

Microsoft

Exam Questions AZ-204

Developing Solutions for Microsoft Azure (beta)



NEW QUESTION 1

- (Exam Topic 1)

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

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

Answer: C

Explanation:

Scenario: The Shipping Logic App requires secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model. You can access to Azure Virtual Network resources from Azure Logic Apps by using integration service environments (ISEs). Sometimes, your logic apps and integration accounts need access to secured resources, such as virtual machines (VMs) and other systems or services, that are inside an Azure virtual network. To set up this access, you can create an integration service environment (ISE) where you can run your logic apps and create your integration accounts. References: <https://docs.microsoft.com/en-us/azure/logic-apps/connect-virtual-network-vnet-isolated-environment-overview>

NEW QUESTION 2

- (Exam Topic 1)

You need to 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 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><div></div><div>StandardPremium</div></div>
Profile	<div><div></div><div>AkamaiMicrosoft</div></div>
Optimization	<div><div></div><div>general web deliverylarge file downloaddynamic site accelerationvideo-on-demand media streaming</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 4

- (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><div>Basic</div><div>Standard</div><div>Premium</div><div>Isolated</div></div>
Instance Count	<div><div>▼</div><div>1</div><div>10</div><div>20</div><div>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 5

- (Exam Topic 1)
You need to secure the Shipping Function app.
How should you configure the app? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Setting	Value
Authorization level	<div><div></div><div>Function</div><div>Anonymous</div><div>Admin</div></div>
User claims	<div><div></div><div>JSON Web Token (JWT)</div><div>Shared Access Signature (SAS) token</div><div>API Key</div></div>
Trigger type	<div><div></div><div>blob</div><div>HTTP</div><div>queue</div><div>timer</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
Scenario: Shipping Function app: Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).
Box 1: Function
Box 2: JSON based Token (JWT)
Azure AD uses JSON based tokens (JWTs) that contain claims Box 3: HTTP
How a web app delegates sign-in to Azure AD and obtains a token
User authentication happens via the browser. The OpenID protocol uses standard HTTP protocol messages. References:
<https://docs.microsoft.com/en-us/azure/active-directory/develop/authentication-scenarios>

NEW QUESTION 6
- (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 7

- (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 8

- (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 ▼

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

create --location centralus - -name \$resourcegroupname
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 ▼

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

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

az ▼

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

az ▼

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

az ▼

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

- 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 9

- (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		
Draft	Deploy solution.	
Brigade		
KubeCtl	View cluster and external IP addressing.	
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 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	<div><div></div><div>▼</div><div>Can be restored</div><div>Cannot be restored</div></div>
Snapshot 1	<div><div></div><div>▼</div><div>Can be restored</div><div>Cannot be restored</div></div>
Snapshot 2	<div><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 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 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 11

- (Exam Topic 3)

You are developing a project management service by using ASP.NET. The service hosts conversations, files, to-do lists, and a calendar that users can interact with at any time.

The application uses Azure Search for allowing users to search for keywords in the project data.

You need to implement code that creates the object which is used to create indexes in the Azure Search service.

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

- A. SearchService
- B. SearchIndexClient
- C. SearchServiceClient
- D. SearchCredentials

Answer: BC

Explanation:

The various client libraries define classes like Index, Field, and Document, as well as operations like Indexes.Create and Documents.Search on the SearchServiceClient and SearchIndexClient classes.

Example:

The sample application we'll be exploring creates a new index named "hotels", populates it with a few documents, then executes some search queries. Here is the main program, showing the overall flow:

/ This sample shows how to delete, create, upload documents and query an index static void Main(string[] args)

```
{
IConfigurationBuilder builder = new ConfigurationBuilder().AddJsonFile("appsettings.json"); IConfigurationRoot configuration = builder.Build();
SearchServiceClient serviceClient = CreateSearchServiceClient(configuration); Console.WriteLine("{0}", "Deleting index...\n");
DeleteHotelsIndexIfExists(serviceClient);
Console.WriteLine("{0}", "Creating index...\n"); CreateHotelsIndex(serviceClient);
ISearchIndexClient indexClient = serviceClient.Indexes.GetClient("hotels"); References:
https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk
```

