with Azure Reserved VM Instances. For automated management of a VM fleet, see the Azure Virtual Machine Scale Sets overview.

QUESTION NO: 71

You have an Azure subscription that contains multiple resource groups.

You need to identify which Azure roles can assign Azure roles to users at the resource group scope.

Which two roles should you identify? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Owner.

B. Contributor.

C. Reader.

D. User Access Administrator.

ANSWER: A D

Explanation:

Owner. and **User Access Administrator.** can assign Azure RBAC roles to users at the resource group scope when the role is assigned at that resource group, or at a parent scope such as the subscription. Azure role assignments require the `Microsoft.Authorization/roleAssignments/write` permission. The Owner built-in role grants full management access to Azure resources and includes the ability to manage access, allowing its holder to create role assignments within the applicable scope. User Access Administrator is the built-in role intended for managing access to Azure resources; it

includes permissions to manage user access and can create role assignments at its assigned scope and any child scopes. Scope inheritance is important: a role assignment made at a resource group applies to resources within that resource group, while a role assigned at subscription scope can be used to manage access across its child resource groups. These roles support Azure's role-based access control model, which lets organizations delegate resource administration and access management without granting unnecessary permissions beyond the required scope. See the Azure built-in roles documentation and Assign Azure roles using the Azure portal.

QUESTION NO: 72

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, white 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.

Your company has an Azure subscription that contains the following unused resources:

* 20 user accounts in Azure Active Directory (Azure AD)

* Five groups in Azure AD

* 10 public !P addresses

* 10 network interfaces

You need to reduce the Azure costs for the company.

Solution: You remove the unused public IP addresses.

Does this meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes is correct because Azure Public IP addresses are billable networking resources, and an allocated public IP can continue to incur charges even when it is not actively used by a workload. Deleting public IP addresses that are no longer required removes those resources from the subscription and stops the applicable Public IP address charges. This directly supports the goal of reducing Azure costs.

Azure pricing distinguishes between public IP configurations and regions, but the central cost-management principle remains the same: resources should be deallocated or deleted when they are no longer needed. A public IP address is a separately managed Azure resource, so simply leaving it unused does not provide a business benefit while it may still create billable consumption. Before deletion, an administrator should confirm that the addresses are not associated with virtual machines, load balancers, VPN gateways, firewalls, DNS records, or allowlists that require stable public addressing. In the stated scenario, the addresses are explicitly unused, so removing them is an appropriate cost-optimization action. See Azure Public IP Address pricing and Azure Public IP addresses documentation.

QUESTION NO: 73

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.

Your company has an Azure subscription that contains the following unused resources:

* 20 user accounts in Azure Active Directory (Azure AD)

* Five groups in Azure AD

* 10 public [P addresses

* 10 network interfaces

You need to reduce the Azure costs for the company.

Solution: You remove the unused groups.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because removing unused Microsoft Entra ID (formerly Azure AD) groups does not, by itself, reduce Azure consumption charges. Groups are directory objects used to organize users, assign access, and support administration; their existence is not a separately metered Azure resource charge. Therefore, deleting the five unused groups does not address a billable resource in the subscription and cannot meet the stated goal of reducing Azure costs.

Cost-reduction actions need to target services or resource configurations that incur charges. Microsoft Entra ID licensing is based on the selected identity edition and licensed users, rather than a per-group price. Group cleanup can still be a sound governance practice because it reduces administrative clutter, helps maintain accurate access management, and may simplify operations. However, it is not a direct Azure cost optimization measure in this scenario. Microsoft's identity pricing information is available in the Microsoft Entra pricing page, and Microsoft describes groups as directory resources for managing users and access in its groups overview.

QUESTION NO: 74

Your company is planning a deployment using Azure Database for PostgreSQ

L. The deployment should meet the following requirements:

- Up to 10 TB storage

- Azure Premium Storage

- Point-in-time-restore for up to 35 days

You need to select the appropriate deployment and pricing tier to meet these requirements and minimize costs. What should you select?

A. Azure Database for PostgreSQL Single Server Memory Optimized tier.

B. Azure Database for PostgreSQL Hyperscale (Citus).

C. Azure Database for PostgreSQL Flexible Server General Purpose tier.

D. Azure Database for PostgreSQL Single Server Basic tier .

E. The deployment should meet the following requirements:
- Up to 10 TB storage
- Azure Premium Storage
- Point-in-time-restore for up to 35 days
You need to select the appropriate deployment and pricing tier to meet these requirements and minimize costs. What should you select?

ANSWER: C

Explanation:

Azure Database for PostgreSQL Flexible Server General Purpose tier. is the appropriate selection because it provides the balanced compute, memory, and storage profile intended for most production and business workloads, helping minimize cost when an application does not require the specialized high-memory configuration of a memory-optimized service tier. Flexible Server General Purpose supports Premium SSD storage and can scale storage beyond the required 10 TB capacity. Therefore, it can accommodate the stated database size while providing the required storage performance characteristics.

Flexible Server also includes automated backups and point-in-time restore capabilities. Its backup retention period is configurable from 7 through 35 days, allowing the deployment to be configured with the full 35-day recovery window required by the scenario. Point-in-time restore creates a new server from backup data at a selected point within that configured retention period, supporting recovery from accidental data changes or deletion.

Configure the Flexible Server with sufficient General Purpose compute, allocate the needed Premium SSD storage capacity, and set backup retention to 35 days. This directly fulfills all stated technical requirements while selecting the standard balanced tier rather than paying for resources intended for more demanding memory-intensive workloads. Microsoft documents the supported tiers and storage capabilities in Azure Database for PostgreSQL Flexible Server service tiers and storage and backup retention and recovery behavior in Backup and restore in Flexible Server.

QUESTION NO: 75

You have an Azure environment.

You need to create a new Azure virtual machine from a tablet that runs the Android operating system.

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

A. Use Bash in Azure Cloud Shell.

B. Use PowerShell in Azure Cloud Shell.

C. Use the PowerApps portal.

D. Use the Security & Compliance admin center.

E. Use the Azure portal.

ANSWER: A B E

Explanation:

Use Bash in Azure Cloud Shell, Use PowerShell in Azure Cloud Shell, and Use the Azure portal are complete ways to create an Azure virtual machine from an Android tablet. The Azure portal is a web-based management interface, so a user can sign in through a supported browser on the tablet and use the virtual machine creation workflow to configure such settings as the subscription, resource group, region, image, VM size, administrator credentials, disks, and networking. Azure Cloud Shell is also browser-accessible and does not require the user to install Azure management tools locally on the Android device. Cloud Shell provides both Bash and PowerShell environments, each authenticated for Azure resource management. In the Bash environment, Azure CLI commands, including `az vm create`, can deploy a virtual machine. In the PowerShell environment, Azure PowerShell cmdlets, including `New-AzVM`, can create and configure a virtual machine. These methods are independently capable of submitting a VM deployment to Azure while being usable from a browser-based Android tablet session. Microsoft documents Cloud Shell as an authenticated browser-based shell with Bash and PowerShell experiences, and documents creating VMs through the Azure portal. See Azure Cloud Shell overview and Create a virtual machine in the Azure portal.

QUESTION NO: 76

Consider the following scenario: You have an Azure environment and need to create a new Azure virtual machine using a tablet that operates on the Android platform.

Proposed Solution: You utilize the Azure portal to accomplish this task.

Does this solution fulfill the requirement?

A. Yes

B. No

ANSWER: A

Explanation:

Yes. The Azure portal is Microsoft's web-based interface for managing Azure resources, rather than a Windows-only application that must be installed on a computer. A user can access it by signing in at https://portal.azure.com through a supported modern web browser with an Azure account that has sufficient permissions. Therefore, an Android tablet with internet access and a compatible browser can be used to initiate and manage the deployment of an Azure virtual machine.

From the portal, the user can open the Virtual machines service and use the creation workflow to select the subscription, resource group, Azure region, operating-system image, virtual machine size, authentication configuration, disks, networking, and management settings. When the configuration is submitted, Azure deploys the virtual machine through Azure Resource Manager. The portal is specifically designed to create, configure, monitor, and administer Azure resources through a graphical web experience. Although a tablet's smaller display and touch-based input may make some configuration steps less convenient than on a desktop device, those usability considerations do not prevent the portal from fulfilling the requirement. Microsoft describes the portal as a web-based, unified console for building, managing, and monitoring Azure resources; see the Azure portal overview.

QUESTION NO: 77

Consider you have an Azure environment composed of several Azure virtual machines. You are planning to set up a system that allows client computers on your on-premises network to connect and communicate with the Azure virtual machines. What Azure resources would you recommend creating to achieve this planned solution?

Please identify two Azure resources that should be part of your recommendation. Each correct selection earns one point.

A. a virtual network gateway

B. a load balancer

C. an application gateway

D. a virtual network

E. a gateway subnet

ANSWER: A D

Explanation:

a virtual network and **a virtual network gateway** are the core Azure resources for providing private connectivity from an on-premises network to Azure virtual machines through a site-to-site VPN. A virtual network supplies the private network boundary in Azure. It contains the virtual-machine subnets, defines the private IP address space, and provides the routing scope to which the on-premises network will connect. The virtual machines must be deployed to, or reachable through, this virtual network.

A virtual network gateway is the managed Azure gateway resource that terminates the VPN connection. For a site-to-site design, it establishes encrypted IPsec/IKE connectivity with a compatible on-premises VPN device. Once the connection is configured, on-premises clients can route traffic to the private addresses of the Azure virtual machines, provided that appropriate routes and security rules permit that traffic. The gateway is deployed into a dedicated subnet named *GatewaySubnet*, which is a necessary configuration component, but the primary resources requested for this architecture are the virtual network and the virtual network gateway.

Microsoft describes VPN Gateway and its site-to-site connectivity model in About VPN Gateway and provides deployment guidance in Create a site-to-site VPN connection.

QUESTION NO: 78

You plan to deploy an internet-facing web application that must distribute HTTP requests across multiple backend servers.

You need to identify Azure services that provide Layer 7 load-balancing capabilities.

Which two services should you identify? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Azure Application Gateway.

B. Azure Load Balancer.

C. Azure Front Door.

D. Azure Traffic Manager.

ANSWER: A C

Explanation:

Azure Application Gateway and **Azure Front Door** provide Layer 7 (application-layer) load balancing for HTTP and HTTPS workloads. Application Gateway is a regional web traffic load balancer that receives client HTTP/S requests and distributes them to a configured backend pool. Because it operates at Layer 7, it can make routing decisions using application-request properties, including the host name, URL path, and HTTP headers. It also supports TLS termination, cookie-based session affinity, and integration with Azure Web Application Firewall.

Azure Front Door provides global Layer 7 application delivery and load balancing. It accepts web traffic at Microsoft's global edge network and routes requests to appropriate backend origins based on health, latency, priority, and routing configuration. This makes it suitable for internet-facing applications with globally distributed users or backends, while retaining HTTP/S-aware capabilities such as path-based routing and TLS termination. Microsoft documents Application Gateway as a web traffic load balancer operating at Layer 7 and Front Door as a global application delivery network for HTTP/S applications. See Azure Application Gateway overview and Azure Front Door overview.

QUESTION NO: 79

Your company hosts an accounting application named App1, which is utilized by all your customers. The application experiences low usage during the first three weeks of every month, followed by significantly increased usage during the final week. Which feature of Azure Cloud Services aids in managing costs effectively given this fluctuating usage pattern?

A. high availability

B. high latency

C. elasticity

D. load balancing

ANSWER: C

Explanation:

Elasticity is the cloud characteristic that enables an application's resources to adjust to changing demand. For App1, the organization can increase the number or size of compute resources during the final week, when customer usage is high, and reduce that capacity throughout the quieter first three weeks. This aligns the resources being consumed with the application's actual workload instead of requiring the company to provision and pay for peak-level capacity for the entire month. Because the usage cycle is predictable, the organization can configure scheduled scaling actions in advance; it can

also use demand-based autoscale rules when appropriate. Azure Autoscale supports this model by adjusting resource capacity according to schedules, metrics, or rules, helping maintain application performance during busy periods while avoiding unnecessary consumption during low-demand periods. This ability to dynamically add or remove resources in response to workload requirements is elasticity, and it is a key cloud-computing benefit for efficient resource use and cost management. For further details, see Microsoft Learn: Scaling and partitioning and Microsoft Learn: Autoscale overview.

QUESTION NO: 80

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

I. 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 .

F. 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.

ANSWER: A C

Explanation:

Command Prompt. and **Windows PowerShell.** are both complete tools for running the Azure CLI after it has been installed on a Windows computer. Azure CLI is a cross-platform command-line tool whose commands use the `az` executable, such as `az login` to authenticate and `az group list` to retrieve resource groups. Because the installer makes the `az` command available through the system command path, it can be invoked interactively from a standard Windows command shell.

