

AZ-204 Dumps

Developing Solutions for Microsoft Azure (beta)

<https://www.certleader.com/AZ-204-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

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

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

References:

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

NEW QUESTION 2

- (Exam Topic 1)

You need to support the message processing for the ocean transport workflow.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Create an integration account in the Azure portal.

Link the custom connector to the Logic App.

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

Create a custom connector for the Logic App.

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

Link the Logic App to the integration account.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Create an integration account in the Azure portal

You can define custom metadata for artifacts in integration accounts and get that metadata during runtime for your logic app to use. For example, you can provide metadata for artifacts, such as partners, agreements, schemas, and maps - all store metadata using key-value pairs.

Step 2: Link the Logic App to the integration account

A logic app that's linked to the integration account and artifact metadata you want to use. Step 3: Add partners, schemas, certificates, maps, and agreements

Step 4: Create a custom connector for the Logic App. References:

<https://docs.microsoft.com/bs-latn-ba/azure/logic-apps/logic-apps-enterprise-integration-metadata>

NEW QUESTION 3

- (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>▼</div> <div>Basic</div> <div>Standard</div> <div>Premium</div> <div>Isolated</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 4

- (Exam Topic 3)

You develop a web application.

You need to register the application with an active Azure Active Directory (Azure AD) tenant.

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

Actions	Answer Area
Select Manifest from the middle-tier service registration.	
In Enterprise Applications, select New application .	
Add a Cryptographic key.	
Create a new application and provide the name, account type, and redirect URL	⬅️
Select the Azure AD instance.	➡️
Use an access token to access the secure resource.	⬆️
In App Registrations, select New registration .	⬇️

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Register a new application using the Azure portal

- Sign in to the Azure portal using either a work or school account or a personal Microsoft account.
- If your account gives you access to more than one tenant, select your account in the upper right corner.

Set your portal session to the Azure AD tenant that you want.

- Search for and select Azure Active Directory. Under Manage, select App registrations.
- Select New registration. (Step 1)
- In Register an application, enter a meaningful application name to display to users.
- Specify who can use the application. Select the Azure AD instance. (Step 2)
- Under Redirect URI (optional), select the type of app you're building: Web or Public client (mobile & desktop). Then enter the redirect URI, or reply URL, for your application. (Step 3)
- When finished, select Register.

NEW QUESTION 5

- (Exam Topic 3)

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

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

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

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

Solution: Configure the Azure Web App for the website to allow only authenticated requests and require Azure AD log on.

Does the solution meet the goal?

- A. Yes
B. No

Answer: B

Explanation:

Instead in the Azure AD application's manifest, set value of the groupMembershipClaims option to All. References:

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

NEW QUESTION 6

- (Exam Topic 3)

You must ensure that the external party cannot access the data in the SSN column of the Person table.

Will each protection method meet the requirement? To answer, drag the appropriate responses to the correct protection methods. Each response may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

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

Box 2: No

Box 3: Yes

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

Box 4: No

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

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

NEW QUESTION 7

- (Exam Topic 3)

You are developing a microservices solution. You plan to deploy the solution to a multinode Azure Kubernetes Service (AKS) cluster.

You need to deploy a solution that includes the following features:

- reverse proxy capabilities
- configurable traffic routing
- TLS termination with a custom certificate

Which components should you use? To answer, drag the appropriate components to the correct requirements. Each component 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

Components	Action	Component
Helm	Deploy solution.	
Draft		
Brigade	View cluster and external IP addressing.	
KubeCtl		
Ingress Controller	Implement a single, public IP endpoint that is routed to multiple microservices.	
CoreDNS		
Virtual Kubelet		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Helm

To create the ingress controller, use Helm to install nginx-ingress. Box 2: kubectl

To find the cluster IP address of a Kubernetes pod, use the kubectl get pod command on your local machine, with the option -o wide .

