

# Microsoft

## Exam Questions AZ-204

Developing Solutions for Microsoft Azure (beta)



**NEW QUESTION 1**

- (Exam Topic 1)

You need to secure the Shipping Logic App. What should you use?

- A. Azure App Service Environment (ASE)
- B. Azure AD B2B integration
- C. Integration Service Environment (ISE)
- D. VNet service endpoint

**Answer: C**

**Explanation:**

Scenario: The Shipping Logic App requires secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.

You can access to Azure Virtual Network resources from Azure Logic Apps by using integration service environments (ISEs).

Sometimes, your logic apps and integration accounts need access to secured resources, such as virtual machines (VMs) and other systems or services, that are inside an Azure virtual network. To set up this access, you can create an integration service environment (ISE) where you can run your logic apps and create your integration accounts.

References:

<https://docs.microsoft.com/en-us/azure/logic-apps/connect-virtual-network-vnet-isolated-environment-overview>

**NEW QUESTION 2**

- (Exam Topic 1)

You need to resolve the Shipping web site error.

How should you configure 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>
      <
        AllowedHeaders
        ExposedHeaders
        AllowedMethods
        AllowedOrigins
      >
      http://*.wideworldimporters.com
      http://test.wideworldimporters.com
      http://test-shippingapi.wideworldimporters.com
      http://www.wideworldimporters.com
    </
    AllowedHeaders
    ExposedHeaders
    AllowedMethods
    AllowedOrigins
  >
  <AllowedMethods>
    GET,PUT
    GET
    POST
    GET,HEAD
  </AllowedMethods>
    </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 3**

- (Exam Topic 1)

You need to migrate on-premises shipping data to Azure. What should you use?

- A. Azure Migrate
- B. Azure Cosmos DB Data Migration tool (dt.exe)
- C. AzCopy
- D. Azure Database Migration service

**Answer:** D

**Explanation:**

Migrate from on-premises or cloud implementations of MongoDB to Azure Cosmos DB with minimal downtime by using Azure Database Migration Service. Perform resilient migrations of MongoDB data at scale and with high reliability.

Scenario: Data migration from on-premises to Azure must minimize costs and downtime.

The application uses MongoDB JSON document storage database for all container and transport information. References:

<https://azure.microsoft.com/en-us/updates/mongodb-to-azure-cosmos-db-online-and-offline-migrations-are-now>

**NEW QUESTION 4**

- (Exam Topic 1)

You need to configure Azure CDN for the Shipping web site.

Which configuration options should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Answer Area**

Option	Value
Tier	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>Standard</p> <p>Premium</p> </div> </div>
Profile	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>Akamai</p> <p>Microsoft</p> </div> </div>
Optimization	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>general web delivery</p> <p>large file download</p> <p>dynamic site acceleration</p> <p>video-on-demand media streaming</p> </div> </div>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Scenario: Shipping website

Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.

Tier: Standard Profile: Akamai

Optimization: Dynamic site acceleration

Dynamic site acceleration (DSA) is available for Azure CDN Standard from Akamai, Azure CDN Standard from Verizon, and Azure CDN Premium from Verizon profiles.

DSA includes various techniques that benefit the latency and performance of dynamic content. Techniques include route and network optimization, TCP optimization, and more.

You can use this optimization to accelerate a web app that includes numerous responses that aren't cacheable. Examples are search results, checkout transactions, or real-time data. You can continue to use core Azure CDN caching capabilities for static data.

Reference:

<https://docs.microsoft.com/en-us/azure/cdn/cdn-optimization-overview>

**NEW QUESTION 5**

- (Exam Topic 1)

You need to configure Azure App Service to support the REST API requirements.

Which values should you use? To answer, select the appropriate options in the answer area.  
 NOTE: Each correct selection is worth one point.

Setting	Value
Plan	<ul style="list-style-type: none"> <li>Basic</li> <li>Standard</li> <li>Premium</li> <li>Isolated</li> </ul>
Instance Count	<ul style="list-style-type: none"> <li>1</li> <li>10</li> <li>20</li> <li>100</li> </ul>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Plan: Standard

Standard support auto-scaling Instance Count: 10

Max instances for standard is 10. Scenario:

The REST API's that support the solution must meet the following requirements:

- > Allow deployment to a testing location within Azure while not incurring additional costs.
- > Automatically scale to double capacity during peak shipping times while not causing application downtime.
- > Minimize costs when selecting an Azure payment model. References:

<https://azure.microsoft.com/en-us/pricing/details/app-service/plans/>

**NEW QUESTION 6**

- (Exam Topic 1)

You need to update the APIs to resolve the testing error.

How should you complete the Azure CLI command? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
az webapp   -g shipping-apis-test-rg -n web
```

cors
config
deployment

add
up
remove

```
--  
```

slot
allowed-origins
name

http://*.wideworldimporters.com
http://test-shippingapi.wideworldimporters.com
http://test.wideworldimporters.com
http://www.wideworldimporters.com

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Enable Cross-Origin Resource Sharing (CORS) on your Azure App Service Web App.  
 Enter the full URL of the site you want to allow to access your WEB API or \* to allow all domains. Box 1: cors  
 Box 2: add  
 Box 3: allowed-origins  
 Box 4: http://testwideworldimporters.com/ References:  
<http://donovanbrown.com/post/How-to-clear-No-Access-Control-Allow-Origin-header-error-with-Azure-App-Service>

**NEW QUESTION 7**

- (Exam 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	<div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; height: 20px; display: flex; justify-content: flex-end; align-items: center; padding-right: 5px;">▼</div> <div style="padding: 2px;"> <p>Azure Site Recovery</p> <p>Azure Backup</p> <p>Azure Data Box</p> <p>Azure Migrate</p> </div> </div>
Performance	<div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; height: 20px; display: flex; justify-content: flex-end; align-items: center; padding-right: 5px;">▼</div> <div style="padding: 2px;"> <p>Azure Network Watcher</p> <p>Azure Traffic Manager</p> <p>ExpressRoute</p> <p>Accelerated Networking</p> </div> </div>

- 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 8**

- (Exam Topic 3)  
 You must ensure that the external party cannot access the data in the SSN column of the Person table.  
 Will each protection method meet the requirement? To answer, drag the appropriate responses to the correct protection methods. Each response 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.

Responses	Protection method	Response
<input type="checkbox"/> Yes	Enable AlwaysOn encryption.	<input type="checkbox"/>
<input type="checkbox"/> No	Set the column encryption setting to disabled.	<input type="checkbox"/>
	Assign users to the Public fixed database role.	<input type="checkbox"/>
	Store column encryption keys in the system catalog view in the database.	<input type="checkbox"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Yes

You can configure Always Encrypted for individual database columns containing your sensitive data. When setting up encryption for a column, you specify the information about the encryption algorithm and cryptographic keys used to protect the data in the column.

Box 2: No

Box 3: Yes

In SQL Database, the VIEW permissions are not granted by default to the public fixed database role. This enables certain existing, legacy tools (using older versions of DacFx) to work properly. Consequently, to work with encrypted columns (even if not decrypting them) a database administrator must explicitly grant the two VIEW permissions.

Box 4: No

All cryptographic keys are stored in an Azure Key Vault. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

**NEW QUESTION 9**

- (Exam Topic 3)

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 SearchIndexClient 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. Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource. 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 10**

- (Exam Topic 3)

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: Add the web applications to Docker containers. Deploy the containers. Deploy the containers to Azure Kubernetes Service (AKS).

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Instead use Azure Cache for Redis.

Note: Azure Cache for Redis provides a session state provider that you can use to store your session state in-memory with Azure Cache for Redis instead of a SQL Server database. To use the caching session state provider, first configure your cache, and then configure your ASP.NET application for cache using the Azure Cache for Redis Session State NuGet package.

References:

<https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-aspnet-session-state-provider>

**NEW QUESTION 10**

- (Exam Topic 3)

You develop a web app that uses tier D1 app service plan by using the Web Apps feature of Microsoft Azure App Service.

Spikes in traffic have caused increases in page load times.

You need to ensure that the web app automatically scales when CPU load is about 85 percent and minimize costs.

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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

**Actions**

**Answer Area**

Configure the web app to the Premium App Service tier.

Configure the web app to the Standard App Service tier.

Enable autoscaling on the web-app.

Add a Scale rule.

Switch to an Azure App Services consumption plan.

Configure a Scale condition.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Configure the web app to the Standard App Service Tier  
 The Standard tier supports auto-scaling, and we should minimize the cost. Step 2: Enable autoscaling on the web app  
 First enable autoscale Step 3: Add a scale rule  
 Step 4: Add a Scale condition Reference:  
<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-autoscale-get-started>

**NEW QUESTION 13**

- (Exam Topic 3)

You develop and deploy a Java RESTful API to Azure App Service. You open a browser and navigate to the URL for the API. You receive the following error message:

```
Failed to load http://api.azurewebsites.net:6000/#/api/Products: No 'Access-Control-Allow-Origin' header is present on the requested resource. Origin 'http://localhost:6000' is therefore not allowed access
```

You need to resolve the error. What should you do?

- A. Bind an SSL certificate
- B. Enable authentication
- C. Enable CORS
- D. Map a custom domain
- E. Add a CDN

**Answer:** C

**Explanation:**

We need to enable Cross-Origin Resource Sharing (CORS). References:  
<https://medium.com/@xinganwang/a-practical-guide-to-cors-51e8fd329a1f>

**NEW QUESTION 16**

- (Exam Topic 3)

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 21**

- (Exam Topic 3)

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 Notification Hub. Register all devices with the hub. 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.

Reference:

<https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services>

**NEW QUESTION 22**

- (Exam Topic 3)

You are developing an Azure Cosmos DB solution by using the Azure Cosmos DB SQL API. The data includes millions of documents. Each document may contain hundreds of properties.

The properties of the documents do not contain distinct values for partitioning. Azure Cosmos DB must scale individual containers in the database to meet the performance needs of the application by spreading the workload evenly across all partitions over time.

You need to select a partition key.

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

- A. a concatenation of multiple property values with a random suffix appended
- B. a single property value that does not appear frequently in the documents
- C. a hash suffix appended to a property value
- D. a value containing the collection name
- E. a single property value that appears frequently in the documents

**Answer:** AC

**Explanation:**

You can form a partition key by concatenating multiple property values into a single artificial partitionKey property. These keys are referred to as synthetic keys. Another possible strategy to distribute the workload more evenly is to append a random number at the end of the partition key value. When you distribute items in this way, you can perform parallel write operations across partitions.

Note: It's the best practice to have a partition key with many distinct values, such as hundreds or thousands. The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key.

References:

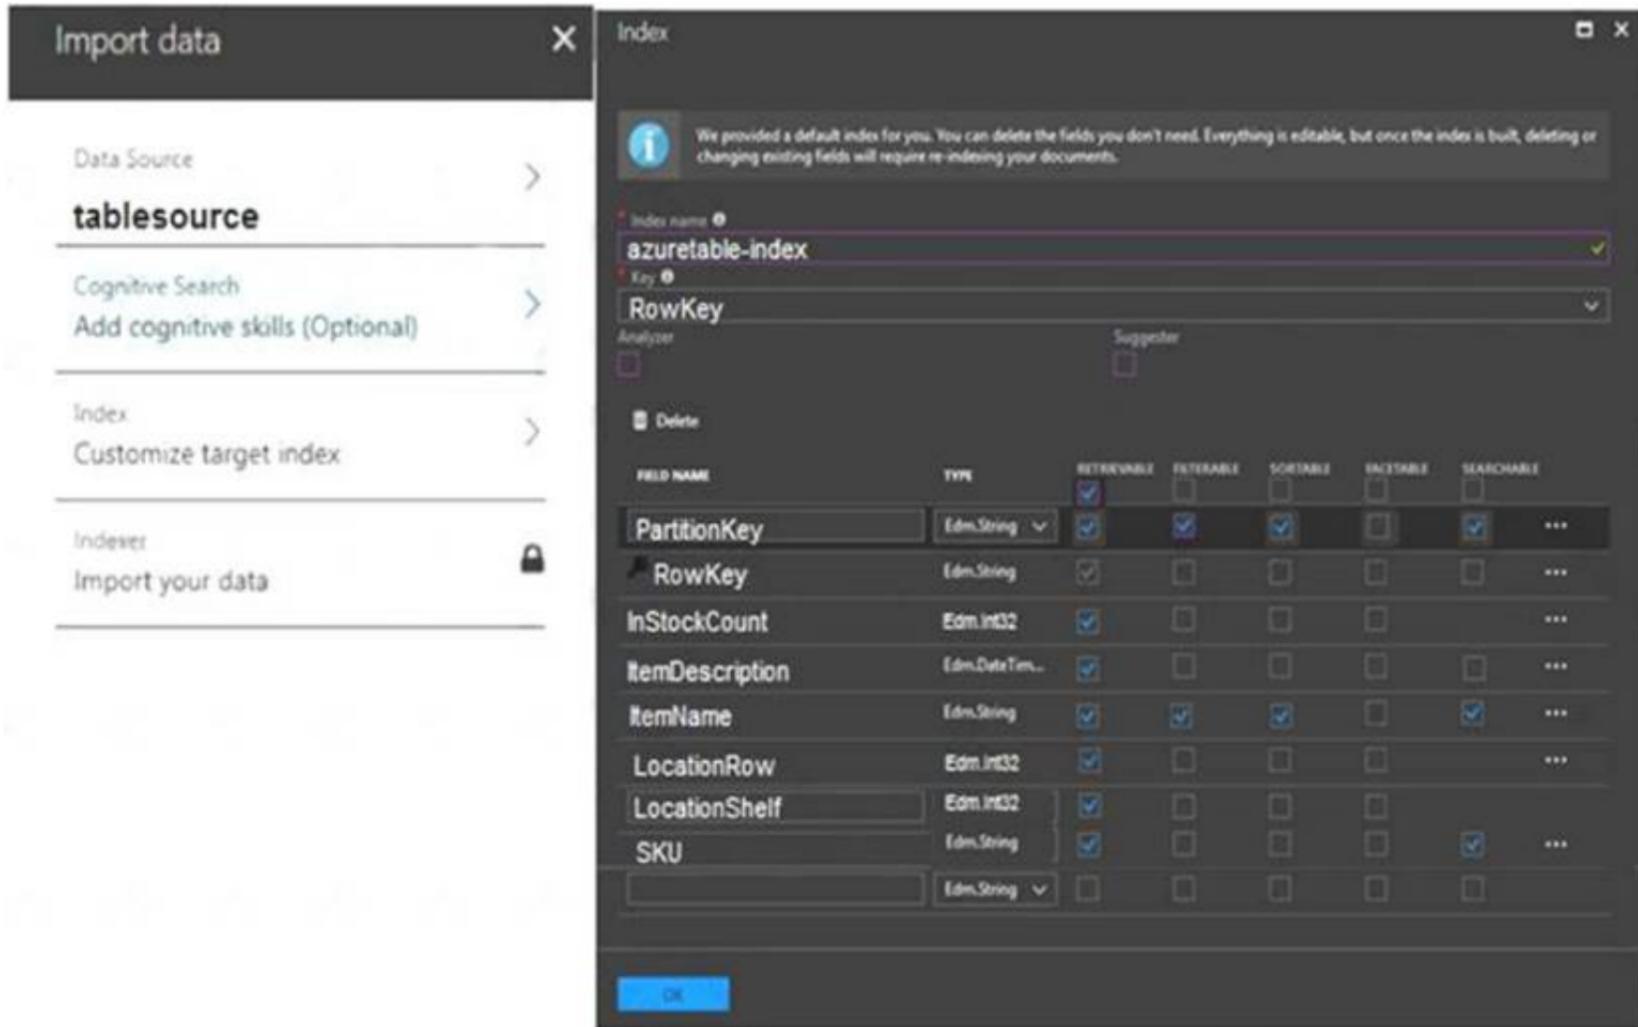
<https://docs.microsoft.com/en-us/azure/cosmos-db/synthetic-partition-keys>

**NEW QUESTION 26**

- (Exam Topic 3)

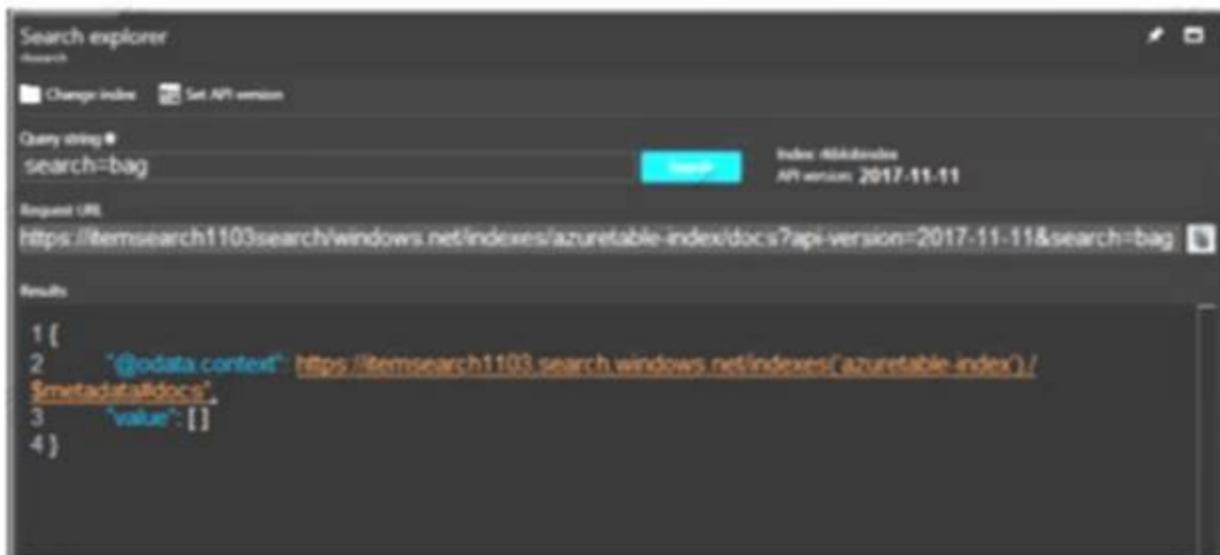
You are validating the configuration of an Azure Search indexer.

The service has been configured with an indexer that uses the Import Data option. The index is configured using options as shown in the Index Configuration exhibit. (Click the Index Configuration tab.)



You use an Azure table as the data source for the import operation. The table contains three records with item inventory data that matches the fields in the Storage data exhibit. These records were imported when the index was created. (Click the Storage Data tab.) When users search with no filter, all three records are displayed.

PartitionKey	RowKey	Timestamp	InStockCount	ItemDescription	ItemName	LocationRow	LocationShelf	SKU
Food	3	2018-08-25T 15:47:29.135Z	32	A box of chocolate candy bars	Choco-bar	5	3	123421
Hardware	2	2018-08-25T 15:46:08.402Z	2	A bag of bolts	Bolts	1	4	67894
Hardware	1	2018-08-25T 15:46:41.402Z	23	A box of nails	Nails	2	1	654365



When users search for items by description, Search explorer returns no records. The Search Explorer exhibit shows the query and results for a test. In the test, a user is trying to search for all items in the table that have a description that contains the word bag. (Click the Search Explorer tab.)

You need to resolve the issue.

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

- |  | Yes                   | No                    |
|--|-----------------------|-----------------------|
| You can resolve the issue by recreating the search index with the same settings for all fields except ItemDescription. Select the SEARCHABLE option for this field | <input type="radio"/> | <input type="radio"/> |
| You can resolve the issue by selecting the index, editing the ItemDescription field, and selecting the SEARCHABLE option for the field.                            | <input type="radio"/> | <input type="radio"/> |
| You can resolve the issue by running the indexer.  | <input type="radio"/> | <input type="radio"/> |
| You can resolve the issue by changing the query string in Search explorer to <code>bag of</code> to return the correct results                                     | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Yes

The ItemDescription field in not searchable. Box 2: No

The ItemDescription field in not searchable, but we would need to recreate the index. Box 3: Yes

An indexer in Azure Search is a crawler that extracts searchable data and metadata from an external Azure data source and populates an index based on field-to-field mappings between the index and your data source. This approach is sometimes referred to as a 'pull model' because the service pulls data in without you having to write any code that adds data to an index.

Box 4: No References:

<https://docs.microsoft.com/en-us/azure/search/search-what-is-an-index> <https://docs.microsoft.com/en-us/azure/search/search-indexer-overview>

**NEW QUESTION 28**

- (Exam Topic 3)

You develop a news and blog content delivery app for Windows devices.

A notification must arrive on a user's device when there is a new article available for them to view. You need to implement push notifications.

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**