NEW QUESTION 13

- (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 are developing an app that manages users for a video game. You plan to store the region, email address, and phone number for the player. Some players may not have a phone number. The player's region will be used to load-balance data.
Data for the app must be stored in Azure Table Storage.
You need to develop code to retrieve data for an individual player.
How should you complete the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```
public class PlayerEntity : TableEntity
{
    public PlayerEntity()
    {
    }
    public PlayerEntity(string region, string email)
    {
        PartitionKey =  ;
        RowKey=  ;
    }
    public string Phone { get; set; }
}
public class Player
{
    protected PlayerEntity player;
    async void GetPlayer(string cs,  table, string pk, string rk)
    {
        
        TableEntity query =TableEntity.Retrieve<PlayerEntity>(pk, rk);
        TableOperation query =TableOperation.Retrieve<PlayerEntity>(pk,rk);
        TableResult query =TableQuery.Retrieve<PlayerEntity>(pk,rk);
        TableResultSegment query =TableResult.Retrieve<PlayerEntity>(pk, rk);

        
        TableEntity data =await table.ExecuteAsync(query);
        TableOperation data =await table.ExeucteAsync(query);
        TableQuery data =await table.ExecuteAsync(query);
        TableResult data =await table.ExecuteAsync(query);
        player=data.Result as PlayerEntity;
    }
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: region
The player's region will be used to load-balance data. Choosing the PartitionKey.
The core of any table's design is based on its scalability, the queries used to access it, and storage operation requirements. The PartitionKey values you choose will dictate how a table will be partitioned and the type of queries that can be used. Storage operations, in particular inserts, can also affect your choice of PartitionKey values.
Box 2: email
Not phone number some players may not have a phone number. Box 3: CloudTable
Box 4 : TableOperation query =.. Box 5: TableResult
References:
<https://docs.microsoft.com/en-us/rest/api/storageservices/designing-a-scalable-partitioning-strategy-for-azure-ta>

NEW QUESTION 19

- (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>Get-AzureRmRoleDefinition-Name"Reader" ConvertTo-Json Out-File C:\SupportRole.json</div><div>Get-AzureRmRoleDefinition-Name"Operator" ConvertTo-Json Out-File C:\SupportRole.json</div><div>Set-AzureRmRoleDefinition-Name"Reader" Input-File C:\SupportRole.json</div><div>Set-AzureRmRoleDefinition Input-File C:\SupportRole.json</div></div>
Actions section	<div><div>"*/read*", "Microsoft.Support/*"</div><div>"*/read"</div><div>"*/read*", "Microsoft.Support/*"</div><div>"*/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 21

- (Exam Topic 3)

You are preparing to deploy a medical records application to an Azure virtual machine (VM). The application will be deployed by using a VHD produced by an on-premises build server.

You need to ensure that both the application and related data are encrypted during and after deployment to Azure.

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

Actions

Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage.

Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension.

Run the Azure PowerShell command Set-AzureRmVMOSDisk.

Encrypt the on-premises VHD by using BitLocker with a TPM. Upload the VM to Azure Storage.

Run the Azure PowerShell command New-AzureRmVm.

Answer area

Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage.

Run the Azure PowerShell command Set-AzureRmVMOSDisk.

Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage Step 2: Run the Azure PowerShell command Set-AzureRmVMOSDisk

To use an existing disk instead of creating a new disk you can use the Set-AzureRmVMOSDisk command. Example:

\$osDiskName = \$vmname+'_osDisk'

\$osDiskCaching = 'ReadWrite'

\$osDiskVhdUri = "https://\$storage.blob.core.windows.net/vhds/" + \$vmname + "_os.vhd"

\$vm = Set-AzureRmVMOSDisk -VM \$vm -VhdUri \$osDiskVhdUri -name \$osDiskName -Create Step 3: Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension

Use the Set-AzVMDiskEncryptionExtension cmdlet to enable encryption on a running IaaS virtual machine in Azure.

Incorrect:

Not TPM: BitLocker can work with or without a TPM. A TPM is a tamper resistant security chip on the system board that will hold the keys for encryption and check the integrity of the boot sequence and allows the most secure BitLocker implementation. A VM does not have a TPM.

