

- **Vendor: Microsoft**
- **Exam Code: AZ-204**
- **Exam Name: Developing Solutions for Microsoft Azure**
- **New Updated Questions from [Braindump2go](#) (Updated in [Nov./2020](#))**

Visit Braindump2go and Download Full Version AZ-204 Exam Dumps

QUESTION 95

Hotspot Question

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.

Answer Area

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

az	<div style="border: 1px solid black; padding: 2px;"> <p>group</p> <p>webapp</p> <p>appservice plan</p> <p>webapp deployment slot</p> <p>webapp deployment source</p> </div>	create --location centralus --name \$resourcegroupname
az	<div style="border: 1px solid black; padding: 2px;"> <p>group</p> <p>webapp</p> <p>appservice plan</p> <p>webapp deployment slot</p> <p>webapp deployment source</p> </div>	create --name \$webappname --resource-group \$resourcegroupname --sku S3
az	<div style="border: 1px solid black; padding: 2px;"> <p>group</p> <p>webapp</p> <p>appservice plan</p> <p>webapp deployment slot</p> <p>webapp deployment source</p> </div>	create --name \$webappname --resource-group \$resourcegroupname --plan \$webappname
az	<div style="border: 1px solid black; padding: 2px;"> <p>group</p> <p>webapp</p> <p>appservice plan</p> <p>webapp deployment slot</p> <p>webapp deployment source</p> </div>	create --name \$webappname --resource-group \$resourcegroupname --slot staging
az	<div style="border: 1px solid black; padding: 2px;"> <p>group</p> <p>webapp</p> <p>appservice plan</p> <p>webapp deployment slot</p> <p>webapp deployment source</p> </div>	config --name \$webappname --resource-group \$resourcegroupname \ --slot staging --repo-uri \$gitrepo --branch master --manual-integration

Answer:

Answer Area

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

az	<ul style="list-style-type: none"> group webapp appservice plan webapp deployment slot webapp deployment source 	create --location centralus --name \$resourcegroupname
az	<ul style="list-style-type: none"> group webapp appservice plan webapp deployment slot webapp deployment source 	create --name \$webappname --resource-group \$resourcegroupname --sku S3
az	<ul style="list-style-type: none"> group webapp appservice plan webapp deployment slot webapp deployment source 	create --name \$webappname --resource-group \$resourcegroupname --plan \$webappname
az	<ul style="list-style-type: none"> group webapp appservice plan webapp deployment slot webapp deployment source 	create --name \$webappname --resource-group \$resourcegroupname --slot staging
az	<ul style="list-style-type: none"> 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

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

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-staging-environment>

QUESTION 96

Hotspot Question

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.

Answer Area

```
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); |   |

  
}
```

Answer:

Answer Area

```
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);
  
}
}
```

QUESTION 97

Drag and Drop Question

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

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

The system has the following requirements for subscriptions:

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

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

Which filter types should you implement? To answer, drag the appropriate filter types to the correct subscriptions. Each filter type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Filter types	Answer Area
SQLFilter	Subscription: FutureOrders, Filter type: filter type
CorrelationFilter	Subscription: HighPriorityOrders, Filter type: filter type
No Filter	Subscription: InternationalOrders, Filter type: filter type
	Subscription: HighQuantityOrders, Filter type: filter type
	Subscription: AllOrders, Filter type: filter type

Answer:

Filter types	Answer Area
SQLFilter	Subscription: FutureOrders, Filter type: SQLFilter
CorrelationFilter	Subscription: HighPriorityOrders, Filter type: CorrelationFilter
No Filter	Subscription: InternationalOrders, Filter type: SQLFilter
	Subscription: HighQuantityOrders, Filter type: SQLFilter
	Subscription: AllOrders, Filter type: No Filter

Explanation:

FutureOrders: SQLFilter

HighPriorityOrders: CorrelationFilter

CorrelationID only

InternationalOrders: SQLFilter

Country NOT USA requires an SQL Filter

HighQuantityOrders: SQLFilter

Need to use relational operators so an SQL Filter is needed.

AllOrders: No Filter

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

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

conditions must match.

Boolean filters - The TrueFilter and FalseFilter either cause all arriving messages (true) or none of the arriving messages (false) to be selected for the subscription.

Reference:

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

QUESTION 98

Drag and Drop Question

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

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

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

Actions	Answer Area
If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.	
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 file from cache if the TTL has not expired.	➤
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:

Actions	Answer Area
	A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.
	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.
	Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the file from cache if the TTL has not expired.

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.

Reference:

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

QUESTION 99

Hotspot Question

A company develops a series of mobile games. All games use a single leaderboard service.