```
string notificationHubName = "contoso_hub";
string notificationHubConnection = "connection_string";
NotificationHubClient hub =
NotificationHubClientSettings
NotificationHubJob
NotificationDetails
NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails
NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails
(notificationHubConnection, notificationHubName);
string windowsToastPayload =
@"<toast><visual><binding template=""ToastText01""><text id=""1"">"+
@"New item to view" + @"</text></binding></visual></toast>";
try
{
var result=
await hub.
SendWindowsNativeNotificationAsync
SubmitNotificationHubJobAsync
ScheduleNotificationAsync
SendAppleNativeNotificationAsync
(notificationHubConnection, notificationHubName);
}
catch (System.Exception ex)
{
}
}
```

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: NotificationHubClient

Box 2: NotificationHubClient

Box 3: CreateClientFromConnectionString

// Initialize the Notification Hub

NotificationHubClient hub = NotificationHubClient.CreateClientFromConnectionString(listenConnString, hubName);

Box 4: SendWindowsNativeNotificationAsync Send the push notification.

var result = await hub.SendWindowsNativeNotificationAsync(windowsToastPayload);

References: <https://docs.microsoft.com/en-us/azure/notification-hubs/notification-hubs-push-notification-registration-manag>

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/app-service-mobile/app-service-mobile-windo>

**NEW QUESTION 33**

- (Exam Topic 3)

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",
  [dropdown] : "All",
  [box 1]
  "optionalClaims"
  "groupMembershipClaims"
  [dropdown] : true
  [box 2]
  "allowPublicClient"
  "oauth2Permissions"
  "requiredResourceAccess"
  "oauth2AllowImplicitFlow"
  ...
}

```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: groupMembershipClaims

Scenario: 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"

Box 2: oauth2Permissions

Scenario: Azure AD users must be able to login to the website.

