Exam

AZ-305

Designing Microsoft Azure Infrastructure Solutions

Free demo	40 sample questions
Full exam	393 questions with answers and explanations
Contents	Questions with answers and explanations

Exam domains

Domain	Exam weight	Questions
D1 Design identity, governance, and monitoring solutions	25–30%	12
D2 Design data storage solutions	20–25%	10
D3 Design business continuity solutions	15–20%	5
D4 Design infrastructure solutions	30–35%	13

This is a free demo from Dumpsgate: 40 of the exam's 393 questions. **Get the full AZ-305 exam** with every question, answer and explanation.

Question 1

Fill in the blank · Domain: D1 Design identity, governance, and monitoring solutions › Design governance

You are configuring Azure Policy to enforce that all virtual machines in your subscription must have a specific tag applied at creation time. What property of the Azure Policy definition should you use to deny resource creation if the tag is missing?

Answer: deny

Explanation

Use the policy definition's `effect` property with the value `deny`. When the policy condition detects that a virtual machine is missing the required tag, the deny effect blocks its creation, enforcing the requirement at creation time.

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/effect-deny
https://learn.microsoft.com/en-us/azure/governance/policy/concepts/definition-structure-basics

✓ **Box 1: deny** deny
`deny` is the policy effect that prevents a resource request from being created when it violates the policy condition. In this scenario, it blocks creation of a virtual machine that lacks the required tag.

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/effect-deny

---

Question 2

Single choice · Domain: D1 Design identity, governance, and monitoring solutions › Design governance

You need to implement a solution where Azure resources are automatically protected from accidental deletion by enforcing a policy that prevents deletion of resources with a specific tag. Which approach implements this requirement?

A. Create an Azure Policy with a deny effect that prevents deletion of resources with the required tag
B. Apply CanNotDelete locks to each resource individually
C. Use Azure RBAC to deny the Delete action for all users except administrators
D. Configure Azure Backup to immediately restore deleted resources

Answer: A. Create an Azure Policy with a deny effect that prevents deletion of resources with the required tag

Explanation

Microsoft documents that the Azure Policy denyAction effect is specifically designed to block actions against matching resources at scale. Currently, the supported action is DELETE, making it appropriate for protecting tagged critical resources from accidental deletion. Microsoft even provides an example where resources with a particular tag are protected from delete requests. Microsoft Learn The ordinary deny effect is primarily evaluated against resource creation or update requests that violate policy conditions; it is not the correct effect for directly blocking an existing resource's delete operation.

Free demo · full exam: 393 questions

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/effect-deny-action

Question 3

Single choice · Domain: D1 Design identity, governance, and monitoring solutions › Design governance

Your company is deploying a multi-region application that requires consistent configuration across all regions. You need a way to define infrastructure, deploy it uniformly, and ensure compliance with organizational standards across all regions. Which Azure service is best suited for this?

A. Azure Deployment Stacks with Template Specs and Azure Policy
B. Terraform scripts managed through Azure DevOps
C. Resource Manager templates (ARM templates) for each region with manual updates
D. Azure Portal UI for deploying each region individually

Answer: A. Azure Deployment Stacks with Template Specs and Azure Policy

Explanation

Azure Blueprints would have been the historical/intended answer A, because it packaged repeatable infrastructure and governance artifacts. However, Microsoft is retiring Azure Blueprints. As of July 31, 2026, new blueprint definitions and versions can no longer be created, and Microsoft recommends migrating to Azure Deployment Stacks and Template Specs.

https://learn.microsoft.com/en-us/azure/governance/blueprints/blueprint-retirement

Question 4

Single choice · Domain: D1 Design identity, governance, and monitoring solutions › Design solutions for logging and monitoring

Your application stores customer documents in Azure Blob Storage and processes them with Azure Functions. You need to implement audit logging to track who accessed which documents and when. Which solution provides comprehensive audit trails with minimal code changes?

A. Enable Storage Analytics logging, configure Azure Monitor diagnostic settings to send logs to Log Analytics, and query for blob read/write operations
B. Implement custom logging code in each Azure Function that writes to a SQL database
C. Use Azure Blob Storage access tiers to track document access history
D. Configure Azure AD Conditional Access to log all access attempts

Answer: A. Enable Storage Analytics logging, configure Azure Monitor diagnostic settings to send logs to Log Analytics, and query for blob read/write operations

Free demo · full exam: 393 questions

Dumpsgate

Question 5

Single choice · Domain: D1 Design identity, governance, and monitoring solutions › Design authentication and authorization solutions

Your company has a requirement to implement role-based access control (RBAC) for an Azure subscription shared by multiple teams. Each team should manage their own resources within their resource group, but cannot modify other teams' resources or modify subscription-level settings. Which approach implements this requirement?

A. Create a custom role that includes only resource-level permissions, assign it at the resource group scope for each team

B. Assign Owner role to all team members at the subscription level

C. Use Azure AD groups with no role assignments and manage access through resource group owners manually

D. Implement Contributor role for all teams at the resource group level

Answer: A. Create a custom role that includes only resource-level permissions, assign it at the resource group scope for each team

Explanation

A custom role can limit permissions to the resource operations the teams need, and assigning it at each team's resource group scope confines those permissions to that group. This supports team-level management while preventing access to other teams' resource groups and avoiding subscription-level permissions. Configure the role's allowed actions carefully so it does not include permissions to change subscription settings.

https://learn.microsoft.com/en-us/Azure/role-based-access-control/role-assignments
https://learn.microsoft.com/azure/role-based-access-control/custom-roles
https://learn.microsoft.com/en-us/azure/role-based-access-control/role-definitions

✓ **A** Create a custom role that includes only resource-level permissions, assign it at the resource group scope for each team
A custom role lets you specify the permissions teams need, and assigning it at the team's resource group scope limits where those permissions apply. With permissions restricted to resource management, teams cannot manage other teams' resources or change subscription-level settings.

https://learn.microsoft.com/en-us/Azure/role-based-access-control/role-assignments
https://learn.microsoft.com/azure/role-based-access-control/custom-roles
https://learn.microsoft.com/en-us/azure/role-based-access-control/role-definitions

✗ **B** Assign Owner role to all team members at the subscription level
The Owner role at subscription scope gives team members broad access across the subscription, including the ability to manage other teams' resources and grant access. That directly violates both isolation requirements.

https://learn.microsoft.com/en-us/Azure/role-based-access-control/role-assignments
https://learn.microsoft.com/en-us/azure/role-based-access-control/built-in-roles/management-and-governance

✗ **C** Use Azure AD groups with no role assignments and manage access through resource group owners manually
Groups alone do not grant Azure resource access; permissions require role assignments. Relying on resource group owners to manage access manually does not implement the requested RBAC permissions for each team.

https://learn.microsoft.com/en-us/Azure/role-based-access-control/role-assignments
https://learn.microsoft.com/en-us/azure/role-based-access-control/rbac-and-directory-admin-roles

Free demo · full exam: 393 questions

✗ **D** Implement Contributor role for all teams at the resource group level
Resource-group scope would keep each assignment from applying to other teams' resource groups, but the built-in Contributor role allows management of all resources within that scope. It does not satisfy the requirement for permissions limited to the specific resource-level actions the teams should perform.

https://learn.microsoft.com/en-us/Azure/role-based-access-control/role-assignments
https://learn.microsoft.com/en-us/azure/role-based-access-control/built-in-roles/management-and-governance

---

Question 6

Single choice · Domain: D1 Design identity, governance, and monitoring solutions › Design authentication and authorization solutions

You have an Azure subscription that is linked to a Microsoft Entra tenant. The tenant contains a user named User1 that is assigned a Microsoft Entra ID P2 license.

The tenant has security defaults disabled.

You need to ensure that User1 is prompted for multifactor authentication (MFA) when the user attempts to access Azure resources from untrusted networks. The solution must minimize administrative effort.

What should you use?

A. a user risk policy
B. an access package
C. an identity governance policy
D. A Conditional Access policy using the Network

Answer: D. A Conditional Access policy using the Network

Explanation

he requirement is specifically to require MFA when User1 accesses Azure resources from untrusted networks. Microsoft recommends using Conditional Access network/location conditions for this scenario. You can define trusted network locations and configure the policy to require MFA when the user is outside those trusted locations. Microsoft Learn A sign-in risk policy is different: it triggers controls such as MFA when Microsoft Entra ID Protection determines that a particular sign-in has a medium/high risk level. Simply signing in from an untrusted network does not automatically mean that the sign-in is risky.

https://learn.microsoft.com/en-us/entra/identity/conditional-access/concept-assignment-network

---

Free demo · full exam: 393 questions

Question 7

Hotspot (dropdowns) · Domain: D1 Design identity, governance, and monitoring solutions › Design authentication and authorization solutions

Your on premises network contains an Active Directory Domain Services (AD DS) forest. The forest contains a top level domain, three child domains, and an on-premises server named Server1.

You have a Microsoft Entra tenant. Server1 uses Microsoft Entra Connect Sync to replicate all the user objects from the three child domains to the tenant.

New contractors and employees are onboarded manually by using the Workday cloud-based human resources (HR) application.

You plan to automatically provision accounts for new users in one of the on-premises child domains and the Microsoft Entra tenant. The provisioning logic for the employees will be distinct from the provisioning logic for the contractors.

You need to identify the following:

- The minimum number of apps to register in the Microsoft Entra tenant
- The minimum number of Microsoft Entra Connect provisioning agents to deploy

The solution must minimize implementation effort.

What should you identify? To answer, select the appropriate options in the answer area.

NOTE Each correct selection is worth one point.

The minimum number of apps to register in the Microsoft Entra tenant: 1 | 2 | 3

The minimum number of Microsoft Entra Connect provisioning agents to deploy: 1 | 2 | 3

Answer

- The minimum number of apps to register in the Microsoft Entra tenant: **2**
- The minimum number of Microsoft Entra Connect provisioning agents to deploy: **1**

Explanation

Register two apps so employees and contractors can use separate provisioning configurations from Workday, as their provisioning logic differs. Deploy one Microsoft Entra Connect provisioning agent: the requirement is to provision accounts in only one child domain, and the question asks for the minimum rather than a high-availability deployment. [6][10][1]

https://learn.microsoft.com/entra/identity/hybrid/cloud-sync/how-to-prerequisites
https://learn.microsoft.com/en-us/azure/active-directory/saas-apps/workday-inbound-tutorial?source=docs
https://learn.microsoft.com/en-us/entra/identity/app-provisioning/plan-cloud-hr-provision

✗ **Box 1: 1** 1 app

Free demo · full exam: 393 questions

A single Workday provisioning app would not provide separate provisioning configurations for employees and contractors, even though the scenario requires distinct logic for those groups. Separate app instances allow each group to have its own provisioning configuration. [6][10]

https://learn.microsoft.com/en-us/azure/active-directory/saas-apps/workday-inbound-tutorial?source=docs
https://learn.microsoft.com/en-us/entra/identity/app-provisioning/plan-cloud-hr-provision

✓ **Box 1: 2** 2 apps
Two apps allow separate Workday provisioning configurations for employees and contractors, matching the requirement for distinct provisioning logic. [6][10]

https://learn.microsoft.com/en-us/azure/active-directory/saas-apps/workday-inbound-tutorial?source=docs
https://learn.microsoft.com/en-us/entra/identity/app-provisioning/plan-cloud-hr-provision

✗ **Box 1: 3** 3 apps
The scenario has two worker groups with distinct provisioning logic, so two apps meet the requirement. A third app is not needed for the described provisioning configurations. [6][10]

https://learn.microsoft.com/en-us/azure/active-directory/saas-apps/workday-inbound-tutorial?source=docs
https://learn.microsoft.com/en-us/entra/identity/app-provisioning/plan-cloud-hr-provision

✓ **Box 2: 1** 1 provisioning agent
The target on-premises location is one child domain, and one provisioning agent is the minimum needed to connect the provisioning service to the on-premises environment. Additional agents can support availability or scale, but those are not required by the minimum-agent question. [1][3]

https://learn.microsoft.com/entra/identity/hybrid/cloud-sync/how-to-prerequisites
https://learn.microsoft.com/en-us/entra/identity/hybrid/cloud-sync/what-is-cloud-sync

✗ **Box 2: 2** 2 provisioning agents
Two agents are not the minimum for provisioning to the one specified child domain. This could provide redundancy, but the scenario asks to minimize implementation effort and does not require high availability. [1][3]

https://learn.microsoft.com/entra/identity/hybrid/cloud-sync/how-to-prerequisites
https://learn.microsoft.com/en-us/entra/identity/hybrid/cloud-sync/what-is-cloud-sync

✗ **Box 2: 3** 3 provisioning agents
Three agents are not required to provision accounts in the single specified child domain. Deploying more agents than the minimum would add unnecessary implementation effort in this scenario. [1][3]

https://learn.microsoft.com/entra/identity/hybrid/cloud-sync/how-to-prerequisites
https://learn.microsoft.com/en-us/entra/identity/hybrid/cloud-sync/what-is-cloud-sync

Question 8

Yes / No · Domain: D1 Design identity, governance, and monitoring solutions › Design governance

You have five Azure subscriptions as shown in the following table.

Name	Management group	Tag
Subscription1	Marketing	InternalCharge: True
Subscription2	IT	InternalCharge: False
Subscription3	Finance	InternalCharge: False
Subscription4	IT	InternalCharge: True
Subscription5	Marketing	InternalCharge: True

You have the resource groups shown in the following table.

Name	Subscription

RG1	Subscription2
RG2	Subscription1
RG3	Subscription4
RG4	Subscription3
RG5	Subscription1

You have the Azure resources shown in the following table.

Name	Type	Resource group
VM1	Virtual machine	RG1
Storage1	Storage account	RG2
VM2	Virtual machine	RG3
KV1	Azure Key Vault	RG4
SQL1	Azure SQL Database	RG5

You create the Azure Policy assignments shown in the following table.

Name	Scope
Resource tagging	Finance
Deny creating virtual machines	Marketing

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

NOTE: Each correct selection is worth one point.

Statement	Yes	No
The resource tagging policy assignment will be applied to KV1.	☐	☐
Cost management data for the InternalCharge = True tag will include VM1.	☐	☐
Moving Subscription4 to the Marketing management group will cause VM2 to shut down.	☐	☐

Answer

- The resource tagging policy assignment will be applied to KV1. **Yes**
- Cost management data for the InternalCharge = True tag will include VM1. **No**
- Moving Subscription4 to the Marketing management group will cause VM2 to shut down. **No**

Explanation

The Resource tagging assignment is scoped to the Finance management group, which contains Subscription3 and its resource KV1. VM1 is in Subscription2, whose InternalCharge tag is False, so its cost data is not associated with InternalCharge=True. Moving Subscription4 under Marketing would bring the deny policy into scope for VM2, but a deny policy blocks creation of matching resources; it does not shut down an existing VM.

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/scope
https://learn.microsoft.com/en-us/azure/cost-management-billing/costs/enable-tag-inheritance
https://learn.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage

✓ **Statement 1** The resource tagging policy assignment will be applied to KV1.
The Resource tagging assignment is scoped to Finance. Subscription3 is in the Finance management group, and KV1 is in a resource group within Subscription3, so the assignment scope includes KV1.

