# Exam Questions AZ-204

Developing Solutions for Microsoft Azure

## https://www.2passeasy.com/dumps/AZ-204/

**NEW QUESTION 1**
- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data. You need to ensure the app does not time out and processes the blob data.
Solution: Update the functionTimeout property of the host.json project file to 10 minutes. Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**Explanation:**
Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.
Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:
Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be
processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices

**NEW QUESTION 2**
- (Topic 8)
You develop and add several functions to an Azure Function app that uses the latest runtime host. The functions contain several REST API endpoints secured by using SSL. The Azure Function app runs in a Consumption plan.
You must send an alert when any of the function endpoints are unavailable or responding too slowly.
You need to monitor the availability and responsiveness of the functions. What should you do?

A. Create a URL ping test.
B. Create a timer triggered function that calls TrackAvailability() and send the results to ApplicationInsights.
C. Create a timer triggered function that calls GetMetric("Request Size") and send the results toApplication Insights.
D. Add a new diagnostic setting to the Azure Function ap
E. Enable the FunctionAppLogs and Send to Log Analytics options.

**Answer:** B

**Explanation:**
You can create an Azure Function with TrackAvailability() that will run periodically
according to the configuration given in TimerTrigger function with your own business logic. The results of this test will be sent to your Application Insights resource, where you will be able to query for and alert on the availability results data. This allows you to create customized tests similar to what you can do via Availability Monitoring in the portal. Customized tests will allow you to write more complex availability tests than is possible using the portal UI, monitor an app inside of your Azure VNET, change the endpoint address, or create an availability test even if this feature is not available in your region.
D18912E1457D5D1DDCBD40AB3BF70D5D
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-azure-functions

**NEW QUESTION 3**
- (Topic 8)
An organization hosts web apps in Azure. The organization uses Azure Monitor You discover that configuration changes were made to some of the web apps. You need to identify the configuration changes. Which Azure Monitor log should you review?

A. AppServiceEnvironmentPlatformLogs
B. AppServiceApplogs
C. AppServiceAuditLogs
D. AppServiceConsoteLogs

**Answer:** C

**NEW QUESTION 4**
DRAG DROP - (Topic 8)
Contoso, Ltd. provides an API to customers by using Azure API Management (APIM). The API authorizes users with a JWT token.
You must implement response caching for the APIM gateway. The caching mechanism must detect the user ID of the client that accesses data for a given location and cache the response for that user ID.
You need to add the following policies to the policies file:
• a set-variable policy to store the detected user identity
• a cache-lookup-value policy
• a cache-store-value policy
• a find-and-replace policy to update the response body with the user profile information To which policy section should you add the policies? To answer, drag the appropriate
sections to the correct policies. Each section 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

| Policy section |
|---|
| Inbound |
| Outbound |

Answer Area

| Policy | Policy section |
|---|---|
| Set-variable | policy section |
| Cache-lookup-value | policy section |
| Cache-store-value | policy section |
| Find-and-replace | policy section |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Inbound.
A set-variable policy to store the detected user identity. Example:
<policies>
<inbound>
<!-- How you determine user identity is application dependent -->
<set-variable name="enduserid"
value="@(context.Request.Headers.GetValueOrDefault("Authorization","").Split(' ')[1].AsJwt()?.Subject)" />
Box 2: Inbound
A cache-lookup-value policy Example:
<inbound>
<base />
<cache-lookup vary-by-developer="true | false" vary-by-developer-groups="true | false" downstream-caching-type="none | private | public" must-revalidate="true | false">
<vary-by-query-parameter>parameter name</vary-by-query-parameter> <!-- optional, can repeated several times -->
</cache-lookup>
</inbound>
Box 3: Outbound
A cache-store-value policy. Example:
<outbound>
<base />
<cache-store duration="3600" />
</outbound>
Box 4: Outbound
A find-and-replace policy to update the response body with the user profile information. Example:
<outbound>
<!-- Update response body with user profile-->
<find-and-replace from="'$userprofile$'"
to="@((string)context.Variables["userprofile"])" />
<base />
</outbound>

**NEW QUESTION 5**
- (Topic 8)
An organization deploys Azure Cosmos DB.
You need to ensure that the index is updated as items are created, updated, or deleted. What should you do?

A. Set the value of the EnableScanInQuery option to True.
B. Set the indexing mode to Consistent.
C. Set the indexing mode to Lazy.
D. Set the value of the automatic property of the indexing policy to False.

**Answer:** B

**NEW QUESTION 6**
- (Topic 8)
You are developing an ASP.NET Core website that uses Azure FrontDoor. The website is used to build custom weather data sets for researchers. Data sets are downloaded by users as Comma Separated Value (CSV) files. The data is refreshed every 10 hours.
Specific files must be purged from the FrontDoor cache based upon Response Header values.
You need to purge individual assets from the Front Door cache. Which type of cache purge should you use?

A. single path
B. wildcard
C. root domain

**Answer:** A

**Explanation:**
These formats are supported in the lists of paths to purge:
? Single path purge: Purge individual assets by specifying the full path of the asset (without the protocol and domain), with the file extension, for example,
/pictures/strasbourg.png;
? Wildcard purge: Asterisk (*) may be used as a wildcard. Purge all folders, subfolders, and files under an endpoint with /* in the path or purge all subfolders and files under a specific folder by specifying the folder followed by /*, for example,
/pictures/*.
? Root domain purge: Purge the root of the endpoint with "/" in the path.
Reference:

https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching

**NEW QUESTION 7**
DRAG DROP - (Topic 8)
You are developing Azure WebJobs.
You need to recommend a WebJob type for each scenario.
Which WebJob type should you recommend? To answer, drag the appropriate WebJob types to the correct scenarios. Each WebJob type 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.

| WebJob types | Scenario | WebJob type |
| --- | --- | --- |
| Triggered | Run on all instances that the web app runs on. Optionally restrict the WebJob to a single instance. | |
| Continuous | Run on a single instance that Azure select for load balancing. | |
| | Supports remote debugging | |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Continuous
Continuous runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.
Box 2: Triggered
Triggered runs on a single instance that Azure selects for load balancing.
Box 3: Continuous
Continuous supports remote debugging.
Note:
The following table describes the differences between continuous and triggered WebJobs.

| Continuous | Triggered |
| --- | --- |
| Starts immediately when the WebJob is created. To keep the job from ending, the program or script typically does its work inside an endless loop. If the job does end, you can restart it. | Starts only when triggered manually or on a schedule. |
| Runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance. | Runs on a single instance that Azure selects for load balancing. |
| Supports remote debugging. | Doesn't support remote debugging. |

References:
https://docs.microsoft.com/en-us/azure/app-service/web-sites-create-web-jobs

**NEW QUESTION 8**
DRAG DROP - (Topic 8)
You manage several existing Logic Apps.
You need to change definitions, add new logic, and optimize these apps on a regular basis. What should you use? To answer, drag the appropriate tools to the correct functionalities.
Each tool 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.

| Tools | Functionality | Tool |
| --- | --- | --- |
| Logic Apps Designer | Edit B2B workflows | |
| Code View Editor | Edit definitions in JSON | |
| Enterprise Integration Pack | Visually add functionality | |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Enterprise Integration Pack
After you create an integration account that has partners and agreements, you are ready to create a business to business (B2B) workflow for your logic app with

the Enterprise Integration Pack.
Box 2: Code View Editor
To work with logic app definitions in JSON, open the Code View editor when working in the Azure portal or in Visual Studio, or copy the definition into any editor that you want.
Box 3: Logical Apps Designer
You can build your logic apps visually with the Logic Apps Designer, which is available in the Azure portal through your browser and in Visual Studio.
References:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-enterprise-integration-b2b https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-overview

**NEW QUESTION 9**
HOTSPOT - (Topic 8)
You develop new functionality in a web application for a company that provides access to seismic data from around the world. The seismic data is stored in Redis Streams within an Azure Cache for Redis instance.
The new functionality includes a real-time display of seismic events as they occur. You need to implement the Azure Cache for Redis command to receive seismic data.
How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.



A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
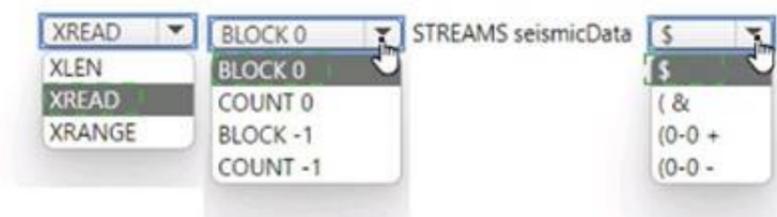


**NEW QUESTION 10**
- (Topic 8)
Your company is designing an application named App1 that will use data from Azure SQL Database. App1 will be accessed over the internet by many users.
You need to recommend a solution for improving the performance ofApp1. What should you include in the recommendation?

A. Azure HPC cache
B. ExpressRoute
C. a CON profile
D. Azure Cache for Redis

**Answer:** D

**NEW QUESTION 10**
DRAG DROP - (Topic 8)
You have an application that provides weather forecasting data to external partners. You use Azure API Management to publish APIs.
You must change the behavior of the API to meet the following requirements:
• Support alternative input parameters.
• Remove formatting text from responses.
• Provide additional context to back-end services.
Which types of policies should you implement? To answer, drag the policy types to the correct scenarios. Each policy type 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.



A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

| Policy types | | Answer Area | | Requirement | Policy type |
|---|---|---|---|---|---|

Policy types: Inbound, Outbound, Backend

Answer Area:

| Requirement | Policy type |
|---|---|
| Support alternative input parameters. | Inbound |
| Remove formatting text from responses. | Outbound |
| Provide additional context to back-end services. | Inbound |
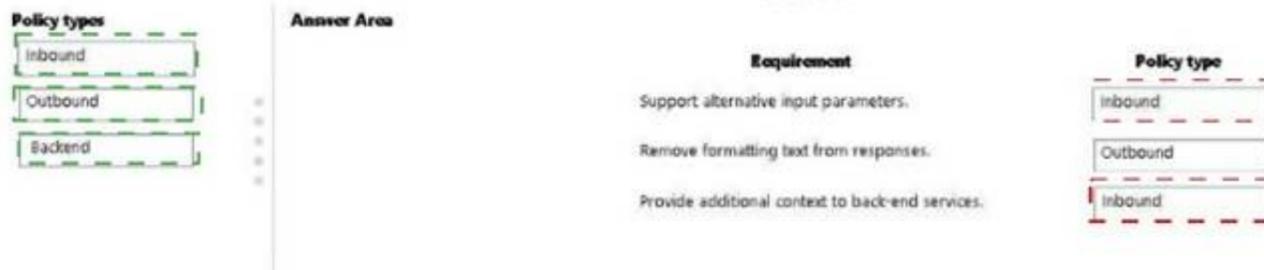
**NEW QUESTION 14**
- (Topic 8)
You are developing an application that uses Azure Blob storage.
The application must read the transaction logs of all the changes that occur to the blobs and the blob metadata in the storage account for auditing purposes. The changes must be in the order in which they occurred, include only create, update, delete, and copy operations and be retained for compliance reasons.
You need to process the transaction logs asynchronously. What should you do?

A. Process all Azure Blob storage events by using Azure Event Grid with a subscriber Azure Function app.
B. Enable the change feed on the storage account and process all changes for available events.
C. Process all Azure Storage Analytics logs for successful blob events.
D. Use the Azure Monitor HTTP Data Collector API and scan the request body for successful blob events.

**Answer:** B

**Explanation:**
Change feed support in Azure Blob Storage
The purpose of the change feed is to provide transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account. The change feed provides ordered, guaranteed, durable, immutable, read-only log of these changes. Client applications can read these logs at any time, either in streaming or in batch mode. The change feed enables you to build efficient and scalable solutions that process change events that occur in your Blob Storage account at a low cost.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed

**NEW QUESTION 16**
- (Topic 8)
You are developing a solution that will use a multi-partitioned Azure Cosmos DB database. You plan to use the latest Azure Cosmos DB SDK for development.
The solution must meet the following requirements:
? Send insert and update operations to an Azure Blob storage account.
? Process changes to all partitions immediately.
? Allow parallelization of change processing.
You need to process the Azure Cosmos DB operations.
What are two possible ways to achieve this goal? Each correct answer presents a complete solution.
NOTE:Each correct selection is worth one point.

