



Microsoft

Exam Questions AZ-204

Developing Solutions for Microsoft Azure (beta)

NEW QUESTION 1

- (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>
  <
  </
  <AllowedMethods>
  </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 2

- (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 3

- (Exam 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

Answer Area

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.



- 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 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;"> ▼ </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;"> ▼ </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;"> ▼ </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	<div style="border: 1px solid black; padding: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</div> <div style="padding: 2px;">Basic</div> <div style="padding: 2px;">Standard</div> <div style="padding: 2px;">Premium</div> <div style="padding: 2px;">Isolated</div> </div>
Instance Count	<div style="border: 1px solid black; padding: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</div> <div style="padding: 2px;">1</div> <div style="padding: 2px;">10</div> <div style="padding: 2px;">20</div> <div style="padding: 2px;">100</div> </div>

- 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 support the requirements for the Shipping Logic App.

What should you use?

- A. Azure Active Directory Application Proxy
- B. Point-to-Site (P2S) VPN connection
- C. Site-to-Site (S2S) VPN connection
- D. On-premises Data Gateway

Answer: D

Explanation:

Before you can connect to on-premises data sources from Azure Logic Apps, download and install the on-premises data gateway on a local computer. The gateway works as a bridge that provides quick data transfer and encryption between data sources on premises (not in the cloud) and your logic apps.

The gateway supports BizTalk Server 2016.

Note: Microsoft have now fully incorporated the Azure BizTalk Services capabilities into Logic Apps and Azure App Service Hybrid Connections.

Logic Apps Enterprise Integration pack bring some of the enterprise B2B capabilities like AS2 and X12, EDI standards support

Scenario: The Shipping Logic app must meet the following requirements:

- > Support the ocean transport and inland transport workflows by using a Logic App.
- > Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
- > Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
- > Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Reference:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-gateway-install>

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

	▼
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 8

- (Exam Topic 3)

You use Azure Table storage to store customer information for an application. The data contains customer details and is partitioned by last name. You need to create a query that returns all customers with the last name Smith. Which code segment should you use?

- A. `TableQuery.GenerateFilterCondition("PartitionKey", Equals, "Smith")`
- B. `TableQuery.GenerateFilterCondition("LastName", Equals, "Smith")`
- C. `TableQuery.GenerateFilterCondition("PartitionKey", QueryComparisons.Equal, "Smith")`
- D. `TableQuery.GenerateFilterCondition("LastName", QueryComparisons.Equal, "Smith")`

Answer: C

Explanation:

Retrieve all entities in a partition. The following code example specifies a filter for entities where 'Smith' is the partition key. This example prints the fields of each entity in the query results to the console.

```
Construct the query operation for all customer entities where PartitionKey="Smith".
TableQuery<CustomerEntity> query = new
TableQuery<CustomerEntity>().Where(TableQuery.GenerateFilterCondition("PartitionKey",
QueryComparisons.Equal, "Smith"));
```

References:

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

NEW QUESTION 9

- (Exam Topic 3)

You are developing an Azure Function App by using Visual Studio. The app will process orders input by an Azure Web App. The web app places the order information into Azure Queue Storage.

You need to review the Azure Function App code shown below. NOTE: Each correct selection is worth one point.

```
public static class OrderProcessor
{
    [FunctionName("ProcessOrders")]
    public static void ProcessOrders([QueueTrigger("incoming-orders")]IInkQueueMessage myQueueItem, [Table
("Orders")]ICollector<Order> tableBindings,
    TraceWriter log)
    {
        log.Info($"Processing Order: {myQueueItem.Id}");
        log.Info($"Queue Insertion Time: {myQueueItem.InsertionTime}");
        log.Info($"Queue Expiration Time: {myQueueItem.ExpirationTime}");
        tableBindings.Add(new Order().DeserializeObject<Order>(myQueueItem.AsString));
    }
    [FunctionName("ProcessOrders-Poison")]
    public static void ProcessFailedOrders([QueueTrigger("incoming-orders-poison")]IInkQueueMessage
myQueueItem, TraceWriter log)
    {
        log.Error($"Failed to process order: {myQueueItem.AsString}");
    }
}
```

Yes **No**

- The code will log the time that the order was processed from the queue. Yes No
- When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try. Yes No
- When there are multiple orders in the queue, a batch of orders will be retrieved from the queue and the ProcessOrders function will run multiple instances concurrently to process the orders. Yes No
- The ProcessOrders function will output the order to an Orders table in Azure Table Storage. Yes No

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

ExpirationTime - The time that the message expires. InsertionTime - The time that the message was added to the queue.

Box 2: Yes

maxDequeueCount - The number of times to try processing a message before moving it to the poison queue.

Default value is 5.

Box 3: Yes

When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them.

By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.

Box 4: Yes References:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue>

NEW QUESTION 10

- (Exam Topic 3)

A company is developing a Java web app. The web app code is hosted in a GitHub repository located at <https://github.com/Contoso/webapp>.

The web app must be evaluated before it is moved to production. You must deploy the initial code release to a deployment slot named staging.

You need to create the web app and deploy the code.

How should you complete the commands? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```
gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup

az [ ] create --location centralus - -name $resourcegroupname
group
webapp
appservice plan
webapp deployment slot
webapp deployment source
create --name $webappname - -resource-group $resourcegroupname
- -sku S3
create --name $webappname - -resource-group $resourcegroupname
\ - -plan $webappname
create --name $webappname - -resource-group $resourcegroupname
\ - -slot staging

az [ ] config - -name $webappname - -resource-group $resourcegroupname
\ - -slot staging - -repo-url
$gitrepo - -branch master - -manual-integration

az [ ]

az [ ]

az [ ]

az [ ]
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: group
 # Create a resource group.
 az group create --location westeurope --name myResourceGroup
 Box 2: appservice plan
 # Create an App Service plan in STANDARD tier (minimum required by deployment slots). az appservice plan create --name \$webappname --resource-group myResourceGroup --sku S1
 Box 3: webapp
 # Create a web app.
 az webapp create --name \$webappname --resource-group myResourceGroup \ --plan \$webappname
 Box 4: webapp deployment slot
 #Create a deployment slot with the name "staging".
 az webapp deployment slot create --name \$webappname --resource-group myResourceGroup \ --slot staging
 Box 5: webapp deployment source
 # Deploy sample code to "staging" slot from GitHub.
 az webapp deployment source config --name \$webappname --resource-group myResourceGroup \ --slot staging --repo-url \$gitrepo --branch master --manual-integration
 References:
<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-staging-environment>

NEW QUESTION 10

- (Exam Topic 3)

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	<input type="text"/>	▼
	Can be restored	
	Cannot be restored	
Snapshot 1	<input type="text"/>	▼
	Can be restored	
	Cannot be restored	
Snapshot 2	<input type="text"/>	▼
	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 14

- (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 17

- (Exam Topic 3)

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.

