DUMPSBOSS.

Microsoft Azure Administrator

Microsoft AZ-104

Version Demo

Total Demo Questions: 95

Total Premium Questions: 957

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Manage Azure identities and governance	198
Topic 2, Implement and manage storage	128
Topic 3, Deploy and manage Azure compute resources	213
Topic 4, Implement and manage virtual networking	293
Topic 5, Monitor and maintain Azure resources	122
Topic 6, Mix Questions	3
Total	957

QUESTION NO: 1

You need to define a custom domain name for Azure AD to support the planned infrastructure.

Which domain name should you use?

A. Join the client computers in the Miami office to Azure AD.

B. Add http://autologon.microsoftazuread-sso.com to the intranet zone of each client computer in the Miami office.

C. Allow inbound TCP port 8080 to the domain controllers in the Miami office.

D. Install Azure AD Connect on a server in the Miami office and enable Pass-through Authentication

E. Install the Active Directory Federation Services (AD FS) role on a domain controller in the Miami office.

F. humongousinsurance.com

G. Use humongousinsurance.com as a verified Azure AD custom domain.

ANSWER: F G

Explanation:

The appropriate custom domain is **humongousinsurance.com**. It is the existing on-premises Active Directory domain name and therefore the user principal name suffix most closely aligned with the organization's current identity infrastructure. Before it can be used in Microsoft Entra ID (formerly Azure AD), Humongous Insurance must add the domain as a custom domain and prove ownership by creating the DNS record that Microsoft provides. Once verification is complete, the domain is available for user sign-in names such as **user@humongousinsurance.com**. This provides a consistent sign-in experience for users whose on-premises identities are synchronized to Microsoft Entra ID and avoids requiring them to use the tenant's default **.onmicrosoft.com** domain. The option that explicitly states that the domain is to be verified in Azure AD describes the same required configuration outcome: the organization must use its owned and verified corporate namespace, humongousinsurance.com. Microsoft documents the custom-domain addition and DNS-verification process at Add a custom domain name in Microsoft Entra ID. For synchronized identities, Microsoft also documents how verified domains and user principal name suffixes are used with directory synchronization at Microsoft Entra Connect prerequisites.

QUESTION NO: 2

You have an Azure subscription that contains a resource group named RG1. RG1 contains existing resources that do not have a tag named CostCenter.

You need to ensure that all existing and future resources in RG1 have a CostCenter tag. The tag value must be inherited from RG1.

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 inherits a tag from the resource group and uses the Modify effect

B. Create a remediation task for the policy assignment

C. Assign an Azure Policy that uses the Deny effect

D. Create a resource lock on RG1

E. Add the CostCenter tag manually to each existing resource

ANSWER: A B

Explanation:

Assigning a policy that uses the **Modify** effect to RG1 can add or replace a tag on resources during create and update operations. The built-in policy definition that inherits a tag and its value from the resource group is appropriate for applying the CostCenter value from RG1 to future resources.

A remediation task is required to apply the Modify policy to resources that already exist. The remediation task evaluates the existing noncompliant resources and applies the policy operation to add the inherited tag. A managed identity configured for the policy assignment must have the permissions required by the policy definition to modify the target resources.

For more information, see Azure Policy modify effect and Remediate non-compliant resources with Azure Policy.

QUESTION NO: 3

You have an Azure subscription,

You have an on-premises virtual machine named VM1, The settings for VM" are shown in the exhibit. (Click the Exhibit tab.)

You need to ensure that you can use the disks attached to VM' as a template for Azure virtual machines,

What should you modify on VM1?

A. Integration Services

B. the processor

C. the hard drive

D. the memory

ANSWER: C

Explanation:

The hard drive must be modified because Azure requires an uploaded operating-system disk to use the supported VHD format. For a Windows virtual machine prepared from an on-premises Hyper-V VM, the OS disk must be a fixed-size VHD; Azure does not support using a VHDX file directly as an uploaded OS disk image. Therefore, if VM1's attached disk is configured as a VHDX or as a dynamically expanding VHD, convert it to a fixed-size VHD before preparing and uploading the image. After the disk is in the required format, prepare the guest operating system for imaging, including generalizing it when creating a reusable image, and upload the resulting VHD to Azure. The uploaded disk can then be used to create a managed image or an Azure Compute Gallery image version, which supports deploying consistent Azure virtual machines from that template. Microsoft's image-preparation guidance explicitly states the fixed-size VHD requirement and provides the supported preparation workflow. See Prepare a Windows VHD or VHDX to upload to Azure and Upload a generalized VHD and use it to create new VMs in Azure.

QUESTION NO: 4 - (SIMULATION)

You have an Azure subscription.

You need to deploy a virtual machine by using an Azure Resource Manager (ARM) template.

How should you complete the template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



ANSWER: See the explanation for the answer

Explanation:

Select the following options:

For the `dependsOn` entry, select `resourceId`.
Under `storageProfile`, select `imageReference`.

The completed relevant template section is:

```
"dependsOn": [
  "[resourceId('Microsoft.Network/networkInterfaces', 'VM1')]"
],
"properties": {
  "storageProfile": {
    "imageReference": {
      "publisher": "MicrosoftWindowsServer",
      "offer": "WindowsServer",
      "sku": "2019-Datacenter",
      "version": "latest"
    }
  }
}
```

`resourceId()` creates the resource ID for the network interface so that the VM deployment waits for that NIC. The VM image details belong in the `imageReference` object.

QUESTION NO: 5

You deploy an Azure Application Gateway.

You need to ensure that all the traffic requesting https://adatum.com/internal resources is directed to an internal server pool and all the traffic requesting https://adatum.com/external resources is directed to an external server pool.

What should you configure on the Application Gateway?

A. URL path-based routing

B. multi-site listeners

C. basic routing

D. SSL termination

ANSWER: A

Explanation:

URL path-based routing is the appropriate Application Gateway configuration because both request types use the same host name, `adatum.com`, but must be sent to different backend pools according to the URL path. A path-based routing rule uses a URL path map to evaluate the path portion of each incoming request and associate matching paths with backend targets. Configure path rules such as `/internal/*` to use the internal server pool and `/external/*` to use the external server pool. The rule also includes a default backend pool, which receives requests that do not match either configured path.

This routing model enables one HTTPS listener for the host name while providing granular, Layer 7 request routing. The listener accepts the HTTPS request, and the request-routing rule applies the path map before forwarding the request to the relevant backend pool and its associated HTTP settings. Microsoft documents URL path maps specifically for routing traffic to different backend pools based on request paths. For configuration details and path-pattern behavior, see Azure Application Gateway URL routing overview. For the components involved in listeners, routing rules, backend pools, and backend settings, see Application Gateway configuration overview.

QUESTION NO: 6

You need to configure an Azure web app named contoso.azurewebsites.net to host wwwcontoso.com.

What should you do first?

A. Create a TXT record named wwsv.contoso.com that has a value of contosoazurewebsites.net.

B. Create a TXT record named asuid that contains the domain verification ID.

C. Create a CNAME record named asuid that contains the domain verification ID.

D. Create A records named www.contoso.com and asuid.contoso.com.

E. Create a CNAME record named www that points to contoso.azurewebsites.net.

ANSWER: E

Explanation:

Create a CNAME record named www that points to contoso.azurewebsites.net. A CNAME record is the appropriate DNS record type for mapping a subdomain such as `www.contoso.com` to an Azure App Service app's default hostname. This establishes the DNS association that directs requests for the custom host name to the web app. After the CNAME record is created and publicly resolvable, the custom domain can be added to the App Service app through the Azure portal, Azure CLI, or an ARM-based deployment. App Service validates the domain mapping as part of the custom-domain binding process.

The target of the CNAME must be the app's assigned `azurewebsites.net` host name, here `contoso.azurewebsites.net`, rather than an IP address. This approach also lets Azure manage infrastructure changes without requiring the DNS record for the custom subdomain to be updated. The DNS zone hosting `contoso.com` should therefore contain a record with host/name `www` and canonical target `contoso.azurewebsites.net`. See Tutorial: Map an existing custom DNS name to Azure App Service and Configure a CNAME record.

QUESTION NO: 7

You have an Azure virtual machine named VM1 and an Azure key vault named Vault1.

On VM1, you plan to configure Azure Disk Encryption to use a key encryption key (KEK)

You need to prepare Vault! for Azure Disk Encryption.

Which two actions should you perform on Vault1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Create a new key.

B. Select Azure Virtual machines for deployment

C. Configure a key rotation policy.

D. Create a new secret.

E. Select Azure Disk Encryption for volume encryption

ANSWER: A E

Explanation:

Create a new key. is required because the planned configuration uses a key encryption key. Azure Disk Encryption uses a key encryption key stored in Azure Key Vault to protect the disk-encryption key or secret. The key must exist in the vault before it can be supplied when enabling disk encryption. A suitable key can be generated in the vault or imported, provided that it meets the supported key requirements and is available to the encryption configuration.

Select Azure Disk Encryption for volume encryption is also required. This Key Vault advanced access-policy setting authorizes the Azure Disk Encryption service to retrieve the keys and secrets needed during encryption and when the virtual machine starts. In PowerShell, this corresponds to enabling the vault for disk encryption, such as through the `EnabledForDiskEncryption` setting. If the vault uses network restrictions, the required trusted-service or network access configuration must additionally permit Azure Disk Encryption to reach the vault. Microsoft documents the Key Vault preparation and the required disk-encryption access setting at Azure Disk Encryption and Azure Key Vault. For the key encryption key workflow and supported parameters when enabling encryption, see Enable Azure Disk Encryption with Azure Key Vault.

QUESTION NO: 8

You have an Azure subscription that contains the virtual networks shown in the following table.

You need to ensure that all the traffic between VNet1 and VNet2 traverses the Microsoft backbone network. What should you configure?

A. ExpressRoute

B. a private endpoint

C. peering

D. a route table

ANSWER: C

Explanation:

peering is the appropriate configuration because Azure Virtual Network peering directly connects two virtual networks, allowing resources in each network to communicate using private IP addresses. Traffic sent between peered virtual networks stays on the Microsoft global network rather than being routed through the public internet. This provides a private, low-latency path for communication between VNet1 and VNet2 while preserving the networks as separate virtual networks with their own address spaces and resources.

VNet peering can be configured for virtual networks in the same Azure region (virtual network peering) or in different regions (global virtual network peering). In either case, Azure carries the traffic over its backbone infrastructure. After peering is established in both directions and the address spaces do not overlap, Azure automatically provides system routes for the remote virtual network address prefixes, enabling direct connectivity. For the relevant configuration requirements and behavior, see Virtual network peering overview and Azure virtual network traffic routing.

QUESTION NO: 9

You have a Microsoft 365 tenant and an Azure Active Directory (Azure AD) tenant named contoso.com.

You plan to grant three users named User1, User2, and User3 access to a temporary Microsoft SharePoint document library named Library1.

You need to create groups for the users. The solution must ensure that the groups are deleted automatically after 180 days.

Which two groups should you create? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. a Security group that uses the Assigned membership type

B. an Office 365 group that uses the Assigned membership type

C. an Office 365 group that uses the Dynamic User membership type

D. a Security group that uses the Dynamic User membership type

E. a Security group that uses the Dynamic Device membership type

ANSWER: B C

Explanation:

An Office 365 group that uses the Assigned membership type and an Office 365 group that uses the Dynamic User membership type are appropriate because Microsoft 365 group expiration policies apply to Microsoft 365 (formerly Office 365) groups. An administrator can configure the expiration interval as 180 days and apply the policy to all Microsoft 365 groups or to selected groups. When a group reaches its expiration date and is not renewed, Microsoft Entra ID deletes it along with the Microsoft 365 services associated with that group, including its SharePoint site and content. This supports automatic cleanup of temporary collaboration resources.

Assigned membership is suitable when User1, User2, and User3 should be explicitly added as members. Dynamic User membership is also supported for a Microsoft 365 group and can populate membership automatically based on a user-based membership rule. In either case, the group type remains a Microsoft 365 group, so it can be governed by the configured group expiration policy. Microsoft documents group expiration and associated-service deletion at Configure the expiration policy for Microsoft 365 groups and documents supported dynamic group types and membership rules at Dynamic membership groups.

QUESTION NO: 10

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 deploy an Azure Kubernetes Service (AKS) cluster named AKS1. You need to deploy a YAML file to AKS1.

Solution: From Azure CLI, you run az aks. Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. `az aks` is the Azure CLI command group for administering Azure Kubernetes Service resources. It provides AKS management commands, including commands to create and update clusters, manage node pools, and retrieve cluster

access credentials. Running the command group name by itself does not submit a Kubernetes manifest or create the resources described in a YAML file.

A typical command-line workflow is to use `az aks get-credentials` for AKS1, which merges the cluster credentials into the local kubeconfig file. Then, use the Kubernetes command-line tool to deploy the manifest, for example `kubectl apply -f deployment.yaml`. The `kubectl apply` operation sends the YAML-defined Kubernetes objects to the AKS cluster API server, creating or updating those objects as appropriate.

Therefore, merely running `az aks` does not meet the requirement to deploy a YAML file. Azure CLI can be part of the authentication and cluster-access setup, but a manifest application command such as `kubectl apply -f` is required for the deployment itself. See Quickstart: Deploy an AKS cluster using the Azure CLI and the Azure CLI az aks reference.

QUESTION NO: 11

You have an Azure subscription that contains a Log Analytics workspace named workspace1 and a virtual machine named VM1.

You need to collect VM1 performance data and send the data to workspace1 by using Azure Monitor Agent.

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

NOTE: Each correct selection is worth one point.

A. Create a data collection rule that has workspace1 as a destination

B. Create a data collection rule association for VM1

C. Create an Azure Automation account

D. Enable boot diagnostics on VM1

E. Create an Azure Service Health alert

ANSWER: A B

Explanation:

You must create a **data collection rule (DCR)** that specifies the performance counters to collect and configures workspace1 as a Log Analytics destination. A DCR defines what data Azure Monitor Agent collects and where the agent sends that data.

You must associate the DCR with VM1. The association instructs Azure Monitor Agent on VM1 to use the rule. VM1 must also have Azure Monitor Agent installed, either directly or through an Azure Policy initiative that deploys the agent and associates the DCR.

For more information, see Collect performance counters with Azure Monitor Agent and Data collection rules in Azure Monitor.

QUESTION NO: 12

You plan to automate the deployment of a virtual machine scale set that uses the Windows Server 2016 Datacenter image.

You need to ensure that when the scale set virtual machines are provisioned, they have web server components installed.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Upload a configuration script

B. Create an automation account

C. Create an Azure policy

D. Modify the extensionProfile section of the Azure Resource Manager template

E. Create a new virtual machine scale set in the Azure portal

ANSWER: A D

Explanation:

Upload a configuration script is required so that the deployment has an accessible script containing the Windows feature installation commands, such as a PowerShell command to install the IIS Web-Server feature. The script can be stored in a location reachable by the scale set instances, commonly an Azure Storage account or a source repository. Supplying the script externally keeps the VM image unchanged while ensuring every newly provisioned instance can run the same installation procedure.

Modify the extensionProfile section of the Azure Resource Manager template is required to configure a VM extension for the scale set model. In this scenario, the Custom Script Extension downloads the uploaded configuration script and executes it during instance provisioning. Because extensions defined in the scale set model are applied to instances, this approach consistently installs the web-server components on the initial instances and on instances added later through scale-out. The extension settings identify the script location and the command to execute, allowing the entire process to be declared and repeated through the deployment template. Microsoft documents this pattern for installing applications on Virtual Machine Scale Sets through the Custom Script Extension and an extension profile. See Install applications in a scale set with an ARM template and Custom Script Extension for Windows.

QUESTION NO: 13

You have an Azure virtual machine named VM1 that is protected by Azure Backup.

You need to restore a single file from a recovery point without restoring the entire virtual machine.

What should you use?

A. File Recovery

B. Restore VM

C. Restore disks

D. Site Recovery failover

ANSWER: A

Explanation:

File Recovery is correct. Azure Backup File Recovery enables an administrator to recover individual files and folders from an Azure VM recovery point. The recovery point is mounted to a machine, allowing the required files to be copied to the intended destination without replacing VM1 or its disks.

Restoring the virtual machine or restoring disks would recover a broader set of data than required and can require additional deployment steps. File Recovery is designed for granular recovery scenarios in which only selected files or folders are needed. See Restore files from Azure virtual machine backup.

QUESTION NO: 14

You have an Azure subscription that contains a virtual network named VNet1 and an Azure private DNS zone named contoso.local.

You deploy virtual machines to VNet1. You need host names for the virtual machines to be registered automatically in contoso.local.

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

NOTE: Each correct selection is worth one point.

A. Create a virtual network link for VNet1

B. Enable auto-registration on the virtual network link

C. Create a private endpoint in VNet1

D. Create a public DNS zone named contoso.local

E. Create a service endpoint for Microsoft.Network

ANSWER: A B

Explanation:

You should create a **virtual network link** between VNet1 and the private DNS zone. A virtual network link makes the private DNS zone available to resources in the linked virtual network for name resolution.

You should enable **auto-registration** on the virtual network link. When auto-registration is enabled, Azure automatically creates and manages DNS records in the private zone for virtual machines that have network interfaces in the linked virtual network. A private endpoint does not provide automatic registration for virtual machine host names, and a public DNS zone is not used for private Azure virtual network name resolution. See Azure private DNS zone auto registration.

QUESTION NO: 15

You plan to deploy three Azure virtual machines named VM1, VM2, and VM3. The virtual machines will host a web app named App1.

You need to ensure that at least two virtual machines are available if a single Azure datacenter becomes unavailable.

What should you deploy?

A. all three virtual machines in a single Availability Zone

B. all virtual machines in a single Availability Set

C. each virtual machine in a separate Availability Zone

D. each virtual machine in a separate Availability Set

ANSWER: C

Explanation:

The correct deployment is **each virtual machine in a separate Availability Zone**. Availability Zones are physically separate datacenter locations within an Azure region. Each zone has independent power, cooling, and networking infrastructure, so distributing the three virtual machines across three zones provides resilience against the failure of one complete datacenter location.

With VM1, VM2, and VM3 each assigned to a different zone, the loss of any one zone affects only the virtual machine hosted in that zone. The two virtual machines in the remaining zones continue to run, satisfying the requirement that at least two virtual machines remain available during a single-datacenter outage. The web application should also be designed to direct traffic only to healthy virtual machine instances, typically by using an Azure Load Balancer or Application Gateway with health probes.

This design depends on deploying in an Azure region that supports Availability Zones and selecting zone assignments during virtual machine creation. Microsoft describes Availability Zones as the appropriate resilience feature for protecting workloads from datacenter-level failures within a region. See Availability Zones overview and Availability options for Azure virtual machines.

QUESTION NO: 16

You have an Azure subscription named Subscription1.

You deploy a Linux virtual machine named VM1 to Subscription1. You need to monitor the metrics and the logs of VM1.

What should you use?

A. Azure HDInsight

B. Linux Diagnostic Extension (LAD) 3.0

C. the AzurePerformanceDiagnostics extension

D. Azure Analysis Services

ANSWER: B

Explanation:

Linux Diagnostic Extension (LAD) 3.0 is the appropriate choice because it is designed to collect both guest-level diagnostic data and performance metrics from a Linux virtual machine. After the extension is installed and configured on VM1, it can collect Linux syslog events and specified log files, as well as metrics such as processor, memory, disk, and network-related performance data. The collected data can be delivered to configured Azure destinations, enabling centralized monitoring and analysis of the VM's operational health and activity.

This is important because Azure platform metrics alone do not provide all information generated within the guest operating system. A guest diagnostics extension supplies the collection mechanism needed for operating-system logs and guest performance counters. LAD configuration defines the metrics, log sources, sampling intervals, and destination settings, allowing monitoring to be tailored to the workload running on VM1.

Microsoft documents the Linux Diagnostic Extension and its supported metrics and log collection configuration in Use the Azure Diagnostics extension for Linux. For broader Azure Monitor guidance on collecting monitoring data from virtual machines, see Azure Monitor VM insights overview.

QUESTION NO: 17

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 deploy an Azure Kubernetes Service (AKS) cluster named AKS1. You need to deploy a YAML file to AKS1.

