DUMPSBOSS.

Microsoft Azure Security Technologies

Microsoft AZ-500

Version Demo

Total Demo Questions: 109

Total Premium Questions: 1093

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Manage identity and access	388
Topic 2, Implement platform protection	224
Topic 3, Manage security operations	269
Topic 4, Secure data and applications	187
Topic 5, Case Study 1	3
Topic 6, Case Study 2	3
Topic 7, Case Study 3	4
Topic 8, Case Study 4	3
Topic 9, Case Study 5	5
Topic 10, Case Study 6	2
Topic 11, Case Study 7	2
Topic 12, Case Study 8	3
Total	1093

QUESTION NO: 1

You are planning to deploy Azure Container Instances and have an application consisting of two containers: an application container and a validation container. The validation container monitors the application container by performing security checks via requests and waiting for responses post every transaction. It is essential to ensure that these two containers are scheduled for deployment together and that they communicate exclusively over ports which are not accessible from outside. What should you include in the deployment to achieve this?

A. application security groups

B. network security groups (NSGs)

C. management groups

D. container groups

ANSWER: D

Explanation:

Include a container group. In Azure Container Instances (ACI), a container group is the fundamental scheduling unit: all containers in the group are deployed together to the same host, start/stop together, and share lifecycle and resources. This is exactly the “sidecar” pattern described in the scenario, where a validation container must run alongside and monitor an application container.

Container groups also provide the required private, internal-only communication. Containers within the same container group share a network namespace, meaning they can communicate with each other over localhost and any ports you choose without exposing those ports externally. Only ports you explicitly publish on the container group (or on specific containers, depending on configuration) are reachable from outside; internal ports used solely between the two containers can remain unexposed, satisfying the requirement that communication occurs exclusively over ports not accessible externally.

This design is the recommended ACI approach for tightly coupled multi-container workloads such as monitoring, logging, proxying, or validation sidecars.

References: Azure Container Instances - container groups, Container groups networking.

QUESTION NO: 2

You have an Azure subscription that is associated with an Azure Active Directory (Azure AD) tenant.

From the Azure portal, you register an enterprise application. What additional resource is automatically created in Azure AD as a result?

A. a service principal

B. an X.509 certificate

C. a managed identity

D. a user account

ANSWER: A

Explanation:

When you register an enterprise application in Azure AD, Azure AD automatically creates a corresponding service principal in your tenant. The service principal is the tenant-specific representation of the application and is what actually gets used for sign-in, authorization, and assigning permissions/roles in that directory. In other words, the application object describes the app definition (often considered the “global” template), while the service principal is the local instance that enables the app to access resources in your tenant and allows administrators to configure things like single sign-on, user assignment, and conditional access. This is why enterprise applications in the Azure portal are fundamentally service principals: they are the security identity that Azure AD uses to grant the app access to resources and to manage the app’s presence in the directory. No certificate, managed identity, or user account is automatically created just by registering the enterprise application; those are separate configurations or identities created under different workflows. For more details on the relationship between application objects and service principals, see Microsoft Learn: Application objects and service principals and Microsoft Learn: How applications are added in Microsoft Entra ID.

QUESTION NO: 3

You have an Azure subscription linked to an Azure AD tenant named contoso.com. Contoso.com contains a user named User1 and an Azure web app named App1.

You plan to enable User1 to perform the following tasks:

• Configure contoso.com to use Microsoft Entra Verified ID.

• Register App1 in contoso.com.

You need to identify which roles to assign to User1. The solution must use the principle of least privilege.

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

A. Authentication Policy Administrator

B. Authentication Administrator

C. Cloud App Security Administrator

D. Application Administrator

E. User Administrator

F. Verified ID Administrator

ANSWER: D F

Explanation:

To configure a tenant to use Microsoft Entra Verified ID, the least-privilege built-in role intended for managing Verified ID settings is the Verified ID Administrator role. This role is specifically designed to let an admin set up and manage Verified ID within the tenant without granting broader directory-wide permissions that come with roles like Global Administrator.

Separately, to register an application in the tenant (create and manage app registrations such as App1 in Microsoft Entra ID), the least-privilege role is Application Administrator, which allows creating and managing app registrations and service principals.

Therefore, the correct combination is assigning Verified ID Administrator for configuring Microsoft Entra Verified ID and Application Administrator for registering App1. These roles align directly to the required tasks while minimizing unnecessary privileges.

References: Microsoft Entra built-in roles (permissions reference), Microsoft Entra Verified ID documentation.

QUESTION NO: 4

You have an Azure subscription that includes a user named User1 and an Azure Container Registry named ConReg1. You have enabled content trust for ContReg1. Your task is to ensure User1 can create trusted images in ConReg1 while adhering to the principle of least privilege. Which two roles should you assign to User1? Each correct answer contributes to the solution.

NOTE: Each correct selection is worth one point.

A. AcrQuarantineReader

B. Contributor

C. AcrPush

D. AcrImageSigner

E. AcrQuarantineWriter

ANSWER: C D

Explanation:

To create trusted images in Azure Container Registry when content trust is enabled, the user must be able to both push image artifacts to the registry and sign those images so that consumers can verify publisher authenticity. The built-in role that provides the minimum permissions required to upload (push) images and related artifacts to an Azure Container Registry is AcrPush. This role grants the data-plane permissions needed for pushing to repositories without granting broad management permissions over the registry resource itself, which aligns with least privilege.

In addition, content trust relies on image signing (Notary v1) so that signed metadata is stored and can be validated during pulls. The built-in role intended for signing operations is AcrImageSigner, which grants the necessary permissions to sign images in the registry without requiring full administrative access. Assigning both AcrPush and AcrImageSigner enables User1 to push images and produce the trusted (signed) metadata required by content trust, while avoiding overly permissive roles such as general resource management roles.

References: Content trust in Azure Container Registry, Azure Container Registry roles and permissions.

QUESTION NO: 5

You have an Azure Active Directory (Azure AD) tenant named contoso.com that contains a user named User1.

You plan to publish several apps in the tenant.

You need to ensure that User1 can grant admin consent for the published apps. Which two possible user roles can you assign to User1 to achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Security administrator

B. Cloud application administrator

C. Application administrator

D. User administrator

E. Application developer

ANSWER: B C

Explanation:

To allow a user to grant tenant-wide admin consent to application permissions (for example, consenting on behalf of the organization to delegated permissions and application permissions requested by an enterprise application), the user must be assigned an Azure AD role that includes the ability to consent to permissions. The built-in roles intended for managing app registrations and enterprise applications include the necessary permissions to grant admin consent. In practice, assigning either the Cloud application administrator role or the Application administrator role enables the user to manage application objects and enterprise applications and to grant admin consent for permissions requested by apps published in the tenant. This aligns with Microsoft guidance that these roles can perform admin consent operations without requiring the Global Administrator role, supporting least-privilege administration for app governance. After assigning one of these roles, User1 can use the admin consent workflow in the Azure portal (or equivalent APIs) to approve permissions for the organization when publishing or configuring apps.

References: Grant tenant-wide admin consent to an application, Microsoft Entra built-in roles and permissions

QUESTION NO: 6

This question is part of a series where each question presents the same scenario, but may have different solutions. You cannot return to questions once they are answered and they will not appear in the review screen.

Your Azure subscription includes 50 virtual machines running either Windows Server 2012 R2 or Windows Server 2016.

Objective: Deploy Microsoft Antimalware to the virtual machines.

Proposed Solution: Add an extension to each virtual machine to install Microsoft Antimalware.

Question: Does this solution achieve the intended goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes, adding an extension to each Azure virtual machine achieves the goal of deploying Microsoft Antimalware. Microsoft Antimalware for Azure is delivered as a VM extension (the Microsoft Antimalware extension for Windows VMs). VM extensions are the supported mechanism to install and configure post-deployment software on Azure VMs, and they can be applied per-VM (manually in the portal, via ARM/Bicep, PowerShell, CLI, or at scale using policies/automation). Once the extension is installed, it deploys the antimalware engine and enables real-time protection and scheduled scanning based on the configuration you provide (or defaults). This approach works for supported Windows Server versions such as Windows Server 2012 R2 and Windows Server 2016, matching the scenario. In practice, you would add the extension to each of the 50 VMs (or use automation to apply it across all VMs), and Azure will provision the extension and install the antimalware components on each machine, fulfilling the deployment objective.

References: Microsoft Antimalware for Azure, Azure virtual machine extensions and features

QUESTION NO: 7

You are troubleshooting a security issue for an Azure Storage account.

You enable Azure Storage Analytics logs and archive it to a storage account. What should you use to retrieve the diagnostics logs?

A. Azure Cosmos DB explorer

B. SQL query editor in Azure

C. AzCopy

D. File Explorer in Windows

ANSWER: C

Explanation:

When you enable Azure Storage Analytics logging and archive the logs to a storage account, the resulting diagnostic log files are written as blobs in a container (for classic Storage Analytics, this is typically the `$logs` container) using a structured path by service and date. To retrieve those log files for troubleshooting, you need a tool that can efficiently list and download blobs from Azure Storage. AzCopy is designed for exactly this purpose: it's a command-line utility optimized for high-performance data transfer to and from Azure Storage, including downloading log blobs from a container to a local machine for analysis. This aligns well with incident response workflows where you may need to pull a large set of log files quickly and reliably, potentially filtering by path/prefix. After downloading, you can parse the log files locally or ingest them into other tools for investigation.

References: Get started with AzCopy, Storage Analytics Logging

QUESTION NO: 8

You have an Azure subscription that contains the resources shown in the following table. You plan to deploy an Azure Private Link service named APL1.

Which resource should you reference during the creation of APL1.

Name	Type
LB1	Azure Standard Load Balancer
VM1	Virtual machine
SQL1	Azure SQL Database
VMSS1	Virtual machine scale set

A. LB1

B. SQL1

C. VMSS1

D. VM1

ANSWER: A

Explanation:

When you create an Azure Private Link service (the provider side of Private Link), you must reference a Standard Load Balancer that fronts the service you want to publish privately. The Private Link service is associated to the load balancer's frontend IP configuration, and the load balancer then distributes traffic to the backend pool (for example, a VM scale set or VMs). This is why the correct resource to reference during creation is the load balancer resource rather than the compute instances or a PaaS resource like Azure SQL Database. In practice, you deploy your service behind an internal Standard Load Balancer, then create the Private Link service pointing at that load balancer so consumers can create private endpoints to it. This is the required architecture for publishing your own service via Private Link.

References: https://learn.microsoft.com/en-us/azure/private-link/private-link-service-overview, https://learn.microsoft.com/en-us/azure/private-link/create-private-link-service-portal

QUESTION NO: 9 - (SIMULATION)

SIMULATION

You need to prevent administrative users from accidentally deleting a virtual network named VNET1. The administrative users must be allowed to modify the settings of VNET1.

To complete this task, sign in to the Azure portal.

ANSWER: See the explanation for the answer

Explanation:

Locking prevents other users in your organization from accidentally deleting or modifying critical resources, such as Azure subscription, resource group, or resource.

Note: In Azure, the term resource refers to an entity managed by Azure. For example, virtual machines, virtual networks, and storage accounts are all referred to as Azure resources.

1. In the Settings blade for virtual network VNET, select Locks.

2. To add a lock, select Add.

3. For Lock type select Delete lock, and click OK

Reference: https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-lock-resources

QUESTION NO: 10

You have an Azure subscription that includes four Azure SQL managed instances. You need to assess the vulnerability of these managed instances to SQL injection attacks. What action should you take first?

A. Create an Azure Sentinel workspace.

B. Enable Advanced Data Security.

C. Add the SQL Health Check solution to Azure Monitor.

D. Create an Azure Advanced Threat Protection (ATP) instance.

ANSWER: B

Explanation:

Enable Advanced Data Security is the right first step because it turns on the Microsoft Defender for SQL capabilities for Azure SQL (including SQL Managed Instance), which include Vulnerability Assessment and Advanced Threat Protection features. Vulnerability Assessment helps you discover, track, and remediate potential database vulnerabilities and misconfigurations by running scans and producing findings and remediation guidance. While SQL injection is primarily an application-layer risk, Defender for SQL's assessment and threat protection features are the Azure-native way to evaluate and monitor for suspicious database activities and common weakness patterns that can increase exposure to injection-style attacks (for example, excessive permissions, unsafe configurations, and anomalous query behavior). Enabling it centrally is also the prerequisite before you can run vulnerability scans and start collecting the relevant security signals across all managed instances in the subscription. After enabling, you would configure and run Vulnerability Assessment scans on each managed instance and review the findings to reduce attack surface and improve detection. See Microsoft Defender for SQL and SQL Vulnerability Assessment.

QUESTION NO: 11

You have an Azure Container Registry named Registry1.

From Azure Security Center, you enable Azure Container Registry vulnerability scanning of the images in Registry1.

You perform the following actions:

- Push a Windows image named Image1 to Registry1.
- Push a Linux image named Image2 to Registry1.
- Push a Windows image named Image3 to Registry1.
- Modify Image1 and push the new image as Image4 to Registry1. Modify Image2 and push the new image as Image5 to Registry1.

Which two images will be scanned for vulnerabilities? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Image4

B. Image2

C. Image1

D. Image3

E. Image5

ANSWER: B E

Explanation:

With Microsoft Defender for Cloud (formerly Azure Security Center) container registry vulnerability assessment enabled for an Azure Container Registry, the service assesses images that are pushed to the registry and that are supported by the underlying scanner. In this scenario, the key support boundary is the operating system: vulnerability assessment for Azure Container Registry images applies to Linux-based container images, while Windows container images aren't scanned by this capability. Therefore, the Linux images that were pushed (including the updated Linux image pushed under a new name) are the ones that will be scanned for vulnerabilities. Concretely, the initial Linux image push results in a scan, and when that Linux image is modified and pushed again as a new image, that new Linux image is also scanned. This aligns with Defender for Cloud's container registry integration behavior, where each newly pushed supported image is assessed and findings are surfaced in Defender for Cloud recommendations.

References: Microsoft Docs: Defender for container registries, Microsoft Docs: Microsoft Defender for Cloud overview

QUESTION NO: 12

You are collecting events from Azure virtual machines to an Azure Log Analytics workspace.

You plan to create alerts based on the collected events.

You need to identify which Azure services can be used to create the alerts.

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

NOTE: Each correct selection is worth one point.

A. Azure Monitor

B. Azure Security Center

C. Azure Analysis Services

D. Azure Sentinel

E. Azure Advisor

ANSWER: A D

Explanation:

Azure Monitor is a correct choice because it provides native alerting over data stored in a Log Analytics workspace. When VM events are collected into the workspace (for example, via the Azure Monitor agent and data collection rules), you can create log search (KQL) alert rules that run on a schedule and trigger notifications or actions through action groups. This is the standard platform service for metric and log-based alerting in Azure, including alerts sourced from Log Analytics.

Azure Sentinel is also a correct choice because it is built on top of Log Analytics and uses the same workspace data to create analytics rules. Sentinel analytics rules can generate incidents and alerts based on KQL queries over collected events, which fits the requirement to alert on events ingested from Azure virtual machines. In security operations scenarios, Sentinel is commonly used to create detections and alerting workflows from Log Analytics data.

References: https://learn.microsoft.com/azure/azure-monitor/alerts/alerts-overview, https://learn.microsoft.com/azure/sentinel/detect-threats-built-in

QUESTION NO: 13

You have a Microsoft Entra tenant that uses Microsoft Entra Permissions Management and contains the accounts shown in the following table:

Name	Role
Admin1	Global Administrator
Admin2	Privileged Role Administrator
Admin3	Privileged Authentication Administrator
Admin4	Exchange Administrator

Which accounts will be listed as assigned to highly privileged roles on the Azure AD insights tab in the Entra Permissions Management portal?

A. Admin1 only

B. Admin2 and Admin3 only

C. Admin2 and Admin4 only

D. Admin1, Admin2, and Admin3 only

E. Admin2, Admin3, and Admin4 only

F. Admin1, Admin2, Admin3, and Admin4

ANSWER: C

Explanation:

In Microsoft Entra Permissions Management, the Azure AD insights view highlights identities that are assigned to highly privileged Microsoft Entra ID (Azure AD) directory roles. This includes both permanently assigned roles and roles assigned through Privileged Identity Management (PIM) where the user is eligible/assigned for activation, because those assignments represent potential high-impact access that should be governed and reviewed. The “highly privileged roles” category focuses on roles that can manage identity, security, and access broadly across the tenant (for example, Global Administrator, Privileged Role Administrator, and similar top-tier roles). Based on the table in the prompt, the accounts that meet the criteria for being assigned to highly privileged roles are the ones holding those high-impact directory role assignments, which are Admin2 and Admin4. These will therefore appear as assigned to highly privileged roles on the Azure AD insights tab, enabling you to prioritize remediation actions such as enforcing just-in-time access, reducing standing privileges, and applying least privilege. For more details on how Entra Permissions Management surfaces identity insights and privileged access risk, see Microsoft Entra Permissions Management documentation and the overview of privileged roles and governance in Entra ID at Entra ID role-based access control.

QUESTION NO: 14

You have an Azure subscription that is linked to an Azure Active Directory (Azure AD) tenant.

From the Azure portal, you register an enterprise application. Which additional resource will be created in Azure AD?

A. a service principal

B. an X.509 certificate

C. a managed identity

D. a user account

ANSWER: A

Explanation:

When you register (add) an enterprise application in Azure AD, Azure AD creates a corresponding service principal in your tenant. The service principal is the identity instance of the application within a specific directory and is what Azure AD uses for authentication, authorization, and assigning permissions/roles for that application in the tenant. In practical terms, the service principal is the security principal you manage for tenant-specific settings such as who can sign in, what permissions are granted, and what conditional access policies apply. This is why enterprise applications are closely associated with service principals: the “application object” describes the app definition, while the “service principal” represents the app’s presence and identity in a given tenant. Creating an enterprise application in the Azure portal results in that service principal being created so the app can be used and governed within the directory.

References: How applications are added to Microsoft Entra ID, Application and service principal objects

QUESTION NO: 15

You have an Azure AD tenant that contains three users named User1, User2, and User3.

You configure Azure AD Password Protection as shown in the following exhibit. The users perform the following tasks:

• User1 attempts to reset her password to C0nt0s0.

• User2 attempts to reset her password to F@brikamHQ.

• User3 attempts to reset her password to Pr0duct123. Which password reset attempts fail?

A. User1 only

B. User2 only

C. User3 only

D. User1 and User 3 only

E. User1, User2, and User3

ANSWER: E

Explanation:

All three password reset attempts fail because Azure AD Password Protection evaluates proposed passwords against both the global banned password list (maintained by Microsoft) and any custom banned password list you configure for the tenant, and then enforces the configured mode. In the exhibit, Azure AD Password Protection is enabled and set to enforce

(not just audit), meaning a password that matches or is too similar to banned terms will be rejected during password change/reset. The attempted passwords are all based on common company/brand-style strings with predictable substitutions and suffixes (for example, “Contoso”, “Fabrikam”, and “Product” patterns), which Azure AD Password Protection is designed to detect using normalization (case-insensitivity, common character substitutions like 0→o, @→a, and appended digits). With enforcement enabled, each of these proposed passwords is considered weak and is blocked, causing each reset attempt to fail. This behavior aligns with Microsoft’s description of how Azure AD Password Protection blocks easily guessable passwords by checking normalized variants against banned lists and applying the tenant’s enforcement setting.

https://learn.microsoft.com/en-us/entra/identity/authentication/concept-password-ban-bad
https://learn.microsoft.com/en-us/entra/identity/authentication/howto-password-ban-bad-on-premises-operations

QUESTION NO: 16

You have 15 Azure virtual machines in a resource group named RG1. All the virtual machines run identical applications.

You need to prevent unauthorized applications and malware from running on the virtual machines.

What should you do?

A. Apply an Azure policy to RG1.

B. From Azure Security Center, configure adaptive application controls.

C. Configure Azure Active Directory (Azure AD) Identity Protection.

D. Apply a resource lock to RG1.

ANSWER: B

Explanation:

From Microsoft Defender for Cloud (formerly Azure Security Center), configuring adaptive application controls is the right approach because it provides application allowlisting (whitelisting) for VMs. Adaptive application controls learns the known-good processes and applications that typically run on your virtual machines and then helps you enforce rules so that only those approved applications can execute. This directly addresses the requirement to prevent unauthorized applications and malware from running, especially in a scenario where all VMs run identical applications—making it well-suited for a consistent allowlist policy across the fleet. The feature is designed to harden both Windows and Linux VMs by reducing the attack surface and blocking unexpected executables, which is a common malware technique. Once enabled and rules are applied, Defender for Cloud can alert on and help block deviations from the approved application set, providing ongoing protection aligned with Azure security best practices.