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.

**NEW QUESTION 36**

- (Exam Topic 3)

You are a developer for a software as a service (SaaS) company that uses an Azure Function to process orders. The Azure Function currently runs on an Azure Function app that is triggered by an Azure Storage queue.

You are preparing to migrate the Azure Function to Kubernetes using Kubernetes-based Event Driven Autoscaling (KEDA).

You need to configure Kubernetes Custom Resource Definitions (CRD) for the Azure Function.

Which CRDs should you configure? To answer, drag the appropriate CRD types to the correct locations. Each CRD 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.

**Answer Area**

CRD types	Setting	CRD type
Secret	Azure Function code	<input type="text"/>
Deployment	Polling interval	<input type="text"/>
ScaledObject	Azure Storage connection string	<input type="text"/>
TriggerAuthentication		

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Box 1: Deployment**

To deploy Azure Functions to Kubernetes use the func kubernetes deploy command has several attributes that directly control how our app scales, once it is deployed to Kubernetes.

**Box 2: ScaledObject**

With --polling-interval, we can control the interval used by KEDA to check Azure Service Bus Queue for messages.

Example of ScaledObject with polling interval apiVersion: keda.k8s.io/v1alpha1

kind: ScaledObject metadata:

name: transformer-fn namespace: tt

labels:

deploymentName: transformer-fn spec:

scaleTargetRef: deploymentName: transformer-fn pollingInterval: 5

minReplicaCount: 0

maxReplicaCount: 100

**Box 3: Secret**

Store connection strings in Kubernetes Secrets. Example: to create the Secret in our demo Namespace:

```
# create the k8s demo namespace kubectl create namespace tt
```

```
# grab connection string from Azure Service Bus KEDA_SCALER_CONNECTION_STRING=$(az servicebus queue authorization-rule keys list \
```

```
-g $RG_NAME \
```

```
--namespace-name $SBN_NAME \
```

```
--queue-name inbound \
```

```
-n keda-scaler \
```

```
--query "primaryConnectionString" \
```

```
-o tsv)
```

```
# create the kubernetes secret
```

```
kubectl create secret generic tt-keda-auth \
```

```
--from-literal KedaScaler=$KEDA_SCALER_CONNECTION_STRING \
```

```
--namespace tt Reference:
```

<https://www.thinktecture.com/en/kubernetes/serverless-workloads-with-keda/>

**NEW QUESTION 37**

- (Exam Topic 3)

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 an Azure Database for PostgreSQL. Update the web applications. Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**NEW QUESTION 41**

- (Exam Topic 3)

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.  ();

        job.Id = jobId,
        job.PoolInformation = 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.  ));

            outputFileList.Add(new OutputFile(fileTask.Output,
                new OutputFileDestination(failedContainer),
                new OutputFileUploadOptions(OutputFileUploadCondition,  ));

            task.  =outputFileList;

            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 43**

- (Exam Topic 3)

Your company is developing an Azure API.

You need to implement authentication for the Azure API. You have the following requirements:

- All API calls must be secure.
- Callers to the API must not send credentials to the API. Which authentication mechanism should you use?

- A. Basic
- B. Anonymous
- C. Managed identity
- D. Client certificate

**Answer: C**

**Explanation:**

Use the authentication-managed-identity policy to authenticate with a backend service using the managed identity of the API Management service. This policy essentially uses the managed identity to obtain an access token from Azure Active Directory for accessing the specified resource. After successfully obtaining the token, the policy will set the value of the token in the Authorization header using the Bearer scheme.

Reference:

<https://docs.microsoft.com/bs-cyrl-ba/azure/api-management/api-management-authentication-policies>

**NEW QUESTION 48**

- (Exam Topic 3)

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: Move photo processing to an Azure Function triggered from the blob upload. Does the solution meet the goal?

- A. Yes
- B. No

**Answer: A**

**Explanation:**

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.

Events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener.

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 52**

- (Exam Topic 3)

You have a web service that is used to pay for food deliveries. The web service uses Azure Cosmos DB as the data store.

You plan to add a new feature that allows users to set a tip amount. The new feature requires that a property named tip on the document in Cosmos DB must be present and contain a numeric value.

There are many existing websites and mobile apps that use the web service that will not be updated to set the tip property for some time.

How should you complete the trigger?

NOTE: Each correct selection is worth one point.

```
function ensureTip() {
var r = 

|                             |
|-----------------------------|
| _.value();                  |
| _.readDocument('item');     |
| getContext().getRequest();  |
| getContext().getResponse(); |


var i = r.getBody();
    

|                                            |
|--------------------------------------------|
| if (!("tip" in i)) {                       |
| if (request.getValue("tip") === null){     |
| if (isNaN(i["tip"])    i["tip"]=== null) { |
| if (typeof_.pluck("tip") == 'number') {    |


        i["tip"] = 0;
    }
    

|                      |
|----------------------|
| r.setBody(i);        |
| r.setValue(i);       |
| _.upsertDocument(i); |
| _.replaceDocument(i) |


}
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: getContext().getRequest(); Box 2: if(isNaN(i) ["tip"] ..

In JavaScript, there are two ways to check if a variable is a number :

isNaN() – Stands for “is Not a Number”, if variable is not a number, it return true, else return false. typeof – If variable is a number, it will returns a string named “number”.

Box 3:r.setBody(i);

// update the item that will be created References:

<https://docs.microsoft.com/bs-latn-ba/azure/cosmos-db/how-to-write-stored-procedures-triggers-udfs>

<https://mkyong.com/javascript/check-if-variable-is-a-number-in-javascript/>

**NEW QUESTION 56**

- (Exam Topic 3)

You are preparing to deploy an ASP.NET Core website to an Azure Web App from a GitHub repository. The website includes static content generated by a script. You plan to use the Azure Web App continuous deployment feature.

You need to run the static generation script before the website starts serving traffic.

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 a file named .deployment in the root of the repository that calls a script which generates the static content and deploys the website.
- B. Add a PreBuild target in the websites csproj project file that runs the static content generation script.
- C. Create a file named run.cmd in the folder /run that calls a script which generates the static content and deploys the website.
- D. Add the path to the static content generation tool to WEBSITE\_RUN\_FROM\_PACKAGE setting in the host.json file.

**Answer:** AD

**Explanation:**

A: To customize your deployment, include a .deployment file in the repository root.

You just need to add a file to the root of your repository with the name .deployment and the content: [config]

command = YOUR COMMAND TO RUN FOR DEPLOYMENT

this command can be just running a script (batch file) that has all that is required for your deployment, like copying files from the repository to the web root directory for example.

D: In Azure, you can run your functions directly from a deployment package file in your function app. The other option is to deploy your files in the d:\home\site\wwwroot directory of your function app (see A above).

To enable your function app to run from a package, you just add a WEBSITE\_RUN\_FROM\_PACKAGE setting to your function app settings.

Note: The host.json metadata file contains global configuration options that affect all functions for a function app.

References:

<https://github.com/projectkudu/kudu/wiki/Custom-Deployment-Script>

<https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment-package>

**NEW QUESTION 58**

- (Exam Topic 3)

A company develops a series of mobile games. All games use a single leaderboard service. You have the following requirements:

- Code should be scalable and allow for growth.
- Each record must consist of a playedId, gameId, score, and time played.
- When users reach a new high score, the system will save the new score using the SaveScore function below.
- Each game is assigned and Id based on the series title.

You have the following code. (Line numbers are included for reference only.)

You store customer information in an Azure Cosmos database. The following data already exists in the database:

```
01 CloudTableClient tableClient = account.CreateCloudTableClient();
02 CloudTable table = tableClient.GetTableReference("people");
03 TableQuery<CustomerEntity> query = new TableQuery<CustomerEntity>()
04     .Where(TableQuery.CombineFilters(
05     TableQuery.Generate.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal, "Smith")
06     TableOperators.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal,
"ssmith@contoso.com")
07     ));
08 await table.ExecuteQuerySegmentedAsync<CustomerEntity>(query, null);
```

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

	Yes	No
The code will work with Cosmos DB.	<input type="radio"/>	<input type="radio"/>
The save score function will update and replace a record if one already exists with the same playerId and gameId.	<input type="radio"/>	<input type="radio"/>
The data for the game will be automatically partitioned.	<input type="radio"/>	<input type="radio"/>
This code will store the values for the gameId and playerId parameters in the database.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Yes

Code for CosmosDB, example:

```
// Parse the connection string and return a reference to the storage account. CloudStorageAccount storageAccount = CloudStorageAccount.Parse(
CloudConfigurationManager.GetSetting("StorageConnectionString"));
```

// Create the table client.

```
CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
```

// Retrieve a reference to the table.

```
CloudTable table = tableClient.GetTableReference("people");
```

// Create the TableOperation object that inserts the customer entity. TableOperation insertOperation = TableOperation.Insert(customer1);

Box 2: No  
A new record will always be added as TableOperation.Insert is used, instead of TableOperation.InsertOrReplace.

Box 3: No

No partition key is used. Box 4: Yes

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet>

**NEW QUESTION 59**

- (Exam Topic 3)

You are developing an ASP.NET Core Web API web service that uses Azure Application Insights to monitor performance and track events.

You need to enable logging and ensure that log messages can be correlated to events tracked by Application Insights.

How should you complete the code? To answer, drag the appropriate code segments to the correct locations. Each code 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.

### Code segments

### Answer Area

```

public class Startup
{
    . . .
    public void ConfigureServices (IServiceCollection services)
    {
        services.AddOptions<ApplicationInsightsLoggerOptions>().
        Configure(o => o.IncludeEventId = true );
        services.AddMvc();
    }
    public void Configure (IApplicationBuilder app,
        IHostingEnvironment env, ILoggerFactory loggerFactory)
    {
        loggerFactory.AddApplicationInsights(app.ApplicationServices, LogLevel.Trace);
        app.UseMvc();
    }
}

```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: ApplicationInsightsLoggerOptions

If you want to include the EventId and EventName properties, then add the following to the ConfigureServices method:  
 services.AddOptions<ApplicationInsightsLoggerOptions>().Configure(o => o.IncludeEventId = true);

Box 2: IncludeEventID

Box 3: ApplicationServices

In Asp.Net core apps it turns out that trace logs do not show up in Application Insights out of the box. We need to add the following code snippet to our Configure method in Startup.cs:

loggerFactory.AddApplicationInsights(app.ApplicationServices, logLevel);

References:  
<https://blog.computecloud.com/enabling-application-insights-trace-logging-in-asp-net-core/>

**NEW QUESTION 62**

- (Exam Topic 3)

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 65**

- (Exam Topic 3)

You provide an Azure API Management managed web service to clients. The back end web service implements HTTP Strict Transport Security (HSTS).

Every request to the backend service must include a valid HTTP authorization header. You need to configure the Azure API Management instance with an authentication policy. Which two policies can you use? Each correct answer presents a complete solution NOTE: Each correct selection is worth one point.

- A. Certificate Authentication
- B. Basic Authentication

- C. OAuth Client Credential Grant
- D. Digest Authentication

**Answer:** AC

**NEW QUESTION 66**

- (Exam Topic 3)

You are creating an app that uses Event Grid to connect with other services. Your app's event data will be sent to a serverless function that checks compliance. This function is maintained by your company. You write a new event subscription at the scope of your resource. The event must be invalidated after 3 specific period of time. You need to configure Event Grid to ensure security.

What should you implement? To answer, select the appropriate options in [he answer area. NOTE: Each correct selection is worth one point

Authentication	Type
WebHook event delivery	<div style="border: 1px solid black; padding: 5px;"> <div style="text-align: right; margin-bottom: 5px;">▼</div> <div style="padding: 5px;"> <p>SAS tokens</p> <p>Key authentication</p> <p>JWT token</p> </div> </div>
Topic publishing	<div style="border: 1px solid black; padding: 5px;"> <div style="text-align: right; margin-bottom: 5px;">▼</div> <div style="padding: 5px;"> <p>ValidationCode handshake</p> <p>ValidationURL handshake</p> <p>Management Access Control</p> </div> </div>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: SAS tokens

Custom topics use either Shared Access Signature (SAS) or key authentication. Microsoft recommends SAS, but key authentication provides simple programming, and is compatible with many existing webhook publishers.

In this case we need the expiration time provided by SAS tokens. Box 2: ValidationCode handshake

Event Grid supports two ways of validating the subscription: ValidationCode handshake (programmatic) and ValidationURL handshake (manual).

If you control the source code for your endpoint, this method is recommended.

**NEW QUESTION 70**

- (Exam Topic 3)

A company is implementing a publish-subscribe (Pub/Sub) messaging component by using Azure Service Bus. You are developing the first subscription application.

In the Azure portal you see that messages are being sent to the subscription for each topic. You create and initialize a subscription client object by supplying the correct details, but the subscription application is still not consuming the messages.

You need to complete the source code of the subscription client What should you do?

- A. `await subscriptionClient.CloseAsync();`
- B. `await subscriptionClient.AddRuleAsync(new RuleDescription(RuleDescription.DefaultRuleName, new TrueFilter()));`
- C. `subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync, messageHandlerOptions);`
- D. `subscriptionClient = new SubscriptionClient(ServiceBusConnectionString, TopicName, SubscriptionName);`

**Answer:** C

**Explanation:**

Using topic client, call RegisterMessageHandler which is used to receive messages continuously from the entity. It registers a message handler and begins a new thread to receive messages. This handler is waited on every time a new message is received by the receiver.

`subscriptionClient.RegisterMessageHandler(ReceiveMessagesAsync, messageHandlerOptions);` References:

<https://www.c-sharpcorner.com/article/azure-service-bus-topic-and-subscription-pub-sub/>

**NEW QUESTION 74**

- (Exam Topic 3)

You develop a gateway solution for a public facing news API. The news API back end is implemented as a RESTful service and uses an OpenAPI specification. You need to ensure that you can access the news API by using an Azure API Management service instance. Which Azure PowerShell command should you run?

- A. `Import-AzureRmApiManagementApi -Context $ApiMgmtContext -SpecificationFormat "Swagger" -SpecificationPath $SwaggerPath -Path $Path`

- B. New-AzureRmApiManagementBackend -Context \$ApiMgmtContext -Url \$Url -Protocol http
- C. New-AzureRmApiManagement -ResourceGroupName \$ResourceGroup -Name \$Name - Location \$Location -Organization \$Org -AdminEmail \$AdminEmail
- D. New-AzureRmApiManagementBackendProxy -Url \$ApiUrl

**Answer:** D

**Explanation:**

New-AzureRmApiManagementBackendProxy creates a new Backend Proxy Object which can be piped when creating a new Backend entity.

Example: Create a Backend Proxy In-Memory Object

```
PS C:\>$secpassword = ConvertTo-SecureString "PlainTextPassword" -AsPlainText -Force
```