Solution: From Azure CLI, you run the kubectl client. Does this meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes. The `kubectl` client is the standard Kubernetes command-line tool for deploying resources that are defined in YAML manifest files, and it can be used with an Azure Kubernetes Service cluster. After authenticating to Azure, an administrator normally obtains AKS1's Kubernetes configuration and credentials with `az aks get-credentials`. This command merges the cluster connection information into the local kubeconfig file and sets up `kubectl` to communicate with the AKS API server. The YAML manifest can then be deployed by running a command such as `kubectl apply -f deployment.yaml`. This creates or updates the Kubernetes resources described in the file, such as deployments, pods,

services, config maps, or ingress resources. Azure CLI is therefore an appropriate environment from which to run `kubectl`, provided that the kubectl client is installed and configured for the target cluster. Microsoft documents the AKS connection workflow, including retrieving credentials before using kubectl, in its AKS Kubernetes deployment quickstart. The Azure CLI reference for az aks get-credentials describes the command used to configure access to the cluster.

QUESTION NO: 18

You have an app named App1 that runs on an Azure web app named webapp1.

The developers at your company upload an update of App1 to a Git repository named GUI.

Webapp1 has the deployment slots shown in the following table.

You need to ensure that the App1 update is tested before the update is made available to users.

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

NOTE: Each correct selection is worth one point.

A. Swap the slots

B. Deploy the App1 update to webapp1-prod, and then test the update

C. Stop webapp1-prod

D. Deploy the App1 update to webapp1-test, and then test the update

E. Stop webapp1-test

ANSWER: A D

Explanation:

Deploy the App1 update to webapp1-test, and then test the update provides the required safe release workflow. An Azure App Service deployment slot is a separate live app environment with its own hostname. Deploying the revision to the test slot allows the team to validate the new application version by using that slot's URL, without changing the version currently served to production users. This supports functional, integration, and smoke testing before release.

After the update has passed testing, Swap the slots promotes the tested version from webapp1-test to webapp1-prod. App Service slot swapping exchanges the app content between the source and target slots and applies the production slot configuration to the app before the final swap, helping the application warm up under its production settings. The tested revision then becomes available through the production endpoint, while the former production revision is retained in the test slot. This approach supports a controlled release with minimal interruption and also leaves the previous version available in the nonproduction slot for validation or rollback planning.

See Set up staging environments in Azure App Service and Swap two slots.

QUESTION NO: 19

You have an Azure subscription that contains a virtual network named VNet1.

You plan to deploy Azure Bastion to provide secure Remote Desktop Protocol (RDP) and Secure Shell Protocol (SSH) access to virtual machines in VNet1. The solution must not expose the virtual machines by using public IP addresses.

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

NOTE: Each correct selection is worth one point.

A. A subnet named AzureBastionSubnet

B. A Standard SKU public IP address

C. A public IP address for each virtual machine

D. A virtual network gateway

E. An Application Gateway

ANSWER: A B

Explanation:

An **AzureBastionSubnet** is required for Azure Bastion. Azure Bastion is deployed into a dedicated subnet that must be named AzureBastionSubnet. The subnet is reserved for the Bastion host and cannot contain virtual machines or other workload resources.

A **Standard SKU public IP address** is also required for a Bastion deployment. The public IP address is assigned to the Bastion host, not to the protected virtual machines. Administrators connect to the Bastion service through the Azure portal or a supported client, and Bastion establishes the RDP or SSH connection to the private IP address of the target VM. See What is Azure Bastion? and Deploy Azure Bastion by using the Azure portal.

QUESTION NO: 20

You have a Microsoft 365 tenant and an Azure Active Directory (Azure AD) tenant named contoso.com.

You plan to grant three users named User1, User2, and User3 access to a temporary Microsoft SharePoint document library named Library1.

You need to create groups for the users. The solution must ensure that the groups are deleted automatically after 180 days.

Which two groups should you create? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. a Microsoft 365 group that uses the Assigned membership type

B. a Security group that uses the Assigned membership type

C. a Microsoft 365 group that uses the Dynamic User membership type

D. a Security group that uses the Dynamic User membership type

E. a Security group that uses the Dynamic Device membership type

ANSWER: A C

Explanation:

A Microsoft 365 group that uses the Assigned membership type supports Microsoft 365 group expiration policies. An administrator can configure a group expiration policy with a 180-day lifetime and apply it to all Microsoft 365 groups or selected groups. Assigned membership is appropriate when the membership must be explicitly controlled, such as manually adding User1, User2, and User3 to provide access to the SharePoint-connected resources. Microsoft 365 groups also provide the collaboration-resource model used by SharePoint, including the associated SharePoint site and document library.

A Microsoft 365 group that uses the Dynamic User membership type is also eligible for the same Microsoft 365 group expiration policy. Dynamic user membership allows Microsoft Entra ID to calculate membership from a configured user-attribute rule, while the group remains a Microsoft 365 group for lifecycle purposes. Therefore, it can be governed by the 180-day expiration configuration while providing automated membership management. When an unrenewed group reaches its expiration date, it is deleted, and its connected Microsoft 365 resources are deleted as part of the group lifecycle. Microsoft documents group expiration policy configuration and the supported Microsoft 365 group scope at Microsoft Entra group expiration. Dynamic membership for Microsoft 365 groups is documented at Dynamic membership groups.

QUESTION NO: 21

You have an Azure subscription that contains three virtual networks named VNet1, VNet2, VNet3. VNet2

contains a virtual appliance named VM2 that operates as a router.

You are configuring the virtual networks in a hub and spoke topology that uses VNet2 as the hub network.

You plan to configure peering between VNet1 and VNet2 and between VNet2 and VNet3.

You need to provide connectivity between VNet1 and VNet3 through VNet2.

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

NOTE: Each correct selection is worth one point.

A. On the peering connections, allow forwarded traffic.

B. On the peering connections, allow gateway transit.

C. Create route tables and assign the table to subnets.

D. Create a route filter.

E. On the peering connections, use remote gateways.

ANSWER: A C

Explanation:

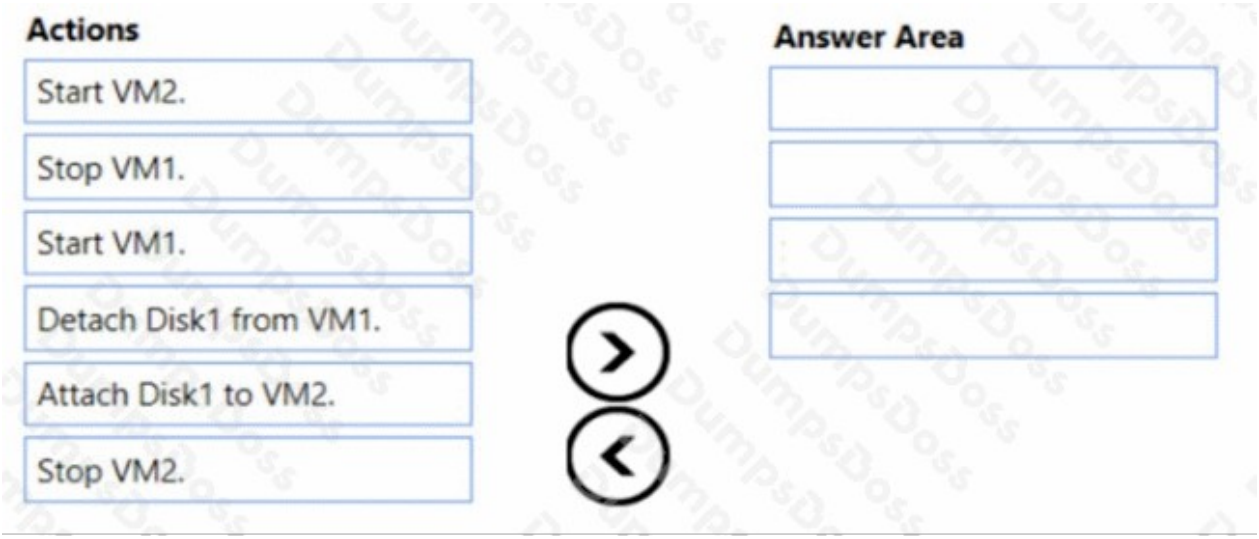
On the peering connections, allow forwarded traffic. is required because VM2 is a network virtual appliance that receives packets from one spoke and forwards them toward the other spoke. Azure virtual network peering is not transitive by default, and forwarded traffic must be enabled on the relevant peering configurations to permit traffic that originated outside the directly peered virtual network to traverse the peering. This setting supports routing through an NVA deployed in the hub.

Create route tables and assign the table to subnets. is also required to explicitly direct traffic between the spoke address spaces through VM2. A user-defined route on the applicable VNet1 subnet should send VNet3 prefixes to VM2's private IP address with a next-hop type of *Virtual appliance*; corresponding routes on applicable VNet3 subnets should send VNet1 prefixes to VM2. Associating these route tables with the workload subnets ensures their effective routes select VM2 rather than relying on the nontransitive system routes created by peering. VM2 must also be configured to forward IP traffic and apply any required routing or firewall policies. See Azure virtual network peering and Virtual network traffic routing.

QUESTION NO: 22 - (DRAG DROP)

You have two Azure virtual machines named VM1 and VM2. VM1 has a single data disk named Disk1. You need to attach Disk1 to VM2. The solution must minimize downtime for both virtual machines.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



ANSWER:



Explanation:

To move Disk1 from VM1 to VM2, the disk must first no longer be attached to VM1. The correct sequence begins by stopping VM1, which provides the required state for safely changing its disk configuration. Disk1 is then detached from VM1. Once the detach operation has completed, VM1 can be started immediately, minimizing the period during which the source virtual machine is unavailable.

After Disk1 has been detached, it becomes an available managed disk that can be attached to another virtual machine. The final action is to attach Disk1 to VM2. VM2 does not need to be stopped as part of this sequence, so the solution avoids unnecessary downtime on the destination virtual machine. The key is that the disk cannot be attached to VM2 until its existing attachment to VM1 has been removed; starting VM1 again before attaching the disk to VM2 restores the source workload as early as possible.

This ordering therefore limits downtime to VM1 only and only for the interval necessary to detach its data disk. It also preserves the disk and its contents because the operation changes the disk attachment rather than deleting or recreating the managed disk. Microsoft's guidance for managing Azure VM disks is available in the detach a data disk documentation and the attach a managed data disk documentation.

QUESTION NO: 23

You need to deploy an Azure virtual machine scale set that contains five instances as quickly as possible. What should you do?

A. Deploy five virtual machines. Modify the Size setting for each virtual machine.

B. Deploy live virtual machines. Modify the Availability Zones setting for each virtual machine.

C. Deploy one virtual machine scale set that is set to ScaleSetVM orchestration mode.

D. Deploy one virtual machine scale set that is set to VM (virtual machines) orchestration mode.

ANSWER: C

Explanation:

Deploy one virtual machine scale set that is set to ScaleSetVM orchestration mode. ScaleSetVM is the Azure portal label for the Uniform orchestration mode. Uniform mode is designed for a set of identical VM instances that are deployed and managed as a single scale-set resource. This model uses a common VM configuration and is optimized for scale-out scenarios in which Azure provisions homogeneous instances together. Consequently, it is the appropriate orchestration mode when the primary requirement is to create five matching scale-set instances as quickly as possible.

For a Uniform scale set, the desired capacity is specified on the scale set, and Azure creates the requested number of instances from the same model. This provides the streamlined deployment behavior needed here while retaining scale-set capabilities such as autoscaling, load-balancing integration, and centralized configuration. Microsoft identifies Uniform orchestration as suitable for large-scale, identical workloads and scenarios that require optimized deployment speed. The initial capacity can be set to five during deployment, allowing Azure to create the five instances as part of the scale-set deployment.

For details, see Azure VM Scale Sets orchestration modes and Virtual Machine Scale Sets overview.

QUESTION NO: 24

You have an Azure Active Directory (Azure AD) tenant that syncs to on-premises Active Directory and contains the users shown in the following table.

[IMAGE_1_OCR_EXTRACTED_TEXT_START] Name Type Source User1 Member Azure AD User2 Member Azure AD User3 Member Windows Server Active Directory User4 Guest Microsoft account
[IMAGE_1_OCR_EXTRACTED_TEXT_END]

You create a group named Group1 and add User1 to the group. You need to configure the ownership of Group1. Which users can you add as owners of Group1?

A. East US, West Europe, and North Europe

B. East US and West Europe only

C. East US only

D. East US and North Europe only

E. User1, User2, User3, and User4

ANSWER: E

Explanation:

User1, User2, User3, and User4 is correct. Microsoft Entra ID group ownership can be assigned to user objects in the directory regardless of whether the member user was created directly in Microsoft Entra ID or synchronized from Windows Server Active Directory. Therefore, the cloud-only member accounts, User1 and User2, and the synchronized member account, User3, are eligible to be designated as owners. A user does not need to be a member of a group before being assigned as an owner, so the fact that only User1 was initially added as a member does not limit the available owners.

A B2B guest user is also represented by a user object in the Microsoft Entra tenant. Consequently, User4, whose identity originates from a Microsoft account and is present as a guest, can be assigned group-owner responsibility. Group owners can manage group membership and applicable group settings, subject to the permissions and policies configured in the tenant. This permits organizations to delegate group management to internal, synchronized, and external collaborator

identities when appropriate. Microsoft documents that owners can be added from users in the directory and describes guest participation in group scenarios. See Manage Microsoft Entra groups and owners and Microsoft Entra group members and guests.

QUESTION NO: 25

You have the Azure virtual machines shown in the following table.

Name	IP address	Virtual network
VM1	10.0.0.4	VNET1
VM2	10.0.0.5	VNET1

VNET1 is linked to a private DNS zone and named contoso.com that contains the records shown in the following table.

Name	Type	TTL	Value	Auto registered
comp1	TXT	3600	10.0.0.5	False
comp2	A	3600	10.0.0.5	False
comp3	CNAME	3600	comp1.contoso.com	False
comp4	PTR	3600	10.0.0.5	False

You need to ping VM2 from VM1.

Which DNS names can you use to ping VM2.

A. comp2 contoso.com only

B. com1.contoso.com and comp2.contoso.com only

C. comp2.contoso.com and comp4.contoso.com only

D. comp1.contoso.com, comp2.contoso.com and comp4.contoso.con only

E. comp1.contoso.com comp2contoso.com.comp3.contoso.com and comp4.contoso.com

F. comp2.contoso.com, comp3.contoso.com, and comp4.contoso.com only

ANSWER: F

Explanation:

comp2.contoso.com, comp3.contoso.com, and comp4.contoso.com only is correct because these names resolve, through the private DNS zone linked to VNET1, to VM2. An A record maps a DNS name directly to an IP address, while a CNAME record provides an alias that ultimately resolves to the target host's A record. Therefore, both direct A-record names assigned to VM2 and an alias whose resolution leads to VM2 can be used when pinging from VM1.

Azure Private DNS zones provide DNS resolution to virtual networks that are linked to the zone. As VM1 is in VNET1 and VNET1 is linked to *contoso.com*, VM1 can query the records in that private zone and resolve the applicable fully qualified domain names to VM2's private IP address. The names must be used as fully qualified names because the question identifies records in the *contoso.com* zone. For more information, see Azure Private DNS overview and Azure DNS zones and records overview.

QUESTION NO: 26

You have an Azure Storage account named storage1 and a virtual network named VNet1.

You need to provide private access to the blob service in storage1 from VNet1. Clients must resolve the storage account blob endpoint to a private IP address.

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

NOTE: Each correct selection is worth one point.

A. Create a private endpoint for the blob service of storage1

B. Create a private DNS zone named privatelink.blob.core.windows.net

C. Link the private DNS zone to VNet1

D. Create a Microsoft.Storage service endpoint on the subnet

E. Create a public IP address for storage1

F. Create a local network gateway

ANSWER: A B C

Explanation:

You must create a private endpoint for the blob subresource of storage1 in VNet1. The private endpoint receives a private IP address from the selected subnet and provides private connectivity to the storage account blob service.

You must create the private DNS zone named `privatelink.blob.core.windows.net` and link the zone to VNet1. The virtual network link allows resources in VNet1 to resolve names in the private DNS zone. Associating the private endpoint with a private DNS zone group creates the required DNS record for the storage account endpoint.

For more information, see Azure private endpoint DNS configuration and Use private endpoints for Azure Storage.

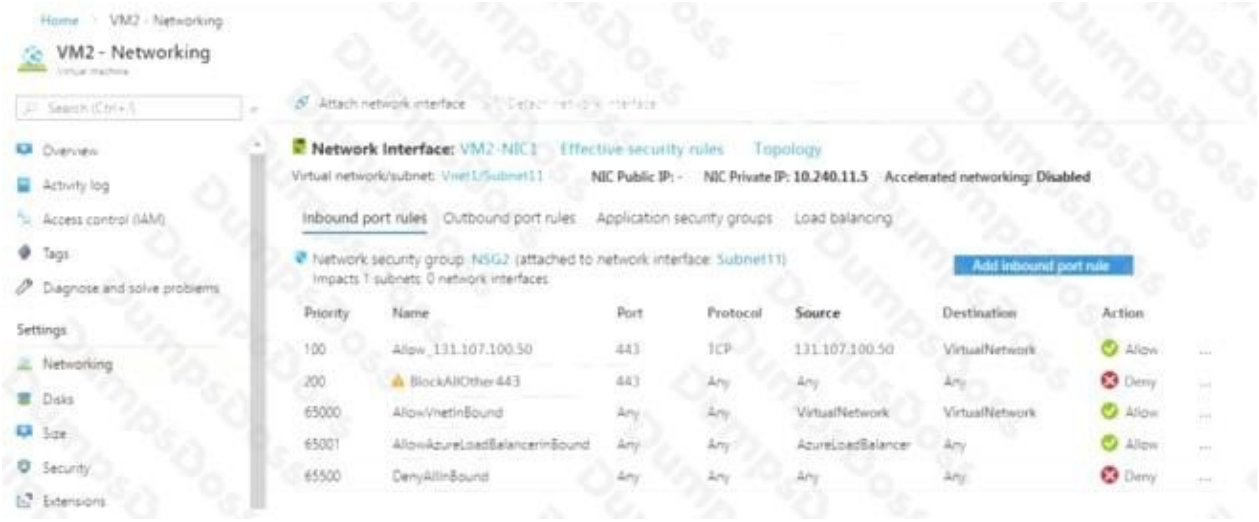
QUESTION NO: 27

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 have an app named App1 that is installed on two Azure virtual machines named VM1 and VM2. Connections to App1 are managed by using an Azure Load Balancer.

The effective network security configurations for VM2 are shown in the following exhibit.



You discover that connections to App1 from 131.107.100.50 over TCP port 443 fail.

You verify that the Load Balancer rules are configured correctly.

You need to ensure that connections to App1 can be established successfully from 131.107.100.50 over TCP port 443.

Solution: You create an inbound security rule that allows any traffic from the AzureLoadBalancer source and has a cost of 150.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. An inbound rule whose source is the *AzureLoadBalancer* service tag does not permit application traffic originating from the public client address 131.107.100.50. The AzureLoadBalancer tag represents Azure Load Balancer infrastructure traffic, including load-balancer health probes. It is useful when a network security group must allow those probes to reach backend virtual machines, but it does not represent all connections that are forwarded through a load balancer.

For an inbound connection to a load-balanced application, the network security group evaluates the connection using the client's original source IP address. Therefore, the rule required to meet the goal must allow TCP port 443 from 131.107.100.50, either explicitly or through an applicable source range/service tag, and must take precedence over any matching deny rule. A rule at priority 150 for AzureLoadBalancer traffic can allow the required platform probe traffic but cannot authorize this client's HTTPS connection. Microsoft documents the AzureLoadBalancer service tag and the default network security group rule that allows it separately from Internet client access. See Azure service tags overview and Network security groups overview.

QUESTION NO: 28

You have an Azure virtual machine named VM1.

You use Azure Backup to create a backup of VM1 named Backup1.

After creating Backup1, you perform the following changes to VM1:

- Modify the size of VM1.
- Copy a file named Budget.xls to a folder named Data.
- Reset the password for the built-in administrator account. - Add a data disk to VM1.

An administrator uses the Replace existing option to restore VM1 from Backup1.

You need to ensure that all the changes to VM1 are restored.

Which change should you perform again?

A. Modify the size of VM1.

B. Reset the password for the built-in administrator account.

C. Add a data disk.

D. Copy Budget.xls to Data.

E. Perform all of the post-backup disk and operating-system changes again: copy Budget.xls to Data, reset the built-in administrator password, and add the data disk.