References: https://learn.microsoft.com/en-us/azure/defender-for-cloud/adaptive-application-controls, https://learn.microsoft.com/en-us/azure/defender-for-cloud/defender-for-cloud-introduction

QUESTION NO: 17

You have an Azure subscription that includes a web application named App1. Users must have the option to choose either a Google identity or a Microsoft identity for authentication to App1. Which two pieces of information must you configure to add Google as an identity provider in Azure Active Directory? Select two options, each correct choice is worth one point.

A. a tenant name

B. a tenant ID

C. the endpoint URL Of an application

D. a client ID

E. a client secret

ANSWER: D E

Explanation:

To add Google as an identity provider for sign-in (so users can choose Google or Microsoft identities), you must register an OAuth 2.0 client application in Google (Google Cloud Console) and then configure Azure with the credentials that Google issues for that app. The two required values are the application (OAuth) client identifier and the corresponding client secret. Azure uses the client ID to identify the relying party application during the OAuth/OpenID Connect flow, and it uses the client secret to authenticate the application when exchanging authorization codes for tokens. Without both values, Azure can't complete the token exchange and validate the sign-in flow from Google. This is the standard configuration pattern for adding Google as an external identity provider in Azure App Service authentication and also aligns with how external identity providers are configured for Azure AD/B2C-style federated sign-in experiences.

References: Microsoft Docs: Configure Google authentication for App Service, Microsoft Docs: OAuth 2.0 authorization code flow

QUESTION NO: 18

You have an Azure subscription.

You create an Azure web app named Contoso1812 that uses an S1 App Service plan.

You plan to

create a CNAME DNS record for www.contoso.com that points to Contoso1812.

You need to ensure that users can access Contoso1812 by using the https://www.contoso.com URL.

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

NOTE: Each correct selection is worth one point.

A. Turn on the system-assigned managed identity for Contoso1812.

B. Add a hostname to Contoso1812.

C. Scale out the App Service plan of Contoso1812.

D. Add a deployment slot to Contoso1812.

E. Scale up the App Service plan of Contoso1812.

F. Upload a PFX file to Contoso1812.

ANSWER: B F

Explanation:

To make an App Service web app reachable at https://www.contoso.com, you must first map the custom DNS name to the app and then bind an SSL/TLS certificate to that hostname. Adding a hostname to Contoso1812 is required because App Service won't serve traffic for a custom domain until the domain is added and validated (typically via a CNAME plus a verification record such as a TXT record). Once the hostname is added, HTTPS requires a certificate that App Service can present for that specific custom domain. Uploading a PFX file to Contoso1812 provides the private key and certificate chain needed to create an App Service certificate binding (SNI SSL) for the custom hostname. Because the app is on an S1 plan, it supports custom domains and TLS/SSL bindings, so no scaling actions are necessary for HTTPS to work. After the hostname is mapped and the PFX is uploaded, you complete the SSL binding to the custom domain so users can browse securely using the https://www.contoso.com URL.

References: Map an existing custom DNS name to Azure App Service, Add a TLS/SSL certificate in Azure App Service

QUESTION NO: 19 - (SIMULATION)

SIMULATION

You need to ensure that the events in the NetworkSecurityGroupRuleCounter log of the VNET01-Subnet0-NSG network security group (NSG) are stored in the logs1234578 Azure Storage account for 30 days.

To complete this task, sign in to the Azure portal.

ANSWER: See the explanation for the answer

Explanation:

You need to configure the diagnostic logging for the NetworkSecurityGroupRuleCounter log.

1. In the Azure portal, type Network Security Groups in the search box, select Network Security Groups from the search results then select VNET01

Subnet0-NSG. Alternatively, browse to Network Security Groups in the left navigation pane.

2. In the properties of the Network Security Group, click on Diagnostic Settings.

3. Click on the Add diagnostic setting link.

4. Provide a name in the Diagnostic settings name field. It doesn't matter what name you provide for the exam.

5. In the Log section, select NetworkSecurityGroupRuleCounter.

6. In the Destination details section, select Archive to a storage account.

7. In the Storage account field, select the logs1234578 storage account.

8. In the Retention (days) field, enter 30.

9. Click the Save button to save the changes.

QUESTION NO: 20

You have an Azure subscription that contains a virtual network named VNet1. VNet1 contains a single subnet. The subscription contains a virtual machine named VM1 that is connected to VNet1.

You plan to deploy an Azure SQL managed instance named SQL1. You need to ensure that VM1 can access SQL1.

Which three components should you create? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a subnet

B. a network security perimeter

C. a virtual network gateway

D. a network security group (NSG)

E. a route table

ANSWER: A D E

Explanation:

To deploy Azure SQL Managed Instance, you must place the managed instance into a dedicated subnet in a virtual network. Because VNet1 currently has only a single subnet (already used by VM1), you should create a subnet specifically for SQL1. Connectivity from VM1 to SQL1 is then private within the virtual network, but you still need to control and allow the required traffic. A network security group (NSG) is used to permit the necessary inbound/outbound flows for SQL Managed Instance and to ensure VM1 can reach SQL1 on the required ports while maintaining least privilege. In addition, a route table (user-defined routes) is commonly required/used with the managed instance subnet to ensure correct routing behavior and to meet SQL Managed Instance networking requirements (for example, controlling next hops and avoiding unintended routing through appliances that could break required service traffic). Together, creating a dedicated subnet, associating an NSG, and associating a route table provides the standard network components needed for a functional and secure SQL Managed Instance deployment that VM1 can access privately.

References: Azure SQL Managed Instance connectivity architecture, Create and configure a subnet for Azure SQL Managed Instance

QUESTION NO: 21

You have an Azure subscription that contains a web app named App1.

Users must be able to select between a Google identity or a Microsoft identity when authenticating to App1.

You need to add Google as an identity provider in Azure AD.

Which two pieces of information should you configure? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a client ID

B. a tenant name

C. the endpoint URL of an application

D. a tenant ID

E. a client secret

ANSWER: A E

Explanation:

To add Google as an external identity provider for Azure AD (Microsoft Entra ID) so users can choose Google or a Microsoft identity at sign-in, you must register an OAuth 2.0 client application in the Google API Console and then configure that provider in Entra ID. The required configuration values are the OAuth client ID and the corresponding client secret issued by Google. Entra ID uses the client ID to identify your app to Google's authorization endpoint and uses the client secret to authenticate your app when exchanging the authorization code for tokens. These two values are the core credentials for the OpenID Connect/OAuth trust between Entra ID and Google. Tenant name/tenant ID are concepts for Microsoft directories and aren't required to establish the Google federation. Likewise, you don't manually supply an "endpoint URL of an application" for Google in this scenario because Entra ID already knows the standard Google OpenID Connect endpoints; you primarily provide the client credentials and any required redirect URI during Google app registration. For details, see Microsoft's guidance on adding Google as an identity provider and the required client credentials: https://learn.microsoft.com/en-us/entra/external-id/google-federation and the identity provider setup overview: https://learn.microsoft.com/en-us/entra/external-id/identity-providers.

QUESTION NO: 22

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

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

You have an Azure subscription named Sub1, which includes an Azure Storage account called sa1 located in a resource group named RG1. Both users and applications access the blob service and the file service within sa1 utilizing multiple shared access signatures (SASs) and stored access policies.

It has come to your attention that unauthorized users have gained access to both the file service and the blob service. You need to revoke all access to sa1. One solution considered is creating a lock on sa1.

Does this solution meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because an Azure resource lock (ReadOnly or CanNotDelete) is an Azure Resource Manager control that prevents management-plane changes such as deleting or modifying the storage account resource. A lock does not affect the data plane for Azure Storage, meaning it won't invalidate existing shared access signatures (SAS), won't block access via storage account keys, and won't stop clients from continuing to read/write blobs or files if they already have valid credentials. Since the scenario requires revoking all access to the blob and file services, you must use data-plane revocation mechanisms such as rotating/regenerating storage account keys (to invalidate any SAS signed with the account key) and/or deleting or changing stored access policies (to invalidate SAS that reference those policies). Resource locks are useful for preventing accidental administrative changes, but they are not a security revocation tool for Storage access.

References: https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/lock-resources, https://learn.microsoft.com/en-us/azure/storage/common/storage-sas-overview

QUESTION NO: 23

You have an Azure subscription. The subscription contains a virtual network named VNet1 that contains the subnets shown in the following table.

The subscription contains the function apps shown in the following table. The outbound traffic of which app is controlled by using NSG1?

Name	Associated network security group (NSG)
Subnet1	NSG1
Subnet2	NSG1
Subnet3	NSG1
Subnet4	NSG1

Name	Description
App1	Uses the Azure Functions Premium plan and has virtual network integration with VNet1/Subnet1
App2	Uses an App Service plan in the Basic pricing tier and has virtual network integration with VNet1/Subnet2
App3	Uses an App Service plan in the Premium pricing tier and has virtual network integration with VNet1/Subnet3
App4	Uses an App Service plan in the isolated pricing tier and is deployed to VNet1/Subnet4

A. App4 only

B. App3 and App4 only

C. App2, App3, and App4 only

D. App1, App2, App3, and App4

ANSWER: C

Explanation:

Outbound traffic is controlled by a subnet's network security group only when the resource's outbound flows actually traverse that subnet. For Azure Functions, this happens when the function app is integrated with a virtual network using Regional VNet Integration, which places the app's outbound traffic onto a delegated integration subnet (typically delegated to Microsoft.Web/serverFarms). Once integrated, the function app's outbound connections to resources in the VNet (and, if configured, to on-premises via VPN/ExpressRoute) are subject to the NSG rules applied to that integration subnet. Therefore, the apps whose configuration indicates they use the subnet associated with NSG1 will have their outbound traffic

governed by NSG1. This is the key distinction: NSGs don't apply to a function app unless its outbound path is injected into a subnet via VNet Integration; simply having a function app in the same subscription or region doesn't make NSG rules apply. Microsoft documents that VNet Integration enables outbound traffic from the app to flow through the selected subnet and that NSG/UDR on that subnet can affect the app's outbound connectivity.

References: Azure Functions networking options, Integrate your app with an Azure virtual network

QUESTION NO: 24

You have an Azure subscription containing 100 virtual machines with Azure Defender enabled.

Your plan is to conduct a vulnerability scan on each virtual machine.

You aim to deploy the vulnerability scanner extension to these virtual machines utilizing an Azure Resource Manager template.

Which two specific values must you include in the template code to ensure the automated deployment of the extension to the virtual machines?

NOTE: Each correct selection is worth one point.

A. the user-assigned managed identity

B. the workspace ID

C. the Azure Active Directory (Azure AD) ID

D. the Key Vault managed storage account key

E. the system-assigned managed identity

F. the primary shared key

ANSWER: B F

Explanation:

For Microsoft Defender for Cloud's integrated vulnerability assessment on Azure VMs (the Qualys-based solution), deploying the vulnerability assessment extension at scale with an ARM template requires you to provide the extension with the information it needs to register and report findings back to Defender for Cloud. In practice, that means configuring the extension to send its data to the Log Analytics workspace used by Defender for Cloud. The template therefore must include the workspace ID (the Log Analytics workspace customerId) and the primary shared key (the workspace sharedKey) so the agent/extension can authenticate to the workspace and ingest vulnerability assessment data. These values are commonly passed in the extension's protectedSettings/settings when deploying security/monitoring extensions via ARM. This aligns with how Azure VM extensions that onboard to Log Analytics are automated, and it's the critical requirement to make deployment and reporting fully automated across many VMs.

References: Deploy vulnerability assessment solutions on Azure VMs, Log Analytics agent overview (workspace ID and key).

QUESTION NO: 25

You have an Azure subscription that contains an Azure SQL database named SQL1 and an Azure key vault named KeyVault1. KeyVault1 stores the keys shown in the following table.

You need to configure Transparent Data Encryption (TDE). TDE will use a customer-managed key for SQL1.

Which keys can you use?

Name	Type	RSA key size	Elliptic curve name
Key1	RSA	2048	Not applicable
Key2	RSA	3072	Not applicable
Key3	RSA	4096	Not applicable
Key4	EC	Not applicable	P-512

A. Key2 only

B. Key1 only

C. Key2 and Key3 only

D. Key1, Key2, Key3, and Key4

E. Key1 and Key2 only

ANSWER: E

Explanation:

To use a customer-managed key (BYOK) for Transparent Data Encryption in Azure SQL Database, the TDE protector must be an asymmetric key stored in Azure Key Vault. Specifically, Azure SQL supports RSA keys (software-protected) and RSA-HSM keys (HSM-protected) as the TDE protector. In addition, the key size must be one of the supported RSA lengths for TDE integration (commonly 2048-bit or 3072-bit). Symmetric keys (such as AES) and secrets/certificates aren't valid as the TDE protector for Azure SQL Database because the service needs an asymmetric Key Vault key to wrap/unwrap the database encryption key (DEK) and to integrate with Key Vault's key operations and access control model.

Therefore, the usable keys are the ones that are asymmetric RSA-based Key Vault keys with a supported size, which corresponds to "Key1 and Key2 only".

References: https://learn.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-byok-overview, https://learn.microsoft.com/en-us/azure/key-vault/keys/about-keys

QUESTION NO: 26

You manage an Azure Active Directory (Azure AD) tenant. There are several deleted objects listed in the table below:

On May 4, 2020, you try to restore the deleted objects using the Azure Active Directory admin center. Which two objects can you successfully restore? Each correct answer is worth one point. Please select two options.

A. Group1

B. Group2

C. User2

D. User1

ANSWER: B C

Explanation:

In Microsoft Entra ID (Azure AD), many directory objects are “soft deleted” and can be restored from the Deleted users or Deleted groups blades within a limited retention window. Deleted users are typically recoverable for 30 days before they’re permanently removed, assuming they haven’t passed the soft-delete retention period. Similarly, Microsoft 365 groups (formerly Office 365 groups) are also soft deleted and can be restored within 30 days, which restores the group and associated resources (like the group mailbox and SharePoint site) as part of the group restore process.

Therefore, the objects that can be successfully restored on May 4, 2020 are the ones that are within the 30-day restore window and are of a restorable type (user objects and Microsoft 365 groups). This aligns with the documented behavior for restoring deleted Microsoft 365 groups and the general 30-day deleted user restore period in Entra ID.

References: Restore a deleted Microsoft 365 group, Restore or remove a recently deleted user

QUESTION NO: 27

You have an Azure subscription that contains an Azure Key Vault and an Azure Storage account. The Key Vault holds customer-managed keys, and the Storage account is configured to utilize these keys stored in the Key Vault.

You plan to store data in Azure using the following services:

Azure Files
Azure Blob Storage
Azure Table Storage
Azure Queue Storage

Which two services support data encryption with the keys stored in the Key Vault? Each correct answer presents a complete solution.

Note: Each correct selection is worth one point.

A. Table storage

B. Azure Files

C. Blob storage

D. Queue storage

ANSWER: B C

Explanation:

Azure Storage encryption at rest supports using customer-managed keys (CMK) stored in Azure Key Vault (or Managed HSM) for supported storage services. When you configure a storage account to use CMK, Azure Storage uses those keys to protect the data encryption keys that encrypt your stored data, giving you control over key rotation, access policies, and key lifecycle management in Key Vault.

Within the listed services, Azure Blob Storage supports CMK integration with Key Vault for service-side encryption, allowing you to meet regulatory and organizational requirements that mandate customer control of encryption keys. Azure Files also supports CMK with Key Vault for encryption at rest, enabling similar key control for file shares stored in the storage account.

These capabilities are configured at the storage account level (encryption settings) and rely on granting the storage account's identity appropriate permissions to the Key Vault key. For details on which Azure Storage services support customer-managed keys and how the integration works, see Customer-managed keys for Azure Storage encryption and Azure Storage encryption for data at rest.

QUESTION NO: 28 - (DRAG DROP)

DRAG DROP

You have an Azure subscription that contains a Microsoft SQL server named Server1 and an Azure key vault named vault1. Server1 hosts a database named DB1. Vault1 contains an encryption key named key1.

You need to ensure that you can enable Transparent Data Encryption (TDE) on DB1 by using key1.

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.

Select and Place:

ANSWER:

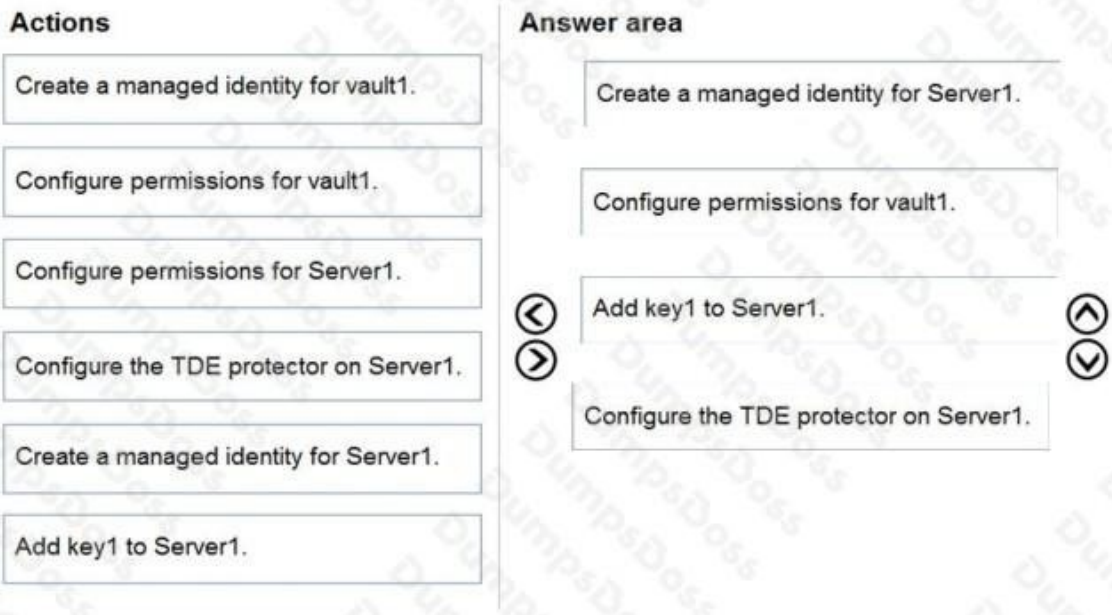

Explanation:

To enable Transparent Data Encryption with a customer-managed key for an Azure SQL Database, the SQL logical server must be able to authenticate to Azure Key Vault and then be authorized to use the key. The standard pattern is to enable a managed identity on the SQL server, then grant that identity the required Key Vault key permissions, and only then associate the Key Vault key with the server and set it as the TDE protector.

First, you create (enable) a managed identity for the SQL server so Azure SQL has an identity in Microsoft Entra ID that can be used to access Key Vault without storing secrets. Next, you configure permissions on the Key Vault (vault1) for that managed identity. Those permissions must allow the server to retrieve the key and perform cryptographic operations needed for TDE, typically including get, wrapKey, and unwrapKey on the key. Once Key Vault authorization is in place, you can add (register) key1 to Server1, which effectively tells Azure SQL where the key lives and establishes the server-level association to that Key Vault key. Finally, you configure the TDE protector on Server1 to use that key; the server-level TDE protector is what databases on the server (including DB1) will use when TDE is enabled with a customer-managed key.

This sequence matters because adding the key and setting the TDE protector will fail if the server identity doesn't already have access to the key in Key Vault. Microsoft's BYOK/TDE guidance follows this flow: enable managed identity on the SQL server, grant Key Vault permissions, then set the server's TDE protector to the Key Vault key. See: Configure TDE with customer-managed key (BYOK) for Azure SQL Database and Key Vault access control concepts at Assign a Key Vault access policy.

QUESTION NO: 29

You have multiple development teams that will create apps in Azure.

You plan to create a standard development environment that will be deployed for each team.

You need to recommend a solution that will enforce resource locks across the development environments and ensure that the locks are applied in a consistent manner.

What should you include in the recommendation?