Box 3: Ingress Controller

An ingress controller is a piece of software that provides reverse proxy, configurable traffic routing, and TLS termination for Kubernetes services. Kubernetes ingress resources are used to configure the ingress rules and routes for individual Kubernetes services.

Reference:

<https://docs.microsoft.com/bs-cyrl-ba/azure/aks/ingress-basic> <https://www.digitalocean.com/community/tutorials/how-to-inspect-kubernetes-networking>

NEW QUESTION 8

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

- 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 9

- (Exam Topic 3)

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

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

Margie's Travel is an international travel and bookings management service. The company is expanding into restaurant bookings. You are tasked with implementing Azure Search for the restaurants listed in their solution.

You create the index in Azure Search.

You need to import the restaurant data into the Azure Search service by using the Azure Search .NET SDK. Solution:

- * 1. Create a SearchIndexClient object to connect to the search index.
- * 2. Create a DataContainer that contains the documents which must be added.
- * 3. Create a DataSource instance and set its Container property to the DataContainer.
- * 4. Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Use the following method:

- * 1. - Create a SearchIndexClient object to connect to the search index
- * 2. - Create an IndexBatch that contains the documents which must be added.
- * 3. - Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch. References:

<https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk>

NEW QUESTION 10

- (Exam Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You are developing and deploying several ASP.Net web applications to Azure App Service. You plan to save session state information and HTML output. You must use a storage mechanism with the following requirements:

- Share session state across all ASP.NET web applications
- Support controlled, concurrent access to the same session state data for multiple readers and a single writer
- Save full HTTP responses for concurrent requests

You need to store the information.

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

Does the solution meet the goal?

- A. Yes

B. No

Answer: B

Explanation:

Instead use Azure Cache for Redis.

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

References:

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

NEW QUESTION 10

- (Exam Topic 3)

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

Spikes in traffic have caused increases in page load times.

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

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

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

Actions		Answer Area
Configure the web app to the Premium App Service tier.		
Configure the web app to the Standard App Service tier.		
Enable autoscaling on the web-app.	⬅	⬆
Add a Scale rule.	➡	⬇
Switch to an Azure App Services consumption plan.		
Configure a Scale condition.		

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Step 1: Configure the web app to the Standard App Service Tier

The Standard tier supports auto-scaling, and we should minimize the cost. Step 2: Enable autoscaling on the web app

First enable autoscale Step 3: Add a scale rule

Step 4: Add a Scale condition Reference:

<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-autoscale-get-started>

NEW QUESTION 13

- (Exam Topic 3)

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 Service Bus. Configure a topic to receive the device data by using a correlation filter.

Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:

A message is raw data produced by a service to be consumed or stored elsewhere. The Service Bus is for high-value enterprise messaging, and is used for order processing and financial transactions.

Reference:

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

NEW QUESTION 17

- (Exam Topic 3)

Your company is migrating applications to Azure. The IT department must allow internal developers to communicate with Microsoft support.

The service agents of the IT department must only have view resources and create support ticket permissions to all subscriptions. A new custom role must be created by reusing a default role definition and changing the permissions.

You need to create the custom role.

To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Item	Value
Powershell command	<div><div></div><div>Get-AzureRmRoleDefinition-Name"Reader" ConvertTo-Json Out-File C:\SupportRole.json Get-AzureRmRoleDefinition-Name"Operator" ConvertTo-Json Out-File C:\SupportRole.json Set-AzureRmRoleDefinition-Name"Reader" Input-File C:\SupportRole.json Set-AzureRmRoleDefinition Input-File C:\SupportRole.json</div></div>
Actions section	<div><div></div><div>"*/read*", "Microsoft.Support/*" "*/read" "*/read*", "Microsoft.Support/*" "*/read"</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Set-AzureRmRoleDefinition Input-File C:\SupportRole.json

The Set-AzureRmRoleDefinition cmdlet updates an existing custom role in Azure Role-Based Access Control. Provide the updated role definition as an input to the command as a JSON file or a PSRoleDefinition object.