ANSWER: E

Explanation:

Perform all of the post-backup disk and operating-system changes again: copy Budget.xls to Data, reset the built-in administrator password, and add the data disk. is the only answer that ensures the requested state after a *Replace existing* restore. Azure Backup restores the VM disks from the selected recovery point. Because Backup1 was created

before the file was copied and before the administrator password was reset, the restored OS disk contains its earlier contents and account state. Those two changes must therefore be reapplied.

The restore also uses the disks represented by the selected recovery point. A data disk added after Backup1 is not part of that recovery point, so it must be added again to recreate the post-backup disk layout. In contrast, the VM size is a VM resource configuration setting rather than disk content; restoring disks with the replace-existing workflow does not require resizing the VM again to retain its current size.

Microsoft documents that the replace-existing restore workflow replaces the VM disks using the recovery point, while the VM resource configuration is retained. Review the restore behavior and prerequisites in Restore Azure VMs by using Azure Backup and the broader recovery-point behavior in Azure VM backup architecture.

QUESTION NO: 29

You have the Azure virtual machines shown in the following table.

Name	IP address	Virtual network
VM1	10.0.0.4	VNET1
VM2	10.0.0.5	VNET1

VNNET1 is linked to a private DNS zone named contoso.com that contains the records shown in the following table.

Name	Type	TTL	Value	Auto registered
comp1	TXT	3600	10.0.0.5	False
comp2	A	3600	10.0.0.5	False
comp3	CNAME	3600	comp1.contoso.com	False
comp4	PTR	3600	10.0.0.5	False

Which DNS names can you use to ping VM2?

A. comp1.contoso.com and comp2.contoso.com only

B. comp2.contoso.com and comp4.contoso.com only

C. comp2.contoso.cam only

D. comp1.contoso.com, comp2.contoso.com, and comp4.contoso.com only

E. comp1.contoso.com, comp2.contoso.com, comp3.contoso.com, and comp4.contoso.com

F. comp2.contoso.com only

ANSWER: F

Explanation:

comp2.contoso.com only is correct because it is the name represented by an A record in the linked Azure Private DNS zone, and that record maps the host name to VM2's private IP address. A ping command begins with a forward DNS query

when a host name is supplied. The DNS resolver must therefore be able to obtain an IP address for the specified name before it can send ICMP echo requests to VM2.

Azure Private DNS zones provide DNS name resolution for resources in linked virtual networks. Because VNNET1 is linked to the **contoso.com** private zone, workloads that use Azure-provided DNS in that virtual network can resolve records held in that zone. The A record for **comp2** supplies the required IPv4 address, so the fully qualified name **comp2.contoso.com** resolves to VM2 and can be used with ping.

In Azure DNS, A records associate a DNS name with an IPv4 address and are the standard record type for forward resolution of a virtual machine name to its private address. See Azure Private DNS overview and Azure DNS zones and records.

QUESTION NO: 30

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result. Establish if the solution satisfies the requirements.

Your company has an Azure Active Directory (Azure AD) subscription.

You want to implement an Azure AD conditional access policy.

The policy must be configured to require members of the Global Administrators group to use Multi-Factor Authentication and an Azure AD-joined device when they connect to Azure AD from untrusted locations.

Solution: You access the Azure portal to alter the session control of the Azure AD conditional access policy.

Does the solution meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because Conditional Access session controls cannot impose the required authentication and device-state conditions. Session controls govern how an already-authorized session is handled, such as sign-in frequency, persistent browser behavior, Conditional Access App Control, or application-enforced restrictions. They do not select the targeted administrators or require multifactor authentication and an Azure AD-joined device as conditions for granting access.

To implement the stated requirement, the policy must target the Global Administrators directory role or its members and use a location condition that applies to untrusted locations, normally by including all locations and excluding configured trusted named locations. The access decision must then be configured in the policy's Grant controls. Grant controls support requiring multifactor authentication and requiring an approved device state. For the Azure AD-joined-device requirement in the terminology used by older Azure AD Conditional Access experiences, this is configured through the domain-joined device grant requirement; current Microsoft Entra documentation describes the related device-based grant controls and device filters.

Because changing only session control leaves the required grant conditions unapplied, it does not meet the goal. See Microsoft Entra Conditional Access grant controls and Microsoft Entra Conditional Access session controls.

QUESTION NO: 31

You have an Azure subscription that contains a storage account named storage1. The storage1 account contains blob data.

You need to assign a role to a user named User1 to ensure that the user can access the blob data in storage1. The role assignment must support conditions.

Which two roles can you assign to User1? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Storage Blob Data Contributor

B. Storage Blob Data Owner

C. Storage Account Backup Contributor

D. Storage Account Contributor

E. Storage Blob Delegator

F. Owner

ANSWER: A B

Explanation:

Storage Blob Data Contributor and **Storage Blob Data Owner** are Azure Storage data-plane roles that can be assigned at the storage-account scope (or a narrower blob scope) to provide access to blob data. Storage Blob Data Contributor grants permissions to read, create, modify, and delete blob containers and blobs. This enables User1 to work directly with the blob data in storage1 using Microsoft Entra ID authorization. Microsoft documents its data-access permissions in the Storage Blob Data Contributor built-in role definition.

Storage Blob Data Owner likewise grants full access to blob containers and data, and additionally permits management of blob data ownership and POSIX access control lists where applicable. Its blob data actions make it suitable when User1 requires broader control over the data. Microsoft documents this role at Storage Blob Data Owner built-in role definition.

Both roles are supported for Azure role assignment conditions (Azure ABAC) for Blob Storage data actions. A condition can further constrain the allowed blob operations, such as limiting access based on blob index tags, container name, or request attributes, while retaining the permissions supplied by the assigned data role. Therefore, each role independently meets both the blob-data access requirement and the requirement for a condition-capable assignment.

QUESTION NO: 32

You have web app in the West US, Central US and East US Azure regions.

You have the App plans shown in the following table.

Name	Operating system	Location	SKU and size
ASP1	Windows	West US	Standard S1
ASP2	Linux	Central US	Premium V2 P1v2
ASP3	Linux	East US	Premium V2 P1v2
ASP4	Linux	East US	Premium V2 P1v2

You plan to create an additional App Service plan named ASPs that will use the Linux operating system.

You need to identify in which of the currently used locations you can deploy ASPs.

What should you recommend?

A. West US only

B. East US only

C. Central US only

D. West US, Central US, or East US

ANSWER: D

Explanation:

West US, Central US, or East US is correct. An Azure App Service plan defines the compute resources and configuration under which App Service apps run, and each plan is created in a specific Azure region. The operating system choice is a property of the new plan: creating a Linux-based App Service plan does not require an existing App Service plan in that region to use Linux, nor does an existing Windows-based plan prevent creation of a separate Linux plan.

Therefore, ASPs can be created as a new Linux App Service plan in any of the listed locations, provided the selected App Service SKU and required capabilities are available there. West US, Central US, and East US support App Service and can host a Linux App Service plan. The plans already shown in the table do not impose a cross-region or cross-operating-system restriction on this new plan. Apps that need to run on Linux should be placed into the newly created Linux plan rather than a Windows plan, because an App Service plan has a single operating-system type.

For details on the role and regional scope of an App Service plan, see Azure App Service plans overview and Configure an App Service app.

QUESTION NO: 33

You need to define a custom domain name for Azure AD to support the planned infrastructure.

Which domain name should you use?

A. Join the client computers in the Miami office to Azure A

B. Add http://autologon.microsoftazuread-sso.com to the intranet zone of each client computer in the Miami office.

C. Allow inbound TCP port 8080 to the domain controllers in the Miami office.

D. Install Azure AD Connect on a server in the Miami office and enable Pass-through Authentication

E. Install the Active Directory Federation Services (AD FS) role on a domain controller in the Miami office.

ANSWER: B D

Explanation:

The supplied answer set addresses Microsoft Entra seamless single sign-on rather than selecting a custom domain name. For that intended configuration, adding `http://autologon.microsoftazuread-sso.com` to the Local intranet zone is required so browsers can treat the seamless SSO endpoint as an intranet resource and automatically send the signed-in user's Kerberos ticket. Microsoft recommends deploying this setting through Group Policy to the applicable domain-joined client computers.

Installing Azure AD Connect on a server in the Miami office and enabling Pass-through Authentication is also appropriate when the design requires Microsoft Entra ID to validate user sign-ins against on-premises Active Directory. Azure AD Connect installs and registers the required Pass-through Authentication agent, and seamless SSO can be enabled with this authentication approach. The agent uses outbound connectivity to Microsoft Entra ID and validates credentials through the local Active Directory infrastructure, allowing users to access cloud resources with their existing organizational identities. Microsoft documents the required intranet-zone setting and the Azure AD Connect setup for this combination at Microsoft Entra seamless SSO quick start and Pass-through Authentication quick start.

QUESTION NO: 34

You have an Azure subscription named Subscription1 that contains the storage accounts shown in the following table:

Name	Account kind	Azure service that contains data
storage1	Storage	File
storage2	StorageV2 (general purpose v2)	File, Table
storage3	StorageV2 (general purpose v2)	Queue
storage4	BlobStorage	Blob

You plan to use the Azure Import/Export service to export data from Subscription1.

You need to identify which storage account can be used to export the data.

What should you identify?

A. storage1

B. storage2

C. storage3

D. storage4

ANSWER: D

Explanation:

storage4

is the storage account that can be used for the Azure Import/Export export job because it has the blob-storage capability required for export operations. Azure Import/Export is an offline data-transfer service: Microsoft copies the selected data from the specified Azure Storage account to encrypted drives, then ships those drives to the customer. For an export job, the source data must be Azure blobs in a supported storage account configuration. The export workflow is designed around blob data, including selecting blob containers and identifying the blobs that Azure should copy to the drives.

Before creating the job, an administrator should also confirm that the account and blobs meet the service requirements, such as using a supported storage account type and ensuring the data is accessible for the export job. The job is created in the same Azure region as the storage account, and the destination shipping address must be in a supported country or region. Microsoft documents the supported account and blob requirements for Import/Export export jobs at Azure Import/Export requirements and describes the export-job process at Export data from Azure Blob Storage.

QUESTION NO: 35

Your on-premises network contains a VPN gateway.

You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
vgw1	Virtual network gateway	Gateway for Site-to-Site VPN to the on-premises network
storage1	Storage account	Standard performance tier
Vnet1	Virtual network	Enabled forced tunneling
VM1	Virtual machine	Connected to Vnet1

You need to ensure that all the traffic from VM1 to storage! travels across the Microsoft backbone network.

What should you configure?

A. service endpoints

B. Azure Active Directory (Azure AD) Application Proxy

C. a network security group (NSG)

D. Azure Virtual WAN

ANSWER: A

Explanation:

Configure **service endpoints**, specifically the **Microsoft.Storage** service endpoint on the subnet that contains VM1. A virtual network service endpoint extends the subnet's identity to the Azure Storage service and provides a direct path from the virtual network to the storage account over the Microsoft backbone network. This is particularly important when the virtual

network uses forced tunneling through the on-premises VPN gateway, because the service endpoint supplies routes for the Azure service that keep eligible Storage traffic on Azure's network rather than sending it toward the default on-premises path.

After enabling the endpoint on VM1's subnet, the storage account can be configured to allow access from that selected virtual network and subnet. This restricts the storage account's network access to the intended private network path while retaining the service's public endpoint architecture; a service endpoint does not create a private IP address for the storage account. Microsoft documents that service endpoints optimize routing to Azure services and keep traffic on the Azure backbone. For implementation details, see Virtual Network service endpoints overview and Configure Azure Storage network security.

QUESTION NO: 36

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 these questions will not appear m the review screen.

You manage a virtual network named VNetl1 that is hosted in the West US Azure region.

VNetl1 hosts two virtual machines named VM1 and VM2 that run Windows Server.

You need to inspect all the network traffic from VM1 to VM2 for a period of three hours.

Solution: From Azure Network Watcher, you create a packet capture.

Does this meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes. Azure Network Watcher packet capture can capture network traffic to and from a virtual machine, enabling inspection of the packets exchanged between VM1 and VM2. A capture can be started on VM1, with filters configured for VM2's IP address so that the captured data is limited to the relevant traffic. Packet capture filtering supports criteria including local or remote IP addresses, ports, and protocol. The capture can be stored locally on the VM or in an Azure Storage account for later download and analysis in a tool such as Wireshark.

The requested three-hour duration is supported because Network Watcher packet capture allows a capture duration of up to 18,000 seconds (five hours). Therefore, configuring the duration as 10,800 seconds captures the required period. Packet capture requires the Network Watcher agent/extension on the target VM; Azure installs the extension when needed, subject to the VM and extension prerequisites. For traffic visibility from VM1 to VM2, creating the capture on VM1 meets the stated requirement. See the Azure Network Watcher packet capture overview and manage packet captures documentation.

QUESTION NO: 37

You have a Microsoft Entra tenant.

You plan to perform a bulk import of users by using the Azure portal.

You need to ensure that imported user objects are added automatically as the members of a specific group based on each user ' s department. The solution must minimize administrative effort.

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

NOTE: Each correct selection is worth one point.

A. Create an Azure Resource Manager (ARM) template.

B. Create groups that use the Dynamic User membership type.

C. Write a PowerShell script that parses an import file.

D. Create groups that use the Assigned membership type.

E. Create a CSV file that contains user information and the appropriate attributes.

F. Create an XML file that contains user information and the appropriate attributes.

ANSWER: B E

Explanation:

Create groups that use the Dynamic User membership type. is required because dynamic membership evaluates a rule against user attributes and automatically maintains group membership. A group rule can use the `department` attribute, such as `user.department -eq "Finance"`. After Microsoft Entra ID evaluates the rule, users whose department value matches are added to the group automatically. This avoids an administrator having to assign memberships individually after the import and continues to maintain membership if a user's department is changed later.

Create a CSV file that contains user information and the appropriate attributes. is also required because the Azure portal bulk-create process uses a downloaded and completed CSV template. The user records must include the relevant department values so that the dynamic membership rule has the attribute data needed to evaluate each imported account. Together, importing users with populated department attributes and using a department-based dynamic group rule provides the requested automated, low-administration solution. Dynamic membership requires an appropriate Microsoft Entra ID license for each user who benefits from the feature. See Dynamic membership rules for groups in Microsoft Entra ID and Bulk add users in the Microsoft Entra admin center.

QUESTION NO: 38

You have an Azure subscription that contains the virtual networks shown in the following table: - Vnet1–Vnet3: East US - Vnet4–Vnet6: UK South - Vnet7–Vnet10: East Asia All the virtual networks are peered. Each virtual network contains nine virtual machines. You need to configure secure RDP connections to the virtual machines by using Azure Bastion. What is the minimum number of Bastion hosts required?

A. 1

B. 3

C. 9

D. 10

ANSWER: B

Explanation:

3 is the minimum required because an Azure Bastion host using the Standard SKU can provide RDP connectivity to virtual machines in its own virtual network and in peered virtual networks within the same Azure region. Therefore, a single Bastion host can serve Vnet1 through Vnet3 in East US, another can serve Vnet4 through Vnet6 in UK South, and a third can serve Vnet7 through Vnet10 in East Asia.

Although all the virtual networks are peered, the networks are distributed across three separate Azure regions. Azure Bastion does not support using global virtual network peering to connect through a Bastion host to virtual machines in a different region. The Bastion deployment must consequently be regional for this scenario. The number of virtual machines in each virtual network does not increase the minimum host count; Bastion is deployed as a managed service into the dedicated AzureBastionSubnet and can serve multiple eligible VMs. Deploying one Standard SKU Bastion host in a suitable hub or selected virtual network in each region provides secure browser-based RDP access to all the listed machines without exposing public IP addresses on those VMs. See Azure Bastion and VNet peering and the Azure Bastion overview.

QUESTION NO: 39

You have an Azure subscription that has the public IP addresses shown in the following table.

Name	IP version	SKU	Tier	IP address assignment
IP1	IPv4	Standard	Regional	Static
IP2	IPv4	Standard	Global	Static
IP3	IPv4	Basic	Regional	Dynamic
IP4	IPv4	Basic	Regional	Static
IP5	IPv6	Basic	Regional	Dynamic

You plan to deploy an Azure Bastion Basic SKU host named Bastion1.

Which IP addresses can you use for Bastion1?

A. IP1 only

B. IP1 and IP2 only

C. IP3, IP4, and IPS only

D. IP1, IP2, IP4, and IP5 only

E. IP1, IP2, IP3, IP4, and IPS

ANSWER: A

Explanation:

IP1 only is correct because an Azure Bastion host requires a public IP address that is compatible with the Bastion service. For a Bastion host using the Basic SKU, the public IP address must use the Standard SKU, be configured with a Static allocation method, and use an IPv4 address. IP1 is the listed address that satisfies these required characteristics and can therefore be associated with Bastion1 when the host is deployed.

The public IP address is assigned to the Bastion resource itself, allowing users to connect securely to virtual machines through the Azure portal or supported native-client features without exposing those virtual machines through individual public IP addresses. The address must be available for association at deployment time and must be in the same Azure region as the Bastion resource and its virtual network. Microsoft documents the Basic SKU Bastion deployment requirements, including the need for a Standard, static public IP address, in the Azure Bastion deployment quickstart. For the SKU and allocation characteristics of Azure public IP resources, see Public IP addresses in Azure.

QUESTION NO: 40 - (HOTSPOT)

HOTSPOT

-

You have an Azure subscription that contains the vaults shown in the following table.

Name	Type
Recovery1	Recovery Services vault
Backup1	Azure Backup vault

You deploy the virtual machines shown in the following table.

Name	Operating system	Security Configuration
VM1	Windows Server	Azure Disk Encryption
VM2	Linux	Trusted launch

You have the backup policies shown in the following table.

Name	Type	In vault
Policy1	Standard	Recovery1
Policy2	Enhanced	Recovery2
Policy3	*Not applicable*	Backup1

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

Answer Area

Statements	Yes	No
VM1 can be backed up by using Policy1.	○	○
VM2 can be backed up by using Policy3.	○	○
VM2 can be backed up by using Policy2.	○	○

ANSWER:

Answer Area

Statements	Yes	No
VM1 can be backed up by using Policy1.	○	○
VM2 can be backed up by using Policy3.	○	○
VM2 can be backed up by using Policy2.	○	○

Explanation:

The correct selections are Yes for backing up VM1 by using Policy1, and No for backing up VM2 by using Policy3 and Policy2. Azure IaaS virtual machine backup is configured in a Recovery Services vault. A backup policy used for this workload must be a policy in the applicable Recovery Services vault, and the vault must be available for the VM protection scenario. Policy1 meets these requirements because it is a Standard policy in Recovery1, which is a Recovery Services vault. Standard Azure VM backup supports virtual machines that use Azure Disk Encryption, subject to the normal encryption-key and access prerequisites. Therefore, VM1 can be protected by Policy1.

An Azure Backup vault is intended for supported newer Azure Backup data-source scenarios, such as Azure Disk Backup, rather than Azure VM backup configured as an IaaS VM workload. Consequently, Policy3 in Backup1 cannot be selected to back up VM2 as an Azure virtual machine. In addition, a backup policy cannot be used from a vault that is not present and available in the subscription for the protection configuration. Policy2 is identified as belonging to Recovery2, while Recovery2 is not included in the listed vaults. Thus, it cannot be used to protect VM2 in this environment. Microsoft describes the distinct supported workloads for vault types in the Azure Backup vault overview. The Azure VM protection process and its use of a Recovery Services vault are documented in Back up Azure VMs.

QUESTION NO: 41

You have an Azure Active Directory (Azure AD) tenant that has Azure AD Privileged Identity Management configured.

You have 10 users who are assigned the Security Administrator role for the tenant.

You need the users to verify whether they still require the Security Administrator role.

What should you do?

A. From Azure AD Identity Protection, configure a user risk policy.

B. From Azure AD Privileged Identity Management, create an access review.

C. From Azure AD Identity Protection, configure the Weekly Digest.

D. From Azure AD Privileged Identity Management, create a conditional access policy.

ANSWER: B

Explanation:

From Azure AD Privileged Identity Management, create an access review. Access reviews are an identity governance capability used to periodically validate whether users still need their existing access, including assignments to Microsoft Entra ID privileged roles such as Security Administrator. When configuring the review, you can select the role assignment scope, set the assigned users to review their own continued need for the role, and define a recurring review schedule if ongoing attestation is required.

