Microsoft

Exam Questions AZ-204

Developing Solutions for Microsoft Azure (beta)

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

```
<?xml version="1.0" encoding="utf-8"?>
<StorageServiceProperties>
    ...
    <Cors>
    <CorsRule>
```

<	▼ >	▼	</	▼ >
AllowedHeaders	http://*.wideworldimporters.com		AllowedHeaders	
ExposedHeaders	http://test.wideworldimporters.com		ExposedHeaders	
AllowedMethods	http://test-shippingapi.wideworldimporters.com		AllowedMethods	
AllowedOrigins	http://www.wideworldimporters.com		AllowedOrigins	

`<AllowedMethods>` | ▼ | `</AllowedMethods>`

GET,PUT
GET
POST
GET,HEAD

```
    ...
    </CorsRule>
    </Cors>
</StorageServiceProperties>
```

A. Mastered
B. Not Mastered

Answer: A

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

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

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

NEW QUESTION 2
- (Exam Topic 1)
You need to migrate on-premises shipping data to Azure. What should you use?

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

Answer: D

Explanation:
Migrate from on-premises or cloud implementations of MongoDB to Azure Cosmos DB with minimal downtime by using Azure Database Migration Service.
Perform resilient migrations of MongoDB data at scale and with high reliability.
Scenario: Data migration from on-premises to Azure must minimize costs and downtime.
The application uses MongoDB JSON document storage database for all container and transport information. References:
https://azure.microsoft.com/en-us/updates/mongodb-to-azure-cosmos-db-online-and-offline-migrations-are-now

NEW QUESTION 3
- (Exam Topic 1)

You need to support the message processing for the ocean transport workflow.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

- 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 4
- (Exam Topic 1)
You need to update the APIs to resolve the testing error.
How should you complete the Azure CLI command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
az webapp  [ ▼ ]  [ ▼ ]  -g shipping-apis-test-rg -n web
```

cors		add
config		up
deployment		remove

```
--  [ ▼ ]          [ ▼ ]
```

slot		http://*.wideworldimporters.com
allowed-origins		http://test-shippingapi.wideworldimporters.com
name		http://test.wideworldimporters.com
		http://www.wideworldimporters.com

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 5
- (Exam Topic 1)
You need to correct the VM issues.
Which tools should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

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

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 6
- (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 7
- (Exam Topic 3)
A company is developing a Java web app. The web app code is hosted in a GitHub repository located at https://github.com/Contoso/webapp.
The web app must be evaluated before it is moved to production. You must deploy the initial code release to a deployment slot named staging.
You need to create the web app and deploy the code.
How should you complete the commands? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

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

az ▼ create --location centralus - -name $resourcegroupname
| group |
| webapp |
| appservice plan |
| webapp deployment slot |
| webapp deployment source |

create --name $webappname - -resource-group $resourcegroupname
- -sku S3

create --name $webappname - -resource-group $resourcegroupname
\ - -plan $webappname

create --name $webappname - -resource-group $resourcegroupname
\ - -slot staging

az ▼ config - -name $webappname - -resource-group $resourcegroupname
| group |
| webapp |
| appservice plan |
| webapp deployment slot |
| webapp deployment source |

\ - -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 8
- (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	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 9
- (Exam Topic 3)
You are creating a script that will run a large workload on an Azure Batch pool. Resources will be reused and do not need to be cleaned up after use.
You have the following parameters:
You need to write an Azure CLI script that will create the jobs, tasks, and the pool.
In which order should you arrange the commands to develop the solution? To answer, move the appropriate commands from the list of command segments to the answer area and arrange them in the correct order.

Command segments

```
az batch pool create
 --id mypool --vm-size Standard_A1_v2
 --target-dedicated-nodes 2
 --image $image
 --node-agent-sku-id $sku
```