Command Prompt provides a direct command-line environment for entering Azure CLI commands. Windows PowerShell can also invoke the same `az` executable and is particularly useful when CLI commands are incorporated into administrative scripts or automation workflows using variables, loops, and pipelines. The choice of shell does not change the Azure CLI command syntax or the Azure resources that can be managed; both are supported ways to execute the installed CLI locally. Microsoft describes Azure CLI as a command-line tool for managing Azure resources and provides Windows installation guidance that includes opening a terminal and using the `az` command. See What is the Azure CLI? and Install Azure CLI on Windows.

QUESTION NO: 81

You need to deploy the same Azure resources repeatedly to multiple resource groups.

The deployments must be consistent and repeatable.

What should you use?

A. Azure Resource Manager templates.

B. Azure Policy initiatives.

C. network security groups (NSGs).

D. Azure Service Health alerts.

ANSWER: A

Explanation:

Azure Resource Manager templates are the correct choice. An ARM template is an infrastructure-as-code file that declaratively defines Azure resources, their configuration, and dependencies. The same template can be deployed repeatedly to create a consistent environment in different resource groups or subscriptions.

Azure Resource Manager processes template deployments as a group and helps ensure that dependent resources are deployed in the appropriate order. Templates can be deployed by using the Azure portal, Azure CLI, Azure PowerShell, or automated pipelines. Azure Policy governs resources, but it does not define and deploy the full resource configuration. For more information, see What are ARM templates?.

QUESTION NO: 82

This question requires that you evaluate the underlined BOLD text to determine if it is correct.

Azure Site Recovery provides fault tolerance for virtual machines.

Instructions: Review the underlined text. If it makes the statement correct, select “No change is needed.” If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed.

B. disaster recovery

C. elasticity

D. high availability

ANSWER: B

Explanation:

disaster recovery makes the statement correct because Azure Site Recovery is designed to support a business continuity and disaster recovery strategy. It continuously replicates supported workloads, including Azure virtual machines and on-premises virtual machines, to a secondary location. If a significant disruption affects the primary environment, such as a datacenter failure or a regional outage, administrators can perform a failover to bring the replicated workloads online in the recovery location. This enables the organization to restore application availability according to its recovery objectives.

Azure Site Recovery also provides capabilities that are central to disaster recovery operations: replication health monitoring, recovery points, test failovers that do not interrupt production, recovery plans for orchestrating multi-tier applications, planned and unplanned failovers, and failback after the primary location is available again. These features help organizations prepare for and recover from events that make the primary workload location unavailable. Microsoft explicitly identifies Azure Site Recovery as a disaster recovery service that helps keep applications running during outages. See the Azure Site Recovery overview and Microsoft disaster recovery guidance.

QUESTION NO: 83 - (HOTSPOT)

HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:



ANSWER:



Explanation:

The correct selections are No, Yes, and Yes. Access to Azure resources is controlled through Azure role-based access control (Azure RBAC). Azure RBAC assignments are made to security principals in Microsoft Entra ID, rather than being restricted to only individual user accounts. A security principal can be a user, a group, a service principal representing an application, or a managed identity representing an Azure workload. This means authorization can be granted to the appropriate identity type for a person, application, automation process, or Azure service. Microsoft documents the supported identity types and role-assignment model in its Azure RBAC overview.

Microsoft Entra ID also supports identity scenarios that span cloud and on-premises environments. Organizations can synchronize identities from on-premises Active Directory Domain Services to Microsoft Entra ID, enabling those users to authenticate to cloud services and receive Azure RBAC permissions. Azure also supports workload identities for applications and services, including integrations involving third-party cloud services where a supported Microsoft Entra security principal is used. The hybrid identity overview explains how organizations connect on-premises identity infrastructure with cloud identity services.

Finally, Azure includes core services for both authentication and authorization. Microsoft Entra ID authenticates identities through sign-in and identity protection capabilities, while Azure RBAC authorizes authenticated principals to perform specific actions at defined scopes, such as a subscription, resource group, or resource. This combination provides the built-in identity and access-management foundation used to secure Azure resources. See What is Microsoft Entra ID? for the identity-service overview.

QUESTION NO: 84

This question requires that you evaluate the underlined text to determine if it is correct. Azure Databricks is an Apache Spark-based analytics service.

Instructions: Review the underlined text. If it makes the statement correct, select "œNo change is needed." If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed.

B. Azure Data Factory.

C. Azure DevOps.

D. Azure HDInsight .

ANSWER: A

Explanation:

No change is needed. Azure Databricks is a fully managed analytics platform built on Apache Spark. It provides an Azure-native environment for processing and analyzing large volumes of data with distributed Spark compute, and supports common data and AI tasks such as data engineering, data science, machine learning, and analytics. Teams can use collaborative notebooks and workflows with familiar languages including Python, SQL, Scala, and R. Azure Databricks also integrates with Azure identity, storage, networking, and governance capabilities, allowing organizations to build data pipelines and analyze data while Azure manages core platform infrastructure. Because Apache Spark is the foundational processing engine for Azure Databricks and analytics is a primary intended use, the statement accurately describes the service. Microsoft describes Azure Databricks as a fast, easy, and collaborative Apache Spark-based analytics platform that is optimized for Microsoft Azure. Therefore, the existing wording is correct and does not need replacement. See What is Azure Databricks? and Azure Databricks product page.

QUESTION NO: 85

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:

Azure Kubernetes Service (AKS) and **Azure Container Instances** are Azure services specifically designed to run and manage containerized workloads without requiring an organization to operate the underlying container-host infrastructure. AKS is Azure's managed Kubernetes offering. It supports Kubernetes orchestration for container applications that require capabilities such as cluster-based deployment, scaling, load balancing, service discovery, rolling updates, and lifecycle management. Azure manages the Kubernetes control plane, reducing the administrative effort required to operate a Kubernetes environment.

Azure Container Instances provides a serverless container-execution service. It enables a team to deploy a container image directly to Azure, specify the needed CPU and memory resources, and run the workload without provisioning virtual machines or administering a Kubernetes cluster. This model is useful for isolated container workloads, short-lived tasks, automation, and situations where full orchestration is unnecessary. Together, these services cover both orchestrated Kubernetes workloads and straightforward on-demand container execution. For details, see Azure Kubernetes Service documentation and the Azure Container Instances overview.

QUESTION NO: 86

Your company plans to migrate to Azure. The company has several departments. All the Azure resources used by each department will be managed by a department administrator.

You need to recommend an Azure deployment that provides the ability to segment Azure for the departments.

What are two possible techniques to segment Azure for the departments? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. multiple subscriptions

B. multiple Azure Active Directory (Azure AD) directories

C. multiple regions

D. multiple resource groups

ANSWER: A D

Explanation:

multiple subscriptions and **multiple resource groups** are both valid ways to segment Azure resources for departmental administration. An Azure subscription is a management, billing, and governance boundary. Assigning a subscription to each department can provide independent cost tracking and budgets, allow department-specific Azure Policy and resource-provider settings, and enable administrators to manage only the subscription assigned to their department. Azure role-based access control can be used to grant each department administrator the appropriate permissions at that subscription scope.

A resource group is a logical container for Azure resources. Resources that share a common lifecycle, such as departmental workloads, can be placed in a department-specific resource group. Role assignments can be applied at the resource-group scope, enabling a department administrator to manage resources in that group without receiving permissions elsewhere in the subscription. Resource groups also support applying tags, policies, and resource locks consistently to the departmental resources they contain. The choice between these approaches depends on the required degree of isolation: subscriptions provide a broader governance and billing boundary, while resource groups provide efficient logical segmentation within a subscription.

See Azure resource organization guidance and Manage Azure resource groups.

QUESTION NO: 87

Which type of failure can Azure Availability Zones protect against to help maintain access to Azure services?

A. an Azure region failure

B. an Azure data center failure

C. a physical server failure

D. a storage failure

ANSWER: B

Explanation:

Azure Availability Zones protect workloads from an Azure data center failure within a single Azure region. An availability zone is a physically separate location made up of one or more datacenters, with independent power, cooling, and networking. This isolation is designed so that an incident affecting the infrastructure in one zone does not also affect resources deployed in another zone. By deploying supported services and application instances across multiple zones, an organization can maintain service access when a single zone becomes unavailable. For example, virtual machines can be distributed across zones and placed behind a load-balancing solution, allowing healthy instances in unaffected zones to continue serving requests. Zone-redundant Azure services can similarly replicate resources across zones automatically, depending on the service and configuration. Availability Zones are therefore a high-availability and fault-isolation capability for datacenter-level outages while retaining low-latency operation within the same region. Microsoft defines availability zones as unique physical locations within an Azure region and states that they are designed to protect applications and data from datacenter failures. Learn more in the Azure Availability Zones overview and Microsoft's Availability Zones reliability guidance.

QUESTION NO: 88

Your company plans to deploy an Artificial Intelligence (AI) solution on Microsoft Azure. Which Azure service should the company use to construct, evaluate, and deploy predictive analytics solutions?

A. Azure Logic Apps

B. Azure Machine Learning Designer

C. Azure Batch

D. Azure Cosmos DB

ANSWER: B

Explanation:

Azure Machine Learning Designer is the appropriate service because it provides a visual, drag-and-drop interface for creating machine learning pipelines in an Azure Machine Learning workspace. Predictive analytics requires a complete workflow: bringing in and preparing data, selecting and configuring training methods, training a model, assessing its performance with evaluation components and metrics, and making the trained model available for inference. Designer supports this workflow through reusable pipeline components that can be assembled without requiring users to write extensive code. The pipelines run as Azure Machine Learning jobs, allowing the organization to use managed compute and retain the associated assets and results in its workspace. After a model has been trained and evaluated, Azure Machine Learning provides model registration and managed endpoint deployment capabilities so applications can submit new data and receive predictions. This combination makes Azure Machine Learning Designer especially suitable for teams that want to construct, evaluate, and operationalize predictive models using a graphical authoring experience while still using the broader Azure Machine Learning platform for governance and deployment. Microsoft documents Designer as a drag-and-drop tool for building machine learning pipelines and explains how Azure Machine Learning endpoints support model deployment for real-time or batch inference. See What is Azure Machine Learning designer? and Azure Machine Learning endpoints.

QUESTION NO: 89

What is an example of vertical scaling in a cloud environment?

A. additional CPU to an existing Azure virtual machine.

B. adding an additional Azure App Service instance automatically.

C. adding an additional Azure virtual machine.

D. adding an additional Azure Virtual Desktop session host .

ANSWER: A

Explanation:

“additional CPU to an existing Azure virtual machine.” is an example of vertical scaling, also known as scaling up. Vertical scaling increases the capacity of a single existing resource rather than increasing the number of resources handling the workload. In Azure Virtual Machines, this is normally accomplished by resizing the VM to a larger supported size. A larger VM size can provide additional virtual CPUs, more memory, and, depending on the series and size selected, greater disk, network, or accelerated-computing capability. The application continues to run as one machine, but that machine has more resources available to process demand.

This approach is appropriate for workloads that require greater capacity on one server, such as an application that is difficult to partition across multiple machines or that depends on substantial memory or CPU in a single operating-system instance. Azure VM sizes define the available combinations of compute, memory, storage, and networking resources, so selecting a larger size is the Azure implementation of adding CPU capacity to an existing VM. Resizing may require a restart or deallocation in some scenarios, so it should be planned with availability requirements in mind. Microsoft architecture guidance distinguishes scaling up an individual resource from scaling out by adding resource instances. See Azure autoscaling guidance and Azure VM sizes documentation.

QUESTION NO: 90

Your Azure environment contains multiple Azure virtual machines.

You need to ensure that a virtual machine named VM1 is accessible from the Internet over HTTP.

What are two possible solutions? Each correct answer presents a complete solution.

A. Modify a DDoS protection plan.

B. Modify an Azure firewall.

C. Modify an Azure Traffic Manager profile.

D. Modify a network security group (NSG)

ANSWER: B D

Explanation:

Modifying an Azure firewall can provide Internet HTTP access when Azure Firewall is configured as the workload's public ingress point. Azure Firewall supports destination network address translation (DNAT), which can map inbound TCP traffic on port 80 received at the firewall's public IP address to VM1's private IP address and HTTP port. This makes it possible to publish the virtual machine through a centralized, stateful firewall while applying the required firewall policy and rule collection. Microsoft documents this inbound-publication pattern in its Azure Firewall DNAT tutorial: Filter inbound traffic with DNAT using Azure Firewall.