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/scope

✗ **Statement 2** Cost management data for the InternalCharge = True tag will include VM1.
VM1 is in RG1 in Subscription2, whose InternalCharge tag is False. Therefore, VM1 is not associated with the InternalCharge=True tag in the scenario. Tags do not automatically propagate from a subscription to its resources; cost tag inheritance is a separate capability.

https://learn.microsoft.com/en-us/azure/cost-management-billing/costs/enable-tag-inheritance
https://learn.microsoft.com/en-us/azure/azure-resource-manager/resource-group-using-tags

✗ **Statement 3** Moving Subscription4 to the Marketing management group will cause VM2 to shut down.
Moving Subscription4 under Marketing would put VM2 within the scope of the Deny creating virtual machines assignment. A deny policy prevents creation of matching virtual machines; it does not shut down VM2, which already exists.

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/scope
https://learn.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage

---

Question 9

Single choice · Domain: D1 Design identity, governance, and monitoring solutions › Design solutions for logging and monitoring

You have an Azure resource group that contains a storage account named storage1.

You need to configure a diagnostic setting for storage1 to automatically route transaction logs directly to a third-party security information and event management (SIEM) solution.

What should you use as the destination?

A. Send to Log Analytics workspace
B. Send to a storage account
C. Stream to an event hub

Answer: C. Stream to an event hub

Explanation

To send storage1 transaction logs directly to a third-party SIEM, configure the diagnostic setting to stream the logs to an event hub. The SIEM can consume the streamed data from Event Hubs, making this the appropriate destination for near-real-time integration. A Log Analytics workspace and a storage account are destinations for analysis and retention, respectively, rather than direct streaming to an external SIEM.

https://learn.microsoft.com/en-us/azure/azure-monitor/data-collection/stream-monitoring-data-event-hubs
https://learn.microsoft.com/en-us/azure/azure-monitor/platform/stream-monitoring-data-event-hubs
https://learn.microsoft.com/en-us/entra/identity/monitoring-health/howto-stream-logs-to-event-hub?tabs=SumoLogic

✗ **A** Send to Log Analytics workspace
This sends diagnostic data to a Log Analytics workspace for querying and analysis, not directly to the third-party SIEM. It does not meet the requirement to stream the transaction logs to an external destination.

https://learn.microsoft.com/hi-in/azure/azure-monitor/platform/diagnostic-settings

Free demo · full exam: 393 questions

https://learn.microsoft.com/en-us/azure/azure-monitor/data-collection/diagnostic-settings-faq

✗ **B** Send to a storage account
This archives diagnostic data in a storage account for retention or later processing. It does not provide the direct streaming destination needed for the SIEM.

https://learn.microsoft.com/hi-in/azure/azure-monitor/platform/diagnostic-settings
https://learn.microsoft.com/en-us/azure/azure-monitor/data-collection/diagnostic-settings-faq

✓ **C** Stream to an event hub
Diagnostic settings can stream logs to an event hub, which a third-party SIEM can consume. This supports routing storage1 transaction logs directly to the external security monitoring solution.

https://learn.microsoft.com/en-us/azure/azure-monitor/data-collection/stream-monitoring-data-event-hubs
https://learn.microsoft.com/en-us/azure/azure-monitor/platform/stream-monitoring-data-event-hubs
https://learn.microsoft.com/en-us/entra/identity/monitoring-health/howto-stream-logs-to-event-hub?tabs=SumoLogic

---

Question 10

Drag and drop · Domain: D1 Design identity, governance, and monitoring solutions › Design solutions for logging and monitoring

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

Name	Description
VM1	A Windows 11 image used as a jump box by multiple administrators
VM2	An Ubuntu image used as a syslog server for network devices

You need to recommend a logging solution for the virtual machines. The solution must meet the following requirements:

- Operating system logs from VM1 must be collected and stored in a Log Analytics workspace.
- Syslog data logs from VM2 must be archived to a storage account.

What solution should you include in the recommendation for each virtual machine? To answer, drag the appropriate solutions to the correct virtual machines.

Each solution may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Name	Description
VM1	A Windows 11 image used as a jump box by multiple administrators
VM2	An Ubuntu image used as a syslog server for network devices

Answer

- VM1 → **Azure Monitor Agent**
- VM2 → **Linux Diagnostic Extension**

Explanation

Use Azure Monitor Agent on VM1 to collect Windows operating system events and send them to a Log Analytics workspace. Use the Linux Diagnostic Extension on VM2 to collect its syslog data and archive it to

a storage account. These choices match each VM's operating system and the required destination for its logs.

https://learn.microsoft.com/en-us/azure/azure-monitor/vm/data-collection-windows-events
https://learn.microsoft.com/en-us/azure/virtual-machines/extensions/diagnostics-linux
https://learn.microsoft.com/en-us/azure/azure-monitor/vm/data-collection

✓ **VM1: Azure Monitor Agent** Azure Monitor Agent
VM1 runs Windows 11, and Azure Monitor Agent can collect Windows event logs for Azure Monitor. Configure data collection to send the operating system logs to the required Log Analytics workspace.

https://learn.microsoft.com/en-us/azure/azure-monitor/vm/data-collection-windows-events
https://learn.microsoft.com/en-us/azure/azure-monitor/vm/data-collection

✓ **VM2: Linux Diagnostic Extension** Linux Diagnostic Extension
VM2 runs Ubuntu and receives syslog data from network devices. The Linux Diagnostic Extension can collect Linux syslog data and archive it to an Azure Storage account, meeting the requirement for VM2.

https://learn.microsoft.com/en-us/azure/virtual-machines/extensions/diagnostics-linux

---

Question 11

Hotspot (dropdowns) · Domain: D1 Design identity, governance, and monitoring solutions › Design governance

You need to design an Azure policy that will implement the following functionality:

- For new resources, assign tags and values that match the tags and values of the resource group to which the resources are deployed.
- For existing resources, identify whether the tags and values match the tags and values of the resource group that contains the resources.
- For any non-compliant resources, trigger auto-generated remediation tasks to create missing tags and values. The solution must use the principle of least privilege.

What should you include in the design? To answer, select the appropriate options in the answer area.

a. NOTE: Each correct selection is worth one point.

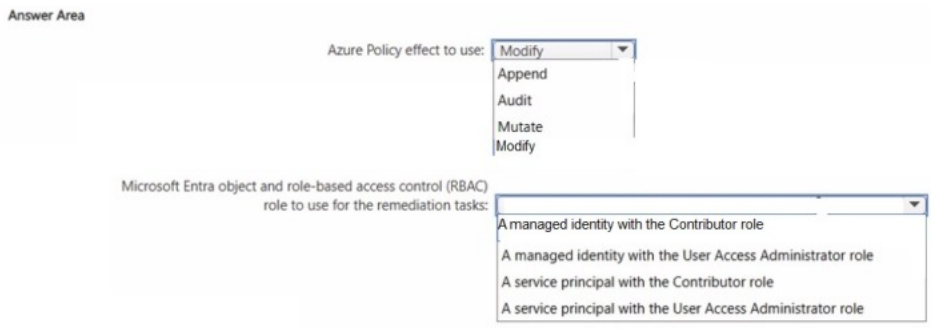

Azure Policy effect to use: Append | Audit | Mutate | Modify

Microsoft Entra object and role-based access control (RBAC) role to use for the remediation tasks: A managed identity with the Contributor role | A managed identity with the User Access Administrator role | A service principal with the Contributor role | A service principal with the User Access Administrator role

Answer

- Azure Policy effect to use: **Modify**
- Microsoft Entra object and role-based access control (RBAC) role to use for the remediation tasks: **A managed identity with the Contributor role**

Explanation

Use the Modify effect to add or update tags on new resources and to remediate existing non-compliant resources. Assign a managed identity the Contributor role so remediation can change resource tags; this provides the needed permissions without granting the broader role-assignment permissions of User Access Administrator.

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/effect-modify
https://learn.microsoft.com/en-us/azure/governance/policy/how-to/remediate-resources?WT.mc_id=AZ-MVP-5004810

✗ **Box 1: Append** Append
Append can add tags during a resource create or update request, but it does not support remediation of existing resources. The scenario requires auto-generated remediation tasks for existing non-compliant resources as well.

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/effect-append
https://learn.microsoft.com/en-us/azure/governance/policy/how-to/remediate-resources?WT.mc_id=AZ-MVP-5004810

✗ **Box 1: Audit** Audit
Audit identifies and reports non-compliant resources but does not add or correct tags. It therefore cannot meet the requirement to create missing tags and values through remediation.

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/evaluate-impact
https://learn.microsoft.com/en-us/azure/governance/policy/concepts/compliance-states

✗ **Box 1: Mutate** Mutate
Mutate is not an Azure Policy effect in the choices' Azure Policy effect model. The supported Modify effect can change tags and remediate existing resources.

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/effect-modify

✓ **Box 1: Modify** Modify
Modify can add or update tags on resources during creation or update and supports remediation tasks for existing resources that are non-compliant. This matches both the new-resource and existing-resource requirements.

https://learn.microsoft.com/en-us/azure/governance/policy/concepts/effect-modify
https://learn.microsoft.com/en-us/azure/governance/policy/how-to/remediate-resources?WT.mc_id=AZ-MVP-5004810

✓ **Box 2: A managed identity with the Contributor role** A managed identity with the Contributor role
Azure Policy remediation uses a managed identity, and the identity needs permissions to modify resource tags. Contributor provides those resource modification permissions without the role-assignment privileges associated with User Access Administrator.

https://learn.microsoft.com/en-us/azure/governance/policy/how-to/remediate-resources?WT.mc_id=AZ-MVP-5004810

✗ **Box 2: A managed identity with the User Access Administrator role** A managed identity with the User Access Administrator role
A managed identity is appropriate for remediation, but User Access Administrator is not the least-privilege role for changing tags. Its role-assignment permissions are unnecessary for this task.

https://learn.microsoft.com/en-us/azure/governance/policy/how-to/remediate-resources?WT.mc_id=AZ-MVP-5004810

✗ **Box 2: A service principal with the Contributor role** A service principal with the Contributor role
Contributor can provide the resource modification permissions needed for tags, but the remediation identity should be a managed identity rather than a service principal. Azure Policy remediation uses a managed identity for the assigned policy or initiative.

https://learn.microsoft.com/en-us/azure/governance/policy/how-to/remediate-resources?WT.mc_id=AZ-MVP-5004810

✗ **Box 2: A service principal with the User Access Administrator role** A service principal with the User

Free demo · full exam: 393 questions

Access Administrator role
This uses a service principal instead of the managed identity used for Azure Policy remediation, and User Access Administrator grants unnecessary role-assignment permissions. It does not follow the least-privilege requirement for changing tags.
https://learn.microsoft.com/en-us/azure/governance/policy/how-to/remediate-resources?WT.mc_id=AZ-MVP-5004810

Question 12

Single choice · Domain: D1 Design identity, governance, and monitoring solutions › Design governance

You have a Microsoft Entra tenant named contoso.com that has a security group named Group1. Group! is configured for assigned membership. Group! has 50 members, including 20 guest users You need to recommend a solution for evaluating the membership of Group1. The solution must meet the following requirements:

- The evaluation must be repeated automatically every three months.
- Every member must be able to report whether they need to be in Group1.
- Users who report that they do not need to be in Group1 must be removed from Group! automatically.
- Users who do not report whether they need to be in Group1 must be removed from Group1 automatically.

What should you include in the recommendation?

A. Create an access review.
B. Implement Microsoft Entra Privileged Identity Management (PIM).
C. Change the Membership type of Group1 to Dynamic User
D. Implement Microsoft Entra ID Protection.

Answer: A. Create an access review.

Explanation

An access review is designed to periodically evaluate whether users still need access to a group. Configure the review to recur every three months, have each member review their own access, and apply the outcome automatically so members who deny needing access—or do not respond—are removed. This directly meets the requirements for Group1, including its guest members. [2, 11]

https://learn.microsoft.com/en-us/entra/id-governance/create-access-review-pim-for-groups
https://learn.microsoft.com/en-us/entra/id-governance/create-access-review

✓ **A** Create an access review.
An access review can be configured to recur every three months and let members review their own group membership. Its review settings can automatically remove users who say they do not need access and apply a specified outcome to nonresponders, such as removing them. [2, 11]

https://learn.microsoft.com/en-us/entra/id-governance/create-access-review-pim-for-groups
https://learn.microsoft.com/en-us/entra/id-governance/create-access-review

✗ **B** Implement Microsoft Entra Privileged Identity Management (PIM).
PIM for groups manages privileged access, such as eligibility and activation, rather than conducting the requested recurring review in which every member reports whether they still need membership. An access review is the feature that supports periodic membership evaluation and decisions for

nonresponders. [2, 5]

https://learn.microsoft.com/en-us/entra/id-governance/create-access-review-pim-for-groups
https://learn.microsoft.com/en-us/entra/id-governance/privileged-identity-management/concept-pim-for-groups

✗ **C** Change the Membership type of Group1 to Dynamic User
Dynamic membership determines group membership from user attributes and rules; it does not ask members to confirm whether they need access. It therefore cannot meet the reporting and nonresponse-removal requirements of this evaluation. [4]

https://learn.microsoft.com/en-us/entra/identity/users/groups-change-type

✗ **D** Implement Microsoft Entra ID Protection.
ID Protection detects and investigates identity-based risks, such as compromised accounts; it is not a process for periodically asking group members to confirm continued need for membership. It does not provide the requested access-review responses and removal handling. [1]

https://learn.microsoft.com/en-us/entra/id-protection/overview-identity-protection

---

Question 13

Single choice · Domain: D2 Design data storage solutions › Design data storage solutions for semi-structured and unstructured data

You are designing an application that requires both relational data (customer information) and unstructured data (document files). The application must query relational data frequently and store documents efficiently.

What is the most appropriate storage combination?