Uploading of videos occurs on an irregular basis.

You need to copy specific blobs from Container1 to Container2 in real time when specific requirements are met, excluding backup blob copies.

What should you do?

- A. Download the blob to a virtual machine and then upload the blob to Container2.
- B. Run the Azure PowerShell command Start-AzureStorageBlobCopy.
- C. Copy blobs to Container2 by using the Put Blob operation of the Blob Service REST API.
- D. Use AzCopy with the Snapshot switch blobs 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.

References:

<https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azuremp>

NEW QUESTION 22

- (Exam Topic 3)

You are a developer for a SaaS company that offers many web services. All web services for the company must meet the following requirements:

- > Use API Management to access the services
- > Use OpenID Connect for authentication
- > Prevent anonymous usage

A recent security audit found that several web services can be called without any authentication. Which API Management policy should you implement?

- A. jsonp
- B. authentication-certificate
- C. check-header
- D. validate-jwt

Answer: D

Explanation:

Add the validate-jwt policy to validate the OAuth token for every incoming request. Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad>

NEW QUESTION 23

- (Exam Topic 3)

You are developing a .NET Core model-view controller (MVC) application hosted on Azure for a health care system that allows providers access to their information.

You develop the following code:

```
services.AddAuthorization (options =>
{
    options.AddPolicy ("ProviderPartner", policy =>
    {
        .policy.AddAuthenticationSchemes ("Cookie, Bearer");
        policy.RequireAuthenticatedUser ();
        policy.RequireRole ("ProviderAdmin", "SysAdmin");
        policy.RequireClaim ("editor", "partner");
    });
});
```

You define a role named SysAdmin.

You need to ensure that the application meets the following authorization requirements:

- > Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.
- > Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.

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.

```
services.AddAuthorization (options =>
{
    options.AddPolicy ("ProviderPartner", .policy.AddAut
    {
        .policy.AddAuthenticationSchemes ("Cookie, Bearer");
        policy.RequireAuthenticatedUser ();
        policy.RequireRole ("ProviderAdmin" });
        policy.RequireClaim ("editor", "pa
    });
})
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:
 Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.
 Box 2:
 Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.

NEW QUESTION 28

- (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.

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 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 29

- (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 32

- (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 33

- (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 38

- (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 15:46:08.405Z	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 43

- (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
GetInstallation
CreateClientFromConnectionString
CreateOrUpdateInstallation
PatchInstallation
(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, windowsToastPayload);
}
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 45

- (Exam Topic 3)

You are developing an ASP.NET Core Web API web service. The web service uses Azure Application Insights for all telemetry and dependency tracking. The web service reads and writes data to a database other than Microsoft SQL Server.

You need to ensure that dependency tracking works for calls to the third-party database.

Which two Dependency Telemetry properties should you store in the database? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Telemetry.Context.Operation.Id
- B. Telemetry.Context.Cloud.RoleInstance
- C. Telemetry.Id

D. Telemetry.ContextSession.Id
 E. Telemetry.Name

Answer: AC

Explanation:

References:
<https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking>

NEW QUESTION 49

- (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",
  [Box 1] : "All",
  [Box 2] : true
  ...
}
    
```

Box 1 options: "optionalClaims", "groupMembershipClaims"

Box 2 options: "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 54

- (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 style="width: 100%;" type="text"/>
Deployment	Polling interval	<input style="width: 100%;" type="text"/>
ScaledObject	Azure Storage connection string	<input style="width: 100%;" 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 58

- (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 59

- (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,
}
```

- GetJob
- GetTask
- EnableJob
- CreateJob

- TaskSuccess
- TaskFailure
- TaskCompletion

- TaskSuccess
- TaskFailure
- TaskCompletion

- OutputFiles
- FilesToStage
- ResourceFiles
- StageFiles

- 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 63

- (Exam Topic 3)

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

Answer Area

Policy section	Policy	Policy section
	Set-variable	<input type="text"/>
<input type="text" value="Inbound"/>	Cache-lookup-value	<input type="text"/>
<input type="text" value="Outbound"/>	Cache-store-value	<input type="text"/>
	Find-and-replace	<input type="text"/>

- 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)" />
</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>
```

Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>

NEW QUESTION 66

- (Exam Topic 3)

You are developing an application that uses Azure Storage Queues. You have the following code:

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse
(CloudConfigurationManager.GetSetting("StorageConnectionString"));
CloudQueueClient queueClient = storageAccount.CreateCloudQueueClient()

CloudQueue queue = queueClient.GetQueueReference("appqueue") ;
await queue.CreateIfNotExistsAsync() ;