The role definition for the updated custom role MUST contain the Id and all other required properties of the role even if they are not updated: DisplayName, Description, Actions, AssignableScope

Box 2: "*/read*", "Microsoft.Support/*" Microsoft.Support/* Create and manage support tickets "Microsoft.Support" role definition azure

NEW QUESTION 22

- (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 24

- (Exam Topic 3)

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

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

You need to select a partition key.

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

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

Answer: AC

Explanation:

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

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

References:

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

NEW QUESTION 26

- (Exam Topic 3)

You are implementing a software as a service (SaaS) ASP.NET Core web service that will run as an Azure Web App. The web service will use an on-premises SQL Server database for storage. The web service also includes a WebJob that processes data updates. Four customers will use the web service.

- Each instance of the WebJob processes data for a single customer and must run as a singleton instance.
- Each deployment must be tested by using deployment slots prior to serving production data.
- Azure costs must be minimized.
- Azure resources must be located in an isolated network. You need to configure the App Service plan for the Web App.

How should you configure the App Service plan? To answer, select the appropriate settings in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

App service plan setting

Value

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Number of VM instances: 4

You are not charged extra for deployment slots. Pricing tier: Isolated

The App Service Environment (ASE) is a powerful feature offering of the Azure App Service that gives network isolation and improved scale capabilities. It is essentially a deployment of the Azure App Service into a subnet of a customer's Azure Virtual Network (VNet).

References:


<https://azure.microsoft.com/sv-se/blog/announcing-app-service-isolated-more-power-scale-and-ease-of-use/>

NEW QUESTION 30

- (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	2018-08-25T15:46:08.405Z	2	A bag of bolts	Bolts	1	4	678564
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 34

- (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",
  "optionalClaims"
  "groupMembershipClaims"
  [Box 2] : true
  "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 35

- (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 38

- (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>
```

Box 2: Inbound

A cache-lookup-value policy Example:

```
<inbound>
<base />
<cache-lookup vary-by-developer="true | false" vary-by-developer-groups="true | false" downstream-caching-type="none | private | public" must-revalidate="true | false">
<vary-by-query-parameter>parameter name</vary-by-query-parameter> <!-- optional, can repeated several times -->
</cache-lookup>
</inbound>
```

Box 3: Outbound

A cache-store-value policy. Example:

```
<outbound>
<base />
<cache-store duration="3600" />
</outbound>
```

Box 4: Outbound

A find-and-replace policy to update the response body with the user profile information. Example:

```
<outbound>
<!-- Update response body with user profile-->
<find-and-replace from="$userprofile$"
to="@((string)context.Variables["userprofile"])" />
<base />
</outbound>
```

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 40

- (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 42

- (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 43

- (Exam Topic 3)

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

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

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

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

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

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

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

Code for CosmosDB, example:

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

```
// Create the table client.
```

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