```
az batch job
create
 --id myjob
 --pool-id mypool
```

```
for i in {1..$numberOfJobs}
do
```

```
az batch task create
 --task-id mytask$i
 --job-id myjob
 --command-line $script
```

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: az batch pool create
Create a new Linux pool with a virtual machine configuration. az batch pool create \
--id mypool \
--vm-size Standard_A1 \
--target-dedicated 2 \
--image canonical:ubuntuserver:16.04-LTS \
--node-agent-sku-id "batch.node.ubuntu 16.04" Step 2: az batch job create
Create a new job to encapsulate the tasks that are added. az batch job create \
--id myjob \
--pool-id mypool
Step 3: az batch task create
Add tasks to the job. Here the task is a basic shell command. az batch task create \
--job-id myjob \
--task-id task1 \
--command-line "/bin/bash -c 'printenv AZ_BATCH_TASK_WORKING_DIR'" Step 4: for i in {1..$numberOfJobs} do
References:
https://docs.microsoft.com/bs-latn-ba/azure/batch/scripts/batch-cli-sample-run-job

NEW QUESTION 10
- (Exam Topic 3)
You are developing an application that use an Azure blob named data to store application data. The application creates blob snapshots to allow application state to be reverted to an earlier state. The Azure storage account has soft deleted enabled.
The system performs the following operations in order:
•The blob is updated
•Snapshot 1 is created.
•Snapshot 2 is created.
•Snapshot 1 is deleted.
A system error then deletes the data blob and all snapshots. You need to determine which application states can be restored.
What is the restorability of the application data? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Application State   Restorability

Data blob

| Can be restored |
| Cannot be restored |

Snapshot 1

| Can be restored |
| Cannot be restored |

Snapshot 2

| Can be restored |
| Cannot be restored |

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 10
- (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 14
- (Exam Topic 3)
You develop an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage
containers named Container1 and Container2. Uploading of videos occurs on an irregular basis.
You need to copy specific blobs from Container1 to Container2 in real time when specific requirements are met, excluding backup blob copies.
What should you do?

A. Download the blob to a virtual machine and then upload the blob to Container2.
B. Run the Azure PowerShell command Start-AzureStorageBlobCopy.

C. Copy blobs to Container2 by using the Put Blob operation of the Blob Service REST API.
D. Use AzCopy with the Snapshot switch blobs to Container2.

Answer: B

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

NEW QUESTION 19
- (Exam Topic 3)
You are developing a .NET Core model-view controller (MVC) application hosted on Azure for a health care system that allows providers access to their information.
You develop the following code:

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

You define a role named SysAdmin.
You need to ensure that the application meets the following authorization requirements:

≫ Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.

≫ Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.
How should you complete the code? To answer, drag the appropriate code segments to the correct locations. Each code segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

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

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 22
- (Exam Topic 3)
You are developing a data storage solution for a social networking app.
The solution requires a mobile app that stores user information using Azure Table Storage.
You need to develop code that can insert multiple sets of user information.
How should you complete the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse(
  CloudConfigurationManager.GetSetting("StorageConnectionString"));
CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
CloudTable table = tableClient.GetTableReference("clients");
Table.CreateIfNotExists();
```

```
                    ▼    op = new                 ▼    ();
TableOperation                          TableOperation
TableBatchOperaton                      TableBatchOperaton
TableEntity                             TableEntity
TableQuery                              TableQuery
```

```
. . .
                 ▼    (op);
table.
     ExecuteBatch
     Execute
     Insert
     InsertOrMerge
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1, Box 2: TableBatchOperation Create the batch operation.
TableBatchOperation op = new TableBatchOperation(); Box 3: ExecuteBatch
/ Execute the batch operation. table.ExecuteBatch(op);
Note: You can insert a batch of entities into a table in one write operation. Some other notes on batch operations:
You can perform updates, deletes, and inserts in the same single batch operation. A single batch operation can include up to 100 entities.
All entities in a single batch operation must have the same partition key.
While it is possible to perform a query as a batch operation, it must be the only operation in the batch. References:
https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet

NEW QUESTION 24
- (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 25
- (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'
$osDiskVhdUri = "https://$stoname.blob.core.windows.net/vhds/"+$vmname+"_os.vhd"
$vm = Set-AzureRmVMOSDisk -VM $vm -VhdUri $osDiskVhdUri -name $osDiskName -Create Step 3: Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension
Use the Set-AzVMDiskEncryptionExtension cmdlet to enable encryption on a running IaaS virtual machine in Azure.
Incorrect:
Not TPM: BitLocker can work with or without a TPM. A TPM is a tamper resistant security chip on the system board that will hold the keys for encryption and check the integrity of the boot sequence and allows the most secure BitLocker implementation. A VM does not have a TPM.
References:
https://www.itprotoday.com/iaaspaas/use-existing-vhd-azurerm-vm

NEW QUESTION 28
- (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 33
- (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		GetInstallation
NotificationHubClientSettings		CreateClientFromConnectionString
NotificationHubJob		CreateOrUpdateInstallation
NotificationDetails		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 37
- (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 40
- (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 42
- (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 45
- (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 46
- (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.

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 47
- (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.	○	○
Edge nodes must be purged of all cache assets.	○	○
The compression type is supported.	○	○

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 52
- (Exam Topic 3)
Note: This question is part of a series of questions that present the same scenario. Each question in the scries 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 a DataContainer that contains the documents which must be added.
* 3. Create a DataSource instance and set its Container property to the DataContamer
* 4 Call the Documents.Suggest method of the SearchlndexClient and pass the DataSource.
Does the solution meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 54
- (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.

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

NEW QUESTION 58
- (Exam Topic 3)
You provide an Azure API Management managed web service lo 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 uses? 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 59
- (Exam Topic 3)
You are creating an app that uses Event Grid to connect with other services. Your app's event data will be sent to a serverless function that checks compliance. This function is maintained by your company.
You write a new event subscription at the scope of your resource. The event must be invalidated after 3 specific period of time. You need to configure Event Grid to ensure security.
What should you implement? To answer, select the appropriate options in [he answer area. NOTE: Each correct selection is worth one point

Authentication   Type

WebHook event delivery

SAS tokens
Key authentication
JWT token

Topic publishing

ValidationCode handshake
ValidationURL handshake
Management Access Control

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: SAS tokens
Custom topics use either Shared Access Signature (SAS) or key authentication. Microsoft recommends SAS, but key authentication provides simple programming, and is compatible with many existing webhook publishers.
In this case we need the expiration time provided by SAS tokens. Box 2: ValidationCode handshake
Event Grid supports two ways of validating the subscription: ValidationCode handshake (programmatic) and ValidationURL handshake (manual).
If you control the source code for your endpoint, this method is recommended.

NEW QUESTION 63
- (Exam Topic 3)
A company is implementing a publish-subscribe (Pub/Sub) messaging component by using Azure Service Bus. You are developing the first subscription application.
In the Azure portal you see that messages are being sent to the subscription for each topic. You create and initialize a subscription client object by supplying the correct details, but the subscription application is still not consuming the messages.
You need to complete the source code of the subscription client What should you do?

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

Answer: C

Explanation:
Using topic client, call RegisterMessageHandler which is used to receive messages continuously from the entity. It registers a message handler and begins a new thread to receive messages. This handler is waited on every time a new message is received by the receiver.
subscriptionClient.RegisterMessageHandler(ReceiveMessagesAsync, messageHandlerOptions); References:
https://www.c-sharpcorner.com/article/azure-service-bus-topic-and-subscription-pub-sub/

NEW QUESTION 66
- (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 conn I 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 69
- (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 confing 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 74
- (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

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.

Answer Area

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 76
- (Exam Topic 3)
You are configuring a development environment for your team. You deploy the latest Visual Studio image from the Azure Marketplace to your Azure subscription.
The development environment requires several software development kits (SDKs) and third-party components to support application development across the organization. You install and customize the deployed virtual machine (VM) for your development team. The customized VM must be saved to allow provisioning of a new team member development environment.
You need to save the customized VM for future provisioning.
Which tools or services should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Action	Tool or service
Generalize the VM.	Azure PowerShell Visual Studio command prompt Azure Migrate Azure Backup
Store images.	Azure Blob Storage Azure Data Lake Storage Azure File Storage Azure Table Storage