CloudQueueMessage peekedMessage = await queue.PeekMessageAsync() ;
if (peekedMessage != null)
{
    Console.WriteLine("The peeked message is: {0}", peekedMessage.AsString);
}
CloudQueueMessage message = await queue.GetMessageAsync() ;
```

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

NOTE: Each correct selection is worth one point.

Answer Area

Statement	Yes	No
The code configures the lock duration for the queue.	<input type="radio"/>	<input type="radio"/>
The last message read remains in the queue after the code runs.	<input type="radio"/>	<input type="radio"/>
The storage queue remains in the storage account after the code runs.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

The QueueDescription.LockDuration property gets or sets the duration of a peek lock; that is, the amount of time that the message is locked for other receivers. The maximum value for LockDuration is 5 minutes; the default value is 1 minute.

Box 2: Yes

You can peek at the message in the front of a queue without removing it from the queue by calling the PeekMessage method.

Box 3: Yes Reference:

<https://docs.microsoft.com/en-us/azure/storage/queues/storage-dotnet-how-to-use-queues> <https://docs.microsoft.com/en-us/dotnet/api/microsoft.servicebus.messaging.queuedescription.lockduration>

NEW QUESTION 67

- (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 69

- (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 70

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

- 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.computedcloud.com/enabling-application-insights-trace-logging-in-asp-net-core/>

NEW QUESTION 75

- (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 80

- (Exam Topic 3)

You are using Azure Front Door Service.

You are expecting inbound files to be compressed by using Brotli compression. You discover that inbound XML files are not compressed. The files are 9 megabytes (MB) in size.

You need to determine the root cause for the issue.

To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Statement	Yes	No
The file MIME type is supported by the service.	<input type="radio"/>	<input type="radio"/>
Edge nodes must be purged of all cache assets.	<input type="radio"/>	<input type="radio"/>
The compression type is supported.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

Front Door can dynamically compress content on the edge, resulting in a smaller and faster response to your clients. All files are eligible for compression. However, a file must be of a MIME type that is eligible for compression list.

Box 2: No

Sometimes you may wish to purge cached content from all edge nodes and force them all to retrieve new updated assets. This might be due to updates to your web application, or to quickly update assets that contain incorrect information.

Box 3: Yes

These profiles support the following compression encodings: Gzip (GNU zip), Brotli Reference:

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

NEW QUESTION 85

- (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	<input type="text" value="SAS tokens"/> <input type="text" value="Key authentication"/> <input type="text" value="JWT token"/>
Topic publishing	<input type="text" value="ValidationCode handshake"/> <input type="text" value="ValidationURL handshake"/> <input type="text" value="Management Access Control"/>

- 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 90

- (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 92

- (Exam Topic 3)

You are developing an application to use Azure Blob storage. You have configured Azure Blob storage to include change feeds.

A copy of your storage account must be created in another region. Data must be copied from the current storage account to the new storage account directly between the storage servers.

You need to create a copy of the storage account in another region and copy the data.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

- Use AZCopy to copy the data to the new storage account.
- Deploy the template to create a new storage account in the target region.
- Export a Resource Manager template.
- Create a new template deployment.
- Modify the template by changing the storage account name and region.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

To move a storage account, create a copy of your storage account in another region. Then, move your data to that account by using AzCopy, or another tool of your choice.

The steps are:

- > Export a template.
- > Modify the template by adding the target region and storage account name.
- > Deploy the template to create the new storage account.
- > Configure the new storage account.
- > Move data to the new storage account.
- > Delete the resources in the source region.

Note: You must enable the change feed on your storage account to begin capturing and recording changes. You can enable and disable changes by using Azure Resource Manager templates on Portal or Powershell.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move> <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed>

NEW QUESTION 96

- (Exam Topic 3)

You develop an ASP.NET Core MVC application. You configure the application to track webpages and custom events.

You need to identify trends in application usage.

Which Azure Application Insights Usage Analysis features should you use? To answer, drag the appropriate features to the correct requirements. Each feature 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.

Requirement	Feature
Which pages visited by users most often correlate to a product purchase?	<input type="text"/>
How does load time of the product display page affect a user's decision to purchase a product?	<input type="text"/>
Which events most influence a user's decision to continue to use the application?	<input type="text"/>
Are there places in the application that users often perform repetitive actions?	<input type="text"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box1: Users Box 2: Impact

One way to think of Impact is as the ultimate tool for settling arguments with someone on your team about how slowness in some aspect of your site is affecting whether users stick around. While users may tolerate a certain amount of slowness, Impact gives you insight into how best to balance optimization and performance to maximize user conversion.

Box 3: Retention

The retention feature in Azure Application Insights helps you analyze how many users return to your app, and how often they perform particular tasks or achieve goals. For example, if you run a game site, you could compare the numbers of users who return to the site after losing a game with the number who return after winning. This knowledge can help you improve both your user experience and your business strategy.

Box 4: User flows

The User Flows tool visualizes how users navigate between the pages and features of your site. It's great for answering questions like:

How do users navigate away from a page on your site? What do users click on a page on your site?

Where are the places that users churn most from your site?

Are there places where users repeat the same action over and over?

NEW QUESTION 98

- (Exam Topic 3)

Your company has several websites that use a company logo image. You use Azure Content Delivery Network (CDN) to store the static image.

You need to determine the correct process of how the CDN and the Point of Presence (POP) server will distribute the image and list the items in the correct order.

In which order do the actions occur? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.	
Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the files from cache if the TTL has not expired.	<div style="text-align: center;"> ⬅ ⬆ </div>
If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.	<div style="text-align: center;"> ➡ ⬇ </div>
The origin server returns the logo image to an edge server in the POP. An edge server in the POP caches the logo image and returns the image to the client.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: A user requests the image..

A user requests a file (also called an asset) by using a URL with a special domain name, such as <endpoint name>.azureedge.net. This name can be an endpoint hostname or a custom domain. The DNS routes the request to the best performing POP location, which is usually the POP that is geographically closest to the user.

Step 2: If no edge servers in the POP have the..

If no edge servers in the POP have the file in their cache, the POP requests the file from the origin server. The origin server can be an Azure Web App, Azure Cloud Service, Azure Storage account, or any publicly accessible web server.

Step 3: The origin server returns the..

The origin server returns the file to an edge server in the POP.

An edge server in the POP caches the file and returns the file to the original requestor (Alice). The file remains cached on the edge server in the POP until the time-to-live (TTL) specified by its HTTP headers expires. If the origin server didn't specify a TTL, the default TTL is seven days.

Step 4: Subsequent requests for..

Additional users can then request the same file by using the same URL that the original user used, and can also be directed to the same POP.

If the TTL for the file hasn't expired, the POP edge server returns the file directly from the cache. This process results in a faster, more responsive user experience.

References:

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

NEW QUESTION 99

- (Exam Topic 3)

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

Time grain statistic 1 minute time grain

Total

Maximum

Average

Count

Greater than

Greater than or equal to

Less than

Less than or equal to

Threshold

1000

Action

Operation

Increase count by

Increase count to

Decrease count by

Decrease count to

Instance count

1

Cool down (minutes)

5

- 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 101

- (Exam Topic 3)

You are configuring a development environment for your team. You deploy the latest Visual Studio image from the Azure Marketplace to your Azure subscription. The development environment requires several software development kits (SDKs) and third-party components to support application development across the organization. You install and customize the deployed virtual machine (VM) for your development team. The customized VM must be saved to allow provisioning of a new team member development environment.

You need to save the customized VM for future provisioning.

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

Answer Area

Action	Tool or service
Generalize the VM.	<ul style="list-style-type: none"> Azure PowerShell Visual Studio command prompt Azure Migrate Azure Backup
Store images.	<ul style="list-style-type: none"> Azure Blob Storage Azure Data Lake Storage Azure File Storage Azure Table Storage

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure Powershell

Creating an image directly from the VM ensures that the image includes all of the disks associated with the VM, including the OS disk and any data disks.

Before you begin, make sure that you have the latest version of the Azure PowerShell module. You use Sysprep to generalize the virtual machine, then use Azure PowerShell to create the image. Box 2: Azure Blob Storage

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource#create-an-image-of-a>

NEW QUESTION 106

- (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 107

- (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 111

- (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 115

- (Exam Topic 3)

You are developing a Docker/Go using Azure App Service Web App for Containers. You plan to run the container in an App Service on Linux. You identify a Docker container image to use.

None of your current resource groups reside in a location that supports Linux. You must minimize the number of resource groups required.

You need to create the application and perform an initial deployment.

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

Azure CLI Commands	Answer Area
az group create	
az group update	
az webapp update	⬅
az webapp create	➡
az appservice plan create	⬆
	⬆

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

You can host native Linux applications in the cloud by using Azure Web Apps. To create a Web App for Containers, you must run Azure CLI commands that create a group, then a service plan, and finally the web app itself.

Step 1: az group create

In the Cloud Shell, create a resource group with the az group create command. Step 2: az appservice plan create

In the Cloud Shell, create an App Service plan in the resource group with the az appservice plan create command.

Step 3: az webapp create

In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command. Don't forget to replace with a unique app name, and <docker-ID> with your Docker ID.

References:

<https://docs.microsoft.com/mt-mt/azure/app-service/containers/quickstart-docker-go?view=sql-server-ver15>

NEW QUESTION 119

- (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 124

- (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			
delegated			
profile			
application			
user_impersonation			
	API	Permission	Type
	Azure Storage	Setting	Setting
	Microsoft Graph	User.Read	Setting

- 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 129

- (Exam Topic 3)

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.

Answer Area

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	<input type="radio"/>	<input type="radio"/>
If the body size is >256k, an error will occur.	<input type="radio"/>	<input type="radio"/>
If the request is <code>http://contoso.com/api/9.2/</code> , the policy will retain the higher version.	<input type="radio"/>	<input type="radio"/>

- 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.

Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies>

NEW QUESTION 134

- (Exam Topic 3)

You are developing a ticket reservation system for an airline.

The storage solution for the application must meet the following requirements:

- > Ensure at least 99.99% availability and provide low latency.
- > Accept reservations event when localized network outages or other unforeseen failures occur.
- > Process reservations in the exact sequence as reservations are submitted to minimize overbooking or selling the same seat to multiple travelers.
- > Allow simultaneous and out-of-order reservations with a maximum five-second tolerance window. You provision a resource group named `airlineResourceGroup` in the Azure South-Central US region. You need to provision a SQL SPI Cosmos DB account to support the app.

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.