A. an Azure policy

B. an Azure Resource Manager template

C. a management group

D. an Azure blueprint

ANSWER: D

Explanation:

Including

an Azure blueprint

is the right recommendation because Azure Blueprints are designed to define and repeatedly deploy a standardized set of Azure resources and governance controls across multiple environments in a consistent, auditable way. A blueprint can package artifacts such as ARM templates, Azure Policy assignments, role assignments, and importantly, resource locks. When resource locks are configured as part of a blueprint assignment, they're applied automatically and consistently to the target scope (typically a subscription) each time the blueprint is assigned, which matches the requirement to enforce locks across each team's standard development environment. Blueprints also support centralized lifecycle management (versioning and updates), helping ensure that every team environment stays aligned with the same baseline controls over time. This makes blueprints particularly suitable for "environment stamping" scenarios where governance must be embedded into the deployment process rather than applied manually after the fact.

References: Azure Blueprints resource locking, Azure Blueprints overview

QUESTION NO: 30

You have an Azure Sentinel workspace. You need to create a playbook.

Which two triggers will start the playbook? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. An Azure Sentinel scheduled query rule is executed.

B. An Azure Sentinel data connector is added.

C. An Azure Sentinel alert is generated.

D. An Azure Sentinel hunting query result is returned.

E. An Azure Sentinel incident is created.

ANSWER: C E

Explanation:

In Microsoft Sentinel, playbooks are Azure Logic Apps that run in response to specific Sentinel events. The most common and supported Sentinel-native triggers are tied to security detections and case management: when an alert is created and when an incident is created. These triggers let you automate response actions such as enriching entities, notifying teams, opening tickets, or containing threats. For example, an "alert is generated" trigger can kick off enrichment and triage steps immediately after a detection fires, while an "incident is created" trigger is ideal for orchestrating end-to-end incident

response workflows (assignment, notifications, ticketing, and remediation) once Sentinel groups alerts into an incident. These are the two standard ways Sentinel initiates automation from within the product's analytics and incident pipeline, and they map directly to the Sentinel connectors/triggers available in Logic Apps for automation rules and playbook execution. See the Sentinel playbook tutorial and automation documentation for the supported triggers and how they're used in incident/alert-driven response workflows: https://learn.microsoft.com/en-us/azure/sentinel/tutorial-respond-threats-playbook and https://learn.microsoft.com/en-us/azure/sentinel/automation-rules.

QUESTION NO: 31

You have an Azure Sentinel workspace, and you need to create a playbook. Which two triggers can initiate the playbook? Each correct answer provides a complete solution. NOTE: Each correct selection is worth one point.

A. An Azure Sentinel scheduled query rule is executed.

B. An Azure Sentinel data connector is added.

C. An Azure Sentinel alert is generated.

D. An Azure Sentinel hunting query result is returned.

E. An Azure Sentinel incident is created.

ANSWER: C E

Explanation:

In Microsoft Sentinel, playbooks are Azure Logic Apps that can be started by Sentinel-specific triggers. Two common built-in triggers are when an alert is generated and when an incident is created. The "Microsoft Sentinel alert" trigger allows automation to run as soon as an analytics rule produces an alert, enabling immediate enrichment, notification, or remediation steps. The "Microsoft Sentinel incident" trigger starts the playbook when Sentinel creates (or updates) an incident, which is useful for incident-centric workflows such as assigning ownership, adding comments, enriching entities, or opening tickets in ITSM tools. These triggers align with how Sentinel automation is designed: analytics rules generate alerts, Sentinel correlates alerts into incidents, and playbooks can be attached to automation rules to run on incident creation (and related incident events) or invoked from alert/incident contexts. This is the documented and supported way to initiate Sentinel playbooks for response automation.

References: Tutorial: Use playbooks to respond to threats in Microsoft Sentinel, Automate threat response with playbooks in Microsoft Sentinel

QUESTION NO: 32

Your organization utilizes Azure DevOps with established branch policies. Which of the following statements accurately describe branch policies? (Select all that apply.)

A. It enforces your team's change management standards.

B. It controls who can read and update the code in a branch.

C. It enforces your team's code quality.

D. It places a branch into a read-only state.

ANSWER: A C

Explanation:

Branch policies in Azure Repos are designed to protect important branches (such as main) by enforcing consistent quality gates and governance on changes before they're merged. They are primarily applied to pull requests and can require specific conditions to be met, such as a minimum number of reviewers, successful build validation, linked work items, or passing status checks. This is why it's accurate to say branch policies enforce a team's change management standards: they standardize the review/approval workflow and ensure changes follow agreed processes before integration. It's also accurate that branch policies enforce code quality, because they can mandate automated builds, tests, and other validations, and can require peer review to catch issues early. In practice, these policies help prevent direct, unreviewed changes from being merged into protected branches and ensure that only changes meeting defined criteria are allowed through the pull request process. For details on what branch policies do and how they're used to enforce quality and governance, see Azure DevOps branch policies and related guidance on pull requests in Azure Repos.

QUESTION NO: 33

You have an Azure subscription with an Azure SQL Database named **sql1**. You are planning to set up auditing for **sql1**.

Your objective is to configure the audit log destination while meeting the following criteria:

Enable querying events using the Kusto query language.
Minimize administrative effort.

What configuration should you implement?

A. an event hub

B. a storage account

C. a Log Analytics workspace

ANSWER: C

Explanation:

Configure auditing to send logs to a Log Analytics workspace. Azure Monitor Logs stores data in a Log Analytics workspace and is queried using the Kusto Query Language (KQL), which directly satisfies the requirement to query audit events with Kusto. This option also minimizes administrative effort because it's a managed, integrated destination: you can use built-in Azure Monitor experiences (Log Analytics queries, alert rules, workbooks, and dashboards) without having to build and operate your own ingestion, storage lifecycle, or query tooling. In contrast to destinations that primarily focus on raw retention or streaming, a Log Analytics workspace is designed for interactive investigation and operational monitoring, making it the most straightforward way to search and analyze SQL auditing events at scale using KQL. Azure SQL auditing supports writing to Azure Monitor Logs (Log Analytics) as a first-class destination, enabling centralized monitoring across resources and easier operationalization through Azure Monitor features. See Azure SQL auditing overview and Log Analytics workspace overview.

QUESTION NO: 34

This question is part of a series that presents a scenario. Each question in the series contains a unique solution that might meet the stated goals. Some series may have more than one correct solution, while others might not have a correct solution.

Once you answer a question in this section, you cannot go back to it. As such, these questions will not appear in the review screen.

You are using Azure Security Center for centralized policy management across three Azure subscriptions. You have several policy definitions intended to manage these subscriptions' security.

Your task is to deploy these policy definitions as a group to all three subscriptions.

Solution: Create a policy initiative and an assignment scoped to the Tenant Root Group management group.

Does this solution meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes, this meets the goal. An Azure Policy initiative (also called a policy set) is specifically designed to group multiple policy definitions into a single, manageable unit so you can deploy and track them together. By creating an initiative that contains the required security-related policy definitions, you can then create one assignment for the entire group rather than assigning each policy definition individually.

Scoping the assignment to the Tenant Root Group management group applies the initiative to all subscriptions that are under that management group hierarchy. This is a common pattern for centralized governance because management group scope enables consistent policy enforcement across multiple subscriptions, including future subscriptions added under the same management group. This approach aligns with how Azure Policy is intended to be used for enterprise-wide policy management and is compatible with Microsoft Defender for Cloud (formerly Azure Security Center) scenarios where policy drives security posture and recommendations.

References: https://learn.microsoft.com/en-us/azure/governance/policy/overview, https://learn.microsoft.com/en-us/azure/governance/management-groups/overview

QUESTION NO: 35

You have an Azure subscription and you reconfigure it to use a different Azure Active Directory (Azure AD) tenant. What are two possible outcomes of this configuration change? Select two options that describe complete solutions. Each correct selection is worth one point.

A. Role assignments at the subscription level are lost.

B. Virtual machine managed identities are lost.

C. Virtual machine disk snapshots are lost.

D. Existing Azure resources are deleted.

ANSWER: A B

Explanation:

When you change (transfer) an Azure subscription to a different Azure AD tenant, the subscription's trust relationship moves to the new directory. A key impact is that access control entries that reference security principals in the old tenant no longer resolve in the new tenant, so role assignments at the subscription scope can effectively be lost and must be recreated for users, groups, and service principals in the new directory. Another common impact is on identities that are issued by Azure AD for Azure resources. Managed identities are represented as service principals in the tenant; after the directory change, those identities can be removed or become unusable because they were created in the original tenant, which breaks workloads that depend on them (for example, VMs accessing Key Vault or Storage using their managed identity). These are expected outcomes called out in Microsoft guidance for changing a subscription's directory association and are important to plan for during tenant moves, including reassigning RBAC and reestablishing identity-dependent access. See How Azure subscriptions are associated with Azure AD and Azure RBAC overview.

QUESTION NO: 36

You manage a network with 10 virtual machines (VMs) located within a single subnet, which is protected by a single Network Security Group (NSG). Your objective is to log the network traffic of these VMs to an Azure Storage account.

What are the two actions you should take to accomplish this task? Each correct answer contributes to the solution.

NOTE: Each correct choice is worth one point.

A. Install the Network Performance Monitor solution.

B. Enable Azure Network Watcher.

C. Enable diagnostic logging for the NSG.

D. Enable NSG flow logs.

E. Create an Azure Log Analytics workspace.

ANSWER: B C D

Explanation:

To log network traffic for the VMs protected by a Network Security Group and store it in an Azure Storage account, you use NSG flow logs, a capability provided by Azure Network Watcher. Network Watcher must be enabled in the Azure region where the NSG exists because it is the service that hosts and runs the flow logging capability. Once Network Watcher is available, you then enable NSG flow logs on the specific NSG and select the target Storage account; Azure will write flow log records (5-tuple flow information and allow/deny decisions) to blobs in that Storage account. This approach captures traffic at the NSG level, which matches the scenario of multiple VMs in a subnet protected by a single NSG, and it meets the requirement to log to Storage directly. For configuration steps and prerequisites, see https://learn.microsoft.com/en-us/azure/network-watcher/nsg-flow-logs-overview and https://learn.microsoft.com/en-us/azure/network-watcher/network-watcher-nsg-flow-logging-portal.

QUESTION NO: 37

You are planning to deploy an application that will modify the properties of users in Azure Active Directory (Azure AD) utilizing Microsoft Graph. To ensure the application can access Azure AD, what is the first configuration step you should take?

A. a custom role-based access control (RBAC) role

B. an external identity

C. an Azure AD Application Proxy

D. an app registration

ANSWER: D

Explanation:

The first required step is creating an app registration because Microsoft Graph calls must be made by a security principal that Azure AD can authenticate and authorize. An app registration establishes the application's identity in Azure AD (tenant-specific or multi-tenant), which then enables you to configure authentication settings (such as redirect URIs and certificates/secrets) and, critically for modifying user properties, request and grant Microsoft Graph permissions (delegated or application permissions like User.ReadWrite.All). After registration, you can grant admin consent where required and use OAuth 2.0/OpenID Connect to obtain tokens for Microsoft Graph. Without an app registration, the application has no identity in Azure AD, cannot be issued tokens, and therefore cannot call Microsoft Graph to update user objects. This is the foundational configuration step before any permission assignment, consent, or credential setup can occur. See Microsoft's guidance on registering applications and how app registrations relate to permissions and token acquisition for Microsoft Graph access: https://learn.microsoft.com/en-us/entra/identity-platform/quickstart-register-app and https://learn.microsoft.com/en-us/graph/permissions-overview.

QUESTION NO: 38

You have an Azure subscription that contains 100 virtual machines and has Azure Defender enabled.

You plan to perform a vulnerability scan of each virtual machine.

You need to deploy the vulnerability scanner extension to the virtual machines by using an Azure Resource Manager template.

Which two values should you specify in the code to automate the deployment of the extension to the virtual machines? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. the user-assigned managed identity

B. the workspace ID

C. the Azure Active Directory (Azure AD) ID

D. the Key Vault managed storage account key

E. the system-assigned managed identity

F. the primary shared key

ANSWER: B E

Explanation:

To automate deployment of the Microsoft Defender for Cloud vulnerability assessment extension to Azure virtual machines via an ARM template, you typically deploy the Log Analytics agent (or Azure Monitor agent, depending on the solution) and the vulnerability assessment extension so findings can be collected and surfaced in Defender for Cloud. In ARM, the agent configuration requires the Log Analytics workspace identifier so the VM can connect to the correct workspace for data

ingestion and correlation. Additionally, the VM must be able to authenticate to Azure to install and operate the extension without embedding secrets in the template; using a managed identity is the recommended approach. A system-assigned managed identity is commonly used because it's created and lifecycle-managed with the VM, enabling secure, passwordless access to required Azure resources and APIs during extension operations and related Defender for Cloud workflows. This aligns with Microsoft guidance to avoid shared keys/secrets in templates and to use managed identities for Azure resource authentication. For more details on Defender for Cloud vulnerability assessment and agent/workspace integration, see Microsoft Docs and managed identity usage patterns at Microsoft Docs.

QUESTION NO: 39

You have an Azure subscription that contains as Azure key vault and an Azure Storage account. The key vault contains customer-managed keys. The storage account is configured to use the customer-managed keys stored in the key vault.

You plan to store data in Azure by using the following services:

- Azure Files
- Azure Blob storage
- Azure Table storage
- Azure Queue storage

Which two services support data encryption by using the keys stored in the key vault? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Table storage

B. Azure Files

C. Blob storage

D. Queue storage

ANSWER: B C

Explanation:

Azure Storage encryption supports customer-managed keys (CMK) stored in Azure Key Vault for specific storage services within a storage account. When you configure a storage account to use CMK, the services that support CMK can use the Key Vault key to wrap and protect the data encryption keys used for encryption at rest. In practice, Azure Blob storage supports CMK for encrypting blob data, and Azure Files supports CMK for encrypting file shares. This allows you to meet compliance requirements where you must control key lifecycle operations such as rotation, disabling, and auditing via Key Vault. The configuration is done at the storage account level, and the supported services inherit that encryption configuration. Microsoft documents CMK support for Azure Storage encryption and explicitly calls out Blob storage and Azure Files as supported services for encryption with customer-managed keys in Azure Key Vault. See Customer-managed keys for Azure Storage encryption and Azure Storage encryption for data at rest for the authoritative service support details and configuration guidance.

QUESTION NO: 40

You have an Azure subscription that includes an Azure Key Vault. You aim to configure the maximum number of days for which newly created keys will remain valid. Your solution should minimize administrative overhead. Which feature or service should you utilize to achieve this?

A. Key Vault properties

B. Azure Policy

C. Azure Purview

D. Azure Blueprints

ANSWER: B

Explanation:

Azure Policy is the right choice because it can enforce a consistent key expiration requirement across your Key Vault(s) automatically, which minimizes ongoing administrative effort. With Azure Policy, you can assign a built-in policy initiative for Key Vault (or create a custom policy) that audits or denies creation of keys that don't have an expiration date set, and you can standardize the expected lifetime by requiring an expiry to be configured at creation time. This approach scales well because it applies governance at the subscription or resource group scope, rather than relying on per-key manual configuration. While Key Vault supports key rotation policies, those policies are configured per key and don't inherently provide a single "maximum validity days for all newly created keys" setting across the environment. Azure Policy is designed specifically for enforcing organizational standards and compliance at scale, including Key Vault key settings such as requiring expiration dates. See Azure Policy overview and Azure Policy built-in definitions for Azure Key Vault for details on governing Key Vault with policy.

QUESTION NO: 41

You have an Azure subscription that includes a user named User1 and an Azure Container Registry named ContReg1. You have enabled content trust for ContReg1. What are the two roles you should assign to User1 to ensure they can create trusted images in ContReg1 while adhering to the principle of least privilege?
Each correct answer represents part of the solution, and each correct selection is worth one point.

A. AcrQuarantineReader

B. Contributor

C. AcrPush

D. AcrImageSigner

E. AcrQuarantineWriter

ANSWER: C D

Explanation:

To create trusted images in an Azure Container Registry with content trust enabled, the user must be able to both push image artifacts to the registry and sign those images so that consumers can verify publisher integrity. The **AcrPush** role provides the minimum permissions required to push images (and related artifacts such as tags and manifests) into the registry without granting broader management rights over the registry resource itself. Separately, content trust in ACR relies

on signing operations (Notary v1) where the signer needs permissions to create and manage signatures/metadata associated with the repository. The **AcrImageSigner** role is designed for this purpose, enabling signing of images in the registry while still keeping permissions scoped to signing rather than full administrative control. Assigning these two roles together satisfies the functional requirement (push + sign) and aligns with least privilege by avoiding overly broad roles like subscription or registry-level Contributor. For details on how ACR content trust works and the required permissions, see https://learn.microsoft.com/en-us/azure/container-registry/container-registry-content-trust and the built-in ACR roles at https://learn.microsoft.com/en-us/azure/container-registry/container-registry-roles.

QUESTION NO: 42

You have an Azure Container Registry named **Registry1**. You have enabled vulnerability scanning of images in **Registry1** through Azure Security Center. You then perform the following actions:

- Upload a Windows image named **Image1** to **Registry1**.
- Upload a Linux image named **Image2** to **Registry1**.
- Upload a Windows image named **Image3** to **Registry1**.
- Modify **Image1** and upload it as **Image4** to **Registry1**.
- Modify **Image2** and upload it as **Image5** to **Registry1**.

Which two images will be analyzed for vulnerabilities? Each correct answer provides a complete solution.

NOTE: Each correct selection is worth one point.

A. Image4

B. Image2

C. Image1

D. Image3

E. Image5

ANSWER: B E

Explanation:

When vulnerability scanning is enabled for an Azure Container Registry via Microsoft Defender for Cloud (formerly Azure Security Center), the service evaluates container images pushed to the registry to identify known CVEs in the OS and package layers. In this scenario, only the Linux-based images are eligible for this type of registry image vulnerability assessment. That means the Linux image uploaded as Image2 will be analyzed, and when you modify that Linux image and push it again as Image5, that new image (new digest/layers) will also be analyzed. The Windows-based images (Image1, Image3, and the modified Windows image Image4) are not included in this specific ACR image vulnerability assessment capability, so they won't be scanned under the described feature set. The key point is that the scanning is triggered by pushing supported images to the registry, and each distinct pushed Linux image is assessed, including subsequent modified versions pushed under a new tag. For more details on Defender for Cloud container image vulnerability assessment and supported image types, see Microsoft Defender for Containers overview and Vulnerability assessment for container images.

QUESTION NO: 43

You manage an Azure Active Directory (Azure AD) tenant that includes users with Azure AD Premium Plan 2 licenses. There is an external partner company with the domain named "fabrikam.com" that includes a user named "User1" with an email

address of user1@fabrikam.com. You need to ensure User1 can access the resources within your Azure AD tenant. The solution should comply with specified requirements. What action should you take?

A. Create a user account for user1.

B. Create an invite for User1.

C. To the tenant add fabrikamcom as a custom domain

D. Set Enable guest self-service sign up via user flows to Yes for the tenant.

ANSWER: B

Explanation:

The correct action is to create an invite for User1 using Azure AD B2B collaboration. Azure AD B2B is the built-in approach for granting external identities access to resources in your tenant (apps, groups, SharePoint sites, etc.) without creating and managing a full internal member account for them. When you invite user1@fabrikam.com, Azure AD creates a guest user object in your directory and sends an invitation. After redemption, the guest signs in using their home identity (typically their organization's Azure AD account or another supported identity provider), and you can then apply access controls such as group membership, app assignments, Conditional Access policies, and entitlement management as needed. This aligns with best practices because it keeps the partner's identity lifecycle primarily managed by their organization while still allowing you to control authorization in your tenant. Microsoft documents this as the standard method for collaborating with users from external organizations and granting them access to your tenant resources. See B2B collaboration overview and Add B2B collaboration users.

QUESTION NO: 44

You have an Azure subscription that contains an app named App1. App1 has the app registration shown in the following table.

You need to ensure that App1 can read all user calendars and create appointments. The solution must use the principle of least privilege.

What should you do?