You have the following requirements:

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

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

PartitionKey	RowKey	Email
Harp	Walter	wharp@contoso.com
Smith	Steve	ssmith@contoso.com
Smith	Jeff	jsmith@contoso.com

You develop the following code to save scores in the data store. (Line numbers are included for reference only.)

```
01 public void SaveScore(string gameId, string playerId, int score, long timePlayed)
02 {
03     CloudStorageAccount storageAccount = CloudStorageAccount.Parse(connectionString);
04     CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
05     CloudTable table = tableClient.GetTableReference("scoreTable");
06     table.CreateIfNotExists();
07     var scoreRecord = new PlayerScore(gameId, playerId, score, timePlayed);
08     TableOperation insertOperation = TableOperation.Insert(scoreRecord);
09     table.Execute(insertOperation);
10 }
```

You develop the following code to query the database. (Line numbers are included for reference only.)

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

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
SaveScore will work with Cosmos DB.	<input type="radio"/>	<input type="radio"/>
SaveScore will update and replace a record if one already exists with the same playerId and gameld.	<input type="radio"/>	<input type="radio"/>
Leader board data for the game will be automatically partitioned using gameld.	<input type="radio"/>	<input type="radio"/>
SaveScore will store the values for the gameld and playerId parameters in the database.	<input type="radio"/>	<input type="radio"/>

Answer:
Answer Area

Statements	Yes	No
SaveScore will work with Cosmos DB.	<input checked="" type="radio"/>	<input type="radio"/>
SaveScore will update and replace a record if one already exists with the same playerId and gameld.	<input type="radio"/>	<input checked="" type="radio"/>
Leader board data for the game will be automatically partitioned using gameld.	<input type="radio"/>	<input type="radio"/>
SaveScore will store the values for the gameld and playerId parameters in the database.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Box 1: Yes

Create a table.

A CloudTableClient object lets you get reference objects for tables and entities. The following code creates a CloudTableClient object and uses it to create a new CloudTable object, which represents a table // Retrieve storage account from connection-string.

```
CloudStorageAccount storageAccount =
CloudStorageAccount.parse(storageConnectionString);
```

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

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

```
// Create the table if it doesn't exist.
```

```
String tableName = "people";
```

```
CloudTable cloudTable = tableClient.getTableReference(tableName); cloudTable.createIfNotExists();
```

Box 2: No

New records are inserted with TableOperation.insert. Old records are not updated.

To update old records TableOperation.insertOrReplace should be used instead.

Box 3: No

Box 4: Yes

Reference:

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

QUESTION 100

Hotspot Question

[AZ-204 Exam Dumps](#) [AZ-204 Exam Questions](#) [AZ-204 PDF Dumps](#) [AZ-204 VCE Dumps](#)

<https://www.braindump2go.com/az-204.html>

You are developing a solution that uses the Azure Storage Client library for .NET. You have the following code: (Line numbers are included for reference only.)

```

01 CloudBlockBlob src = null;
02 try
03 {
04     src = container.ListBlobs().OfType<CloudBlockBlob>().FirstOrDefault();
05     var id = await src.AcquireLeaseAsync(null);
06     var dst = container.GetBlockBlobReference(src.Name);
07     string cpid = await dst.StartCopyAsync(src);
08     await dst.FetchAttributeAsync();
09     return id;
10 }
11 catch (Exception e)
12 {
13     throw;
14 }
15 finally
16 {
17     if (src != null)
18         await src.FetchAttributesAsync();
19     if (src.Properties.LeaseState != LeaseState.Available)
20         await src.BreakLeaseAsync(new TimeSpan(0));
21 }

```

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

Answer Area

Statement	Yes	No
The code creates an infinite lease	<input type="radio"/>	<input type="radio"/>
The code at line 06 always creates a new blob	<input type="radio"/>	<input type="radio"/>
The finally block releases the lease	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statement	Yes	No
The code creates an infinite lease	<input checked="" type="radio"/>	<input type="radio"/>
The code at line 06 always creates a new blob	<input type="radio"/>	<input checked="" type="radio"/>
The finally block releases the lease	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Box 1: Yes

AcquireLeaseAsync does not specify leaseTime.

leaseTime is a TimeSpan representing the span of time for which to acquire the lease, which will be rounded down to seconds. If null, an infinite lease will be acquired. If not null, this must be 15 to 60 seconds.

Box 2: No

The GetBlockBlobReference method just gets a reference to a block blob in this container.

Box 3: Yes

The BreakLeaseAsync method initiates an asynchronous operation that breaks the current lease on this container.