References:

<https://www.itprotoday.com/iaaspaas/use-existing-vhd-azurerm-vm>

NEW QUESTION 23

- (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 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";
```

▼ hub=

NotificationHubClient
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.


▼ (windowsToastPayload);



SendWindowsNativeNotificationAsync  
SubmitNotificationHubJobAsync  
ScheduleNotificationAsync  
SendAppleNativeNotificationAsync


    . . .
}
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 29

- (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",
  : "All",
  "optionalClaims"
  "groupMembershipClaims"
  : 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 30

- (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 32

- (Exam Topic 3)

A company is developing a solution that allows smart refrigerators to send temperature information to a central location. You have an existing Service Bus. The solution must receive and store messages until they can be processed. You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location. You need to complete the configuration. Which Azure CLI or PowerShell command should you run?

- A. `az servicebus namespace create`
 `- -resource-group fridge-rg`
 `- -name fridge-ns`
 `- -location fridge-loc`
- B. `az servicebus queue create`
 `--resource-group fridge-rg`
 `--namespace-name fridge-ns`
 `--name fridge-q`
- C. `connectionString=$(az servicebus namespace authorization-rule keys list`
 `--resource-group fridge-rg`
 `--fridge-ns fridge-ns`
 `--name RootManageSharedAccessKey`
 `--query primaryConnectionString --output tsv)`
- D. `az group create`
 `--name fridge-rg`
 `--location fridge-log`

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

Answer: B

Explanation:

A service bus instance has already been created (Step 2 below). Next is step 3, Create a Service Bus queue. Note: Steps:

Step 1: # Create a resource group resourceGroupName="myResourceGroup"

`az group create --name $resourceGroupName --location eastus`

Step 2: # Create a Service Bus messaging namespace with a unique name namespaceName=myNameSpace\$RANDOM

`az servicebus namespace create --resource-group $resourceGroupName --name $namespaceName --location eastus`

Step 3: # Create a Service Bus queue

`az servicebus queue create --resource-group $resourceGroupName --namespace-name $namespaceName`
`--name BasicQueue`

Step 4: # Get the connection string for the namespace

`connectionString=$(az servicebus namespace authorization-rule keys list --resource-group`

`$resourceGroupName --namespace-name $namespaceName --name RootManageSharedAccessKey --query primaryConnectionString --output tsv)`

References:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quickstart-cli>

NEW QUESTION 35

- (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 40

- (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	<div><div></div><div>2</div><div>4</div><div>8</div><div>16</div></div>
Pricing tier	<div><div></div><div>Isolated</div><div>Standard</div><div>Premium</div><div>Consumption</div></div>

- 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 45

- (Exam Topic 3)
You develop software solutions for a mobile delivery service. You are developing a mobile app that users can use to order from a restaurant in their area. The app uses the following workflow:
* 1. A driver selects the restaurants for which they will deliver orders.
* 2. Orders are sent to all available drivers in an area.
* 3. Only orders for the selected restaurants will appear for the driver.
* 4. The first driver to accept an order removes it from the list of available orders. You need to implement an Azure Service Bus solution.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer area
Create a Service Bus topic for each restaurant for which a driver can receive messages.	
Create a single Service Bus topic.	
Create a single Service Bus subscription.	
Create a single Service Bus Namespace.	
Create a Service Bus Namespace for each restaurant for which a driver can receive messages.	
Create a Service Bus subscription for each restaurant for which a driver can receive orders.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Create a single Service Bus Namespace
To begin using Service Bus messaging entities in Azure, you must first create a namespace with a name that is unique across Azure. A namespace provides a scoping container for addressing Service Bus resources within your application.
Box 2: Create a Service Bus Topic for each restaurant for which a driver can receive messages. Create topics.
Box 3: Create a Service Bus subscription for each restaurant for which a driver can receive orders. Topics can have multiple, independent subscriptions.
References:
<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messaging-overview>

NEW QUESTION 46

- (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 51

- (Exam Topic 3)

You are developing a web app that is protected by Azure Web Application Firewall (WAF). All traffic to the web app is routed through an Azure Application Gateway instance that is used by multiple web apps. The web app address is contoso.azurewebsites.net.

All traffic must be secured with SSL. The Azure Application Gateway instance is used by multiple web apps. You need to configure the Azure Application Gateway for the app.

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