A. Create an Azure App Service API and implement the change feed estimator of the SD
B. Scale the API by using multiple Azure App Service instances.
C. Create a background job in an Azure Kubernetes Service and implement the change feed feature of the SDK.
D. Create an Azure Function to use a trigger for Azure Cosmos D
E. Configure the trigger toconnect to the container.
F. Create an Azure Function that uses a FeedIterator object that processes the change feed by using the pull model on the containe
G. Use a FeedRange objext to parallelize the processing of the change feed across multiple functions.

**Answer:** CD

**Explanation:**
Azure Functions is the simplest option if you are just getting started using the change feed. Due to its simplicity, it is also the recommended option for most change feed use cases. When you create an Azure Functions trigger for Azure Cosmos DB, you select the container to connect, and the Azure Function gets triggered whenever there is a change in the container. Because Azure Functions uses the change feed processor behind the scenes, it automatically parallelizes change processing across your container's partitions.
Note: You can work with change feed using the following options:
? Using change feed with Azure Functions
? Using change feed with change feed processor
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed
https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed#azure-functions
https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model#using-feedrange-for-parallelization

**NEW QUESTION 18**
- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure solution to collect point-of-sale (POS) device data from
2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the

future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Event Grid. Configure the machine identifier as the partition key and enable capture.
Does the solution meet the goal?

A. Yes
B. No

**Answer:** A

**Explanation:**
Reference:
https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

**NEW QUESTION 19**
DRAG DROP - (Topic 8)
You develop and deploy an Azure App Service ---- app. The web app accesses data in an Azure SQL database
You must update the web app to store frequently used data m a new Azure Cache for Redis Premium instance.
You need to implement the Azure Cache for Redis features.
Which feature should you implement? To answer, drag the appropriate feature to the correct requirements Each feature may be used once, more than once, or not at all You may need to ------------ between panes or scroll to view content.
NOTE Each correct selection is worth one point



A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**



**NEW QUESTION 21**
HOTSPOT - (Topic 8)
You are building a website to access project data related to terms within your organization. The website does not allow anonymous access. Authentication performed using an Azure Active Directory (Azure AD) app named internal.
The website has the following authentication requirements:
• Azure AD users must be able to login to the website.
• Personalization of the website must be based on membership in Active Directory groups. You need to configure the application's manifest to meet the authentication requirements.
How should you configure the manifest? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
{
    ...
    "appId": "d61126e3-089b-4adb-b721-
d5023213df7d",
```

|  ▼ | : "All", |
|---|---|
| "optionalClaims" | |
| "groupMembershipClaims" | |

|  ▼ | : true |
|---|---|
| "allowPublicClient" | |
| "oauth2Permissions" | |
| "requiredResourceAccess" | |
| "oauth2AllowImplicitFlow" | |

```
    ...
}
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: groupMembershipClaims
Personalization of the website must be based on membership in Active Directory groups. Group claims can also be configured in the Optional Claims section of the Application Manifest. Enable group membership claims by changing the groupMembershipClaim The valid values are:
- "All"
- "SecurityGroup"
- "DistributionList"
- "DirectoryRole"
Here we need to mention that we want to get the groups for the users. Hence we need to mention to set the groupMembershipClaims property to All.
Box 2: oauth2AllowImplicitFlow
Azure AD users must be able to login to the website.
auth2Permissions can only accept collections value like an array, not a boolean. oauth2AllowImplicitFlow accepts boolean value.
Here from the list of options given, if we want the application to fetch the required tokens , we would need to allow Implicit Flow.

**NEW QUESTION 23**
HOTSPOT - (Topic 8)
You are developing an online game that allows players to vote for their favorite photo that illustrates a word. The game is built by using Azure Functions and uses durable entities to track the vote count
The voting window is 30 seconds. You must minimize latency. You need to implement the Azure Function for voting.
How should you complete the code? To answer, select the appropriate options in the answer area.

**Answer Area**

```
[FunctionName("Vote")]
public static async Task<HttpResponseMessage> Run(
    [HttpTrigger("POST", Route = "pic/{id}")] HttpRequestMessage req,
```

| SignalEntityAsync ▼ | c, |
|---|---|
| CallEntityAsync | |
| SignalEntityAsync | |
| [DurableClient] IDurableEntityClient | |
| [DurableClient] IDurableOrchestrationClient | |

```
{
```

```
    return req.CreateResponse(HttpStatusCode.OK);
{
    var eid = new EntityId("pic", id);

    await c. [[DurableClient] IDurableEntityClient ▼] (eid, "vote");
    return req.Cr | CallEntityAsync
}               | SignalEntityAsync
                | [DurableClient] IDurableEntityClient
                | [DurableClient] IDurableOrchestrationClient
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

```
Answer Area

    [FunctionName("Vote")]
    public static async Task<HttpResponseMessage> Run(
        [HttpTrigger("POST", Route = "pic/{id}")] HttpRequestMessage req,
        ┌─────────────────────────────────────────┬───┐
        │ SignalEntityAsync                        │ ▼ │ c,
        ├─────────────────────────────────────────┴───┤
        │ CallEntityAsync                              │
    {   │ SignalEntityAsync                            │
        │ [DurableClient] IDurableEntityClient         │
        │ [DurableClient] IDurableOrchestrationClient  │
        └──────────────────────────────────────────────┘

        return req.CreateResponse(HttpStatusCode.OK);
    {
        var eid = new EntityId("pic", id);

        await c.  ┌──────────────────────────────────────────┬───┐
                  │ [DurableClient] IDurableEntityClient      │ ▼ │ (eid, "vote");
        return req.Cr├────────────────────────────────────────┴───┤
                  │ CallEntityAsync                              │
    }             │ SignalEntityAsync                            │
                  │ [DurableClient] IDurableEntityClient         │
                  │ [DurableClient] IDurableOrchestrationClient  │
                  └──────────────────────────────────────────────┘
```
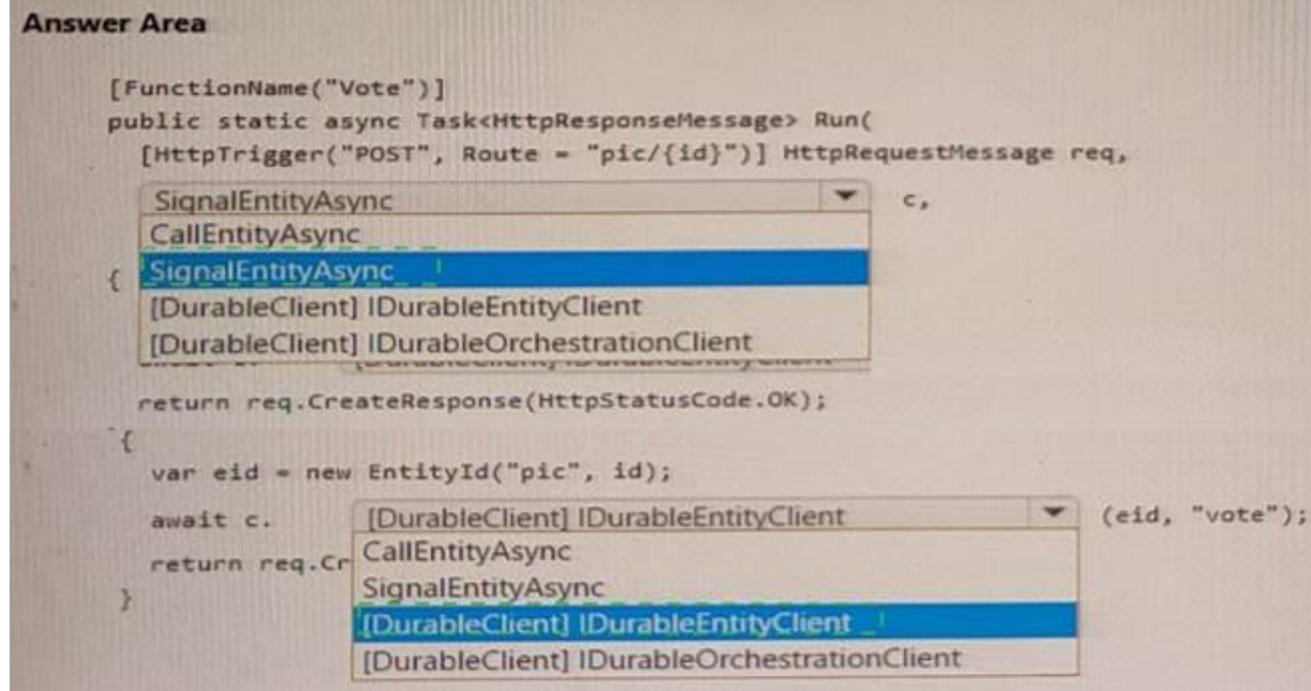
## NEW QUESTION 24

- (Topic 8)

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

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

You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is General-purpose V2.

When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.

You need to design the process that starts the photo processing.

Solution: Convert the Azure Storage account to a BlockBlobStorage storage account. Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**Explanation:**

Not necessary to convert the account, instead move photo processing to an Azure Function triggered from the blob upload..

Azure Storage events allow applications to react to events. Common Blob storage event scenarios include image or video processing, search indexing, or any file-oriented workflow.

Note: Only storage accounts of kind StorageV2 (general purpose v2) and BlobStorage support event integration. Storage (general purpose v1) does not support integration with Event Grid.

Reference:

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview

## NEW QUESTION 28

- (Topic 8)

You are developing a .Net web application that stores data is Azure Consmos DB. The application must use the Core API and allow millions of reads and writes. The Azure Cosmos DII account has been created with multiple write region enabled. The application has been deployed to the East US2 and Central US region.

You need to update the application to support multi-region writes.

What are two possible ways to achieve this goal? Each correct answer presents parts of the solutions.

NOTE: Each correct selection is worth one point.

A. Update the ConnectionPolicy class for the Cosmos client and populate the PreferredLocations property based on the geo-proximity of the application.
B. Update Azure Cosmos DB to use the Strong consistency leve
C. Add indexed properties to the container to indicate region.
D. Update the ConnectionPolicy class for the Cosmos client and set the UseMultipleWriteLocations property to true.
E. Create and deploy a custom conflict resolution policy.
F. Update Azure Cosmos DB to use the Session consistency leve
G. Send the SessionToken property value from the FeedResponse object of the write action to the end- user by using a cookie.

**Answer:** CD

## NEW QUESTION 31

- (Topic 8)

You are a developing a SaaS application that stores data as key value pairs.

You must make multiple editions of the application available. In the lowest cost edition, the performance must be best-effort, and there is no regional failover.

In higher cos! editions customers must be able to select guaranteed performance and support for multiple regions. Azure costs must be minimized.

Which Azure Cosmos OB API should you use for the application?

A. Core
B. MongoDB
C. Cassandra
D. Table API

**Answer:** D

**NEW QUESTION 34**
- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Event Grid. Configure event filtering to evaluate the device identifier.
Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**Explanation:**
Instead use an Azure Service Bus, which is used order processing and financial transactions.
Note: An event is a lightweight notification of a condition or a state change. Event hubs is usually used reacting to status changes.
Reference:
https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

**NEW QUESTION 39**
- (Topic 8)
You are developing an e-commerce solution that uses a microservice architecture.
You need to design a communication backplane for communicating transactional messages between various parts of the solution. Messages must be communicated in first- in-first-out (FIFO) order.
What should you use?

A. Azure Storage Queue
B. Azure Event Hub
C. Azure Service Bus
D. Azure Event Grid

**Answer:** C

**Explanation:**
As a solution architect/developer, you should consider using Service Bus queues when:
? Your solution requires the queue to provide a guaranteed first-in-first-out (FIFO) ordered delivery.
Reference:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted

**NEW QUESTION 40**
- (Topic 8)
You are designing a multi-tiered application that will be hosted on Azure virtual machines. The virtual machines will run Windows Server. Front-end servers will be accessible from the Internet over port 443. The other servers will NOT be directly accessible over the internet
You need to recommend a solution to manage the virtual machines that meets the following requirement
• Allows the virtual machine to be administered by using Remote Desktop.
• Minimizes the exposure of the virtual machines on the Internet Which Azure service should you recommend?

A. Azure Bastion
B. Service Endpoint
C. Azure Private Link
D. Azure Front Door

**Answer:** C

**NEW QUESTION 42**
HOTSPOT - (Topic 8)
You are developing a solution to store documents in Azure Blob storage. Customers upload documents to multiple containers. Documents consist of PDF, CSV, Microsoft Office format, and plain text files.
The solution must process millions of documents across hundreds of containers. The solution must meet the following requirements:
* Document must the categorized by a customer identifier as they are uploaded to the storage account.
* Allow filtering by the customer identifier.
* Allow searching of information contained within a document.
* Minimize costs.
You created and configure a standard general-purpose v2 storage account to support the solution.
You need to implement the solution.
NOTE: Each correct selection is worth one point.

Answer Area

| Requirement | Solution |
|---|---|
| Search and filter by customer identifier. | ▼ |
| | Azure Cognitive Search |
| | Azure Blob index tags |
| | Azure Blob inventory policy |
| | Azure Blob metadata |
| Search information inside documents. | ▼ |
| | Azure Cognitive Search |
| | Azure Blob index tags |
| | Azure Blob inventory policy |
| | Azure Blob metadata |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Azure Blob Index tags: https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob- index-how-to?tabs=azure-portal
Azure Cognitive Search: Search inside documents

**NEW QUESTION 45**
HOTSPOT - (Topic 8)
You are developing an application to store and retrieve data in Azure Blob storage. The application will be hosted in an on-premises virtual machine (VM). The VM is connected to Azure by using a Site-to-Site VPN gateway connection. The application is secured by using Azure Active Directory (Azure AD) credentials.
The application must be granted access to the Azure Blob storage account with a start time, expiry time, and read permissions. The Azure Blob storage account access must use the Azure AD credentials of the application to secure data access. Data access must be able to be revoked if the client application security is breached.
You need to secure the application access to Azure Blob storage.
Which security features should you use? To answer select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

| Component | Security Feature |
|---|---|
| Application (Client) | ▼ |
| | Storage Account Access Key |
| | System-assigned Managed Identity |
| | Shared access signature (SAS) token |
| Azure Storage (Server) | ▼ |
| | Stored Access Policy |
| | User-assigned Managed Identity |
| | Cross-Origin Resource Sharing (CORS) |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Shared access signature (SAS) token
When your application design requires shared access signatures for access to Blob storage, use Azure AD credentials to create a user delegation SAS when possible for superior security.
Box 2: Stored access policy
Stored access policies give you the option to revoke permissions for a service SAS without
having to regenerate the storage account keys.
A shared access signature can take one of the following two forms:
? Service SAS with stored access policy. A stored access policy is defined on a resource container, which can be a blob container, table, queue, or file share. The stored access policy can be used to manage constraints for one or more service shared access signatures. When you associate a service SAS with a stored access policy, the SAS inherits the constraints – the start time, expiry time, and permissions – defined for the stored access policy.
? Ad hoc SAS.

**NEW QUESTION 46**
HOTSPOT - (Topic 8)
You are developing a solution that uses the Azure Storage Client library for .NET. You have the following code: (Line numbers are included for reference only.)