```
resourceGroupName- +airlineResourceGroup'
name- +docdb-airline-reservations'
databaseName- 'docdb-tickets-database'
collectionName- 'docdb-tickets-collection'
consistencyLevel-
```

▼

Strong

Eventual

ConsistentPrefix

BoundedStaleness

```
az cosmosdb create \
--name $name \
```

▼

--enable-virtual-network true\

--enable-automatic-failover true\

--kind 'GlobalDocumentDB' \

--kind 'MongoDB'\

```
--resource group $resourceGroupName \
--max interval 5 \
```

▼

--locations 'southcentralus'

--locations 'eastus'

--locations'southcentralus=0 eastus=1 westus=2'

--locations 'southcentralus=0'

```
--default-consistency-level - $consistencylevel
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: BoundedStaleness

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. In other words, 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 Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels> <https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/cosmos-db/manage-with-cli.md>

NEW QUESTION 136

- (Exam Topic 3)

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
<div style="border: 1px solid black; padding: 2px; width: fit-content;">Triggered</div>	Run on all instances that the web app runs on. Optionally restrict the WebJob to a single instance.	<div style="border: 1px solid black; width: 100px; height: 20px;"></div>
<div style="border: 1px solid black; padding: 2px; width: fit-content;">Continuous</div>	Run on a single instance that Azure select for load balancing.	<div style="border: 1px solid black; width: 100px; height: 20px;"></div>
	Supports remote debugging	<div style="border: 1px solid black; width: 100px; height: 20px;"></div>

- 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 137

- (Exam Topic 3)

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.

Answer Area

Setting	Action
Caching behavior	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Bypass cache</p> <p>Override</p> <p>Set if missing</p> </div> </div>
Cache expiration duration	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>1 second</p> <p>1 minute</p> <p>1 hour</p> <p>1 day</p> </div> </div>
Query string caching behavior	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Ignore query strings</p> <p>Bypass caching for query strings</p> <p>Cache every unique URL</p> </div> </div>

- 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.

Reference:

<https://docs.microsoft.com/en-us/azure/cdn/cdn-query-string>

NEW QUESTION 142

- (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 Azure Cache for Redis. Update the web applications.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 145

- (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 149

- (Exam Topic 3)

You develop a solution that uses an Azure SQL Database to store user information for a mobile app. The app stores sensitive information about users.

You need to hide sensitive information from developers that query the data for the mobile app.

Which three items must you identify when configuring dynamic data masking? Each correct answer presents a part of the solution.


NOTE: Each correct selection is worth one point.

- A. Column
- B. Table
- C. Trigger
- D. Index
- E. Schema

Answer: ABE

Explanation:

In the Dynamic Data Masking configuration page, you may see some database columns that the recommendations engine has flagged for masking. In order to accept the recommendations, just click Add Mask for one or more columns and a mask is created based on the default type for this column. You can change the masking function by clicking on the masking rule and editing the masking field format to a different format of your choice.



Dynamic Data Masking
demo_database

Save Discard Add Mask

Downlevel clients require the use of Security Enabled Connection Strings.

Masking Rules

MASK NAME	MASK FUNCTION
You haven't created any masking rules.	