```
PS C:\>$proxyCreds = New-Object System.Management.Automation.PSCredential ("foo", $secpassword) PS C:\>$credential = New-
```

```
AzureRmApiManagementBackendProxy -Url "http://12.168.1.1:8080"
```

```
-ProxyCredential $proxyCreds
```

```
PS C:\>$apimContext = New-AzureRmApiManagementContext -ResourceGroupName "Api-Default-WestUS" -ServiceName "contoso"
```

```
PS C:\>$backend = New-AzureRmApiManagementBackend -Context $apimContext -BackendId 123 -Url 'https://contoso.com/awesomeapi' -Protocol http -Title
```

```
"first backend" -SkipCertificateChainValidation $true
```

```
-Proxy $credential -Description "backend with proxy server"
```

Creates a Backend Proxy Object and sets up Backend

**NEW QUESTION 78**

- (Exam Topic 3)

Fourth Coffee has an ASP.NET Core web app that runs in Docker. The app is mapped to the www.fourthcoffee.com domain.

Fourth Coffee is migrating this application to Azure.

You need to provision an App Service Web App to host this docker image and map the custom domain to the App Service web app.

A resource group named FourthCoffeePublicWebResourceGroup has been created in the WestUS region that contains an App Service Plan named AppServiceLinuxDockerPlan.

Which order should the CLI commands be used to develop the solution? To answer, move all of the Azure CLI command from the list of commands to the answer area and arrange them in the correct order.

**Azure CLI commands**

```
az webapp config hostname add
--webapp-name $appName
--resource-group fourthCoffeePublicWebResourceGroup
--hostname $fqdn
```

```
#!/bin/bash
appName="FourthCoffeePublicWeb$random".
location "WestUS"
dockerHubContainerPath="FourthCoffee/publicweb:v1"
fqdn=http://www.fourthcoffee.com>www.fourthcoffee.com
```

```
az webapp create
--name $appName
--plan AppServiceLinuxDockerPlan
--resource-group fourthCoffeePublicWebResourceGroup
```

```
az webapp config container set
--docker-custom-image-name $dockerHibContainerPath
--name $appName
--resource-group fourthCoffeePublicWebResourceGroup
```

**Answer area**



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: #bin/bash

The appName is used when the webapp-name is created in step 2. Step 2: az webapp config hostname add

The webapp-name is used when the webapp is created in step 3. Step 3: az webapp create

Create a web app. In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command.

Step : az webapp config container set

In Create a web app, you specified an image on Docker Hub in the az webapp create command. This is good enough for a public image. To use a private image, you need to configure your Docker account ID and password in your Azure web app.

In the Cloud Shell, follow the az webapp create command with az webapp config container set.

References:

<https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image>

### NEW QUESTION 79

- (Exam Topic 3)

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.

You are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.

You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level. You need to configure authorization.

Solution:

- Create a new Azure AD application's manifest, set value of the groupMembershipClaims option to All.
- In the website, use the value of the groups claim from the JWT for the user to determine permissions. Does the solution meet the goal?

A. Yes

B. No

**Answer: A**

#### Explanation:

To configure Manifest to include Group Claims in Auth Token

\* 1. Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to App registrations to find your application:

\* 2. Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.

\* 3. Locate the "groupMembershipClaims" setting. Set its value to either "SecurityGroup" or "All". To help you decide which:

"SecurityGroup" - groups claim will contain the identifiers of all security groups of which the user is a member.

"All" - groups claim will contain the identifiers of all security groups and all distribution lists of which the user is a member

Now your application will include group claims in your manifest and you can use this fact in your code. References:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

### NEW QUESTION 81

- (Exam Topic 3)

You develop a website. You plan to host the website in Azure. You expect the website to experience high traffic volumes after it is published. You must ensure that the website remains available and responsive while minimizing cost. You need to deploy the website. What should you do?

- A. Deploy the website to an App Service that uses the Shared service tie
- B. Configure the App Service plan to automatically scale when the CPU load is high.
- C. Deploy the website to a virtual machine
- D. Configure the virtual machine to automatically scale when the CPU load is high.
- E. Deploy the website to an App Service that uses the Standard service tie
- F. Configure the App Service plan to automatically scale when the CPU load is high.
- G. Deploy the website to a virtual machine
- H. Configure a Scale Set to increase the virtual machine instance count when the CPU load

**Answer: C**

#### Explanation:

Windows Azure Web Sites (WAWS) offers 3 modes: Standard, Free, and Shared.

Standard mode carries an enterprise-grade SLA (Service Level Agreement) of 99.9% monthly, even for sites with just one instance.

Standard mode runs on dedicated instances, making it different from the other ways to buy Windows Azure Web Sites.

### NEW QUESTION 85

- (Exam Topic 3)

You are developing a new page for a website that uses Azure Cosmos DB for data storage. The feature uses documents that have the following format:

You must display data for the new page in a specific order. You create the following query for the page:

You need to configure a Cosmos DB policy to support the query.

How should you configure the policy? 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**

**Answer Area**

- 
- 
- 
- 
- 

```

{
  "automatic": true,
  "indexingMode": "Consistent",
  "includedPaths": [
    {
      "path": "/*"
    }
  ],
  "excludedPaths": [],
  "compositeIndexes": [
    {
      "path": "/name", "order": "descending"
    },
    {
      "path": "/city", "order": "
  ]
}

```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: compositeIndexes

You can order by multiple properties. A query that orders by multiple properties requires a composite index. Box 2: descending

Example: Composite index defined for (name ASC, age ASC):

It is optional to specify the order. If not specified, the order is ascending.

```

{
  "automatic":true, "indexingMode":"Consistent", "includedPaths":[
  {
  "path":"/*"
  }
  ],
  "excludedPaths":[], "compositeIndexes":[ [
  {
  "path":"/name",
  },
  {
  "path":"/age",
  }
  ]
  ]
}

```

**NEW QUESTION 87**

- (Exam Topic 3)

You are developing a solution that will use Azure messaging services.

You need to ensure that the solution uses a publish-subscribe model and eliminates the need for constant polling.

What are two possible ways to achieve the goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Service Bus
- B. Event Hub
- C. Event Grid
- D. Queue

**Answer:** AC

**Explanation:**

It is strongly recommended to use available messaging products and services that support a publish-subscribe model, rather than building your own. In Azure, consider using Service Bus or Event Grid. Other technologies that can be used for pub/sub messaging include Redis, RabbitMQ, and Apache Kafka.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber>

**NEW QUESTION 92**

- (Exam Topic 3)

You develop a serverless application using several Azure Functions. These functions connect to data from within the code.

You want to configure tracing for an Azure Function App project. You need to change configuration settings in the hostjson file. Which tool should you use?

- A. Azure portal
- B. Azure PowerShell
- C. Azure Functions Core Tools (Azure CLI)
- D. Visual Studio

**Answer: A**

**Explanation:**

The function editor built into the Azure portal lets you update the function.json file and the code file for a function. The host.json file, which contains some runtime-specific configurations, is in the root folder of the function app.

References:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-reference#fileupdate>

**NEW QUESTION 94**

- (Exam Topic 3)

You are writing code to create and run an Azure Batch job. You have created a pool of compute nodes.

You need to choose the right class and its method to submit a batch job to the Batch service. Which method should you use?

- A. JobOperations.CreateJobO
- B. CloudJob.Enable(IEnumerable<BatchClientBehavior>)
- C. CloudJob.CommitAsync(IEnumerable<BatchClientBehavior>, CancellationToken)
- D. JobOperations.EnableJob(String, IEnumerable<BatchClientBehavior>)
- E. JobOperations.EnableJobAsync(Strin
- F. IEnumerable<BatchClientBehavior>. CancellationToken)

**Answer: C**

**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.

The Commit method submits the job to the Batch service. Initially the job has no tasks.