Modifying a network security group (NSG) is also part of a complete solution when VM1 has an appropriate public endpoint, such as a public IP address associated with its network interface. An inbound NSG security rule can allow TCP port 80 from Internet source addresses to the VM. NSGs can be associated with a subnet or network interface and evaluate inbound traffic against configured rules, enabling the network-layer permission required for HTTP connectivity. Microsoft's Network security groups overview describes NSG associations, security rules, and traffic evaluation.

QUESTION NO: 91

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:

Cloud computing enables the rapid provisioning of resources because cloud services make compute, storage, networking, databases, and other capabilities available on demand. Instead of acquiring, installing, and configuring physical hardware before a workload can run, an organization can deploy the required Azure resources through the portal, command-line tools, APIs, templates, or automated deployment pipelines. This supports faster experimentation, development, scaling, and response to business needs. Microsoft identifies agility and the ability to provision resources quickly as important cloud benefits.

Cloud computing also shifts capital expenditures (CAPEX) to operating expenditures (OPEX). Traditional datacenters typically require substantial upfront purchases of servers, storage, networking equipment, facilities, and supporting infrastructure. Cloud customers can instead consume services as needed and pay according to their usage or selected service plan. This consumption-based approach reduces the need to purchase maximum anticipated capacity in advance and helps align technology spending with operational demand. These economic and agility benefits are foundational concepts for Azure Fundamentals. See Microsoft Learn: Business justification and Microsoft Learn: Describe cloud computing.

QUESTION NO: 92

At which OSI layer does ExpressRoute operate?

A. Layer 5.

B. Layer 7.

C. Layer 3.

D. Layer 2 .

ANSWER: C

Explanation:

Layer 3 is correct because Azure ExpressRoute provides private, routed IP connectivity between an on-premises network and Microsoft cloud services. ExpressRoute uses Border Gateway Protocol (BGP) to exchange routing information between the customer or connectivity provider edge and Microsoft's enterprise edge. BGP advertises and learns IP prefixes, enabling each network to determine how IP traffic should be forwarded. IP addressing, packet routing, and path selection are core functions of the Network layer, which is OSI Layer 3.

With an ExpressRoute circuit, an organization extends its network to Azure over a private connection rather than sending the traffic across the public internet. The circuit supports BGP peerings that provide routes to Azure virtual networks and, where applicable, Microsoft public services. Although the physical connection can involve provider infrastructure, Ethernet handoffs, and VLAN tagging, ExpressRoute's service-level connectivity is routed IP networking. Therefore, its classification for Azure Fundamentals purposes is Layer 3. Microsoft's documentation describes ExpressRoute as a private connection to Microsoft cloud services and identifies BGP as the routing protocol used for its peerings. For details, see the Azure ExpressRoute introduction and ExpressRoute routing requirements.

QUESTION NO: 93

What are two benefits of cloud computing? Select each correct answer as it represents a full solution. Note: Each chosen option 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:

Cloud computing enables the rapid provisioning of resources because organizations can create and configure services such as virtual machines, storage, databases, and networking through the Azure portal, command-line tools, APIs, or infrastructure-as-code templates. Resources can be deployed in minutes when demand arises, rather than waiting for procurement, delivery, installation, and configuration of physical datacenter hardware. This agility helps teams respond quickly to business needs, experiment with solutions, and scale capacity as workloads change.

Cloud computing also shifts capital expenditures (CAPEX) to operating expenditures (OPEX). Instead of purchasing and maintaining servers, storage, networking equipment, and datacenter capacity upfront, an organization can consume cloud services and pay according to usage. This consumption-based model provides financial flexibility because spending can more closely align with actual demand. It also reduces the need to estimate and purchase peak capacity long before it is needed. These benefits—agility through on-demand provisioning and a consumption-based economic model—are foundational Azure cloud concepts. See Microsoft's cloud benefits guidance and its cloud economic model guidance.

QUESTION NO: 94

A team has an Azure Cosmos DB account. A solution needs to be in place to generate an alert from Azure Log Analytics when a query request charge exceeds 40 units more than 10 times during a 10-minute window.

Which of the following would you recommend? (Choose two)

A. Create a search query to identify when the requestCharge_s exceeds 10.

B. Configure a period of 10 and a frequency of 10.

C. Create a search query to identify when the requestCharge_s exceeds 40.

D. Create a search query to identify when the duration_s exceeds 10.

ANSWER: B C

Explanation:

A scheduled query rule in Azure Monitor can evaluate data collected in a Log Analytics workspace and trigger an alert when the query result meets a configured threshold. To detect the required Cosmos DB behavior, the query must filter records based on the request-charge field, `requestCharge_s`, with a value greater than 40. It should then count the matching records, so the alert condition can fire when the count is greater than 10. For example, the query logic can filter `requestCharge_s > 40` and use `summarize count()` over the selected time range. Azure Monitor alert rules support a lookback period and an evaluation frequency. Configuring both the period and frequency as 10 minutes evaluates the preceding 10-minute window every 10 minutes, matching the stated monitoring requirement. The alert condition should be configured against the aggregated count produced by the query. Azure Monitor scheduled query rules are designed for this type of log-based monitoring and can send notifications or invoke automated actions through an action group. For configuration concepts and supported alert-rule behavior, see Log alerts in Azure Monitor. Cosmos DB diagnostic logging and its integration with Azure Monitor are described in Monitor Azure Cosmos DB data by using diagnostic settings.

QUESTION NO: 95

You have an Azure web app.

You need to manage the settings of the web app from an iPhone.

What are two Azure management tools that you can use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point

A. Azure CLI.

B. the Azure portal.

C. Azure Cloud Shell.

D. Windows PowerShell.

E. Azure Storage Explorer.

ANSWER: B C

Explanation:

The Azure portal is a complete solution because it is Azure's web-based management console and can be opened in a browser on an iPhone. After signing in, an administrator can navigate to the web app's App Service resource and manage its configuration, including application settings, connection strings, deployment settings, networking-related settings, and scaling configuration where supported by the portal experience. The portal provides a centrally authenticated interface for administering Azure resources without requiring a desktop operating system or a locally installed management tool.

Azure Cloud Shell is also a complete solution because it is available through a mobile browser, either within the Azure portal or directly from the Cloud Shell website. Cloud Shell provides authenticated Azure CLI and Azure PowerShell environments hosted by Azure, allowing administrators to run commands that inspect and update App Service and web app configuration. For example, an administrator can use the hosted command environment to manage app settings without installing command-line software on the iPhone. Microsoft documents the Azure portal as a web-based console for managing Azure services and Cloud Shell as a browser-accessible shell for Azure administration. See Azure portal overview and Azure Cloud Shell overview.

QUESTION NO: 96

Fill in the blank

To save on the costs of an unused Azure virtual machine that runs Windows 10, you should ( )

A. Select Stop for the virtual machine from the Azure portal.

ANSWER: A

Explanation:

Select **Stop for the virtual machine from the Azure portal.** This Azure management-plane action stops and deallocates the virtual machine, placing it in the *Stopped (deallocated)* state. Deallocation releases the physical compute resources that Azure assigned to the VM, so billing for the VM's compute usage stops. This is the appropriate choice when the Windows 10 VM is not needed temporarily but may need to be started again later.

The VM's operating system disk, data disks, network configuration, and other saved settings are retained, allowing the machine to be restarted without rebuilding it. Although compute charges cease after deallocation, charges can still apply to resources that remain allocated independently of the VM, such as managed disks and certain public IP address configurations. Therefore, stopping the VM through the Azure portal provides the intended compute-cost savings while preserving the workload for future use. Microsoft documents the relationship between VM power states and billing, including that the *Stopped (deallocated)* state does not incur VM compute charges, at Azure VM states and billing. Operational guidance is available at Start and stop Azure Windows VMs.

QUESTION NO: 97

Please note: This question is a part of a multi-question set that addresses the same scenario. Each question within this set may contain unique solutions to meet the outlined objectives. Some questions in the set may have more than one correct solution, while others may not have a correct solution at all.

Once you answer a question in this section, you will not have the opportunity to revisit it. As such, these questions will not appear in the review screen.

Your company possesses an Azure subscription that encompasses the following unused resources:

- 20 user accounts in Azure Active Directory (Azure AD)
- Five groups in Azure AD
- 10 public IP addresses
- 10 network interfaces

Your task is to reduce the Azure costs for your company.

Proposed Solution: You decide to remove the unused groups.

Does this solution fulfill the objective of cost reduction?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because Microsoft Entra ID groups, formerly called Azure Active Directory groups, are directory objects used to organize users, manage access assignments, and apply licenses or policies. Azure does not impose a consumption-based charge merely for creating or retaining groups. Therefore, deleting five unused groups does not reduce the Azure subscription's resource-consumption costs. Although removing obsolete groups can be a sensible identity-governance measure that improves directory hygiene and makes access management easier, it is not a direct cost-optimization action.

To reduce Azure costs, the company should focus on resources that can incur Azure charges while provisioned or allocated. Public IP addresses are particularly relevant because their pricing depends on factors such as SKU and whether an address is associated with a resource; allocated Standard public IP addresses can be billable. Cost optimization should be based on usage and billing analysis, including reviewing public IP addresses and any dependent resources before deletion. Azure Cost Management and Azure Advisor can help identify cost data and optimization recommendations. Microsoft documents Entra licensing by user and feature tier, rather than as a per-group charge. See Microsoft Entra licensing and Azure public IP address pricing and lifecycle.

QUESTION NO: 98

A support engineer intends to perform multiple Azure management tasks utilizing the Azure CLI.

After installing the CLI on a computer, you are required to advise the support engineer on which tools can be used to execute the CLI commands.

Which two tools should you recommend to the support engineer? Each correct answer represents a complete solution.

NOTE: Each correct choice 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:

Command Prompt and **Windows PowerShell** can both execute Azure CLI commands on a Windows computer after the Azure CLI is installed. The installation provides the `az` command-line executable and adds it to the system path, enabling an engineer to open either supported shell and run commands such as `az login`, `az account show`, and resource-management commands. Command Prompt is the standard Windows command-line environment, while Windows PowerShell supports interactive command execution as well as scripting and automation. Azure CLI command syntax remains consistent between these shells because both invoke the same `az` executable. Microsoft's Windows installation instructions explicitly state that Azure CLI commands can be run from Command Prompt or PowerShell after installation. This makes each tool a complete solution for authenticating to Azure, managing subscriptions and resource groups, and administering Azure resources through Azure CLI. See Install Azure CLI on Windows and the Azure CLI documentation.

QUESTION NO: 99

You have an Azure virtual network that contains several virtual machines.

You need to enable the virtual machines to access an Azure Storage account by using a private IP address. The Storage account must not be accessible from the public Internet.

Which two resources should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a private endpoint.

B. a private DNS zone.

C. a public IP address.

D. an Azure Traffic Manager profile.

E. a virtual network gateway.

ANSWER: A B

Explanation:

An Azure private endpoint provides a private IP address from the virtual network for an Azure PaaS resource, including Azure Storage. Traffic from the virtual machines to the Storage account can then remain on the Microsoft network rather than traversing the public Internet.

A private DNS zone supports name resolution of the Storage account's public service name to the private endpoint IP address. The private DNS zone is linked to the virtual network so that clients can resolve the service name privately. The Storage account public network access can be disabled to prevent public connectivity. See Azure Private Endpoint overview and Azure Private Endpoint DNS configuration.

QUESTION NO: 100

Your company plans to migrate all its network resources to Azure. You need to start the planning process by exploring Azure. What should you create first?

A. a virtual network.

B. a management group.

C. a subscription.

D. a resource group .

ANSWER: C

Explanation:

a subscription. is the correct starting point because an Azure subscription is the logical container that enables an organization to provision and use Azure resources. It establishes the essential boundaries for billing, access control, and resource governance. The company needs this context before it can explore services in the Azure portal, deploy proof-of-concept workloads, estimate and monitor consumption, or assign permissions to administrators and contributors.

Within Azure’s resource hierarchy, subscriptions provide the scope in which resources are created and managed. A subscription can contain resource groups, and those resource groups can contain networking services and other Azure resources. It is also a key scope for Azure role-based access control, Azure Policy assignments, budgets, and cost analysis. This makes it suitable for an initial migration evaluation as well as for ongoing operational management after the migration proceeds.

Microsoft recommends organizing Azure resources and subscriptions according to business, governance, and operational needs. An organization can begin with an appropriate subscription for evaluation and then add or structure subscriptions as environments, workloads, departments, or billing requirements evolve. See Organize Azure resources and subscriptions and Azure Resource Manager management scope.