```
01 CloudBlockBlob src = null;
02 try
03 {
04     src = container.ListBlobs().OfType<CloudBlockBlob>().FirstOrDefault();
05     var id = await src.AcquireLeaseAsync(null);
06     var dst = container.GetBlockBlobReference(src.Name);
07     string cpid = await dst.StartCopyAsync(src);
08     await dst.FetchAttributeAsync();
09     return id;
10 }
11 catch (Exception e)
12 {
13     throw;
14 }
15 finally
16 {
17     if (src != null)
18     await src.FetchAttributesAsync();
19     if (src.Properties.LeaseState != LeaseState.Available)
20     await src.BreakLeaseAsync(new TimeSpan(0));
21 }
```

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 code creates an infinite lease | ○ | ○ |
| The code at line 06 always creates a new blob | ○ | ○ |
| The finally block releases the lease | ○ | ○ |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Yes
AcquireLeaseAsync does not specify leaseTime.
leaseTime is a TimeSpan representing the span of time for which to acquire the lease, which will be rounded down to seconds. If null, an infinite lease will be acquired. If not null, this must be 15 to 60 seconds.
Box 2: No
The GetBlockBlobReference method just gets a reference to a block blob in this container.
Box 3: Yes
The BreakLeaseAsync method initiates an asynchronous operation that breaks the current lease on this container.

**NEW QUESTION 51**
- (Topic 8)
You are developing a software solution for an autonomous transportation system. The solution uses large data sets and Azure Batch processing to simulate navigation sets for entire fleets of vehicles.
You need to create compute nodes for the solution on Azure Batch. What should you do?

A. In Python, implement the class: TaskAddParameter
B. In Python, implement the class: JobAddParameter
C. In the Azure portal, create a Batch account
D. In a .NET method, call the method: BatchClient.PoolOperations.CreateJob

**Answer:** D

**Explanation:**
A Batch job is a logical grouping of one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. The app uses the BatchClient.JobOperations.CreateJob method to create a job on your pool.
Note:
Step 1: Create a pool of compute nodes. When you create a pool, you specify the number of compute nodes for the pool, their size, and the operating system. When each task in your job runs, it's assigned to execute on one of the nodes in your pool.
Step 2 : Create a job. A job manages a collection of tasks. You associate each job to a specific pool where that job's tasks will run.
Step 3: Add tasks to the job. Each task runs the application or script that you uploaded to process the data files it downloads from your Storage account. As each task completes, it can upload its output to Azure Storage.

**NEW QUESTION 56**
FILL IN THE BLANK - (Topic 8)

You are developing a web application by using the Azure SDK. The web application accesses data m a zone-redundant BlockBlobStorage storage account
The application must determine whether the data has changed since the application last reao the data. Update operations must use the latest data changes when writing data to the storages..................
You need to implement the update operations.
Which values should you use? To answer, select the appropriate option m the answer area.
NOTE Each correct selection is worth one point.

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

Answer Area

| Code evaluation | Value |
| --- | --- |
| HTTP Header value | VersionId ▼ |
| Conditional header | If-Match ▼ |
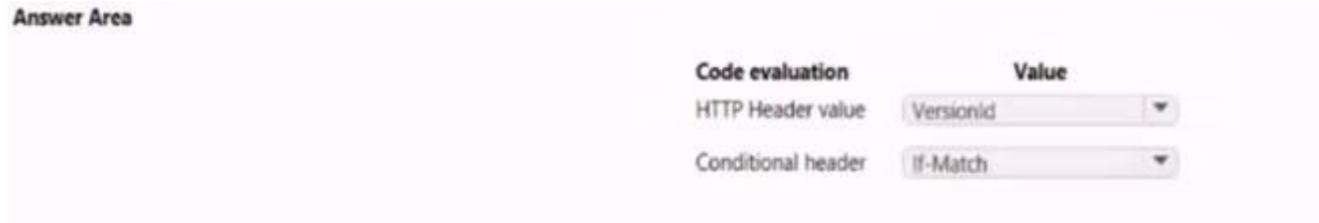
**NEW QUESTION 57**
HOTSPOT - (Topic 8)
You are building a traffic monitoring system that monitors traffic along six highways. The system produces time series analysis-based reports for each highway.
Data from traffic sensors are stored in Azure Event Hub.
Traffic data is consumed by four departments. Each department has an Azure Web App that displays the time-series-based reports and contains a WebJob that processes the incoming data from Event Hub. All Web Apps run on App Service Plans with three instances.
Data throughout must be maximized. Latency must be minimized. You need to implement the Azure Event Hub.
Which settings should you use? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

**Setting** | **Value**

Number of partitions ▼
3
4
6
12

Partition Key ▼
Highway
Department
Timestamp
VM name

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: 6
The number of partitions is specified at creation and must be between 2 and 32. There are 6 highways.
Box 2: Highway References:
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features

**NEW QUESTION 62**
DRAG DROP - (Topic 8)
You are developing a serverless Java application on Azure. You create a new Azure Key Vault to work with secrets from a new Azure Functions application.
The application must meet the following requirements:
? Reference the Azure Key Vault without requiring any changes to the Java code.
? Dynamically add and remove instances of the Azure Functions host based on the number of incoming application events.
? Ensure that instances are perpetually warm to avoid any cold starts.
? Connect to a VNet.
? Authentication to the Azure Key Vault instance must be removed if the Azure Function application is deleted.
You need to grant the Azure Functions application access to the Azure Key Vault.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

**Actions**

Create a user-assigned managed identity for the application.

Create the Azure Functions app with a Premium plan type.

Create an access policy in Azure Key Vault for the application identity.

Create an SSL certification in Azure Key Vault for the application identity.

Create the Azure Functions app with an App Service plan type.

Create the Azure Functions app with a Consumption plan type.

Create a system-assigned managed identity for the application.

**Answer Area**

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Step 1: Create the Azure Functions app with a Consumption plan type. Use the Consumption plan for serverless.
Step 2: Create a system-assigned managed identity for the application. Create a system-assigned managed identity for your application.
Key Vault references currently only support system-assigned managed identities. User- assigned identities cannot be used.
Step 3: Create an access policy in Key Vault for the application identity.
Create an access policy in Key Vault for the application identity you created earlier. Enable the "Get" secret permission on this policy. Do not configure the "authorized application" or
applicationId settings, as this is not compatible with a managed identity.


**NEW QUESTION 63**
HOTSPOT - (Topic 8)
You are working for Contoso, Ltd.
You define an API Policy object by using the following XML markup:

```
<set-variable name= "bodySize" value="@(context.Request.Headers["Content-Length"] [0])"/>
<choose>
 <when condition= "@(int.Parse(context.Variables.GetValueOrDefault<string> ("bodySize"))<512000)">
</when>
<otherwise>
    <rewrite-uri template= "/put"/>
    <set-backend-service base-url= "http://contoso.com/api/9.1/"/>
</otherwise>
</choose>
```

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 XML segment belongs in the \<inbound> section of the policy. | ○ | ○ |
| If the body size is >256k, an error will occur. | ○ | ○ |
| If the request is http://contoso.com/api/9.2/, the policy will retain the higher version. | ○ | ○ |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Yes
Use the set-backend-service policy to redirect an incoming request to a different backend than the one specified in the API settings for that operation. Syntax: <set-backend-service base-url="base URL of the backend service" />
Box 2: No

The condition is on 512k, not on 256k.
Box 3: No
The set-backend-service policy changes the backend service base URL of the incoming
request to the one specified in the policy.


**NEW QUESTION 64**
- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the
stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage
Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version
of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Use the Azure Blob Storage change feed to trigger photo processing. Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**Explanation:**
The change feed is a log of changes that are organized into hourly segments but appended to and updated every few minutes. These segments are created only
when there are blob change events that occur in that hour.
Instead catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview


**NEW QUESTION 67**
DRAG DROP - (Topic 8)
You have an application that provides weather forecasting data to external partners. You use Azure API Management to publish APIs.
You must change the behavior of the API to meet the following requirements:
• Support alternative input parameters.
• Remove formatting text from responses.
• Provide additional context to back-end services.
Which types of policies should you implement? To answer, drag the policy types to the correct scenarios. Each policy type 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.

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

**NEW QUESTION 70**
DRAG DROP - (Topic 8)
A company has multiple warehouse. Each warehouse contains IoT temperature devices which deliver temperature data to an Azure Service Bus queue.
You need to send email alerts to facility supervisors immediately if the temperature at a warehouse goes above or below specified threshold temperatures.
Which five actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the
correct order.

**Actions**

| Add a logic app trigger that fires when one or more messages arrive in the queue. |
| Add a Recurrence trigger that schedules the app to run every 15 minutes. |
| Add an action that sends an email to specified personnel if the temperature is outside of those thresholds. |
| Add a trigger that reads IoT temperature data from a Service Bus queue. |
| Add a logic app action that fires when one or more messages arrive in the queue. |
| Add a condition that compares the temperature against the upper and lower thresholds. |
| Create a blank Logic app. |
| Add an action that reads IoT temperature data from the Service Bus queue. |

**Answer Area**

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Step 1: Create a blank Logic app. Create and configure a Logic App.
Step 2: Add a logical app trigger that fires when one or more messages arrive in the queue. Configure the logic app trigger.
Under Triggers, select When one or more messages arrive in a queue (auto-complete). Step 3: Add an action that reads IoT temperature data from the Service Bus queue
Step 4: Add a condition that compares the temperature against the upper and lower thresholds.
Step 5: Add an action that sends an email to specified personnel if the temperature is outside of those thresholds

**NEW QUESTION 73**
- (Topic 8)
You are developing an application to store business-critical data in Azure Blob storage. The application must meet the following requirements:
• Data must not be modified or deleted for a user-specified interval.
• Data must be protected from overwrites and deletes.
• Data must be written once and allowed to be read many times. You need to protect the data fen the Azure Blob storage account.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Enable version-level immutability support for the storage account.
B. Create an account shared-access signature (SAS).
C. Enable point-in-time restore for containers in the storage account.
D. Create a service shared-access signature (SAS).
E. Enable the blob change feed for the storage account.
F. Configure a time-based retention policy for the storage account.

**Answer:** DF

**NEW QUESTION 78**
- (Topic 8)
A development team is creating a new REST API. The API will store data in Azure Blob storage. You plan to deploy the API to Azure App Service.
Developers must access the Azure Blob storage account to develop the API for the next two months. The Azure Blob storage account must not be accessible by the developers after the two-month time period.
You need to grant developers access to the Azure Blob storage account. What should you do?

A. Generate a shared access signature (SAS) for the Azure Blob storage account and provide the SAS to all developers.
B. Create and apply a new lifecycle management policy to include a last accessed date valu
C. Apply the policy to the Azure Blob storage account.
D. Provide all developers with the access key for the Azure Blob storage accoun
E. Update the API to include the Coordinated Universal Time (UTC) timestamp for the request header.
F. Grant all developers access to the Azure Blob storage account by assigning role-based access control (RBAC) roles.

**Answer:** A

**Explanation:**
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

**NEW QUESTION 83**
- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.
You are developing and deploying several ASP.Net web applications to Azure App Service. You plan to save session state information and HTML output. You must use a storage mechanism with the following requirements:

•Share session state across all ASP.NET web applications
•Support controlled, concurrent access to the same session state data for multiple readers and a single writer
•Save full HTTP responses for concurrent requests You need to store the information.
Proposed Solution: Deploy and configure Azure Cache for Redis. Update the web applications.
Does the solution meet the goal?

A. Yes
B. No

**Answer:** A

**Explanation:**
The session state provider for Azure Cache for Redis enables you to share session information between different instances of an ASP.NET web application.
The same connection can be used by multiple concurrent threads. Redis supports both read and write operations.
The output cache provider for Azure Cache for Redis enables you to save the HTTP responses generated by an ASP.NET web application.
Note: Using the Azure portal, you can also configure the eviction policy of the cache, and control access to the cache by adding users to the roles provided. These roles, which define the operations that members can perform, include Owner, Contributor, and Reader. For example, members of the Owner role have complete control over the cache (including security) and its contents, members of the Contributor role can read and write information
in the cache, and members of the Reader role can only retrieve data from the cache.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching


**NEW QUESTION 86**
HOTSPOT - (Topic 8)
You develop a containerized application. You plan to deploy the application to a new Azure Container instance by using a third-party continuous integration and continuous delivery (CI/CD) utility.
The deployment must be unattended and include all application assets. The third-party utility must only be able to push and pull images from the registry. The authentication must be managed by Azure Active Directory (Azure AD). The solution must use the principle of least privilege.
You need to ensure that the third-party utility can access the registry.
Which authentication options should you use? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.



A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Service principal
Applications and container orchestrators can perform unattended, or "headless," authentication by using an Azure Active Directory (Azure AD) service principal.
Box 2: AcrPush
AcrPush provides pull/push permissions only and meets the principle of least privilege.


**NEW QUESTION 91**
HOTSPOT - (Topic 8)
You are developing a web application that uses the Microsoft identify platform for user and resource authentication. The web application calls several REST APIs.
You are implementing various authentication and authorization flows for the web application.
You need to validate the claims in the authentication token.
Which token type should use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

| Requirement | Token type |
|---|---|
| Identify users for the application by using a JWT token that contains claims. | ▼ |
| | Access |
| | **ID** |
| | Refresh |
| | SAML |
| Provide XML representations of claims that can be consumed by applications that use WS-Federation. | Access ▼ |
| Provide the web application with long-term access to resources on behalf of users without requiring interaction with those users. | ▼ |
| | Access |
| | ID |
| | Refresh |
| | SAML |
| Provide XML representations of claims that can be consumed by applications that use WS-Federation. | ▼ |
| | Access |
| | ID |
| | **Refresh** |
| | SAML |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

Answer Area

| Requirement | Token type |
|---|---|
| Identify users for the application by using a JWT token that contains claims. | ▼ |
| | Access |
| | **ID** |
| | Refresh |
| | SAML |
| Provide XML representations of claims that can be consumed by applications that use WS-Federation. | Access ▼ |
| Provide the web application with long-term access to resources on behalf of users without requiring interaction with those users. | ▼ |
| | Access |
| | ID |
| | Refresh |
| | SAML |
| Provide XML representations of claims that can be consumed by applications that use WS-Federation. | ▼ |
| | Access |
| | ID |
| | **Refresh** |
| | SAML |

**NEW QUESTION 96**
HOTSPOT - (Topic 8)
You are developing an application that use an Azure blob named data to store application data. The application creates blob snapshots to allow application state to be reverted to an earlier state. The Azure storage account has soft deleted enabled.
The system performs the following operations in order:
•The blob is updated
•Snapshot 1 is created.
•Snapshot 2 is created.
•Snapshot 1 is deleted.
A system error then deletes the data blob and all snapshots.
You need to determine which application states can be restored.
What is the restorability of the application data? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

**Application State**       **Restorability**

Data blob

Can be restored
Cannot be restored

Snapshot 1

Can be restored
Cannot be restored

Snapshot 2

Can be restored
Cannot be restored

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Can be restored
When enabled, soft delete enables you to save and recover your data when blobs or blob snapshots are deleted. This protection extends to blob data that is erased as the result of an overwrite.
Box 2: Cannot be restored It has been deleted.
Box 3: Can be restored It has not been deleted.
References:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-soft-delete

**NEW QUESTION 101**
- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data. You need to ensure the app does not time out and processes the blob data.
Solution: Use the Durable Function async pattern to process the blob data. Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**Explanation:**
Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.
Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:
Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices

**NEW QUESTION 103**
- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
Margie's Travel is an international travel and bookings management service. The company
is expanding into restaurant bookings. You are tasked with implementing Azure Search tor the restaurants listed in their solution.
You create the index in Azure Search.
You need to import the restaurant data into the Azure Search service by using the Azure Search NET SDK.
Solution:
* 1. Create a SearchServiceClient object to connect to the search index.
* 2. Create a DataContainer that contains the documents which must be added.
* 3. Create a DataSource instance and set its Container property to the DataContainer.
* 4. Set the DataSource property of the SearchServiceClient
Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**Explanation:**
Use the following method:
* 1.Create a SearchIndexClient object to connect to the search index
* 2.Create an IndexBatch that contains the documents which must be added.
* 3.Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.
References:
https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk


**NEW QUESTION 108**
- (Topic 8)
You are developing a software solution for an autonomous transportation system. The solution uses large data sets and Azure Batch processing to simulate navigation sets for entire fleets of vehicles.
You need to create compute nodes for the solution on Azure Batch. What should you do?

A. In the Azure portal, create a Batch account.
B. In a .NET method, call the method:BatchClient.PoolOperations.CreatePool
C. In Python, implement the class:JobAddParameter
D. In Python, implement the class:TaskAddParameter

**Answer:** B

**Explanation:**
A Batch job is a logical grouping of one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. The app uses the BatchClient.JobOperations.CreateJob method to create a job on your pool.


**NEW QUESTION 110**
- (Topic 8)
You are developing applications for a company. You plan to host the applications on Azure App Services.
The company has the following requirements:
? Every five minutes verify that the websites are responsive.
? Verify that the websites respond within a specified time threshold. Dependent requests such as images and JavaScript files must load properly.
? Generate alerts if a website is experiencing issues.
? If a website fails to load, the system must attempt to reload the site three more times.
You need to implement this process with the least amount of effort. What should you do?
D18912E1457D5D1DDCBD40AB3BF70D5D

A. Create a Selenium web test and configure it to run from your workstation as a scheduled task.
B. Set up a URL ping test to query the home page.
C. Create an Azure function to query the home page.
D. Create a multi-step web test to query the home page.
E. Create a Custom Track Availability Test to query the home page.

**Answer:** D

**Explanation:**
You can monitor a recorded sequence of URLs and interactions with a website via multi- step web tests.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-multistep


**NEW QUESTION 112**
HOTSPOT - (Topic 8)
You are developing an Azure App Service hosted ASP.NET Core web app to deliver video on-demand streaming media. You enable an Azure Content Delivery Network (CDN) Standard for the web endpoint. Customer videos are downloaded from the web app by using the following example URL.:
http://www.contoso.com/content.mp4?quality=1
All media content must expire from the cache after one hour. Customer videos with varying quality must be delivered to the closest regional point of presence (POP) node.
You need to configure Azure CDN caching rules.
Which options should you use? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

| Setting | Action |
|---|---|
| Caching behavior | ▼ |
| | Bypass cache |
| | Override |
| | Set if missing |
| Cache expiration duration | ▼ |
| | 1 second |
| | 1 minute |
| | 1 hour |
| | 1 day |
| Query string caching behavior | ▼ |
| | Ignore query strings |
| | Bypass caching for query strings |
| | Cache every unique URL |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Override
Override: Ignore origin-provided cache duration; use the provided cache duration instead. This will not override cache-control: no-cache.
Set if missing: Honor origin-provided cache-directive headers, if they exist; otherwise, use the provided cache duration.
Incorrect:
Bypass cache: Do not cache and ignore origin-provided cache-directive headers.
Box 2: 1 hour
All media content must expire from the cache after one hour.
Box 3: Cache every unique URL
Cache every unique URL: In this mode, each request with a unique URL, including the query string, is treated as a unique asset with its own cache. For example, the response from the origin server for a request for example.ashx?q=test1 is cached at the POP node and returned for subsequent caches with the same query string. A request for example.ashx?q=test2 is cached as a separate asset with its own time-to-live setting.

**NEW QUESTION 116**
- (Topic 8)
D18912E1457D5D1DDCBD40AB3BF70D5D
You are building a website that uses Azure Blob storage for data storage. You configure Azure Blob storage lifecycle to move all blobs to the archive tier after 30 days.
Customers have requested a service-level agreement (SLA) for viewing data older than 30 days.
You need to document the minimum SLA for data recovery. Which SLA should you use?

A. at least two days
B. between one and 15 hours
C. at least one day
D. between zero and 60 minutes

**Answer:** B

**Explanation:**
The archive access tier has the lowest storage cost. But it has higher data retrieval costs compared to the hot and cool tiers. Data in the archive tier can take several hours to retrieve depending on the priority of the rehydration. For small objects, a high priority rehydrate may retrieve the object from archive in under 1 hour.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal

**NEW QUESTION 119**
DRAG DROP - (Topic 8)
You are preparing to deploy a medical records application to an Azure virtual machine (VM). The application will be deployed by using a VHD produced by an on-premises build server.
You need to ensure that both the application and related data are encrypted during and after deployment to Azure.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

**Actions**

| |
|---|
| Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage. |
| Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension. |
| Run the Azure PowerShell command Set-AzureRmVMOSDisk. |
| Encrypt the on-premises VHD by using BitLocker with a TPM. Upload the VM to Azure Storage. |
| Run the Azure PowerShell command New-AzureRmVM. |

**Answer area**

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Step 1: Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage
Step 2: Run the Azure PowerShell command Set-AzureRMVMOSDisk
To use an existing disk instead of creating a new disk you can use the Set- AzureRMVMOSDisk command.
Example:
$osDiskName = $vmname+'_osDisk'
$osDiskCaching = 'ReadWrite'
$osDiskVhdUri = "https://$stoname.blob.core.windows.net/vhds/"+$vmname+"_os.vhd"
$vm = Set-AzureRmVMOSDisk -VM $vm -VhdUri $osDiskVhdUri -name $osDiskName - Create
Step 3: Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension Use the Set-AzVMDiskEncryptionExtension cmdlet to enable encryption on a running IaaS virtual machine in Azure.
Incorrect:
Not TPM: BitLocker can work with or without a TPM. A TPM is a tamper resistant security chip on the system board that will hold the keys for encryption and check the integrity of the boot sequence and allows the most secure BitLocker implementation. A VM does not have a TPM.
References:
https://www.itprotoday.com/iaaspaas/use-existing-vhd-azurerm-vm

**NEW QUESTION 122**
HOTSPOT - (Topic 8)
You are developing a back-end Azure App Service that scales based on the number of
messages contained in a Service Bus queue.
A rule already exists to scale up the App Service when the average queue length of unprocessed and valid queue messages is greater than 1000.
You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.
How should you configure the Scale rule? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

**Answer Area**

Scale rule ✕

Metric source

| |
|---|
| Storage queue |
| Service Bus queue |
| Current resource |
| Storage queue (classic) |

Resource type
Service Bus Namespaces ▾

Resource
MessageQueue1103 ▾

* Queues
itemqueue ▾

Criteria
* Metric name

| |
|---|
| Message Count |
| Active Message Count |

1 minute time grain

* Time grain statistic ⓘ

| |
|---|
| Total |
| Maximum |
| Average |
| Count |

* Operator

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Service bus queue
You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.
Box 2: ActiveMessage Count
ActiveMessageCount: Messages in the queue or subscription that are in the active state and ready for delivery.
Box 3: Count
Box 4: Less than or equal to
You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.
Box 5: Decrease count by

**NEW QUESTION 125**
- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Create an Azure Function app that uses the Consumption hosting model and that is triggered from the blob upload.
Does the solution meet the goal?

A. Yes
B. No

**Answer:** A

**Explanation:**
In the Consumption hosting plan, resources are added dynamically as required by your functions.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob- triggered-function

**NEW QUESTION 129**
HOTSPOT - (Topic 8)
You are developing an Azure-hosted e-commerce web application. The application will use Azure Cosmos DB to store sales orders. You are using the latest SDK to manage the sales orders in the database.
You create a new Azure Cosmos DB instance. You include a valid endpoint and valid authorization key to an appSettings.json file in the code project.
You are evaluating the following application code: (Line number are included for reference only.)