The review process asks each user to approve or deny their continued assignment. A reviewer decision is recorded for audit purposes, and the review can be configured to automatically apply results when it completes. This enables removal of role assignments for users who no longer require privileged access, supporting least-privilege administration and regular access recertification. Privileged Identity Management access reviews are specifically designed for reviewing both active and eligible role assignments, making this the appropriate administrative action for the stated requirement.

For implementation details, see Microsoft Entra access reviews overview and Create an access review of Microsoft Entra roles in PIM.

QUESTION NO: 42

You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Location
VNET1	Virtual network	East US
IP1	Public IP address	West Europe
RT1	Route table	North Europe

You need to create a network interface named NIC1. In which location can you create NIC1?

A. East US and North Europe only

B. East US only

C. East US, West Europe, and North Europe

D. East US and West Europe only

ANSWER: B

Explanation:

East US only is correct because an Azure network interface must be created in the same Azure region as the virtual network and subnet to which it will be connected. A network interface is not an independently deployable cross-region resource: its IP configuration is associated with a subnet, and that subnet is part of a virtual network in a specific region. Therefore, the available location for NIC1 is determined by the location of the virtual network shown in the resource table, rather than by the locations of unrelated resources such as virtual machines or storage accounts.

The table shows that the relevant virtual network is located in East US. NIC1 can consequently be created in East US and attached to a subnet in that virtual network. This regional co-location requirement also ensures that the NIC can be used by compatible compute resources in the same Azure region. Microsoft documents that a network interface is connected to a virtual network subnet and that Azure resources involved in that configuration must satisfy the applicable regional requirements. See Virtual network interfaces in Azure and Manage network interfaces.

QUESTION NO: 43

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 need to ensure that an Azure Active Directory (Azure AD) user named Admin1 is assigned the required role to enable Traffic Analytics for an Azure subscription.

Solution: You assign the Traffic Manager Contributor role at the subscription level to Admin1.

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. Traffic Analytics is a Network Watcher capability that analyzes network-security-group flow logs in a Log Analytics workspace. Enabling it requires Azure permissions for the relevant Network Watcher and network resources, including the ability to configure flow logs and associate the required Log Analytics workspace. The Traffic Manager Contributor role has a narrowly defined purpose: it permits management of Azure Traffic Manager profiles and endpoints. Assigning that role at subscription scope expands the scope at which those Traffic Manager permissions apply, but it does not add the Network Watcher, network-security-group flow-log, or Log Analytics permissions needed for Traffic Analytics.

Admin1 should instead receive a role assignment that includes the required actions at the appropriate scopes. Depending on the deployment and existing assignments, this can include Contributor or Network Contributor for the subscription or relevant network resources, plus sufficient access to the Log Analytics workspace used to store and analyze flow-log data. Azure RBAC role assignments should follow least privilege while still covering all resources involved in configuring the feature. Microsoft documents Traffic Analytics as a Network Watcher flow-log analysis feature and documents the permissions included in the Traffic Manager Contributor built-in role. See Traffic Analytics documentation and Traffic Manager Contributor role documentation.

QUESTION NO: 44

You have an Azure subscription that contains two virtual machines named VM1 and VM2

You create an Azure load balancer.

You plan to create a load balancing rule that will load balance HTTPS traffic between VM1 and VM2.

Which two additional load balance resources should you create before you can create the load balancing rule? Each correct answer presents part of the solution

MOTL Each correct selection 5 worth one point.

A. a frontend IP address

B. a backend pool

C. a health probe

D. an inbound NAT rule

E. a virtual network

ANSWER: B C

Explanation:

A backend pool is required because it defines the set of targets to which the Azure Load Balancer distributes connections. For this scenario, the backend pool must include VM1 and VM2, typically by associating each virtual machine network interface or IP configuration with the pool. The HTTPS load-balancing rule then uses that pool as its backend destination.

A health probe is also required for a load-balancing rule. Azure Load Balancer uses the probe to determine the availability of each backend instance. For HTTPS traffic, the probe can be configured with an appropriate protocol, port, and optionally a request path. Only virtual machines that respond successfully to the configured probe are considered healthy and can receive new traffic through the rule. This health-based distribution helps maintain availability when one VM or its hosted application is unavailable.

The load-balancing rule connects an existing frontend configuration to the selected backend pool and health probe. A frontend IP configuration is part of configuring the load balancer itself, whereas the backend pool and health probe are the additional resources that must be prepared for this rule configuration. See Azure Load Balancer components and Manage Azure Load Balancer rules.

QUESTION NO: 45

You have an Azure subscription named Subscription1 that has the following providers registered:

- Authorization
- Automation
- Resources
- Compute

➪ KeyVault

➪ Network

➪ Storage

➪ Billing

➪ Web

Subscription1 contains an Azure virtual machine named VM1 that has the following configurations:

➪ Private IP address: 10.0.0.4 (dynamic)

➪ Network security group (NSG): NSG1

➪ Public IP address: None

➪ Availability set: AVSet

➪ Subnet: 10.0.0.0/24

➪ Managed disks: No

➪ Location: East US

You need to record all the successful and failed connection attempts to VM1.

Which three actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Enable Azure Network Watcher in the East US Azure region.

B. Add an Azure Network Watcher connection monitor.

C. Register the MicrosoftLogAnalytics provider.

D. Create an Azure Storage account.

E. Register the Microsoft.Insights resource provider.

F. Enable Azure Network Watcher flow logs.

ANSWER: A D F

Explanation:

To record connection attempts that are permitted and denied for VM1, enable NSG flow logs for NSG1. NSG flow logs capture traffic evaluated by an NSG and include the flow's source and destination addresses, ports, protocol, direction, and flow state. In particular, the log data identifies whether the NSG allowed or denied the flow, which provides the successful and failed connection-attempt information required here. Because NSG1 is associated with VM1, logging that NSG covers traffic processed by its rules.

Flow logs are configured through Azure Network Watcher, which must be available in the same region as the monitored NSG. Therefore, enabling Azure Network Watcher in East US establishes the regional Network Watcher capability needed to configure NSG flow logging for NSG1.

An Azure Storage account is also required because NSG flow-log records are written to a storage account. The account provides the durable destination for the generated log files, enabling the organization to retain and later inspect allowed and denied flows. The subscription already has the Microsoft.Network and Microsoft.Storage resource providers registered, which support these resources. See NSG flow logs overview and Manage NSG flow logs in the Azure portal.

QUESTION NO: 46

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result. Establish if the solution satisfies the requirements.

Your company has an azure subscription that includes a storage account, a resource group, a blob container and a file share.

A colleague named Jon Ross makes use of a solitary Azure Resource Manager (ARM) template to deploy a virtual machine and an additional Azure Storage account.

You want to review the ARM template that was used by Jon Ross.

Solution: You access the Container blade.

Does the solution meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because the Container blade is a management interface for an individual Azure Blob Storage container. It provides access to container-specific operations and settings, such as viewing blobs and managing container properties, but it does not expose the Azure Resource Manager deployment record or the ARM template used to create resources. The relevant deployment details are retained in the deployment history of the resource group targeted by the ARM deployment. To review the template, an administrator should open the appropriate resource group in the Azure portal, select **Deployments**, select the completed deployment, and view its template and parameters. Azure also supports exporting a template from a resource group when an administrator needs to inspect the resource configuration in template form. This approach is appropriate for reviewing the infrastructure definition associated with Jon Ross's virtual machine and storage-account deployment. For portal steps and template-export behavior, see Export Azure Resource Manager templates from the Azure portal and View deployment history with Azure Resource Manager.

QUESTION NO: 47 - (SIMULATION)

You have the App Service plan shown in the following exhibit.

The scale-in settings for the App Service plan are configured as shown in the following exhibit.

The scale out rule is configured with the same duration and cool down tile as the scale in rule.

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

ANSWER: See the explanation for the answer

Explanation:

Cannot determine the drop-down selections from the supplied JSON.

The question refers to two exhibits (the App Service plan configuration and the scale-in rule) and to statement text/drop-down options, but none of those images or statements are present. The only available detail is that the scale-out rule uses the same duration and cooldown as the scale-in rule, which is insufficient to calculate or select an answer.

Please provide the missing exhibits and the statements/drop-down choices. In particular, the following values are required: minimum/default/maximum instance count, metric thresholds, aggregation/duration, scale increment, and cooldown period.

QUESTION NO: 48

You have an app named App1 that runs on two Azure virtual machines named VM1 and VM2.

You plan to implement an Azure Availability Set for Appl. The solution must ensure that App1 is available during planned maintenance of the hardware hosting VM1 and VM2.

What should you include in the Availability Set?

A. one update domain

B. two update domains

C. one fault domain

D. two fault domains

ANSWER: B

Explanation:

Two update domains should be included in the availability set. An update domain is a logical group of virtual machines that Azure can reboot together when it performs planned maintenance, such as host operating-system, platform, or hardware updates. Azure processes planned maintenance one update domain at a time and waits for the affected resources to become available before proceeding to the next domain. With VM1 and VM2 assigned to separate update domains, planned maintenance that restarts the host for one virtual machine does not restart the other virtual machine at the same time. App1 can therefore remain available from the VM that is not undergoing maintenance. Availability sets distribute their member VMs across update domains automatically; configuring at least two update domains provides the required separation for the two application instances. This design applies when the application is deployed redundantly across both VMs and can continue serving requests when one instance is temporarily unavailable. Microsoft documents update domains as the availability-set feature intended specifically to reduce downtime from planned maintenance events. For implementation details, see Availability sets overview and Maintenance and updates for Azure virtual machines.

QUESTION NO: 49

You have an Azure web app named webapp1 that is hosted in an App Service plan that supports custom domains.

You need to make webapp1 available by using https://www.contoso.com. The solution must ensure that clients use HTTPS.

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

NOTE: Each correct selection is worth one point.

A. Add www.contoso.com as a custom domain for webapp1

B. Create a TLS/SSL binding for www.contoso.com

C. Create an Azure Load Balancer rule for port 443

D. Enable a service endpoint on the App Service plan subnet

E. Create an inbound NAT rule for webapp1

ANSWER: A B

Explanation:

You must add **www.contoso.com** as a custom domain for webapp1 after validating domain ownership. Validation normally requires a DNS record, such as a CNAME record for the www subdomain and, where applicable, a TXT record containing the App Service domain-verification ID.

You must also configure a TLS/SSL binding for the custom hostname. The binding associates a certificate with www.contoso.com and enables the web app to accept HTTPS connections for that hostname. The certificate can be an App Service managed certificate where supported, or a certificate obtained from another certificate authority.

For more information, see Map an existing custom DNS name to Azure App Service and Configure a custom domain name in Azure App Service.

QUESTION NO: 50

Name	Function
webapp1-prod	Production
webapp1-test	Staging

You have an Azure web app named App1. App1 has the deployment slots shown in the following table:

In webapp1-test, you test several changes to App1. You back up App1.

You swap webapp1-test for webapp1-prod and discover that App1 is experiencing performance issues. You need to revert to the previous version of App1 as quickly as possible.

What should you do?

A. Redeploy App1

B. Swap the slots

C. Clone App1

D. Restore the backup of App1

ANSWER: B

Explanation:

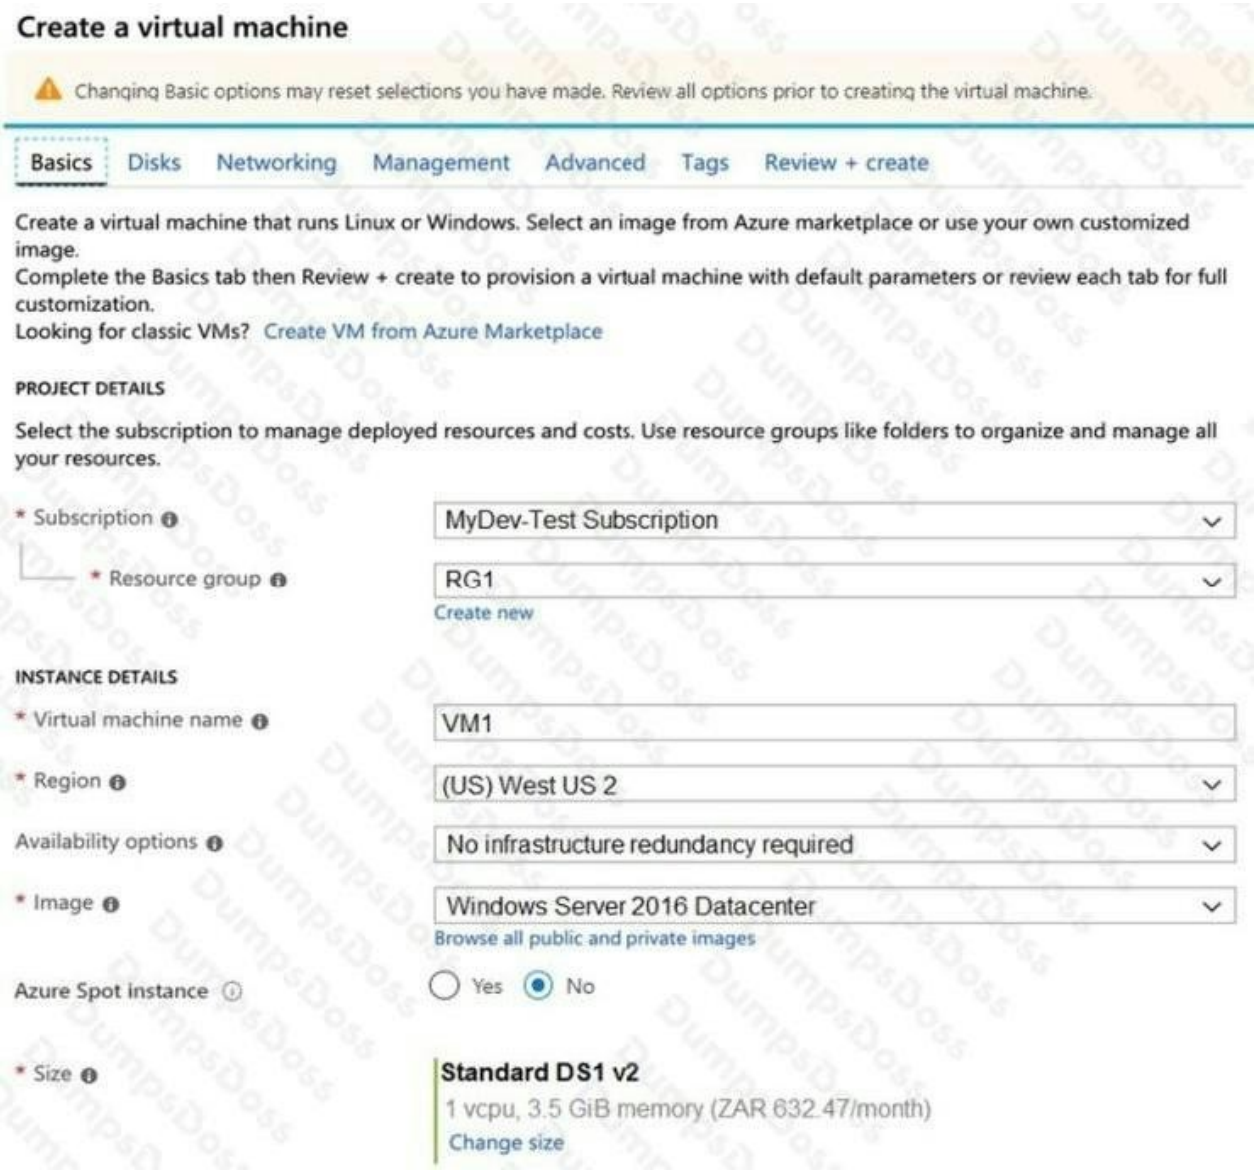
Swap the slots is the quickest way to return App1 to its prior production state. Before the initial swap, the production slot contained the known-good production application version, while the test slot contained the newly tested changes. A deployment-slot swap exchanges the slot content and applicable configuration between the source and target slots. Therefore, performing the swap again places the former production version back into the production slot and directs production traffic to it without waiting for a new deployment or a backup restoration operation.

This rollback approach is a key operational benefit of Azure App Service deployment slots: teams can validate a release in a nonproduction slot, promote it through a swap, and rapidly reverse the promotion if an issue is discovered. App Service also warms up the target slot during a swap, helping minimize disruption when returning to the earlier version. The existing backup can remain available for disaster recovery purposes, but it is not needed for this immediate rollback scenario.

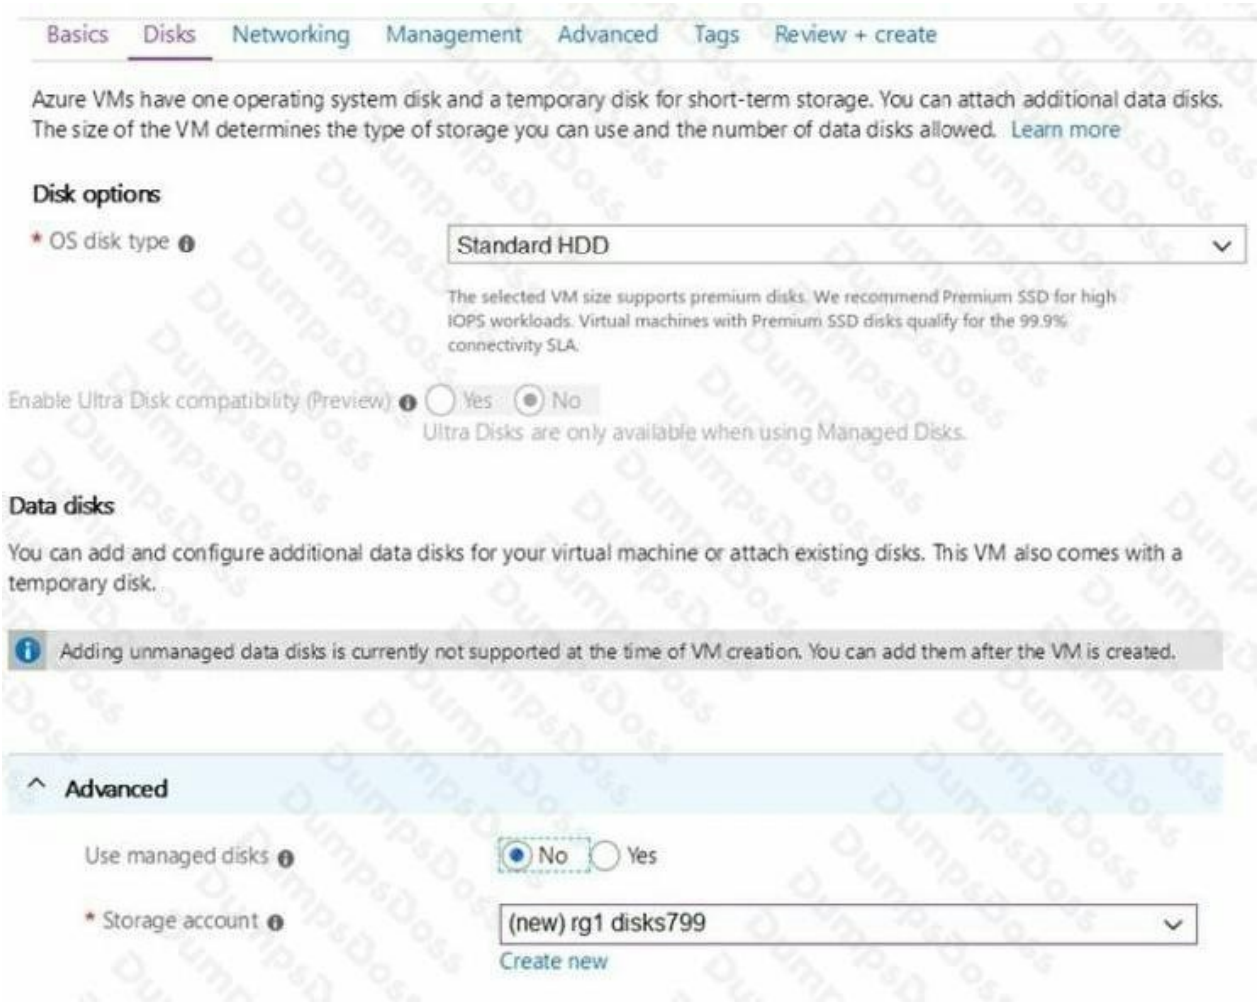
Microsoft describes swapping slots and reversing a swap in its Set up staging environments in Azure App Service documentation. See also Swap deployment slots.