API	Permission	Type	Admin consent required	Status
Microsoft.Graph	User.Read	Delegated	No	**None**
Microsoft.Graph	Calendars.Read	Delegated	No	**None**

A. Add a new Delegated API permission for Microsoft.Graph Calendars.ReadWrite.

B. Add a new Application API permission for Microsoft.Graph Calendars.ReadWrite.

C. Select Grant admin consent.

D. Add new Delegated API permission for Microsoft.Graph Calendars.ReadWrite.Shared.

ANSWER: B

Explanation:

To allow App1 to read and write calendar data for *all users* in the tenant, the app must use an application permission (app-only access) rather than a delegated permission. Delegated permissions only apply when a signed-in user is present and are constrained to what that user can access, which doesn't meet the requirement to access every user's calendar across the organization. The least-privilege application permission that enables both reading calendars and creating/updating appointments across mailboxes is Microsoft Graph **Calendars.ReadWrite** under **Application permissions**. This permission grants the app the ability to read and write calendars in all mailboxes without requiring a user context, which is exactly what's needed for organization-wide calendar management scenarios such as scheduling services or background automation.

After adding the permission, an administrator must grant tenant-wide admin consent for the permission to take effect, but the key action that satisfies the requirement is selecting the correct least-privilege permission type and scope: application-level Calendars.ReadWrite.

References: Microsoft Graph permissions reference (Calendars), Microsoft Entra ID permissions and consent

QUESTION NO: 45

You have configured your Azure subscription to use an alternative Azure Active Directory (Azure AD) tenant. What are two possible impacts of this change? Each correct answer presents a complete solution.

Note: Each correct selection is worth one point.

A. Role assignments at the subscription level are lost.

B. Virtual machine managed identities are lost.

C. Virtual machine disk snapshots are lost.

D. Existing Azure resources are deleted.

ANSWER: A B

Explanation:

When you change the Azure AD tenant (directory) associated with an Azure subscription, the most immediate impact is on identity-based access and identities that are anchored to the original tenant. Subscription-level Azure RBAC role assignments can effectively be lost because the assignments reference security principals (users, groups, service principals) from the old directory; after the move, those object IDs no longer exist in the new tenant, so access must be re-established in the new directory. In addition, managed identities (including virtual machine system-assigned managed identities) are created as service principals in the subscription's associated Azure AD tenant. After changing the directory, those managed identities are no longer valid in the context of the new tenant and typically must be recreated and reauthorized for any downstream resources they access. These impacts are specifically tied to the subscription-to-tenant association and how Azure RBAC and managed identities depend on Azure AD objects. For details, see Microsoft's guidance on changing a subscription's directory and the resulting access implications, and how managed identities are represented in Azure AD.

Microsoft Docs: Associate or add an Azure subscription to your Azure AD tenant
Microsoft Docs: Managed identities for Azure resources overview

QUESTION NO: 46

Within an Azure subscription, you manage an Azure SQL Database logical server named SQL1 and a virtual machine named VM1, which only uses a private IP address. Below are the current Firewall and virtual network settings for SQL1 as shown in the following image.

To ensure that VM1 can connect to SQL1 while adhering to the principle of least privilege, what action should you take?

A. Add an existing virtual network.

B. Set Connection Policy to Proxy.

C. Create a new firewall rule.

D. Set Allow Azure services and resources to access this server to Yes.

ANSWER: A

Explanation:

The correct action is to add an existing virtual network (by creating a virtual network rule) so that SQL1 allows traffic from the subnet where VM1 resides. Azure SQL Database is a PaaS service and its server-level firewall rules are based on public IP addresses; a VM that “only uses a private IP address” typically reaches Azure SQL over the Azure backbone using service endpoints or Private Link rather than by presenting a stable public source IP. By enabling the Microsoft.Sql service endpoint on the VM’s subnet and then adding that virtual network/subnet to SQL1, you restrict access to only that specific subnet, which aligns well with least privilege. This approach avoids broad allowances and provides a scoped network boundary at the subnet level for the SQL logical server.

References: Azure SQL Database virtual network rules (service endpoints), Configure Azure SQL Database firewall settings.

QUESTION NO: 47

Your network includes an Active Directory forest with the domain contoso.com, and an Azure Active Directory (Azure AD) tenant also named contoso.com. You plan to set up synchronization using the Express Settings installation option in Azure AD Connect. You need to determine the necessary roles and groups required to perform this configuration while following the principle of least privilege. Which two roles and groups should you select? Each correct choice contributes to the solution and is valued at one point.

A. the Domain Admins group in Active Directory

B. the Security administrator role in Azure AD

C. the Global administrator role in Azure AD

D. the User administrator role in Azure AD

E. the Enterprise Admins group in Active Directory

ANSWER: C E

Explanation:

For an Azure AD Connect installation using Express Settings, the wizard needs credentials that can both create/configure objects in Azure AD and create/configure the required synchronization service account and permissions in the on-premises Active Directory forest. On the Azure AD side, Express Settings requires an account with the Global administrator role so the setup can create the Azure AD Connector account, configure directory synchronization features, and register the necessary settings in the tenant. On the on-premises AD side, Express Settings requires Enterprise Admins because the installer must be able to create the AD DS connector account (MSOL_*) and set permissions at the directory/forest scope used for synchronization; Enterprise Admins provides the required forest-wide rights in a least-effort way for the Express path. While more granular delegation is possible with custom settings, Express Settings specifically calls for these elevated roles to complete the end-to-end configuration successfully.

References: Azure AD Connect: Accounts and permissions, Install Azure AD Connect using express settings

QUESTION NO: 48

You have an Azure subscription that includes two virtual machines, named VM1 and VM2, both running Windows Server 2019. You are in the process of implementing Update Management within Azure Automation.

Your goal is to establish a new update deployment named Update1 with the following objectives:

Automatically apply updates to VM1 and VM2.
Automatically include any newly added Windows Server 2019 virtual machines in the Update1 deployment.

Which component should you incorporate into Update1 to achieve these objectives?

A. a security group that has a Membership type of Assigned

B. a security group that has a Membership type of Dynamic Device

C. a dynamic group query

D. a Kusto query language query

ANSWER: C

Explanation:

To meet the requirement of updating VM1 and VM2 now and also automatically targeting any future Windows Server 2019 VMs, you should use a dynamic group query in the update deployment's target selection. In Azure Automation Update Management, update deployments can target machines by selecting saved searches (dynamic groups) that evaluate membership at deployment runtime based on criteria you define (for example, OS type/version, tags, or other properties). This means that when new Windows Server 2019 virtual machines are added and match the query conditions, they are automatically included in the next run of the same update deployment without you having to manually edit the deployment targets. This is the intended mechanism for "evergreen" targeting in Update Management, ensuring consistent patch coverage as your environment grows. See Microsoft's guidance on using dynamic groups/saved searches for Update Management targeting at https://learn.microsoft.com/en-us/azure/automation/update-management/configure-groups and the overview of Update Management deployments and targeting at https://learn.microsoft.com/en-us/azure/automation/update-management/overview.

QUESTION NO: 49

You have an Azure subscription named Sub1. Sub1 contains a virtual network named VNet1 that contains one subnet named Subnet1.

Subnet1 contains an Azure virtual machine named VM1 that runs Ubuntu Server 18.04. You create a service endpoint for Microsoft.Storage in Subnet1.

You need to ensure that when you deploy Docker containers to VM1, the containers can access Azure Storage resources by using the service endpoint.

What should you do on VM1 before you deploy the container?

A. Create an application security group and a network security group (NSG).

B. Edit the docker-compose.yml file.

C. Install the container network interface (CNI) plug-in.

ANSWER: C

Explanation:

Install the container network interface (CNI) plug-in is the correct action because Azure Virtual Network service endpoints apply to traffic sourced from the virtual network/subnet. By default, Docker containers typically use a bridge/overlay network and NAT through the host, meaning the container traffic may not be seen as originating directly from the subnet in a way that reliably leverages subnet-based capabilities. Using Azure's container networking approach with a CNI plug-in attaches containers directly to the Azure virtual network and assigns them IP addresses from the subnet. This makes container egress to Azure Storage originate from the subnet address space, allowing the Microsoft.Storage service endpoint on Subnet1 to be used as intended. In other words, the CNI integration ensures the container network identity is part of the VNet, aligning with how service endpoints enforce access based on VNet/subnet association. This is the key prerequisite on VM1 before deploying containers when you need subnet-level features like service endpoints to apply to container traffic.

References: https://learn.microsoft.com/en-us/azure/virtual-network/container-networking-overview, https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-service-endpoints-overview

QUESTION NO: 50

This question is part of a series that presents the same scenario. Each question 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 answering a question in this section, you will NOT be able to return to it, as these questions will not appear in the review screen.

You have an Azure subscription named Sub1 and an Azure Storage account named sa1 within a resource group named RG1. Several users and applications access both the blob service and the file service of sa1 using multiple shared access signatures (SASs) and stored access policies.

Recently, unauthorized users managed to access both the file and blob services. You are required to revoke all access to sa1.

Solution: You regenerate the Azure storage account access keys.

Does this meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes—regenerating the Azure Storage account access keys meets the goal of revoking access to the storage account when access is being granted via SAS tokens that are signed with the account keys. Account SAS tokens (and any clients using the account keys directly) depend on the current storage account key to authenticate requests. When you regenerate (rotate) the keys, any previously issued SAS tokens that were signed with the old key can no longer be validated by the service, so requests using those SAS tokens will fail. This is a common emergency response step when you suspect key compromise because it immediately cuts off access for any party relying on the compromised key material. After rotation, only clients updated to use the new key (or new SAS tokens signed with the new key) will regain access. Microsoft documents key rotation as the mechanism to invalidate SAS tokens that were created with the account key and to restore control after potential exposure. See Manage storage account access keys and Shared access signatures (SAS) overview.

QUESTION NO: 51

After enabling Azure Container Registry vulnerability scanning via Azure Security Center for images contained in Registry1, you perform specific operations. Which two images will be subjected to vulnerability scanning? Each correct selection represents a complete answer.

NOTE: Each correct answer is worth one point.

A. Image4

B. Image2

C. Image1

D. Image3

E. Image5

ANSWER: B C

Explanation:

When vulnerability scanning for Azure Container Registry is enabled through Microsoft Defender for Cloud (formerly Azure Security Center), the scanning is triggered for images as they are pushed to the registry (or otherwise newly introduced/updated in the registry). In other words, Defender for Cloud evaluates container images that are newly pushed to the protected registry so it can generate vulnerability findings for those images. This behavior aligns with the design goal of continuously assessing the security posture of container images as they enter the supply chain, rather than retroactively scanning every historical image already present before the feature was enabled.

Therefore, the images that will be subjected to vulnerability scanning are the ones that are pushed/updated after enabling the capability (the “newly pushed” images). This is the key operational trigger for image assessment in Azure Container Registry when integrated with Defender for Cloud.

References: Microsoft Docs: Defender for container registries, Microsoft Docs: Microsoft Defender for Cloud overview

QUESTION NO: 52

You have an Azure subscription that uses Microsoft Defender for Cloud. The subscription contains an instance of Azure Database for PostgreSQL.

You need to ensure that an email alert is triggered when a suspected brute force attack on the database is detected. The solution must minimize administrative effort.

What should you configure?

A. the Azure Monitor activity log

B. an Azure Monitor alert rule

C. Microsoft Defender for open-source relational databases

D. the PostgreSQL Audit extension (pgAudit)

ANSWER: C

Explanation:

Configuring Microsoft Defender for open-source relational databases is the right approach because it provides built-in threat detection for Azure Database for PostgreSQL (and MySQL/MariaDB) and can generate security alerts for suspicious activities such as abnormal authentication patterns that may indicate brute-force attempts. Since the requirement is to trigger an email alert with minimal administrative effort, using Defender for Cloud's database protection is preferable to building custom log collection and alerting pipelines. Once enabled, Defender for Cloud continuously analyzes relevant signals and raises security alerts in Defender for Cloud. Those alerts can then be routed to email via Defender for Cloud's integrated alert notifications (or via the standard workflow of sending Defender for Cloud alerts to an action group), without needing to author custom detection logic. This aligns with Microsoft's recommended best practice of using Defender plans for managed threat detection on PaaS databases rather than relying on manual auditing and bespoke alert rules.

References: Microsoft Defender for databases in Defender for Cloud, Configure email notifications for security alerts (Defender for Cloud)

QUESTION NO: 53

You have an Azure SQL database. You implement Always Encrypted.

You need to ensure that application developers can retrieve and decrypt data in the database.

Which two pieces of information should you provide to the developers? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a stored access policy

B. a shared access signature (SAS)

C. the column encryption key

D. user credentials

E. the column master key

ANSWER: C E

Explanation:

To retrieve and decrypt data protected by Always Encrypted, client applications must be able to perform encryption and decryption operations on the client side. That requires the application (and therefore the developers who build/configure it) to have access to the Always Encrypted key metadata: the column encryption key and the column master key. The column encryption key is the symmetric key that actually encrypts and decrypts the values stored in encrypted columns. The column master key is the key-protecting key that encrypts (wraps) the column encryption key and is stored in an external key store such as Azure Key Vault, Windows Certificate Store, or an HSM. With both pieces of information available to the client driver (for example, Microsoft.Data.SqlClient with Always Encrypted enabled), the driver can locate the column master key, unwrap the column encryption key, and transparently decrypt result sets for authorized users. Providing these key details is essential for developers to configure their connection strings, client libraries, and access to the underlying key store so decryption can occur outside SQL Database. See Always Encrypted (Database Engine) and Always Encrypted keys.

QUESTION NO: 54

You have an Azure subscription that contains an Azure key vault named Vault1 and a virtual machine named VM1.

VM1 is connected to a virtual network named VNet1. You need to allow access to Vault1 only from VM1.

What should you do in the Networking settings of Vault1?

A. From the Firewalls and virtual networks tab, add the IP address of VM1.

B. From the Private endpoint connections tab, create a private endpoint for VM1.

C. From the Firewalls and virtual networks tab, add VNet1.

D. From the Firewalls and virtual networks tab, set Allow trusted Microsoft services to bypass this firewall to Yes for Vault1.

ANSWER: C

Explanation:

To allow access to an Azure Key Vault only from a specific VM, you should restrict network access to the virtual network that VM uses by configuring the Key Vault firewall to allow selected virtual networks. Key Vault network rules support allowing traffic from specific VNets/subnets via service endpoints (or, alternatively, via private endpoints), but they do not provide a native way to allow “only this one VM” by its private IP. In the portal’s Networking settings, the practical control you can apply is to allow the virtual network (and subnet) that contains VM1, then ensure only VM1 can reach the vault within that subnet (for example, by using NSGs on the subnet/NIC). This aligns with how Key Vault firewall rules are designed: they accept virtual network rules rather than per-VM rules. Microsoft documents that you can configure Key Vault to allow access from selected networks, including specific virtual networks, using the firewall and virtual network settings. See Azure Key Vault network security and Secure your key vault.

QUESTION NO: 55

Your company plans to create separate subscriptions for each department. Each subscription will be associated to the same Azure Active Directory (Azure AD) tenant. You need to configure each subscription to have the same role assignments.

What should you use?

A. Azure Security Center

B. Azure Policy

C. Azure AD Privileged Identity Management (PIM)

D. Azure Blueprints

ANSWER: D

Explanation:

Azure Blueprints is the right choice because it lets you define a repeatable, governed “package” of subscription configuration that can be applied consistently across many subscriptions. A blueprint definition can include Azure role assignments as first-class artifacts, so when you assign the blueprint to each department’s subscription, the same RBAC role assignments are deployed in a standardized way. This is specifically designed for scenarios where an organization creates multiple subscriptions (often aligned to departments, projects, or environments) and needs each one to start with identical governance and access controls. Blueprints also supports bundling other governance components (like policy assignments, resource group structure, and ARM templates) so the subscription baseline is consistent and auditable over time. This makes it a better fit than tools that focus only on policy evaluation or just-in-time privileged access, because the requirement is to stamp out the same role assignments across subscriptions reliably and repeatedly.

References: Azure Blueprints overview, Blueprint definition and artifacts

QUESTION NO: 56

You have an Azure subscription that contains an Azure key vault named Vault1 and a virtual machine named VM1. VM1 has the Key Vault VM extension installed.

For Vault1, you rotate the keys, secrets, and certificates. What will be updated automatically on VM1?

A. the keys only

B. the secrets only

C. the certificates only

D. the keys and secrets only

E. the secrets and certificates only

F. the keys, secrets, and certificates

ANSWER: C

Explanation:

The certificates only will be updated automatically on VM1. The Key Vault VM extension (commonly used via the Key Vault certificates/secret management extension) is designed to retrieve and periodically refresh certificates from Azure Key Vault onto a virtual machine. When a certificate in Key Vault is renewed/rotated (for example, a new version is created), the extension can detect the newer version and automatically download and install it on the VM (and, depending on configuration, place it into the local certificate store and/or a file path). This is the primary “auto-rotation” scenario supported end-to-end for VMs: keeping the VM’s local certificate material in sync with the latest certificate version in the vault. By contrast, Key Vault keys are not “installed” onto a VM by this extension, and secrets are typically consumed by applications through direct Key Vault access (or other mechanisms) rather than being automatically updated on the VM by the extension in the same way as certificates. For details on how the Key Vault VM extension handles certificate retrieval and rotation behavior, see Key Vault virtual machine extension for Windows and About Azure Key Vault certificates.

QUESTION NO: 57

You are tasked with gathering events from Azure virtual machines into an Azure Log Analytics workspace, intending to create alerts based on these collected events. It is crucial to determine which Azure services are capable of creating these alerts. Which two Azure services should you choose? Each correct answer contributes to the complete solution.

Note: Each accurate selection grants one point.

A. Azure Monitor

B. Azure Security Center

C. Azure Analysis Services

D. Azure Sentinel

E. Azure Advisor

ANSWER: A D

Explanation:

Azure Monitor is a correct choice because Log Analytics workspaces are part of Azure Monitor Logs, and Azure Monitor alert rules can be created directly from Log Analytics queries (log search alerts). This lets you trigger notifications or action groups when specific events appear in collected VM logs, including Windows Event Logs and Syslog data ingested into the workspace. Azure Sentinel is also a correct choice because it is built on top of a Log Analytics workspace and provides analytics rules that generate incidents/alerts from the same ingested log data, enabling security-focused detections and alerting workflows based on events collected from Azure VMs. Together, these services cover both general operational alerting (Azure Monitor) and SIEM/SOAR-style security alerting (Microsoft Sentinel) using the events stored in the Log Analytics workspace.

References: Azure Monitor alerts overview, Microsoft Sentinel overview

QUESTION NO: 58

You have an Azure subscription named Sub1 that contains the virtual machines shown in the following table.

You need to ensure that the virtual machines in RG1 have the Remote Desktop port closed until an authorized user requests access.

What should you configure?

Name	Resource group
VM1	RG1
VM2	RG2
VM3	RG1
VM4	RG2

A. Azure Active Directory (Azure AD) Privileged Identity Management (PIM)

B. an application security group

C. Azure Active Directory (Azure AD) conditional access

D. just in time (JIT) VM access

ANSWER: D

Explanation:

Configure just in time (JIT) VM access to keep the Remote Desktop (RDP) port closed by default and only open it temporarily when an authorized user requests access. JIT VM access (in Microsoft Defender for Cloud) is designed specifically to reduce exposure from always-open management ports like TCP 3389 by creating and managing restrictive inbound rules (typically on the VM's NSG, and it can also integrate with Azure Firewall). With JIT enabled, inbound access to selected ports is denied until a user submits an access request. Defender for Cloud then validates the requester's permissions (via Azure RBAC) and, if allowed, automatically opens the port only for the approved source IP/range and only for the requested time window. After the time expires, Defender for Cloud reverts the rules to the locked-down state, restoring the "closed until requested" posture required by the scenario. This provides controlled, auditable, time-bound RDP access without leaving the port exposed continuously.

References: https://learn.microsoft.com/en-us/azure/defender-for-cloud/just-in-time-access-overview, https://learn.microsoft.com/en-us/azure/defender-for-cloud/enable-just-in-time-access

QUESTION NO: 59

You are tasked with collecting event logs from your Azure virtual machines into an Azure Log Analytics workspace. Subsequently, you intend to generate alerts based on these events. Your objective is to determine which Azure services can facilitate the creation of these alerts. Can you identify the two services that would fulfill this requirement? Note that each correct selection is worth one point, and you should select two answers to form a complete solution.