```
01 using System;
02 using System.Threading.Tasks;
03 using Microsoft.Azure.Cosmos;
04 using Microsoft.Extensions.Configuration;
05 using Newtonsoft.Json;
06 namespace SalesOrders
07 {
08   public class SalesOrder
09   {
10     . . .
11   }
12   internal class ManageSalesOrders
13   {
14     private static async Task GenerateSalesOrders()
15     {
16       IConfigurationRoot configuration = new ConfigurationBuilder().AddJsonFile("appSettings.json").Build();
17       string endpoint = configuration["EndPointUrl"];
18       string authKey = configuration["AuthorizationKey"];
19       using CosmosClient client = new CosmosClient(endpoint, authKey);
20       Database database = null;
21       using (await client.GetDatabase("SalesOrders").DeleteStreamAsync()) { }
22       database = await client.CreateDatabaseIfNotExistsAsync("SalesOrders");
23       Container container1 = await database.CreateContainerAsync(id: "Container1", partitionKeyPath: "/AccountNumber");
24       Container container2 = await database.CreateContainerAsync(id: "Container2", partitionKeyPath: "/AccountNumber");
25       SalesOrder salesOrder1 = new SalesOrder() { AccountNumber = "123456" };
26       await container1.CreateItemAsync(salesOrder1, new PartitionKey(salesOrder1.AccountNumber));
27       SalesOrder salesOrder2 = new SalesOrder() { AccountNumber = "654321" };
28       await container1.CreateItemAsync(salesOrder2, new PartitionKey(salesOrder2.AccountNumber));
29       SalesOrder salesOrder3 = new SalesOrder() { AccountNumber = "109876" };
30       await container2.CreateItemAsync(salesOrder3, new PartitionKey(salesOrder3.AccountNumber));
31       _ = await database.CreateUserAsync("User1");
32       User user1 = database.GetUser("User1");
33       _ = await user1.ReadAsync();
34     }
35   }
36 }
```

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

| Statements | Yes | No |
|---|---|---|
| A database named SalesOrders is created. The database will include two containers. | ○ | ○ |
| Container1 will contain two items. | ○ | ○ |
| Container2 will contain one item. | ○ | ○ |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: Yes
The createDatabaseIfNotExistsAsync method checks if a database exists, and if it doesn't, create it.
The Database.CreateContainerAsync method creates a container as an asynchronous operation in the Azure Cosmos service.
Box 2: Yes
The CosmosContainer.CreateItemAsync method creates an item as an asynchronous operation in the Azure Cosmos service.
Box 3: Yes

**NEW QUESTION 131**
HOTSPOT - (Topic 8)
You have a single page application (SPA) web application that manages information based on data returned by Microsoft Graph from another company's Azure Active Directory (Azure AD) instance.
Users must be able to authenticate and access Microsoft Graph by using their own company's Azure AD instance.
You need to configure the application manifest for the app registration.
How should you complete the manifest? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
{
    "oauth2AllowImplicitFlow": [          ▼] ,
                               add
                               false
                               spa
                               true

    "[                ▼]": [{
        addIns
        orgRestrictions
        availableToOtherTenants
        requiredResourceAccess

    "resourceAppId": "00000003-0000-0000-c000-000000000000",
    "resourceAccess": [{
        "id" : "24a6cdd6-fab1-4aaf-91b8-3cc8225e90d0",
        "type": "Scope"
    }]}],
    "signInAudience": "[                ▼]"
                        All
                        AzureADMyOrg
                        AzureADMultipleOrgs
                        AzureADandPersonalMicrosoftAccount

}
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: true
The oauth2AllowImplicitFlow attribute Specifies whether this web app can request OAuth2.0 implicit flow access tokens. The default is false. This flag is used for browser- based apps, like JavaScript single-page apps.
In implicit flow, the app receives tokens directly from the Azure Active Directory (Azure AD) authorize endpoint, without any server-to-server exchange. All authentication logic and session handling is done entirely in the JavaScript client with either a page redirect or a pop-up box.
Box 2: requiredResourceAccess

With dynamic consent, requiredResourceAccess drives the admin consent experience and the user consent experience for users who are using static consent. However, this parameter doesn't drive the user consent experience for the general case.

resourceAppId is the unique identifier for the resource that the app requires access to. This value should be equal to the appId declared on the target resource app.

resourceAccess is an array that lists the OAuth2.0 permission scopes and app roles that the app requires from the specified resource. Contains the id and type values of the specified resources.

Example: "requiredResourceAccess": [
{
"resourceAppId": "00000002-0000-0000-c000-000000000000",
"resourceAccess": [
{
"id": "311a71cc-e848-46a1-bdf8-97ff7156d8e6", "type": "Scope"
}
]
}
],

Box 3: AzureADMyOrg
The signInAudience attribute specifies what Microsoft accounts are supported for the current application. Supported values are:
AzureADMyOrg - Users with a Microsoft work or school account in my organization's Azure AD tenant (for example, single tenant)
AzureADMultipleOrgs - Users with a Microsoft work or school account in any organization's Azure AD tenant (for example, multi-tenant)
AzureADandPersonalMicrosoftAccount - Users with a personal Microsoft account, or a work or school account in any organization's Azure AD tenant


**NEW QUESTION 133**
- (Topic 8)
You develop and deploy an ASP.NET web app to Azure App Service. You use Application Insights telemetry to monitor the app.
You must test the app to ensure that the app is available and responsive from various points around the world and at regular intervals. If the app is not responding, you must send an alert to support staff.
You need to configure a test for the web app.
Which two test types can you use? Each correct answer presents a complete solution.
NOTE:Each correct selection is worth one point.

A. integration
B. multi-step web
C. URL ping
D. unit
E. load

**Answer:** BC

**Explanation:**
There are three types of availability tests:
? URL ping test: a simple test that you can create in the Azure portal.
? Multi-step web test: A recording of a sequence of web requests, which can be played back to test more complex scenarios. Multi-step web tests are created in Visual Studio Enterprise and uploaded to the portal for execution.
? Custom Track Availability Tests: If you decide to create a custom application to run availability tests, the TrackAvailability() method can be used to send the results to Application Insights.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability


**NEW QUESTION 137**
HOTSPOT - (Topic 8)
You develop two Python scripts to process data.
The Python scripts must be deployed to two, separate Linux containers running in an Azure Container Instance container group. The containers must access external data by using the Server Message Block (SMB) protocol. Containers in the container group must run only once
You need to configure the Azure Container Instance.
Which configuration value should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

**Answer Area**

| Configuration Setting | Configuration Value |
| --- | --- |
| External data volume | Secret<br>Empty directory<br>Cloned git repo<br>Azure file share |
| Container restart policy | Never<br>Always<br>OnFailure |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

## Answer Area

| Configuration Setting | Configuration Value |
|---|---|
| External data volume | Secret / Empty directory / Cloned git repo / **Azure file share** |
| Container restart policy | **Never** / Always / OnFailure |

---

**NEW QUESTION 138**
HOTSPOT - (Topic 8)
An organization deploys a blob storage account. Users take multiple snapshots of the blob storage account over time. You need to delete all snapshots of the blob storage account. You must not delete the blob storage account itself. How should you complete the code segment? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Answer Area**

```
Delete (Azure.Storage.Blobs.Models.DeleteSnapshotsOption
snapshotsOption = Azure.Storage.Blobs.Models.    [DeleteSnapshotsOption ▼]  .  [OnlySnapshots ▼]  )
                                                    DeleteIfExists                IncludeSnapshots
                                                    DeleteSnapshotsOption         None
                                                    WithSnapshot                  OnlySnapshots
                                                    WithSnapshotCore
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

```
Delete (Azure.Storage.Blobs.Models.DeleteSnapshotsOption
snapshotsOption = Azure.Storage.Blobs.Models.    [DeleteSnapshotsOption ▼]  .  [OnlySnapshots ▼]  )
                                                    DeleteIfExists                IncludeSnapshots
                                                    DeleteSnapshotsOption         None
                                                    WithSnapshot                  OnlySnapshots
                                                    WithSnapshotCore
```

---

**NEW QUESTION 141**
HOTSPOT - (Topic 8)
You are developing an Azure Web App. You configure TLS mutual authentication for the web app.
You need to validate the client certificate in the web app. To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

| Property | Value |
|---|---|

Client certificate location ▼

HTTP request header
Client cookie
HTTP message body
URL query string

Encoding type ▼

HTML
URL
Unicode
Base64

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Accessing the client certificate from App Service.
If you are using ASP.NET and configure your app to use client certificate authentication, the certificate will be available through the HttpRequest.ClientCertificate property. For other application stacks, the client cert will be available in your app through a base64 encodedvalue in the "X-ARR-ClientCert" request header. Your application can create a certificate from this value and then use it for authentication and authorization purposes in your application.
References:
https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual- auth

**NEW QUESTION 143**
- (Topic 8)
You develop and deploy an Azure App Service web app. The app is deployed to multiple regions and uses Azure Traffic Manager. Application Insights is enabled for the app.
You need to analyze app uptime for each month.
Which two solutions win achieve the goal? Each correct answer presents a complete solution
NOTE: Each correct selection is worth one point

A. Application Insights alerts
B. Application Insights web tests
C. Azure Monitor logs
D. Azure Monitor metrics

**Answer:** AC

**Explanation:**
Reference:
https://azure.microsoft.com/en-us/blog/creating-a-web-test-alert-programmatically-with-application-insights/

**NEW QUESTION 144**
DRAG DROP - (Topic 8)
You provision virtual machines (VMs) as development environments. One VM does not have host.
The VM is stuck in a Windows update process. You attach the OS disk for the affected VM to a recovery VM.
You need to correct the issue.
In which order should you perform the actions' To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

| Actions | | Answer area |
|---|---|---|
| Open C:\temp\Patch.txt file and locate the update that is in a pending s | | 1 |
| Run the following command at an elevated command prompt: | | |
| dism /Image:<Attached OS disk>:\ /Remove-Package /PackageName= ... delete> | | 2 |
| Run the following command at an elevated command prompt: | | |
| dism /Image:\ /get-packages > C:\temp\Patch.txt | | 3 |
| Detach the OS disk and recreate the VM. | | 4 |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Remove the update that causes the problem
? Take a snapshot of the OS disk of the affected VM as a backup.
? Attach the OS disk to a recovery VM.
? Once the OS disk is attached on the recovery VM, run diskmgmt.msc to open Disk Management, and ensure the attached disk is ONLINE.
? (Step 1) Open an elevated command prompt instance (Run as administrator). Run the following command to get the list of the update packages that are on the

attached OS disk:
dism /image:<Attached OS disk>:\ /get-packages > c:\temp\Patch_level
? (Step 2) Open the C:\temp\Patch_level.txt file, and then read it from the bottom up.
Locate the update that's in Install Pending or Uninstall Pending state.
? Remove the update that caused the problem:
dism /Image:<Attached OS disk>:\ /Remove-Package /PackageName:<PACK
? (Step 4) Detach the OS disk and recreate the VM. Then check whether the issue is resolved.