A. Azure SQL Database for relational data and Azure Blob Storage for unstructured data with a SAS token link stored in SQL
B. Azure Cosmos DB for all data types with custom serialization
C. Azure Data Lake Storage Gen2 for both relational and unstructured data using Parquet format
D. Azure SQL Database with FILESTREAM for all documents

Answer: A. Azure SQL Database for relational data and Azure Blob Storage for unstructured data with a SAS token link stored in SQL

Explanation

Azure SQL Database is designed for relational data, making it appropriate for frequently queried customer information. Azure Blob Storage is suited to storing unstructured documents efficiently, while SQL can retain a SAS-token link to provide access to each document. This combination uses a storage service suited to each data type rather than forcing both into one system.

https://learn.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql-db
https://learn.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/understand-data-store-models
https://learn.microsoft.com/en-us/azure/storage/blobs/authorize-data-operations-cli
https://learn.microsoft.com/en-us/azure/storage/common/storage-sas-overview

✓ **A** Azure SQL Database for relational data and Azure Blob Storage for unstructured data with a SAS token link stored in SQL
Azure SQL Database supports the relational customer data that the application queries frequently, and Blob Storage is designed for unstructured files. A SAS token can grant scoped access to a blob, so storing a link or reference in SQL separates document storage from relational records.

https://learn.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql-db
https://learn.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/understand-data-store-models
https://learn.microsoft.com/en-us/azure/storage/blobs/authorize-data-operations-cli
https://learn.microsoft.com/en-us/azure/storage/common/storage-sas-overview

✗ **B** Azure Cosmos DB for all data types with custom serialization
Cosmos DB is a NoSQL database, so it is not the most appropriate choice for frequently queried relational customer data when a relational database is required. Custom serialization also does not provide the distinct, efficient document storage service called for by the scenario.

https://learn.microsoft.com/en-us/azure/cosmos-db/modeling-data
https://learn.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/understand-data-store-models
https://learn.microsoft.com/en-us/azure/cosmos-db/faq

✗ **C** Azure Data Lake Storage Gen2 for both relational and unstructured data using Parquet format
Data Lake Storage is suited to storing files and analytics data, but it does not provide the relational database capabilities needed for frequent customer-data queries. Parquet is a file format, not a substitute for a relational database engine.

https://learn.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/understand-data-store-models

✗ **D** Azure SQL Database with FILESTREAM for all documents
This puts both relational records and documents in the database rather than using a storage service suited to unstructured files. The scenario calls for efficient document storage, for which Blob Storage is the more appropriate fit; FILESTREAM is also not a feature of Azure SQL Database as offered in this option.

https://learn.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql-db
https://learn.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/understand-data-store-models

Question 14

Single choice · Domain: D2 Design data storage solutions › Design data storage solutions for semi-structured and unstructured data

Your organization needs to implement a cost-effective solution for storing historical data that is rarely accessed (cold data). The data must be retrieved within 24 hours if needed, retained for 10 years for compliance, and stored on Azure. Which storage solution optimizes for this use case?

A. Azure Blob Storage Archive access tier with lifecycle management policies to automatically transition data
B. Azure Blob Storage Hot access tier with manual archival to Data Box
C. Azure File Shares with snapshot retention
D. Azure SQL Database with long-term backup retention

Answer: A. Azure Blob Storage Archive access tier with lifecycle management policies to automatically transition data

Explanation

Azure Blob Storage Archive is designed for rarely accessed data and has the lowest storage cost among the blob access tiers, making it suitable for historical cold data. Lifecycle management policies can automatically transition blobs into the Archive tier, reducing manual work. Archived data must be rehydrated before it can be read, and that process can take hours, aligning with the requirement to retrieve it within 24 hours. The Archive tier and lifecycle policies support this cost-conscious approach to

long-term retention; the organization should ensure its retention configuration meets the 10-year compliance requirement.

https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-policy-access-tiers
https://learn.microsoft.com/en-us/azure/storage/blobs/archive-blob
https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview
https://learn.microsoft.com/en-us/azure/storage/blobs/archive-rehydrate-to-online-tier

✓ **A** Azure Blob Storage Archive access tier with lifecycle management policies to automatically transition data
The Archive tier is intended for data that is rarely accessed and offers low storage costs, while lifecycle management policies can automatically move data into that tier. Retrieving archived data requires rehydration, which can take hours, so the 24-hour retrieval window is compatible with this option. Configure retention to preserve the data for the required 10 years.

https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-policy-access-tiers
https://learn.microsoft.com/en-us/azure/storage/blobs/archive-blob
https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview
https://learn.microsoft.com/en-us/azure/storage/blobs/archive-rehydrate-to-online-tier

✗ **B** Azure Blob Storage Hot access tier with manual archival to Data Box
The Hot tier is optimized for frequently accessed data, so keeping cold historical data there is not the cost-effective choice. Manual transfer to Data Box also does not provide the stated automatic tier-transition approach for ongoing storage.

https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-policy-access-tiers
https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview

✗ **C** Azure File Shares with snapshot retention
Azure Files snapshots are not the low-cost archival access tier described for rarely accessed historical data. This option does not best match the requirement to optimize cost for cold data retained over a long period.

https://learn.microsoft.com/en-us/azure/backup/azure-file-share-support-matrix

✗ **D** Azure SQL Database with long-term backup retention
Long-term retention in Azure SQL is for retaining database backups, rather than storing general historical cold data as an archival dataset. It therefore does not fit the stated data-storage use case as well as Blob Storage Archive.

https://learn.microsoft.com/en-us/azure/azure-sql/database/long-term-retention-overview?view=azuresql
https://learn.microsoft.com/en-us/azure/azure-sql/database/long-term-backup-retention-configure?view=azuresql

Question 15

Single choice · Domain: D2 Design data storage solutions › Design data storage solutions for relational data

You have an Azure subscription that contains 10 Azure SQL databases in an elastic pool. Each database is 600 GB. You need to ensure that you can increase the size of each database to 15 TB. The solution must minimize database downtime. What should you do?

A. Enable geo-replication for each database.
B. Remove the elastic pool. Update each database to the Hyperscale service tier
C. Update the elastic pool to the Premium service tier.
D. Remove the elastic poo1. Update each database to the Business Critical service tier.

Answer: B. Remove the elastic pool. Update each database to the Hyperscale service tier

Dumpsgate

Question 16

Single choice · Domain: D2 Design data storage solutions › Design data storage solutions for semi-structured and unstructured data

You have an Azure subscription.

You plan to deploy an app that requires an Azure Data Lake Storage Gen2 account.

You need to recommend which settings to modify during the storage account deployment. The solution must ensure that the data in the account is protected if an attacker gains access to the physical disk that stores the data.

Which two settings should you recommend? To answer, select the appropriate settings in the answer area.

NOTE: Each correct selection is worth one point.

.

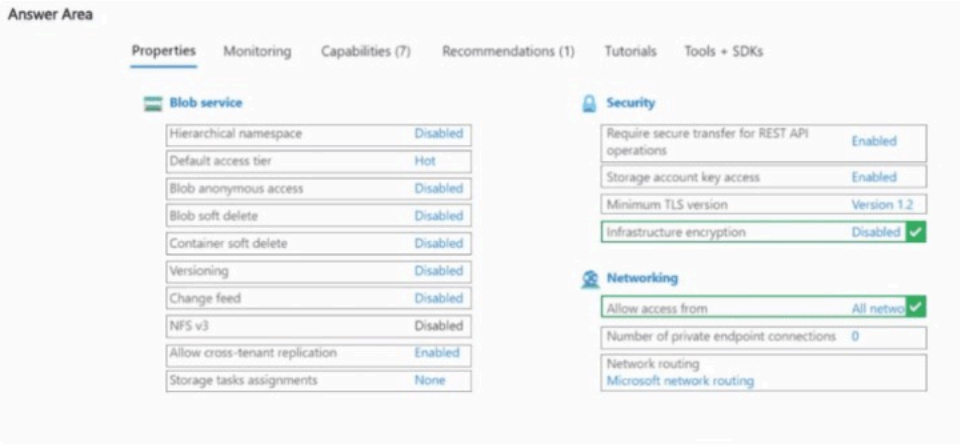

A. Infrastructure encryption and Allow access from
B. Infrastructure encryption and Minimum TLS version
C. Storage account key access and Allow access from
D. Minimum TLS version and Storage account key access

Answer: A. Infrastructure encryption and Allow access from

Explanation

Enable infrastructure encryption to add an additional encryption layer at the storage infrastructure level, protecting data at rest even if someone obtains access to the physical disk. Also restrict the storage account's network access by changing Allow access from, which limits who can reach the account over the network; it does not itself encrypt the disk. Together, these are the two settings specified by the correct answer.

https://learn.microsoft.com/en-us/azure/storage/common/infrastructure-encryption-enable?tabs=portal
https://learn.microsoft.com/en-us/azure/storage/common/storage-network-security-set-default-access

✓ **A** Infrastructure encryption and Allow access from
Infrastructure encryption adds an additional layer of encryption at rest, which addresses the requirement involving access to the physical disk. Allow access from can be restricted to limit network reachability to the storage account, providing a separate access-control measure.

https://learn.microsoft.com/en-us/azure/storage/common/infrastructure-encryption-enable?tabs=portal
https://learn.microsoft.com/en-us/azure/storage/common/storage-network-security-set-default-access

✗ **B** Infrastructure encryption and Minimum TLS version
Infrastructure encryption addresses protection of data at rest on the physical disk. Minimum TLS version controls the security of connections in transit, so it does not provide the network-access restriction specified by the correct setting.

https://learn.microsoft.com/en-us/azure/storage/common/infrastructure-encryption-enable?tabs=portal
https://learn.microsoft.com/azure/storage/common/transport-layer-security-configure-minimum-version

✗ **C** Storage account key access and Allow access from
Allow access from can restrict network reachability, but changing storage account key access does not enable the additional at-rest encryption layer needed for the physical-disk scenario. Infrastructure encryption is the relevant encryption setting.

https://learn.microsoft.com/en-us/azure/storage/common/infrastructure-encryption-enable?tabs=portal
https://learn.microsoft.com/en-us/azure/storage/common/storage-network-security-set-default-access
https://learn.microsoft.com/en-us/azure/storage/common/secure-storage

✗ **D** Minimum TLS version and Storage account key access
Minimum TLS version concerns data in transit, and storage account key access controls an authentication method. Neither setting enables infrastructure-level encryption to protect data at rest if the physical disk is accessed.

https://learn.microsoft.com/en-us/azure/storage/common/infrastructure-encryption-enable?tabs=portal
https://learn.microsoft.com/azure/storage/common/transport-layer-security-configure-minimum-version
https://learn.microsoft.com/en-us/azure/storage/common/secure-storage

Dumpsgate

Question 17

Hotspot (dropdowns) · Domain: D2 Design data storage solutions › Design data storage solutions for relational data

You have an on-premises app named App1 that stores data in a Microsoft SQL Server database named DB1.

App1 uses cross-database transactions.

You plan to move App1 to Azure.

You need to recommend which services to use to host DB1 in Azure and migrate DB1 to Azure. The solution must meet the following requirements:

- Ensure that once App1 moves to Azure, DB1 users can authenticate by using their Microsoft Entra credentials.
- Minimize administrative effort to patch and update the database host once App1 moves to Azure.
- Minimize administrative effort during the migration of DB1 to Azure.
- Minimize costs.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

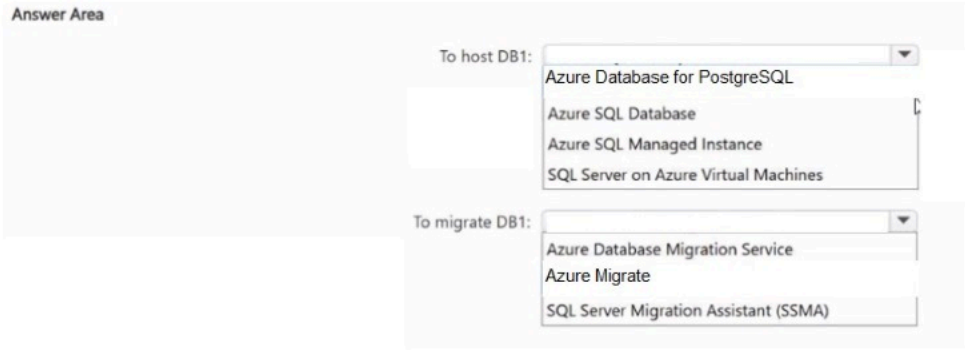

To host DB1: Azure Database for PostgreSQL | Azure SQL Database | Azure SQL Managed Instance | SQL Server on Azure Virtual Machines

To migrate DB1: Azure Database Migration Service | Azure Migrate | SQL Server Migration Assistant (SSMA)

Answer

- To host DB1: **Azure SQL Managed Instance**
- To migrate DB1: **Azure Database Migration Service**

Explanation

Use Azure SQL Managed Instance to preserve SQL Server compatibility for App1's cross-database transactions while reducing the effort to manage and patch the database host compared with running SQL Server on a virtual machine. Managed Instance also supports Microsoft Entra authentication. Use Azure Database Migration Service to move DB1 with a purpose-built database migration service, minimizing migration administration; the other choices either do not fit the SQL Server workload and requirements or are less suitable for this migration.

https://learn.microsoft.com/en-us/data-migration/sql-server/managed-instance/guide
https://learn.microsoft.com/azure/dms/
https://learn.microsoft.com/azure/azure-sql/migration-guides/managed-instance/sql-server-to-managed-instance-overview?view=azuresql

✗ **Box 1: Azure Database for PostgreSQL** Azure Database for PostgreSQL
This is a PostgreSQL service, but DB1 is a Microsoft SQL Server database. Choosing it would require changing database platforms and would not preserve the SQL Server compatibility needed for App1's cross-database transactions.

✗ **Box 1: Azure SQL Database** Azure SQL Database
Azure SQL Database is a managed service, but it does not provide the SQL Server instance-level compatibility needed for this workload's cross-database transactions. Managed Instance is the Azure SQL option that better matches the SQL Server application requirements while reducing host maintenance.

https://learn.microsoft.com/en-us/data-migration/sql-server/managed-instance/guide
https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-azure-ad-only-authentication?view=azuresql

✓ **Box 1: Azure SQL Managed Instance** Azure SQL Managed Instance
Managed Instance is the appropriate SQL Server-compatible managed service for an application that uses cross-database transactions. As a managed service, it reduces the administrative effort of patching and updating the database host, and it supports Microsoft Entra authentication.

https://learn.microsoft.com/en-us/data-migration/sql-server/managed-instance/guide
https://learn.microsoft.com/azure/azure-sql/migration-guides/managed-instance/sql-server-to-managed-instance-overview?view=azuresql
https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/frequently-asked-questions-faq?view=azuresql

✗ **Box 1: SQL Server on Azure Virtual Machines** SQL Server on Azure Virtual Machines
A SQL Server virtual machine can provide SQL Server compatibility, but you remain responsible for more of the operating-system and database-host maintenance, including patching. That conflicts with the requirement to minimize administrative effort, and it is not the preferred managed option for minimizing costs.

https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/frequently-asked-questions-faq?view=azuresql-vm

✓ **Box 2: Azure Database Migration Service** Azure Database Migration Service
Azure Database Migration Service is designed to migrate databases to Azure SQL targets, including Managed Instance. It provides a database-focused migration path that helps minimize administrative effort during the move.

https://learn.microsoft.com/en-us/data-migration/sql-server/managed-instance/guide
https://learn.microsoft.com/azure/dms/
https://learn.microsoft.com/azure/azure-sql/migration-guides/managed-instance/sql-server-to-managed-instance-overview?view=azuresql

✗ **Box 2: Azure Migrate** Azure Migrate
Azure Migrate is not the best choice for this scenario because the task is to migrate a SQL Server database to Managed Instance with minimal database-migration administration. Azure Database Migration Service is the purpose-built service for that database migration.

https://learn.microsoft.com/azure/dms/

✗ **Box 2: SQL Server Migration Assistant (SSMA)** SQL Server Migration Assistant (SSMA)
SSMA is not the appropriate tool for migrating this SQL Server database to Managed Instance; it is primarily used to assess and migrate databases from other database platforms to SQL Server. Azure Database Migration Service is the matching service for this SQL Server-to-Managed Instance migration.

https://learn.microsoft.com/en-us/ssms/migrate/migrate-sql-server-azure-sql
https://learn.microsoft.com/azure/dms/

Question 18

Single choice · Domain: D2 Design data storage solutions › Design data integration

You have multiple on-premises locations. The locations host IoT endpoints that generate real-time telemetry data.

You have an Azure subscription.