SQL users excluded from masking (administrators are always excluded)

SQL users excluded from masking (administrators are always excluded)

Recommended fields to mask

SCHEMA	TABLE	COLUMN	
SalesLT	Customer	FirstName	ADD MASK
SalesLT	Customer	LastName	ADD MASK
SalesLT	Customer	EmailAddress	ADD MASK
SalesLT	Customer	Phone	ADD MASK
SalesLT	CustomerAddress	AddressID	ADD MASK

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started-portal>

NEW QUESTION 151

- (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 156

.....

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (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
  </
  <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 2

- (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 3

- (Exam 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

Answer Area

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.



- 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 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;"> ▼ </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;"> ▼ </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;"> ▼ </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	<div style="border: 1px solid black; padding: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</div> <div style="padding: 2px;">Basic</div> <div style="padding: 2px;">Standard</div> <div style="padding: 2px;">Premium</div> <div style="padding: 2px;">Isolated</div> </div>
Instance Count	<div style="border: 1px solid black; padding: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</div> <div style="padding: 2px;">1</div> <div style="padding: 2px;">10</div> <div style="padding: 2px;">20</div> <div style="padding: 2px;">100</div> </div>

- 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 support the requirements for the Shipping Logic App.

What should you use?

- A. Azure Active Directory Application Proxy
- B. Point-to-Site (P2S) VPN connection
- C. Site-to-Site (S2S) VPN connection
- D. On-premises Data Gateway

Answer: D

Explanation:

Before you can connect to on-premises data sources from Azure Logic Apps, download and install the on-premises data gateway on a local computer. The gateway works as a bridge that provides quick data transfer and encryption between data sources on premises (not in the cloud) and your logic apps.

The gateway supports BizTalk Server 2016.

Note: Microsoft have now fully incorporated the Azure BizTalk Services capabilities into Logic Apps and Azure App Service Hybrid Connections.

Logic Apps Enterprise Integration pack bring some of the enterprise B2B capabilities like AS2 and X12, EDI standards support

Scenario: The Shipping Logic app must meet the following requirements:

- > Support the ocean transport and inland transport workflows by using a Logic App.
- > Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
- > Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
- > Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Reference:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-gateway-install>

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

	▼
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 8

- (Exam Topic 3)

You use Azure Table storage to store customer information for an application. The data contains customer details and is partitioned by last name. You need to create a query that returns all customers with the last name Smith. Which code segment should you use?

- A. `TableQuery.GenerateFilterCondition("PartitionKey", Equals, "Smith")`
- B. `TableQuery.GenerateFilterCondition("LastName", Equals, "Smith")`
- C. `TableQuery.GenerateFilterCondition("PartitionKey", QueryComparisons.Equal, "Smith")`
- D. `TableQuery.GenerateFilterCondition("LastName", QueryComparisons.Equal, "Smith")`

Answer: C

Explanation:

Retrieve all entities in a partition. The following code example specifies a filter for entities where 'Smith' is the partition key. This example prints the fields of each entity in the query results to the console.

Construct the query operation for all customer entities where PartitionKey="Smith". `TableQuery<CustomerEntity> query = new`

`TableQuery<CustomerEntity>().Where(TableQuery.GenerateFilterCondition("PartitionKey",`

`QueryComparisons.Equal, "Smith"));`

References:

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

NEW QUESTION 9

- (Exam Topic 3)

You are developing an Azure Function App by using Visual Studio. The app will process orders input by an Azure Web App. The web app places the order information into Azure Queue Storage.

You need to review the Azure Function App code shown below. NOTE: Each correct selection is worth one point.

```
public static class OrderProcessor
{
    [FunctionName("ProcessOrders")]
    public static void ProcessOrders([QueueTrigger("incoming-orders")]IInkQueueMessage myQueueItem, [Table
("Orders")]ICollector<Order> tableBindings,
    TraceWriter log)
    {
        log.Info($"Processing Order: {myQueueItem.Id}");
        log.Info($"Queue Insertion Time: {myQueueItem.InsertionTime}");
        log.Info($"Queue Expiration Time: {myQueueItem.ExpirationTime}");
        tableBindings.Add(new Order().DeserializeObject<Order>(myQueueItem.AsString));
    }
    [FunctionName("ProcessOrders-Poison")]
    public static void ProcessFailedOrders([QueueTrigger("incoming-orders-poison")]IInkQueueMessage
myQueueItem, TraceWriter log)
    {
        log.Error($"Failed to process order: {myQueueItem.AsString}");
    }
}
```

Yes **No**

- The code will log the time that the order was processed from the queue. Yes No
- When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try. Yes No
- When there are multiple orders in the queue, a batch of orders will be retrieved from the queue and the ProcessOrders function will run multiple instances concurrently to process the orders. Yes No
- The ProcessOrders function will output the order to an Orders table in Azure Table Storage. Yes No

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

ExpirationTime - The time that the message expires. InsertionTime - The time that the message was added to the queue.

Box 2: Yes

maxDequeueCount - The number of times to try processing a message before moving it to the poison queue.

Default value is 5.

Box 3: Yes

When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them.

By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.

Box 4: Yes References:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue>

NEW QUESTION 10

- (Exam Topic 3)

A company is developing a Java web app. The web app code is hosted in a GitHub repository located at <https://github.com/Contoso/webapp>.

The web app must be evaluated before it is moved to production. You must deploy the initial code release to a deployment slot named staging.

You need to create the web app and deploy the code.

How should you complete the commands? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```
gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup

az [ ] create --location centralus - -name $resourcegroupname
group
webapp
appservice plan
webapp deployment slot
webapp deployment source
create --name $webappname - -resource-group $resourcegroupname
- -sku S3
create --name $webappname - -resource-group $resourcegroupname
\ - -plan $webappname
create --name $webappname - -resource-group $resourcegroupname
\ - -slot staging

az [ ] config - -name $webappname - -resource-group $resourcegroupname
\ - -slot staging - -repo-url
$gitrepo - -branch master - -manual-integration

az [ ]

az [ ]

az [ ]

az [ ]
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: group
 # Create a resource group.
 az group create --location westeurope --name myResourceGroup
 Box 2: appservice plan
 # Create an App Service plan in STANDARD tier (minimum required by deployment slots). az appservice plan create --name \$webappname --resource-group myResourceGroup --sku S1
 Box 3: webapp
 # Create a web app.
 az webapp create --name \$webappname --resource-group myResourceGroup \ --plan \$webappname
 Box 4: webapp deployment slot
 #Create a deployment slot with the name "staging".
 az webapp deployment slot create --name \$webappname --resource-group myResourceGroup \ --slot staging
 Box 5: webapp deployment source
 # Deploy sample code to "staging" slot from GitHub.
 az webapp deployment source config --name \$webappname --resource-group myResourceGroup \ --slot staging --repo-url \$gitrepo --branch master --manual-integration
 References:
<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-staging-environment>

NEW QUESTION 10

- (Exam Topic 3)

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 14

- (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 17

- (Exam Topic 3)

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.