**NEW QUESTION 147**
- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
Margie's Travel is an international travel and bookings management service. The company is expanding into restaurant bookings. You are tasked with implementing Azure Search for the restaurants listed in their solution.
You create the index in Azure Search.
You need to import the restaurant data into the Azure Search service by using the Azure Search .NET SDK.
Solution:
* 1. Create a SearchServiceClient object to connect to the search index.
* 2. Create a DataContainer that contains the documents which must be added.
* 3. Create a DataSource instance and set its Container property to the DataContainer.
* 4. Set the DataSources property of the SearchServiceClient.
Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**Explanation:**
Use the following method:
* 1.- Create a SearchIndexClient object to connect to the search index
* 2.- Create an IndexBatch that contains the documents which must be added.
* 3.- Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.
References:
https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk

**NEW QUESTION 152**
- (Topic 8)
You develop an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage containers named Container1 and Container2.
Uploading of videos occurs on an irregular basis.
You need to copy specific blobs from Container1 to Container2 when a new video is uploaded.
What should you do?

A. Copy blobs to Container2 by using thePut Bloboperation of the Blob Service REST API
B. Create anEvent Gridtopic that uses theStart-AzureStorageBlobCopycmdlet
C. UseAzCopywith theSnapshotswitch to copy blobs to Container2
D. Download the blob to a virtual machine and then upload the blob to Container2

**Answer:** B

**Explanation:**
The Start-AzureStorageBlobCopy cmdlet starts to copy a blob. Example 1: Copy a named blob
C:\PS>Start-AzureStorageBlobCopy -SrcBlob "ContosoPlanning2015" -DestContainer "ContosoArchives" -SrcContainer "ContosoUploads"
This command starts the copy operation of the blob named ContosoPlanning2015 from the container named ContosoUploads to the container named ContosoArchives.
Reference:
https://docs.microsoft.com/en-us/powershell/module/azure.storage/start- azurestorageblobcopy?view=azurermps-6.13.0

**NEW QUESTION 157**
HOTSPOT - (Topic 8)
You have an Azure Batch project that processes and converts files and stores the files in Azure storage. You are developing a function to start the batch job.
You add the following parameters to the function.

| Parameter name | Description |
|---|---|
| fileTasks | a list of tasks to be run |
| jobId | the identifier that must be assigned to the job |
| outputContainerSasUrl | a storage SAS URL to store successfully converted files |
| failedContainerSasUrl | a storage SAS URL to store copies of files that failed to convert. |

You must ensure that converted files are placed in the container referenced by the outputContainerSasUrl parameter. Files which fail to convert are places in the container referenced by the failedContainerSasUrl parameter.
You need to ensure the files are correctly processed.
How should you complete the code segment? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

## Answer Area

```
public list<CloudTasks> StartTasks(List<FileTask> fileTasks, string jobId,
 string outputContainerSasUrl, string failedContainerSasUrl)
{
  BatchSharedKeyCredentials sharedKeyCredentials =
   new BatchSharedKeyCredentials(batchAccountUrl, batchAccountName,
batchAccountKey);
  List<CloudTask> tasks = new List<CloudTask>();
  using (BatchClient batchClient = BatchClient.Open(sharedKeyCredentials))
  {
     CloudJob = batchClient.JobOperations. [▼] ();
                                            GetJob
                                            GetTask
                                            EnableJob
                                            CreateJob

          job.Id = jobId,
          job.PoolInfromation = new PoolInformation { PoolId = poolId };
          job.Commit();
          fileTasks.ForEach((fileTask) =>
          {
            string taskId = $"Task{DateTime.Now.ToFileTimeUtc().ToString()}";
            CloudTask task = new CloudTask (taskId, fileTask.Command);
            List<OutputFile> outputFileList = new List<OutputFile>();
            OutputFileBlobContainerDestination outputContainer =
              new OutputFileBlobContainerDestination(outputContainerSasUrl);
          OutputFileBlobContainerDestination failedContainer =
    new OutputFileBlobContainerDestination (failedContainerSasUrl);
  outputFileList.Add(new OutputFile(fileTask.Output,
    new OutputFileDestination(outputContainer),
    new OutputFileUploadOptions(OutputFileUploadCondition. [▼] )));
                                                          TaskSuccess
                                                          TaskFailure
                                                          TaskCompletion

      outputFileList.Add(new OutputFile(fileTask.Output,
       new OutputFileDestination(failedContainer),
       new OutputFileUploadOptions(OutputFileUploadCondition, [▼] )));
                                                          TaskSuccess
                                                          TaskFailure
                                                          TaskCompletion

      task [▼] =outputFileList;
           OutputFiles
           FilesToStage
           ResourceFiles
           StageFiles
      task.Add(task);
     });
     ]
     return tasks,
}
```
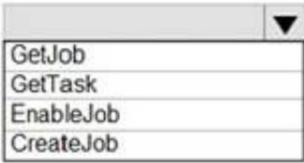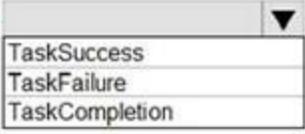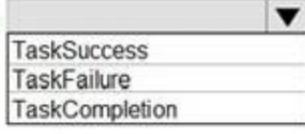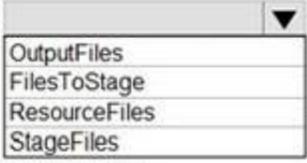
A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: CreateJob
Box 2: TaskSuccess
TaskSuccess: Upload the file(s) only after the task process exits with an exit code of 0.
Incorrect: TaskCompletion: Upload the file(s) after the task process exits, no matter what the exit code was.
Box 3: TaskFailure
TaskFailure:Upload the file(s) only after the task process exits with a nonzero exit code.
Box 4: OutputFiles
To specify output files for a task, create a collection of OutputFile objects and assign it to
the CloudTask.OutputFiles property when you create the task.
References: https://docs.microsoft.com/en-
us/dotnet/api/microsoft.azure.batch.protocol.models.outputfileuploadcondition https://docs.microsoft.com/en-us/azure/batch/batch-task-output-files

**NEW QUESTION 160**
HOTSPOT - (Topic 8)
You are authoring a set of nested Azure Resource Manager templates to deploy Azure resources. You author an Azure Resource Manager template named
mainTemplate.json that contains the following linked templates: linkedTemplate1.json, linkedTemplate2json.
You add parameters to a parameters template file named mainTemplate.parameters.json. You save all templates on a local device in the C:\templates\ folder.
You have the following requirements:
• Store the templates in Azure for later deployment.
• Enable versioning of the templates.
• Manage access to the templates by using Azure RBAC
You need to store the templates in Azure.
How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

**Answer Area**

```
az   ts create                    ▼   \
     ts create
   - storage account create
   - storage account update
   - blueprint artifact template create

--template-file      mainTemplate.json          ▼   " \
"C:\templates\        mainTemplate.json
   --tags Dept=HumanResource: linkedTemplate1.json
                      linkedTemplate2.json
                      mainTemplate.parameters.json
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

```
az   ts create                    ▼   \
     ts create _ _ _
   - storage account create
   - storage account update
   - blueprint artifact template create

--template-file      mainTemplate.json          ▼   " \
"C:\templates\        mainTemplate.json
   --tags Dept=HumanResource: linkedTemplate1.json
                      linkedTemplate2.json
                      mainTemplate.parameters.json
```
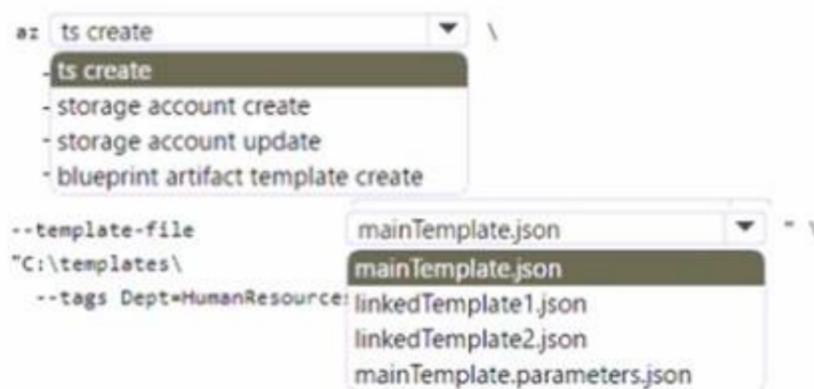
**NEW QUESTION 162**
- (Topic 8)
You are developing a complex workflow by using Azure Durable Functions.
During testing you observe that the results of the workflow differ based on how many instances of the Azure Function are running.
You need to resolve the issue. What should you do?

A. Ensure that all Orchestrator code is deterministic.
B. Read all state data from the durable function context
C. Configure the Azure Our able f unction to run on an App Service Plan with one instance.
D. Implement the monitor pattern within the workflow.

**Answer:** A

**NEW QUESTION 163**
HOTSPOT - (Topic 8)
You are configuring a new development environment for a Java application.
The environment requires a Virtual Machine Scale Set (VMSS), several storage accounts, and networking components.
The VMSS must not be created until the storage accounts have been successfully created and an associated load balancer and virtual network is configured.
How should you complete the Azure Resource Manager template? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Answer Area

```
{
  . . .
  "resources": [
    {                                           pyln
      "apiVersion": "2016-01-01",
      "type": "Microsoft.Storage/storageAccounts",
      "name": "[concat(              (), 'storage', uniqueString(resourceGroup().id))]",
                    copy
                    copyIndex
                    priority
                    dependsOn

      "location": "[resourceGroup().location]",
      . . .
      "sku": {
        "name": "Standard_LRS"
      },
      "kind": "Storage",
      "properties": {},
      "              ": {
            copy
            copyIndex
            priority
            dependsOn

        "name": "storagesetup",
        "count": 3
      }
    },
    {
      "apiVersion": "2015-06-15",
      "type": "Microsoft.Compute/virtualMachines",
      "name": "[concat('VM', uniqueString(resourceGroup().id))]",
      "              ": [
            copy
            copyIndex
            priority
            dependsOn

        "[variables('loadBalancerName')]",
        "[variables('virtualNetworkName')]",
        "storagesetup",
      ],
      . . .
    }
  ],
  "outputs": {}
}
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: copyIndex
Notice that the name of each resource includes the copyIndex() function, which returns the current iteration in the loop. copyIndex() is zero-based.
Box 2: copy
By adding the copy element to the resources section of your template, you can dynamically set the number of resources to deploy.
Box 3: dependsOn Example:
"type": "Microsoft.Compute/virtualMachineScaleSets", "apiVersion": "2020-06-01",
"name": "[variables('namingInfix')]",
"location": "[parameters('location')]", "sku": {
"name": "[parameters('vmSku')]", "tier": "Standard",
"capacity": "[parameters('instanceCount')]"
},
"dependsOn": [
"[resourceId('Microsoft.Network/loadBalancers', variables('loadBalancerName'))]", "[resourceId('Microsoft.Network/virtualNetworks',
variables('virtualNetworkName'))]"
],

**NEW QUESTION 168**
DRAG DROP - (Topic 8)
You are developing an application to store millions of images in Azure blob storage. The images are uploaded to an Azure blob storage container named companyimages contained in an Azure blob storage account named companymedia. The stored images are uploaded with multiple blob index tags across multiple blobs in the container.
You must find all blobs whose tags match a search expression in the container. The search expression must evaluate an index tag named status with a value of final.
You need to construct the GET method request URL
How should you complete the URI? To answer, drag the appropriate parameters to the correct request URI targets. Each parameter 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.