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.getblockblobreference>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync>

QUESTION 101

Hotspot Question

You are preparing to deploy a Python website to an Azure Web App using a container. The solution will use multiple containers in the same container group. The Dockerfile that builds the container is as follows:

```
FROM python:3
ADD website.py
CMD [ "python", "./website.py" ]
```

You build a container by using the following command. The Azure Container Registry instance named images is a private registry.

```
docker build -t images.azurecr.io/website:v1.0.0
```

The user name and password for the registry is admin.

The Web App must always run the same version of the website regardless of future builds.

You need to create an Azure Web App to run the website.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
az configure --defaults web=website
az configure --defaults group=website
az appservice plan create --name websitePlan
```

--sku SHARED
--tags container
--sku B1 --hyper-v
--sku B1 --is-linux

```
az webapp create --plan websitePlan
```

--deployment-source-url images.azurecr.io/website:v1.0.0
--deployment-source-url images.azurecr.io/website:latest
--deployment-container-image-name images.azurecr.io/website:v1.0.0
--deployment-container-image-name images.azurecr.io/website:latest

```
az webapp config
```

set --python-version 2.7 --generic-configurations user=admin password=admin
set --python-version 3.6 --generic-configurations user=admin password=admin
container set --docker-registry-server-url https://images.azurecr.io --u admin -p admin
container set --docker-registry-server-url https://images.azurecr.io/wsebsite -u admin -p admin

Answer:

Answer Area

```

az configure --defaults web=website
az configure --defaults group=website
az appservice plan create --name websitePlan

```

--sku SHARED
--tags container
--sku B1 --hyper-v
--sku B1 --is-linux

```

az webapp create --plan websitePlan

```

--deployment-source-url images.azurecr.io/website:v1.0.0
--deployment-source-url images.azurecr.io/website:latest
--deployment-container-image-name images.azurecr.io/website:v1.0.0
--deployment-container-image-name images.azurecr.io/website:latest

```

az webapp config

```

set --python-version 2.7 --generic-configurations user=admin password=admin
set --python-version 3.6 --generic-configurations user=admin password=admin
container set --docker-registry-server-url https://images.azurecr.io -u admin -p admin
container set --docker-registry-server-url https://images.azurecr.io/wsebsite -u admin -p admin

Explanation:

Box 1: --SKU B1 --hyper-v
--hyper-v

Host web app on Windows container.

Box 2: --deployment-source-url images.azurecr.io/website:v1.0.0 --deployment-source-url -u
Git repository URL to link with manual integration.

The Web App must always run the same version of the website regardless of future builds.

Incorrect:

--deployment-container-image-name -i

Linux only. Container image name from Docker Hub, e.g. publisher/image-name:tag.

Box 3: az webapp config container set -url https://images.azurecr.io -u admin -p admin

az webapp config container set

Set a web app container's settings.

Parameter: --docker-registry-server-url -r

The container registry server url.

The Azure Container Registry instance named images is a private registry.

Example:

az webapp config container set --docker-registry-server-url https://{azure-container-registry-name}.azurecr.io

Reference:

<https://docs.microsoft.com/en-us/cli/azure/appservice/plan>

QUESTION 102

Hotspot Question

You are building a website to access project data related to teams within your organization. The website does not allow anonymous access. Authentication is 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 configuration in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```

{
  ....
  "appld": "d61126e3-089b-4adb-b721-d5023213df7d",
  "displayName": "internal",
  [dropdown] : "All",
  "optionalClaims"
  "groupMembershipClaims"
  [dropdown] : true
  "allowPublicClient"
  "oauth2Permissions"
  "requiredResourceAccess"
  "oauth2AllowImplicitFlow"
}
  
```

Answer:

Answer Area

```

{
  ....
  "appld": "d61126e3-089b-4adb-b721-d5023213df7d",
  "displayName": "internal",
  [dropdown] : "All",
  "optionalClaims"
  "groupMembershipClaims"
  [dropdown] : true
  "allowPublicClient"
  "oauth2Permissions"
  "requiredResourceAccess"
  "oauth2AllowImplicitFlow"
}
  
```

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.

Incorrect Answers:

oauth2AllowImplicitFlow. oauth2AllowImplicitFlow specifies whether this web app can request OAuth2.0 implicit flow access tokens. The default is false. This flag is used for browser-based apps, like Javascript single-page apps.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-fed-group-claims>

QUESTION 103

Hotspot Question

You are debugging an application that is running on Azure Kubernetes cluster named cluster1. The cluster uses Azure Monitor for containers to monitor the cluster.

The application has sticky sessions enabled on the ingress controller.

Some customers report a large number of errors in the application over the last 24 hours.