You need to process the telemetry data and provide real-time insights. The solution must minimize development effort.

What should you use?

A. Azure Data Factory
B. Azure Data Lake Analytics
C. Log Analytics
D. Azure Stream Analytics

Answer: D. Azure Stream Analytics

Explanation

Azure Stream Analytics is designed to process streaming data and produce real-time insights, which matches the IoT telemetry scenario. It supports low-effort stream processing, including no-code options, helping minimize development effort. Therefore, it best meets the requirements.

https://learn.microsoft.com/en-us/azure/stream-analytics/no-code-stream-processing

✗ **A** Azure Data Factory
Azure Data Factory is primarily for data integration and orchestrating data movement and transformation, including moving data between on-premises locations and Azure. It is not the best fit for continuously processing telemetry streams to provide real-time insights with minimal development effort.

https://learn.microsoft.com/en-us/azure/data-factory/data-factory-move-data-between-onprem-and-cloud
https://learn.microsoft.com/en-us/azure/data-factory/introduction

✗ **B** Azure Data Lake Analytics
Azure Data Lake Analytics is not the service identified for low-effort real-time stream processing. The scenario calls for processing incoming telemetry as a stream and generating insights promptly, which is the role of Azure Stream Analytics.

✗ **C** Log Analytics
Log Analytics is used to query and analyze log data in Azure Monitor, rather than serving as the stream-processing service for IoT telemetry. It does not best satisfy the requirement to process the live data stream with minimal development effort.

https://learn.microsoft.com/en-us/azure/azure-monitor/logs/log-analytics-overview?WT.mc_id=AZ-MVP-4033351
https://learn.microsoft.com/en-us/azure/azure-monitor/overview

✓ **D** Azure Stream Analytics
Azure Stream Analytics processes streaming data and supports real-time analytics, making it appropriate for telemetry generated by IoT endpoints. Its no-code stream-processing capability can reduce the development effort required.

https://learn.microsoft.com/en-us/azure/stream-analytics/no-code-stream-processing

Question 19

Single choice · Domain: D2 Design data storage solutions › Design data integration

You have an Azure subscription.

You are designing a data solution that will process millions of items nightly. The solution must meet the following requirements:

- Support integration with Microsoft Dataverse.
- Support dynamic data masking.
- Support auto scaling.

What should you include in the solution?

A. Microsoft Fabric
B. Azure Databricks
C. Microsoft Power Bl
D. Azure Synapse Analytics

Answer: D. Azure Synapse Analytics

Explanation

Azure Synapse Analytics is the best fit because it combines SQL data warehousing capabilities for large-scale workloads with dynamic data masking and scalable compute. Its SQL pools support dynamic data masking, and compute resources can be managed to scale for nightly processing. Synapse also supports integration with Dataverse, meeting all three requirements.

https://learn.microsoft.com/en-us/azure/synapse-analytics/sql/dynamic-data-masking-overview
https://learn.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-manage-compute-overview
https://learn.microsoft.com/en-us/azure/synapse-analytics/sql/overview-features
https://learn.microsoft.com/en-us/azure/synapse-analytics/overview-faq

✗ **A** Microsoft Fabric
Fabric supports dynamic data masking in its data warehouse, but this question's intended solution is Azure Synapse Analytics for the Azure subscription workload. The listed references do not establish that Fabric meets the full combination of Dataverse integration and the required auto-scaling for this scenario.

https://learn.microsoft.com/en-us/fabric/data-warehouse/howto-dynamic-data-masking

✗ **B** Azure Databricks
Azure Databricks is designed for large-scale data processing, but the cited material does not show that it provides the required dynamic data masking and Dataverse integration as a unified fit. Synapse has documented SQL dynamic data masking and scalable SQL pool compute.

https://learn.microsoft.com/en-us/azure/synapse-analytics/sql/dynamic-data-masking-overview
https://learn.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-manage-compute-overview

✗ **C** Microsoft Power BI
Power BI is primarily an analytics and reporting platform, rather than the nightly large-scale data processing solution described. It does not meet the complete set of processing, dynamic masking, and auto-scaling requirements as well as Synapse does.

✓ **D** Azure Synapse Analytics
Synapse SQL supports dynamic data masking, and its SQL pools provide scalable compute for high-volume data workloads. Synapse is also an Azure analytics service that can be integrated with Dataverse,

Free demo · full exam: 393 questions

so it meets the scenario's requirements.

https://learn.microsoft.com/en-us/azure/synapse-analytics/sql/dynamic-data-masking-overview
https://learn.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-manage-compute-overview
https://learn.microsoft.com/en-us/azure/synapse-analytics/sql/overview-features
https://learn.microsoft.com/en-us/azure/synapse-analytics/overview-faq

---

Question 20

Hotspot (dropdowns) · Domain: D2 Design data storage solutions › Design data integration

You have an on-premises datacenter and an Azure subscription. The environment contains 100 databases of various types, including Azure SQL Database, Azure SQL Managed Instances, Azure Cosmos DB, Microsoft SQL Server, MySQL and Oracle. Multiple apps access the databases.

You need to recommend a data solution. The solution must meet the following requirements

- Provide a single searchable catalog of the data in all the databases. The catalog must contain metadata of the data.

a. The solution must minimize storage requirements.

- Provide a central data lake that contains specific data from multiple databases. The solution must include a big data engine that can create machine learning models.

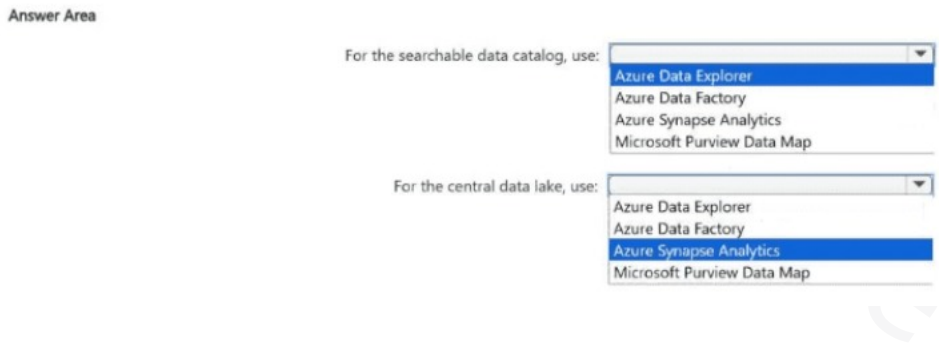

For the searchable data catalog, use: Azure Data Explorer | Azure Data Factory | Azure Synapse Analytics | Microsoft Purview Data Map

For the central data lake, use: Azure Data Explorer | Azure Data Factory | Azure Synapse Analytics | Microsoft Purview Data Map

Answer

- For the searchable data catalog, use: **Microsoft Purview Data Map**
- For the central data lake, use: **Azure Synapse Analytics**

Explanation

Use Microsoft Purview Data Map for the searchable catalog: it scans supported database sources and captures metadata, rather than requiring a separate stored copy of all the database data. Use Azure Synapse Analytics for the central data lake because it can bring together selected data from multiple sources and provides Spark capabilities for creating machine learning models. This pairing addresses the catalog, storage, data lake, and machine learning requirements.

https://learn.microsoft.com/en-us/purview/data-map-scanning-best-practices
https://learn.microsoft.com/en-gb/purview/data-map-scan-data-sources
https://learn.microsoft.com/en-us/azure/purview/overview
https://learn.microsoft.com/en-us/azure/purview/how-to-search-catalog
https://learn.microsoft.com/en-us/azure/synapse-analytics/spark/apache-spark-machine-learning-training
https://learn.microsoft.com/en-us/azure/architecture/data-guide/scenarios/data-lake?source=recommendations
https://learn.microsoft.com/hr-hr/azure/synapse-analytics/machine-learning/what-is-machine-learning

✗ **Box 1: Azure Data Explorer** For the searchable data catalog, use Azure Data Explorer.
Azure Data Explorer is primarily an analytics service for ingesting and querying data, not the searchable metadata catalog required across the varied database sources in this scenario. Microsoft Purview Data Map is designed to catalog data sources and their metadata. [4]

https://learn.microsoft.com/en-us/azure/purview/overview

✗ **Box 1: Azure Data Factory** For the searchable data catalog, use Azure Data Factory.
Azure Data Factory is a data integration service for moving and transforming data; it is not the single searchable metadata catalog requested. Purview Data Map scans sources to populate a catalog, so it avoids treating full data movement as the solution to the catalog requirement. [3][4]

https://learn.microsoft.com/en-gb/purview/data-map-scan-data-sources
https://learn.microsoft.com/en-us/azure/purview/overview

✗ **Box 1: Azure Synapse Analytics** For the searchable data catalog, use Azure Synapse Analytics.
Synapse Analytics can provide analytics and data-lake capabilities, but it is not the dedicated cross-source searchable catalog of data metadata required here. Microsoft Purview Data Map is the appropriate catalog service. [4][6]

https://learn.microsoft.com/en-us/azure/purview/overview
https://learn.microsoft.com/en-us/azure/purview/how-to-search-catalog

✓ **Box 1: Microsoft Purview Data Map** For the searchable data catalog, use Microsoft Purview Data Map.
Purview Data Map discovers and catalogs data sources and their metadata, and the catalog can be searched. Scanning catalogs metadata rather than requiring a full copy of every database, which fits the storage-minimization requirement. [1][3][4][6]

https://learn.microsoft.com/en-us/purview/data-map-scanning-best-practices
https://learn.microsoft.com/en-gb/purview/data-map-scan-data-sources
https://learn.microsoft.com/en-us/azure/purview/overview
https://learn.microsoft.com/en-us/azure/purview/how-to-search-catalog

✗ **Box 2: Azure Data Explorer** For the central data lake, use Azure Data Explorer.
Azure Data Explorer is focused on analytics over ingested data, not on providing the central data lake and big-data machine-learning engine specified in the question. Synapse Analytics offers lake-oriented analytics together with Spark machine-learning capabilities. [2][10]

https://learn.microsoft.com/en-us/azure/synapse-analytics/spark/apache-spark-machine-learning-training
https://learn.microsoft.com/en-us/azure/architecture/data-guide/scenarios/data-lake?source=recommendations

✗ **Box 2: Azure Data Factory** For the central data lake, use Azure Data Factory.
Azure Data Factory can orchestrate data movement and transformation, but it is not itself the central data lake or the big-data engine for building machine-learning models. Synapse Analytics is the suitable choice for those combined requirements. [2][10]

https://learn.microsoft.com/en-us/azure/synapse-analytics/spark/apache-spark-machine-learning-training
https://learn.microsoft.com/en-us/azure/architecture/data-guide/scenarios/data-lake?source=recommendations

✓ **Box 2: Azure Synapse Analytics** For the central data lake, use Azure Synapse Analytics.
Synapse Analytics can support a central lake-based analytics solution using selected data from multiple databases. Its Spark capabilities support training machine-learning models, satisfying the requirement for a big data engine. [2][10][11]

https://learn.microsoft.com/en-us/azure/synapse-analytics/spark/apache-spark-machine-learning-training
https://learn.microsoft.com/en-us/azure/architecture/data-guide/scenarios/data-lake?source=recommendations
https://learn.microsoft.com/hr-hr/azure/synapse-analytics/machine-learning/what-is-machine-learning

✗ **Box 2: Microsoft Purview Data Map** For the central data lake, use Microsoft Purview Data Map.
Purview Data Map catalogs and describes data; it is not a data lake or a big-data processing engine for building machine-learning models. Synapse Analytics provides the lake and Spark-based machine-learning capabilities needed here. [2][4][10]

https://learn.microsoft.com/en-us/azure/synapse-analytics/spark/apache-spark-machine-learning-training
https://learn.microsoft.com/en-us/azure/purview/overview
https://learn.microsoft.com/en-us/azure/architecture/data-guide/scenarios/data-lake?source=recommendations

---

Question 21

Single choice · Domain: D2 Design data storage solutions › Design data storage solutions for relational data

Your company has 50 business units across the globe. The business units operate from 08:00 AM to 06:00 PM from Monday to Friday in their local time zone. Transactions are only processed during business hours.

You have an Azure subscription.

You plan to deploy an app named App1 that will manage the transactions for the business units. App1 will use a separate Azure SQL database for each business unit.

You need to recommend an Azure SQL Database configuration for App1. The solution must meet the following requirements:

- Support Azure Hybrid Benefit licensing.
- Minimize costs.

What should you recommend?

A. a vCore purchasing model and multiple single database instances
B. a DTU purchasing model and multiple single database instances
C. a vCore purchasing model and multiple database instances in an elastic pool
D. a DTU purchasing model and multiple database instances in an elastic pool

Answer: C. a vCore purchasing model and multiple database instances in an elastic pool

Explanation

A vCore purchasing model is required to use Azure Hybrid Benefit, which can reduce licensing costs. An elastic pool lets the 50 separate databases share provisioned compute resources, helping minimize costs when their transaction loads vary rather than peak at the same time; the different business-hour windows across time zones support that pooling approach.

https://learn.microsoft.com/en-us/azure/azure-sql/azure-hybrid-benefit?view=azuresql
https://learn.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-resource-management?view=azuresql
https://learn.microsoft.com/en-ca/azure/azure-sql/database/purchasing-models?view=azuresql

✗ **A** a vCore purchasing model and multiple single database instances
The vCore model supports Azure Hybrid Benefit, but separate single databases do not share compute resources. That misses the cost-saving opportunity of pooling the 50 databases with varying business-hour workloads.

https://learn.microsoft.com/en-us/azure/azure-sql/azure-hybrid-benefit?view=azuresql
https://learn.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-resource-management?view=azuresql

✗ **B** a DTU purchasing model and multiple single database instances
The separate databases do not share compute resources, so this does not take advantage of pooling to minimize costs. In addition, Azure Hybrid Benefit is associated with the vCore purchasing model, not the DTU model.

https://learn.microsoft.com/en-us/azure/azure-sql/azure-hybrid-benefit?view=azuresql

https://learn.microsoft.com/en-us/azure/sql-database/sql-database-what-is-a-dtu
https://learn.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-resource-management?view=azuresql
https://learn.microsoft.com/en-ca/azure/azure-sql/database/purchasing-models?view=azuresql

✓ **C** a vCore purchasing model and multiple database instances in an elastic pool
The vCore model supports Azure Hybrid Benefit, meeting the licensing requirement. The elastic pool allows the separate business-unit databases to share resources, which can lower costs when their transaction demands vary with their local business hours.

https://learn.microsoft.com/en-us/azure/azure-sql/azure-hybrid-benefit?view=azuresql
https://learn.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-resource-management?view=azuresql
https://learn.microsoft.com/en-ca/azure/azure-sql/database/purchasing-models?view=azuresql

✗ **D** a DTU purchasing model and multiple database instances in an elastic pool
An elastic pool can help reduce costs by sharing resources across the databases, but the DTU model does not meet the Azure Hybrid Benefit requirement. The vCore purchasing model is needed for that licensing benefit.

https://learn.microsoft.com/en-us/azure/azure-sql/azure-hybrid-benefit?view=azuresql
https://learn.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-resource-management?view=azuresql
https://learn.microsoft.com/en-ca/azure/azure-sql/database/purchasing-models?view=azuresql

Dumpsgate

Question 22

Hotspot (dropdowns) · Domain: D2 Design data storage solutions › Design data integration

You have a production line that is monitored by using IoT devices.