Uploading of videos occurs on an irregular basis.

You need to copy specific blobs from Container1 to Container2 in real time when specific requirements are met, excluding backup blob copies.

What should you do?

- A. Download the blob to a virtual machine and then upload the blob to Container2.
- B. Run the Azure PowerShell command Start-AzureStorageBlobCopy.
- C. Copy blobs to Container2 by using the Put Blob operation of the Blob Service REST API.
- D. Use AzCopy with the Snapshot switch blobs 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.

References:

<https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azuremp>

NEW QUESTION 22

- (Exam Topic 3)

You are a developer for a SaaS company that offers many web services. All web services for the company must meet the following requirements:

- > Use API Management to access the services
- > Use OpenID Connect for authentication
- > Prevent anonymous usage

A recent security audit found that several web services can be called without any authentication. Which API Management policy should you implement?

- A. jsonp
- B. authentication-certificate
- C. check-header
- D. validate-jwt

Answer: D

Explanation:

Add the validate-jwt policy to validate the OAuth token for every incoming request. Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad>

NEW QUESTION 23

- (Exam Topic 3)

You are developing a .NET Core model-view controller (MVC) application hosted on Azure for a health care system that allows providers access to their information.

You develop the following code:

```
services.AddAuthorization (options =>
{
    options.AddPolicy ("ProviderPartner", policy =>
    {
        .policy.AddAuthenticationSchemes ("Cookie, Bearer");
        policy.RequireAuthenticatedUser ();
        policy.RequireRole ("ProviderAdmin", "SysAdmin");
        policy.RequireClaim ("editor", "partner");
    });
});
```

You define a role named SysAdmin.

You need to ensure that the application meets the following authorization requirements:

- > Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.
- > Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.

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.

```
services.AddAuthorization (options =>
{
    options.AddPolicy ("ProviderPartner", .policy.AddAut
    {
        .policy.AddAuthenticationSchemes ("Cookie, Bearer");
        policy.RequireAuthenticatedUser ();
        policy.RequireRole ("ProviderAdmin" });
        policy.RequireClaim ("editor", "pa
    });
})
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:
 Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.
 Box 2:
 Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.

NEW QUESTION 28

- (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.

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 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 29

- (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 32

- (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 33

- (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 38

- (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	2018-08-25T15:47:29.135Z	32	A box of chocolate candy bars	Choco-bar	5	3	123421
Hardware	2	2018-08-25T15:46:08.409Z	2	A bag of bolts	Bolts	1	4	67894
Hardware	1	2018-08-25T15:46:41.402Z	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 43

- (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 =
    NotificationHubClient.CreateClientFromConnectionString(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(windowsToastPayload);
}
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 45

- (Exam Topic 3)

You are developing an ASP.NET Core Web API web service. The web service uses Azure Application Insights for all telemetry and dependency tracking. The web service reads and writes data to a database other than Microsoft SQL Server.

You need to ensure that dependency tracking works for calls to the third-party database.

Which two Dependency Telemetry properties should you store in the database? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Telemetry.Context.Operation.Id
- B. Telemetry.Context.Cloud.RoleInstance
- C. Telemetry.Id

D. Telemetry.ContextSession.Id
 E. Telemetry.Name

Answer: AC

Explanation:

References:
<https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking>

NEW QUESTION 49

- (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",
  [Box 1] : "All",
  [Box 2] : true
  ...
}
    
```

Box 1 options: "optionalClaims", "groupMembershipClaims"

Box 2 options: "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 54

- (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 58

- (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 59

- (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. [dropdown] ();

        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. [dropdown] ));

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

            task [dropdown] =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 63

- (Exam Topic 3)

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

Answer Area

Policy section	Policy	Policy section
	Set-variable	<input type="text"/>
<input type="text" value="Inbound"/>	Cache-lookup-value	<input type="text"/>
<input type="text" value="Outbound"/>	Cache-store-value	<input type="text"/>
	Find-and-replace	<input type="text"/>

- 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)" />
</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>
```

Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>

NEW QUESTION 66

- (Exam Topic 3)

You are developing an application that uses Azure Storage Queues. You have the following code:

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse
(CloudConfigurationManager.GetSetting("StorageConnectionString"));
CloudQueueClient queueClient = storageAccount.CreateCloudQueueClient()

CloudQueue queue = queueClient.GetQueueReference("appqueue") ;
await queue.CreateIfNotExistsAsync() ;

CloudQueueMessage peekedMessage = await queue.PeekMessageAsync() ;
if (peekedMessage != null)
{
    Console.WriteLine("The peeked message is: {0}", peekedMessage.AsString);
}
CloudQueueMessage message = await queue.GetMessageAsync() ;
```

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

NOTE: Each correct selection is worth one point.

Answer Area

Statement	Yes	No
The code configures the lock duration for the queue.	<input type="radio"/>	<input type="radio"/>
The last message read remains in the queue after the code runs.	<input type="radio"/>	<input type="radio"/>
The storage queue remains in the storage account after the code runs.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

The QueueDescription.LockDuration property gets or sets the duration of a peek lock; that is, the amount of time that the message is locked for other receivers. The maximum value for LockDuration is 5 minutes; the default value is 1 minute.

Box 2: Yes

You can peek at the message in the front of a queue without removing it from the queue by calling the PeekMessage method.

Box 3: Yes Reference:

<https://docs.microsoft.com/en-us/azure/storage/queues/storage-dotnet-how-to-use-queues> <https://docs.microsoft.com/en-us/dotnet/api/microsoft.servicebus.messaging.queuedescription.lockduration>

NEW QUESTION 67

- (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 69

- (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 70

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

- 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.computedcloud.com/enabling-application-insights-trace-logging-in-asp-net-core/>

NEW QUESTION 75

- (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 80

- (Exam Topic 3)

You are using Azure Front Door Service.