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Azure Powershell

Creating an image directly from the VM ensures that the image includes all of the disks associated with the VM, including the OS disk and any data disks.
Before you begin, make sure that you have the latest version of the Azure PowerShell module. You use Sysprep to generalize the virtual machine, then use Azure PowerShell to create the image. Box 2: Azure Blob Storage
References:
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource#create-an-image-of-a

NEW QUESTION 79
- (Exam Topic 3)
You are developing a solution that will use Azure messaging services.
You need to ensure that the solution uses a publish-subscribe model and eliminates the need for constant polling.
What are two possible ways to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

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

Answer: AC

Explanation:
It is strongly recommended to use available messaging products and services that support a publish-subscribe model, rather than building your own. In Azure, consider using Service Bus or Event Grid. Other technologies that can be used for pub/sub messaging include Redis, RabbitMQ, and Apache Kafka.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber

NEW QUESTION 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 84
- (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.

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 85
- (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 88
- (Exam Topic 3)
A company backs up all manufacturing data to Azure Blob Storage. Admins move blobs from hot storage to archive tier storage every month.
You must automatically move blocks to Archive tier after they have not been accessed for 180 days. The path for any item that is not archived must be placed in an existing queue. This operation must be performed automatically once a month. You set the value of TierAgeInDays to 180.
How should you configure the Logic App? To answer, drag the appropriate triggers or action blocks to the correct trigger or action slots. Each trigger or action block may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Triggers and Action Blocks

Insert Entity

| *Table | processing | ✕ |
| *Entity | 🔷 Path ✕ | |

Show advanced options ∨

Tier blob

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

| *Blob path | 🔷 Path ✕ | 🗔 |
| *Blob Tier | Archive | ∨ |

When there are messages in a queue

| *Queue Name | processing | ∨ |

Show advanced options ∨

Connected to tableStorageAccountConnection. Change connection.

Recurrence

| *Interval | | *Frequency | |
| 1 | | Month | ∨ |

Show advanced options ∨

Answer Area

↓

{x} Set tier age variable ⋯

↓

Set tier age variable ⋯

↓

For each ⋯

⋯ Scan all blobs in this folder

*Select an output from previous steps 🔷 value ✕

When there are messages in a queue

| *Queue Name | processing | ∨ |

Show advanced options ∨

Connected to tableStorageAccountConnection. Change connection.

✅ **If true**

❌ **if false**

⬚ Add an action   ⬚ Add an action

⬚ Add an action

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Recurrence Box 2: Insert Entity
Box 3 (if true): Tier Blob Box 4: (if false):
Leave blank. References:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-perform-data-operations

NEW QUESTION 91
- (Exam Topic 3)
You are working for a company that designs mobile applications. They maintain a server where player records are assigned to their different games. The tracking system is new and in development.
The application uses Entity Framework to connect to an Azure Database. The database holds a Player table and Game table.
When adding a player, the code should insert a new player record, and add a relationship between an existing game record and the new player record.
The application will call CreatePlayerWithGame with the correct gameIdand 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 DmModelCreating{DBModelBuilder modelBuilder}
09. {
10.   modelBuilder.Entity<Player>().MesMany(x => x.Games). WithMany (x => x Players);
11. }
12. }
13.   internal series 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. public class Player
42. {
43. public int PlayerId { get ; set; }
44. public string PlayerName { get ; set; }
45. public virtual List<Game> Games { get ; set; }
46. }
47. public class Game
48. {
49. public int GameIs { get ; set }
50. public string Title { get ; set; }
51.   public string Platform { get ; set; ]
52.   public virtual List<Player> Players { get ; set; }
53.      }
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.	○	○
The code has a bug and will insert an additional copy of the Game record with a new Id.	○	○
The code has a bug and will insert the wrong gameId value.	○	○
There is a valid many-to-many relationship between Players and Games.	○	○

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 92
- (Exam Topic 3)
A company runs an international travel and bookings management service. The company plans to begin offering restaurant bookings. You must develop a solution that uses Azure Search and meets the following requirements:
• Users must be able to search for restaurants by name, description, location, and cuisine.
• Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.
• All words in descriptions must be included in searches. You need to add annotations to the restaurant class.
How should you complete the code segment? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```
[SerializePropertyNameAsCamelCase]
public class Restaurant
{
  [Key, IsFilterable]
  public int RestaurantId { get; set; }
  [IsSearchable, IsFilterable, IsSortable]
  public string Name { get; set; }
```