QUESTION NO: 101

Which two capabilities are provided by Azure Key Vault? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. storing secrets.

B. managing cryptographic keys.

C. hosting virtual machines.

D. distributing network traffic.

E. creating virtual networks.

ANSWER: A B

Explanation:

Storing secrets is a core Azure Key Vault capability. Secrets can include passwords, connection strings, API keys, and other sensitive values. Applications can retrieve secrets at runtime rather than storing them in source code or configuration files.

Managing cryptographic keys is also provided by Azure Key Vault. The service supports storing and controlling keys used for cryptographic operations such as encryption, decryption, signing, and verification. Key Vault can also store and manage certificates.

Azure Key Vault does not host virtual machines or provide load balancing. Access to vault contents is controlled through Microsoft Entra ID authentication and authorization mechanisms, including Azure RBAC. For more information, see Azure Key Vault overview.

QUESTION NO: 102

In which Azure support plans can you open a new support request?

A. Premier and Professional Direct only.

B. Premier, Professional Direct, and Standard only.

C. Premier, Professional Direct, Standard, and Developer only.

D. Premier, Professional Direct, Standard, Developer, and Basic .

ANSWER: D

Explanation:

Premier, Professional Direct, Standard, Developer, and Basic is correct because Azure customers can create support requests under each of these support-plan entitlements, although the scope of available support differs by plan. Basic support is included with an Azure account and allows customers to submit requests for billing and subscription-management issues. Thus, Basic does not prevent use of the Azure portal's support-request workflow; instead, it limits the request categories and support coverage available under that plan.

The paid support plans add technical-support benefits. Developer is intended for non-production workloads, while Standard and Professional Direct provide broader technical support and defined response-time targets. Premier support provides enterprise-level support arrangements. These plans determine the types of technical incidents that can be submitted, applicable severity levels, response targets, and any additional support services, rather than determining whether a customer can open a request at all.

Microsoft lists Basic among Azure support-plan offerings and specifies its billing and subscription support coverage. The Azure portal documentation also describes the process for creating a support request and selecting an issue type. See Azure support plans and Create an Azure support request.

QUESTION NO: 103

This question requires that you evaluate the underlined text to determine if it is correct.

From Azure Monitor, you can view which user turned off a specific virtual machine during the last 14 days.

Instructions: Review the underlined text. If it makes the statement correct, select "œNo change is needed" . If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed.

B. Azure Event Hubs.

C. Azure Activity Log.

D. Azure Service Health .

ANSWER: C

Explanation:

Azure Activity Log is the correct service because it records subscription-level control-plane events for Azure resources. Turning off, stopping, or deallocating a virtual machine is a management operation, and its associated Activity Log record can be located by filtering for the affected virtual machine and the relevant operation. The event details provide an audit trail that includes the event time, operation status, target resource, and caller identity. The caller field identifies the user, service principal, or managed identity that initiated the management action, allowing an administrator to determine who turned off the virtual machine.

The required 14-day lookback period is available without additional configuration because Azure retains Activity Log events for 90 days by default. Azure Monitor provides access to this activity data, and organizations can additionally export it to a Log Analytics workspace, storage account, or event hub when they need longer retention, centralized analysis, or custom alerts. For this scenario, the built-in Azure Activity Log contains the necessary administrative audit information. See Azure Activity Log and Activity Log retention period.

QUESTION NO: 104

Which two Azure services can you use if you want to run a container in Microsoft Azure? Each correct answer presents a complete solution.

A. Azure Container Instances

B. Azure Virtual Desktop

C. Azure Functions

D. Azure virtual machines

E. Azure Kubernetes Service (AKS)

ANSWER: A E

Explanation:

Azure Container Instances is a fully managed Azure service designed to run individual containers or container groups without requiring you to provision or manage virtual machines or a container orchestrator. You specify the container image and the required CPU, memory, networking, and restart settings, and Azure runs the workload on demand. This makes it well suited to simple containerized applications, jobs, and burst workloads.

Azure Kubernetes Service (AKS) is Azure's managed Kubernetes offering. It provides a managed control plane for deploying, scaling, networking, and operating containerized applications that need Kubernetes orchestration. AKS is appropriate when workloads consist of multiple containers or services and require capabilities such as declarative deployments, service discovery, scaling, and rolling updates.

Both services therefore provide direct Azure-hosted solutions for running container images: Azure Container Instances focuses on serverless container execution, while AKS provides managed Kubernetes orchestration. For more information, see Azure Container Instances overview and Azure Kubernetes Service overview.

QUESTION NO: 105

What should you use to evaluate whether your company ' s Azure environment meets regulatory requirements?

A. Compliance Manager from the Security Trust Portal

B. the Advisor blade from the Azure policy

C. the Knowledge Center website

D. the Microsoft Defender for Cloud regulatory compliance dashboard in the Azure portal

ANSWER: D

Explanation:

Use the Microsoft Defender for Cloud regulatory compliance dashboard in the Azure portal to evaluate whether an Azure environment aligns with regulatory requirements. This dashboard continuously assesses Azure resources and organizes the resulting security assessments against compliance standards and regulatory frameworks that are enabled for the subscription. It provides an overall compliance posture, shows controls and assessments associated with each standard, and helps identify resources that require remediation. This makes it practical for a company to review its current state, prioritize compliance gaps, and track progress toward satisfying the controls applicable to its industry or regulatory obligations. The dashboard is part of Microsoft Defender for Cloud, the current name for the Azure service previously called Azure Security Center. Its compliance information is based on Defender for Cloud recommendations and assessments, so the organization can move from a high-level compliance score into the specific controls, affected resources, and recommended actions that support improvement. Microsoft notes that the dashboard is intended to help assess and manage compliance posture; an organization remains responsible for determining and demonstrating its own compliance. See Regulatory compliance dashboard in Microsoft Defender for Cloud and Microsoft Defender for Cloud overview.

QUESTION NO: 106

This question requires that you evaluate the underlined BOLD text to determine if it is correct.

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

Instructions: Review the underlined text. If it makes the statement correct, select â€œNo change is needed.â€ If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed.

B. Access control (IAM)

C. Budget alerts

D. Compliance

ANSWER: C

Explanation:

Budget alerts are the Azure Cost Management feature designed to notify stakeholders when a subscription's costs or usage reach defined budget thresholds. A budget is configured for a selected scope, such as a subscription, resource group, or management group, with an amount and a recurring time period. Alert conditions can be set at percentages of the budgeted amount, including an actual-cost threshold that sends notifications after incurred charges reach the configured level. Notification recipients can include specified email addresses, and an action group can also be used to trigger additional automated responses. This supports proactive cost governance by giving teams visibility during the billing period instead of requiring them to wait for an invoice or manually monitor accumulated charges. Azure Cost Management evaluates budget thresholds and sends the configured alert when the threshold is met. Microsoft notes that budget alerts are intended for monitoring and do not stop resource consumption or prevent charges; they provide notification so the organization can investigate and manage spending. For configuration details, see Create and manage Azure budgets and Use cost alerts to monitor usage and spending.

QUESTION NO: 107

Fill in the blanks

You can access Compliance Manager from the ( )

A. Microsoft365admincenter

B. Microsoft Purview compliance portal (Microsoft 365 compliance center)

ANSWER: B

Explanation:

Microsoft Purview compliance portal (Microsoft 365 compliance center) is correct because Compliance Manager is a Microsoft Purview solution and is accessed through the Microsoft Purview portal. After an authorized user signs in at purview.microsoft.com, they can select Compliance Manager from the portal's Solutions navigation area. Compliance Manager provides a compliance score, assessments, improvement actions, and supporting evidence-management capabilities that help an organization understand and improve its compliance posture against regulatory, industry, and internal requirements. Access to this information is controlled through Microsoft Entra ID and the appropriate Purview or Compliance Manager permissions, so users must be assigned a suitable role before they can use relevant features. The Microsoft Purview portal is the current name and central experience for these compliance capabilities. Older training material, administrative guidance, and interface references may use the previous name, Microsoft 365 compliance center; this is why the parenthetical wording accurately identifies the same administrative destination in a current-context question. Microsoft's Compliance Manager documentation identifies it as a Microsoft Purview solution and directs users to the Purview portal for access and management. See Microsoft Purview Compliance Manager documentation and Microsoft Purview portal documentation.

QUESTION NO: 108

You plan to collect and analyze event details for five Azure virtual machines. You need to run queries to compare the event details collected from all the virtual machines. Which two tools should you

use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Azure Service Bus.

B. Azure Monitor.

C. Azure Service Health.

D. Log Analytics.

E. Azure Advisor .

ANSWER: B D

Explanation:

Azure Monitor and Log Analytics provide the Azure-native solution for centralizing virtual machine event data and querying it across multiple machines. Azure Monitor is Azure's observability platform and can collect telemetry from Azure virtual machines. In practice, the Azure Monitor agent collects configured Windows Event Log or Linux Syslog data under a data collection rule and sends that data to a Log Analytics workspace. This creates a centralized repository for event records from all five virtual machines rather than requiring separate review of each machine.

Log Analytics is the query and analysis experience for data stored in a Log Analytics workspace. It uses Kusto Query Language (KQL), which can filter logs by time, event type, severity, or computer name, and can summarize or compare results across the virtual machines. For example, a query can group event records by computer and count matching events to identify trends or differences across the environment. Together, Azure Monitor supplies the monitoring and collection capabilities, while Log Analytics enables the required cross-machine analysis and querying of the collected event details. See the Azure Monitor overview and Log Analytics overview.

QUESTION NO: 109

Your company currently hosts multiple servers within an on-premises network environment. They aim to decrease their administrative workload in several key areas:

- Backing up application data
- Replacing failed server hardware
- Managing physical server security
- Updating server operating systems
- Managing permissions to shared documents

To achieve this goal, the company plans to migrate their servers to Azure virtual machines. You need to determine which administrative responsibilities will be eliminated following the planned migration. Which two responsibilities should be identified? Each correct selection is worth one point.

A. Replacing failed server hardware

B. Backing up application data

C. Managing physical server security

D. Updating server operating systems

E. Managing permissions to shared documents

ANSWER: A C

Explanation:

Azure virtual machines are an infrastructure-as-a-service offering, so the customer manages the guest operating system, applications, data, identities, and configuration within each virtual machine. Microsoft operates the underlying Azure cloud infrastructure. Consequently, **Replacing failed server hardware** is eliminated from the company's workload. Azure maintains the physical host servers and associated datacenter infrastructure, handling physical component failures and host maintenance without requiring the customer to procure, access, or replace hardware.

Managing physical server security is also eliminated. Microsoft is responsible for the physical security of Azure datacenters and the infrastructure they contain. This includes protections for facilities and physical computing equipment, such as controlled access, monitoring, and environmental safeguards. The company still needs to apply appropriate security controls to its virtual machines and workloads, but it no longer administers the security of the physical servers hosting them.

This division of responsibilities is central to the Azure shared responsibility model for IaaS. Migrating to Azure VMs transfers management of the physical datacenter, network, storage, and compute infrastructure to Microsoft while retaining customer responsibility for the VM operating system and the resources deployed in it. See Microsoft's shared responsibility in the cloud guidance and Azure physical security overview.

QUESTION NO: 110

You plan to store 20 TB of data in Azure. The data will be accessed infrequently and visualized by using Microsoft Power BI.

You need to recommend a storage solution for the data.

Which two solutions should you recommend? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Azure SQL Database

B. Azure Cosmos DB

C. Azure SQL Data Warehouse

D. Azure Database for PostgreSQL

E. Azure Data Lake

ANSWER: C E

Explanation:

Azure SQL Data Warehouse and **Azure Data Lake** are appropriate solutions for a 20-TB analytics workload that will be queried infrequently and consumed through Microsoft Power BI. Azure SQL Data Warehouse, now provided as dedicated SQL pool capabilities in Azure Synapse Analytics, is a distributed, massively parallel processing data warehouse service. It is designed to store and analyze very large volumes of relational data, and its compute resources can be scaled to match reporting demand. This model is well suited to occasional Power BI reporting because capacity can be managed around the periods when analytical queries are needed. Microsoft documents the dedicated SQL pool architecture and its use of distributed query processing for large-scale warehousing workloads at Microsoft Learn.

Azure Data Lake, implemented with Azure Data Lake Storage Gen2, is also a complete fit for storing data at this scale. It provides highly scalable, low-cost storage for large analytics datasets and supports hierarchical namespaces and file-system semantics that are useful for data engineering and analytics. Power BI can connect to Azure Data Lake Storage Gen2 data, enabling reports and visualizations over data held in the lake. Microsoft provides connection guidance at Microsoft Learn. Its storage-oriented design is especially suitable when data is retained for analytics but accessed only occasionally.