You are expecting inbound files to be compressed by using Brotli compression. You discover that inbound XML files are not compressed. The files are 9 megabytes (MB) in size.

You need to determine the root cause for the issue.

To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Statement	Yes	No
The file MIME type is supported by the service.	<input type="radio"/>	<input type="radio"/>
Edge nodes must be purged of all cache assets.	<input type="radio"/>	<input type="radio"/>
The compression type is supported.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

Front Door can dynamically compress content on the edge, resulting in a smaller and faster response to your clients. All files are eligible for compression. However, a file must be of a MIME type that is eligible for compression list.

Box 2: No

Sometimes you may wish to purge cached content from all edge nodes and force them all to retrieve new updated assets. This might be due to updates to your web application, or to quickly update assets that contain incorrect information.

Box 3: Yes

These profiles support the following compression encodings: Gzip (GNU zip), Brotli Reference:

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

NEW QUESTION 85

- (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	<input type="text" value="SAS tokens"/> <input type="text" value="Key authentication"/> <input type="text" value="JWT token"/>
Topic publishing	<input type="text" value="ValidationCode handshake"/> <input type="text" value="ValidationURL handshake"/> <input type="text" value="Management Access Control"/>

- 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 90

- (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 92

- (Exam Topic 3)

You are developing an application to use Azure Blob storage. You have configured Azure Blob storage to include change feeds.

A copy of your storage account must be created in another region. Data must be copied from the current storage account to the new storage account directly between the storage servers.

You need to create a copy of the storage account in another region and copy the data.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

- Use AZCopy to copy the data to the new storage account.
- Deploy the template to create a new storage account in the target region.
- Export a Resource Manager template.
- Create a new template deployment.
- Modify the template by changing the storage account name and region.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

To move a storage account, create a copy of your storage account in another region. Then, move your data to that account by using AzCopy, or another tool of your choice.

The steps are:

- > Export a template.
- > Modify the template by adding the target region and storage account name.
- > Deploy the template to create the new storage account.
- > Configure the new storage account.
- > Move data to the new storage account.
- > Delete the resources in the source region.

Note: You must enable the change feed on your storage account to begin capturing and recording changes. You can enable and disable changes by using Azure Resource Manager templates on Portal or Powershell.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move> <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed>

NEW QUESTION 96

- (Exam Topic 3)

You develop an ASP.NET Core MVC application. You configure the application to track webpages and custom events.

You need to identify trends in application usage.

Which Azure Application Insights Usage Analysis features should you use? To answer, drag the appropriate features to the correct requirements. Each feature 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.

Requirement	Feature
Which pages visited by users most often correlate to a product purchase?	<input type="text"/>
How does load time of the product display page affect a user's decision to purchase a product?	<input type="text"/>
Which events most influence a user's decision to continue to use the application?	<input type="text"/>
Are there places in the application that users often perform repetitive actions?	<input type="text"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box1: Users Box 2: Impact

One way to think of Impact is as the ultimate tool for settling arguments with someone on your team about how slowness in some aspect of your site is affecting whether users stick around. While users may tolerate a certain amount of slowness, Impact gives you insight into how best to balance optimization and performance to maximize user conversion.

Box 3: Retention

The retention feature in Azure Application Insights helps you analyze how many users return to your app, and how often they perform particular tasks or achieve goals. For example, if you run a game site, you could compare the numbers of users who return to the site after losing a game with the number who return after winning. This knowledge can help you improve both your user experience and your business strategy.

Box 4: User flows

The User Flows tool visualizes how users navigate between the pages and features of your site. It's great for answering questions like:

How do users navigate away from a page on your site? What do users click on a page on your site?

Where are the places that users churn most from your site?

Are there places where users repeat the same action over and over?

NEW QUESTION 98

- (Exam Topic 3)

Your company has several websites that use a company logo image. You use Azure Content Delivery Network (CDN) to store the static image.

You need to determine the correct process of how the CDN and the Point of Presence (POP) server will distribute the image and list the items in the correct order.

In which order do the actions occur? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.	
Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the files from cache if the TTL has not expired.	<div style="text-align: center;"> ⬅ ⬆ </div>
If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.	<div style="text-align: center;"> ➡ ⬇ </div>
The origin server returns the logo image to an edge server in the POP. An edge server in the POP caches the logo image and returns the image to the client.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: A user requests the image..

A user requests a file (also called an asset) by using a URL with a special domain name, such as <endpoint name>.azureedge.net. This name can be an endpoint hostname or a custom domain. The DNS routes the request to the best performing POP location, which is usually the POP that is geographically closest to the user.

Step 2: If no edge servers in the POP have the..

If no edge servers in the POP have the file in their cache, the POP requests the file from the origin server. The origin server can be an Azure Web App, Azure Cloud Service, Azure Storage account, or any publicly accessible web server.

Step 3: The origin server returns the..

The origin server returns the file to an edge server in the POP.

An edge server in the POP caches the file and returns the file to the original requestor (Alice). The file remains cached on the edge server in the POP until the time-to-live (TTL) specified by its HTTP headers expires. If the origin server didn't specify a TTL, the default TTL is seven days.

Step 4: Subsequent requests for..

Additional users can then request the same file by using the same URL that the original user used, and can also be directed to the same POP.

If the TTL for the file hasn't expired, the POP edge server returns the file directly from the cache. This process results in a faster, more responsive user experience.

References:

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

NEW QUESTION 99

- (Exam Topic 3)

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

Time grain statistic 1 minute time grain

Total

Maximum

Average

Count

Greater than

Greater than or equal to

Less than

Less than or equal to

Threshold

1000

Action

Operation

Increase count by

Increase count to

Decrease count by

Decrease count to

Instance count

1

Cool down (minutes)

5

- 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 101

- (Exam Topic 3)

You are configuring a development environment for your team. You deploy the latest Visual Studio image from the Azure Marketplace to your Azure subscription. The development environment requires several software development kits (SDKs) and third-party components to support application development across the organization. You install and customize the deployed virtual machine (VM) for your development team. The customized VM must be saved to allow provisioning of a new team member development environment.