You are building a project that will process data streams from the IoT devices.

You need to recommend a pipeline solution that will perform the following activities:

- Ingest the data streams from the IoT devices.
- Analyze the data and identify manufacturing defects among items on the production line.
- When a manufacturing defect is identified, flag the item for removal from the production line.

What should you include in the recommendation for each activity? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

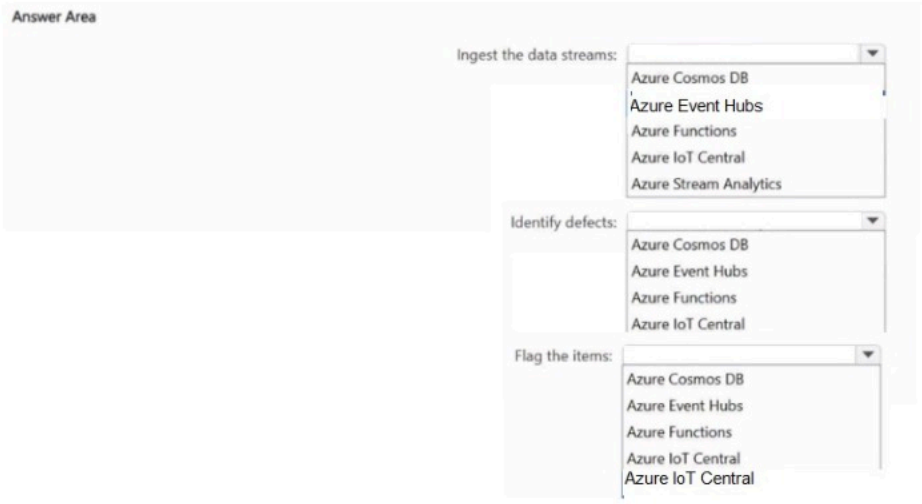

Ingest the data streams: Azure Cosmos DB | Azure Event Hubs | Azure Functions | Azure IoT Central | Azure Stream Analytics

Identify defects: Azure Cosmos DB | Azure Event Hubs | Azure Functions | Azure IoT Central | Azure Stream Analytics

Flag the items: Azure Cosmos DB | Azure Event Hubs | Azure Functions | Azure IoT Central

Answer

- Ingest the data streams: **Azure Event Hubs**
- Identify defects: **Azure Stream Analytics**
- Flag the items: **Azure Functions**

Explanation

Azure Event Hubs is designed for high-throughput, low-latency ingestion of streaming telemetry, including data generated by IoT devices. Microsoft specifically identifies IoT telemetry ingestion as a primary Event Hubs scenario. Microsoft Learn Azure Stream Analytics processes and analyzes streaming data in real time. It can consume the events from Event Hubs and apply queries, filtering, aggregation, and pattern/condition detection to identify manufacturing defects as telemetry arrives. Microsoft Learn Azure Functions provides event-driven code execution. Once the analytics stage identifies a defective

item, a function can execute the required business action—for example, flagging the item or invoking another system that removes it from the production line. Microsoft Learn

https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-about

Question 23

Single choice · Domain: D3 Design business continuity solutions › Design solutions for backup and disaster recovery

You need to implement a disaster recovery solution for a database-heavy application where the secondary region must be ready to handle full production traffic within 30 minutes. The data must be synchronized in near-real-time. Which Azure service combination best suits this requirement?

A. Azure SQL Database with active geo-replication and a pre-deployed scaled-down secondary environment that auto-scales during failover
B. Automated backups to a secondary region with manual restore during disaster
C. Azure Site Recovery for the application tier only, with daily database backups
D. Manual database synchronization using Azure Data Factory jobs every hour

Answer: A. Azure SQL Database with active geo-replication and a pre-deployed scaled-down secondary environment that auto-scales during failover

Explanation

Option A best meets both requirements: active geo-replication keeps a secondary Azure SQL Database updated asynchronously with low replication lag, and a pre-deployed secondary application environment avoids having to build the production environment during an outage. Scaling that environment during failover supports readiness to handle full traffic within the 30-minute target. The other options rely on restores or synchronization intervals that are too slow for near-real-time data protection and rapid recovery.

https://learn.microsoft.com/en-us/azure/azure-sql/database/disaster-recovery-guidance?view=azuresql
https://learn.microsoft.com/en-us/azure/reliability/reliability-sql-database
https://learn.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-configure-portal?view=azuresql

✓ **A** Azure SQL Database with active geo-replication and a pre-deployed scaled-down secondary environment that auto-scales during failover
Active geo-replication continuously copies database changes to a secondary, making it suitable for near-real-time disaster recovery. Keeping the secondary environment deployed in advance, then scaling it during failover, helps meet the 30-minute target to serve full production traffic.

https://learn.microsoft.com/en-us/azure/azure-sql/database/disaster-recovery-guidance?view=azuresql
https://learn.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-configure-portal?view=azuresql
https://learn.microsoft.com/en-us/azure/azure-sql/database/high-availability-disaster-recovery-checklist?view=azuresql

✗ **B** Automated backups to a secondary region with manual restore during disaster
Backups provide a recovery path, but restoring a database during a disaster takes time and does not keep a ready-to-fail-over secondary database synchronized near-real-time. That makes this a poor fit for the 30-minute recovery requirement.

https://learn.microsoft.com/en-us/azure/azure-sql/database/disaster-recovery-guidance?view=azuresql
https://learn.microsoft.com/en-us/azure/azure-sql/database/recovery-using-backups?view=azuresql
https://learn.microsoft.com/en-ca/azure/azure-sql/database/automated-backups-overview?view=azuresql

✗ **C** Azure Site Recovery for the application tier only, with daily database backups

Recovering only the application tier leaves the database dependent on daily backups, so the database could be substantially out of date and would need restoration. This does not provide near-real-time database synchronization or a suitably ready database for the 30-minute target.

https://learn.microsoft.com/en-us/azure/azure-sql/database/disaster-recovery-guidance?view=azuresql
https://learn.microsoft.com/en-us/azure/azure-sql/database/recovery-using-backups?view=azuresql

✗ **D** Manual database synchronization using Azure Data Factory jobs every hour
Hourly jobs can leave the secondary database up to an hour behind, which does not satisfy the near-real-time synchronization requirement. Manual synchronization also adds operational steps during recovery, making the 30-minute target less dependable.

Question 24

Single choice · Domain: D3 Design business continuity solutions › Design for high availability

You are designing a highly available web application using Azure App Service. The application must remain available during Azure region failures and application updates. Currently, it uses a single App Service Plan in one region. Which design pattern ensures high availability with minimal complexity?

A. Deploy App Service Plans in multiple regions using Azure Front Door as the global load balancer
B. Configure auto-scaling rules to increase instance count during failures
C. Use Azure Traffic Manager with weighted routing to distribute traffic
D. Deploy a single App Service Plan across multiple availability zones

Answer: A. Deploy App Service Plans in multiple regions using Azure Front Door as the global load balancer

Explanation

A multi-region design protects the application from a failure affecting its current Azure region: deploy App Service Plans in separate regions and use Azure Front Door to route users to the available application endpoint. Front Door also provides global traffic distribution, making this a straightforward pattern for regional availability while supporting updates. A single-region scaling or zone-based design does not provide the same protection from a region failure.

https://learn.microsoft.com/en-us/azure/app-service/tutorial-multi-region-app
https://learn.microsoft.com/en-us/azure/well-architected/service-guides/azure-front-door
https://learn.microsoft.com/en-us/azure/architecture/web-apps/guides/multi-region-app-service/multi-region-app-service?tabs=paired-regions
https://learn.microsoft.com/en-us/azure/frontdoor/high-availability

✓ **A** Deploy App Service Plans in multiple regions using Azure Front Door as the global load balancer
Separate regional App Service deployments provide a destination if one region fails, and Azure Front Door directs global user traffic to the available region. This directly addresses regional failures and supports a highly available multi-region application design.

https://learn.microsoft.com/en-us/azure/app-service/tutorial-multi-region-app
https://learn.microsoft.com/en-us/azure/well-architected/service-guides/azure-front-door
https://learn.microsoft.com/en-us/azure/architecture/web-apps/guides/multi-region-app-service/multi-region-app-service?tabs=paired-regions
https://learn.microsoft.com/en-us/azure/frontdoor/high-availability

✗ **B** Configure auto-scaling rules to increase instance count during failures
Auto-scaling can add capacity within the existing App Service Plan, but the plan remains in one region. It

Free demo · full exam: 393 questions

therefore does not protect the application from an outage of that region.

https://learn.microsoft.com/en-us/azure/reliability/reliability-app-service
https://learn.microsoft.com/en-us/azure/well-architected/design-guides/disaster-recovery

✗ **C** Use Azure Traffic Manager with weighted routing to distribute traffic
Traffic Manager can distribute DNS traffic across endpoints, but weighted routing alone does not ensure that an unhealthy endpoint is avoided or provide the integrated global application delivery pattern specified here. Azure Front Door with multi-region App Service deployments is the intended pattern for this scenario.

https://learn.microsoft.com/en-us/azure/app-service/web-sites-traffic-manager
https://learn.microsoft.com/en-us/azure/traffic-manager/traffic-manager-configure-weighted-routing-method
https://learn.microsoft.com/en-us/azure/traffic-manager/traffic-manager-load-balancing-azure

✗ **D** Deploy a single App Service Plan across multiple availability zones
Availability zones can improve resilience to failures within a region, but they do not protect against an Azure region failure. The requirement explicitly includes regional failures, so the application needs deployments in multiple regions.

https://learn.microsoft.com/en-us/azure/reliability/reliability-app-service
https://learn.microsoft.com/en-us/azure/well-architected/design-guides/disaster-recovery

Dumpsgate

Question 25

Hotspot (dropdowns) · Domain: D3 Design business continuity solutions › Design for high availability

Your on-premises network contains an Active Directory Domain Services (AD DS) forest. The forest contains servers that run Windows Server. The servers contain file shares that replicate by using Distributed File System (DFS) Replication.

You have multiple remote Windows 11 devices that connect to the on-premises network by using Point-to-Site (P2S) VPN and access the file shares by using drive mappings.

You have an Azure subscription. The subscription contains virtual machines that run Windows Server and are configured as domain controllers. The virtual machines replicate to the on-premises forest. The forest syncs with a Microsoft Entra tenant that is linked to the Azure subscription.

You plan to migrate the on-premises file shares to the Azure virtual machines.

You need to recommend a solution for replicating the file shares to Azure and providing the Windows 11 devices with access to the shares. The solution must ensure that the devices can access the shares by using mapped drive letters via TCP port 445.

What should you use to replicate the shares, and what should the devices use to the access the shares? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

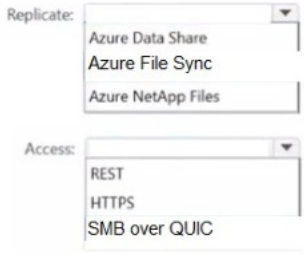

Replicate: Azure Data Share | Azure File Sync | Azure NetApp Files

Access: REST | HTTPS | SMB over TCP

Answer

- Replicate: **Azure File Sync**
- Access: **SMB over TCP**

Explanation

Azure File Sync is appropriate because it can synchronize Windows Server file-server data with Azure Files, and Microsoft supports deploying an additional Azure File Sync server endpoint on a Windows Server VM in Azure. That fits migrating/synchronizing the existing on-premises shares to Azure-hosted Windows file servers. Microsoft Learn For access, the question specifically requires mapped drives over TCP port 445. Standard SMB directly hosted over TCP uses TCP 445. Microsoft Learn SMB over QUIC is not correct for that requirement because Microsoft documents that SMB over QUIC uses UDP 443 by default, not TCP 445. Microsoft Learn Because the Windows 11 clients already use a Point-to-Site VPN, standard SMB access can be carried securely through the VPN while maintaining the mapped-drive experience. Microsoft Learn

https://learn.microsoft.com/en-us/azure/storage/file-sync/file-sync-deployment-guide

Question 26

Hotspot (dropdowns) · Domain: D3 Design business continuity solutions › Design solutions for backup and disaster recovery

You have the Azure subscriptions shown in the following table.

Name	Location	Microsoft Entra tenant
Sub1	East US	contoso.onmicrosoft.com
Sub2	East US	contoso-recovery.onmicrosoft.com

Contoso.onmicrosoft.com contains a user named User1.

You need to deploy a solution to protect against ransomware attacks. The solution must meet the following requirements:

- Ensure that all the resources in Sub1 are backed up by using Azure Backup.
- Require that User1 first be assigned a role for Sub2 before the user can make major changes to the backup configuration.

What should you create in each subscription? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

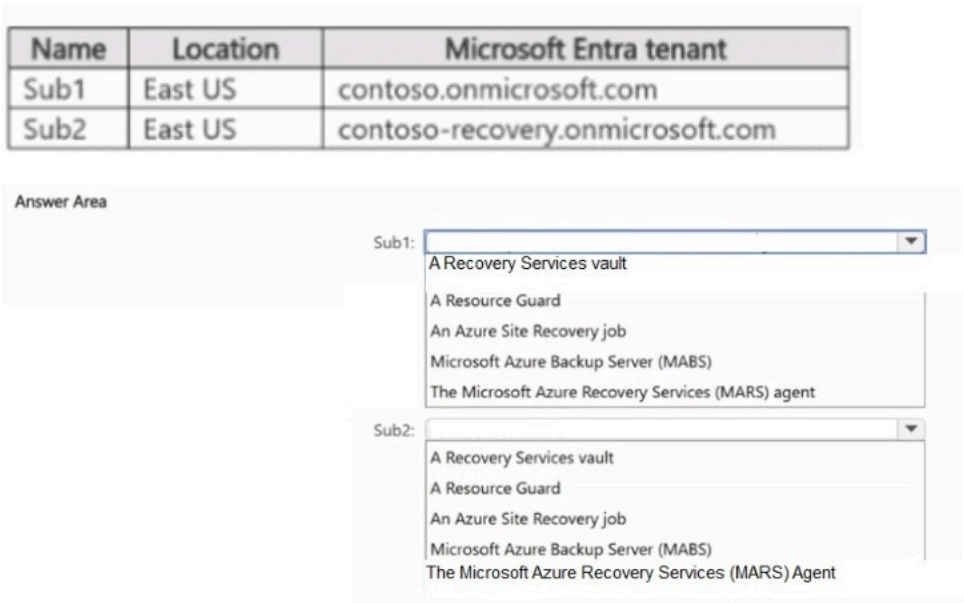

Sub1: A Recovery Services vault | A Resource Guard | An Azure Site Recovery job | Microsoft Azure Backup Server (MABS) | The Microsoft Azure Recovery Services (MARS) agent

Sub2: A Recovery Services vault | A Resource Guard | An Azure Site Recovery job | Microsoft Azure Backup Server (MABS) | The Microsoft Azure Recovery Services (MARS) Agent