QUESTION NO: 111

Which Azure service can you use as a security information and event management (SIEM) solution?

A. Azure Analysis Services.

B. Azure Cognitive Services.

C. Azure Sentinel.

D. Azure information Protection .

ANSWER: C

Explanation:

Azure Sentinel is the former name of Microsoft Sentinel, Microsoft's cloud-native security information and event management solution. Microsoft Sentinel centralizes security data from Azure resources, Microsoft security products, on-premises systems, and third-party services. It uses built-in data connectors and supported ingestion methods to collect logs, alerts, and other security telemetry into a workspace for analysis.

As a SIEM, Microsoft Sentinel helps security operations teams detect, investigate, and respond to threats. Its analytics capabilities correlate events and identify suspicious activity, creating incidents that analysts can investigate through incident-management and investigation experiences. It also integrates threat intelligence to enrich detections and supports automation through automation rules and playbooks. These playbooks can trigger repeatable response actions, such as notifying teams or initiating remediation workflows.

The Azure Sentinel name is commonly found in older Azure training material and exam questions, while current Microsoft documentation uses the Microsoft Sentinel product name. The underlying service is the same SIEM offering, so Azure Sentinel correctly identifies the Azure service that provides SIEM capabilities. See the Microsoft Sentinel overview and What is Microsoft Sentinel?.

QUESTION NO: 112

You have an Azure subscription. Where can you find information about the personal data collected by Microsoft, how Microsoft uses this data, and the purposes for which it is used?

A. the Data Protection Addendum

B. the Microsoft Online Services Terms

C. the Microsoft Privacy Statement

D. Azure Security Center

ANSWER: C

Explanation:

The Microsoft Privacy Statement is the central public notice that describes Microsoft's handling of personal data. It identifies categories of personal data that Microsoft collects through its products, services, websites, and interactions, and explains the purposes for which that information is processed. These purposes include delivering and operating services, improving and developing products, personalizing experiences, communicating with customers, maintaining security, and meeting legal obligations. The statement also describes the privacy choices, controls, and rights that may be available to individuals, such as managing communications preferences and accessing privacy-related resources. This makes it the appropriate source when the question is specifically about what personal data Microsoft collects, how Microsoft uses it, and why it is used. Microsoft publishes and maintains this notice so that individuals can review its privacy practices in one authoritative location. The full notice is available in the Microsoft Privacy Statement. For related Azure guidance on Microsoft's commitments and practices for safeguarding customer data, see Protect customer data in Azure.

QUESTION NO: 113

You have a virtual machine named VM1 that runs Windows Server 2016. VM1 is in the East US Azure region.

Which Azure service should you use from the Azure portal to view service failure notifications that can affect the availability of VM1?

A. Azure Service Fabric

B. Azure Monitor

C. Azure virtual machines

D. Azure Advisor

E. Azure Service Health

ANSWER: E

Explanation:

Azure Service Health is the correct service because it provides a personalized view of Azure platform events that may affect resources in a specific subscription and region, such as VM1 in East US. From the Azure portal, Service Health displays active service issues, planned maintenance, health advisories, and security advisories relevant to the services and regions an organization uses. This lets an administrator determine whether an Azure-side incident or scheduled platform event could affect virtual-machine availability. The service also provides event details, impact information, updates from Microsoft, and recommended actions when they are available. Administrators can configure Service Health alerts for selected event types, services, regions, and subscriptions, enabling proactive notification rather than requiring someone to manually check the portal. This subscription-aware, region-aware view is particularly useful for identifying platform conditions that could affect a deployed workload. Microsoft documents Service Health as the personalized Azure event experience and explains how it communicates incidents and maintenance affecting subscribed services. See Azure Service Health overview and Create Service Health alerts in the Azure portal.

QUESTION NO: 114

Note: The question is part of a series of questions that share the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated requirements.

Your company's infrastructure comprises multiple business units, each requiring a substantial number of Azure resources for their daily operations.

The resources needed by each business unit are uniform.

You need to establish a strategy to automate the creation of Azure resources.

Proposed Solution: You suggest including management groups in the strategy.

Does the solution meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because Azure management groups do not automate the deployment or creation of Azure resources. Management groups provide a governance scope above subscriptions, allowing an organization to arrange subscriptions into a hierarchy. They are valuable for consistently applying Azure Policy, role-based access control, and other governance settings across multiple subscriptions that may align with business units or organizational structures.

The requirement is to automate creation of a uniform set of resources for each business unit. This calls for an infrastructure-as-code approach that declares the required resources and configuration in a reusable deployment definition. Azure Resource Manager templates or Bicep can be parameterized and deployed repeatedly, enabling each business unit to receive the same standardized resource set while allowing values such as names, regions, or sizing to differ where necessary. Management groups can support the broader design by applying governance to the subscriptions receiving those deployments, but including management groups alone does not establish resource-creation automation.

Microsoft documents management groups as an organizational and governance mechanism for subscriptions in the Azure management groups overview. For declarative, repeatable Azure resource provisioning, see What is Bicep?.

QUESTION NO: 115

This question requires that you evaluate the underlined text to determine if it is correct.

You have several virtual machines in an Azure subscription. You create a new subscription. The virtual machines cannot be moved to the new subscription .

Instructions: Review the underlined text. If it makes the statement correct, select “No change is needed”. If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed

B. The virtual machines can be moved to the new subscription

C. The virtual machines can be moved to the new subscription only if they are all in the same resource group

D. The virtual machines can be moved to the new subscription only if they run Windows Server 2016.

ANSWER: B

Explanation:

The virtual machines can be moved to the new subscription because Azure Resource Manager supports moving virtual machines and their associated resources between subscriptions, provided that the move meets Azure’s resource-move requirements. Typically, the source and destination subscriptions must belong to the same Microsoft Entra tenant. When moving a virtual machine, all dependent resources that are required for it to function—such as managed disks, network interfaces, and possibly public IP addresses or virtual networks—must be included in the move where applicable. Azure validates the selected resources before the move is started, helping identify dependencies or configuration constraints that must be addressed. A virtual machine does not need to be recreated simply because organizational ownership, billing, or subscription administration changes. The move operation transfers the supported resources to the destination subscription while retaining their configuration. Microsoft documents the supported resource-move scenarios and the prerequisites for moving Azure virtual machines in its move support for resources guidance and its move resources to a new resource group or subscription documentation.

QUESTION NO: 116

You need to select Azure storage services for the following requirements:

• Store unstructured data such as images and video files.

• Store messages that can be retrieved and processed asynchronously by an application.

Which two Azure Storage services should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure Blob Storage.

B. Azure Queue Storage.

C. Azure Files.

D. Azure Table Storage.

E. Azure managed disks.

ANSWER: A B

Explanation:

Azure Blob Storage is the appropriate Azure Storage service for unstructured object data, including images, video, documents, backups, and log files. Blobs are stored in containers within a storage account and are designed to scale for very large quantities of object data. Applications can access blob data through standard HTTP or HTTPS interfaces, making it suitable for serving and managing media files.

Azure Queue Storage is the appropriate service for asynchronous messaging between application components. A sending component places a message in a queue, where it remains available until a receiving component retrieves and processes it. This pattern decouples the sender from the processor, supports workload buffering, and enables applications to process queued work independently and asynchronously. Queue Storage is therefore a direct fit when messages must be held for later processing by an application.

These services address the two distinct requirements: scalable object storage for unstructured media and durable queue-based messaging for asynchronous processing. For details, see Introduction to Azure Blob Storage and Introduction to Azure Queue Storage.

QUESTION NO: 117

You have an Azure application that uses the services shown in the following table.

Service	Service Level Agreement (SLA)
Azure virtual machines	99.9 %
Azure SQL Database	99.99%

How should you calculate the composite SLA for the application?

A. Max(0.999, 0.9999) = 0.9999 = 99.99%.

B. 0.999 / 0.9999 = 0.9991 = 99.91%.

C. 0.999 * 0.9999 = 0.9989001 = 99.89001%.

D. Min(0.999, 0.9999) = 0.999 = 99.9%.

ANSWER: C

Explanation:

The composite SLA is calculated as **0.999 * 0.9999 = 0.9989001 = 99.89001%**. A composite SLA estimates the availability commitment for an application that depends on multiple Azure services. When every required service must be available for the application to successfully process a request, all dependency availability probabilities must occur together. The probability of that combined event is found by multiplying the individual SLA values expressed as decimals.

In this case, the service commitments are 99.9% and 99.99%, which convert to 0.999 and 0.9999. Their product is 0.9989001, or 99.89001%. This calculation illustrates an important reliability-design principle: adding required, serial dependencies can reduce an application's overall availability target even when each individual service has a high SLA. The resulting value helps teams assess whether the application meets its business availability objective and whether resiliency measures, such as redundancy or an alternative architecture, may be needed.

Microsoft documents composite SLA calculations for dependent components in its Reliability metrics and targets guidance. Applicable service commitments are published in the Microsoft Azure Service Level Agreements.

QUESTION NO: 118

You are considering migrating several servers from an on-premises network to Microsoft Azure. What is a key advantage of leveraging a public cloud service for these servers compared to maintaining them within your own on-premises infrastructure?

A. The public cloud is owned by the public, NOT a private corporation

B. The public cloud is a crowd-sourcing solution that provides corporations with the ability to enhance the cloud

C. All public cloud resources can be freely accessed by every member of the public

D. The public cloud is a shared entity whereby multiple corporations each use a portion of the resources in the cloud

ANSWER: D

Explanation:

The public cloud is a shared entity whereby multiple corporations each use a portion of the resources in the cloud. This accurately describes the multitenant nature of Azure public cloud services. Microsoft operates large-scale datacenters and provides shared underlying physical infrastructure, while each customer consumes logically isolated compute, storage, and networking resources through its own Azure subscription. Logical isolation means an organization can run and manage its workloads independently while benefiting from the provider’s shared platform.

For a server migration, this model offers a significant operational advantage over owning equivalent on-premises infrastructure. The organization does not need to purchase, house, power, cool, or replace the physical servers and datacenter equipment. Instead, it can provision virtual machines and related services on demand, adjust capacity as workload requirements change, and use consumption-based pricing for many services. Microsoft manages the physical facilities and hardware lifecycle, while the customer remains responsible for configuring its cloud resources, protecting its data, and managing workload access according to the shared responsibility model. This combination of shared provider infrastructure, logical tenant isolation, elasticity, and reduced capital expenditure is a core benefit of public cloud computing. See Microsoft Learn: What is cloud computing? and Microsoft Learn: Shared responsibility in the cloud.

QUESTION NO: 119

You need to store data in an Azure storage account. The data must remain available if a single Azure datacenter fails.

Which two redundancy options meet the requirement? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. locally redundant storage (LRS).

B. zone-redundant storage (ZRS).

C. geo-redundant storage (GRS).

D. geo-zone-redundant storage (GZRS).

ANSWER: B D

Explanation:

Zone-redundant storage (ZRS) replicates data synchronously across multiple availability zones within a single Azure region. Because the zones are physically separate locations with independent power, cooling, and networking, ZRS can maintain data availability when one datacenter location fails.

Geo-zone-redundant storage (GZRS) also meets the requirement. GZRS uses ZRS replication in the primary region and asynchronously replicates data to a secondary geographic region. It provides protection from a datacenter failure in the primary region and adds resilience for a regional outage.

Locally redundant storage keeps copies within a single datacenter and therefore does not provide datacenter-level fault isolation. Geo-redundant storage replicates to another region, but its primary-region copies are locally redundant rather than zone redundant. See Azure Storage redundancy.

QUESTION NO: 120

You need to identify the types of Azure platform events that are displayed in Azure Service Health.

Which two event types should you identify? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Application exceptions.

B. Service issues.

C. Planned maintenance.

D. Operating system event logs.

ANSWER: B C

Explanation:

Service issues. and **Planned maintenance.** are event types shown in Azure Service Health. Service Health provides personalized information about Azure events and maintenance that affect the services and regions used by an organization's Azure subscriptions. Service issues communicate active Azure outages, degradations, or other service-impacting incidents, helping administrators assess potential impact on their deployed resources and track restoration updates. Planned maintenance provides advance notification of maintenance activities that could affect availability or performance, enabling teams to prepare operationally, schedule work around the maintenance window, or apply recommended actions where appropriate.