A. Azure Monitor

B. Azure Security Center

C. Azure Analytics Services

D. Azure Sentinel

E. Azure Advisor

ANSWER: A D

Explanation:

Azure Monitor is a primary service for creating alert rules from data stored in a Log Analytics workspace. After VM event logs are ingested into the workspace (for example via the Azure Monitor Agent and a Data Collection Rule), you can use log search alert rules (scheduled query rules) to run KQL queries against the workspace and trigger notifications or action groups when conditions are met. This is the standard way to alert on collected log events in Azure.

Azure Sentinel (Microsoft Sentinel) also fulfills the requirement because it is built on top of Log Analytics workspaces and provides analytics rules that run scheduled KQL queries over ingested data to generate alerts (and optionally incidents). When your VM event logs are in the connected Log Analytics workspace, Sentinel analytics rules can detect patterns and raise security alerts based on those events, which is commonly used for security-focused alerting and correlation.

References: Azure Monitor alerts overview, Microsoft Sentinel overview

QUESTION NO: 60

You have been tasked with creating an Azure key vault using PowerShell. You have been informed that objects deleted from the key vault must be kept for a set period of 90 days.

Which two of the following parameters must be used in conjunction to meet the requirement? (Choose two.)

A. EnabledForDeployment

B. EnablePurgeProtection

C. EnabledForTemplateDeployment

D. EnableSoftDelete

ANSWER: B D

Explanation:

To ensure deleted Key Vault objects are retained for a defined period (90 days), you must use soft delete with an appropriate retention window. Soft delete is the feature that keeps deleted keys, secrets, and certificates in a recoverable “deleted” state for the configured retention duration, rather than removing them immediately. In PowerShell, this is enabled at vault creation by using the EnableSoftDelete-related parameter (noting that in newer Az.KeyVault versions, soft delete is enabled by default and the retention is controlled by the vault’s soft-delete retention settings).

In addition, enabling purge protection is commonly paired with soft delete to prevent permanent deletion (purge) of deleted objects before the retention period elapses. With purge protection enabled, even a privileged user cannot purge a deleted object until the retention period has passed, which enforces the “must be kept” requirement in practice. Together, soft delete provides recoverability for the retention window, and purge protection ensures the retention window can’t be bypassed by purging.

References: Microsoft Docs: Azure Key Vault soft-delete overview, Microsoft Docs: New-AzKeyVault

QUESTION NO: 61

You are managing a group of 15 Azure virtual machines within a designated resource group called RG1, each running the same application. Your task is to ensure that unauthorized applications and malware are blocked from executing on these virtual machines. What action should you take to achieve this?

A. Apply an Azure policy to RG1.

B. From Azure Security Center, configure adaptive application controls.

C. Configure Azure Active Directory (Azure AD) Identity Protection.

D. Apply a resource lock to RG1.

ANSWER: B

Explanation:

From Microsoft Defender for Cloud (formerly Azure Security Center), configuring adaptive application controls is the right action because it provides application allowlisting (whitelisting) for virtual machines. Adaptive application controls uses observed process behavior to recommend a set of allowed applications for a group of similar VMs, and then enforces those rules so that only approved applications can run. This directly addresses the requirement to block unauthorized applications and malware from executing, which is an execution control problem rather than an identity, governance, or resource protection problem.

In practice, you enable adaptive application controls in Defender for Cloud, review the recommended allowlist for the VM group (such as the VMs in RG1 running the same workload), and then apply/enforce the rules. This hardens the VMs by preventing unknown or unapproved executables from running, reducing the attack surface and limiting the impact of malware that relies on launching new processes.

References: https://learn.microsoft.com/en-us/azure/defender-for-cloud/adaptive-application-controls, https://learn.microsoft.com/en-us/azure/defender-for-cloud/defender-for-cloud-introduction

QUESTION NO: 62

You are managing an Azure Active Directory (Azure AD) tenant named contoso.com, which includes a user named User1. You plan to publish multiple applications within the tenant. To accomplish this, you need to ensure that User1 has the ability to grant admin consent for these published applications. Which two user roles can be assigned to User1 to fulfill this requirement? Note: Each correct selection is worth one point, and each correct answer fully satisfies the requirement.

A. Security administrator

B. Cloud application administrator

C. Application administrator

D. User administrator

E. Application developer

ANSWER: B C

Explanation:

To grant admin consent for applications in Microsoft Entra ID (Azure AD), the user must be assigned a directory role that includes permissions to consent to application permissions on behalf of the organization. The built-in roles that satisfy this requirement are “Cloud application administrator” and “Application administrator”. These roles are designed for managing enterprise applications and app registrations, and they include the ability to grant tenant-wide admin consent to delegated permissions and application permissions (where applicable) so that end users aren’t prompted individually and the organization can centrally approve access. This aligns with the common operational model for publishing multiple applications: the app admin roles can manage app objects and consent flows without requiring full Global Administrator privileges. Microsoft documents that admin consent can be granted by Global Administrator and by specific application administration roles, including the two listed here, which are intended for application lifecycle management and consent administration in the tenant. See Grant admin consent to apps and Microsoft Entra built-in roles (permissions reference) for the role capabilities and consent requirements.

QUESTION NO: 63

You have ten virtual machines situated on a single subnet which utilizes a singular Network Security Group (NSG). You are tasked with the requirement to log the network traffic to an Azure Storage account. What is the appropriate course of action to achieve this?

A. Install the Network Performance Monitor solution.

B. Create an Azure Log Analytics workspace.

C. Enable diagnostic logging for the NSG.

D. Enable NSG flow logs.

ANSWER: D

Explanation:

To log network traffic for virtual machines that are governed by a Network Security Group and store those logs in an Azure Storage account, you should use NSG flow logs. NSG flow logs are a feature of Azure Network Watcher that records information about IP traffic flowing through an NSG, including details such as source/destination IP, ports, protocol, direction, and whether the traffic was allowed or denied. When you enable NSG flow logs, you explicitly choose a Storage account as the destination for the log files, which satisfies the requirement to log traffic to Azure Storage. This approach is designed for capturing NSG-level traffic visibility across all resources associated with that NSG (such as a subnet with multiple VMs), and it is the standard Azure capability for NSG traffic logging. After enabling flow logs, you can optionally integrate with Traffic Analytics or other tooling, but the core action required to write the logs to Storage is enabling NSG flow logs in Network Watcher. See https://learn.microsoft.com/en-us/azure/network-watcher/nsg-flow-logs-overview and https://learn.microsoft.com/en-us/azure/network-watcher/network-watcher-nsg-flow-logging-portal.

QUESTION NO: 64

You are investigating a security issue related to an Azure Storage account. To assist you in this process, you have enabled diagnostic logging for the storage account. Which tool or service should you use to access the diagnostic logs?

A. Azure Security Center

B. Azure Monitor

C. the Security admin center

D. Azure Storage Explorer

ANSWER: B

Explanation:

Azure Monitor is the correct service to access diagnostic logs for an Azure Storage account because Storage diagnostic settings route platform logs and metrics into Azure Monitor's logging pipeline. When you enable diagnostic logging (via diagnostic settings) for a storage account, you choose one or more destinations such as a Log Analytics workspace, an Azure Storage account, or an Event Hub. Azure Monitor then provides the experiences to view and query those logs—most commonly through Log Analytics using Kusto Query Language (KQL), where you can filter by operation type, caller IP, authentication method, and status codes to investigate suspicious activity. This is the standard, Microsoft-recommended approach for collecting and analyzing resource logs across Azure services, including Storage, and it integrates with alerting, workbooks, and SIEM integrations. For details on configuring and consuming Storage logs through Azure Monitor, see https://learn.microsoft.com/en-us/azure/storage/common/storage-monitoring-diagnosing-troubleshooting and the overview of diagnostic settings in Azure Monitor at https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/diagnostic-settings.

QUESTION NO: 65 - (DRAG DROP)

DRAG DROP

You have an Azure subscription named Sub1 that contains an Azure Log Analytics workspace named LAW1.

You have 500 Azure virtual machines that run Windows Server 2016 and are enrolled in LAW1.

You plan to add the System Update Assessment solution to LAW1.

You need to ensure that System Update Assessment-related logs are uploaded to LAW1 from 100 of the virtual machines only.

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.

Select and Place:

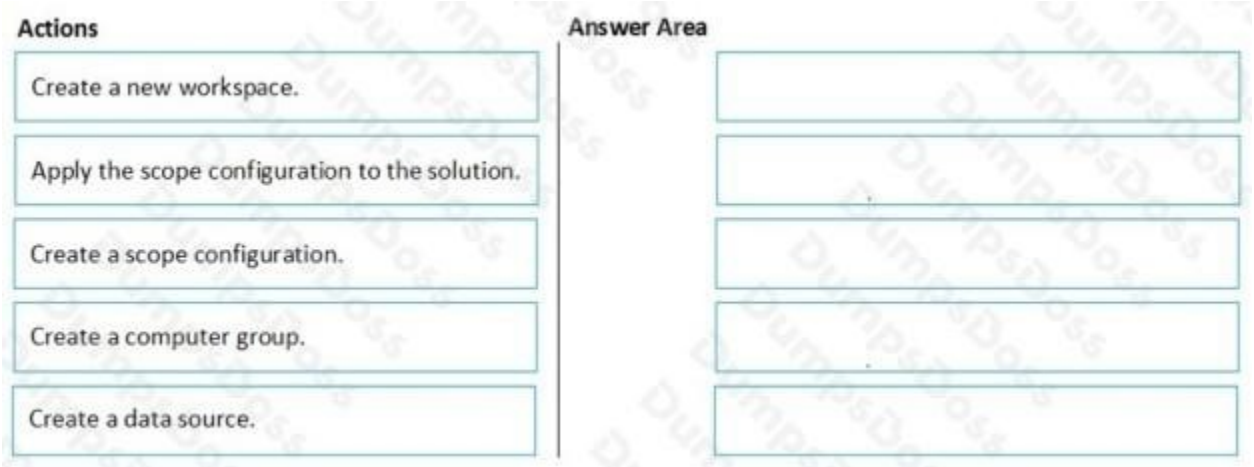

ANSWER:

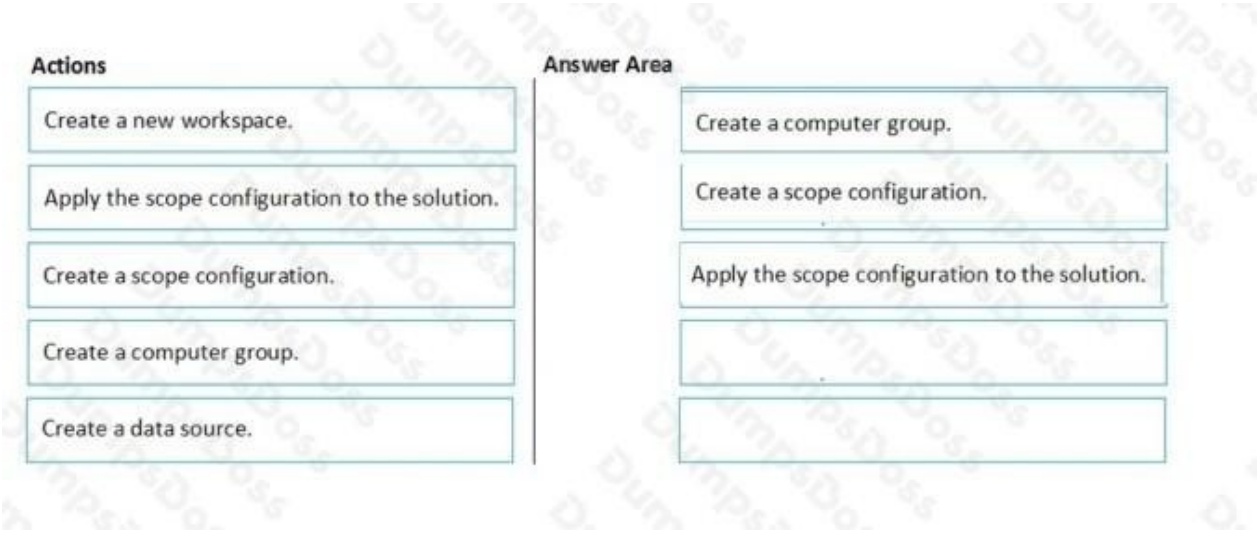

Explanation:

To upload System Update Assessment-related data from only 100 out of 500 enrolled Windows Server VMs, you use *solution targeting* in the Log Analytics workspace. The key idea is that solutions can be configured to collect data only from a defined subset of agents, rather than from every machine connected to the workspace.

The first step is to **create a computer group** that represents exactly the 100 virtual machines you want to include. Computer groups are the mechanism Log Analytics uses to define a dynamic or static set of computers (for example, by name patterns, Azure resource properties, or other criteria). Once that group exists, you then **create a scope configuration**. A scope configuration is the object that ties a solution's data collection scope to one or more computer groups, effectively describing "this solution should apply to these machines." Finally, you **apply the scope configuration to the solution** so that the System Update Assessment solution uses that targeting definition when collecting and uploading its related logs.

This approach avoids creating a separate workspace and avoids configuring per-machine data sources for the solution. Instead, it centrally controls which agents participate in the solution's collection behavior, ensuring only the intended 100 VMs send System Update Assessment data into LAW1.

References: Solution targeting in Azure Monitor (Log Analytics), Computer groups in Azure Monitor logs

QUESTION NO: 66

You have an Azure subscription that includes multiple Azure SQL databases and an Azure Sentinel workspace. Your task is to create a saved query in the Sentinel workspace to identify events generated by Azure Defender for SQL. Which approach should you take?

A. From Azure CLI, run the Get-AzOperationalInsightsWorkspace cmdlet.

B. From the Azure SQL Database query editor, create a Transact-SQL query.

C. From the Azure Sentinel workspace, create a Kusto query language query.

D. From Microsoft SQL Server Management Studio (SSMS), create a Transact-SQL query.

ANSWER: C

Explanation:

From the Azure Sentinel workspace, create a Kusto query language query is correct because Microsoft Sentinel (formerly Azure Sentinel) stores and analyzes security data in a Log Analytics workspace, where data is queried using Kusto Query Language (KQL). Alerts and recommendations produced by Microsoft Defender for SQL (part of Microsoft Defender for Cloud) can be ingested into Sentinel via the Microsoft Defender for Cloud data connector, and once ingested they appear in Log Analytics tables that are designed to be queried with KQL. Sentinel “saved queries” (and analytics rules, hunting queries, workbooks, etc.) are all authored in KQL within the Sentinel/Log Analytics query experience. This makes KQL the supported and intended approach to identify and filter Defender for SQL-generated events across multiple databases and to persist that logic as a saved query in the Sentinel workspace for reuse and automation.

References: https://learn.microsoft.com/en-us/azure/sentinel/kusto-overview, https://learn.microsoft.com/en-us/azure/sentinel/connect-defender-for-cloud

QUESTION NO: 67

You are managing a Microsoft 365 tenant that employs Azure Active Directory (Azure AD), which is synchronized with your on-premises Active Directory domain using Azure AD Connect. Recently, you created a new Azure subscription. However, you notice that the synchronized on-premises user accounts are unable to receive roles in this new subscription. To enable the assignment of Azure and Microsoft 365 roles to these synced Azure AD user accounts, what is the first step you should take?

A. Change the Azure AD tenant used by the new subscription.

B. Configure the Azure AD tenant used by the new subscription to use pass-through authenticate

C. Configure the Azure AD tenant used by the new subscription to use federated authentication.

D. Configure a second instance of Azure AD Connect.

ANSWER: A

Explanation:

Change the Azure AD tenant used by the new subscription is the first step because Azure role assignments (Azure RBAC) and Microsoft 365/Azure AD directory roles can only be granted to identities that exist in the same Microsoft Entra ID (Azure AD) tenant that the subscription is associated with. If the new subscription was created under a different tenant (for example, a default tenant tied to a different account or directory), the synced users from your Azure AD Connect–managed tenant won’t be available for role assignment in that subscription. The correct remediation is to associate (transfer) the subscription to the tenant that contains the synchronized users, so the subscription trusts that directory for authentication and authorization. After the subscription is in the correct tenant, you can assign Azure RBAC roles at subscription/resource scopes and assign directory roles in Microsoft 365/Entra as needed. This is a tenant/subscription association issue, not an authentication method (PTA/federation) issue, and it doesn’t require deploying another Azure AD Connect instance.

References: https://learn.microsoft.com/en-us/azure/role-based-access-control/overview, https://learn.microsoft.com/en-us/azure/active-directory/fundamentals/active-directory-how-subscriptions-associated-directory

QUESTION NO: 68

You have an active Azure subscription and have created an Azure web application named Contoso1812, which is utilizing an S1 App Service plan. You intend to set up a CNAME DNS record for the domain www.contoso.com to direct it towards Contoso1812.

Your objective is to allow users to access Contoso1812 through the URL https://www.contoso.com. What two actions must you undertake? Each correct action constitutes a part of the solution.

Note: Each correct answer selection is worth one point.

A. Turn on the system-assigned managed identity for Contoso1812.

B. Add a hostname to Contoso1812.

C. Scale out the App Service plan of Contoso1812.

D. Add a deployment slot to Contoso1812.

E. Scale up the App Service plan of Contoso1812.

F. Upload a PFX file to Contoso1812.

ANSWER: B F

Explanation:

To make an Azure App Service app reachable at **https://www.contoso.com** when you're using a CNAME, you must first associate the custom domain name with the web app and then enable TLS for that hostname. The required step to associate the domain is to **add a hostname to Contoso1812** (add the custom domain in App Service). This validates the DNS mapping and tells App Service to accept requests for that host header.

After the custom domain is added, HTTPS requires an X.509 certificate bound to that custom hostname. One supported way is to **upload a PFX file to Contoso1812** (upload/import the certificate into App Service) and then create an SSL binding for the custom domain. With an S1 plan, custom domains and TLS/SSL bindings are supported, so no additional scaling actions are required for this scenario.

References: Tutorial: Map an existing custom DNS name to Azure App Service, Secure a custom DNS name with a TLS/SSL binding in Azure App Service.

QUESTION NO: 69

You are responsible for gathering logs from a substantial number of Windows Server 2016 computers using Azure Log Analytics.

Currently, you are configuring an Azure Resource Manager template to deploy the Microsoft Monitoring Agent across all servers automatically.

Which of the following elements should be incorporated into the template for successful deployment? (Choose all that apply.)

A. WorkspaceID

B. AzureADApplicationID

C. WorkspaceKey

D. StorageAccountKey

ANSWER: A C

Explanation:

To deploy the Microsoft Monitoring Agent (MMA) via an Azure Resource Manager template and have it successfully connect to an Azure Log Analytics workspace, the template must provide the workspace identity and a credential that authorizes the agent to register with that workspace. In practice, this means supplying the Log Analytics *WorkspaceID* (also called the Workspace ID or customer ID) so the agent knows which workspace to report to, and the *WorkspaceKey* (the primary or secondary shared key) so the agent can authenticate and establish trust with the workspace during onboarding. These two values are the core parameters used by the MMA/Log Analytics agent configuration and are required for automated, at-scale deployment scenarios. Once installed and configured with these values, the agent can begin sending Windows event logs, performance counters, and other configured data sources into Log Analytics for querying and alerting.

References: Log Analytics agent overview, Install and configure the Log Analytics agent for Windows

QUESTION NO: 70

This question is part of a series that utilizes the same scenario. Each question in the series contains a distinct solution that may or may not achieve the stated objectives. Note that some question sets might have one or more viable solutions, while others might have none.

Once you have answered a question in this section, you cannot revisit it. Consequently, these questions will be omitted from the review screen.

You are utilizing Azure Security Center for centralized policy management across three Azure subscriptions.

You have established multiple policy definitions to ensure the security of the subscriptions.

Your goal is to deploy these policy definitions as a cohesive group to all three subscriptions.

Proposed Solution: You create a policy initiative and assignments scoped to resource groups.

Does this solution meet the stated objective?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because scoping policy initiative assignments to resource groups cannot deploy a cohesive group of policy definitions across multiple subscriptions in one action. Azure Policy initiatives (policy set definitions) are designed to bundle multiple policy definitions together, but the key to applying them consistently across several subscriptions is choosing a scope that spans those subscriptions. The Azure governance construct intended for this is a management group: you can create a management group, place the three subscriptions under it, and then assign the initiative at the management group scope so the assignment is inherited by all child subscriptions. This aligns with centralized policy management goals commonly implemented alongside Microsoft Defender for Cloud (formerly Azure Security Center) recommendations and governance workflows.