- A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting.
- B. Convert the web app to run in an Azure App service environment (ASE).
- C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application gateway.
- D. In the Azure Application Gateway's HTTP setting, set the value of the Override backend path option to contoso22.azurewebsites.net.

Answer: AD

Explanation:

D: The ability to specify a host override is defined in the HTTP settings and can be applied to any back-end pool during rule creation.

The ability to derive the host name from the IP or FQDN of the back-end pool members. HTTP settings also provide an option to dynamically pick the host name from a back-end pool member's FQDN if configured with the option to derive host name from an individual back-end pool member.

A (not C): SSL termination and end to end SSL with multi-tenant services.

In case of end to end SSL, trusted Azure services such as Azure App service web apps do not require whitelisting the backends in the application gateway.


Therefore, there is no need to add any authentication certificates.


Add HTTP setting


saiappgw-appgw

* Protocol

☐ HTTP ☒ HTTPS


 Authentication certificates are not required for trusted Azure certificates for end to end ssl to work

* Port 


443 

* Request timeout (seconds)

20

Override backend path 

☒ Use for App service

☒ Use custom probe 

OK

Reference:

<https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-web-app-overview>

NEW QUESTION 53

- (Exam Topic 3)

You provide an Azure API Management managed web service to clients. The back end web service implements HTTP Strict Transport Security (HSTS). Every request to the backend service must include a valid HTTP authorization header. You need to configure the Azure API Management instance with an authentication policy. Which two policies can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Certificate Authentication
- B. Basic Authentication
- C. OAuth Client Credential Grant
- D. Digest Authentication

Answer: AC

NEW QUESTION 58

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

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

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

Solution:

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

- A. Yes
- B. No

Answer: A

Explanation:

To configure Manifest to include Group Claims in Auth Token

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

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

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

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

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

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

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

NEW QUESTION 64

- (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 68

- (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 70

- (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 72

- (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 75

- (Exam Topic 3)
You are building a traffic monitoring system that monitors traffic along six highways. The system produces time series analysis-based reports for each highway. Data from traffic sensors are stored in Azure Event Hub.
Traffic data is consumed by four departments. Each department has an Azure Web App that displays the time-series-based reports and contains a WebJob that processes the incoming data from Event Hub. All Web Apps run on App Service Plans with three instances.
Data throughout must be maximized. Latency must be minimized. You need to implement the Azure Event Hub.
Which settings should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Setting	Value
Number of partitions	<div><div></div><div>▼</div><div>3</div><div>4</div><div>6</div><div>12</div></div>
Partition Key	<div><div></div><div>▼</div><div>Highway</div><div>Department</div><div>Timestamp</div><div>VM name</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

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

NEW QUESTION 76

- (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 80

- (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 83

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

Answer: B

Explanation:

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

NEW QUESTION 84

- (Exam Topic 3)

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

The screenshot displays the Azure Logic App Designer interface for a workflow. The workflow is structured as follows:

- For each** (Loop):
 - Set tier age variable** (Action)
 - Set tier age variable** (Action)
- Scan all blobs in this folder** (Action)
- When there are messages in a queue** (Condition):
 - Queue Name**: processing
 - Show advanced options**: [Show advanced options](#)
 - Connected to**: tableStorageAccountConnection [Change connection](#)
- If true** (Branch):
 - Add an action** (Button)
- If false** (Branch):
 - Add an action** (Button)

Arrows indicate the flow from the 'For each' loop to the 'Scan all blobs in this folder' action, and then to the 'When there are messages in a queue' condition. The 'If true' and 'If false' branches are shown below the condition.

- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 89

The application will call `CreatePlayerWithGame` with the correct `gameId` and the `playerId` to start the process. (Line numbers are included for reference only.)