These events can be viewed in the Azure portal and can also be routed through Service Health alerts to action groups, allowing timely notification through channels such as email, SMS, push notifications, webhooks, or IT service-management integrations. This supports proactive monitoring of Azure platform conditions relevant to a specific environment. Microsoft's Azure Service Health overview describes Service Health event categories and their subscription-specific scope. For guidance on configuring notifications, see Create Service Health alerts.

QUESTION NO: 121

Which cloud computing model includes on-premises and cloud-based resources?

A. private.

B. public.

C. hybrid.

ANSWER: C

Explanation:

hybrid. is correct because hybrid cloud computing combines an organization's on-premises infrastructure with cloud-based resources, enabling applications, data, and services to operate across both environments. Rather than requiring an immediate, complete migration to the cloud, this model allows an organization to retain selected workloads in its datacenter while using Azure services where cloud scale, flexibility, or managed capabilities are beneficial. The environments are connected and used together to support business operations, such as extending local capacity into Azure, maintaining data locally to meet regulatory or latency needs, or using cloud services for backup and disaster recovery. In Azure, hybrid capabilities can include services and technologies such as Azure Arc, Azure Stack HCI, and Azure ExpressRoute, which help organizations manage, run, or connect resources spanning on-premises locations and Azure. Microsoft's Azure Fundamentals cloud-computing guidance defines hybrid cloud as a model that uses both public-cloud and private-cloud or on-premises resources. Learn more in Microsoft Cloud Adoption Framework: Hybrid and multicloud and Azure Architecture Center: Hybrid compute.

QUESTION NO: 122

Your Azure environment comprises multiple virtual machines. You are tasked with ensuring that the virtual machine named VM1 is accessible via the Internet using the HTTP protocol. What are two potential solutions? Each correct answer represents a full solution.

NOTE: Each correct selection is worth one point.

A. Modify an Azure Traffic Manager profile

B. Modify a network security group (NSG)

C. Modify a DDoS protection plan

D. Modify an Azure firewall

ANSWER: B D

Explanation:

Modify a network security group (NSG) is a valid solution because an NSG can be applied to VM1's network interface or to the subnet containing VM1. An inbound NSG rule can allow TCP traffic on destination port 80, which is the default port for HTTP. Provided that VM1 is reachable through a public IP address or another configured public ingress path, the rule permits Internet clients to connect to the web service. NSGs filter network traffic at Layers 3 and 4 and are stateful, meaning response traffic for an established allowed connection is automatically allowed. Microsoft documents NSG rule processing and associations in the Network security groups overview.

Modify an Azure firewall is also a valid solution when Azure Firewall is used as the Internet-facing entry point. A DNAT rule can receive TCP traffic addressed to port 80 on the firewall's public IP address and translate it to VM1's private IP address and port 80. This makes the HTTP service reachable without assigning a public IP directly to VM1, while allowing centrally managed inspection, policy, and logging. The routing and relevant security rules must permit the translated traffic to reach the virtual machine. Microsoft provides this publishing pattern in its Azure Firewall DNAT tutorial.

QUESTION NO: 123 - (HOTSPOT)

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statements	Yes	No
A Platform as a Service (PaaS) solution provides full control of operating systems that host applications.	○	○
A Platform as a Service (PaaS) solution provides additional memory to apps by changing pricing tiers.	○	○
A Platform as a Service (PaaS) solution can automatically scale the number of instances.	○	○

ANSWER:

Statements	Yes	No
A Platform as a Service (PaaS) solution provides full control of operating systems that host applications.	○	○
A Platform as a Service (PaaS) solution provides additional memory to apps by changing pricing tiers.	○	○
A Platform as a Service (PaaS) solution can automatically scale the number of instances.	○	○

Explanation:

Platform as a Service (PaaS) lets an organization deploy and operate its applications without administering the underlying operating system and core platform infrastructure. Azure manages the datacenter facilities, physical hardware, networking foundations, host infrastructure, virtualization, and operating system components used to provide the managed service. The customer retains responsibility for its application code, application settings, identities and access as applicable, and data. Therefore, a PaaS offering does not give the customer full control of the operating system that hosts the application. Microsoft describes this division of management responsibilities in its shared responsibility in the cloud guidance.

PaaS services can provide additional compute and memory capacity by moving an application to a higher service plan or pricing tier. This is referred to as scaling up, or scaling vertically. For example, Azure App Service plans provide different resource capabilities, and selecting a higher tier can allocate more capable infrastructure to the app while keeping the platform managed. Microsoft documents this process in Scale up an app in Azure App Service.

PaaS can also support automatic scale-out, where the platform increases or decreases the number of application instances based on demand, schedules, or monitored metrics. This enables an application to respond to changing workload levels while avoiding manual instance management. Azure Autoscale provides this type of capacity adjustment for supported resources, as described in the Azure Monitor Autoscale overview.

QUESTION NO: 124

What should a desktop application use to interact with Azure and manage resources?

A. Azure Resource Manager (ARM) templates.

B. APIs.

C. Azure Cloud Shell.

D. Azure Command-Line Interface (CLI) .

ANSWER: B

Explanation:

APIs are the appropriate mechanism for a desktop application to interact programmatically with Azure and manage Azure resources. Azure resource management is provided through the Azure Resource Manager control plane, which exposes REST APIs for operations such as creating, reading, updating, deleting, and configuring resources. A desktop application can authenticate users or workloads through Microsoft Entra ID, obtain an access token, and call the relevant management API endpoints. In practice, developers commonly use an Azure SDK for their application language, because SDKs provide

typed client libraries that simplify authentication, request construction, paging, retries, and responses while using the underlying Azure APIs.

This approach supports application-driven management workflows and integrates directly into the desktop application's code. It also aligns with Azure's standard management architecture: Azure portal, Azure PowerShell, Azure CLI, ARM/Bicep deployments, and custom applications ultimately use Azure Resource Manager APIs to manage resources. The required permissions are assigned through Azure role-based access control, allowing the application identity or signed-in user to perform only authorized operations. See the Azure Resource Manager overview and the Azure REST API reference.

QUESTION NO: 125 - (SIMULATION)

You need to view a list of planned maintenance events that can affect the availability of an Azure subscription.

Which blade should you use from the Azure portal? To answer, select the appropriate blade in the answer

area.



ANSWER: See the explanation for the answer

Explanation:

Select the **Service Health** blade (Azure Service Health).

In Azure Service Health, select **Planned maintenance** to view maintenance events that may affect resources in the subscription.

QUESTION NO: 126

You need to identify Azure services that enable you to run event-driven workloads without managing virtual machines.

Which two services should you identify? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Azure Functions.

B. Azure virtual machines.

C. Azure Logic Apps.

D. Azure Kubernetes Service (AKS).

ANSWER: A C

Explanation:

Azure Functions is an event-driven serverless compute service. It runs small units of code in response to triggers, including HTTP requests, timer schedules, messages placed on queues, changes to data, and Azure Event Grid events. With the serverless hosting model, Azure supplies and manages the compute infrastructure, while the platform can scale execution according to demand. This makes it suitable when an application needs custom code to process an event without provisioning or administering virtual machines.

Azure Logic Apps is a cloud workflow service that supports event-driven automation and integration. A workflow can begin when an event occurs, such as receiving a message, a scheduled time being reached, or a resource generating an event. Logic Apps provides built-in operations and connectors for Azure services, Microsoft services, SaaS applications, and enterprise systems. Azure runs and manages the workflow infrastructure, allowing teams to focus on defining workflow steps rather than maintaining servers or operating systems. These two services therefore provide complete serverless approaches for event-triggered workloads. For details, see Azure Functions overview and Azure Logic Apps overview.

QUESTION NO: 127

To manage containers effectively, which two Azure services can you use? Each correct answer represents 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:

Azure Kubernetes Service (AKS) and Azure Container Instances are Azure services specifically designed to run containerized workloads. Azure Kubernetes Service provides a managed Kubernetes environment, including a managed control plane, for deploying and operating container applications that require orchestration. Kubernetes capabilities enable organizations to manage application scaling, networking, updates, and the lifecycle of workloads composed of multiple containers. This makes AKS appropriate for more complex, distributed, and production-scale applications.

Azure Container Instances provides a serverless container platform that runs containers without requiring you to provision or manage virtual machines, a container host, or a Kubernetes cluster. Containers can be deployed rapidly in container groups,

where they can share a lifecycle, network configuration, storage volumes, and resources. This approach is well suited to isolated applications, automated tasks, and short-lived workloads where full container orchestration is unnecessary.

Both services are core Azure container offerings and provide managed solutions for deploying and operating containers. For additional details, see the Azure Kubernetes Service overview and the Azure Container Instances overview.

QUESTION NO: 128

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, **maintaining the hardware** is a cloud service provider responsibility. Microsoft operates the underlying physical infrastructure that delivers Azure services, including datacenter servers, physical storage hardware, physical networking equipment, and host infrastructure. The provider is responsible for operating, repairing, replacing, and securing this foundational platform so that customers can deploy and manage their virtualized resources.

Physical security of the datacenter infrastructure is also managed by the cloud service provider. This includes protections for Azure datacenter facilities and the physical assets within them, such as controlled access to buildings, monitoring, environmental safeguards, and protection of physical compute and network equipment. These provider-managed controls form the base layer of the cloud environment.

Under IaaS, the customer retains responsibility for the resources and workload configuration they deploy on that platform, while the provider manages the physical datacenter and physical infrastructure. Microsoft documents this division in its shared responsibility in the cloud guidance. Additional information about Azure facility protections is available in Azure physical security.

QUESTION NO: 129

An Azure administrator plans to execute a PowerShell script that is designed to create various Azure resources. You need to advise which computer configurations are suitable to run this script successfully. Which three computers are capable of running this script? Each correct answer presents a complete solution. **Note:** Each correct selection is worth one point.

A. a computer that runs macOS and has PowerShell Core 6.0 installed.

B. a computer that runs Windows 10 and has the Azure PowerShell module installed.

C. a computer that runs Linux and has the Azure PowerShell module installed.

D. a computer that runs Linux and has the Azure CLI tools installed.

E. a computer that runs Chrome OS and uses Azure Cloud Shell.

ANSWER: B C E

Explanation:

A computer that runs Windows 10 and has the Azure PowerShell module installed can run the script because Windows provides a PowerShell environment, while the Azure PowerShell module supplies the Azure-specific cmdlets used to authenticate, deploy, and manage Azure resources. Azure PowerShell is supported on Windows and is installed as the Az PowerShell module. See the Install Azure PowerShell guidance.

A computer that runs Linux and has the Azure PowerShell module installed is also a complete solution. PowerShell is cross-platform, and Azure PowerShell can be installed and used in supported Linux environments. Once PowerShell and the Azure module are available, the script can run its Azure cmdlets to create the required resources.

A computer that runs Chrome OS and uses Azure Cloud Shell is suitable because Azure Cloud Shell is a browser-based, Microsoft-hosted shell environment. Its PowerShell experience includes Azure PowerShell tooling, so the Chrome OS device only needs a supported browser and internet access. The script can be created or uploaded to Cloud Shell and then executed in the hosted PowerShell session. Microsoft documents both Bash and PowerShell experiences in Azure Cloud Shell overview.

QUESTION NO: 130

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, white 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.

Your company has an Azure subscription that contains the following unused resources:

* 20 user accounts in Azure Active Directory (Azure AD)

* Five groups in Azure AD

* 10 public [P addresses

* 10 network interfaces

* You need to reduce the Azure costs for the company.

Solution: You remove the unused user accounts.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because deleting unused Azure AD user accounts does not, by itself, reduce Azure consumption costs. Microsoft Entra ID (formerly Azure Active Directory) provides user and group directory objects as part of the tenant directory; creating or retaining ordinary user accounts is not billed as an Azure resource usage charge. Although an organization might assign paid Microsoft Entra ID Premium licenses to users, cost control in that situation is achieved by removing or reassigning the licenses, not simply by deleting the user objects. The scenario only states that the accounts are unused and does not identify paid licenses assigned to them.

To reduce the Azure bill, the company should instead focus on resources that can incur charges when allocated but unused. For example, public IP addresses can be billable depending on their SKU and allocation configuration. Azure networking charges and pricing vary according to the specific resource type, SKU, region, and whether an address is associated with a running resource. Removing an unused directory account therefore does not meet the stated cost-reduction goal. Microsoft's guidance on Microsoft Entra licensing is available in Microsoft Entra licensing fundamentals, and public IP pricing details are available on the Azure IP Addresses pricing page.

QUESTION NO: 131

You have an accounting application named App1 that uses a legacy database.

You plan to move App1 to the cloud.

Which service model should you use?