QUESTION NO: 51

You plan to create an Azure virtual machine named VM1 that will be configured as shown in the following exhibit.



The planned disk configurations for VM1 are shown in the following exhibit.



You need to ensure that VM1 can be created in an Availability Zone.

Which two settings should you modify? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Use managed disks

B. OS disk type

C. Availability options

D. Size

E. Image

ANSWER: A C

Explanation:

To create VM1 in an Availability Zone, modify **Availability options** and select an Availability Zone during deployment. This deployment setting determines the physical zone to which the virtual machine is assigned. Azure Availability Zones are separate datacenter groups within a supported region, each with independent power, cooling, and networking. Selecting a zone makes the VM a zonal resource and is the setting that explicitly establishes its zone placement.

You must also enable **Use managed disks**. Azure virtual machines deployed in Availability Zones require managed disks because Azure manages disk placement and aligns zonal disks with the selected VM zone. Managed disks abstract the underlying storage account management and provide the disk resource model used for zonal virtual-machine deployments. When a zonal VM is created, its managed OS and data disks are placed in the same zone as the VM to provide the required locality.

These two settings together establish a zonal VM deployment: the availability configuration selects the zone, and managed disks provide zone-aligned storage for the VM. For more information, see Availability zones for Azure virtual machines and Introduction to Azure managed disks.

QUESTION NO: 52

You are configuring Azure Active Directory (Azure AD) authentication for an Azure Storage account named storage1.

You need to ensure that the members of a group named Group1 can upload files by using the Azure portal. The solution must use the principle of least privilege.

Which two roles should you configure for storage!? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point

A. Reader

B. Storage Blob Data Contributor

C. Storage Account Contributor

D. Storage Blob Data Reader

E. Contributor

ANSWER: A B

Explanation:

Reader and **Storage Blob Data Contributor** provide the required combination of Azure Resource Manager and blob data-plane permissions while maintaining least privilege. Reader grants management-plane visibility to the storage account. This enables Group1 members to locate and open storage1 in the Azure portal and navigate its resources without granting permissions to modify the storage account's configuration.

Storage Blob Data Contributor grants the required Azure RBAC data actions for blobs, including creating and writing blobs. Therefore, members of Group1 can upload files through the portal while Azure AD authenticates and authorizes their data access. This role should be assigned at the narrowest scope that meets the requirement, such as the target container when uploads are limited to a particular container; otherwise, it can be assigned at the storage-account scope.

Azure Storage distinguishes management-plane access to the storage resource from data-plane access to the blobs themselves. For portal-based blob operations using Azure AD credentials, Microsoft documents that users need both a role permitting portal navigation, such as Reader, and an appropriate Azure Storage blob data role. See Assign an Azure role for access to blob data and Azure built-in roles for Storage.

QUESTION NO: 53

You have an Azure subscription that contains a storage account named account1.

You plan to upload the disk files of a virtual machine to account1 from your on-premises network. The on-premises network uses a public IP address space of 131.107.1.0/24.

You plan to use the disk files to provision an Azure virtual machine named VM1. VM1 will be attached to a virtual network named VNet1. VNet1 uses an IP address space of 192.168.0.0/24.

You need to configure account1 to meet the following requirements:

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

NOTE: Each correct selection is worth one point.

A. From the Service endpoints blade of VNet1, add a service endpoint

B. From the Networking blade of account1, add the 131.107.1.0/24 IP address range.

C. From the Networking blade of account1, add VNet1.

D. From the Networking blade of account1, select Selected networks.

E. From the Networking blade of account1, select Allow trusted Microsoft services to access this storage account

ANSWER: B D

Explanation:

To permit upload of the virtual machine disk files from the on-premises environment while restricting public access to account1, configure the storage account firewall to use **Selected networks**. This changes the public network access rule from permitting all networks to permitting only explicitly configured network sources. The storage account can then continue to expose its public endpoint for the upload operation, but only authorized traffic will be accepted.

After selecting this firewall mode, add the public source range **131.107.1.0/24** on the Networking blade of account1. Azure Storage IP network rules use the public IP address observed by the storage service, so the on-premises public address range is the relevant range to allow. This rule permits upload tools and clients originating from that defined on-premises egress range to reach the storage account.

These settings provide controlled public-endpoint access for the upload workflow and do not depend on the private address space assigned to the future virtual network. Microsoft documents that storage firewalls can be limited to selected networks and that specific public IPv4 address ranges can be allowed through IP network rules. See Configure Azure Storage firewalls and virtual networks and Grant access from an internet IP range.

QUESTION NO: 54

You have a general purpose v1 storage account named storageaccount1 that has a private container named container1. You need to allow read access to the data inside container1, but only within a 14 day window. How do you accomplish this?

A. Create a stored access policy

B. Create a service SAS

C. Create a shared access signatures

D. Upgrade the storage account to general purpose v2

ANSWER: A B C

Explanation:

A Shared Access Signature (SAS) is the Azure Storage mechanism for delegating limited, time-bound access to data in a private blob container without changing the container's public access level. A service SAS can be scoped to the Blob service and to *container1*, assigned only the read permission, and configured with an expiration time 14 days from its creation. Anyone using the resulting SAS URI can read the permitted blobs only while the SAS is valid.

Creating a stored access policy on the container is also appropriate when the SAS should be governed centrally. The policy defines the read permission and the access interval, and a service SAS can reference that policy instead of embedding those settings directly in the URI. This approach supports changing the policy or revoking associated SAS access before the planned expiration when necessary. “Create a shared access signatures” describes the same overarching access-delegation mechanism; the service SAS is the specific SAS type used for blob-container data. General-purpose v1 storage accounts support service SAS and stored access policies, so an account upgrade is not required for this scenario.

See Grant limited access to Azure Storage resources using shared access signatures and Define a stored access policy.

QUESTION NO: 55

You deploy Azure virtual machines to three Azure regions

Each region contains a virtual network. Each virtual network contains multiple subnets peered in a full mesh topology.

Each subnet contains a network security group (NSG) that has defined rules.

A user reports that he cannot use port 33000 to connect from a virtual machine in one region to a virtual machine in another region.

Which two options can you use to diagnose the issue? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. Azure Virtual Network Manager

B. IP flow verify

C. Azure Monitor Network Insights

D. Connection troubleshoot

E. Effective security rules

ANSWER: B D

Explanation:

IP flow verify and **Connection troubleshoot** are appropriate Azure Network Watcher diagnostic tools for investigating a failed connection to TCP or UDP port 33000 between virtual machines in peered virtual networks. IP flow verify evaluates a specified traffic flow for a selected VM by using the source IP address, destination IP address, protocol, direction, and port. It reports whether Azure allows or denies that flow. When a network security group causes a denial, the result identifies the applied rule, which is particularly useful when multiple subnet- and NIC-level NSGs and their effective rules must be considered.

Connection troubleshoot performs an end-to-end connectivity test from a source VM to a destination VM and port. Its results identify the reachability status and provide diagnostic information for the path, including network security group evaluations, route issues, and other relevant connection checks. This makes it suitable for testing the actual cross-region workload path rather than inspecting a single policy in isolation. Together, these tools provide both a focused NSG flow decision and a broader connection-path diagnosis. For details, see IP flow verify in Azure Network Watcher and Connection troubleshoot overview.

QUESTION NO: 56

Your VMware vSphere on-premises infrastructure hosts 600 virtual machines (VMs).

Your company is planning to move all of these VMs to Azure. You are asked to provide information about

the resources that will be needed in Azure to host all of the VMs.

All VMs hosted in your on-premise infrastructure are based on Windows Server 2012 R2 or newer and RedHat Enterprise Linux 7.0 or newer.

You conduct the initial migration assessment and get a message that some virtual machines are conditionally ready for Azure.

You need to find the cause of this message.

What are two reasons why are you might get this message on some VMs? (Choose two)

Each correct answer presents part of the solution.

A. The vCenter user does not have enough permissions on affected VMs.

B. The operating system is configured as Windows Server 2003 in vCenter Server.

C. The operating system is configured as Others in vCenter Server.

D. The VMs are configured with the BIOS boot type.

E. The VMs are configured with the UEFI boot type.

ANSWER: B E

Explanation:

The operating system is configured as Windows Server 2003 in vCenter Server is a valid cause because Azure Migrate uses vCenter inventory metadata as part of discovery and assessment. When a VM is identified in vCenter as Windows Server 2003, the assessment evaluates it against the support and compatibility requirements associated with that reported guest OS. This can produce a conditional readiness result even when the actual installed operating system is newer. Correct guest OS metadata in vCenter is therefore important before relying on assessment readiness results and sizing recommendations.

The VMs are configured with the UEFI boot type is also a valid cause. Azure Migrate can assess UEFI-based VMware VMs, but migration readiness can be conditional because the Azure target must support the corresponding Generation 2 configuration and the guest OS must meet the applicable UEFI and Generation 2 requirements. The conditional status identifies that this firmware configuration needs validation or a compatible target selection before migration. Azure Migrate readiness is intended to surface such required conditions early, so they can be addressed during migration planning rather than after replication begins.

For readiness definitions and assessment behavior, see Azure Migrate assessment calculations. For VMware migration platform, operating-system, and boot requirements, see the VMware migration support matrix.

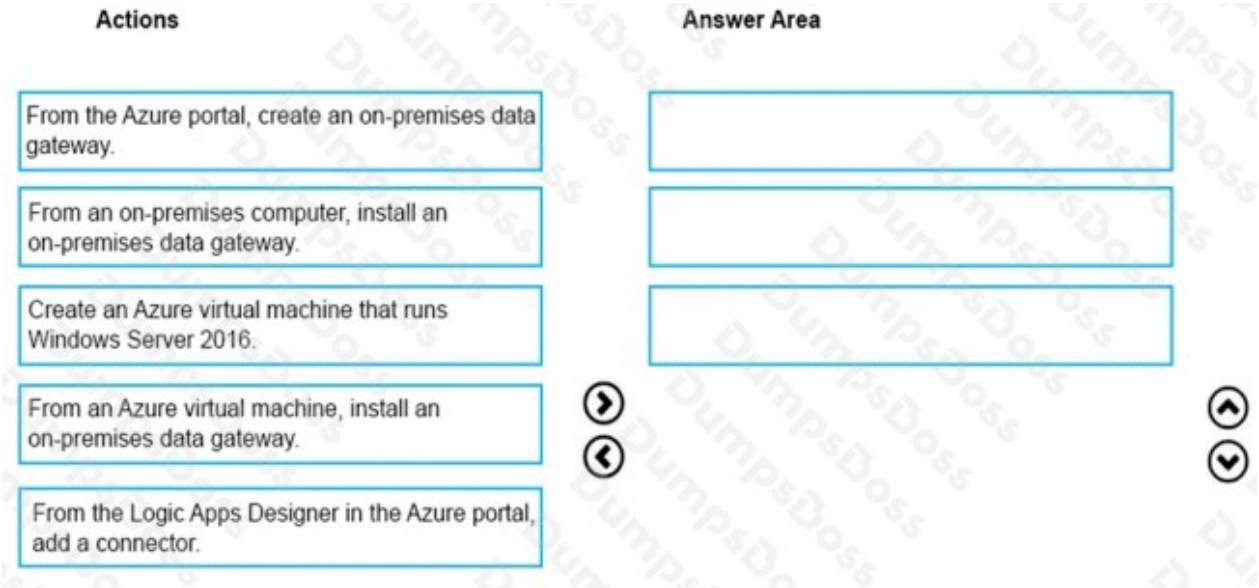
QUESTION NO: 57 - (DRAG DROP)

You have an on-premises network that includes a Microsoft SQL Server instance named SQL1.

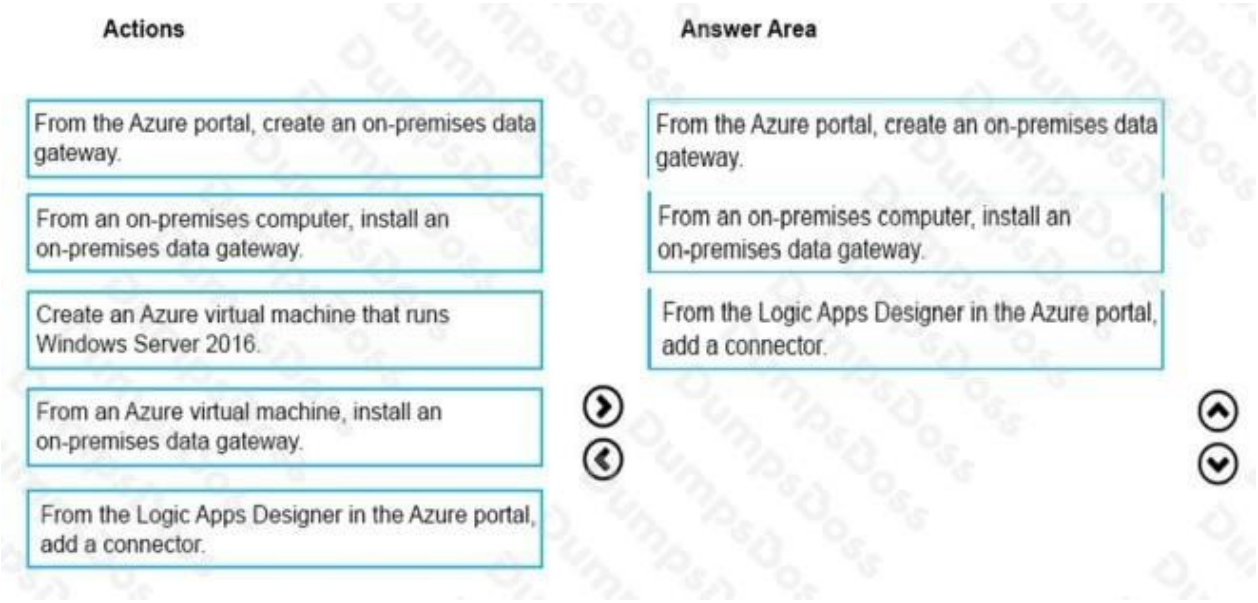
You create an Azure Logic App named App1.

You need to ensure that App1 can query a database on SQL1.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



ANSWER:



Explanation:

Azure Logic Apps access an on-premises SQL Server through an on-premises data gateway. The gateway provides the secure communication path between Azure services and resources that remain inside the local network. The correct sequence begins by creating the on-premises data gateway resource from the Azure portal. This establishes the gateway resource in Azure so that it can be selected and used by connections created for the Logic App.

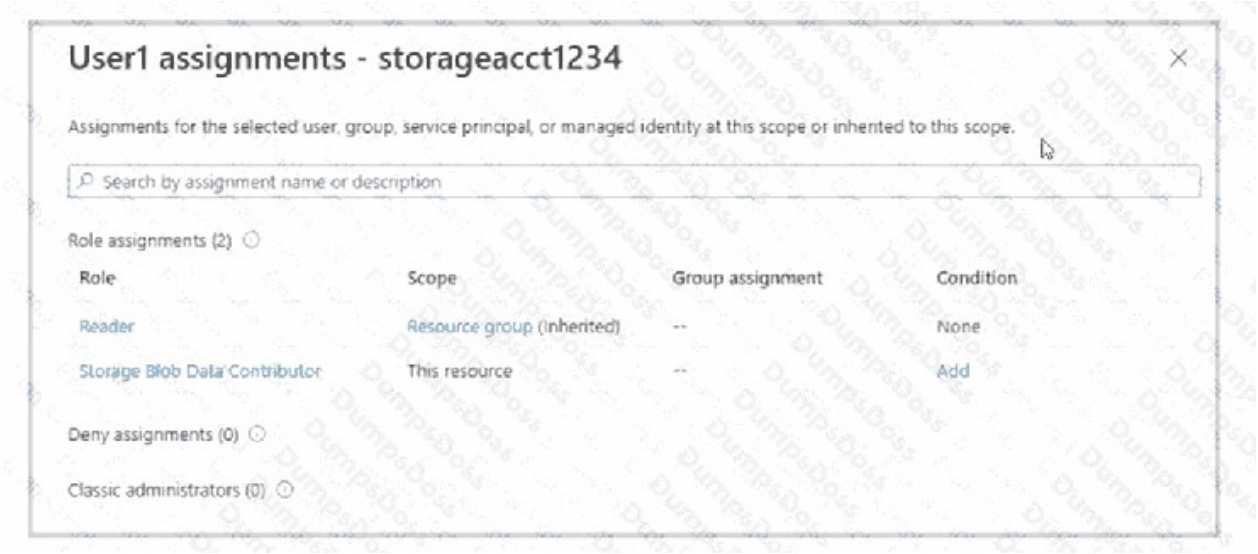
Next, the gateway software is installed on an on-premises computer. That computer must have network connectivity to SQL1, because it is the component that receives requests from the Logic App connection and forwards them to the local SQL Server. During setup, the gateway is registered and configured with its Azure identity, allowing Azure to recognize it as an available gateway for the subscription and region. Installing it on an on-premises computer also ensures the traffic can reach the database without exposing SQL1 directly to the internet.

After the gateway is available and online, a connector is added from Logic Apps Designer. The SQL Server connector connection can then be configured to use the on-premises data gateway and the SQL Server connection details. Once that connection is created, actions in the workflow can query the database on SQL1 through the gateway. This sequence follows the standard Azure integration pattern: provision the Azure gateway resource, install and register its local gateway service, and then create the Logic Apps connection that consumes it. Microsoft documents this workflow in Connect to on-premises data sources from Azure Logic Apps and in the on-premises data gateway documentation.

QUESTION NO: 58

You have an Azure subscription that contains a storage account named storageacct1234 and two users named User1 and User2.

You assign User1 the roles shown in the following exhibit.



Which two actions can User1 perform? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. View file shares in storageacct1234.

B. Upload blob data to storageacct1234.

C. Assign roles to User2 for storageacctl234.

D. View blob data in storageacctl234.

E. Modify the firewall of storageacct1234.

ANSWER: A B D

Explanation:

User1 can upload blob data to storageacct1234 and view blob data in storageacctl234 because the Storage Blob Data Contributor role grants data-plane permissions to read, write, and delete blob containers and blob data. These permissions are separate from Azure Resource Manager management-plane permissions and permit the user to work with the contents of blobs at the scope where the role is assigned. Uploading is a blob write operation, while viewing blob content is a blob read operation; both are explicitly included in this built-in role.

User1 can also view file shares in storageacct1234 because the Reader role provides management-plane read access to resource configuration and child resources. A file share can therefore be viewed as an Azure Storage resource, although Reader does not grant permissions to read the files and directories stored within that share. The distinction is important: viewing the share resource is a management operation, whereas accessing file-share data requires an appropriate Azure Storage Files data role.

Microsoft documents the permissions of Storage Blob Data Contributor, including blob read and write data actions. The Reader role documentation confirms that it permits viewing Azure resources without changing them.

QUESTION NO: 59

You need to add VM1 and VM2 to the backend poo! of LB1. What should you do first?

A. Create a new NSG and associate the NSG to VNET1/Subnet1.

B. Connect VM2 to VNET1/Subnet1.

C. Redeploy VM1 and VM2 to the same availability zone.

D. Redeploy VM1 and VM2 to the same availability set.

ANSWER: D

Explanation:

Redeploy VM1 and VM2 to the same availability set. LB1 is identified as a Basic internal Azure Load Balancer, and Basic Load Balancer backend-pool membership has a legacy availability-set requirement for standalone virtual machines. Before their network interfaces can be used together in the backend pool, the virtual machines must be placed in the same availability set. An availability set is an Azure logical grouping that distributes VMs across fault domains and update domains, supporting availability during planned maintenance and underlying infrastructure failures.

For this scenario, the relevant prerequisite is not that both VMs use the frontend subnet. Azure Load Balancer backend instances can be reached through their NICs within the virtual network; the frontend IP configuration's subnet does not require every backend VM NIC to be in that same subnet. Because the load balancer is the Basic SKU described in the question, meeting its availability-set placement requirement is the first action needed before adding both VMs as backend members. This is a legacy-SKU design consideration; Microsoft recommends Standard Load Balancer for current deployments and has announced retirement of the Basic SKU. See Azure Load Balancer overview and Availability sets overview.

QUESTION NO: 60

You have an Azure subscription that contains two virtual machines named VM1 and VM2. You create an Azure load balancer.

You plan to create a load balancing rule that will load balance HTTPS traffic between VM1 and VM2.