Management groups provide a hierarchy above subscriptions specifically for organizing and applying policy and access controls at scale, enabling consistent governance across many subscriptions without duplicating assignments per subscription or per resource group. For details, see Azure management groups overview and Azure Policy scope.

QUESTION NO: 71

You have an Azure subscription that contains two virtual machines named VM1 and VM2 that run Windows Server 2019.

You are implementing Update Management in Azure Automation. You plan to create a new update deployment named Update1.

You need to ensure that Update1 meets the following requirements:

 Automatically applies updates to VM1 and VM2.

 Automatically adds any new Windows Server 2019 virtual machines to Update1. What should you include in Update1?

A. a security group that has a Membership type of Assigned

B. a security group that has a Membership type of Dynamic Device

C. a dynamic group query

D. a Kusto query language query

ANSWER: C

Explanation:

The correct choice is to include a dynamic group query. In Azure Automation Update Management, update deployments can target machines using “dynamic groups,” which are defined by a query that evaluates machine properties and automatically maintains the target set over time. By creating a dynamic group query that matches Windows Server 2019 VMs (and optionally additional constraints like resource group, tags, or naming conventions), Update1 will automatically include VM1 and VM2 today and will also automatically include any newly created Windows Server 2019 virtual machines in the future without manual intervention. This meets both requirements: updates are applied automatically to the current VMs and the scope expands automatically as new qualifying VMs are added. This approach is purpose-built for Update Management’s targeting model and avoids the need to manually manage membership lists. You define the query once, and Update Management continuously evaluates it to determine which machines are in scope for the deployment schedule.

References: Configure dynamic groups for Update Management, Update Management overview

QUESTION NO: 72

You are managing an Azure subscription and aim to set up a workflow automation in Azure Security Center for automatic remediation of a security vulnerability. What should be your first action in this process?

A. a managed identity

B. an automation account

C. an Azure function app

D. an alert rule

E. an Azure logic app

ANSWER: E

Explanation:

To set up workflow automation in Microsoft Defender for Cloud (formerly Azure Security Center) for automatic remediation, the first action is to create the playbook that will run when the workflow automation is triggered. In Defender for Cloud, workflow automation is implemented by connecting security findings (such as alerts or recommendations) to a Logic App playbook. The Logic App defines the remediation steps (for example, creating a ticket, notifying a team, or invoking a remediation action via Azure Resource Manager, Azure Automation, or an Azure Function). Once the Logic App exists, you then create the workflow automation rule in Defender for Cloud to specify the trigger conditions and point to that Logic App. This is why creating an Azure logic app is the correct first step: it provides the executable workflow that Defender for Cloud will call. Afterward, you can configure permissions (often via managed identity) and any downstream remediation components, but the playbook itself is foundational to the automation setup.

References: Microsoft Learn: Workflow automation in Microsoft Defender for Cloud, Microsoft Learn: Use Logic Apps playbooks with Defender for Cloud

QUESTION NO: 73

You create a new Azure subscription.

You need to ensure that you can create custom alert rules in Azure Security Center. Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Onboard Azure Active Directory (Azure AD) Identity Protection.

B. Create an Azure Storage account.

C. Implement Azure Advisor recommendations.

D. Create an Azure Log Analytics workspace.

E. Upgrade the pricing tier of Security Center to Standard.

ANSWER: D E

Explanation:

To create custom alert rules in Microsoft Defender for Cloud (formerly Azure Security Center), you need the capability to author analytics-based detections and have a place to store and run the underlying queries. Custom alert rules are created on top of Log Analytics data (Azure Monitor Logs) using scheduled query rule logic, so a Log Analytics workspace is required to collect and query the relevant security data. In addition, the subscription must be enabled for the enhanced Defender for Cloud features (the paid plan), because advanced threat protection and alerting capabilities are part of the upgraded plan rather than the free foundational posture management features. With the upgraded plan enabled, Defender for Cloud can generate and surface security alerts, and you can build custom detections that leverage the workspace's data. Practically, you also need appropriate permissions (such as write permissions) on the selected workspace to save the rule and configure alerting. This aligns with Microsoft guidance that custom detections/alerts are built from Log Analytics queries and depend on Defender for Cloud's enhanced protections and integrations.

References: Microsoft Defender for Cloud overview, Azure Monitor log alerts (scheduled query rules)

QUESTION NO: 74

You have 10 on-premises servers that are running Windows Server 2019, and you plan to implement vulnerability scanning on these servers using Azure Security Center. What is the first component you should install on the servers to facilitate this integration?

A. the Azure Arc enabled servers Connected Machine agent

B. the Microsoft Defender for Endpoint agent

C. the Security Events data connector in Azure Sentinel

D. the Microsoft Endpoint Configuration Manager client

ANSWER: A

Explanation:

To use Microsoft Defender for Cloud (formerly Azure Security Center) capabilities such as vulnerability assessment with on-premises Windows Server machines, the machines must first be onboarded into Azure so Defender for Cloud can manage and apply security recommendations and extensions. The supported way to connect non-Azure servers is Azure Arc-enabled servers. Installing the Azure Arc enabled servers Connected Machine agent registers each on-premises server as an Azure resource (a “hybrid machine”), which then allows Defender for Cloud to deploy and manage security features and extensions, including vulnerability assessment solutions, from Azure. Without the Arc Connected Machine agent, the servers won’t appear in Azure Resource Manager in a way that Defender for Cloud can target for VA deployment and policy enforcement. After onboarding via Arc, you can enable Defender for Cloud plans and then deploy the appropriate vulnerability assessment extension/solution to those Arc-enabled servers.

References: Azure Arc-enabled servers agent overview, Deploy a vulnerability assessment solution (Defender for Cloud)

QUESTION NO: 75

You have an Azure subscription that contains 100 virtual machines with the Azure Security Center's Standard tier enabled.

To assess the security posture of these virtual machines, you plan to execute a vulnerability scan on each of them.

In order to automate the deployment of the vulnerability scanner extension across all virtual machines, you intend to utilize an Azure Resource Manager template.

Which two parameters must you specify in the template to ensure successful deployment of the extension to the virtual machines? Each correct answer constitutes part of the total solution.

Note: Each correct selection is worth one point.

A. the user assigned managed identity

B. the Key Vault managed storage account Key

C. the Azure Active Directory (Azure AD) ID

D. the system-assigned managed identity

E. the primary shared key

F. the workspace ID

ANSWER: C F

Explanation:

When you deploy the Microsoft Defender for Cloud (formerly Azure Security Center) vulnerability assessment solution to Azure virtual machines via an ARM template, the extension must be able to register with the vulnerability assessment service and report findings back to Defender for Cloud. For the built-in Qualys-based vulnerability assessment extension, the ARM deployment requires you to provide identifiers that let the extension associate itself with the correct tenant and subscription context. In practice, this means specifying the Azure Active Directory (Azure AD) ID (tenant ID) so the extension can authenticate/associate to the correct directory, and specifying the workspace ID used by Defender for Cloud/Log Analytics integration so collected assessment data can be routed and correlated properly for posture reporting. These values are commonly required inputs in automated deployments of security/monitoring extensions and are the key parameters that enable the extension to onboard and send results without embedding secrets. See Microsoft's guidance on vulnerability assessment in Defender for Cloud and extension-based deployment patterns for VMs: Deploy a vulnerability assessment solution on Azure VMs and ARM template parameterization concepts: ARM template parameters.

QUESTION NO: 76

Your network contains an Active Directory forest named contoso.com. You have an Azure Active Directory (Azure AD) tenant named contoso.com.

You plan to configure synchronization by using the Express Settings installation option in Azure AD Connect.

You need to identify which roles and groups are required to perform the planned configuration. The solution must use the principle of least privilege.

Which two roles and groups should you identify? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. the Domain Admins group in Active Directory

B. the Security administrator role in Azure AD

C. the Global administrator role in Azure AD

D. the User administrator role in Azure AD

E. the Enterprise Admins group in Active Directory

ANSWER: C E

Explanation:

When you install Azure AD Connect using Express Settings, the wizard performs both Azure AD and on-premises Active Directory configuration tasks automatically. For Azure AD, Express Settings requires an account that can create and configure the Azure AD Connect synchronization service principal and set up directory synchronization in the tenant. Under least privilege guidance for Azure AD Connect, this is satisfied by the Global administrator role in Azure AD during installation (you can reduce privileges after setup, but the initial configuration requires Global Admin).

On-premises, Express Settings also creates the required AD DS connector account and configures permissions needed to read from the directory and (depending on features) write back attributes. For the initial setup across the forest, the installer needs sufficient rights to create objects and set permissions in the directory. Microsoft's documented least-privilege approach for the installation phase is to use an account that is a member of the Enterprise Admins group in Active Directory, which has the necessary forest-wide rights to perform the required configuration during the wizard-driven setup.

References: https://learn.microsoft.com/en-us/entra/identity/hybrid/connect/reference-connect-accounts-permissions, https://learn.microsoft.com/en-us/entra/identity/hybrid/connect/how-to-connect-install-express

QUESTION NO: 77

You have an Azure subscription with 100 virtual machines and Azure Security Center Standard tier enabled. You plan to conduct a vulnerability scan on each virtual machine. What are the two values you need to include in your Azure Resource Manager template to automate the deployment of the vulnerability scanner extension onto the virtual machines? Select the correct options from the following choices. Each option you select is a part of the solution and worth one point.

A. the user-assigned managed identity

B. the workspace ID

C. the Azure Active Directory (Azure AD) ID

D. the Key Vault managed storage account key

E. the system-assigned managed identity

F. the primary shared key

ANSWER: B F

Explanation:

To automate deployment of the built-in vulnerability assessment solution for Azure VMs (Microsoft Defender for Cloud vulnerability assessment), your ARM template must provide the Log Analytics workspace details that the agent/extension uses to report assessment data. Specifically, the template needs the workspace ID (also called the customer ID) and the workspace primary shared key so the VM can authenticate and send data to the workspace. These two values are the standard pair used when onboarding VM agents/extensions to Log Analytics and are required in many ARM-based deployments that connect a VM to a workspace. Once the VM is connected, Defender for Cloud can orchestrate vulnerability assessment and surface findings in the security recommendations and inventory views. This aligns with Microsoft guidance for deploying and configuring vulnerability assessment and for connecting machines to a Log Analytics workspace via templates. See Deploy a vulnerability assessment solution on Azure VMs and Install Log Analytics agent using Azure Resource Manager templates.

QUESTION NO: 78

You have an Azure subscription configured with various resources, detailed as follows:

Name	Type	Description
RG1	Resource group	Used to store virtual machines
RG2	Resource group	Used to store virtual networks
ServerAdmins	Security group	Used to manage virtual machines

Your task is to enable the ServerAdmins to carry out specific operations, adhering strictly to the principle of least privilege. Which two role-based access control (RBAC) roles should you assign to the ServerAdmins group? Each correct selection contributes to part of the complete solution.

NOTE: Every correct response is worth one point.

A. a custom RBAC role for RG2

B. the Network Contributor role for RG2

C. the Contributor role for the subscription

D. a custom RBAC role for the subscription

E. the Network Contributor role for RG1

F. the Virtual Machine Contributor role for RG1

ANSWER: B F

Explanation:

To follow least privilege while enabling the ServerAdmins group to manage the environment described, scope permissions to the resource groups that contain the relevant resource types. For virtual machine administration in RG1, the built-in Virtual Machine Contributor role is designed specifically to manage virtual machines (including start/stop, configuration, extensions, and related VM operations) without granting broad rights across unrelated resources. For virtual network administration in RG2, the built-in Network Contributor role provides the necessary permissions to manage networking resources such as virtual networks, subnets, network interfaces, public IPs, and load balancers within that resource group. Assigning these two built-in roles at the resource-group scope aligns with least privilege by avoiding subscription-wide Contributor access and by using Microsoft-maintained role definitions rather than introducing unnecessary custom roles when an appropriate built-in role exists. This approach is consistent with Azure RBAC best practices: use built-in roles when possible and scope assignments as narrowly as practical. See Azure built-in roles and Azure RBAC overview.

QUESTION NO: 79

You have an Azure Active Directory (Azure AD) tenant. You have the deleted objects shown in the following table.

On May 4, 2020, you attempt to restore the deleted objects by using the Azure Active Directory admin center.

Which two objects can you restore? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

Name	Type	Deleted on
Group1	Security group	April 5, 2020
Group2	Office 365 group	April 5, 2020
User1	User	March 25, 2020
User2	User	April 30, 2020

A. Group1

B. Group2

C. User2

D. User1

ANSWER: B C

Explanation:

In Azure AD, many directory objects are “soft deleted” and can be restored from the Deleted users or Deleted groups areas in the Azure Active Directory admin center, as long as the restore is attempted within the soft-delete retention window. For users, Azure AD retains deleted user accounts for 30 days by default, during which time an administrator can restore the user and recover key properties. For Microsoft 365 groups (formerly Office 365 groups), Azure AD also supports restoring a deleted group within a 30-day retention period, allowing the group and its associated resources to be recovered. Therefore, the objects that were deleted within that 30-day window as of May 4, 2020 are the ones that can be restored using the admin center.

This aligns with Microsoft’s documented behavior for restoring deleted Microsoft 365 groups and users from the Azure AD recycle bin experience. After the retention period expires, the objects are permanently deleted and can no longer be restored through the portal.

References: Restore a deleted Microsoft 365 group in Microsoft Entra ID, Restore a deleted user in Microsoft Entra ID

QUESTION NO: 80

You have an Azure subscription called Sub1, which includes a virtual network named VNet1 containing a single subnet named Subnet1. Within Subnet1, there is an Azure virtual machine named VM1 operating on Ubuntu Server 18.04.

A service endpoint for Microsoft.Storage has been created in Subnet1. To ensure that Docker containers deployed on VM1 can utilize the service endpoint to access Azure Storage resources, what action must be taken on VM1 prior to deploying the container?

A. Create an application security group and a network security group (NSG).

B. Edit the docker-compose.yml file.

C. Install the container network interface (CNI) plug-in.

ANSWER: C

Explanation:

Install the container network interface (CNI) plug-in.

is correct because Azure Virtual Network service endpoints apply to traffic that originates from a subnet in a virtual network. By default, Docker containers typically use a bridge/overlay network (NAT) on the host, so their outbound traffic may not be seen by Azure as coming directly from the subnet IP space in a way that reliably satisfies service endpoint expectations. Using an Azure VNet-integrated container networking approach (via a CNI plugin) attaches container network interfaces directly to the Azure virtual network and assigns IP addresses from the subnet to the containers. This makes container traffic originate from the subnet where the Microsoft.Storage service endpoint is enabled, allowing the storage account’s virtual network rules/service endpoint policy to recognize and permit the traffic as intended. In practice, this is the same principle used by Azure CNI for Kubernetes workloads: pods/containers get VNet IPs and communicate as first-class VNet endpoints, enabling features like service endpoints and private access patterns to work consistently at the network layer. See Azure

container networking overview and service endpoint behavior in Microsoft documentation: https://learn.microsoft.com/en-us/azure/virtual-network/container-networking-overview and https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-service-endpoints-overview.

QUESTION NO: 81

You have an Azure Active Directory (Azure AD) tenant named contoso.com that includes a user called User1. You are planning to publish several applications in the tenant. You need to ensure that User1 can grant admin consent for these published applications. Which two user roles can you assign to User1 to achieve this goal? Each correct answer provides a complete solution.

NOTE: Each correct selection is worth one point.

A. Application developer

B. Security administrator

C. Application administrator

D. User administrator

E. Cloud application administrator

ANSWER: C E

Explanation:

To allow User1 to grant admin consent for published applications, you must assign an Azure AD role that includes permissions to manage enterprise applications/app registrations and to consent to delegated and application permissions on behalf of the organization. The built-in roles that are designed for this purpose are application-focused administrator roles. Both the Application administrator role and the Cloud application administrator role can manage application registrations and enterprise applications, and they can grant tenant-wide admin consent to permissions requested by apps (without requiring the Global Administrator role). This aligns with Microsoft's guidance that these roles can perform admin consent operations for apps, enabling centralized approval of permissions for users in the tenant while maintaining least privilege compared to Global Administrator.

For details on the admin consent workflow and who can grant admin consent, see Grant tenant-wide admin consent to an application. For role capabilities and least-privilege role selection, see Microsoft Entra built-in roles (permissions reference).

QUESTION NO: 82 - (HOTSPOT)

HOTSPOT (Drag and Drop is not supported) (Drag and Drop is not supported) You have an Azure subscription that contains the storage accounts shown in the following table.

You need to configure authorization access.

Which authorization types can you use for each storage account? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.
Hot Area:

Name	Type
storage1	Azure Blob storage
storage2	Azure Files SMB
storage3	Azure Table storage

Answer Area

storage1:

- Shared Key only
- Shared access signature (SAS) only
- Azure Active Directory (Azure AD) only
- Shared Key and shared access signature (SAS) only
- Shared Key, shared access signature (SAS), and Azure Active Directory (Azure AD)

storage2:

- Shared Key only
- Shared access signature (SAS) only
- Shared Key and shared access signature (SAS)

storage3:

- Shared Key only
- Shared access signature (SAS) only
- Azure Active Directory (Azure AD) only
- Shared Key and shared access signature (SAS) only
- Shared Key, shared access signature (SAS), and Azure Active Directory (Azure AD)

ANSWER:

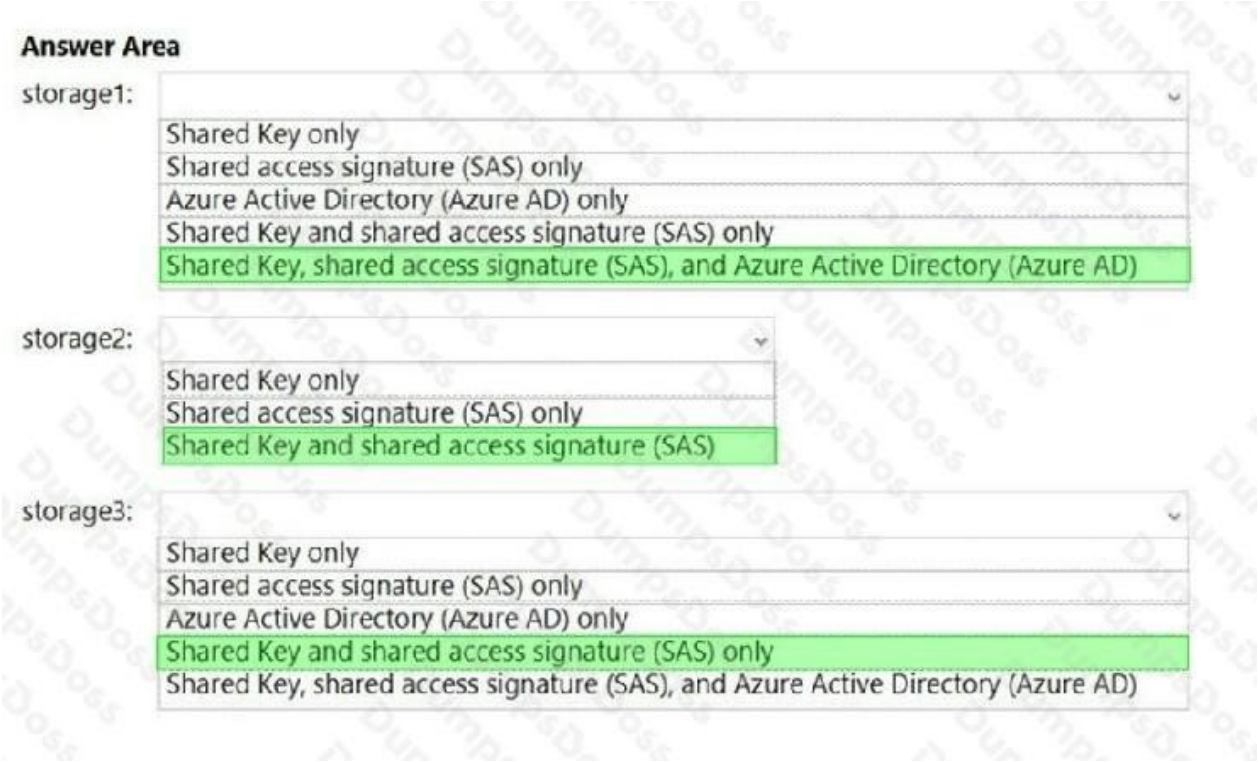

Explanation:

For Azure Storage, the available authorization methods depend on the specific storage service you're accessing. Azure Blob storage supports all three common data-plane authorization models: using the storage account access key (Shared Key), delegating access with a shared access signature (SAS), and using Microsoft Entra ID (Azure AD) with OAuth tokens and Azure RBAC for fine-grained access to containers and blobs. That's why the correct selection for the Blob storage account includes Shared Key, SAS, and Azure AD.