```
// Retrieve a reference to the table.
```

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

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

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

Box 3: No

No partition key is used. Box 4: Yes

References:

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

NEW QUESTION 46

- (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 47

- (Exam Topic 3)

You are developing a mobile instant messaging app for a company. The mobile app must meet the following requirements:

- Support offline data sync.
- Update the latest messages during normal sync cycles. You need to implement Offline Data Sync.

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

- A. Retrieve records from Offline Data Sync on every call to the PullAsync method.
B. Retrieve records from Offline Data Sync using an Incremental Sync.
C. Push records to Offline Data Sync using an Incremental Sync.
D. Return the updatedAt column from the Mobile Service Backend and implement sorting by using the column.
E. Return the updatedAt column from the Mobile Service Backend and implement sorting by the message id.

Answer: BE

Explanation:

B: Incremental Sync: the first parameter to the pull operation is a query name that is used only on the client. If you use a non-null query name, the Azure Mobile SDK performs an incremental sync. Each time a pull operation returns a set of results, the latest updatedAt timestamp from that result set is stored in the SDK local system tables. Subsequent pull operations retrieve only records after that timestamp.

E (not D): To use incremental sync, your server must return meaningful updatedAt values and must also support sorting by this field. However, since the SDK adds its own sort on the updatedAt field, you cannot use a pull query that has its own orderBy clause.

References:

<https://docs.microsoft.com/en-us/azure/app-service-mobile/app-service-mobile-offline-data-sync>

NEW QUESTION 51

- (Exam Topic 3)

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

Fourth Coffee is migrating this application to Azure.

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

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

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

Azure CLI commands

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

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

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

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

Answer area

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Step 1: #/bin/bash
The appName is used when the webapp-name is created in step 2. Step 2: az webapp config hostname add
The webapp-name is used when the webapp is created in step 3. Step 3: az webapp create
Create a web app. In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command.
Step : az webapp config container set
In Create a web app, you specified an image on Docker Hub in the az webapp create command. This is good enough for a public image. To use a private image, you need to configure your Docker account ID and password in your Azure web app.
In the Cloud Shell, follow the az webapp create command with az webapp config container set.
References:
<https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image>

NEW QUESTION 56

- (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?	<div></div>
How does load time of the product display page affect a user's decision to purchase a product?	<div></div>
Which events most influence a user's decision to continue to use the application?	<div></div>
Are there places in the application that users often perform repetitive actions?	<div></div>

The Leader of IT Certification

visit - <https://www.certleader.com>

- 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 60

- (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 Hub. Configure the machine identifier as the partition key and enable capture.

- A. Yes
- B. No

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-programming-guide>

NEW QUESTION 62

- (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.	⬅
If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.	➡
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 65

- (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 70

- (Exam Topic 3)

You are developing a .NET Core MVC application for customers to research hotels. The application will use Azure Search. The application will search the index by using various criteria to locate documents related to hotels. The index will include search fields for rate, a list of amenities, and distance to the nearest airport.

The application must support the following scenarios for specifying search criteria and organizing results:

- Search the index by using regular expressions.
- Organize results by counts for name-value pairs.
- List hotels within a specified distance to an airport and that fall within a specific price range. You need to configure the SearchParameters class.

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

Scenario	Property
Search the index by using regular expressions.	<input type="checkbox"/> QueryType <input type="checkbox"/> OrderBy <input type="checkbox"/> SearchMode
Organize results by counts for name-value pairs.	<input type="checkbox"/> Facets <input type="checkbox"/> Filter <input type="checkbox"/> SearchMode
List hotels within a specified distance to an airport and that fall within a specific price range.	<input type="checkbox"/> Order by <input type="checkbox"/> Top <input type="checkbox"/> Filter

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: QueryType

The SearchParameters.QueryType Property gets or sets a value that specifies the syntax of the search query. The default is 'simple'. Use 'full' if your query uses the Lucene query syntax.

You can write queries against Azure Search based on the rich Lucene Query Parser syntax for specialized query forms: wildcard, fuzzy search, proximity search, regular expressions are a few examples.

Box 2: Facets

The facets property gets or sets the list of facet expressions to apply to the search query. Each facet expression contains a field name, optionally followed by a comma-separated list of name:value pairs.

Box 3: Filter

The Filter property gets or sets the OData \$filter expression to apply to the search query. References:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters> <https://docs.microsoft.com/en-us/azure/search/query-lucene-syntax>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters.querytype>

NEW QUESTION 75

- (Exam Topic 3)

A company is developing a mobile app for field service employees using Azure App Service Mobile Apps as the backend.

The company's network connectivity varies throughout the day. The solution must support offline use and synchronize changes in the background when the app is online app.

You need to implement the solution.

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.