A. software as a service (SaaS)

B. infrastructure as a service (IaaS)

C. platform as a service (PaaS)

ANSWER: B

Explanation:

Infrastructure as a service (IaaS) is the appropriate model for moving an existing accounting application that depends on a legacy database to the cloud. IaaS provides virtualized compute, storage, and networking resources while leaving the customer responsible for the operating system, application, data, runtime, and configuration. This level of control is valuable when an application has dependencies on a database version, operating-system configuration, drivers, networking settings, or other components that must be retained for compatibility. App1 and its legacy database can be hosted on Azure virtual machines, allowing the organization to perform a lift-and-shift migration with minimal architectural change. The organization can select and administer the required VM operating system and install, configure, patch, back up, and secure the database according to its existing operational requirements. Azure remains responsible for the physical datacenter infrastructure, physical servers, and underlying virtualization layer. Microsoft describes IaaS as an on-demand infrastructure model that supplies compute, storage, and networking resources, while customers manage the software running on those resources. See Microsoft's IaaS definition and Azure Virtual Machines overview.

QUESTION NO: 132 - (DRAG DROP)

Match the Azure Services service to the correct descriptions.

Instructions: To answer, drag the appropriate 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 completed service mapping is correct because it pairs each security requirement with the Azure service designed for that purpose. Azure Sentinel, now called Microsoft Sentinel, is a cloud-native security information and event management service. It can collect, correlate, and analyze security data from Azure resources, including log data associated with virtual machines. This makes it suitable for analyzing security log files from Azure virtual machines, detecting suspicious activity, and supporting investigations. Microsoft describes these monitoring, analytics, detection, and response capabilities in the Microsoft Sentinel overview.

Azure Security Center is the former product name for the cloud security posture management capabilities now provided by Microsoft Defender for Cloud. A central capability is Secure Score, which gives an Azure subscription a measurable security posture score and provides recommendations for strengthening resource configurations. Displaying the secure score for an Azure subscription is therefore the intended use of Azure Security Center in the terminology used by this question. The current documentation explains Secure Score and its security recommendations in Microsoft Defender for Cloud Secure Score.

Azure Key Vault is the appropriate managed service for storing sensitive application values, including passwords, secrets, tokens, and connection strings. An Azure Function application can retrieve secrets from Key Vault at runtime by using managed identities and appropriate authorization, rather than placing passwords directly in source code or ordinary application settings. This improves secret protection and makes secret rotation and access management more secure. See the Azure Key Vault overview for its supported secret-management capabilities.

QUESTION NO: 133

You have several Azure virtual machines that do not have public IP addresses.

You need to provide administrators with secure Remote Desktop Protocol (RDP) access to the virtual machines from the Azure portal.

What should you use?

A. Azure Bastion.

B. Azure Firewall.

C. Azure DDoS Protection.

D. Azure Traffic Manager.

ANSWER: A

Explanation:

Azure Bastion provides secure, browser-based RDP and SSH access to Azure virtual machines directly through the Azure portal. It is designed for management connectivity when virtual machines do not have public IP addresses. Bastion is deployed into the same virtual network as the target virtual machines (or can connect through supported virtual-network connectivity configurations), and it reaches those machines by using their private IP addresses. Administrators authenticate to Azure and then initiate an RDP session from the portal without exposing the VM's RDP port to the public internet.

This design reduces the attack surface because inbound management ports, including TCP 3389 for RDP, do not need to be opened on public IP addresses or exposed through network security rules for internet-based administration. Azure Bastion is therefore the Azure service purpose-built for secure portal-based remote administration of private Azure VMs. Microsoft documents that Azure Bastion lets users connect to virtual machines by using RDP or SSH without exposing the VMs through public IP addresses. See What is Azure Bastion? and Connect to a VM using RDP through Azure Bastion.

QUESTION NO: 134

This question requires that you evaluate the underlined text to determine if it is correct.

The Microsoft Online Services Privacy Statement explains what data Microsoft processes, how Microsoft processes the data, and the purpose of processing the data.

Instructions: Review the underlined text. If it makes the statement correct, select â€œNo change is needed.â€ If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed.

B. Microsoft Online Services Terms

C. Microsoft Online Service Level Agreement

D. Online Subscription Agreement for Microsoft Azure

E. Microsoft Privacy Statement

ANSWER: E

Explanation:

Microsoft Privacy Statement is the correct text. This statement describes Microsoft's approach to personal data: the types of personal data Microsoft collects and processes, the purposes for which that data is used, the circumstances in which data may be shared, and the controls available to individuals. It applies to Microsoft consumer products, services, websites, apps, devices, and interactions that reference the statement. Therefore, it is the Microsoft privacy document that directly explains what personal data is processed, how and why Microsoft uses it, and the choices available to users regarding that processing. The statement also addresses important privacy topics such as account data, diagnostics, communications, cookies, advertising, data retention, and ways to access or manage personal data. These disclosures support transparency obligations and help customers understand Microsoft's data-handling practices. For the current authoritative statement, see the Microsoft Privacy Statement. Microsoft also provides an overview of its privacy principles and governance at Microsoft Trust Center Privacy.

QUESTION NO: 135

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:

The correct solution is **a virtual network for FinServer and another virtual network for all the other servers**. An Azure virtual network (VNet) is a logically isolated private network in Azure and acts as a fundamental network boundary for

resources such as virtual machines. Placing FinServer in a dedicated VNet gives the organization a separate address space and an independent network segment, which directly supports the stated compliance requirement for isolation from the other migrated servers.

Separate VNets do not have automatic connectivity with one another. This means that traffic between FinServer and the other servers is unavailable unless the organization intentionally configures a connection. Where a justified business requirement exists, connectivity can be explicitly established through services such as VNet peering, a VPN gateway, or ExpressRoute. The organization can then apply additional controls, including network security groups, user-defined routes, and Azure Firewall, to restrict and monitor permitted traffic. This approach provides clear segmentation and enables deliberate governance of any cross-network communication involving the sensitive financial server. Microsoft identifies VNets as the core building block for private Azure networks and documents their isolation and connectivity capabilities in the Azure Virtual Network overview and the virtual network peering overview.

QUESTION NO: 136

Which two features of services can be integrated with Azure Monitor? Each correct answer presents part of the solution.

NOTE Each correct answer is worth one point:

A. Azure Service Health.

B. Azure Advisor.

C. Azure status.

D. Application Insights.

E. Log Analytics .

ANSWER: D E

Explanation:

Application Insights is integrated with Azure Monitor to provide application performance monitoring. It collects telemetry from web applications and services, including requests, dependencies, exceptions, traces, availability-test results, and performance data. Azure Monitor can correlate this telemetry so that teams can assess application health, investigate failures, and identify performance issues across application components. Application Insights data is stored in a Log Analytics workspace for modern workspace-based resources, enabling unified analysis with other Azure Monitor data.

Log Analytics is the Azure Monitor experience for storing, querying, and analyzing log and telemetry data in a Log Analytics workspace. Data collected from Azure resources, virtual machines, applications, and supported external sources can be queried by using Kusto Query Language (KQL). These queries support investigation, dashboards, workbooks, alert rules, and automated responses. Together, Application Insights supplies detailed application telemetry, while Log Analytics provides the workspace and query capabilities used to analyze Azure Monitor logs across monitored resources. Microsoft describes both as key parts of the Azure Monitor platform. See the Azure Monitor overview and the Application Insights overview.

QUESTION NO: 137

This question is part of a series of questions that present the same scenario. Each question in the series provides a unique solution that might achieve the stated objectives. Some question sets might have more than one correct solution, while others may not have any correct solutions. Once you answer a question in this section, you will NOT have the option to return to it. Thus, 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 these virtual machines remain available in the event of a single data center failure.

Solution: You deploy the virtual machines across two or more availability zones.

Does this solution satisfy the requirement?

A. Yes

B. No

ANSWER: A

Explanation:

Yes. Deploying multiple virtual machines across two or more Availability Zones satisfies the requirement to maintain service availability during a single datacenter failure. An Availability Zone is a physically separate location within an Azure region. Each zone has independent power, cooling, and networking infrastructure, so a failure affecting one datacenter or zone is isolated from the virtual machines running in the other zones. This provides the required in-region resiliency against a datacenter-level outage.

To realize service availability, the workload must run redundant instances in separate zones and route requests to healthy instances. A zone-redundant Azure Load Balancer or an appropriate application-level traffic-routing mechanism can distribute traffic across those instances. If one zone becomes unavailable, the service instances in the remaining zone or zones can continue handling requests. Availability Zones are therefore the appropriate Azure capability when the objective is to protect virtual-machine workloads from the failure of one datacenter while remaining within the same Azure region.

Microsoft documents Availability Zones as unique physical locations within an Azure region that provide high availability through independent infrastructure. See the Availability Zones overview and availability options for Azure Virtual Machines.

QUESTION NO: 138

Which cloud computing model incorporates both on-premises and cloud-based resources?

A. private

B. public

C. hybrid

ANSWER: C

Explanation:

Hybrid is correct because a hybrid cloud model combines on-premises infrastructure with cloud-based resources and services in a connected operating environment. This model lets an organization place workloads, applications, and data in the location that best meets its operational needs while still using cloud capabilities. For example, an organization can retain systems with specific data-residency, latency, regulatory, or legacy-integration requirements in its datacenter, while using Azure for scalable computing, backup, disaster recovery, analytics, and modernization initiatives.

Hybrid cloud does not require all resources to move to the cloud at once. Instead, it supports an integrated approach in which local infrastructure and Azure services work together. Azure hybrid offerings, including Azure Arc and Azure Stack HCI, help organizations manage, govern, and run selected services across infrastructure outside Azure and resources hosted in Azure. This flexibility is especially useful for organizations that need a gradual cloud migration or must maintain some on-premises capabilities.

Microsoft defines hybrid scenarios as using a combination of cloud and on-premises resources, which directly matches the requirement in the question. See Microsoft Learn: Hybrid and multicloud scenarios and Microsoft Learn: Hybrid compute.

QUESTION NO: 139

Which two features are provided by Microsoft Entra ID? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. single sign-on (SSO).

B. multi-factor authentication (MFA).

C. physical server maintenance.

D. automatic replacement of network hardware.

E. guest operating system patching.

ANSWER: A B

Explanation:

Single sign-on (SSO) is provided by Microsoft Entra ID. After users authenticate, SSO enables them to access multiple authorized cloud applications without entering their credentials separately for each application. This improves the user experience while allowing administrators to centrally manage identities and application access.

Multi-factor authentication (MFA) is also supported by Microsoft Entra ID. MFA requires users to provide an additional verification factor, such as an authenticator app approval, in addition to a password. This reduces the risk that a compromised password alone can be used to access organizational resources. Microsoft Entra ID is a cloud-based identity and access management service that provides authentication, authorization, and identity protection capabilities. See What is Microsoft Entra ID? and How Microsoft Entra MFA works.

QUESTION NO: 140

You have a set of Azure policies. You need to group related policies together. What should you use?

A. exemptions.

B. metadata.

C. parameters.

D. initiatives.

ANSWER: D

Explanation:

Initiatives are the Azure Policy feature designed to group related policy definitions and manage them as a single unit. An initiative is also called a policy set definition. It supports a broader governance objective, such as meeting a regulatory compliance standard, enforcing an organizational security baseline, or applying consistent resource-tagging requirements. Instead of assigning each policy separately, an administrator can assign the initiative at an appropriate scope, including a management group, subscription, resource group, or individual resource. Azure then evaluates compliance for all policy definitions included in that initiative. Initiatives can also define parameters that are passed to included policy definitions at assignment time, allowing a common value—such as an approved list of Azure regions or a required tag name—to be consistently applied across the set. Microsoft uses built-in initiatives, including those aligned with security and compliance frameworks, and organizations can create custom initiatives for their own governance needs. This grouping capability makes initiatives the appropriate choice when related Azure policies must be organized and assigned together. For more information, see the Azure Policy overview and Microsoft's policy initiative definition documentation.

QUESTION NO: 141

What is the Service Level Agreement (SLA) provided for Azure services that are in public preview?

A. Each service defines its own SLA.

B. The SLA will be 99%.

C. The SLA will be 1% less than the general availability (GA) SLA.

D. The SLA will be 99.95%.

E. There is no SLA for Azure services in public preview.

ANSWER: E

Explanation:

There is no SLA for Azure services in public preview. Public preview offerings are made available to let customers evaluate functionality, test integrations, and provide feedback before the service or feature becomes generally available. Because preview capabilities can change, have limitations, or be discontinued, they are not intended for production workloads that require a defined availability commitment. Microsoft's Azure Preview Supplemental Terms state that preview products and services are provided "as-is," and service-level commitments do not apply unless Microsoft specifically states otherwise. Consequently, customers cannot rely on an uptime percentage, service-credit policy, or other guarantee normally documented in an Azure Service Level Agreement while a service remains in public preview. For workloads that need contractual availability targets, organizations should use generally available Azure services and review the SLA applicable to the selected service and deployment configuration. Once a service reaches general availability, Microsoft may publish the SLA terms and conditions that govern it. See the Microsoft Azure Preview Supplemental Terms and the Azure preview features documentation.