```
{
CloudJob job = batchClient.JobOperations.CreateJob(); job.Id = JobId;
job.PoolInformation = new PoolInformation { PoolId = PoolId }; job.Commit();
}
```

References:

<https://docs.microsoft.com/en-us/azure/batch/quick-run-dotnet>

**NEW QUESTION 95**

- (Exam Topic 3)

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.

Policy types	Requirement	Policy type
Inbound	Rewrite the request URL to match to the format expected by the web service.	policy type
Outbound	Remove formatting text from responses.	policy type
Backend	Forward the user ID that is associated with the subscription key for the original request to the back-end service.	policy type

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Policy types	Requirement	Policy type
Inbound	Rewrite the request URL to match to the format expected by the web service.	Outbound
Outbound	Remove formatting text from responses.	Inbound
Backend	Forward the user ID that is associated with the subscription key for the original request to the back-end service.	Backend

**NEW QUESTION 97**

- (Exam Topic 3)

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 99**

- (Exam Topic 3)

You have an Azure App Services Web App. Azure SQL Database instance. Azure Storage Account and an Azure Redis Cache instance in a resource group. A developer must be able to publish code to the web app. You must grant the developer the Contributor role to the web app

You need to grant the role.

What two commands can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. New-AzureRmRoleAssignment
- B. az role assignment create
- C. az role definition create
- D. New-AzureRmRoleDefinition

**Answer: AB**

**Explanation:**

References:

<https://docs.microsoft.com/en-us/cli/azure/role/assignment?view=azure-cli-latest#az-role-assignment-create> <https://docs.microsoft.com/en-us/powershell/module/azurerm.resources/new-azureroleassignment?view=azur>

**NEW QUESTION 101**

- (Exam Topic 3)

A company backs up all manufacturing data to Azure Blob Storage. Admins move blobs from hot storage to archive tier storage every month.

You must automatically move blocks to Archive tier after they have not been accessed for 180 days. The path for any item that is not archived must be placed in an existing queue. This operation must be performed automatically once a month. You set the value of TierAgeInDays to 180.

How should you configure the Logic App? To answer, drag the appropriate triggers or action blocks to the correct trigger or action slots. Each trigger or action block 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.

### Triggers and Action Blocks

**Insert Entity**

\*Table: processing

\*Entity: Path

Show advanced options

**Tier blob**

If blob is older than the defined value, tier it to Cool or Archive tier

\*Blob path: Path

\*Blob Tier: Archive

**When there are messages in a queue**

\*Queue Name: processing

Connected to tableStorageAccountConnection

**Recurrence**

\*Interval: 1

\*Frequency: Month

### Answer Area

Set tier age variable

Set tier age variable

For each

Scan all blobs in this folder

Select an output from previous steps: value

**When there are messages in a queue**

\*Queue Name: processing

Connected to tableStorageAccountConnection

If true

If false

Add an action

Add an action

Add an action

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Recurrence  
 Box 2: Insert Entity  
 Box 3 (if true): Tier Blob  
 Box 4: (if false):  
 Leave blank. References:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-perform-data-operations>

**NEW QUESTION 104**

- (Exam Topic 3)

A company runs an international travel and bookings management service. The company plans to begin offering restaurant bookings. You must develop a solution that uses Azure Search and meets the following requirements:

- Users must be able to search for restaurants by name, description, location, and cuisine.
- Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.
- All words in descriptions must be included in searches. You need to add annotations to the restaurant class.

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.

```
[SerializePropertyNameAsCamelCase]
public class Restaurant
{
    [Key, IsFilterable]
    public int RestaurantId { get; set; }
    [IsSearchable, IsFilterable, IsSortable]
    public string Name { get; set; }
    [IsSearchable.IsFilterable.IsSortable, IsFacetable]
    [IsFilterable.IsFacetable, Required]
    [IsSearchable]
    [IsSearchable, Required]
    public string location { get; set; }
    public string Phone { get; set; }
    [Required]
    [IsSearchable]
    [IsFilterable, IsFacetable, Required]
    [IsFilterable, IsFacetable, IsSortable]
    public string Description { get; set; }
    [IsFilterable, IsSortable, IsSearchable]
    [IsFilterable, IsSortable, IsFacetable]
    [IsFilterable, IsSortable, Key]
    [IsFilterable, IsSortable, IsSearchable, Required]
    public double Rating { get; set; }
    [IsSearchable, IsFilterable, IsFacetable]
    [IsFilterable, IsSortable, Key]
    [IsFilterable, IsSortable, IsSearchable]
    [IsFilterable, IsSortable, Key, Required]
    public List<string> Cuisines { get; set; }
    [IsFilterable, IsSortable, Key, Required]
    [IsSearchable, IsSortable, IsFacetable]
    [IsFilterable, IsSortable, Key, IsSearchable]
    [IsFilterable, IsFacetable]
    public bool FamilyFriendly { get; set; }
}
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: [IsSearchable.IsFilterable.IsSortable,IsFacetable] Location

Users must be able to search for restaurants by name, description, location, and cuisine.

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. Box 2: [IsSearchable.IsFilterable.IsSortable,Required]

Description

Users must be able to search for restaurants by name, description, location, and cuisine. All words in descriptions must be included in searches.

Box 3: [IsFilterable,IsSortable,IsFaceTable] Rating

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. Box 4: [IsSearchable.IsFilterable,IsFacetable]

Cuisines

Users must be able to search for restaurants by name, description, location, and cuisine.

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. Box 5: [IsFilterable,IsFacetable]

FamilyFriendly

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. References:

<https://www.henkboelman.com/azure-search-the-basics/>

**NEW QUESTION 109**

- (Exam Topic 3)

You are developing an ASP.NET Core website that can be used to manage photographs which are stored in Azure Blob Storage containers.

Users of the website authenticate by using their Azure Active Directory (Azure AD) credentials.

You implement role-based access control (RBAC) role permission on the containers that store photographs. You assign users to RBAC role.

You need to configure the website's Azure AD Application so that user's permissions can be used with the Azure Blob containers.

How should you configure the application? To answer, drag the appropriate setting to the correct location. Each setting 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.

Settings	Answer Area		
client_id	API	Permission	Type
delegated	Azure Storage	Setting	Setting
profile	Microsoft Graph	User.Read	Setting
application			
user_impersonation			

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: user\_impersonation

Box 2: delegated Example:

- \* 1. Select the API permissions section
- \* 2. Click the Add a permission button and then: Ensure that the My APIs tab is selected
- \* 3. In the list of APIs, select the API TodoListService-aspnetcore.
- \* 4. In the Delegated permissions section, ensure that the right permissions are checked: user\_impersonation. 5. Select the Add permissions button.