```

01. namespace ContosoCradt
02. {
03.     public class PlayerDbContext : DbContext
04.     {
05.         public PlayerDbContext() : base ("name-dBConnString") { }
06.         public DbSet<Player> Players { get ; set ; }
07.         public DbSet<Game> Games { get ; set ; }
08.         protected override void OnModelCreating(ModelBuilder modelBuilder)
09.         {
10.             modelBuilder.Entity<Player>().HasMany(x => x.Games).WithMany (x => x.Players);
11.         }
12.     }
13.     internal class dbConfiguration : DbMigrationConfiguration<PlayerDbContext>
14.     {
15.         public dbConfiguration() { AutomaticMigrationsEnabled = true ; }
16.     }
17.     public class mp
18.     {
19.         public void CreatePlayerWithGame(int playerId, int gameId) => AddPlayer(playerId, GetGame(gameId));
20.         public Game GetGame(int gameId)
21.         {
22.             using (var db = new PlayerDbContext())
23.             {
24.                 return db.Games.FirstOrDefault(x => x.GameId == gameId);
25.             }
26.         }
27.         public Player AddPlayer (int playerId, Game game)
28.         {
29.             using (var db = new PlayerDbContext())
30.             {
31.                 var player = new Player
32.                 {
33.                     PlayerId = playerId,
34.                     Games = new List <Game> {game },
35.                 };
36.                 db.Players.Add(player);
37.                 db.SaveChanges();
38.                 return player;
39.             }
40.         }
41.     }
42.     public class Player
43.     {
44.         public int PlayerId { get ; set; }
45.         public string PlayerName { get ; set; }
46.         public virtual List<Game> Games { get ; set; }
47.     }
48.     public class Game
49.     {
50.         public int GameId { get ; set; }
51.         public string Title { get ; set; }
52.         public string Platform { get ; set; }
53.         public virtual List<Player> Players { get ; set; }
54.     }

```

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 successfully insert a player record. | <input type="radio"/> | <input type="radio"/> |
| The code has a bug and will insert an additional copy of the Game record with a new Id. | <input type="radio"/> | <input type="radio"/> |
| The code has a bug and will insert the wrong gameId value. | <input type="radio"/> | <input type="radio"/> |
| There is a valid many-to-many relationship between Players and Games. | <input type="radio"/> | <input type="radio"/> |

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Many-to-many relationships without an entity class to represent the join table are not yet supported. However, you can represent a many-to-many relationship by including an entity class for the join table and mapping two separate one-to-many relationships.

```
protected override void OnModelCreating(ModelBuilder modelBuilder)
```

```
{
    modelBuilder.Entity<PostTag>() HasKey(t => new { t.PostId, t.TagId }); modelBuilder.Entity<PostTag>() HasOne(pt => pt.Post)
    WithMany(p => p.PostTags) HasForeignKey(pt => pt.PostId); modelBuilder.Entity<PostTag>() HasOne(pt => pt.Tag) WithMany(t => t.PostTags) HasForeignKey(pt
    => pt.TagId);
}
```

NEW QUESTION 93

- (Exam Topic 3)

You manage several existing Logic Apps.

You need to change definitions, add new logic, and optimize these apps on a regular basis.

What should you use? To answer, drag the appropriate tools to the correct functionalities. Each tool may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer Area

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

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Box 1: Enterprise Integration Pack

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

Box 2: Code View Editor

To work with logic app definitions in JSON, open the Code View editor when working in the Azure portal or in Visual Studio, or copy the definition into any editor that you want.

Box 3: Logical Apps Designer

You can build your logic apps visually with the Logic Apps Designer, which is available in the Azure portal through your browser and in Visual Studio.

References:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-enterprise-integration-b2b> <https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions> <https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-overview>

NEW QUESTION 97

- (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 99

- (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 103

- (Exam Topic 3)

You have a web app named MainApp. You are developing a triggered App Service background task by using the WebJobs SDK. This task automatically invokes a function code whenever any new data is received in a queue.

You need to configure the services.

Which service should you use for each scenario? To answer, drag the appropriate services to the correct scenarios. Each service 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	Scenario	Service
Logic Apps	Process a queue data item.	
WebJobs	Manage all code segments from the same DevOps environment.	
Flow		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: WebJobs

A WebJob is a simple way to set up a background job, which can process continuously or on a schedule. WebJobs differ from a cloud service as it gives you get less fine-grained control over your processing environment, making it a more true PaaS service.

Box 2: Flow

NEW QUESTION 104

- (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.	
Continuous	Run on a single instance that Azure select for load balancing.	
	Supports remote debugging	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Continuous

Continuous runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.

Box 2: Triggered

Triggered runs on a single instance that Azure selects for load balancing. Box 3: Continuous

Continuous supports remote debugging. Note:

The following table describes the differences between continuous and triggered WebJobs.

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

References:

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

NEW QUESTION 107

- (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><div></div><div><div>Bypass cache</div><div>Override</div><div>Set if missing</div></div></div>
Cache expiration duration	<div><div></div><div><div>1 second</div><div>1 minute</div><div>1 hour</div><div>1 day</div></div></div>
Query string caching behavior	<div><div></div><div><div>Ignore query strings</div><div>Bypass caching for query strings</div><div>Cache every unique URL</div></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 112

- (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 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..

Does the solution meet the goal?

- A. Yes
B. No

Answer: A

Explanation:

* 1. The index needs to be populated. To do this, we will need a SearchIndexClient. There are two ways to obtain one: by constructing it, or by calling Indexes.GetClient on the SearchServiceClient. Here we will use the first method.

* 2. Create the indexBatch with the documents Something like:

```
var hotels = new Hotel[];
```

```
{
new Hotel()
{
HotelId = "3",
BaseRate = 129.99,
Description = "Close to town hall and the river"
}
};
...
var batch = IndexBatch.Upload(hotels);
* 3. The next step is to populate the newly-created index Example:
var batch = IndexBatch.Upload(hotels); try
{
indexClient.Documents.Index(batch);
}
References:
https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk
```

NEW QUESTION 113

- (Exam Topic 3)
ASP.NET Core API app by using C#. The API app will allow users to authenticate by using Twitter and Azure Active Directory (Azure AD). Users must be authenticated before calling API methods. You must log the user's name for each method call. You need to configure the API method calls. Which values should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Code segment	Value
Attribute	<div><div></div><div>Authorize</div><div>AllowAnonymous</div><div>AutoValidateAntiforgeryToken</div></div>
Request Header	<div><div></div><div>X-MS-CLIENT-PRINCIPAL-NAME</div><div>Proxy-Authorization</div><div>X-Forwarded-For</div><div>X-MS-CLIENT-PRINCIPAL-ID</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Authorize
Box 2: X-MS-CLIENT-PRINCIPAL-NAME
App Service passes user claims to your application by using special headers. External requests aren't allowed to set these headers, so they are present only if set by App Service. Some example headers include:
X-MS-CLIENT-PRINCIPAL-NAME X-MS-CLIENT-PRINCIPAL-ID
Here's the set of headers you get from Easy Auth for a Twitter authenticated user:
{
"cookie": "AppServiceAuthSession=Lx43...xHDTA==", "x-ms-client-principal-name": "evilSnobu",
"x-ms-client-principal-id": "35....", "x-ms-client-principal-idp": "twitter",
"x-ms-token-twitter-access-token": "35...Dj",
"x-ms-token-twitter-access-token-secret": "OK3...Jx",
}
References:
https://docs.microsoft.com/en-us/azure/app-service/app-service-authentication-how-to

NEW QUESTION 118

- (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 119

- (Exam Topic 3)

You must implement Application Insights instrumentation capabilities utilizing the Azure Mobile Apps SDK to provide meaningful analysis of user interactions with a mobile app.

You need to capture the data required to implement the Usage Analytics feature of Application Insights. Which three data values should you capture? Each correct answer presents part of the solution

NOTE: Each correct selection is worth one point.

- A. Trace
- B. Session Id
- C. Exception
- D. User Id
- E. Events

Answer: ADE

Explanation:

Application Insights is a service for monitoring the performance and usage of your apps. This module allows you to send telemetry of various kinds (events, traces, etc.) to the Application Insights service where your data can be visualized in the Azure Portal.

Application Insights manages the ID of a session for you. References: <https://github.com/microsoft/ApplicationInsights-Android>

NEW QUESTION 124

- (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 126

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

AZ-204 Practice Exam Features:

- * AZ-204 Questions and Answers Updated Frequently
- * AZ-204 Practice Questions Verified by Expert Senior Certified Staff
- * AZ-204 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AZ-204 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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