QUESTION NO: 142

You have a virtual machine named VM1 that runs Windows Server 2016. VM1 is in the East US Azure region.

Which Azure service should you use from the Azure portal to view service failure notifications that can affect the availability of VM1?

A. Azure Service Fabric.

B. Azure Monitor.

C. Azure virtual machines.

D. Azure Advisor .

ANSWER: B

Explanation:

Azure Monitor is correct because the Azure Service Health experience is accessed through Azure Monitor in the Azure portal and provides personalized notifications about Azure platform events that could affect subscribed resources. For VM1 in East US, Service Health can show active service issues, planned maintenance, health advisories, and security advisories relevant to the selected subscription, region, and Azure service. When an Azure-side failure affects Azure Virtual Machines in East US, the incident details include the affected service and region, status updates, mitigation information, and resolution status. This information helps determine whether reduced availability is caused by an Azure platform event rather than by the Windows Server guest operating system or the virtual machine configuration. Service Health can also be filtered to the relevant subscription and region, and administrators can create Service Health alerts in Azure Monitor to proactively send notifications through action groups. These notifications can use channels such as email, SMS, push notifications, voice calls, webhooks, or ITSM integrations. Microsoft documents Service Health as the personalized Azure resource-health notification capability, while Azure Monitor provides the portal experience and alerting integration used to view and act on these events. See the Azure Service Health overview and Create Service Health alerts in the Azure portal.

QUESTION NO: 143

You have an Azure subscription that contains several resource groups.

Which two statements are true about Azure resource groups? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. A resource can belong to only one resource group.

B. All resources in a resource group must be in the same Azure region.

C. Azure role-based access control can be assigned at the resource-group scope.

D. A resource group can contain resources from multiple subscriptions.

E. A resource group can contain only one type of Azure resource.

ANSWER: A C

Explanation:

A resource can belong to only one resource group. Azure Resource Manager uses resource groups as logical containers for resources that share a lifecycle, such as components of the same application or solution. Although a resource can be moved to another resource group when the resource type supports moving, its association at any given time is with one resource group. This organization enables resources to be deployed, managed, and deleted together where appropriate.

Azure role-based access control can be assigned at the resource-group scope. Azure RBAC supports role assignments at management group, subscription, resource group, and individual resource scopes. A role assigned to a resource group applies to the resources within that group, subject to inheritance and any more-specific assignments. This allows administrators to delegate access to a collection of related resources without creating separate role assignments for every individual resource. Resource groups are therefore an important scope for both organizing Azure resources and managing access consistently. For details, see Manage Azure resource groups and Understand scope for Azure RBAC.

QUESTION NO: 144

Your company uses a software as a service (SaaS) application.

Which two responsibilities remain with the company under the shared responsibility model? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Managing user identities and access.

B. Managing data.

C. Patching the application operating system.

D. Replacing failed physical servers.

E. Securing the physical datacenter.

ANSWER: A B

Explanation:

Managing user identities and access remains the company’s responsibility when it uses a SaaS application. Although the cloud provider operates the service, the company determines which employees, partners, and other users are authorized to use the application. This includes assigning appropriate roles and permissions, removing access when it is no longer needed, and configuring identity protections such as multifactor authentication and conditional access when those controls are available. Proper access management helps ensure that the SaaS service is used only by approved identities and according to the organization’s security requirements.

Managing data also remains the company’s responsibility. The company owns the decisions about the data it places into the SaaS application, including data classification, acceptable use, sharing permissions, retention requirements, and governance. It must ensure that sensitive information is handled in accordance with internal policies and applicable compliance obligations. The provider supplies and operates the SaaS application, but the customer remains accountable for the content it creates, stores, uploads, and shares through that service. Microsoft’s shared responsibility guidance identifies customer data and identities as customer responsibilities across cloud service models, including SaaS. See Shared responsibility in the cloud and the SaaS shared-responsibility model.

QUESTION NO: 145

This question requires that you evaluate the underlined text to determine if it is correct.

You deploy an Azure resource. The resource becomes unavailable for an extended period due to a service outage. Microsoft will automatically refund your bank account .

Instructions: Review the underlined text. If it makes the statement correct, select “No change is needed”. If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed.

B. automatically migrate the resource to another subscription

C. automatically credit your account

D. send you a coupon code that you can redeem for Azure credits

E. provide a Service Credit, if eligible, after you submit a valid claim

ANSWER: E

Explanation:

Microsoft’s Azure Service Level Agreements provide a potential *Service Credit* when an Azure service does not meet its stated service level, subject to the terms, eligibility requirements, exclusions, and calculation rules in the applicable SLA. A Service Credit is applied to the customer’s Azure account or billing arrangement as a credit against future payments; it is not a direct refund to the customer’s bank account. Importantly, this process is not automatic: the customer must submit a claim with the required supporting information within the claim period defined by the SLA. If Microsoft validates the claim and determines that the customer is eligible, it issues the Service Credit. Therefore, the statement is made accurate by describing an eligible Service Credit after a valid claim, rather than an automatic bank refund. Azure SLAs differ by service and may also depend on the deployment configuration, such as whether availability zones or multiple instances are used. Review the applicable product SLA and its service-credit terms for the specific resource involved. See Azure Service Level Agreements and Microsoft Azure terms and service-level conditions.

QUESTION NO: 146

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.

Your Azure environment contains multiple Azure virtual machines.

You need to ensure that a virtual machine named VM1 is accessible from the Internet over HTTP.

Solution: You modify a DDoS protection plan.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. A DDoS Protection plan is a security service intended to help protect Azure public IP resources from distributed denial-of-service attacks. It provides detection, mitigation, telemetry, and protection-related capabilities for resources associated with the plan. Modifying this plan can affect DDoS protection settings, but it does not publish a virtual machine application to the internet or create an HTTP access path to that virtual machine.

For VM1 to be reachable from the internet through HTTP, Azure must have an inbound network route to the workload, typically using a public IP address directly on the VM’s network interface or through a load balancer, application gateway, or

similar service. The relevant network security group must also allow inbound TCP traffic on port 80. These connectivity and traffic-filtering configurations determine whether HTTP requests can reach VM1; DDoS Protection is an additional protection layer for eligible public IP resources once that exposure exists.

Microsoft distinguishes DDoS mitigation from network access control and application publishing. See Azure DDoS Protection overview and Network security groups overview.

QUESTION NO: 147

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

A. networking layer

ANSWER: A

Explanation:

Azure DDoS Protection is implemented at the networking layer. Its purpose is to defend Azure workloads that are exposed through public IP addresses from distributed denial-of-service attacks, in which large volumes of malicious traffic attempt to exhaust network capacity or overwhelm resources so legitimate users cannot access an application. Azure continuously monitors inbound traffic and applies automated, adaptive mitigations when attack patterns are detected.

The service is a networking-layer protection because it operates on traffic directed at public endpoints associated with Azure virtual networks and public IP resources. Mitigation occurs before malicious traffic can consume connectivity and network resources needed by the hosted workload. Azure provides infrastructure-level DDoS protection automatically for Azure services, while DDoS Network Protection adds capabilities for protected virtual networks, including enhanced telemetry, attack-mitigation reports, and cost protection. These characteristics place Azure DDoS Protection squarely within network security controls rather than application-level protections.

Microsoft documents the service's scope, traffic-monitoring approach, and mitigation capabilities in the Azure DDoS Protection overview. For the capabilities available across protection tiers, see Azure DDoS Protection service tiers.

QUESTION NO: 148

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:

Azure IoT Hub. and **Azure Data Lake.** are the appropriate resources for this sensor-ingestion architecture. Azure IoT Hub is a managed service designed for secure, bidirectional communication between Azure and Internet of Things devices. It supports device identities, authentication, device-to-cloud telemetry ingestion, and management capabilities needed to connect a very large fleet of sensors reliably. IoT Hub is intended to scale device messaging and can route telemetry to downstream Azure services for storage and processing.

Azure Data Lake, typically implemented with Azure Data Lake Storage Gen2, provides massively scalable and durable storage for the high-volume data emitted by several million sensors. Its hierarchical namespace and integration with Azure

analytics services make it appropriate for retaining raw telemetry and historical data that will later be queried, transformed, or used in data engineering and machine-learning workloads. Combining IoT Hub for device connectivity and telemetry ingestion with Data Lake storage for long-term analytical storage establishes the core collection and retention layers of the planned solution. See What is Azure IoT Hub? and Introduction to Azure Data Lake Storage Gen2.

QUESTION NO: 149

Your Azure environment contains multiple Azure virtual machines.

You need to ensure that a virtual machine named VM1 is accessible from the Internet over HTTP.

What are two possible solutions? Each correct answer presents a complete solution.

A. Modify a DDoS protection plan.

B. Modify an Azure firewall.

C. Modify an Azure Traffic Manager profile.

D. Modify a network security group (NSG).

ANSWER: B D

Explanation:

Modifying an Azure firewall can provide HTTP access to VM1 by configuring an application rule or network rule as appropriate and, for inbound publishing, a destination network address translation (DNAT) rule. A DNAT rule maps traffic arriving at the firewall's public IP address and HTTP port to VM1's private IP address and port. Azure Firewall is therefore able to act as the controlled, publicly reachable entry point for the virtual machine while applying centralized firewall policy.

Modifying a network security group (NSG) can also enable the required access. An NSG can be associated with VM1's network interface or its subnet, and it can contain an inbound security rule that allows TCP traffic on destination port 80 from the required Internet source range. With public IP connectivity configured for the VM, this rule permits HTTP traffic to reach the workload. NSGs are Azure's fundamental distributed network traffic-filtering mechanism, using five-tuple rules to allow or deny inbound and outbound flows.

For implementation details, see Tutorial: Deploy and configure Azure Firewall using the Azure portal and Network security groups.

QUESTION NO: 150

You need to select Azure services for a solution that has the following requirements:

Store unstructured object data such as documents and images.

Provide a managed shared file system that can be accessed by using the Server Message Block (SMB) protocol.

Which two Azure Storage services should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure Blob Storage.

B. Azure Files.

C. Azure Queue Storage.

D. Azure Table Storage.

E. Azure Managed Disks.

ANSWER: A B

Explanation:

Azure Blob Storage is designed for storing massive amounts of unstructured object data, including documents, images, video, backups, and application data. Applications access blobs by using HTTP or HTTPS and Azure Storage APIs.

Azure Files provides managed file shares that can be mounted by using the SMB protocol. This makes Azure Files appropriate for workloads that require a shared file system and compatibility with Windows, Linux, and macOS clients that support SMB. Azure Queue Storage is intended for message storage, while Azure Table Storage is intended for NoSQL key-value data. See the Azure Blob Storage introduction and the Azure Files overview.

QUESTION NO: 151

This question requires that you evaluate the underlined text to determine if it is correct.

One of the benefits of Azure SQL Data Warehouse is that high availability is built into the platform.

Instructions: Review the underlined text. If it makes the statement correct, select “No change is needed”. If the

statement is incorrect, select the answer choice that makes the statement correct.

A. No change is needed

B. automatic scaling

C. data compression

D. versioning

ANSWER: A

Explanation:

“No change is needed” is correct. Azure SQL Data Warehouse, now named a dedicated SQL pool in Azure Synapse Analytics, is a fully managed platform service. Microsoft operates and maintains the underlying infrastructure, including the mechanisms that provide service resiliency and availability. This means an organization does not need to build, configure, or manage separate database servers, clustering, storage replication, and failover infrastructure to obtain the normal high-availability capabilities of the service.

Dedicated SQL pools also use an architecture that separates compute from storage. Data is persisted in Azure Storage, while compute resources can be managed independently. This design contributes to resilience and enables the platform to maintain data independently from compute operations. Microsoft provides automated backups and recovery capabilities for dedicated SQL pools, supporting restoration of data when required.

Therefore, built-in high availability is an appropriate benefit of this managed Azure data warehousing offering. Although the former Azure SQL Data Warehouse product name is used in the statement, Microsoft renamed the service to dedicated SQL pool within Azure Synapse Analytics; the underlying managed-service benefit remains applicable. See the Azure Synapse dedicated SQL pool overview and Microsoft’s dedicated SQL pool backup and restore guidance.