Answer

- Sub1: **A Recovery Services vault**
- Sub2: **A Resource Guard**

Explanation

Create a Recovery Services vault in Sub1 because Azure Backup uses a vault to protect supported subscription resources. Create a Resource Guard in Sub2 so that major backup-configuration operations can require authorization through a separately controlled resource; User1 must be assigned an appropriate role for Sub2 before performing those operations. This separation helps protect backups from ransomware-related changes.

https://learn.microsoft.com/en-us/azure/backup/backup-create-recovery-services-vault
https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization-concept?tabs=backup-vault
https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization
https://learn.microsoft.com/azure/backup/azure-backup-data-protection-best-practices

✓ **Sub1: A Recovery Services vault** A Recovery Services vault
Azure Backup uses a Recovery Services vault to store and manage backup data and configuration for protected resources. Creating the vault in Sub1 meets the requirement to back up resources in that subscription.

https://learn.microsoft.com/en-us/azure/backup/backup-create-recovery-services-vault
https://learn.microsoft.com/en-us/azure/backup/backup-azure-manage-vms

✗ **Sub1: A Resource Guard** A Resource Guard
A Resource Guard provides multi-user authorization for protected operations; it is not the vault used to back up Sub1's resources. The Resource Guard belongs in Sub2 for the separate-authorization requirement.

https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization-concept?tabs=backup-vault
https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization

✗ **Sub1: An Azure Site Recovery job** An Azure Site Recovery job
Azure Site Recovery jobs are for disaster-recovery replication and failover operations, not for backing up all resources in Sub1 by using Azure Backup. A Recovery Services vault is the required backup resource.

https://learn.microsoft.com/en-us/azure/backup/backup-create-recovery-services-vault
https://learn.microsoft.com/en-us/azure/backup/backup-support-matrix

✗ **Sub1: Microsoft Azure Backup Server (MABS)** Microsoft Azure Backup Server (MABS)
MABS is backup-server software used for particular workloads, not the Azure Backup vault needed to protect the subscription's resources. It does not meet the requirement as stated for all resources in Sub1.

https://learn.microsoft.com/en-us/azure/backup/backup-support-matrix

✗ **Sub1: The Microsoft Azure Recovery Services (MARS) agent** The Microsoft Azure Recovery Services (MARS) agent
The MARS agent protects supported files, folders, and system state on individual machines. It is not the subscription-level vault required to back up Sub1's resources.

https://learn.microsoft.com/en-us/azure/backup/backup-support-matrix

✗ **Sub2: A Recovery Services vault** A Recovery Services vault
The backup vault for the resources in Sub1 must be created in Sub1, not Sub2. A Recovery Services vault in Sub2 would not provide the separate authorization control requested for major changes.

https://learn.microsoft.com/en-us/azure/backup/backup-create-recovery-services-vault
https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization-concept?tabs=backup-vault

✓ **Sub2: A Resource Guard** A Resource Guard
A Resource Guard enables multi-user authorization for critical backup operations. Placing it in Sub2, which is associated with a different Microsoft Entra tenant, separates control from the backup vault so User1 must first receive the required role for Sub2.

https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization-concept?tabs=backup-vault

Free demo · full exam: 393 questions

https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization
https://learn.microsoft.com/azure/backup/azure-backup-data-protection-best-practices

✗ **Sub2: An Azure Site Recovery job** An Azure Site Recovery job
An Azure Site Recovery job is used for disaster-recovery actions, not for requiring separate authorization before major Azure Backup configuration changes. A Resource Guard provides that authorization control.

https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization-concept?tabs=backup-vault
https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization

✗ **Sub2: Microsoft Azure Backup Server (MABS)** Microsoft Azure Backup Server (MABS)
MABS is backup-server software and does not provide the separate authorization control required for changes to the backup configuration. That role is served by a Resource Guard.

https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization-concept?tabs=backup-vault
https://learn.microsoft.com/en-us/azure/backup/backup-support-matrix

✗ **Sub2: The Microsoft Azure Recovery Services (MARS) Agent** The Microsoft Azure Recovery Services (MARS) Agent
The MARS agent is used to back up supported data on individual machines; it does not control authorization for critical operations on the Sub1 backup configuration. A Resource Guard in Sub2 meets that requirement.

https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization-concept?tabs=backup-vault
https://learn.microsoft.com/en-us/azure/backup/backup-support-matrix

Dumpsgate

Question 27

Single choice · Domain: D3 Design business continuity solutions › Design solutions for backup and disaster recovery

You have an Azure subscription that contains an Azure Data Lake Storage Gen2 account named storage1 and an Azure Synapse Analytics workspace named Workspace1. Workspace! contains a serverless SQL pool and an Apache Spark pool.

You need to recommend a disaster recovery solution for storage1.

What should you include in the recommendation?

A. vaulted backup of Azure blobs
B. operational backup of Azure blobs
C. geo-backup in Azure Synapse Analytics
D. asynchronous replication

Answer: D. asynchronous replication

Explanation

Recommend asynchronous replication for storage1. It replicates the storage data to a secondary region, providing a recovery path if the primary region is unavailable; this addresses disaster recovery for the Data Lake Storage Gen2 account itself. Synapse pool backup features do not provide disaster recovery for that storage account.

https://learn.microsoft.com/en-us/azure/storage/common/storage-disaster-recovery-guidance

✗ **A** vaulted backup of Azure blobs
A vaulted backup provides a backup copy for recovery, but it is not the asynchronous cross-region replication solution recommended for regional disaster recovery of storage1. The requirement is to protect the storage account against a regional outage.

Free demo · full exam: 393 questions

https://learn.microsoft.com/en-us/azure/backup/blob-backup-overview

✗ **B** operational backup of Azure blobs
Operational backup supports recovery of blob data from backup points, rather than providing the asynchronous replication-based regional recovery path needed here. The recommendation should address disaster recovery for the storage account.

https://learn.microsoft.com/en-us/azure/backup/blob-backup-overview
https://learn.microsoft.com/en-us/azure/backup/blob-backup-configure-manage

✗ **C** geo-backup in Azure Synapse Analytics
Synapse backup and restore concerns Synapse data warehouses; it does not provide a geo-backup for the separate Data Lake Storage Gen2 account. The serverless SQL and Spark pools in the scenario do not change the storage account's DR requirement.

https://learn.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/backup-and-restore

✓ **D** asynchronous replication
Asynchronous replication copies storage data to a secondary region, supporting recovery if the primary region becomes unavailable. It is the storage-level disaster recovery approach suited to storage1.

https://learn.microsoft.com/en-us/azure/storage/common/storage-disaster-recovery-guidance

---

Question 28

Fill in the blank · Domain: D4 Design infrastructure solutions › Design network solutions

Your company needs to configure a VPN gateway to accept connections from multiple on-premises sites simultaneously. What must you create on the VPN gateway to support multiple concurrent Site-to-Site VPN connections?

Answer: local network gateways

Explanation

To connect multiple on-premises sites to one Azure VPN gateway, define each site as a local network gateway, which represents that site's VPN device and address spaces. You can then create a separate Site-to-Site connection for each local network gateway, allowing the gateway to accept concurrent connections from the sites.

https://learn.microsoft.com/en-in/azure/vpn-gateway/add-remove-site-to-site-connections

✓ **Box 1: local network gateways** local network gateways
A local network gateway represents an on-premises site, including its VPN device and network address spaces. Creating one for each site lets you configure a separate Site-to-Site connection from the Azure VPN gateway to each site.

https://learn.microsoft.com/en-in/azure/vpn-gateway/add-remove-site-to-site-connections

---

Free demo · full exam: 393 questions

Question 29

Single choice · Domain: D4 Design infrastructure solutions › Design network solutions

Your organization uses Azure Virtual Machines for legacy applications that require static IP addresses. You need to ensure that if a VM is deallocated and restarted, it retains the same IP address. How should you configure this?

A. Assign a static private IP address to the VM's network interface and use a static public IP resource
B. Use dynamic IP assignment and configure a DNS CNAME record that updates automatically
C. Store the VM's IP address in Key Vault and update it manually after restart
D. Deploy the VM in an availability set and rely on Azure to maintain IP assignment

Answer: A. Assign a static private IP address to the VM's network interface and use a static public IP resource

Explanation

Configure the VM's network interface with a static private IP address and, if the VM needs a public address, associate a static public IP resource. Static assignment keeps the address reserved for the VM across deallocation and restart, meeting the requirement for a stable IP. A DNS record, manual tracking, or an availability set does not provide that address-retention guarantee.

https://learn.microsoft.com/en-us/azure/virtual-network/ip-services/virtual-network-static-public-ip
https://learn.microsoft.com/en-us/azure/virtual-network/ip-services/virtual-networks-static-private-ip

✓ **A** Assign a static private IP address to the VM's network interface and use a static public IP resource
A static private IP is configured on the network interface, and a static public IP resource retains its assigned public address. This directly provides stable addressing when the VM is deallocated and restarted.

https://learn.microsoft.com/en-us/azure/virtual-network/ip-services/virtual-network-static-public-ip
https://learn.microsoft.com/en-us/azure/virtual-network/ip-services/virtual-networks-static-private-ip

✗ **B** Use dynamic IP assignment and configure a DNS CNAME record that updates automatically
Dynamic assignment does not reserve the same IP address after deallocation, so the VM's address may change. A CNAME can refer to a DNS name, but it does not make the underlying IP address static.

https://learn.microsoft.com/en-us/azure/virtual-network/ip-services/virtual-networks-static-private-ip

✗ **C** Store the VM's IP address in Key Vault and update it manually after restart
Storing an address in Key Vault does not control Azure IP assignment. Manually updating a stored value after restart also does not ensure the VM retains the same address.

✗ **D** Deploy the VM in an availability set and rely on Azure to maintain IP assignment
An availability set helps distribute VMs across fault and update domains for availability; it does not configure a static IP address. The IP must be made static on the network interface or through a static public IP resource.

https://learn.microsoft.com/en-us/azure/virtual-machines/availability-set-overview
https://learn.microsoft.com/en-us/azure/virtual-network/ip-services/virtual-networks-static-private-ip

Free demo · full exam: 393 questions

Question 30

Single choice · Domain: D4 Design infrastructure solutions › Design migrations

You have an Azure subscription that contains an Azure HDInsight cluster named Cluster! and an Azure Data Lake Storage Gen2 account named account1 account1 uses read-access geo-redundant storage (RA-GRS) and blob snapshots. Cluster! stores 2 TB of data in account1.

The primary Azure region of account! experiences an outage.

You create a new Data Lake Storage Gen2 account named account2 in a different Azure region and register the account as the default storage for Cluster 1.

You need to populate account2 with the data from account1. The solution must meet the following requirements:

- Minimize how long it takes to restore access to the data.
- Minimize administrative effort.

What should you use?

A. a restore from a blob snapshot
B. AzCopy
C. Azure Site Recovery
D. Azure Data Factory

Answer: B. AzCopy

Explanation

Use AzCopy to copy the data from account1 to account2. Because account1 uses RA-GRS, its secondary endpoint can provide read access during a primary-region outage, and AzCopy can copy blobs between storage accounts. This is a direct way to transfer the 2 TB without setting up a separate data orchestration or disaster-recovery system, helping minimize restore time and administrative effort.

https://learn.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy?WT.mc_id=MVP_319025

✗ **A** a restore from a blob snapshot
Restoring a blob snapshot is intended to recover blob data in its storage account; it does not directly populate the new account2. It therefore does not meet the requirement to copy the data to the new default storage account. [6]

https://learn.microsoft.com/en-us/azure/backup/blob-restore

✓ **B** AzCopy
AzCopy supports copying blobs between storage accounts, making it a direct option for transferring account1's data into account2. Read access to account1's RA-GRS secondary endpoint can be used while its primary region is unavailable, minimizing delay without requiring a separate data-copy service to be configured. [1]

https://learn.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy?WT.mc_id=MVP_319025

✗ **C** Azure Site Recovery
Azure Site Recovery is intended for disaster recovery and migration of supported workloads, such as Azure virtual machines; it is not the appropriate tool for copying this storage account's blob data into account2. [2]

https://learn.microsoft.com/en-us/azure/site-recovery/azure-to-azure-enable-global-disaster-recovery

✗ **D** Azure Data Factory
Data Factory can copy data, but creating and managing a pipeline adds orchestration setup that is unnecessary for this direct storage-account copy. AzCopy is the simpler, lower-administration option for

the stated requirement. [1][12]

https://learn.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy?WT.mc_id=MVP_319025
https://learn.microsoft.com/en-us/azure/data-factory/quickstart-create-data-factory-copy-data-tool

---

Question 31

Single choice · Domain: D4 Design infrastructure solutions › Design an application architecture

You have an Azure subscription that contains an Azure DevOps pipeline and an Azure Repos repository. The pipeline performs the following actions:

- Builds a custom container image by using a base image stored in a public Docker Hub repository
- Stores the custom container image in Azure Container Registry
- Deploys the custom container image to Azure Container Apps

The pipeline is triggered by commits to the Azure Repos repository.

You need to ensure that when the base image in Docker Hub is updated, the custom image in Azure Container Registry is updated automatically. The solution must minimize administrative effort.

What should you include in the solution?

A. an Azure Automation runbook
B. ACR Tasks
C. an Azure function hosted in Container Apps
D. an Azure DevOps pipeline task

Answer: B. ACR Tasks

Explanation

ACR Tasks can monitor a base image in a public Docker Hub repository and automatically rebuild the dependent custom image in Azure Container Registry when the base image is updated. This handles the required update independently of commits to the Azure Repos repository, avoiding the need to maintain a separate polling or orchestration solution. [4][6]

https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tasks-base-images
https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tutorial-base-image-update

✗ **A** an Azure Automation runbook
A runbook could be programmed and scheduled to check for changes, but that would require you to create and maintain the checking and rebuild logic. ACR Tasks provides a purpose-built base-image update trigger with less administrative effort. [2][4]

https://learn.microsoft.com/en-us/azure/automation/shared-resources/schedules
https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tasks-base-images

✓ **B** ACR Tasks
ACR Tasks supports triggering a rebuild when an upstream base image changes, including a base image hosted in Docker Hub. The rebuilt custom image can then be pushed to Azure Container Registry without relying on a commit to the Azure Repos repository. [4][6]

https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tasks-base-images
https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tutorial-base-image-update

✗ **C** an Azure function hosted in Container Apps

A function could be written to detect changes and start a rebuild, but it adds custom code and operational components to manage. ACR Tasks has built-in support for monitoring base-image updates, so it better meets the requirement to minimize administrative effort. [4]

https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tasks-base-images

✗ **D** an Azure DevOps pipeline task
The pipeline is triggered by commits to Azure Repos, so a pipeline task alone would not run when the Docker Hub base image changes without a repository commit. ACR Tasks can trigger the rebuild based on the upstream base-image update instead. [4]

https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tasks-base-images

---

Question 32

Single choice · Domain: D4 Design infrastructure solutions › Design network solutions

You have an Azure subscription that contains a virtual network named VNet1. VNet1 contains three virtual machines that host a web app named Appl. The virtual machines are assigned only private IP addresses.

You need to recommend a solution to evenly distribute the traffic for App1 across the virtual machines. The solution must meet the following requirements:

- Ensure that App1 is accessible only from VNet1.
- Minimize administrative effort.

What should you include in the recommendation?