▼
[IsSearchable.IsFilterable.IsSortable, IsFacetable]
[IsFilterable.IsFacetable,Required]
[IsSearchable]
[IsSearchable, Required]

```
  public string location { get; set; }
  public string Phone { get; set; }
```

▼
[Required]
[IsSearchable]
[IsFilterable, IsFacetable, Required]
[IsFilterable, IsFacetable, IsSortable]

```
  public string Description { get; set; }
```

▼
[IsFiltrable, IsSortable, IsSearchable]
[IsFilterable, IsSortable, IsFacetable]
[IsFiltrable, IsSortable, Key]
[IsFilteranle, IsSortable, IsSearchable, Required]

```
  public double Rating { get; set; }
```

▼
[IsSearchable, IsFilterable, IsFacetable]
[IsFilterbale, IsSortable, Key]
[IsFilterable, IsSortable, IsSearchable]
[IsFilterable, IsSortable, Key, Required]

```
  public List<string> Cuisines { get; set; }
```

▼
[IsFilterable, IsSortable, Key, Required]
[IsSearchable, IsSortable, IsFacetable]
[IsFilterable, IsSortable, Key, IsSearchable]
[IsFilterable, IsFacetable]

```
  public bool FamilyFriendly { get; set; }
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: [IsSearchable.IsFilterable.IsSortable,IsFacetable] Location
Users must be able to search for restaurants by name, description, location, and cuisine.
Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. Box 2: [IsSearchable.IsFilterable.IsSortable,Required]
Description
Users must be able to search for restaurants by name, description, location, and cuisine. All words in descriptions must be included in searches.
Box 3: [IsFilterable,IsSortable,IsFaceTable] Rating

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. Box 4: [IsSearchable.IsFilterable,IsFacetable]
Cuisines
Users must be able to search for restaurants by name, description, location, and cuisine.
Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. Box 5: [IsFilterable,IsFacetable]
FamilyFriendly
Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. References:
https://www.henkboelman.com/azure-search-the-basics/

NEW QUESTION 94
- (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 Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.
You have the following requirements:

≫ Queue size must not grow larger than 80 gigabytes (GB).

≫ Use first-in-first-out (FIFO) ordering of messages.

≫ Minimize Azure costs.
You need to implement the messaging solution.
Solution: Use the .Net API to add a message to an Azure Service Bus Queue from the mobile application. Create an Azure Function App that uses an Azure Service Bus Queue trigger.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
You can create a function that is triggered when messages are submitted to an Azure Storage queue.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function

NEW QUESTION 99
- (Exam Topic 3)
You plan to create a Docker image that runs as ASP.NET Core application named ContosoApp. You have a setup script named setupScript.ps1 and a series of application files including ContosoApp.dll.
You need to create a Dockerfile document that meets the following requirements:
•Call setupScript.ps1 when the container is built.
•Run ContosoApp.dll when the container starts.
The Docker document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored.
Which four 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.

Commands	Answer Area
RUN powershell ./setupScript.ps1 CMD ["dotnet", "ContosoApp.dll"]	
EXPOSE ./ContosoApp/ /apps/ContosoApp	
COPY ./ .	(<) (>) (^) (v)
FROM microsoft/aspnetcore:2.0	
WORKDIR /apps/ContosoApp	
CMD powershell ./setupScript.ps1 ENTRYPOINT ["dotnet", "ContosoApp.dll"]	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: WORKDIR /apps/ContosoApp Step 2: COPY ./ .The Docker document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are

stored.
Step 3: EXPOSE ./ContosoApp/ /app/ContosoApp Step 4: CMD powershell ./setupScript.ps1 ENTRYPOINT ["dotnet", "ContosoApp.dll"]
You need to create a Dockerfile document that meets the following requirements:

≫ Call setupScript.ps1 when the container is built.

≫ Run ContosoApp.dll when the container starts. References:
https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image

NEW QUESTION 102
- (Exam Topic 3)
You are developing an application. You have an Azure user account that has access to two subscriptions. You need to retrieve a storage account key secret from Azure Key Vault.
In which order should you arrange the PowerShell commands to develop the solution? To answer, move all commands from the list of commands to the answer area and arrange them in the correct order.

Powershell commands   **Answer Area**