| Parameters | Answer Area |
| --- | --- |
| Status='Final' | https:// [____] .blob.core.windows.net/ [____] ?restype=container&comp=blobs&where= [____] |
| Status<='Final' | |
| companymedia | |
| companyimages | |
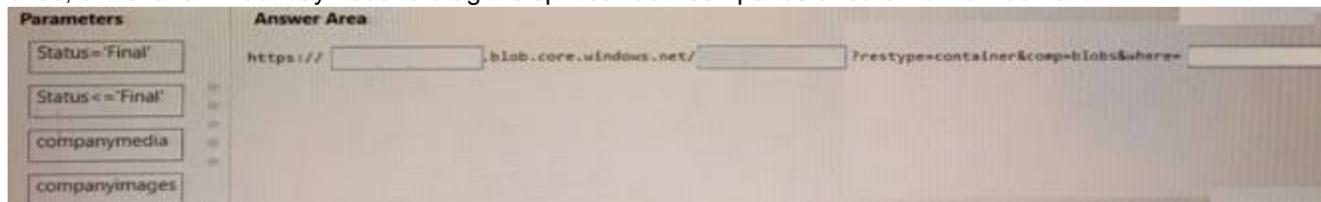
A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
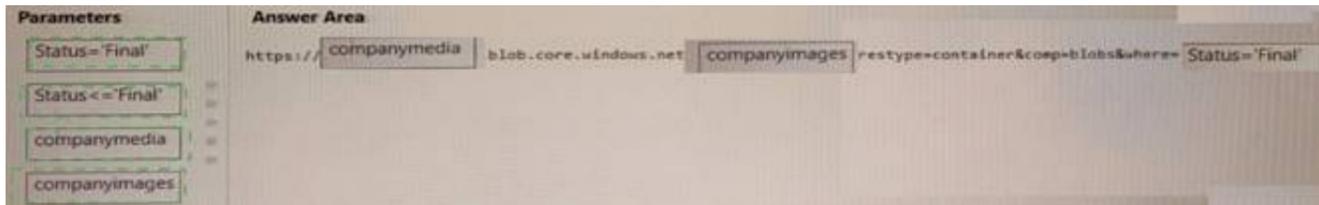
**NEW QUESTION 169**
- (Topic 7)
You need to secure the Azure Functions to meet the security requirements.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Store the RSA-HSM key in Azure Cosmos D
B. Apery the built-in policies for customer-managed keys and allowed locations.
C. Create a free tier Azure App Configuration instance with a new Azure AD service principal.
D. Store the RSA-HSM key in Azure Key Vault with soft-delete and purge-protection features enabled.
E. Store the RSA-HSM key in Azure Blob storage with an Immutability policy applied to the container.
F. Create a standard tier Azure App Configuration instance with an assigned Azure AD managed identity.

**Answer:** CE

**Explanation:**
Scenario: All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Microsoft Azure Key Vault is a cloud-hosted management service that allows users to encrypt keys and small secrets by using keys that are protected by hardware security modules (HSMs).
You need to create a managed identity for your application. Reference:
https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references

**NEW QUESTION 173**
- (Topic 7)
You need to secure the Azure Functions to meet the security requirements.
Which two actions should you perform? Each correct answer presents part of the solution NOTE: Each correct selection is worth one point.

A. Store the RSA-HSM key in Azure Key Vault with soft-delete and purge-protection features enabled
B. Store the RSA-HSM key in Azure Blob storage with an immutability policy applied to the container.
C. Store the RSA-HSM key in Azure Cosmos D
D. Apply the built-in policies for customer- managed Keys and allowed locations
E. Create a standard tier Azure App Configuration instance with an assigned Azure AD managed identity.
F. Create a free tier Azure App Configuration instance with a new Azure AD service principal.

**Answer:** BC

**NEW QUESTION 177**
HOTSPOT - (Topic 7)
You need to implement the corporate website. How should you configure the solution?



A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
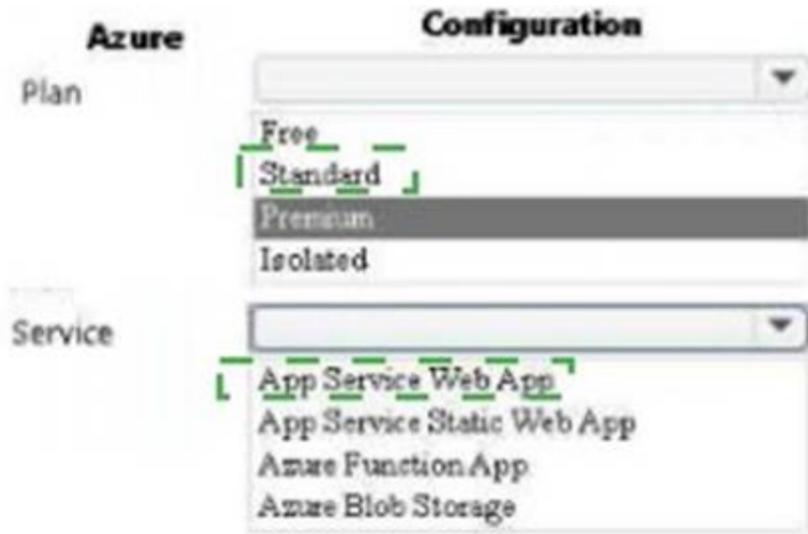
**NEW QUESTION 178**
- (Topic 6)
You need to access data from the user claim object in the e-commerce web app. What should you do first?

A. Write custom code to make a Microsoft Graph API call from the e-commerce web app.
B. Assign the Contributor RBAC role to the e-commerce web app by using the Resource Manager create role assignment API.
C. Update the e-commerce web app to read the HTTP request header values.
D. Using the Azure CLI, enable Cross-origin resource sharing (CORS) from the e- commerce checkout API to the e-commerce web app.

**Answer:** C

**Explanation:**
Methods to Get User Identity and Claims in a .NET Azure Functions App include: ClaimsPrincipal from the Request Context
The ClaimsPrincipal object is also available as part of the request context and can be extracted from the HttpRequest.HttpContext.
User Claims from the Request Headers.
App Service passes user claims to the app by using special request headers. Reference:
https://levelup.gitconnected.com/four-alternative-methods-to-get-user-identity-and-claims-in-a-net-azurefunctions-app-df98c40424bb

**NEW QUESTION 180**
- (Topic 5)
You need to resolve the log capacity issue. What should you do?

A. Create an Application Insights Telemetry Filter
B. Change the minimum log level in the host.json file for the function
C. Implement Application Insights Sampling
D. Set a LogCategoryFilter during startup

**Answer:** C

**Explanation:**
Scenario, the log capacity issue: Developers report that the number of log message in the trace output for the processor is too high, resulting in lost log messages.
Sampling is a feature in Azure Application Insights. It is the recommended way to reduce telemetry traffic and storage, while preserving a statistically correct analysis of application data. The filter selects items that are related, so that you can navigate between items when you are doing diagnostic investigations. When metric counts are presented to you in the portal, they are renormalized to take account of the sampling, to minimize any effect on the statistics.
Sampling reduces traffic and data costs, and helps you avoid throttling. Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling

**NEW QUESTION 181**
DRAG DROP - (Topic 5)
You need to ensure disaster recovery requirements are met. What code should you add at line PC16?
To answer, drag the appropriate code fragments to the correct locations. Each code fragment 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.



A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Scenario: Disaster recovery. Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.
Box 1: DirectoryTransferContext We transfer all files in the directory.
Note: The TransferContext object comes in two forms: SingleTransferContext and DirectoryTransferContext. The former is for transferring a single file and the latter is for transferring a directory of files.
Box 2: ShouldTransferCallbackAsync
The DirectoryTransferContext.ShouldTransferCallbackAsync delegate callback is invoked to tell whether a transfer should be done.
Box 3: False
If you want to use the retry policy in Copy, and want the copy can be resume if break in the middle, you can use SyncCopy (isServiceCopy = false).
Note that if you choose to use service side copy ('isServiceCopy' set to true), Azure (currently) doesn't provide SLA for that. Setting 'isServiceCopy' to false will download the source blob loca

**NEW QUESTION 185**
- (Topic 4)
You need to resolve a notification latency issue.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

A. Set Always On totrue.
B. Ensure that the Azure Function is using an App Service plan.
C. Set Always On tofalse.
D. Ensure that the Azure Function is set to use a consumption plan.

**Answer:** AB

**Explanation:**
Azure Functions can run on either a Consumption Plan or a dedicated App Service Plan. If you run in a dedicated mode, you need to turn on the Always On setting for your Function App to run properly. The Function runtime will go idle after a few minutes of inactivity, so only HTTP triggers will actually "wake up" your functions. This is similar to how WebJobs must have Always On enabled.
Scenario: Notification latency: Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.
Anomaly detection service: You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.
Reference:
https://github.com/Azure/Azure-Functions/wiki/Enable-Always-On-when-running-on- dedicated-App-Service-Plan

**NEW QUESTION 188**
DRAG DROP - (Topic 4)
You need to implement the Log policy.
How should you complete the Azure Event Grid subscription? To answer, drag the appropriate JSON segments to the correct locations. Each JSON segment may be used once, more than once, or not at all. You may need to drag the split bar between panes to view content.
NOTE:Each correct selection is worth one point.

| Code segment |
| --- |
| All |
| WebHook |
| EventHub |
| subjectEndsWith |
| Mictosoft.Storage |
| subjectBeginsWith |
| Microsoft.Storage.BlobCreated |

**Answer Area**

```
{
  "name": "newlogs",
  "properties": {
    "topic": "/subscriptions/. . ./providers/Microsoft.EventGrid/topics/. . .",
    "destination": {
      "endpointType" : "        code segment        " },
    "filter": {
              code segment        ": "/blobServices/default/containers/logdrop/",
      "includedEventTypes": [ "        code segment        " ] },
    },
    "labels": [],
    "eventDeliverySchema": "EventGridSchema"
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1:WebHook
Scenario: If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.
endpointType: The type of endpoint for the subscription (webhook/HTTP, Event Hub, or queue).
Box 2: SubjectBeginsWith
Box 3: Microsoft.Storage.BlobCreated Scenario: Log Policy
All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.
Example subscription schema
{
"properties": { "destination": {
"endpointType": "webhook", "properties": { "endpointUrl":
"https://example.azurewebsites.net/api/HttpTriggerCSharp1?code=VXbGWce53l48Mt8wuo tr0GPmyJ/nDT4hgdFj9DpBiRt38qqnnm5OFg=="
}
},
"filter": {
"includedEventTypes": [ "Microsoft.Storage.BlobCreated", "Microsoft.Storage.BlobDeleted"
```

],
"subjectBeginsWith": "blobServices/default/containers/mycontainer/log", "subjectEndsWith": ".jpg",
"isSubjectCaseSensitive ": "true"
}
}
}

**NEW QUESTION 192**
- (Topic 3)
You need to correct the RequestUserApproval Function app error. What should you do?

A. Update line RA13 to use theasynckeyword and return an HttpRequest object value.
B. Configure the Function app to use an App Service hosting pla
C. Enable theAlways Onsetting of the hosting plan.
D. Update the function to be stateful by using Durable Functions to process the request payload.
E. Update the functionTimeout property of the host.json project file to 15 minutes.