Which two additional load balancer resources should you create before you can create the load balancing rule? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a frontend IP address

B. an inbound NAT rule

C. a virtual network

D. a backend pool

E. a health probe

ANSWER: D E

Explanation:

A **backend pool** is required to define the resources that receive traffic distributed by the load-balancing rule. For VM1 and VM2, the backend pool contains the virtual machines' network interfaces or IP-based backend addresses. The HTTPS load-balancing rule associates its backend configuration with this pool, enabling Azure Load Balancer to select a healthy target for each new flow.

A **health probe** provides the availability signal used by Azure Load Balancer when distributing connections. For an HTTPS workload, the probe can be configured to test an appropriate TCP port, HTTP endpoint, or HTTPS endpoint, depending on the application and load balancer SKU capabilities. Azure Load Balancer evaluates probe responses for every backend instance and sends new traffic only to instances that are considered healthy. This allows VM1 or VM2 to be removed automatically from load-balancing rotation when its application is unavailable and returned when it recovers.

After the backend pool and health probe are available, the rule can associate the frontend listener, HTTPS protocol and port, backend pool, and probe to distribute the intended traffic. See Azure Load Balancer components and Manage Azure Load Balancer rules.

QUESTION NO: 61

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 have an Azure virtual machine named VM1. VM1 was deployed by using a custom Azure Resource Manager template named ARM1.json.

You receive a notification that VM1 will be affected by maintenance. You need to move VM1 to a different host immediately.

Solution: From the VM1 Redeploy + reapply blade, you select Redeploy. Does this meet the goal?

A.Yes B.No

Answer: A

Explanation:

A. Yes

B. No

ANSWER: A

Explanation:

Selecting **Redeploy** meets the goal because Azure redeploys the virtual machine to a new Azure host. This operation stops the VM, moves it to a new node in the Azure datacenter, and starts it again. Therefore, it is the appropriate immediate action when the requirement is to move VM1 away from its current host due to maintenance or a suspected host-level issue.

Redeployment retains the VM's configured disks and VM configuration; it does not require recreating the VM from ARM1.json. The custom Azure Resource Manager template used for the initial deployment is therefore not relevant to performing this host move. The operation does cause downtime while Azure deallocates and restarts the VM, so workloads should be designed for that interruption where availability is required. Azure also notes that data stored on the temporary disk is lost during a redeploy, so durable application data must reside on managed disks or another persistent storage service.

For the operational behavior and considerations of this action, see Redeploy Azure virtual machines to a new node and Maintenance and updates for Azure virtual machines.

QUESTION NO: 62

You plan to use the Azure Import/Export service to copy files to a storage account.

Which two files should you create before you prepare the drives for the import job? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. an XML manifest file

B. a dataset CSV file

C. a JSON configuration file

D. a PowerShell PS1 file

E. a driveset CSV file

ANSWER: B E

Explanation:

The correct files are **a dataset CSV file** and **a driveset CSV file**. Before preparing disks for an Azure Import job with the WAImportExport tool, you define both the data set to be copied and the physical drives that will hold that data. The dataset CSV file supplies the data-copy instructions, including the source paths and the destination details within the target Azure Storage account. This enables the tool to identify the files or folders that must be copied to the encrypted drives.

The driveset CSV file provides the drive-preparation information. It identifies the disks being prepared and associates the disk-related configuration needed by the tool while it formats, BitLocker-encrypts, and loads the drives. Together, these CSV input files allow WAImportExport to generate the journal information required for the import job and ensure that the shipped disks contain data mapped to the intended storage destination.

Microsoft documents creation of Dataset.csv and Driveset.csv as prerequisites for preparing drives for an Azure Import job. See Use the Azure Import/Export service to transfer data to Azure Blob Storage and Prepare hard drives for an import job.

QUESTION NO: 63

Your company has an Azure subscription that includes a Recovery Services vault.

You want to use Azure Backup to schedule a backup of your company's virtual machines (VMs) to the Recovery Services vault.

Which of the following VMs can you back up? Choose all that apply.

A. VMs that run Windows 10.

B. VMs that run Windows Server 2012 or higher.

C. VMs that have NOT been shut down.

D. VMs that run Debian 8.2+.

E. VMs that have been shut down.

ANSWER: A B C D E

Explanation:

Azure Backup can protect supported Azure IaaS virtual machines by storing recovery points in a Recovery Services vault. Windows 10 is included among supported Windows client operating systems, and Windows Server 2012 and later are supported Windows Server operating systems. Debian is also supported for Azure VM backup; Debian 8.2+ meets the stated supported-version requirement.

VM power state does not by itself determine whether the VM is eligible for Azure Backup. A VM that has not been shut down can be protected through normal scheduled backups. A VM that has been shut down can also be backed up. Azure Backup creates recovery points by using VM snapshots, and Microsoft documents support for backing up stopped or deallocated Azure VMs. For a stopped VM, the backup is crash-consistent because the guest operating system is not running to coordinate application-consistent processing. The VM must still be configured for backup and meet the applicable Azure Backup and VM support requirements.

Therefore, each listed VM is eligible based on its supported operating system or because its running versus shut-down state does not prevent Azure VM backup. See the Azure VM backup support matrix and Azure VM backup overview.

QUESTION NO: 64

You have an Azure subscription that contains a user account named User1.

You need to ensure that User1 can assign a policy to the tenant root management group.

What should you do?

A. Assign the Global administrator role to User1, and then instruct User1 to configure access management for Azure resources.

B. Assign the Global administrator role to User1, and then modify the default conditional access policies.

C. Assign the Owner role to User1. and then modify the default conditional access policies.

D. Assign the Owner role to User1. and then instruct User1 to configure access management for Azure resources.

ANSWER: A

Explanation:

Assign the Global administrator role to User1, and then instruct User1 to configure access management for Azure resources. This is required because the tenant root management group is above all subscriptions and management groups in the Microsoft Entra tenant. Azure RBAC permissions assigned only within a subscription do not automatically provide access at this root scope. A user who is a Microsoft Entra Global Administrator can enable *Access management for Azure resources*. This elevation assigns the User Access Administrator role at the tenant root management group scope, allowing the administrator to manage Azure role access at that scope. User1 can then establish the Azure RBAC access required to create and manage Azure Policy assignments at the tenant root management group. Applying policy at this scope is useful when a policy must govern resources consistently across the management-group hierarchy and all associated subscriptions. The elevation applies to Azure resource management access; it does not replace the need to use appropriate Azure RBAC governance for ongoing administration. For the elevation process and its root-scope effect, see Elevate access to manage all Azure subscriptions and management groups. For management-group scope and hierarchy behavior, see What are Azure management groups?.

QUESTION NO: 65

You are the global administrator for an Azure Active Directory (Azure AD) tenet named adatum.com.

You need to enable two-step verification for Azure users.

What should you do?

A. Create a sign-in risk policy in Azure AD Identity Protection

B. Enable Azure AD Privileged Identity Management.

C. Create and configure the Identity Hub.

D. Configure a security policy in Azure Security Center.

E. Create a Conditional Access policy that requires multi-factor authentication (two-step verification) for the users.

ANSWER: E

Explanation:

Create a Conditional Access policy that requires multi-factor authentication (two-step verification) for the users. Conditional Access is the Microsoft Entra ID control plane used to evaluate sign-in conditions and enforce access requirements, including multifactor authentication (MFA). A Global Administrator can create a policy, select the intended users or groups, target the relevant cloud applications, and set the Grant control to require multifactor authentication. When the policy applies, users must complete their registered second-factor method in addition to their password before access is granted.

This approach enables two-step verification as an enforceable sign-in requirement and supports scoped, gradual rollout before broader deployment. Before enforcing the policy, administrators should ensure users have appropriate authentication methods registered and should exclude emergency access (break-glass) accounts to prevent accidental administrative lockout. Microsoft recommends Conditional Access as the strategic approach for requiring MFA because it allows requirements to be applied based on users, applications, devices, locations, and risk signals. For configuration guidance, see Require multifactor authentication for all users and the Microsoft Entra Conditional Access overview.

QUESTION NO: 66

You have an Azure subscription that contains a storage account named account1.

You plan to upload the disk files of a virtual machine to account1 from your on-premises network. The on-premises network uses a public IP address space of 131.107.1.0/24.

You plan to use the disk files to provision an Azure virtual machine named VM1. VM1 will be attached to a virtual network named VNet1. VNet1 uses an IP address space of 192.168.0.0/24.

You need to configure account1 to meet the following requirements:

- Ensure that you can upload the disk files to account1.
- Ensure that you can attach the disks to VM1.
- Prevent all other access to account1.

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

NOTE: Each correct selection is worth one point.

A. From the Networking blade of account1, select Selected networks.

B. From the Networking blade of account1, select Allow trusted Microsoft services to access this storage account.

C. From the Networking blade of account1, add the 131.107.1.0/24 IP address range.

D. From the Networking blade of account1, add VNet1.

E. From the Service endpoints blade of VNet1, add a service endpoint.

ANSWER: A C

Explanation:

Selecting **Selected networks** configures the storage account firewall to deny public-endpoint traffic by default, rather than allowing access from all networks. This establishes the required restricted-access posture while permitting explicitly configured network exceptions. Storage account network rules apply to requests to the public endpoint and can be used to limit access to approved IP address ranges, virtual networks, or trusted-service exceptions.

Adding the **131.107.1.0/24 IP address range** creates the required IP network rule for the on-premises public address space. Upload requests originating from addresses in that range can therefore reach the Blob service in account1, while requests from other public IP addresses remain blocked by the selected-networks firewall configuration. The uploaded VHD can then be used by the Azure platform when provisioning and attaching the disk to VM1; the upload client itself does not require VNet1 access to perform the upload.

Microsoft documents that IP network rules support IPv4 public-address ranges in CIDR notation and that selected networks restrict storage public-endpoint access to configured rules. See Configure Azure Storage firewalls and virtual networks and Configure network security for Azure Storage.

QUESTION NO: 67

You have an app named App1 that runs on an Azure web app named webapp1.

The developers at your company upload an update of App1 to a Git repository named Git1.

Webapp1 has the deployment slots shown in the following table.

Name	Function
webapp1-prod	Production
webapp1-test	Staging

You need to ensure that the App1 update is tested before the update is made available to users.

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

NOTE: Each correct selection is worth one point.

A. Swap the slots

B. Deploy the App1 update to webapp1-prod, and then test the update

C. Stop webapp1-prod

D. Deploy the App1 update to webapp1-test, and then test the update

E. Stop webapp1-test

ANSWER: A D

Explanation:

Deploy the App1 update to webapp1-test, and then test the update is correct because an Azure App Service deployment slot provides an isolated, live app environment with its own hostname. Deploying the repository update to the test slot allows the development or operations team to validate the new application version, including functionality, integrations, and slot-specific configuration, before users of the production app receive it. The production slot remains on the currently validated version throughout this testing period.

Swap the slots is correct after testing succeeds. An App Service slot swap exchanges the content and applicable configuration between the source and target slots, promoting the tested version into the production slot. App Service warms up the destination slot before completing the swap, which supports a low-downtime release process. The formerly production version moves to the test slot, providing a practical rollback path if a post-release issue is discovered. This deployment-slot workflow is specifically intended to separate validation from production exposure while enabling controlled releases. For details, see Set up staging environments in Azure App Service and Swap deployment slots.

QUESTION NO: 68

You have an Azure subscription named Subscription1 that contains an Azure virtual network named VNet1. VNet1 connects to your on-premises network by using

Azure ExpressRoute.

You plan to prepare the environment for automatic failover in case of ExpressRoute failure.

You need to connect VNet1 to the on-premises network by using a site-to-site VPN. The solution must minimize cost.

Which three actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Create a connection

B. Create a local site VPN gateway

C. Create a VPN gateway that uses the VpnGw1 SKU

D. Create a gateway subnet

E. Create a VPN gateway that uses the Basic SKU

ANSWER: A B C

Explanation:

Creating a local site VPN gateway, creating a VPN gateway that uses the VpnGw1 SKU, and creating a connection provide the required components for a site-to-site VPN backup path alongside ExpressRoute. The local network gateway represents the on-premises VPN device and declares the on-premises address prefixes that Azure must route through the VPN tunnel. The Azure VPN gateway provides the VPN endpoint in VNet1. For ExpressRoute and VPN gateway coexistence, Azure

requires a route-based VPN gateway and does not support the Basic SKU. VpnGw1 is therefore the lowest-cost eligible VPN gateway SKU among the available choices. Finally, the connection resource associates the Azure VPN gateway with the local network gateway and defines the site-to-site IPsec/IKE connection. When ExpressRoute and VPN coexist in the same virtual network, they use the existing GatewaySubnet; a separate gateway subnet is not needed. With the VPN connection configured as the backup route, routing can use it if the ExpressRoute path becomes unavailable. See Configure ExpressRoute and VPN gateway coexistence and About Azure VPN Gateway.

QUESTION NO: 69

You have an Azure subscription that contains a policy-based virtual network gateway named GW1 and a virtual network named VNet1.

You need to ensure that you can configure a point-to-site connection from an on-premises computer to VNet1. Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Add a service endpoint to VNet1

B. Reset GW1

C. Create a route-based virtual network gateway

D. Add a connection to GW1

E. Delete GW1

F. Add a public IP address space to VNet1

ANSWER: C E

Explanation:

Point-to-site VPN connectivity requires a route-based virtual network gateway. Azure supports point-to-site configurations on route-based VPN gateways because this gateway type uses routes to direct traffic between VPN clients and virtual-network resources, and it supports the protocols and authentication methods used for point-to-site connections. A policy-based gateway cannot be converted in place to a route-based gateway. Therefore, the existing policy-based gateway must first be removed so that a replacement gateway of the required type can be deployed in the gateway subnet of VNet1.

The required actions are **Delete GW1** and **Create a route-based virtual network gateway**. After the route-based gateway has been provisioned, an administrator can configure its point-to-site settings, including a VPN client address pool and an appropriate authentication method, such as Microsoft Entra ID, certificates, or RADIUS. The client configuration can then be downloaded and installed on the on-premises computer to establish connectivity to resources in VNet1.

Microsoft documents that point-to-site VPN connections are supported by route-based VPN gateways and describes the available point-to-site configuration choices in About Point-to-Site VPN. Gateway types and their capabilities are also covered in About VPN Gateway settings.

QUESTION NO: 70

You have an Azure subscription that contains a virtual network named VNet1. VNet1 contains four subnets named Gateway, Perimeter, NVA, and Production.

The NVA subnet contains two network virtual appliances (NVAs) that will perform network traffic inspection between the Perimeter subnet and the Production subnet.

You need to implement an Azure load balancer for the NVAs. The solution must meet the following requirements:

☞ The NVAs must run in an active-active configuration that uses automatic failover.

☞ The load balancer must load balance traffic to two services on the Production subnet. The services have different IP addresses.

Which three actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Deploy a basic load balancer

B. Deploy a standard load balancer

C. Add two load balancing rules that have HA Ports and Floating IP enabled

D. Add two load balancing rules that have HA Ports enabled and Floating IP disabled

E. Add two frontend IP configurations, a backend pool, and a health probe

F. Add a frontend IP configuration, two backend pools, and a health probe

ANSWER: B C E

Explanation:

Deploy a standard load balancer because Azure Load Balancer HA Ports is a Standard Load Balancer capability and is intended for highly available network virtual appliance deployments. A Standard Load Balancer health probe monitors the NVAs and removes an unhealthy appliance from the available backend targets, providing automatic failover while both healthy NVAs can process traffic in an active-active design.

Add two frontend IP configurations, a backend pool, and a health probe. The backend pool contains the two NVA instances, while the health probe provides the availability signal used for traffic distribution and failover. Two frontend IP configurations are required because the inspection path has two directions: traffic entering the NVA tier from Perimeter and return traffic entering it from Production. User-defined routes can direct each direction to the applicable frontend address, helping maintain a symmetric traffic path through an NVA.

Add two load balancing rules that have HA Ports and Floating IP enabled. HA Ports permits all ports and protocols to be handled by each rule, which is appropriate when the NVAs transparently inspect diverse traffic flows. Floating IP enables direct server return behavior needed by this NVA high-availability pattern. Together, the rules associate each frontend direction with the shared NVA backend pool. See HA Ports overview and Floating IP for Azure Load Balancer.

QUESTION NO: 71

You have an Azure subscription that contains a storage account named storageacct1234 and two users named User1 and User2.

You assign User1 the roles shown in the following exhibit.



Which two actions can User1 perform? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. Modify the firewall of storageacct1234.

B. View blob data in storageacct1234.

C. View file shares in storageacct1234.

D. Upload blob data to storageacct1234.

E. Assign roles to User2 for storageacct1234.

ANSWER: B D

Explanation:

“View blob data in storageacct1234” and “Upload blob data to storageacct1234” are permitted because User1 has the **Storage Blob Data Contributor** role scoped to the storage account. This is an Azure RBAC *data-plane* role, meaning that it authorizes operations on the contents of Azure Blob Storage rather than merely allowing the user to view the storage-account resource in Azure Resource Manager. The role includes permissions to read, list, create, write, and delete blob containers and blob data. Consequently, User1 can browse and read blob content and can upload new blobs or overwrite existing blob content in storageacct1234.

The Reader assignment provides complementary management-plane read access at its assigned scope. It allows User1 to see the relevant Azure resources and their configuration in the portal, while the Storage Blob Data Contributor assignment supplies the separate authorization needed to access the actual blob data. Azure Storage deliberately distinguishes management-plane permissions from data-plane permissions, so a Reader role by itself would not enable blob-content access; the Blob Data Contributor role is what grants these data operations. Microsoft documents the role’s included blob-data actions in the Azure Storage built-in roles reference and describes Azure RBAC authorization for Blob Storage in Assign an Azure role for access to blob data.

QUESTION NO: 72

You have an Azure subscription that contains multiple virtual machines in the West US Azure region. You need to use Traffic Analytics in Azure Network Watcher to monitor virtual machine traffic.

Which two resources should you create? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A.a Log Analytics workspace B.an Azure Monitor workbook

C.a storage account

D.a Microsoft Sentinel workspace

E.a Data Collection Rule (DCR) in Azure Monitor

Answer: AE

Explanation:

A. A Log Analytics workspace - Traffic Analytics requires a Log Analytics workspace to store and analyze network traffic data.

B. A Data Collection Rule (DCR) in Azure Monitor - You need to create a Data Collection Rule within Azure Monitor to specify what data should be collected and sent to the Log Analytics workspace, including the network traffic data for Traffic Analytics.

C. A storage account - NSG flow logs (used by Traffic Analytics) are written to a storage account before Traffic Analytics processes them.

ANSWER: A C

Explanation:

Traffic Analytics uses network flow-log data to provide aggregated views and insights into traffic traversing network security groups, such as communicating IP addresses, ports, protocols, traffic volume, and potentially malicious flows. A **Log Analytics workspace - Traffic Analytics requires a Log Analytics workspace to store and analyze network traffic data.** is required because Traffic Analytics sends its processed analytics results to a Log Analytics workspace. The

workspace provides the Azure Monitor Logs data store that supports Traffic Analytics queries, reports, and the Traffic Analytics experience.

A **storage account - NSG flow logs (used by Traffic Analytics) are written to a storage account before Traffic Analytics processes them.** is also required. NSG flow logs are the data source for Traffic Analytics and are first written as raw log files to an Azure Storage account. When Traffic Analytics is enabled as part of flow-log configuration, Azure processes the relevant flow data and publishes the resulting analytics records to the selected Log Analytics workspace. Together, the storage account supplies the raw NSG flow-log destination, while the Log Analytics workspace supplies the destination for analyzed traffic data and queries. See Traffic Analytics documentation and NSG flow logs overview.

QUESTION NO: 73

You plan to move services from your on-premises network to Azure.

You identify several virtual machines that you believe can be hosted in Azure. The virtual machines are shown in the following table.

Name	Role	Operating system (OS)	Environment
Sea-DC01	Domain controller	Windows Server 2016	Hyper-V on Windows Server 2016
NYC-FS01	File server	Windows Server 2012 R2	VMware vCenter Server 5.1
BOS-DB01	Microsoft SQL server	Windows Server 2016	VMware vCenter Server 6
Sea-CA01	Certification authority (CA)	Windows Server 2012 R2	Hyper-V on Windows Server 2016
Hou-NW01	DHCP/DNS	Windows Server 2008 R2	VMware vCenter Server 5.5

Which two virtual machines can you access by using Azure migrate? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Sea-CA0l

B. Hou-NW01

C. NYC-FS01

D. Sea-DC01

E. BOS-DB01

ANSWER: C E

Explanation:

NYC-FS01 and **BOS-DB01** can be accessed by Azure Migrate because their virtualization environments meet the supported discovery and assessment requirements shown in the table. Azure Migrate uses an Azure Migrate appliance to connect to the supported on-premises virtualization-management infrastructure and discover server configuration, performance, and dependency information. That discovered inventory can then be assessed for Azure readiness, recommended Azure VM sizing, and estimated costs before migration.