```
var client = new MobileServiceClient("MOBILE_APP_URL");
var store = new MobileServiceSQLiteStore
(Constants.OfflineDbPath);
store.DefineTable<ToDoItem>();
await client.SyncContext.InitializeAsync(store);
```

var todoTable = client.GetSyncTable<ToDoItem>();	<input type="checkbox"/>
var todoTable = client.GetTable<ToDoItem>();	<input type="checkbox"/>
var todoTable = client.SyncTable;	<input type="checkbox"/>
var todoTable = client.Table;	<input type="checkbox"/>
await client.SyncContext.PushAsync();	<input type="checkbox"/>

await todoTable.PullAsync("allToDoItems",todoTable.CreateQuery());	<input type="checkbox"/>
await todoTable.UpdateAsync();	<input type="checkbox"/>
todoTable.PullAsync("allToDoItems", todoTable.CreateQuery());	<input type="checkbox"/>
todoTable.UpdateAsync();	<input type="checkbox"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: var todoTable = client GetSyncTable<TodoItem>()

To setup offline access, when connecting to your mobile service, use the method GetSyncTable instead of GetTable (example):

IMobileServiceSyncTable todoTable = App.MobileService.GetSyncTable(); / Box 2: await todoTable.PullAsync("allTodoItems",todo.Table.CreateQuery());

Your app should now use IMobileServiceSyncTable (instead of IMobileServiceTable) for CRUD operations. This will save changes to the local database and also keep a log of the changes. When the app is ready to synchronize its changes with the Mobile Service, use the methods PushAsync and PullAsync (example):

await App.MobileService.SyncContext.PushAsync(); await todoTable.PullAsync();

References:

<https://azure.microsoft.com/es-es/blog/offline-sync-for-mobile-services/>

NEW QUESTION 77

- (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 81

- (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 82

- (Exam Topic 3)

You are developing a software solution for an autonomous transportation system. The solution uses large data sets and Azure Batch processing to simulate navigation sets for entire fleets of vehicles.

You need to create compute nodes for the solution on Azure Batch. What should you do?

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

Answer: D

Explanation:

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

Note:

Step 1: Create a pool of compute nodes. When you create a pool, you specify the number of compute nodes for the pool, their size, and the operating system.

When each task in your job runs, it's assigned to execute on one of the nodes in your pool.

Step 2 : Create a job. A job manages a collection of tasks. You associate each job to a specific pool where that job's tasks will run.

Step 3: Add tasks to the job. Each task runs the application or script that you uploaded to process the data files it downloads from your Storage account. As each task completes, it can upload its output to Azure Storage.

NEW QUESTION 87

- (Exam Topic 3)

You have an Azure App Services Web App. Azure SQL Database instance. Azure Storage Account and an Azure Redis Cache instance in a resource group.

A developer must be able to publish code to the web app. You must grant the developer the Contributor role to the web app

You need to grant the role.

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

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

Answer: AB

Explanation:

References:

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

NEW QUESTION 88

- (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 92

- (Exam Topic 3)

A company uses Azure SQL Database to store data for an app. The data includes sensitive information.

You need to implement measures that allow only members of the managers group to see sensitive information. Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Include the managers group.
- B. Exclude the managers group.
- C. Exclude the administrators group.
- D. Navigate to the following URL:
`PUT https://management.azure.com/subscriptions/00000000-1111-2222-3333-444444444444
/resourceGroups/rg01/providers/Microsoft.Sql/servers/server01/databases/customers
/transparentDataEncryption/current?api-version=2014-04-01`
- E. Run the following Azure PowerShell command:
`New-AzureRmSqlDatabaseDataMaskingRule -SchemaName "dbo" -TableName "customers" ` -ColumnName "ssn" -MaskingFunction "Default"`

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

Answer: BE

Explanation:

Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to designate how much of the sensitive data to reveal with minimal impact on the application layer.

SQL users excluded from masking - A set of SQL users or AAD identities that get unmasked data in the SQL query results.

Note: The New-AzureRmSqlDatabaseDataMaskingRule cmdlet creates a data masking rule for an Azure SQL database.