Azure Files accessed over SMB does not use Azure AD OAuth for authorization in the same way Blob does. In exam terms, the supported authorization types you can select in this hotspot are limited to Shared Key and SAS (SAS applies to Azure Files over the REST API scenarios, and Shared Key covers key-based access). Therefore, the correct selection for the Azure Files SMB account is the option that includes Shared Key and SAS.

Azure Table storage (the classic Table service in a storage account) supports Shared Key authorization and SAS for delegated access, but it does not support Azure AD OAuth authorization for the Table service. So the correct selection for the Table storage account is the option that includes Shared Key and SAS only (and not Azure AD).

References: Authorize access to data in Azure Storage, Authorize access to blobs using Microsoft Entra ID.

QUESTION NO: 83

Your network environment includes an Active Directory forest named contoso.com, consisting of a single domain. You have an Azure subscription, Sub1, linked to an Azure Active Directory (Azure AD) tenant also named contoso.com. You are planning to deploy Azure AD Connect to integrate your on-premises Active Directory with the Azure AD tenant.

Your integration solution must meet the following criteria:

Ensure that password policies and user logon restrictions apply to user accounts synchronized to the tenant.
Minimize the number of servers required for the solution.

Which authentication method should you recommend?

▪

A. federated identity with Active Directory Federation Services (AD FS)

B. password hash synchronization with seamless single sign-on (SSO)

C. pass-through authentication with seamless single sign-on (SSO)

ANSWER: B

Explanation:

password hash synchronization with seamless single sign-on (SSO) is the best fit because it requires the least on-premises infrastructure while still enabling users to authenticate to Azure AD using credentials derived from their on-premises Active Directory. With password hash synchronization, Azure AD Connect synchronizes a hash of the on-premises password to Azure AD, allowing cloud authentication without deploying additional authentication servers. This directly supports the requirement to minimize the number of servers, since you typically only need the Azure AD Connect server itself.

In addition, password hash synchronization supports applying Azure AD sign-in controls (such as Conditional Access) to synchronized users, and it aligns with the goal of keeping authentication simple and highly available without relying on extra on-prem components. Seamless SSO can be enabled alongside password hash synchronization to provide a smoother domain-joined experience for users on the corporate network without introducing the operational overhead of federation infrastructure.

For more details, see Password hash synchronization and Seamless single sign-on.

QUESTION NO: 84

You have an Azure subscription named Subscription1 that contains a resource group named RG1 and the users shown in the following table.

You perform the following tasks:

□ Assign User1 the Network Contributor role for Subscription1.

□ Assign User2 the Contributor role for RG1.

To Subscription1 and RG1, you assign the following policy definition: External accounts with write permissions should be removed from your subscription.

What is the Compliance State of the policy assignments?

Name	User principal name (UPN)	Type
User1	User1@outlook.com	Guest
User2	User2@outlook.com	Guest

A. The Compliance State of both policy assignments is Non-compliant.

B. The Compliance State of the policy assignment to Subscription1 is Compliant, and the Compliance State of the policy assignment to RG1 is Non-compliant.

C. The Compliance State of the policy assignment to Subscription1 is Non-compliant, and the Compliance State of the policy assignment to RG1 is Compliant.

D. The Compliance State of both policy assignments is Compliant.

ANSWER: A

Explanation:

The correct outcome is that the Compliance State of both policy assignments is Non-compliant. The built-in Azure Policy definition “External accounts with write permissions should be removed from your subscription” evaluates role assignments and flags any identities that are considered external (typically guest users from another Azure AD tenant) that have write-capable permissions at the evaluated scope. “Write permissions” in this context are commonly represented by RBAC roles that include Microsoft.Authorization/*/write or broad management actions, such as Contributor or specialized contributor roles like Network Contributor, because they allow creating/updating resources. Since User1 is granted Network Contributor at the subscription scope, the subscription-level assignment will detect an external principal with write permissions and mark the subscription scope non-compliant. Separately, User2 is granted Contributor at the resource group scope, so the resource-group-level assignment will also detect an external principal with write permissions within RG1 and mark that scope non-compliant. Azure Policy compliance is evaluated per assignment and scope, so having external write access at each respective scope results in non-compliance for both assignments.

References: https://learn.microsoft.com/en-us/azure/governance/policy/overview, https://learn.microsoft.com/en-us/azure/role-based-access-control/built-in-roles

QUESTION NO: 85

You possess an Azure subscription known as Subscription1. Your task is to identify the default security settings that have been applied to Subscription1. To obtain this information, which specific Azure policy or initiative definition should you examine?

A. the Audit diagnostic setting policy definition

B. the Enable Monitoring in Azure Security Center initiative definition

C. the Enable Azure Monitor for VMs initiative definition

D. the Azure Monitor solution ‘Security and Audit’ must be deployed policy definition

ANSWER: B

Explanation:

The default security settings applied at the subscription level by Microsoft Defender for Cloud (formerly Azure Security Center) are implemented through an Azure Policy initiative assignment. The built-in initiative that represents these baseline, default security configurations is the one that enables Defender for Cloud monitoring and related policy effects (such as auditing and deploying required agents/extensions for data collection). By reviewing the initiative definition named “Enable Monitoring in Azure Security Center,” you can see the collection of policy definitions that Defender for Cloud uses to assess resources and, where applicable, drive automatic provisioning/monitoring behavior. This is the most direct way to understand what “default security settings” are being enforced or evaluated for the subscription via policy, because Defender for Cloud security policy is expressed and managed through Azure Policy initiatives and assignments at the management group/subscription scope. For more details on how Defender for Cloud uses Azure Policy initiatives to implement security policy, see Microsoft Docs: Security policy in Microsoft Defender for Cloud and the related Microsoft Docs: Defender for Cloud policy reference.

QUESTION NO: 86

You have an Azure subscription, and you create a new virtual network named VNet1. You plan to deploy an Azure web application, referred to as App1, which will utilize VNet1 and be accessible using private IP addresses. The solution must accommodate both inbound and outbound network traffic. What action should you take to achieve this?

A. Create an Azure App Service Hybrid Connection.

B. Configure regional virtual network integration.

C. Create an App Service Environment

D. Create an Azure application gateway.

ANSWER: C

Explanation:

Create an App Service Environment is the correct action because it places the App Service plan and web apps into your virtual network, enabling private IP addressing and private inbound access directly from VNet1. An App Service Environment (ASE) is a single-tenant deployment of Azure App Service that supports internal (ILB) mode, where the app is reachable only via private endpoints inside the VNet, satisfying the “accessible using private IP addresses” requirement. Because the app is deployed into the VNet, it also supports outbound connectivity to resources in the VNet (and beyond, subject to routing/NSGs), meeting the requirement to accommodate both inbound and outbound traffic. In contrast, regional virtual network integration primarily addresses outbound connectivity from the app to the VNet and does not by itself provide private inbound access via a private IP for the app. For private inbound access to App Service, Microsoft documents using an App Service Environment (or Private Endpoint for inbound only), and ASE is the option that natively supports both directions with private addressing.

References: https://learn.microsoft.com/en-us/azure/app-service/environment/overview, https://learn.microsoft.com/en-us/azure/app-service/networking-features

QUESTION NO: 87

You have an Azure subscription that contains a user named User1 and an Azure Container Registry named ContReg1.

You enable content trust for ContReg1.

You need to ensure that User1 can create trusted images in ContReg1. The solution must use the principle of least privilege.

Which two roles should you assign to User1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. AcrQuarantineReader

B. Contributor

C. AcrPush

D. AcrImageSigner

E. AcrQuarantineWriter

ANSWER: C D

Explanation:

To create *trusted* images in an Azure Container Registry with content trust enabled, a user must be able to both push image artifacts to the registry and sign those images so that consumers can verify publisher authenticity and integrity. The least-privilege approach is to grant only the specific data-plane permissions required for these two actions.

The **AcrPush** role provides the ability to push images (and related artifacts) to the registry, which is required to publish the image content that will later be trusted. Separately, content trust relies on signing operations; the **AcrImageSigner** role grants the permissions needed to sign images in the registry so that the resulting images are considered trusted when pulled by clients enforcing signature verification.

Assigning broader roles like subscription or registry-level Contributor would exceed least privilege because it includes management-plane permissions unrelated to pushing and signing images. Therefore, combining the push capability with the image signing capability is the minimal set needed for User1 to create trusted images in ContReg1.

References: Content trust in Azure Container Registry, Azure Container Registry roles and permissions

QUESTION NO: 88

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

Name	Type	Description
RG1	Resource Group	Used to store virtual machines
RG2	Resource Group	Used to store virtual networks
ServerAdmins	Security Group	Used to manage virtual machines

You need to allow ServerAdmins to perform the following task: Create a virtual machine in the existing virtual network located in RG2 only. The solution must adhere to the principle of least privilege.
Which two role-based access control (RBAC) roles should you assign to the ServerAdmins group? Each correct answer is part of the needed solution.
NOTE: Each correct selection is worth one point.

A. the Contributor role for the subscription

B. the Network Contributor role for RG2

C. A custom RBAC role for the subscription

D. a custom RBAC role for RG2

E. the Network Contributor role for RG1.

F. the Virtual Machine Contributor role for RG1.

ANSWER: B F

Explanation:

To create a virtual machine, the group needs permissions to create and manage the VM resource itself in the resource group where VMs are stored, and also permissions to join the VM's network interface to an existing virtual network/subnet.

Assigning the Virtual Machine Contributor role for RG1 provides the required compute permissions to create and manage virtual machines (including related compute operations) within RG1 without granting broad rights across the subscription. Separately, because the virtual network is in RG2, the group must be able to read and use that network (for example, to attach a NIC to a subnet and potentially interact with network-related resources required during VM provisioning). Assigning the Network Contributor role for RG2 grants the necessary network scope permissions on the virtual network resources in RG2, enabling VM creation that targets that existing network while still limiting access to only the networking resource group. This combination follows least privilege by scoping compute permissions to RG1 and network permissions to RG2, rather than using a broad Contributor assignment at subscription scope. For details on built-in roles and their permissions, see Azure built-in roles and role assignment guidance at Assign Azure roles using the Azure portal.

QUESTION NO: 89

You have an Azure Active Directory (Azure AD) tenant named contoso.com. Your task is to configure diagnostic settings for contoso.com while meeting the following requirements:

Retain logs for two years.
Query logs using the Kusto Query Language.
Minimize administrative effort.

Considering these requirements, where should you store the logs?

A. an Azure event hub

B. an Azure Log Analytics workspace

C. an Azure Storage account

ANSWER: B

Explanation:

Storing the logs in an Azure Log Analytics workspace best meets all the stated requirements. Log Analytics is part of Azure Monitor Logs and is designed for interactive log analytics using the Kusto Query Language (KQL), which directly satisfies the need to query logs with KQL without building or maintaining a separate query pipeline. It also supports configurable data retention, allowing you to keep data for extended periods (including two years) by setting the workspace retention appropriately, which aligns with the retention requirement. From an operational standpoint, sending Azure AD diagnostic logs to Log Analytics is a native, first-class integration through diagnostic settings, and it centralizes collection, retention, and querying in one managed service—typically the lowest administrative effort compared to streaming to an event hub (which often implies downstream consumers) or storing raw files in a storage account (which would require additional tooling to query effectively with KQL). Log Analytics also enables built-in experiences like Log Analytics queries, workbooks, and alerting on the collected data.

References: https://learn.microsoft.com/en-us/azure/azure-monitor/logs/log-analytics-workspace-overview, https://learn.microsoft.com/en-us/azure/azure-monitor/logs/data-retention-configure

QUESTION NO: 90

You have an Azure subscription that includes multiple virtual machines with Just-in-Time (JIT) VM access enabled. You need to establish a connection to one of these virtual machines using Remote Desktop Protocol (RDP). What is the first action you should take to accomplish this?

A. From Azure Directory (Azure AD) Privileged Identity Management (PIM), activate the Security administrator user role.

B. From Azure Active Directory (Azure AD) Privileged Identity Management (PIM), activate the Owner role for the virtual machine.

C. From the Azure portal, select the virtual machine, select Connect, and then select Request access.

D. From the Azure portal, select the virtual machine and add the Network Watcher Agent virtual machine extension.

ANSWER: C

Explanation:

When Just-in-Time (JIT) VM access is enabled, inbound management ports such as RDP (3389) are not left open continuously. Instead, you must explicitly request temporary access so that Microsoft Defender for Cloud (or the JIT policy) can create a time-bound NSG (or firewall) rule allowing your source IP to reach the VM on the required port. Therefore, the first step before attempting an RDP connection is to initiate a JIT access request from the VM's Connect experience in the Azure portal. Selecting the VM, choosing Connect, and then selecting Request access triggers the workflow to open RDP for a limited duration and (optionally) a specific IP range, after which access is automatically revoked. This is the intended operational flow for connecting to JIT-protected VMs and aligns with Azure security best practices by reducing exposure of management ports.

References: Use just-in-time VM access in Microsoft Defender for Cloud, Just-in-time VM access overview

QUESTION NO: 91

You have an Azure SQL database configured with Always Encrypted. You need to ensure that application developers have the necessary information to retrieve and decrypt the data within the database.

Which two pieces of information should you provide to the developers? Each correct answer represents a part of the solution.

NOTE: Each correct selection is worth one point.

A. a stored access policy

B. a shared access signature (SAS)

C. the column encryption key

D. user credentials

E. the column master key

ANSWER: C E

Explanation:

To retrieve and decrypt data protected by Always Encrypted, developers need the key material and metadata that enables the client driver to perform encryption and decryption on the client side. Always Encrypted relies on a two-tier key hierarchy: the column encryption key and the column master key. The column encryption key is the symmetric key that actually encrypts and decrypts values in encrypted columns. The column master key is the key-protecting key that encrypts (wraps)

the column encryption key and is stored in an external trusted key store such as Azure Key Vault or a certificate store. When an application queries encrypted columns using an Always Encrypted-enabled driver, the driver uses the column master key to unwrap the column encryption key, and then uses the column encryption key to decrypt the returned ciphertext. Providing developers with access to (and details about) the column encryption key and the column master key is therefore essential for building and configuring applications that can transparently read and write encrypted data.

References: Always Encrypted (Database Engine), Always Encrypted keys

QUESTION NO: 92

You have an Azure subscription.

You configure the subscription to use a different Azure Active Directory (Azure AD) tenant.

What are two possible effects of the change? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Role assignments at the subscription level are lost.

B. Virtual machine managed identities are lost.

C. Virtual machine disk snapshots are lost.

D. Existing Azure resources are deleted.

ANSWER: A B

Explanation:

When you change (transfer) an Azure subscription to a different Azure AD tenant, the subscription's trust boundary for identity changes. As a result, Azure role-based access control (RBAC) assignments that were granted to users, groups, and service principals in the original tenant no longer resolve in the new tenant, so subscription-level role assignments are effectively lost and must be recreated for identities in the destination tenant. In addition, managed identities are tenant-bound: a system-assigned managed identity is represented by a service principal in the tenant, and user-assigned managed identities are Azure resources tied to a tenant as well. After moving the subscription to a new tenant, those identity objects aren't available in the destination tenant in the same way, so workloads that relied on virtual machine managed identities can lose access until identities and permissions are re-established in the new tenant. Importantly, the underlying Azure resources (compute, storage, disks, snapshots, etc.) remain in the subscription; the main impact is on identity, access, and tenant-scoped objects. See Microsoft's guidance on associating subscriptions with a directory and the effects of changing the directory: Microsoft Learn and Azure RBAC overview.

QUESTION NO: 93

You have an Azure subscription that contains an Azure key vault. You create a storage account named storage1.

You plan to store data in the following storage1 services:

• Azure Files

• Azure Blob storage

• Azure Table storage

• Azure Queue storage

For which two services can you configure data encryption by using the keys stored in the key vault? Each correct answer presents a complete solution,

NOTE: Each correct selection is worth one point.

A. Blob storage

B. Table storage

C. Queue storage

D. Azure Files

ANSWER: A D

Explanation:

In Azure Storage, you can use customer-managed keys (CMK) stored in Azure Key Vault (or Managed HSM) to control the encryption keys used by Storage Service Encryption (SSE). CMK support applies to the core data services where SSE is available and integrated with Key Vault for key management and rotation workflows. Specifically, Azure Blob storage supports configuring encryption with customer-managed keys from a key vault, allowing you to meet stricter compliance requirements and maintain control over key lifecycle operations (such as rotation and revocation) while still using platform-managed encryption at rest. Azure Files also supports customer-managed keys for encryption at rest, enabling you to use a key vault key to protect file share data and align with organizational key custody policies. Configuration is done at the storage account level (encryption settings) and requires appropriate permissions for the storage account to access the key vault key (for example via managed identity and Key Vault access policies/RBAC). For details, see Microsoft's documentation on Storage encryption with customer-managed keys and the supported services: https://learn.microsoft.com/en-us/azure/storage/common/customer-managed-keys-overview and https://learn.microsoft.com/en-us/azure/storage/common/storage-service-encryption.

QUESTION NO: 94

You have an Azure subscription that includes an Azure SQL database named SQL1 and an Azure Key Vault named KeyVault1. KeyVault1 contains the keys listed in the table below.

You need to configure Transparent Data Encryption (TDE) so that it uses a customer-managed key for SQL1. Which keys can be used for this purpose?

A. Key1. Key2 Key3. and Key4

B. Key1 only

C. Key2 only

D. Key1 and key2 only

E. Key2 and Key3 only

ANSWER: E

Explanation:

For Azure SQL Database Transparent Data Encryption with customer-managed keys (often called TDE with BYOK), the TDE protector must be an asymmetric RSA key stored in Azure Key Vault (or Managed HSM) that Azure SQL can access. Microsoft documents that supported key types for TDE customer-managed keys are RSA keys with sizes 2048-bit, 3072-bit, or 4096-bit. This requirement exists because Azure SQL uses the Key Vault key as the TDE protector (wrapping the database encryption key), and only those RSA key sizes are supported for this integration. Therefore, the usable keys are the ones in KeyVault1 that match the RSA type and one of the supported sizes, which corresponds to the set described by "Key2 and Key3 only".

References: https://learn.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-byok-overview, https://learn.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-byok-azure-sql

QUESTION NO: 95

Your company has an Azure subscription that includes two virtual machines named VirMac1 and VirMac2, both of which have the status Stopped (Deallocated). These virtual machines belong to separate resource groups, called ResGroup1 and ResGroup2, respectively.

There are two Azure policies in place, each set up for the virtualMachines resource type. The policy configured for ResGroup1 has a policy definition of Not allowed resource types, whereas the policy for ResGroup2 has a policy definition of Allowed resource types.

Additionally, a Read-only resource lock is applied to VirMac1, and another Read-only resource lock is applied to ResGroup2. Considering this scenario, which of the following statements are TRUE? (Choose all that apply.)

A. You will be able to start VirMac1.

B. You will NOT be able to start VirMac1.

C. You will be able to create a virtual machine in ResGroup2.

D. You will NOT be able to create a virtual machine in ResGroup2.

ANSWER: B C

Explanation:

"You will NOT be able to start VirMac1." is true because a Read-only (CanNotDelete/ReadOnly family) management lock prevents any operation that would modify the resource. Starting a VM changes its runtime state and is treated as a write action against the virtual machine resource, so the platform blocks the start operation while the Read-only lock is in place. This is exactly the kind of change Read-only locks are designed to prevent: you can view the resource, but you can't perform updates that alter it.

"You will be able to create a virtual machine in ResGroup2." is also true because a Read-only lock applied at the resource group scope prevents changes to existing resources in that scope, but it doesn't inherently prevent creating new resources (creation is not an update to an existing resource). Creation is governed primarily by RBAC permissions and Azure Policy. With an "Allowed resource types" policy assigned for the virtualMachines resource type, VM creation is permitted as long as the VM resource type is included in the allowed list and the caller has the required permissions.

References: Lock resources to prevent unexpected changes, Azure Policy definition structure (effects and resource type controls).

QUESTION NO: 96

While troubleshooting a security issue for an Azure Storage account, you have enabled diagnostic logs for the account. Which tool should you use to retrieve these diagnostic logs?