For VMware-based discovery, the appliance communicates with vCenter Server and requires a supported vCenter Server and ESXi environment. For Hyper-V-based discovery, Azure Migrate supports the applicable Hyper-V deployment and host-management configuration. The platform and version details associated with NYC-FS01 and BOS-DB01 are supported for Azure Migrate access, allowing the appliance to inventory and assess those workloads as candidates for migration to Azure.

Microsoft maintains the authoritative requirements, including supported VMware versions and required vCenter permissions, in the Azure Migrate VMware support matrix. Supported Hyper-V environments and appliance prerequisites are documented in the Azure Migrate Hyper-V support matrix.

QUESTION NO: 74

You have an Azure subscription that contains the identifies shown in the following table.

Name	Type	Member of
User1	User	None
User2	User	Group1
Principal1	Managed identity	None
Principal2	Managed identity	Group1

User1, Principle, and Group1 are assigned the Monitoring Reader role.

An action an alert rule named Alert1 that uses AG1.

You need to identify who will receive an email notification when Alert1 is triggered.

Who should you identity?

A. User1, User2, Principle, and Principle2

B. User1 and Principle only

C. User1 only

D. User1 and User2 only

ANSWER: D

Explanation:

User1 and User2 only will receive the notification. AG1 is configured to notify the members of the Azure Resource Manager role used by the alert action. Azure Monitor evaluates the Monitoring Reader assignments at the subscription scope and sends the email notification to Microsoft Entra user accounts represented by those assignments. User1 is directly assigned the role, so User1 is a recipient. Group1 also has the role assignment; therefore, its eligible user member, User2, is included when Azure Monitor resolves the role recipients.

This behavior is useful when operations teams manage alert recipients through role-based access control rather than maintaining a separate static list of email addresses in every action group. Membership and role assignments are evaluated so that the appropriate Microsoft Entra users receive operational notifications. The relevant configuration is the action group's *Email Azure Resource Manager Role* notification type, which targets users associated with the selected Azure RBAC role. Microsoft documents this action type and its role-recipient behavior in Azure Monitor action groups. For background on assigning Azure RBAC roles to users and groups, see Assign Azure roles using the Azure portal.

QUESTION NO: 75

You have an Azure web application that is hosted on several virtual machines.

You plan to publish the application by using Azure Application Gateway. You need to protect the application from common web exploits by using managed OWASP rules. The solution must block malicious requests.

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

NOTE: Each correct selection is worth one point.

A. Deploy an Application Gateway that uses the WAF_v2 SKU

B. Create a WAF policy in Prevention mode

C. Configure an NSG rule that denies inbound port 443 traffic

D. Create an Azure DDoS Protection plan

E. Configure a WAF policy in Detection mode

ANSWER: A B

Explanation:

You should deploy an **Application Gateway that uses the WAF_v2 SKU**. The Web Application Firewall SKU provides the web application firewall capability that evaluates HTTP and HTTPS requests before they reach the backend pool. The v2 SKU supports WAF policies and managed rule sets.

You should also create and associate a **WAF policy in Prevention mode**. Prevention mode blocks requests that match enabled WAF rules. In contrast, Detection mode logs matching requests but allows them to continue to the backend application. Microsoft-managed rules include rule sets based on OWASP Core Rule Set protections. See Web Application Firewall on Azure Application Gateway and Create Web Application Firewall policies for Application Gateway.

QUESTION NO: 76

You plan to deploy several Azure virtual machines that will run Windows Server 2019 in a virtual machine scale set by using an Azure Resource Manager template.

You need to ensure that NGINX is available on all the virtual machines after they are deployed. What should you use?

A. the Publish-AzVMDscConfiguration cmdlet

B. Azure Application Insights

C. Azure Custom Script Extension

D. the New-AzConfigurationAssignement cmdlet

ANSWER: C

Explanation:

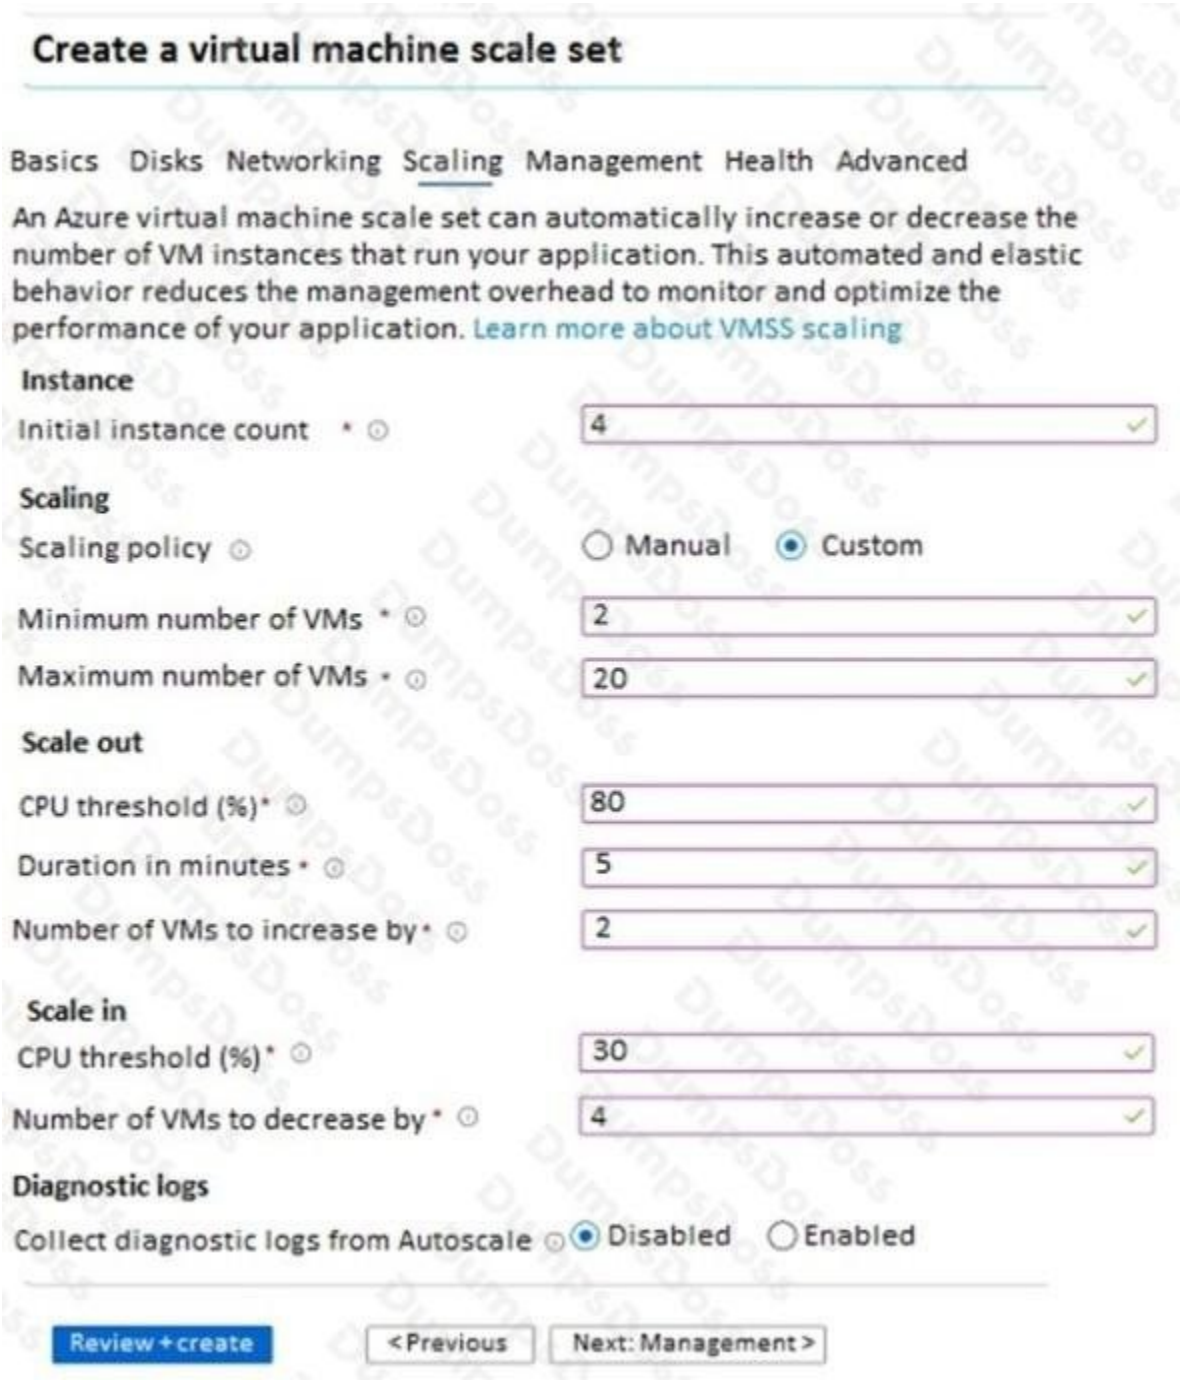
Azure Custom Script Extension is the appropriate choice because a VM extension can be included in the Azure Resource Manager template for the virtual machine scale set and runs on each scale-set VM instance. The extension downloads specified script files from Azure Storage, GitHub, or another accessible location, and then executes a command on the instance. For Windows Server 2019, the deployment can run a PowerShell script that downloads or installs NGINX, applies any required configuration, and starts the service. This makes the NGINX installation part of the repeatable, declarative scale-set deployment rather than a manual post-deployment task.

VM Scale Sets support extensions at the scale-set model level, so instances created from that model receive the extension configuration. This is especially useful when instances are added later through scaling, because the same installation workflow is consistently applied to newly provisioned instances. The script should be written to be idempotent where possible, allowing it to safely handle retries or reimaging scenarios. Microsoft documents the Custom Script Extension as a mechanism for post-deployment configuration and scripting on Windows VMs, including use with scale sets. See Custom Script Extension for Windows and Virtual machine scale set extensions.

QUESTION NO: 77 - (HOTSPOT)

HOTSPOT

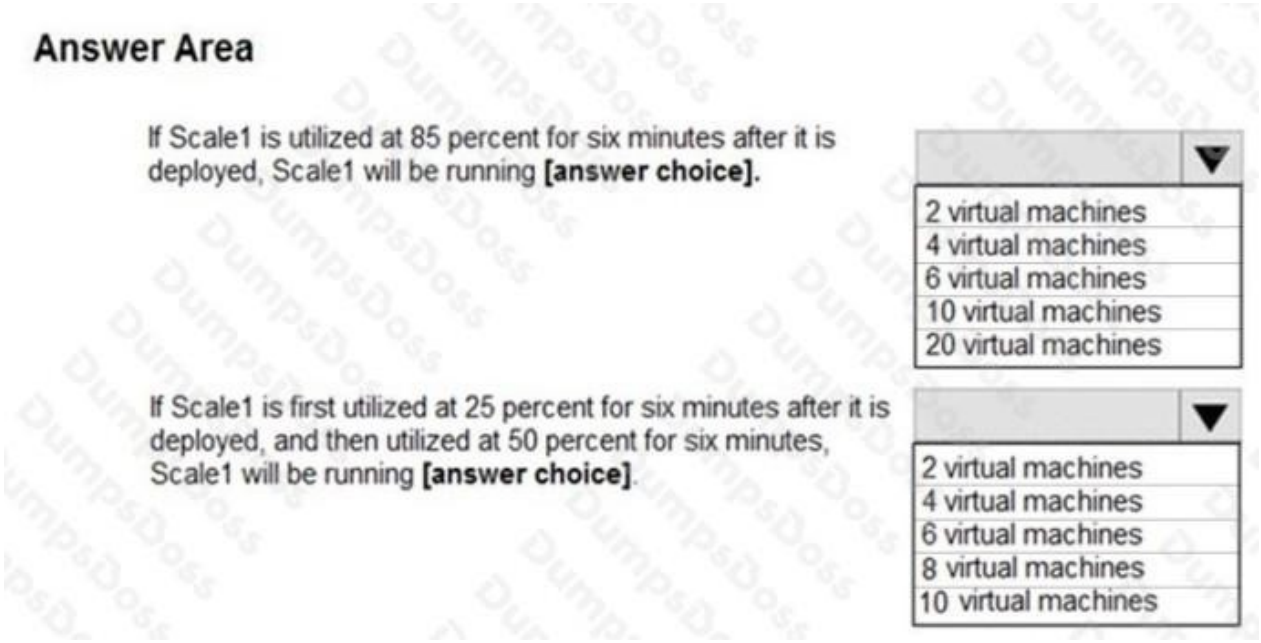
You create a virtual machine scale set named Scale1. Scale1 is configured as shown in the following exhibit.



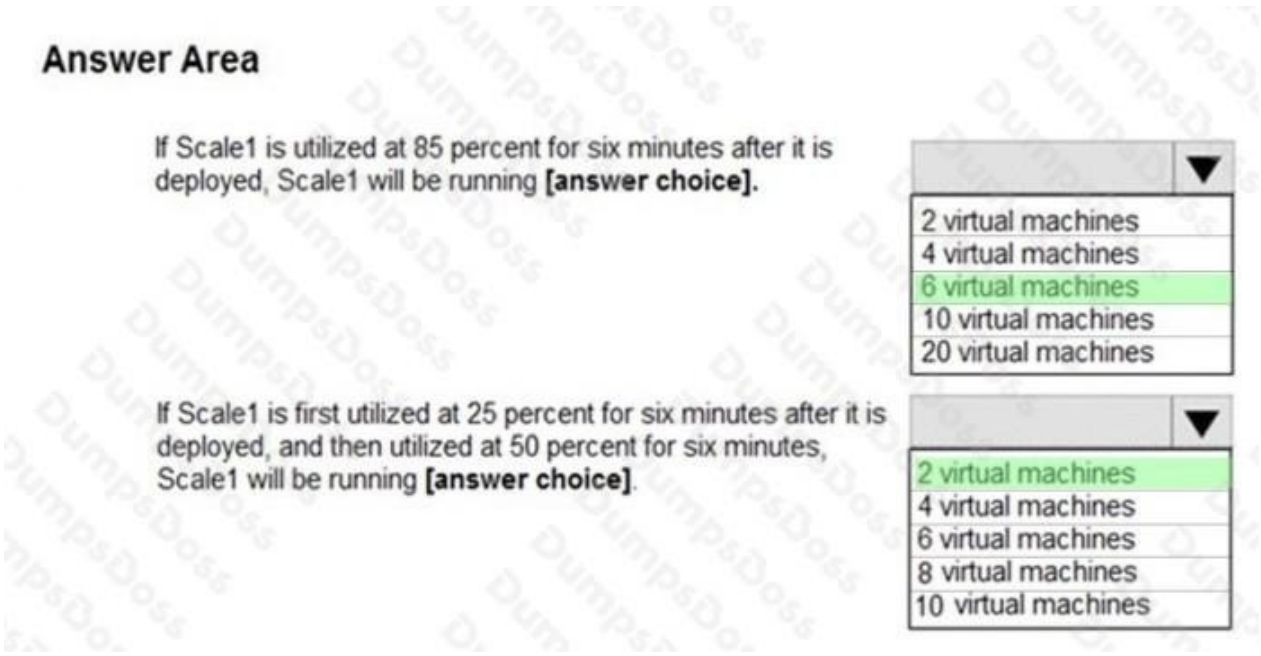
Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Hot Area:



ANSWER:



Explanation:

Azure Monitor autoscale evaluates the metric values over the configured duration and then adjusts the virtual machine scale set according to its active profile, subject to the profile's minimum and maximum instance limits. Scale1 begins with four virtual machine instances. When CPU utilization is 85 percent for six minutes, the configured scale-out threshold has been sustained for the required period. The scale-out action adds two instances to the current capacity, so the scale set runs six virtual machines.

For the second condition, CPU utilization is initially 25 percent for six minutes. This satisfies the scale-in condition for the configured evaluation period. The scale-in action is configured to reduce capacity by four instances, but autoscale does not reduce an instance count below the profile's minimum capacity. Since the scale set began with four instances and has a minimum of two, the resulting capacity is two virtual machines. The following six-minute interval at 50 percent CPU does not trigger another scale action, because it does not satisfy the configured scale-out or scale-in metric threshold. The scale set therefore remains at two instances.

This behavior illustrates that scale actions modify the current instance count, while capacity limits are enforced as guardrails on every action. Microsoft documents that autoscale uses metric rules and profile capacity limits to calculate the resulting

resource capacity. See Azure Monitor autoscale overview and Understand autoscale settings for details on metric-duration evaluation, scale actions, and minimum instance limits.

QUESTION NO: 78

You need to create an Azure Storage account named storage1. The solution must meet the following requirements:

• Support Azure Data Lake Storage.

• Minimize costs for infrequently accessed data.

• Automatically replicate data to a secondary Azure region.

Which three options should you configure for storage1? Each correct answer presents part of the solution.

NOTE: Each correct answer is worth one point.

A. the Cool access tier

B. the Hot access tier

C. hierarchical namespace

D. zone-redundant storage (ZRS)

E. geo-redundant storage (GRS)

ANSWER: A C E

Explanation:

Configure **hierarchical namespace**, **the Cool access tier**, and **geo-redundant storage (GRS)** for storage1. Azure Data Lake Storage Gen2 capabilities are enabled on an Azure Storage account by enabling the hierarchical namespace feature. This provides the directory-oriented namespace and atomic directory operations used by Data Lake Storage workloads. The setting is selected when the storage account is created and is appropriate for a general-purpose v2 storage account.

The Cool access tier is intended for data that is accessed infrequently but must remain available online. It has lower storage charges than the Hot tier, with higher access and transaction charges, making it the relevant account-level access-tier choice when reducing storage costs for infrequently accessed blob data.

Geo-redundant storage (GRS) asynchronously replicates data from the primary region to a paired secondary Azure region. It therefore provides the required automatic cross-region replication while retaining local redundancy in both the primary and secondary regions. Read access to the secondary region is not required by the stated requirements, so standard GRS satisfies the replication requirement.

See Azure Data Lake Storage hierarchical namespace and Azure Storage redundancy.

QUESTION NO: 79

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 have a Microsoft Entra tenant named Adatum.com and an Azure Subscription named Subscription1. Adatum.com contains a group named Developers. Subscription1 contains a resource group named Dev.

You need to provide the Developers group with the ability to create Azure Logic Apps in the Dev resource group.

Solution: On Subscription1, you assign the DevTest Labs User role to the Developers group.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. The DevTest Labs User built-in role is intended for users working with resources within an Azure DevTest Lab. Its permissions focus on lab-oriented activities, such as using virtual machines made available through a lab, rather than granting general Azure resource deployment permissions in a resource group. Assigning this role at the Subscription1 scope does not give the Developers group the required Microsoft.Logic permissions to create a Logic App resource in the Dev resource group.

To create Azure Logic Apps, the group needs a role assignment that includes write permissions for the applicable Logic Apps resource type at the Dev resource-group scope (or a broader scope). For example, the Contributor role can manage and create resources in the assigned scope, including Logic Apps, while not allowing the assignment of Azure RBAC roles. Azure RBAC built-in roles have defined permission sets and their effective access is limited to actions explicitly included by the assigned role. Microsoft documents the DevTest Labs User role separately from general resource-management roles, confirming that it is not a suitable role for creating Logic Apps in a standard resource group.

See Azure built-in roles and Assign Azure roles using the Azure portal.

QUESTION NO: 80

You have an Azure subscription that contains the resources in the following table.

Name	Type
RG1	Resource group
store1	Azure Storage account
Sync1	Azure File Sync

Store1 contains a file share named data. Data contains 5,000 files.

You need to synchronize the files in the file share named data to an on-premises server named Server1.

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

NOTE: Each correct selection is worth one point.