Box 3: delegated Example

- \* 1. Select the API permissions section
- \* 2. Click the Add a permission button and then, Ensure that the Microsoft APIs tab is selected
- \* 3. In the Commonly used Microsoft APIs section, click on Microsoft Graph
- \* 4. In the Delegated permissions section, ensure that the right permissions are checked: User.Read. Use the search box if necessary.
- \* 5. Select the Add permissions button

References:  
<https://docs.microsoft.com/en-us/samples/azure-samples/active-directory-dotnet-webapp-webapi-openidconnect>

**NEW QUESTION 113**

- (Exam Topic 3)

You are implementing an order processing system. A point of sale application publishes orders to topics in an Azure Service Bus queue. The label property for the topic includes the following data:

Property	Description
ShipLocation	the country/region where the order will be shipped
CorrelationId	a priority value for the order
Quantity	a user-defined field that stores the quantity of items in an order
AuditedAt	a user-defined field that records the date an order is audited

The system has the following requirements for subscriptions

Subscription type	Comments
FutureOrders	This subscription is reserved for future use and must not receive any orders.
HighPriorityOrders	Handle all high priority orders and International orders.
InternationalOrders	Handle orders where the country/region is not United States.
HighQuantityOrders	Handle only orders with quantities greater than 100 units.
AllOrders	This subscription is used for auditing purposes. This subscription must receive every single order. AllOrders has an Action defined that updates the AuditedAt property to include the date and time it was received by the subscription.

You need to implement filtering and maximize throughput while evaluating filters.

Which filter types should you implement? To answer, drag the appropriate filter types to the correct subscriptions. Each filter 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.

**Filter types**

- SQLFilter
- CorrelationFilter
- No Filter

**Answer Area**

Subscription	Filter type
FutureOrders	
HighPriorityOrders	
InternationalOrders	
HighQuantityOrders	
AllOrders	

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

FutureOrders: SQLFilter HighPriorityOrders: CorrelationFilter CorrelationID only InternationalOrders: SQLFilter Country NOT USA requires an SQL Filter HighQuantityOrders: SQLFilter Need to use relational operators so an SQL Filter is needed. AllOrders: No Filter

SQL Filter: SQL Filters - A SqlFilter holds a SQL-like conditional expression that is evaluated in the broker against the arriving messages' user-defined properties and system properties. All system properties must be prefixed with sys. in the conditional expression. The SQL-language subset for filter conditions tests for the existence of properties (EXISTS), as well as for null-values (IS NULL), logical NOT/AND/OR, relational operators, simple numeric arithmetic, and simple text pattern matching with LIKE.

Correlation Filters - A CorrelationFilter holds a set of conditions that are matched against one or more of an arriving message's user and system properties. A common use is to match against the CorrelationId property, but the application can also choose to match against ContentType, Label, MessageId, ReplyTo, ReplyToSessionId, SessionId, To, and any user-defined properties. A match exists when an arriving message's value for a property is equal to the value specified in the correlation filter. For string expressions, the comparison is case-sensitive. When specifying multiple match properties, the filter combines them as a logical AND condition, meaning for the filter to match, all conditions must match.

Boolean filters - The TrueFilter and FalseFilter either cause all arriving messages (true) or none of the arriving messages (false) to be selected for the subscription. References: <https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters>

**NEW QUESTION 116**

- (Exam Topic 3)

You plan to deploy a new application to a Linux virtual machine (VM) that is hosted in Azure.

The entire VM must be secured at rest by using industry-standard encryption technology to address organizational security and compliance requirements.

You need to configure Azure Disk Encryption for the VM.

How should you complete the Azure Cli commands? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

```

az provider register -n Microsoft.KeyVault
resourcegroup="myResourceGroup"
az group create --name $resourcegroup --location westus
keyvault_name=myvaultname$RANDOM
az vm create \
  --enable-for-disk-encryption True
  
```

The image shows a screenshot of the Azure CLI interface with several dropdown menus open for the 'vm create' command. The 'vm encryption' dropdown is highlighted in blue, and the 'data' option is selected. Other dropdowns for 'vm encryption' are also visible, showing options like 'all', 'data', and 'OS'.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: keyvault

Create an Azure Key Vault with az keyvault create and enable the Key Vault for use with disk encryption. Specify a unique Key Vault name for keyvault\_name as follows:

```
keyvault_name=myvaultname$RANDOM az keyvault create \
--name $keyvault_name \
--resource-group $resourcegroup \
--location eastus \
--enabled-for-disk-encryption True
```

Box 2: keyvault key

The Azure platform needs to be granted access to request the cryptographic keys when the VM boots to decrypt the virtual disks. Create a cryptographic key in your Key Vault with az keyvault key create. The following example creates a key named myKey:

```
az keyvault key create \
--vault-name $keyvault_name \
--name myKey \
--protection software
```

Box 3: vm

Create a VM with az vm create. Only certain marketplace images support disk encryption. The following example creates a VM named myVM using an Ubuntu 16.04 LTS image:

```
az vm create \
--resource-group $resourcegroup \
--name myVM \
--image Canonical:UbuntuServer:16.04-LTS:latest \
--admin-username azureuser \
--generate-ssh-keys
```

Box 4: vm encryption

Encrypt your VM with az vm encryption enable: az vm encryption enable \

```
--resource-group $resourcegroup \
--name myVM \
--disk-encryption-keyvault $keyvault_name \
--key-encryption-key myKey \
--volume-type all
```

Note: seems to an error in the question. Should have enable instead of create.

Box 5: all

Encrypt both data and operating system.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/virtual-machines/linux/encrypt-disks>

**NEW QUESTION 117**

- (Exam Topic 3)

You are deploying an Azure Kubernetes Services (AKS) cluster that will use multiple containers.

You need to create the cluster and verify that the services for the containers are configured correctly and available.

Which four commands should you use to develop the solution? To answer, move the appropriate command segments from the list of command segments to the answer area and arrange them in the correct order.

**Command segments**

**Answer Area**

- az aks get-credentials
- az appservice plan create
- az aks create
- az group create
- kubectl apply



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: az group create

Create a resource group with the az group create command. An Azure resource group is a logical group in which Azure resources are deployed and managed.

Example: The following example creates a resource group named myAKSCluster in the eastus location. az group create --name myAKSCluster --location eastus

Step 2 : az aks create

Use the az aks create command to create an AKS cluster. Step 3: kubectl apply

To deploy your application, use the kubectl apply command. This command parses the manifest file and creates the defined Kubernetes objects.

Step 4: az aks get-credentials

Configure it with the credentials for the new AKS cluster. Example:  
 az aks get-credentials --name aks-cluster --resource-group aks-resource-group References:  
<https://docs.bitnami.com/azure/get-started-aks/>

**NEW QUESTION 122**

- (Exam Topic 3)

You are developing a solution for a hospital to support the following use cases:

- The most recent patient status details must be retrieved even if multiple users in different locations have updated the patient record.
- Patient health monitoring data retrieved must be the current version or the prior version.
- After a patient is discharged and all charges have been assessed, the patient billing record contains the final charges.

You provision a Cosmos DB NoSQL database and set the default consistency level for the database account to Strong. You set the value for Indexing Mode to Consistent.

You need to minimize latency and any impact to the availability of the solution. You must override the default consistency level at the query level to meet the required consistency guarantees for the scenarios.

Which consistency levels should you implement? To answer, drag the appropriate consistency levels to the correct requirements. Each consistency level 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.

**Consistency levels**

**Answer Area**

Strong	Bounded Staleness	Return the most recent patient status.	
Consistent Prefix	Eventual	Return health monitoring data that is no less than one version behind.	
		After patient is discharged and all changes are assessed, retrieve the correct billing data with the final charges	

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Strong

Strong: Strong consistency offers a linearizability guarantee. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

Box 2: Bounded staleness

Bounded staleness: The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions (that is "updates") of an item or by "t" time interval. When you choose bounded staleness, the "staleness" can be configured in two ways:

The number of versions (K) of the item

The time interval (t) by which the reads might lag behind the writes

Box 3: Eventual  
 Eventual: There's no ordering guarantee for reads. In the absence of any further writes, the replicas eventually converge.

**NEW QUESTION 126**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **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

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The AZ-204 Practice Test Here](#)**