You need to save the customized VM for future provisioning.

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

Answer Area

Action	Tool or service
Generalize the VM.	<input type="checkbox"/> Azure PowerShell <input type="checkbox"/> Visual Studio command prompt <input type="checkbox"/> Azure Migrate <input type="checkbox"/> Azure Backup
Store images.	<input type="checkbox"/> Azure Blob Storage <input type="checkbox"/> Azure Data Lake Storage <input type="checkbox"/> Azure File Storage <input type="checkbox"/> Azure Table Storage

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure Powershell

Creating an image directly from the VM ensures that the image includes all of the disks associated with the VM, including the OS disk and any data disks.

Before you begin, make sure that you have the latest version of the Azure PowerShell module. You use Sysprep to generalize the virtual machine, then use Azure PowerShell to create the image. Box 2: Azure Blob Storage

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource#create-an-image-of-a>

NEW QUESTION 106

- (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 107

- (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 111

- (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 115

- (Exam Topic 3)

You are developing a Docker/Go using Azure App Service Web App for Containers. You plan to run the container in an App Service on Linux. You identify a Docker container image to use.

None of your current resource groups reside in a location that supports Linux. You must minimize the number of resource groups required.

You need to create the application and perform an initial deployment.

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

Azure CLI Commands	Answer Area
az group create	
az group update	
az webapp update	⬅
az webapp create	➡
az appservice plan create	⬆
	⬆

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

You can host native Linux applications in the cloud by using Azure Web Apps. To create a Web App for Containers, you must run Azure CLI commands that create a group, then a service plan, and finally the web app itself.

Step 1: az group create

In the Cloud Shell, create a resource group with the az group create command. Step 2: az appservice plan create

In the Cloud Shell, create an App Service plan in the resource group with the az appservice plan create command.

Step 3: az webapp create

In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command. Don't forget to replace with a unique app name, and <docker-ID> with your Docker ID.

References:

<https://docs.microsoft.com/mt-mt/azure/app-service/containers/quickstart-docker-go?view=sql-server-ver15>

NEW QUESTION 119

- (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 124

- (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			
delegated			
profile			
application			
user_impersonation			
	API	Permission	Type
	Azure Storage	Setting	Setting
	Microsoft Graph	User.Read	Setting

- 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 129

- (Exam Topic 3)

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.

Answer Area

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	<input type="radio"/>	<input type="radio"/>
If the body size is >256k, an error will occur.	<input type="radio"/>	<input type="radio"/>
If the request is <code>http://contoso.com/api/9.2/</code> , the policy will retain the higher version.	<input type="radio"/>	<input type="radio"/>

- 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.

Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies>

NEW QUESTION 134

- (Exam Topic 3)

You are developing a ticket reservation system for an airline.

The storage solution for the application must meet the following requirements:

- > Ensure at least 99.99% availability and provide low latency.
- > Accept reservations event when localized network outages or other unforeseen failures occur.
- > Process reservations in the exact sequence as reservations are submitted to minimize overbooking or selling the same seat to multiple travelers.
- > Allow simultaneous and out-of-order reservations with a maximum five-second tolerance window. You provision a resource group named `airlineResourceGroup` in the Azure South-Central US region. You need to provision a SQL SPI Cosmos DB account to support the app.

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.

```
resourceGroupName- +airlineResourceGroup'
name- +docdb-airline-reservations'
databaseName- 'docdb-tickets-database'
collectionName- 'docdb-tickets-collection'
consistencyLevel-
```

▼

Strong

Eventual

ConsistentPrefix

BoundedStaleness

```
az cosmosdb create \
--name $name \
```

▼

--enable-virtual-network true\

--enable-automatic-failover true\

--kind 'GlobalDocumentDB' \

--kind 'MongoDB'\

```
--resource group $resourceGroupName \
--max interval 5 \
```

▼

--locations 'southcentralus'

--locations 'eastus'

--locations'southcentralus=0 eastus=1 westus=2'

--locations 'southcentralus=0'

```
--default-consistency-level - $consistencylevel
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: BoundedStaleness

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. In other words, 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 Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels> <https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/cosmos-db/manage-with-cli.md>

NEW QUESTION 136

- (Exam Topic 3)

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.	<input type="text"/>
Continuous	Run on a single instance that Azure select for load balancing.	<input type="text"/>
	Supports remote debugging	<input type="text"/>

- 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 137

- (Exam Topic 3)

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.

Answer Area

Setting	Action
Caching behavior	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Bypass cache</p> <p>Override</p> <p>Set if missing</p> </div> </div>
Cache expiration duration	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>1 second</p> <p>1 minute</p> <p>1 hour</p> <p>1 day</p> </div> </div>
Query string caching behavior	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Ignore query strings</p> <p>Bypass caching for query strings</p> <p>Cache every unique URL</p> </div> </div>

- 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.

Reference:

<https://docs.microsoft.com/en-us/azure/cdn/cdn-query-string>

NEW QUESTION 142

- (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 Azure Cache for Redis. Update the web applications.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 145

- (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 149

- (Exam Topic 3)

You develop a solution that uses an Azure SQL Database to store user information for a mobile app. The app stores sensitive information about users.

You need to hide sensitive information from developers that query the data for the mobile app.

Which three items must you identify when configuring dynamic data masking? Each correct answer presents a part of the solution.


NOTE: Each correct selection is worth one point.

- A. Column
- B. Table
- C. Trigger
- D. Index
- E. Schema

Answer: ABE

Explanation:

In the Dynamic Data Masking configuration page, you may see some database columns that the recommendations engine has flagged for masking. In order to accept the recommendations, just click Add Mask for one or more columns and a mask is created based on the default type for this column. You can change the masking function by clicking on the masking rule and editing the masking field format to a different format of your choice.



SCHEMA	TABLE	COLUMN	
SalesLT	Customer	FirstName	ADD MASK
SalesLT	Customer	LastName	ADD MASK
SalesLT	Customer	EmailAddress	ADD MASK
SalesLT	Customer	Phone	ADD MASK
SalesLT	CustomerAddress	AddressID	ADD MASK

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started-portal>

NEW QUESTION 151

- (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 156

.....

Relate Links

100% Pass Your AZ-204 Exam with ExamBible Prep Materials

<https://www.exambible.com/AZ-204-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>