A. Create a container instance

B. Register Server1

C. Install the Azure File Sync agent on Server1

D. Download an automation script

E. Create a sync group

ANSWER: B C E

Explanation:

Azure File Sync is the Azure service designed to synchronize an Azure file share with a Windows Server file system. To make Server1 available as an Azure File Sync server, first install the Azure File Sync agent on Server1. The agent supplies the synchronization components and the Server Registration UI or PowerShell capabilities required to connect the server to Azure. After installation, register Server1 with the Storage Sync Service. Registration establishes the server as a managed

server resource in that Storage Sync Service and permits it to participate in synchronization relationships. Finally, create a sync group. A sync group defines the synchronization topology: it contains one cloud endpoint, which is the *data* Azure file share, and one or more server endpoints, including a local path on Server1. After the cloud endpoint and server endpoint are configured, Azure File Sync can synchronize the existing 5,000 files from the Azure file share to the specified server location. This arrangement also supports ongoing bidirectional synchronization according to Azure File Sync behavior and configuration. See the Azure File Sync deployment guide and Azure File Sync planning documentation.

QUESTION NO: 81

You need to prepare the environment to meet the authentication requirements.

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

NOTE: Each correct selection is worth one point.

A. Allow inbound TCP port 8080 to the domain controllers in the Miami office.

B. Add https://autologon.microsoftazuread-sso.com to the intranet zone of each client computer in the Miami office.

C. Join the client computers in the Miami office to Azure AD.

D. Install the Active Directory Federation Services (AD FS) role on a domain controller in the Miami office.

E. Install Azure AD Connect on a server in the Miami office and enable Pass-through Authentication.

ANSWER: B E

Explanation:

Adding `https://autologon.microsoftazuread-sso.com` to the Local intranet zone on the Miami client computers prepares browsers to use Microsoft Entra seamless single sign-on. When a user accesses the service from a domain-joined computer on the corporate network, this zone assignment allows the browser to automatically submit the user's Kerberos ticket rather than prompting for credentials. Microsoft recommends deploying this setting centrally through Active Directory Group Policy so that it is applied consistently to the appropriate computers and users.

Installing Azure AD Connect on a suitable server in the Miami office and enabling pass-through authentication provides the required hybrid authentication path. Azure AD Connect installs and registers the pass-through authentication agent, which validates sign-in attempts against on-premises Active Directory while avoiding the need to expose inbound authentication ports to the internet. Seamless SSO can be enabled alongside pass-through authentication, allowing users on the corporate network to receive a streamlined sign-in experience while their credentials continue to be validated against the on-premises directory. Microsoft documents both the intranet-zone configuration and Azure AD Connect setup requirements in the Seamless SSO quick start and the Pass-through Authentication quick start.

QUESTION NO: 82

You have an Azure subscription. The subscription contains virtual machines that run Windows Server.

You have a data collection rule (DCR) named Rule1

You plan to use the Azure Monitor Agent to collect events from Windows System event logs.

You only need to collect system events that have an ID of 1001.

Which type of query should you use for the data source in Rulel?

A. SQL

B. KQL

C. XPath

ANSWER: C

Explanation:

XPath is the appropriate query type for a Windows Event Log data source in an Azure Monitor data collection rule. Windows event records are XML-based, and the Azure Monitor Agent uses XPath expressions in the DCR to select the events to collect before they are sent to Azure Monitor. An XPath filter can specify both the Windows event channel and conditions on fields within the event, such as the event ID. For this requirement, the System channel can be filtered with an expression such as `System!*[System[(EventID=1001)]]`. This ensures that the agent collects only events whose EventID is 1001 from the Windows System event log.

This source-side filtering is useful because it limits collection to the required records, reducing unnecessary ingestion and keeping the configured data stream focused on the stated monitoring need. The DCR's Windows Event Log data source accepts one or more XPath queries, with each query identifying a channel and an XPath selection expression. Microsoft documents the required syntax and examples for event ID filtering in its Windows event collection with Azure Monitor Agent guidance. For the DCR data-source schema and XPath-query property, see the Microsoft.Insights data collection rule template reference.

QUESTION NO: 83

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

You need to configure scheduled backups for the virtual machines. The backups must be retained for 30 days.

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

NOTE: Each correct selection is worth one point.

A. a Recovery Services vault

B. a backup policy

C. an Azure Backup Server

D. an Azure Site Recovery replication policy

E. an availability set

ANSWER: A B

Explanation:

A **Recovery Services vault** is required to store and manage recovery points for Azure virtual machines protected by Azure Backup. The vault provides the management boundary for backup configuration, recovery operations, and backup data retention.

A **backup policy** defines the backup schedule and retention settings. For this scenario, the policy can be configured to create recovery points on the required schedule and retain them for 30 days. The policy is then assigned while backup is enabled for each virtual machine. See Azure VM backup overview and Back up Azure virtual machines.

QUESTION NO: 84

You have an Azure subscription named Subscription1.

You have 5 TB of data that you need to transfer to Subscription1.

You plan to use an Azure Import/Export job.

What can you use as the destination of the imported data?

A. an Azure Cosmos DB database

B. Azure Blob storage

C. Azure Data Lake Store

D. the Azure File Sync Storage Sync Service

ANSWER: B

Explanation:

Azure Blob storage is a supported destination for an Azure Import/Export import job. This service is intended for transferring large volumes of data to Azure by preparing and encrypting compatible physical drives, shipping them to a Microsoft datacenter, and having Microsoft copy the data into the specified Azure Storage account. For a blob import job, the job configuration identifies the target storage account and the blob containers to which the contents of each drive are copied. This makes the service well suited to an offline transfer of 5 TB when network transfer is impractical, too slow, or would consume excessive bandwidth. The imported data is stored as block blobs in the selected storage account, after which it can be accessed and managed through standard Azure Storage tools and APIs. Azure Import/Export supports imports to Azure Blob Storage and Azure Files; Blob storage is the applicable supported destination among the listed choices. For the service overview and supported import workflow, see Azure Import/Export service overview and Import data to Azure Blob Storage with Azure Import/Export.

QUESTION NO: 85

You have an Azure Active Directory (Azure AD) tenant.

All administrators must enter a verification code to access the Azure portal.

You need to ensure that the administrators can access the Azure portal only from your on-premises network.

What should you configure?

A. an Azure AD Identity Protection user risk policy.

B. the multi-factor authentication service settings.

C. the default for all the roles in Azure AD Privileged Identity Management

D. an Azure AD Identity Protection sign-in risk policy

E. an Azure AD Conditional Access policy that requires MFA and restricts access to a named location (on-premises IP ranges)

ANSWER: E

Explanation:

An Azure AD Conditional Access policy that requires MFA and restricts access to a named location (on-premises IP ranges) is the appropriate configuration because Conditional Access can evaluate both the identity of the administrator and the network location of the sign-in. Define a named location containing the organization’s trusted public, outbound IP address ranges. Then create a policy scoped to the relevant administrator accounts or directory roles and the Azure management application, including Azure portal access.

To enforce on-premises-only access, configure the policy so that sign-ins from locations outside the named location are blocked. For sign-ins from the permitted named location, require multifactor authentication as the grant control. This ensures administrators must provide the required verification code while also preventing portal access from home networks, mobile networks, or other untrusted internet locations.

Conditional Access is evaluated during authentication and supports location-based conditions, named locations, targeted users or directory roles, cloud applications, and grant controls such as requiring MFA. Microsoft documents named locations and network-based policy assignments at Named locations and Conditional Access network assignments. For the overall policy model and grant controls, see Microsoft Entra Conditional Access overview.

QUESTION NO: 86

You have an Azure subscription that contains a storage account. The account stores website data.

You need to ensure that inbound user traffic uses the Microsoft point-of-presence (POP) closest to the user's location.

What should you configure?

A. private endpoints

B. Azure Firewall rules

C. Routing preference

D. load balancing

ANSWER: C

Explanation:

Routing preference is the appropriate setting because it controls how clients reach the public endpoints of an Azure Storage account. When the storage account uses the Microsoft global network routing preference, client traffic enters Microsoft's global network at the point-of-presence nearest to the client and then travels across Microsoft's backbone to the storage account endpoint. This is designed to improve network performance, reliability, and routing consistency for publicly accessible storage-hosted website content.

The setting applies to Azure Storage public endpoints and can be configured for the account's default endpoint as well as selected custom domain scenarios. The Microsoft global network option is the relevant choice when the requirement specifically concerns using a nearby Microsoft edge location for inbound traffic. Azure documents that this routing method causes traffic to enter the Microsoft global network as close to the client as possible, before it is routed to the storage service.

For configuration details and routing behavior, see Configure network routing preference for Azure Storage and Configure Azure Storage network security.

QUESTION NO: 87

You have an Azure virtual machine named VM1 and an Azure Key Vault named vault1.

An application that runs on VM1 must retrieve a secret from vault1. You must avoid storing credentials in the application code or configuration files.

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

NOTE: Each correct selection is worth one point.

A. Enable a system-assigned managed identity on VM1

B. Assign the Key Vault Secrets User role to the VM1 managed identity

C. Create a shared access signature for vault1

D. Assign the Owner role to the VM1 managed identity

E. Enable a user-assigned public IP address on VM1

ANSWER: A B

Explanation:

Enabling a system-assigned managed identity on VM1 creates an identity in Microsoft Entra ID that is associated with the virtual machine. The application can use this identity to obtain tokens without requiring a stored password, certificate, or client secret.

Assigning the **Key Vault Secrets User** role to the managed identity at the vault scope grants the identity permission to read secret values when the vault uses the Azure RBAC permission model. The application can then authenticate by using the managed identity and retrieve only the required secrets from vault1.

For more information, see Authenticate to Azure Key Vault and Provide access to Key Vault keys, certificates, and secrets with an Azure role-based access control.

QUESTION NO: 88

You have an Azure subscription that contains several resource groups.

You need to retain Azure Activity Log events for 180 days in an Azure Storage account. The solution must collect events from all resource groups in the subscription.

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

NOTE: Each correct selection is worth one point.

A. Create an Azure Storage account

B. Create a diagnostic setting at the subscription scope

C. Create a diagnostic setting for each virtual machine

D. Create a data collection rule

E. Create an Azure Monitor metric alert rule

ANSWER: A B

Explanation:

You should create an **Azure Storage account** to provide a durable destination for the exported Activity Log events. The account can retain the exported logs according to its configured data-management and retention requirements.

You should create a **diagnostic setting at the subscription scope** and select the storage account as a destination. The Azure Activity Log is collected at subscription scope, and a diagnostic setting exports the selected Activity Log categories to the configured destination. A diagnostic setting on an individual virtual machine collects resource diagnostic logs and metrics, not the subscription-wide Activity Log. See Azure Activity Log and Diagnostic settings in Azure Monitor.

QUESTION NO: 89

You have an Azure subscription that contains a virtual network named VNet1 and an Azure Storage account named storage1.

You need to ensure that resources in a subnet of VNet1 can access storage1 by using the Azure Storage public endpoint. Access from all other networks must be denied.

What should you configure first?

A. A Microsoft.Storage service endpoint on the subnet

B. A private endpoint for storage1

C. A public IP address on the subnet

D. A virtual network gateway in VNet1

ANSWER: A

Explanation:

You should configure a **Microsoft.Storage service endpoint** on the subnet. A service endpoint provides a direct and secure route from the subnet to the Azure Storage service over the Azure backbone network. After the endpoint is enabled, the subnet can be selected in the storage account network rules.

You must then configure storage1 to allow access from selected networks and add the subnet to the storage account virtual network rules. A private endpoint is an alternative design that provides a private IP address for the storage account, but it is not required when the stated requirement is to use the Azure Storage public endpoint with service endpoint-based restrictions.

For more information, see Configure Azure Storage firewalls and virtual networks and Virtual network service endpoints.

QUESTION NO: 90

Name	Type	Azure region	Resource group
VNet1	Virtual network	West US	RG2
VNet2	Virtual network	West US	RG1
VNet3	Virtual network	East US	RG1
NSG1	Network security group (NSG)	East US	RG2

You have an Azure subscription that contains the resources in the following table.

To which subnets can you apply NSG1?

A. the subnets on VNet1 only

B. the subnets on VNet2 and VNet3 only

C. the subnets on VNet2 only

D. the subnets on VNet3 only

E. the subnets on VNet1, VNet2, and VNet3

ANSWER: D

Explanation:

the subnets on VNet3 only is correct because an Azure network security group can be associated with either a subnet or a network interface, but the network security group must be deployed in the same Azure region as the virtual network that contains the target subnet. Network security groups are regional Azure resources; their rules can be evaluated for traffic to resources in the virtual network where they are associated, but an NSG cannot be attached across regional boundaries.

Based on the resource table, NSG1 and VNet3 are in the same region. Therefore, NSG1 is eligible for association with any subnet that belongs to VNet3. This association applies the NSG's inbound and outbound security rules to traffic flowing through that subnet, subject to Azure's NSG rule-processing behavior and any applicable rules at individual network interfaces. Before associating an NSG with a subnet in administration workflows, verify that the NSG and the virtual network have matching locations and that the subnet is not subject to a conflicting design requirement. Microsoft documents NSG association and regional requirements in the Network security groups overview and describes subnet-level associations in Manage network security groups.

QUESTION NO: 91

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 have an Azure virtual machine named VM1 that runs Windows Server 2016.

You need to create an alert in Azure when more than two error events are logged to the System event log on VM1 within an hour.

Solution: You create an Azure storage account and configure shared access signatures (SASs). You install the Microsoft Monitoring Agent on VM1. You create an alert in Azure Monitor and specify the storage account as the source.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. To generate an Azure Monitor alert based on entries in a Windows System event log, the virtual machine's event data must be collected into an Azure Monitor Logs workspace, also called a Log Analytics workspace. The Microsoft Monitoring Agent can collect Windows Event Log entries when its data collection settings are configured to send the relevant System log events to that workspace. A scheduled query alert can then query the collected event data over a one-hour evaluation period, count error events, and trigger when the count exceeds two.

A storage account and shared access signature do not provide the Azure Monitor Logs data source required to query Windows event records for this type of alert. SAS configuration only delegates constrained access to resources in the storage account; it does not make the account a Log Analytics workspace or enable Azure Monitor to evaluate System event-log entries from it. Microsoft documents Windows event collection through Azure Monitor agents and Log Analytics, as well as log search alert rules that run scheduled queries against a workspace. See Collect Windows and Sysmon event data sources with the Azure Monitor Agent and Log search alert rules.

QUESTION NO: 92

You have an Azure subscription that contains an Azure virtual machine named VM1. VM1 runs a financial reporting app named App1 that does not support multiple active instances.

At the end of each month, CPU usage for VM1 peaks when App1 runs.

You need to create a scheduled runbook to increase the processor performance of VM1 at the end of each month. What task should you include in the runbook?

A. Add the Azure Performance Diagnostics agent to VM1.

B. Modify the VM size property of VM1.

C. Add VM1 to a scale set.

D. Increase the vCPU quota for the subscription.

E. Add a Desired State Configuration (DSC) extension to VM1.

ANSWER: B

Explanation:

Modify the VM size property of VM1. Resizing a virtual machine changes its Azure VM size (SKU), allowing VM1 to use a size that provides additional virtual CPUs and, depending on the selected SKU, more memory and other performance characteristics. This increases compute capacity while App1 remains deployed as a single active instance, which meets the application's stated constraint. A scheduled Azure Automation runbook can use Azure PowerShell or Azure CLI commands to resize the VM at the required month-end period. In practice, the runbook should account for the VM's operational state and planned interruption: changing to a size that is not available on the current hardware cluster generally requires the VM to be stopped and deallocated before the resize, then started afterward. The target size must also be available in the VM's region, and the subscription must have sufficient quota for the selected VM family. Azure documents the supported resize process and notes when deallocation is required. Azure Automation supports scheduled runbooks, enabling the resize

workflow to execute on the required recurring schedule. See Resize a virtual machine in Azure and Schedule a runbook in Azure Automation.

QUESTION NO: 93

You have two Azure subscriptions named Sub1 and Sub2.

Sub1 contains a virtual machine named VM1 and a storage account named storage1. VM1 is associated to the resources shown in the following table.

Name	Type
Disk1	Operating system disk
NetInt1	Network interface
VNet1	Virtual network

You need to move VM1 to Sub2.

Which resources should you move to Sub2?

A. VM1, Disk1, and NetInt1 only

B. VM1, Disk1, and VNet1 only

C. VM1, Disk1, and storage1 only

D. VM1, Disk1, NetInt1, and VNet1

ANSWER: A

Explanation:

VM1, Disk1, and NetInt1 only is the correct resource set. A virtual machine move between subscriptions must include the VM and the resources that are directly dependent on it and must remain associated with it after the move. For a VM that uses a managed disk, the managed OS disk is a required dependency because it contains the operating system and boot data. The network interface is also required because it provides the VM's network attachment and configuration. Moving these resources together preserves the VM's disk and NIC resource associations in the target subscription.

The virtual network does not have to be moved merely because the NIC is connected to one of its subnets. Azure supports moving the VM, its managed disks, and its network interface while the virtual network remains in its current subscription, subject to normal subscription and tenant requirements. Likewise, a storage account is not part of this move unless it is specifically used by a resource that must move, such as a legacy unmanaged disk scenario. Microsoft's resource-move guidance identifies virtual machines, managed disks, and network interfaces as resources that should be evaluated and moved together to maintain the deployment's dependencies. See Move resources to a new resource group or subscription and Move operation support for resources.

QUESTION NO: 94

You need to prepare the environment to meet the authentication requirements.

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

NOTE: Each correct selection is worth one point.

A. Allow inbound TCP port 8080 to the domain controllers in the Miami office.

B. Add https://autologon.microsoftazuread-sso.com to the intranet zone of each client computer in the Miami office.

C. Join the client computers in the Miami office to Azure AD.

D. Install the Active Directory Federation Services (AD FS) role on a domain controller in the Miami office.

E. Install Azure AD Connect on a server in the Miami office and enable Pass-through Authentication.

ANSWER: B E

Explanation:

Seamless single sign-on can be used with pass-through authentication, and Azure AD Connect is the component used to configure both the cloud authentication method and Seamless SSO for a hybrid identity environment. Therefore, installing Azure AD Connect in the Miami office and enabling Pass-through Authentication provides the required authentication architecture. During configuration, Azure AD Connect creates the required computer account in on-premises Active Directory and configures the tenant for Seamless SSO. For resilient production authentication, additional pass-through authentication agents can also be installed on other servers.

Clients must treat the Seamless SSO endpoint as part of the local intranet so supported browsers can automatically submit the currently signed-in user's Kerberos ticket. Adding `https://autologon.microsoftazuread-sso.com` to the intranet zone, normally through Active Directory Group Policy, enables this browser behavior for the applicable domain-joined client computers. This lets users access supported Microsoft Entra ID applications without being prompted again for credentials when they are on the corporate network. Microsoft documents the intranet-zone requirement and deployment approach in Microsoft Entra Seamless SSO, and documents the associated sign-in method in pass-through authentication.

QUESTION NO: 95

You have an Azure subscription that contains the virtual networks shown in the following table.

Name	Azure region	Resource group
VNET1	West US	RG1
VNET2	Central US	RG1
VNET3	Central US	RG2
VNET4	West US	RG2

You need to deploy an Azure firewall named AF1 to RG1 in the West US Azure region.

To which virtual networks can you deploy AF1?

A. VNET1 only

B. VNET1 and VNET2 only

C. VNET1 and VNET4 only

D. VNET1, VNET2, and VNET4 only

E. VNET1, VNET2, VNET3, and VNET4

ANSWER: C

Explanation:

VNET1 and VNET4 only is correct because an Azure Firewall must be deployed into a virtual network located in the same Azure region as the firewall resource. AF1 is being created in West US, so it can be associated with either West US virtual network: VNET1 or VNET4. The resource group selected for the firewall, RG1, does not require the target virtual network to be in that same resource group. Azure resources in a subscription can be organized into different resource groups while still being used together, provided that the relevant service requirements, including regional requirements, are met.

Before deployment, the selected virtual network must also contain a dedicated subnet named `AzureFirewallSubnet`. This subnet is reserved for Azure Firewall and must meet Azure Firewall subnet sizing requirements. The firewall is deployed as a resource in RG1, while its network placement is determined by the selected West US virtual network. This allows AF1 to be deployed to VNET4 even though that virtual network is in RG2.

For deployment prerequisites and the required dedicated subnet, see Deploy and configure Azure Firewall using the Azure portal. For Azure Firewall architecture and regional deployment concepts, see What is Azure Firewall?.