References:

<https://docs.microsoft.com/en-us/powershell/module/azurerm.sql/new-azurermsqldatabasedatamaskingrule?view>

NEW QUESTION 95

- (Exam Topic 3)

You develop an Azure web app. You monitor performance of the web app by using Application Insights. You need to ensure the cost for Application Insights does not exceed a preset budget. What should you do?

- A. Implement ingestion sampling using the Azure portal.
- B. Set a daily cap for the Application Insights instance.
- C. Implement adaptive sampling using the Azure portal.
- D. Implement adaptive sampling using the Application Insights SDK.
- E. Implement ingestion sampling using the Application Insights SDK.

Answer: D

Explanation:

Sampling is an effective way to reduce charges and stay within your monthly quota.

You can set sampling manually, either in the portal on the Usage and estimated costs page; or in the ASP.NET SDK in the .config file; or in the Java SDK in the ApplicationInsights.xml file, to also reduce the network traffic.

Adaptive sampling is the default for the ASP.NET SDK. Adaptive sampling automatically adjusts to the volume of telemetry that your app sends. It operates automatically in the SDK in your web app so that telemetry traffic on the network is reduced.

References:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling>

NEW QUESTION 98

- (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 http://contoso.com/api/9.2/, 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 100

- (Exam Topic 3)
You plan to deploy a web app to App Service on Linux. You create an App Service plan. You create and push a custom Docker image that image that contains the web app to Azure Container Registry.
You need to access the console logs generated from inside the container in real-time.
How should you complete the Azure CLI command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

az webapp log

▼

--name ContosoWeb --resource-group ContosoDevRG

config

download

show

tail

filesystem

--web-server-logging

--docker-container-logging

--application-logging

az

▼

log

▼

--name ContosoWeb --resource-group ContosoDevRG

webapp

acr

aks

config

download

show

tail

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: config

To Configure logging for a web app use the command: az webapp log config

Box 2: --docker-container-logging Syntax include:

az webapp log config [--docker-container-logging {filesystem, off}] Box 3: webapp

To download a web app's log history as a zip file use the command: az webapp log download

Box 4: download References:

<https://docs.microsoft.com/en-us/cli/azure/webapp/log>

NEW QUESTION 101

- (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 105

- (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 110

- (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
<input type="text" value="Strong"/>	<input type="text" value="Bounded Staleness"/>	Return the most recent patient status. <input type="text"/>
<input type="text" value="Consistent Prefix"/>	<input type="text" value="Eventual"/>	Return health monitoring data that is no less than one version behind. <input type="text"/>
		After patient is discharged and all changes are assessed, retrieve the correct billing data with the final charges <input type="text"/>

- 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 112

- (Exam Topic 3)

You are creating a hazard notification system that has a single signaling server which triggers audio and visual alarms to start and stop.

You implement Azure Service Bus to publish alarms. Each alarm controller uses Azure Service Bus to receive alarm signals as part of a transaction. Alarm events must be recorded for audit purposes. Each transaction record must include information about the alarm type that was activated.

You need to implement a reply trail auditing solution.

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

- A. Assign the value of the hazard message SessionID property to the SequenceNumber property.
- B. Assign the value of the hazard message SequenceNumber property to the DeliveryCount property.
- C. Assign the value of the hazard message MessageID property to the DeliveryCount property.
- D. Assign the value of the hazard message SessionID property to the ReplyToSessionId property.
- E. Assign the value of the hazard message MessageID property to the SequenceNumber property.
- F. Assign the value of the hazard message MessageID property to the CorrelationId property.

Answer: AB

NEW QUESTION 115

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AZ-204 Exam with Our Prep Materials Via below:

<https://www.certleader.com/AZ-204-dumps.html>