A. Azure Traffic Manager
B. Azure Application Gateway
C. Azure Firewall
D. Azure Front Door

Answer: B. Azure Application Gateway

Explanation

Azure Application Gateway can distribute application traffic across the three VM backends and can use a private frontend IP address. With that private frontend in VNet1, the app can be reached from within VNet1 without exposing a public endpoint. It therefore meets the access requirement while avoiding the extra components needed to provide private-only access through an internet-facing service.

https://learn.microsoft.com/en-us/azure/application-gateway/application-gateway-components
https://learn.microsoft.com/en-us/azure/application-gateway/configure-application-gateway-with-private-frontend-ip

✗ **A** Azure Traffic Manager
Traffic Manager provides DNS-based traffic routing rather than acting as a private, in-VNet load balancer for these VMs. It does not itself make App1 accessible only from VNet1.

https://learn.microsoft.com/en-us/azure/networking/design-guide/internet-ingress

✓ **B** Azure Application Gateway
Application Gateway can distribute requests among backend VMs and can be configured with a private frontend IP. Using that frontend in VNet1 allows access from within the VNet without a public frontend, meeting both requirements.

https://learn.microsoft.com/en-us/azure/application-gateway/application-gateway-components

Free demo · full exam: 393 questions

https://learn.microsoft.com/en-us/azure/application-gateway/configure-application-gateway-with-private-frontend-ip

✗ **C** Azure Firewall
Azure Firewall is a network security service, not the application load-balancing solution needed to distribute App1 traffic evenly across the VMs. It would not meet the core traffic-distribution requirement by itself.

https://learn.microsoft.com/en-us/azure/firewall/firewall-known-issues

✗ **D** Azure Front Door
Azure Front Door is an internet-facing global entry point, so it does not by itself satisfy the requirement that App1 be accessible only from VNet1. Private connectivity to an origin does not make the Front Door endpoint private to VNet1.

https://learn.microsoft.com/en-us/azure/frontdoor/front-door-faq
https://learn.microsoft.com/en-us/azure/frontdoor/origin-security?pivots=front-door-standard-premium&tabs=app-service-functions
https://learn.microsoft.com/azure/frontdoor/scenarios?WT.mc_id=modinfra-wfh-pierre&trk=article-ssr-frontend-pulse_little-text-block

---

Question 33

Single choice · Domain: D4 Design infrastructure solutions › Design an application architecture

You have an Azure subscription. The subscription contains three Azure functions and an Azure Kubernetes Service (AKS) cluster that hosts three containers. You have an app named App1 that is hosted on the functions and the containers. App1 supports feature flags for A/B testing.

You need to implement A/8 testing for App1. The solution must ensure that multiple scenarios can be tested by using the existing deployment of App1. What should you use to manage the flags?

A. Azure DevOps feature flags
B. Azure Key Vault secrets
C. Azure App Service app settings
D. Azure App Configuration feature flags

Answer: D. Azure App Configuration feature flags

Explanation

Use Azure App Configuration feature flags to manage App1's flags centrally across the Azure Functions and AKS-hosted containers. App Configuration supports feature management, including configuring flags so that different scenarios or variants can be tested without deploying a different version of App1. This meets the requirement to test with the existing deployment. [2][5]

https://learn.microsoft.com/en-us/azure/azure-app-configuration/concept-feature-management
https://learn.microsoft.com/en-us/azure/azure-app-configuration/howto-variant-feature-flags

✗ **A** Azure DevOps feature flags
Azure DevOps is not the suitable centralized runtime feature-flag service for this app. App1 runs across Azure Functions and AKS containers, so its flags need to be managed in a configuration service those workloads can use while the existing deployment remains in place.

✗ **B** Azure Key Vault secrets
Key Vault is designed to store and protect secrets such as credentials and keys, not to manage feature flags and A/B test variants. It does not provide the feature-management capabilities needed to test

multiple scenarios in the existing deployment.

✗ **C** Azure App Service app settings
App Service app settings configure applications hosted in Azure App Service; they are not a shared feature-flag management solution for this mix of Azure Functions and AKS containers. They also do not provide the feature-management capabilities needed for A/B testing scenarios.

✓ **D** Azure App Configuration feature flags
Azure App Configuration provides feature-flag management that application code can use across the Functions and AKS-hosted containers. Its variant feature flags support assigning different values or experiences for testing, allowing multiple scenarios to be evaluated without deploying separate versions of App1.

https://learn.microsoft.com/en-us/azure/azure-app-configuration/concept-feature-management
https://learn.microsoft.com/en-us/azure/azure-app-configuration/howto-variant-feature-flags

Dumpsgate

Question 34

Hotspot (dropdowns) · Domain: D4 Design infrastructure solutions › Design network solutions

Create a new Azure load balancer named LB1.

Create two backend pools named Pool1 and Pool2.

To LB1, add the inbound NAT rules shown in the following table.

Name	Protocol	Port	Backend port	Floating IP	Idle timeout (minutes)
Rule1	TCP	22	22	Disabled	4
Rule2	TCP	443	443	Disabled	4
Rule3	TCP	3389	3389	Disabled	4
Rule4	TCP	5986	5986	Disabled	4
Rule5	TCP	8080	80	Disabled	4

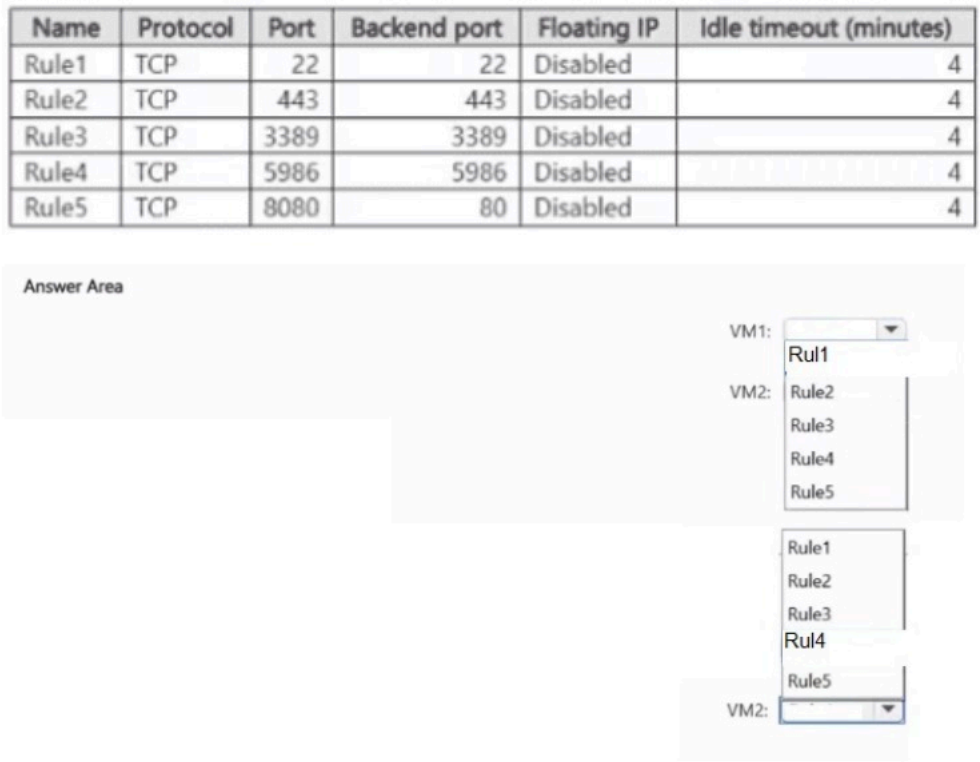

You need to configure LB1 to meet the following requirements:

Enable remote desktop access to VM1 from the internet by using standard TCP ports.

Enable remote Windows PowerShell access to VM2 from the internet by using standard TCP ports.

Which rule should you assign to each virtual machine? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

VM1: Rule1 | Rule2 | Rule3 | Rule4 | Rule5

VM2: Rule1 | Rule2 | Rule3 | Rule4 | Rule5

Answer

- VM1: **Rule3**
- VM2: **Rule4**

Explanation

Assign Rule3 to VM1 because it forwards TCP port 3389 to backend port 3389, the standard RDP port. Assign Rule4 to VM2 because it forwards TCP port 5986 to backend port 5986, the standard HTTPS WinRM port for remote PowerShell access. The inbound NAT rules let the load balancer direct internet traffic on the selected frontend port to the corresponding VM backend port.

https://learn.microsoft.com/azure/load-balancer/inbound-nat-rules
https://learn.microsoft.com/en-us/powershell/scripting/security/remoting/winrm-security?view=powershell-7.6

✗ **Box 1: Rule1** VM1 — Rule1 (TCP 22 to backend port 22)
Rule1 forwards TCP port 22, which is not the standard port for Remote Desktop Protocol. VM1 requires the RDP rule, which uses TCP port 3389.

https://learn.microsoft.com/azure/load-balancer/inbound-nat-rules

✗ **Box 1: Rule2** VM1 — Rule2 (TCP 443 to backend port 443)
Rule2 forwards TCP port 443, not the standard RDP port. It would not meet the requirement to enable standard-port remote desktop access to VM1.

https://learn.microsoft.com/azure/load-balancer/inbound-nat-rules

✓ **Box 1: Rule3** VM1 — Rule3 (TCP 3389 to backend port 3389)
Rule3 forwards TCP port 3389 to backend port 3389. TCP 3389 is the standard port for Remote Desktop Protocol, so this rule meets the VM1 requirement.

https://learn.microsoft.com/azure/load-balancer/inbound-nat-rules

✗ **Box 1: Rule4** VM1 — Rule4 (TCP 5986 to backend port 5986)
Rule4 uses TCP port 5986, the standard HTTPS port for Windows Remote Management, rather than the standard RDP port. It is therefore appropriate for the PowerShell requirement, not VM1's remote desktop requirement.

https://learn.microsoft.com/en-us/powershell/scripting/security/remoting/winrm-security?view=powershell-7.6

✗ **Box 1: Rule5** VM1 — Rule5 (TCP 8080 to backend port 80)
Rule5 maps frontend TCP port 8080 to backend port 80. Neither port is the standard RDP port required for VM1.

https://learn.microsoft.com/azure/load-balancer/inbound-nat-rules

✗ **Box 2: Rule1** VM2 — Rule1 (TCP 22 to backend port 22)
Rule1 forwards TCP port 22, not the standard TCP port for remote Windows PowerShell access. VM2 needs the rule for HTTPS WinRM on TCP port 5986.

https://learn.microsoft.com/en-us/powershell/scripting/security/remoting/winrm-security?view=powershell-7.6

✗ **Box 2: Rule2** VM2 — Rule2 (TCP 443 to backend port 443)
Rule2 forwards TCP port 443, which is not the standard HTTPS WinRM port used here for remote PowerShell access. Rule4 forwards the required TCP port 5986 to the VM.

https://learn.microsoft.com/en-us/powershell/scripting/security/remoting/winrm-security?view=powershell-7.6

✗ **Box 2: Rule3** VM2 — Rule3 (TCP 3389 to backend port 3389)
Rule3 uses TCP port 3389, the standard RDP port, rather than the standard HTTPS WinRM port. It fits the VM1 remote desktop requirement, not VM2's remote PowerShell requirement.

https://learn.microsoft.com/en-us/powershell/scripting/security/remoting/winrm-security?view=powershell-7.6

✓ **Box 2: Rule4** VM2 — Rule4 (TCP 5986 to backend port 5986)
Rule4 forwards TCP port 5986 to backend port 5986. TCP 5986 is the standard HTTPS port for Windows Remote Management, so it meets VM2's remote PowerShell requirement.

https://learn.microsoft.com/en-us/powershell/scripting/security/remoting/winrm-security?view=powershell-7.6

Free demo · full exam: 393 questions

✗ **Box 2: Rule5** VM2 — Rule5 (TCP 8080 to backend port 80)
Rule5 maps frontend TCP port 8080 to backend port 80, rather than forwarding the standard HTTPS WinRM port 5986. It does not meet VM2's requirement for standard-port remote PowerShell access.
https://learn.microsoft.com/azure/load-balancer/inbound-nat-rules
https://learn.microsoft.com/en-us/powershell/scripting/security/remoting/winrm-security?view=powershell-7.6

Question 35

Single choice · Domain: D4 Design infrastructure solutions › Design migrations

You have an on-premises server named Server1 that runs Windows Server and contains a Microsoft SQL Server database named DB1.

You have an Azure subscription.

You plan to migrate D81 to Azure.

You use the Azure Database Migration Service to assess DB1 and receive the assessment results shown in the following table.

Migration target	Number of issues that must be remediated before the migration
SQL Server on Azure Virtual Machines	3
Azure SQL Database	51
Azure SQL Managed Instance	27

You need to recommend a migration target for DB1. The migration target must meet the following requirements:

- Eliminate the need for post-migration operating system maintenance.
- Minimize administrative effort to prepare for the migration.

What should you recommend?

A. Azure SQL Oatabase
B. Azure Cosmos DB for NoSQL
C. SQL Server on Azure Virtual Machines
D. Azure Cosmos DB for PostgreSQL
E. Azure SQL Managed Instance

Answer: E. Azure SQL Managed Instance

Explanation
Azure SQL Managed Instance meets both requirements: it is a managed database service, so you do not maintain the guest operating system after migration, and its broad SQL Server compatibility makes it a lower-effort target than Azure SQL Database for this assessment. The assessment reports 27 issues for Managed Instance, compared with 51 for Azure SQL Database, indicating fewer issues to remediate before migrating DB1. References [4] and [9] support the service comparison and SQL Server-to-Managed

Instance migration context.

https://learn.microsoft.com/en-us/azure/azure-sql/database/features-comparison?view=azuresql
https://learn.microsoft.com/azure/azure-sql/migration-guides/managed-instance/sql-server-to-managed-instance-overview?view=azuresql

✗ **A** Azure SQL Database
Azure SQL Database removes the need to maintain an operating system, but the assessment reports 51 issues that must be remediated before migration. That is more preparation work than the 27 issues reported for Azure SQL Managed Instance, so it does not best meet the requirement to minimize administrative effort.

https://learn.microsoft.com/en-us/azure/azure-sql/database/features-comparison?view=azuresql

✗ **B** Azure Cosmos DB for NoSQL
Cosmos DB for NoSQL is not a like-for-like target for this Microsoft SQL Server relational database; moving DB1 to it would involve changing database technology rather than directly migrating SQL Server. Although it is a managed service, it does not minimize the effort to prepare this SQL Server migration.

✗ **C** SQL Server on Azure Virtual Machines
This option requires managing the virtual machine's operating system after migration, so it fails the requirement to eliminate post-migration operating system maintenance. The assessment also reports only three issues for this target, but that does not overcome the OS-maintenance requirement.

https://learn.microsoft.com/en-us/azure/azure-sql/migration-guides/virtual-machines/sql-server-to-sql-on-azure-vm-migration-overview?view=azuresql

✗ **D** Azure Cosmos DB for PostgreSQL
Cosmos DB for PostgreSQL is not a like-for-like destination for DB1, which is a Microsoft SQL Server database. Changing database platforms would require more than preparing a SQL Server migration, and the assessment provides no results for this target.