**Answer:** C

**Explanation:**
Async operation tracking
The HTTP response mentioned previously is designed to help implement long-running HTTP async APIs with Durable Functions. This pattern is sometimes referred to as the polling consumer pattern.
Both the client and server implementations of this pattern are built into the Durable Functions HTTP APIs.
Function app
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:
FunctionAppLogs
| where FunctionName = = "RequestUserApproval"
References:
https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-http-features

**NEW QUESTION 193**
- (Topic 3)
You need to authenticate the user to the corporate website as indicated by the architectural diagram.
Which two values should you use? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

A. ID token signature
B. ID token claims
C. HTTP response code
D. Azure AD endpoint URI
E. Azure AD tenant ID

**Answer:** AD

**Explanation:**
Claims in access tokens
JWTs (JSON Web Tokens) are split into three pieces:
? Header - Provides information about how to validate the token including information about the type of token and how it was signed.
? Payload - Contains all of the important data about the user or app that is attempting to call your service.
? Signature - Is the raw material used to validate the token.
Your client can get an access token from either the v1.0 endpoint or the v2.0 endpoint using a variety of protocols.
Scenario: User authentication (see step 5 below)
The following steps detail the user authentication process:
? The user selects Sign in in the website.
? The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
? The user signs in.
? Azure AD redirects the user's session back to the web application. The URL includes an access token.
? The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
? The back-end API validates the access token.
Reference:
https://docs.microsoft.com/en-us/azure/api-management/api-management-access- restriction-policies

**NEW QUESTION 196**
HOTSPOT - (Topic 3)
You need to retrieve the database connection string.
Which values should you use? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

REST API Endpoint:

https:// [▼] .vault.azure.net/secrets/ [▼] /

| cpandlkeyvault |
| PostgreSQLConn |
| 80df3e46ffcd4f1cb187f79905e9a1e8 |

| cpandlkeyvault |
| PostgreSQLConn |
| 80df3e46ffcd4f1cb187f79905e9a1e8 |

Variable type to access Azure Key Vault secret values:

[▼]

| Environment |
| Session |
| ViewState |
| Querystring |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Azure database connection string retrieve REST API vault.azure.net/secrets/ Box 1: cpandlkeyvault
We specify the key vault, cpandlkeyvault.
Scenario: The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
Box 2: PostgreSQLConn
We specify the secret, PostgreSQLConn
Example, sample request: https://myvault.vault.azure.net//secrets/mysecretname/4387e9f3d6e14c459867679a90fd0f 79?api-version=7.1
Box 3: Querystring

**NEW QUESTION 198**
HOTSPOT - (Topic 3)
You need to configure the Account Kind, Replication, and Storage tier options for the corporate website's Azure Storage account.
How should you complete the configuration? To answer, select the appropriate options in the dialog box in the answer area.
NOTE:Each correct selection is worth one point.

## Create storage account                                                    ✕

Basics    Advanced    Tags    Review + create

Azure Storage is a Microsoft-managed service providing cloud storage that is highly available, secure, durable, scalable, and redundant. Azure Storage includes Azure Blobs (objects), Azure Data Lake Storage Gen2, Azure Files, Azure Queues, and Azure Tables. The cost of your storage account depends on the usage and the options you choose below.  Learn more

### PROJECT DETAILS

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

* Subscription         | Visual Studio Enterprise            ⌄ |

    * Resource group   | (New) cplcorporatesite              ⌄ |
                         Create new

### INSTANCE DETAILS

The default deployment model is Resource Manager, which supports the latest Azure features. You may choose to deploy using the classic deployment model instead.   Choose classic deployment model

* Storage account name ⓘ   | corporatewebsitecontent            ✓ |

* Location                  | (US) East US                       ⌄ |

Performance ⓘ              ⦿ Standard    ◯ Premium

Account kind ⓘ             |                                    ⌄ |
                            | StorageV2 (general purpose v2) |
                            | Storage (general purpose v1) |
                            | BlobStorage |

Replication ⓘ              |                                    ⌄ |
                            | Locally-redundant storage (LRS) |
                            | Zone-redundant storage (ZRS) |
                            | Geo-redundant storage (GRS) |
                            | Read-access geo-redundant storage (RA-GRS) |
                            | Geo-zone-redundant storage (GZRS) |
                            | Read-access geo-zone-redundant storage (RA-GZRS) |

Access tier (default) ⓘ    ◯ Cool    ◯ Hot

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Account Kind: StorageV2 (general-purpose v2)
Scenario: Azure Storage blob will be used (refer to the exhibit). Data storage costs must be minimized.
General-purpose v2 accounts: Basic storage account type for blobs, files, queues, and tables. Recommended for most scenarios using Azure Storage.

**NEW QUESTION 203**
HOTSPOT - (Topic 2)
You need to add code at line AM09 to ensure that users can review content using ContentAnalysisService.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

| ▼ | , |
|---|---|
| "allowPublicClient":true |
| "oauth2Permissions": ["login"] |
| "oauth2AllowUrlPathMatching":true |
| "oauth2AllowIdTokenImplicitFlow":true |

| ▼ | , |
|---|---|
| "oauth2AllowImplicitFlow": true |
| "oauth2RequiredPostResponse":true |
| "preAuthorizedApplications":["SPA"] |
| "knownClientApplications":["ContentAnalysisService"] |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: "oauth2Permissions": ["login"]
oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.
Box 2: "oauth2AllowImplicitFlow":true
For applications (Angular, Ember.js, React.js, and so on), Microsoft identity platform supports the OAuth 2.0 Implicit Grant flow.

**NEW QUESTION 204**
- (Topic 2)
You need to configure the ContentUploadService deployment.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Add the following markup to line CS23: types: Private
B. Add the following markup to line CS24: osType: Windows
C. Add the following markup to line CS24: osType: Linux
D. Add the following markup to line CS23: types: Public

**Answer:** C

**Explanation:**
Scenario: All Internal services must only be accessible from Internal Virtual Networks (VNets)
There are three Network Location types – Private, Public and Domain Reference:
https://devblogs.microsoft.com/powershell/setting-network-location-to-private/

**NEW QUESTION 209**
- (Topic 2)
You need to monitor ContentUploadService accourding to the requirements. Which command should you use?

A. az monitor metrics alert create –n alert –g … - -scopes … - -condition "avg Percentage CPU > 8"
B. az monitor metrics alert create –n alert –g … - -scopes … - -condition "avg Percentage CPU > 800"
C. az monitor metrics alert create –n alert –g … - -scopes … - -condition "CPU Usage > 800"
D. az monitor metrics alert create –n alert –g … - -scopes … - -condition "CPU Usage > 8"

**Answer:** B

**Explanation:**
Scenario: An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU-cores

Reference:
https://docs.microsoft.com/sv-se/cli/azure/monitor/metrics/alert


**NEW QUESTION 213**
DRAG DROP - (Topic 2)
You need to add markup at line AM04 to implement the ContentReview role.
How should you complete the markup? To answer, drag the appropriate json segments to the correct locations. Each json segment 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.

**Json segments**

```
User
value
role
Application
allowedMemberTypes
allowedAccountTypes
```

**Answer Area**

```
"appRoles" : [
{
    "              ": [
    "              "
    ],
    "displayName": "ContentReviewer",
    "id": "e1c2ade8-98f8-45fd-aa4a-6d24b512c22a",
    "isEnabled" : true,
    "              " :"ContentReviewer"
}
],
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: allowedMemberTypes
allowedMemberTypes specifies whether this app role definition can be assigned to users and groups by setting to "User", or to other applications (that are accessing this application in daemon service scenarios) by setting to "Application", or to both.
Note: The following example shows the appRoles that you can assign to users. "appId": "8763f1c4-f988-489c-a51e-158e9ef97d6a",
"appRoles": [
{
"allowedMemberTypes": [ "User"
],
"displayName": "Writer",
"id": "d1c2ade8-98f8-45fd-aa4a-6d06b947c66f", "isEnabled": true,
"description": "Writers Have the ability to create tasks.", "value": "Writer"
}
],
"availableToOtherTenants": false,
Box 2: User
Scenario: In order to review content a user must be part of a ContentReviewer role.
Box 3: value
value specifies the value which will be included in the roles claim in authentication and access tokens.


**NEW QUESTION 214**
HOTSPOT - (Topic 2)
You need to ensure that validation testing is triggered per the requirements.
How should you complete the code segment? To answer, select the appropriate values in the answer area.
NOTE:Each correct selection is worth one point.

```
var event = getEvent();
if (event.eventType === '         ▼  '
                    ┌─────────────────────┐
                    │ ImagePushed         │
                    │ RepositoryItem      │
                    │ ImageDeployed       │
                    │ RepositoryUpdated   │
                    └─────────────────────┘

    && event.data.target.        ▼   === 'contentanalysisservice'
                    ┌─────────────┐
                    │ aci         │
                    │ image       │
                    │ service     │
                    │ repository  │
                    └─────────────┘

    && event.        ▼   .contains('contosoimages'))
            ┌─────────────────┐
            │ topic           │
            │ service         │
            │ repository      │
            │ imageCollection │
            └─────────────────┘
{
    startValidationTesting();
}
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: RepositoryUpdated
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.
Box 2: service
Box 3: imageCollection


**NEW QUESTION 219**
HOTSPOT - (Topic 1)
You need to resolve the Shipping web site error.
How should you configre the Azure Table Storage service? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

```
<?xml version="1.0" encoding="utf-8"?>
<StorageServiceProperties>
    ...
    <Cors>
    <CorsRule>
    <        ▼  >  http://*.wideworldimporters.com        ▼  </        ▼  >
    ┌────────────────┐  http://test.wideworldimporters.com      ┌────────────────┐
    │ AllowedHeaders │  http://test-shippingapi.wideworldimporters.com │ AllowedHeaders │
    │ ExposedHeaders │  http://www.wideworldimporters.com        │ ExposedHeaders │
    │ AllowedMethods │                                           │ AllowedMethods │
    │ AllowedOrigins │                                           │ AllowedOrigins │
    └────────────────┘                                           └────────────────┘
    <AllowedMethods>        ▼   </AllowedMethods>
                  ┌──────────┐
                  │ GET,PUT  │
                  │ GET      │
                  │ POST     │
                  │ GET,HEAD │
                  └──────────┘
    ...
    </CorsRule>
    </Cors>
</StorageServiceProperties>
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Box 1: AllowedOrigins
A CORS request will fail if Access-Control-Allow-Origin is missing.
Scenario:
The following error message displays while you are testing the website:

Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin' header is present on the requested resource. Origin 'http://testwideworldimporters.com/' is therefore not allowed access.

Box 2: http://test-shippingapi.wideworldimporters.com Syntax: Access-Control-Allow-Origin: *
Access-Control-Allow-Origin: <origin> Access-Control-Allow-Origin: null
<origin> Specifies an origin. Only a single origin can be specified. Box 3: AllowedOrigins
Box 4: POST
The only allowed methods are GET, HEAD, and POST. In this case POST is used. "<Corsrule>" "allowedmethods" Failed to load no "Access-control-Origin" header is present
References:
https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Access-Control-Allow-Origin


**NEW QUESTION 222**
HOTSPOT - (Topic 1)
You need to correct the VM issues.
Which tools should you use? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

| Issue | Tool |
|---|---|
| Backup and Restore | ▼ |
| | Azure Site Recovery |
| | Azure Backup |
| | Azure Data Box |
| | Azure Migrate |
| Performance | ▼ |
| | Azure Network Watcher |
| | Azure Traffic Manager |
| | ExpressRoute |
| | Accelerated Networking |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Backup and Restore: Azure Backup
Scenario: The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.
In-Place restore of disks in IaaS VMs is a feature of Azure Backup. Performance: Accelerated Networking
Scenario: The VM shows high network latency, jitter, and high CPU utilization.
Accelerated networking enables single root I/O virtualization (SR-IOV) to a VM, greatly improving its networking performance. This high-performance path bypasses the host from the datapath, reducing latency, jitter, and CPU utilization, for use with the most demanding network workloads on supported VM types.
References:
https://azure.microsoft.com/en-us/blog/an-easy-way-to-bring-back-your-azure-vm-with-in- place-restore/


**NEW QUESTION 224**
DRAG DROP - (Topic 1)
You need to support the message processing for the ocean transport workflow.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

**Actions**

Create an integration account in the Azure portal.

Link the custom connector to the Logic App.

Update the Logic App to use the partners, schemas, certificates, maps, and agreements.

Create a custom connector for the Logic App.

Add partners, schemas, certificates, maps, and agreements.

Link the Logic App to the integration account.

**Answer Area**

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Step 1: Create an integration account in the Azure portal
You can define custom metadata for artifacts in integration accounts and get that metadata during runtime for your logic app to use. For example, you can provide metadata for artifacts, such as partners, agreements, schemas, and maps - all store metadata using
key-value pairs.
Step 2: Link the Logic App to the integration account
A logic app that's linked to the integration account and artifact metadata you want to use. Step 3: Add partners, schemas, certificates, maps, and agreements
Step 4: Create a custom connector for the Logic App.
References:
https://docs.microsoft.com/bs-latn-ba/azure/logic-apps/logic-apps-enterprise-integration- metadata


**NEW QUESTION 229**
......

# THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AZ-204 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AZ-204 Product From:

## https://www.2passeasy.com/dumps/AZ-204/

# Money Back Guarantee

## AZ-204 Practice Exam Features:

* AZ-204 Questions and Answers Updated Frequently

* AZ-204 Practice Questions Verified by Expert Senior Certified Staff

* AZ-204 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AZ-204 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year