```
$secretvalue = ConvertTo-SecureString
$storAcctkey -AsPlainText
-Force
     Set-AzKeyVaultSecret -VaultName
$vaultName -Name $secretName
-SecretValue $secretvalue
```

```
Get-AzStorageAccountKey -
ResourceGroupName $resGroup -Name
$storAcct
```

```
Set-AzContext -SubscriptionId
$subscriptionID
```

```
Get-AzKeyVaultSecret -VaultName
$vaultName
```

```
Get-AzSubscription
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Get-AzSubscription
If you have multiple subscriptions, you might have to specify the one that was used to create your key vault. Enter the following to see the subscriptions for your account:
Get-AzSubscription
Step 2: Set-AzContext -SubscriptionId
To specify the subscription that's associated with the key vault you'll be logging, enter: Set-AzContext -SubscriptionId <subscriptionID>
Step 3: Get-AzStorageAccountKey You must get that storage account key.
Step 4: $secretvalue = ConvertTo-SecureString <storageAccountKey> -AsPlainText -Force
Set-AzKeyVaultSecret -VaultName <vaultName> -Name <secretName> -SecretValue $secretvalue After retrieving your secret (in this case, your storage account key), you must convert that key to a secure
string, and then create a secret with that value in your key vault.
Step 5: Get-AzKeyVaultSecret
Next, get the URI for the secret you created. You'll need this URI in a later step to call the key vault and retrieve your secret. Run the following PowerShell command and make note of the ID value, which is the secret's URI:
Get-AzKeyVaultSecret –VaultName <vaultName> Reference:
https://docs.microsoft.com/bs-latn-ba/Azure/key-vault/key-vault-key-rotation-log-monitoring

NEW QUESTION 105
- (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 107
- (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 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. 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 115
- (Exam Topic 3)
You develop Azure solutions.
You must connect to a No-SQL globally-distributed database by using the .NET API. You need to create an object to configure and execute requests in the database. Which code segment should you use?

A. new Container(EndpointUri, PrimaryKey);
B. new Database(Endpoint, PrimaryKey);
C. new CosmosClient(EndpointUri, PrimaryKey);

Answer: C

Explanation:
 Example:
// Create a new instance of the Cosmos Client
this.cosmosClient = new CosmosClient(EndpointUri, PrimaryKey)
//ADD THIS PART TO YOUR CODE
await this.CreateDatabaseAsync(); Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/sql-api-get-started

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 121
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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](https://www.surepassexam.com/AZ-204-exam-dumps.html)