✓ **E** Azure SQL Managed Instance
Managed Instance is a managed service, eliminating the need to maintain the operating system after migration. The assessment reports 27 issues to remediate, fewer than the 51 for Azure SQL Database, and Managed Instance provides a SQL Server-compatible destination that supports a lower-effort migration than a more substantial database-platform change.

https://learn.microsoft.com/en-us/azure/azure-sql/database/features-comparison?view=azuresql
https://learn.microsoft.com/azure/azure-sql/migration-guides/managed-instance/sql-server-to-managed-instance-overview?view=azuresql

Question 36

Single choice · Domain: D4 Design infrastructure solutions › Design compute solutions

You plan to develop, deploy, and maintain 100 React-based web apps that will NOT require server-side rendering. You need to recommend a hosting solution for the apps. The solution must meet the following requirements:

- Automatically deploy new versions of the apps when changes to the apps' source code are committed in

Azure Repos.

- Automatically stage a preview version of an app when a pull request is processed.
- Support Azure API Management by using the /api.
- Minimize costs.

What should you include in the recommendation?

A. an Azure function
B. Azure App Service web apps
C. Azure Static Web Apps
D. Azure Blob Storage static sites

Answer: C. Azure Static Web Apps

Explanation

Azure Static Web Apps is designed to host static web applications such as React apps that do not need server-side rendering. It can deploy from Azure Repos, create preview environments for pull requests, and integrate Azure API Management through the /api route. It is also a cost-effective fit for hosting many static apps compared with general-purpose app hosting.

https://learn.microsoft.com/en-us/azure/static-web-apps/preview-environments
https://learn.microsoft.com/en-us/azure/static-web-apps/apis-api-management
https://learn.microsoft.com/azure/static-web-apps/

✗ **A** an Azure function
Azure Functions can host APIs or serverless code, but they are not the best fit as the primary hosting solution for 100 React static web apps. This option does not directly address the requested source-code deployments and pull-request preview environments for the apps.

https://learn.microsoft.com/en-us/azure/static-web-apps/preview-environments
https://learn.microsoft.com/en-us/azure/static-web-apps/apis-overview

✗ **B** Azure App Service web apps
App Service can host web applications, but it is a general-purpose web-hosting service and is not the most direct, cost-conscious solution for apps that do not require server-side rendering. Static Web Apps specifically supports pull-request preview environments and API Management integration for this scenario.

https://learn.microsoft.com/en-us/azure/static-web-apps/preview-environments
https://learn.microsoft.com/en-us/azure/static-web-apps/apis-api-management
https://learn.microsoft.com/en-us/azure/static-web-apps/apis-app-service

✓ **C** Azure Static Web Apps
Azure Static Web Apps is suited to static React apps and supports deployment workflows connected to Azure Repos, as well as staging preview environments for pull requests. It also supports connecting Azure API Management through the /api route, meeting the integration requirement while avoiding unnecessary server-side hosting.

Free demo · full exam: 393 questions

https://learn.microsoft.com/en-us/azure/static-web-apps/preview-environments
https://learn.microsoft.com/en-us/azure/static-web-apps/apis-api-management
https://learn.microsoft.com/azure/static-web-apps/

✗ **D** Azure Blob Storage static sites
Blob Storage static website hosting can serve static files, but it does not provide the requested built-in pull-request preview environments or the Azure API Management integration through /api. As a result, it does not meet all the deployment and API requirements for this recommendation.

https://learn.microsoft.com/en-us/azure/static-web-apps/apis-api-management
https://learn.microsoft.com/azure/storage/blobs/storage-blob-static-website
https://learn.microsoft.com/en-us/azure/storage/blobs/static-website-hosting-faq

---

Question 37

Single choice · Domain: D4 Design infrastructure solutions › Design migrations

You have a network-attached storage device that hosts a file share. The share contains 1 TB of JSON files.

You have an Azure subscription that contains a storage account named storage1, storage1 has hierarchical namespace enabled. You need to recommend a solution to migrate the files to storage1. What should you include in the recommendation?

A. AzCopy
B. Robocopy
C. Azure Storage Mover
D. Azure File Sync

Answer: C. Azure Storage Mover

Explanation

Azure Storage Mover is the suitable recommendation for migrating the 1 TB of files from the network-attached file share to the Azure storage account. The destination has hierarchical namespace enabled, so it is an Azure Data Lake Storage Gen2 account, and Storage Mover supports migrations from network file shares to Azure storage destinations. This addresses the stated source and destination without proposing a tool intended for a different migration scenario.

https://learn.microsoft.com/en-us/azure/storage-mover/migration-basics
https://learn.microsoft.com/en-us/azure/storage-mover/job-definition-create?tabs=portal

✗ **A** AzCopy
AzCopy can transfer data to Azure Storage, but this scenario calls for migrating a large file share to a hierarchical-namespace-enabled account. Storage Mover is the more appropriate managed migration solution for this source and destination.

https://learn.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-migrate-data
https://learn.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-files
https://learn.microsoft.com/en-us/azure/storage-mover/migration-basics

✗ **B** Robocopy
Robocopy is commonly used to copy files between file shares, including in Azure Files migration scenarios. It is not the recommended migration service for moving the network share's files into this hierarchical-namespace-enabled Azure storage account.

https://learn.microsoft.com/en-us/azure/storage/files/storage-files-migration-robocopy

Free demo · full exam: 393 questions

✓ **C** Azure Storage Mover
Azure Storage Mover is designed to migrate data from on-premises or network file shares to Azure Storage. Its migration workflow fits the stated file-share source and the Azure Data Lake Storage Gen2 destination represented by the enabled hierarchical namespace.

https://learn.microsoft.com/en-us/azure/storage-mover/resource-hierarchy
https://learn.microsoft.com/en-us/azure/storage-mover/job-definition-create?tabs=portal
https://learn.microsoft.com/en-us/azure/storage-mover/migration-basics

✗ **D** Azure File Sync
Azure File Sync is for synchronizing and tiering files with Azure file shares, and its documented migration use case concerns Azure file-share-to-file-share moves. The target here is a hierarchical-namespace-enabled storage account, not an Azure file share.

https://learn.microsoft.com/en-us/azure/storage/file-sync/file-sync-share-to-share-migration

---

Question 38

Multiple choice · Domain: D4 Design infrastructure solutions › Design migrations

You have an on-premises application that consumes data from multiple databases. The application code references database tables by using a combination of the server, database, and table name.

You need to migrate the application data to Azure.

To which two services can you migrate the application data to achieve the goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Azure SQL Managed Instance
B. Azure SQL Database
C. SQL Server Stretch Database
D. SQL Server on an Azure virtual machine

Dumpsgate

Answer

- A. **Azure SQL Managed Instance**
- D. **SQL Server on an Azure virtual machine**

Explanation

The application uses server-, database-, and table-qualified names, so its target must support the cross-database behavior the application relies on. Azure SQL Managed Instance provides SQL Server instance-level compatibility, including cross-database access, and SQL Server on an Azure virtual machine retains the SQL Server database engine behavior. Azure SQL Database does not provide the same general instance-level cross-database references, and Stretch Database is not a suitable destination for migrating the application's databases.

https://learn.microsoft.com/data-migration/sql-server/virtual-machines/overview
https://learn.microsoft.com/en-us/sql/database-engine/?view=sql-server-ver16
https://learn.microsoft.com/en-us/azure/sql-database/sql-database-transact-sql-information
https://learn.microsoft.com/en-us/sql/database-engine/deprecated-database-engine-features-in-sql-server-2022?view=sql-server-ver17

✓ **A** Azure SQL Managed Instance
Managed Instance supports cross-database queries within the instance, which fits an application that

Free demo · full exam: 393 questions

names databases as part of its table references. It is a managed Azure destination that retains relevant SQL Server instance-level capabilities.

https://learn.microsoft.com/en-us/sql/database-engine/?view=sql-server-ver16
https://learn.microsoft.com/en-us/answers/questions/5625549/is-it-possible-to-set-up-sql-server-cross-database

✗ **B** Azure SQL Database
Azure SQL Database does not provide the same general cross-database references available in a SQL Server instance or Managed Instance. Its cross-database options have limitations, so it does not meet the requirement for the application's existing database-qualified references as stated.

https://learn.microsoft.com/en-us/azure/sql-database/sql-database-transact-sql-information

✗ **C** SQL Server Stretch Database
Stretch Database is a feature for extending selected SQL Server tables to Azure, not a destination service for migrating multiple application databases while preserving their cross-database references. It is also deprecated in SQL Server.

https://learn.microsoft.com/en-us/sql/database-engine/deprecated-database-engine-features-in-sql-server-2022?view=sql-server-ver17

✓ **D** SQL Server on an Azure virtual machine
SQL Server on an Azure VM runs the SQL Server database engine, so it supports database-qualified table references across databases hosted by that instance. This makes it a suitable migration target when preserving the application's existing SQL Server behavior is required.

https://learn.microsoft.com/data-migration/sql-server/virtual-machines/overview
https://learn.microsoft.com/en-us/sql/database-engine/?view=sql-server-ver16

Question 39

Single choice · Domain: D4 Design infrastructure solutions › Design network solutions

You have an on-premises datacenter named Site1 that contains multiple servers.

You have an Azure subscription that contains a virtual network named VNet1. VNet1 contains multiple virtual machines. VNet1 uses the default Azure DNS service. Site1 connects to VNet1 by using a Site-to-Site (S2S) VPN.

You need to recommend a DNS solution that will forward DNS queries from the virtual machines that target the on-premises DNS namespace to the on-premises DNS servers. The solution must minimize administrative effort and costs.

What should you include in the recommendation?

A. Private DNS zone
B. Azure Firewall DNS Proxy
C. an Azure virtual machine that runs the DNS service
D. Private DNS resolver

Answer: D. Private DNS resolver

Explanation
Azure DNS Private Resolver can forward queries for the on-premises DNS namespace from VNet1 to the on-premises DNS servers by using an outbound endpoint and forwarding ruleset. The S2S VPN provides connectivity to those servers, while the managed resolver avoids deploying and maintaining a DNS virtual

machine. This meets the forwarding requirement with less administrative effort than a self-managed solution. [2][4]

https://learn.microsoft.com/en-us/azure/dns/dns-private-resolver-overview/
https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/ready/azure-best-practices/dns-for-on-premises-and-azure-resources

✗ **A** Private DNS zone
A private DNS zone stores DNS records for names in the zone; it does not forward queries for an on-premises DNS namespace to on-premises DNS servers. It therefore does not provide the required query forwarding. [1][8]

https://learn.microsoft.com/en-us/azure/dns/private-dns-overview
https://learn.microsoft.com/en-us/azure/dns/private-dns-scenarios

✗ **B** Azure Firewall DNS Proxy
Azure Firewall DNS Proxy lets clients send DNS queries through the firewall, which can then forward them to configured DNS servers. However, it requires Azure Firewall and configuring VMs to use the proxy, adding infrastructure and cost when the requirement is simply to forward queries through the existing S2S connection. [9]

https://learn.microsoft.com/azure/firewall/dns-settings

✗ **C** an Azure virtual machine that runs the DNS service
A DNS server VM could forward queries to the on-premises DNS servers, but it must be deployed, configured, secured, and maintained. That creates more administrative effort than using the managed Azure DNS Private Resolver for this forwarding scenario. [2][4]

https://learn.microsoft.com/en-us/azure/dns/dns-private-resolver-overview/
https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/ready/azure-best-practices/dns-for-on-premises-and-azure-resources

✓ **D** Private DNS resolver
Azure DNS Private Resolver supports forwarding DNS queries from Azure to on-premises DNS servers by using an outbound endpoint and forwarding ruleset. With the S2S VPN connecting VNet1 to Site1, it provides the required path without requiring a DNS server VM to manage. [2][4]

https://learn.microsoft.com/en-us/azure/dns/dns-private-resolver-overview/
https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/ready/azure-best-practices/dns-for-on-premises-and-azure-resources

Dumpsgate

Question 40

Hotspot (dropdowns) · Domain: D4 Design infrastructure solutions › Design an application architecture

You are building a mobile app for users to upload data files to an Azure Storage account and set properties for data processing. You need to recommend a messaging solution for the app The solution must meet the following requirements:

- Use a messaging resource that supports 128KB messages and guarantees first in, first out (FIFO) delivery.
- Use an event resource that automatically triggers a webhook when a file is added to the storage account container.

What should you include in the recommendation? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

For the messaging resource: Azure Relay | Azure Service bus queues | Azure Queue Storage

For the event resource: Azure Event Grid | Azure Event Hubs | Azure IoT Hub

Answer

- For the messaging resource: **Azure Service bus queues**
- For the event resource: **Azure Event Grid**

Explanation

Use Azure Service Bus queues for the messaging resource: the Service Bus option supports messages of at least 128 KB and can provide FIFO processing through sessions. Use Azure Event Grid for the event resource because it can receive blob-storage events, such as a file being added to a container, and deliver them to a webhook endpoint. Together, these choices satisfy the messaging and automatic event-trigger requirements.

https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-queues-topics-subscriptions
https://learn.microsoft.com/en-us/azure/service-bus-messaging/message-sessions
https://learn.microsoft.com/en-us/azure/event-grid/event-schema-blob-storage
https://learn.microsoft.com/en-us/azure/event-grid/blob-event-quickstart-portal
https://learn.microsoft.com/en-us/azure/event-grid/namespace-handler-webhook

✗ **Box 1: Azure Relay** Azure Relay
Azure Relay is for connecting applications across network boundaries, not a queue providing the requested FIFO message delivery. It therefore does not meet the messaging requirement in this scenario.

✓ **Box 1: Azure Service bus queues** Azure Service bus queues
Azure Service Bus queues support the required message size, and sessions can be used to process

related messages in FIFO order. This makes the Service Bus option the match for the app's messaging resource.

https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-queues-topics-subscriptions
https://learn.microsoft.com/en-us/azure/service-bus-messaging/message-sessions

✗ **Box 1: Azure Queue Storage** Azure Queue Storage
Azure Queue Storage does not provide guaranteed FIFO delivery; messages are processed with best-effort ordering. It therefore does not satisfy the requirement for guaranteed first-in, first-out delivery.

https://learn.microsoft.com/en-us/rest/api/storageservices/queue-service-rest-api
https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted

✓ **Box 2: Azure Event Grid** Azure Event Grid
Event Grid supports events from Blob Storage, including when a blob is created, and can route events to a webhook handler. It meets the requirement to trigger a webhook when a file is added to the storage container.

https://learn.microsoft.com/en-us/azure/event-grid/event-schema-blob-storage
https://learn.microsoft.com/en-us/azure/event-grid/blob-event-quickstart-portal
https://learn.microsoft.com/en-us/azure/event-grid/namespace-handler-webhook

✗ **Box 2: Azure Event Hubs** Azure Event Hubs
Event Hubs is designed for ingesting and streaming event data, rather than routing a Blob Storage file-added event to a webhook as requested here. Event Grid is the appropriate resource for this storage-event-to-webhook scenario.

✗ **Box 2: Azure IoT Hub** Azure IoT Hub
IoT Hub is intended for communication with IoT devices and is not the resource needed to detect a file added to a storage container and trigger a webhook. Event Grid provides that storage-event routing capability.

https://learn.microsoft.com/en-us/azure/iot-hub/iot-hub-event-grid

Dumpsgate