A. the Security & Compliance admin center

B. Azure Security Center

C. Azure Cosmos DB explorer

D. AzCopy

ANSWER: D

Explanation:

AzCopy is the right tool to retrieve Azure Storage diagnostic logs because those logs are written into a Storage account (for example, into a blob container when using Azure Storage logging via Azure Monitor diagnostic settings). Once the logs are in Storage, the practical way to pull them down for investigation is to copy them from the container to a local machine or another destination. AzCopy is Microsoft's supported command-line utility designed specifically for high-performance data transfer to and from Azure Storage, including downloading blobs and entire virtual directories recursively, which fits the common "retrieve logs for analysis" workflow. In troubleshooting scenarios, you can target the container/path where the diagnostic logs are stored and download them quickly, script the retrieval, and integrate it into incident response runbooks. This aligns with Microsoft guidance that you use storage client tools to access and download data stored in Azure Storage, and AzCopy is one of the primary tools Microsoft recommends for moving data in and out of Storage accounts.

References: https://learn.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10, https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/diagnostic-settings

QUESTION NO: 97

You have an Azure Active Directory (Azure AD) tenant with several deleted objects, as shown in the table below:

On May 4, 2020, you attempt to restore the deleted objects using the Azure Active Directory admin center.

Which two objects can you restore? Each correct answer represents a complete solution.

Note: Each correct selection is worth one point.

A. Group1

B. Group2

C. User2

D. User1

ANSWER: B C

Explanation:

In Azure AD, most deleted directory objects are first placed into a soft-deleted state and can be restored from the “Deleted users” or “Deleted groups” areas for a limited retention period. Deleted users can typically be restored within 30 days of deletion, after which they’re permanently deleted and no longer recoverable from the Azure AD admin center. Similarly, Microsoft 365 groups (and their associated resources) are soft-deleted and can be restored within the same 30-day window. Therefore, the objects that were deleted fewer than 30 days before May 4, 2020 and are of a restorable type (user objects and Microsoft 365 groups) can be restored. This aligns with Azure AD’s group and user restore behavior and the documented 30-day soft-delete retention for these object types in the admin experience.

References: https://learn.microsoft.com/en-us/azure/active-directory/fundamentals/active-directory-users-restore, https://learn.microsoft.com/en-us/azure/active-directory/enterprise-users/groups-restore-deleted

QUESTION NO: 98

You have an Azure subscription that includes an Azure key vault and an Azure Storage account. The key vault contains keys that are managed by customers. Additionally, the storage account is set up to utilize these customer-managed keys stored within the key vault.

You intend to store data in Azure using the following services:

Azure Files
Azure Blob Storage
Azure Log Analytics
Azure Table Storage
Azure Queue Storage

Which two of these services support data encryption utilizing the keys stored in the key vault? Each correct selection counts as one point.

NOTE: Each correct selection is worth one point.

A. Queue storage

B. Table storage

C. Azure Files

D. Blob storage

ANSWER: A C D

Explanation:

Azure Storage supports customer-managed keys (CMK) stored in Azure Key Vault for Storage Service Encryption (SSE). When you configure CMK at the storage account level, the encryption settings apply to supported storage services in that account, allowing you to control key rotation and revoke access by managing the key in Key Vault. In this scenario, both Azure Files and Azure Queue Storage support SSE with customer-managed keys from Key Vault, so data written to file shares and queues can be encrypted at rest using those keys. This capability is part of Azure Storage’s CMK integration and is configured on the storage account’s encryption settings, referencing a Key Vault key (or managed HSM key) and using a managed identity to access it. For details on which Azure Storage services support CMK and how the integration works, see Customer-managed keys for Azure Storage encryption and the broader overview of Azure Storage encryption for data at rest.

QUESTION NO: 99

You have initiated a new Azure subscription.

To enable the creation of custom alert rules within Azure Security Center, what are the two necessary steps you must undertake? Each correct choice is a component of the complete solution.

Note: Each accurate selection will earn you a point.

A. Onboard Azure Active Directory (Azure AD) Identity Protection.

B. Create an Azure Storage account.

C. Implement Azure Advisor recommendations.

D. Create an Azure Log Analytics workspace.

E. Upgrade the pricing tier of Security Center to Standard.

ANSWER: D E

Explanation:

Creating custom alert rules in Microsoft Defender for Cloud (formerly Azure Security Center) is done by authoring analytics rules in Microsoft Sentinel, which runs on top of an Azure Log Analytics workspace. That makes creating an Azure Log Analytics workspace a required foundational step, because the rules, queries (KQL), and alert generation are all executed against data stored in that workspace. In addition, you must enable the appropriate paid plan so that the security data and integrations needed for advanced detections and alerting are available; in classic exam wording this is expressed as upgrading the Security Center pricing tier to Standard (now represented as enabling Defender plans in Defender for Cloud). With a workspace in place and the paid tier enabled, you can connect data sources and create scheduled query analytics rules that generate alerts and incidents.

References: https://learn.microsoft.com/en-us/azure/sentinel/quickstart-onboard, https://learn.microsoft.com/en-us/azure/defender-for-cloud/enable-enhanced-security

QUESTION NO: 100

You have an Azure subscription named Sub1 that utilizes Microsoft Defender for Cloud. The subscription is organized in a management group hierarchy as depicted in the following exhibit.

![Management Group Hierarchy](IMAGE_1_URL)

The following table lists the definitions that have been created:

![Definitions Table](IMAGE_2_URL)

You aim to establish a security policy using Defender for Cloud. Which of the following definitions can be used as a security policy?

A. Policy1 only

B. Policy1 and Initiative1 only

C. Initiative1 and Initiative2 only

D. Initiative1, Initiative2, and Initiatives only

E. Policy1, Initiative1, Initiative2, and Initiative3

ANSWER: C

Explanation:

In Microsoft Defender for Cloud, a “security policy” is implemented through Azure Policy assignments (typically initiatives) at a supported scope. Defender for Cloud uses Azure Policy initiatives to evaluate resources and surface recommendations; when you configure a security policy in Defender for Cloud, you’re effectively working with Azure Policy/initiative assignments that apply to a management group or subscription scope. Therefore, any definition that is an Azure Policy initiative (policy set definition) can be used as the basis of a Defender for Cloud security policy, provided it’s assignable at the relevant scope in your management group hierarchy. In the scenario, the definition that is applicable across the hierarchy and can be used to establish the security policy is the one that is defined/assigned at the Tenant Root Group scope, because that scope can govern all child management groups and subscriptions beneath it, including Sub1. This aligns with how Defender for Cloud security policies are centrally managed and inherited through management group scopes using Azure Policy initiatives.

References: Security policy in Microsoft Defender for Cloud, Azure Policy initiative definitions.

QUESTION NO: 101

You have a web application hosted on an on-premises server that is accessed through the URL https://www.contoso.com.

You plan to migrate this web application to Azure while maintaining the existing URL, https://www.contoso.com.

What is the first step you should take to enable HTTPS for the Azure-hosted web application?

A. Export the public key from the on-premises server and save the key as a P7b file.

B. Export the private key from the on-premises server and save the key as a PFX file that is encrypted by using TripleDES.

C. Export the public key from the on-premises server and save the key as a CER file.

D. Export the private key from the on-premises server and save the key as a PFX file that is encrypted by using AES256.

ANSWER: B

Explanation:

To enable HTTPS on the Azure-hosted version of https://www.contoso.com while keeping the same hostname, you must be able to bind a TLS/SSL certificate for that custom domain in Azure (for example, in Azure App Service). The prerequisite for that binding is having the certificate in a format Azure can import that includes the private key. In practice, that means exporting the existing certificate from the on-premises server as a password-protected PFX that contains the private key. Azure App Service’s private certificate requirements specify that the certificate must be provided as a PFX and that it must contain the private key; public-key-only exports (CER/P7B) are insufficient for HTTPS binding because they cannot complete the server-side TLS handshake. Additionally, App Service requires the PFX to be encrypted using TripleDES (not AES256) for import compatibility. Once you have the correct PFX, you can upload it to the App Service and then create the TLS/SSL binding for the custom domain.

References: Configure an SSL certificate in Azure App Service, Private certificate requirements

QUESTION NO: 102 - (HOTSPOT)

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

Name	Type	Attached to	NSG
NSG1	Network security group (NSG)	VM5	*Not applicable*
NSG2	Network security group (NSG)	Subnet1	*Not applicable*
Subnet1	Subnet	*Not applicable*	*Not applicable*
VM5	Virtual machine	Subnet1	NSG1

An IP address of 10.1.0.4 is assigned to VM5. VM5 does not have a public IP address.

VM5 has just in time (JIT) VM access configured as shown in the following exhibit.

You enable JIT VM access for VM5.

NSG1 has the inbound rules shown in the following exhibit.

Priority	Name	Port	Protocol	Source	Destination	Action
100	SecurityCenter-JITRule-...	3389	Any	Any	10.1.0.4	Allow
1000	SecurityCenter-JITRule_341...	3389	Any	Any	10.1.0.4	Deny
1001	RDP	3389	TCP	Any	Any	Allow
65000	AllowVnetInBound	Any	Any	VirtualNetwork	VirtualNetwork	Allow
65001	AllowAzureLoadBalancerIn...	Any	Any	AzureLoadBalancer	Any	Allow
65500	DenyAllInBound	Any	Any	Any	Any	Deny

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
Deleting the security rule that has a priority of 100 will revoke the approved JIT access request.	○	○
Remote Desktop access to VM5 is blocked.	○	○
An Azure Bastion host will enable Remote Desktop access to VM5 from the internet.	○	○

ANSWER:

Statements	Yes	No
Deleting the security rule that has a priority of 100 will revoke the approved JIT access request.	○	○
Remote Desktop access to VM5 is blocked.	○	○
An Azure Bastion host will enable Remote Desktop access to VM5 from the internet.	○	○

Explanation:

Just-in-time (JIT) VM access in Microsoft Defender for Cloud works by dynamically modifying network security group rules to open a management port (like RDP 3389) only for a limited time and only from the approved source IP range. When a JIT request is approved, Defender for Cloud creates a temporary inbound NSG allow rule (often at a high priority such as 100) that permits the requested traffic for the approved duration. If you delete that temporary rule while the approval window is still active, the NSG no longer contains the allow entry that was granting the time-bound access, so the effective result is that the approved JIT access is revoked because the network path is blocked again at the NSG layer. This is consistent with the JIT design of “open only when needed” by adding and removing NSG rules.

Also, VM5 has no public IP address. Without a public IP, there is no direct internet-reachable endpoint on the VM’s NIC, so inbound RDP from the public internet is effectively blocked regardless of whether an NSG rule exists, because external clients cannot route directly to a private IP like 10.1.0.4. Finally, Azure Bastion provides secure RDP/SSH connectivity to VMs without public IPs by brokering the session through the Bastion host inside the virtual network and using the Azure portal over HTTPS. It does not “enable Remote Desktop access to the VM from the internet” in the sense of exposing port 3389 publicly to the VM; instead it avoids public exposure of RDP entirely while still allowing administrators to connect.

References: Just-in-time VM access (Defender for Cloud), Azure Bastion overview.

QUESTION NO: 103

You are troubleshooting a security issue for an Azure Storage account.

You enable Azure Storage Analytics logs and archive it to a storage account.

What should you use to retrieve the diagnostics logs?

A. Azure Cosmos DB explorer

B. SQL query editor in Azure

C. AzCopy

D. the Security admin center

ANSWER: C

Explanation:

AzCopy is the right tool to retrieve Azure Storage Analytics diagnostics logs when they're archived into a storage account. Storage Analytics writes logs as blob objects into a special container (commonly `$logs`) within the storage account, using a path structure based on service type and date/time. Because these logs are stored as standard blobs, you retrieve them the same way you would retrieve any blob data: by listing and downloading/copying the blob objects from the container. AzCopy is Microsoft's supported command-line utility optimized for high-performance data transfer to and from Azure Storage, and it works well for pulling down large sets of log files efficiently (including recursive downloads and filtering by path/prefix). This makes it a practical choice during incident response or troubleshooting when you need to quickly collect the raw log files for offline analysis or ingestion into another tool. See AzCopy documentation at https://learn.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10 and Storage Analytics logging details at https://learn.microsoft.com/en-us/rest/api/storageservices/storage-analytics-logging.

QUESTION NO: 104

You have an Azure subscription that includes a resource group named RG1 and a security group called ServerAdmins. RG1 consists of 10 virtual machines, a virtual network named VNET1, and a network security group (NSG) known as NSG1. ServerAdmins have the capability to access the virtual machines through RDP.

Your task is to configure NSG1 in such a way that it permits RDP access to the virtual machines for a period no longer than 60 minutes, specifically when any member of ServerAdmins requests access.

What configuration should you implement?

A. an Azure policy assigned to RG1

B. a just in time (JIT) VM access policy in Azure Security Center

C. an Azure Active Directory (Azure AD) Privileged Identity Management (PIM) role assignment

D. an Azure Bastion host on VNET1

ANSWER: B

Explanation:

a just in time (JIT) VM access policy in Azure Security Center is the right configuration because JIT is designed to keep management ports like RDP (3389) closed by default in the NSG and only open them on-demand for a limited, approved time window. When a user requests access, JIT updates the relevant NSG rules (such as in NSG1) to allow inbound RDP for the requested duration (for example, up to 60 minutes), and then automatically removes/reverts the rule when the time expires. This directly matches the requirement to permit RDP only when a member requests it and to enforce a maximum

access duration. JIT also integrates with identity and approval workflows and can scope access by source IP and port, reducing exposure compared to permanently open RDP. In Microsoft Defender for Cloud (formerly Azure Security Center), you configure JIT policies per VM (or at scale) and specify the allowed ports and maximum time, which is exactly the control needed here.

References: https://learn.microsoft.com/en-us/azure/defender-for-cloud/just-in-time-access-overview, https://learn.microsoft.com/en-us/azure/defender-for-cloud/just-in-time-access-usage

QUESTION NO: 105

You manage an Azure subscription that includes a web application named App1. Users need to have the option to authenticate using either a Google identity or a Microsoft identity. How can you add Google as an identity provider in Azure Active Directory (Azure AD)?

Select the two pieces of information you need to configure. Each correct choice contributes to part of the solution.

Each correct selection is worth one point.

A. a tenant name

B. a tenant ID

C. the endpoint URL Of an application

D. a client ID

E. a client secret

ANSWER: D E

Explanation:

To add Google as an identity provider for sign-in, you must register/configure an application in Google's developer console and then provide Azure with the OAuth 2.0 credentials that represent that Google app. In practice, Azure needs the *client ID* to identify the Google application during the authorization flow, and the *client secret* to securely redeem authorization codes and obtain tokens from Google. These two values are the core pieces of configuration that enable the trust relationship and token exchange between Azure (or the Azure-hosted authentication layer) and Google as the external identity provider. Once configured, users can choose Google for authentication alongside Microsoft identities, and the platform can validate and process the resulting tokens appropriately. This is the standard OAuth/OpenID Connect pattern used by Azure when integrating with social identity providers such as Google.

References: Configure Google authentication for App Service, OAuth 2.0 authorization code flow (Microsoft identity platform concepts)

QUESTION NO: 106

You must configure WebApp1 to align with the specified data and application requirements.

Which two actions should you take? Each correct answer contributes to part of the overall solution.

Note: Each correct selection is worth one point.

A. Upload a public certificate.

B. Turn on the HTTPS Only protocol setting.

C. Set the Minimum TLS Version protocol setting to 1.2

D. Change the pricing tier of the App Service plan.

ANSWER: A C

Explanation:

Upload a public certificate is correct because Azure App Service supports uploading public certificates (CER) to the app, which can then be used by the application for scenarios such as validating remote endpoints, establishing trust chains, or enabling certificate-based features within the app code. This aligns with common “data and application requirements” where an app must trust a specific public CA or partner certificate without requiring a private key.

Set the Minimum TLS Version protocol setting to 1.2 is correct because App Service allows you to enforce a minimum inbound TLS version for the app’s HTTPS endpoints. Requiring TLS 1.2 is a widely accepted security baseline and helps meet compliance/security requirements by preventing clients from negotiating older, weaker protocol versions. This setting directly hardens the application’s transport security posture without requiring code changes.

References: https://learn.microsoft.com/en-us/azure/app-service/configure-ssl-certificate and https://learn.microsoft.com/en-us/azure/app-service/configure-ssl-bindings

QUESTION NO: 107

You need to create an Azure Key Vault using PowerShell and ensure that deleted objects are retained for a specified period of 90 days. Which two parameters should you use to meet this requirement? (Choose two.)

A. EnabledForDeployment

B. EnablePurgeProtection

C. EnabledForTemplateDeployment

D. EnableSoftDelete

ANSWER: B D

Explanation:

To retain deleted Key Vault objects for a defined period, you must use the Key Vault soft-delete capability and configure the retention window. Soft delete ensures that when keys, secrets, or certificates are deleted, they remain recoverable for the configured retention duration rather than being immediately removed. In PowerShell, this is enabled via the soft-delete parameter, and the retention behavior is governed by the vault’s soft-delete settings (including the retention period). Additionally, purge protection is used to prevent permanent deletion (purging) of soft-deleted objects until the retention period expires, which helps enforce the retention requirement and protects against accidental or malicious purges. Therefore, using the parameters that enable soft delete and enable purge protection is the correct approach to ensure deleted objects are retained for the required 90 days and cannot be purged early. For details on Key Vault soft delete and purge protection behavior and configuration, see Soft-delete overview and the PowerShell cmdlet documentation for creating a vault, including relevant parameters, at New-AzKeyVault.

QUESTION NO: 108

You need to configure WebApp1 to meet the data and application requirements. 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 public certificate.

B. Turn on the HTTPS Only protocol setting.

C. Set the Minimum TLS Version protocol setting to 1.2.

D. Change the pricing tier of the App Service plan.

E. Turn on the Incoming client certificates protocol setting.

ANSWER: A C

Explanation:

To meet common Azure App Service security requirements around protecting data in transit and enforcing modern cryptography, you should configure the app to require a current TLS version and ensure the app can present a certificate for HTTPS. Setting the Minimum TLS Version protocol setting to 1.2 enforces that clients negotiate TLS 1.2 or later, which helps prevent the use of older, weaker protocols and aligns with Microsoft guidance for App Service TLS configuration. In addition, uploading a public certificate is required when you need the web app to use a specific TLS/SSL certificate for inbound HTTPS (for example, a custom domain certificate), enabling the app to serve HTTPS traffic with that certificate.

These two actions together address the typical “data” requirement (encrypt traffic with strong TLS) and “application” requirement (configure the web app with the appropriate certificate material for TLS). For more details, see Azure App Service TLS settings and certificate management guidance in Microsoft documentation: Configure an SSL certificate in Azure App Service and Secure a custom DNS name with a TLS/SSL binding in Azure App Service.

QUESTION NO: 109

You have an Azure subscription that uses Microsoft Defender for Cloud. You have accounts for the following cloud services:

• Alibaba Cloud

• Amazon Web Services (AWS)

• Google Cloud Platform (GCP)

What can you add to Defender for Cloud?

A. AWS only

B. Alibaba Cloud and AWS only

C. Alibaba Cloud and GCP only

D. AWS and GCP only

E. Alibaba Cloud, AWS, and GCP

ANSWER: D

Explanation:

Microsoft Defender for Cloud supports connecting non-Azure environments so you can get centralized security posture management and threat protection across clouds. Specifically, Defender for Cloud provides native multicloud connectors for both Amazon Web Services and Google Cloud Platform. After you create the AWS or GCP connector, Defender for Cloud can ingest resource inventory and security signals, assess configurations against recommendations, and surface findings in the Defender for Cloud portal alongside your Azure resources. This multicloud onboarding is part of Defender for Cloud's Cloud Security Posture Management (CSPM) and workload protections, enabling unified visibility and governance across supported clouds. At the time of writing, Alibaba Cloud isn't a supported cloud connector in Defender for Cloud in the same way AWS and GCP are, so you can't add it as a first-class multicloud environment. For implementation details and the supported multicloud connectors, see Microsoft's documentation on connecting AWS and GCP accounts to Defender for Cloud and managing multicloud security posture.

References: Microsoft Docs: Connect your AWS accounts to Microsoft Defender for Cloud, Microsoft Docs: Connect your GCP project to Microsoft Defender for Cloud