You need to determine on which virtual machines (VMs) the errors are occurring.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
let startTimestamp =  ;
| 

|                  |
|------------------|
| ago(1d)          |
| since(1d)        |
| totimespan(1d)   |
| date(now() - 1d) |

 ;

let ContainerIDs = KubePodInventory
| where ClusterName == "Cluster1"
| 

|                      |
|----------------------|
| top ContainerID      |
| union ContainerID    |
| sample ContainerID   |
| distinct ContainerID |

 ;

ContainerLog
| 

|                                             |
|---------------------------------------------|
| fork containerIDs                           |
| where ContainerID in (ContainerIDs)         |
| restrict ContainerID in (ContainerIDs)      |
| join ContainerID = ContainerIDs.ContainerID |


| where TimeGenerated > startTimestamp
| where LogEntrySource == "stderr"
| 

|                               |
|-------------------------------|
| project by Computer           |
| summarize by Computer         |
| partition count() by Computer |
| summarize count() by Computer |


```

Answer:

Answer Area

```

let startTimestamp = [dropdown]
    ago(1d)
    since(1d)
    totimespan(1d)
    date(now() - 1d)

let ContainerIDs = KubePodInventory
| where ClusterName == "Cluster1"
| [dropdown]
    top ContainerID
    union ContainerID
    sample ContainerID
    distinct ContainerID

ContainerLog
| [dropdown]
    fork containerIDs
    where ContainerID in (ContainerIDs)
    restrict ContainerID in (ContainerIDs)
    join ContainerID == ContainerIDs.ContainerID

| where TimeGenerated > startTimestamp
| where LogEntrySource == "stderr"
| [dropdown]
    project by Computer
    summarize by Computer
    partition count() by Computer
    summarize count() by Computer
  
```

Explanation:

Box 1: ago(1d)

Box 2: distinct containerID

Box 3: where ContainerID in (ContainerIDs)

Box 4: summarize Count by Computer

Summarize: aggregate groups of rows

Use summarize to identify groups of records, according to one or more columns, and apply aggregations to them. The most common use of summarize is count, which returns the number of results in each group.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/get-started-queries>

<https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/query-optimization>

QUESTION 104

Hotspot Question

You plan to deploy a web app to App Service on Linux. You create an App Service plan. You create and push a custom Docker image that contains the web app to Azure Container Registry.

You need to access the console logs generated from inside the container in real-time.

[AZ-204 Exam Dumps](#) [AZ-204 Exam Questions](#) [AZ-204 PDF Dumps](#) [AZ-204 VCE Dumps](#)

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.

Answer Area

az webapp log --name ContosoWeb --resource-group ContosoDevRG

config
download
show
tail

filesystem

--web-server-logging
--docker-container-logging
--application-logging

az log --name ContosoWeb --resource-group ContosoDevRG

webapp
acr
aks

config
download
show
tail

Answer:

Answer Area

az webapp log --name ContosoWeb --resource-group ContosoDevRG

config
download
show
tail

filesystem

--web-server-logging
--docker-container-logging
--application-logging

az log --name ContosoWeb --resource-group ContosoDevRG

webapp
acr
aks

config
download
show
tail

Explanation:

Box 1: config

To Configure logging for a web app use the command:

az webapp log config

Box 2: --docker-container-logging

Syntax include:

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

Box 3: webapp

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

az webapp log download

Box 4: download

Reference:

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

QUESTION 105

Drag and Drop Question

A web service provides customer summary information for e-commerce partners. The web service is implemented as an Azure Function app with an HTTP trigger. Access to the API is provided by an Azure API Management instance. The API Management instance is configured in consumption plan mode. All API calls are authenticated by using OAuth. API calls must be cached. Customers must not be able to view cached data for other customers.

[AZ-204 Exam Dumps](#) [AZ-204 Exam Questions](#) [AZ-204 PDF Dumps](#) [AZ-204 VCE Dumps](#)

You need to configure API Management policies for caching.
How should you complete the policy statement?

Targets	Answer Area
Expect	<pre> <policies> <inbound> <base /> <cache-lookup caching-type=" Target " downstream-caching-type = " Target " > <vary-by-header> Target </vary-by-header> </cache-lookup> </inbound> </policies> </pre>
Public	
Private	
Internal	
External	
Authorization	

Answer:

Targets	Answer Area
Expect	<pre> <policies> <inbound> <base /> <cache-lookup caching-type=" Internal " downstream-caching-type = " Target " > <vary-by-header> Authorization </vary-by-header> </cache-lookup> </inbound> </policies> </pre>
Public	
Private	
External	
Authorization	

Explanation:

Box 1: internal
caching-type

Choose between the following values of the attribute:

internal to use the built-in API Management cache, external to use the external cache as Azure Cache for Redis prefer-external to use external cache if configured or internal cache otherwise.

Box 2: private

downstream-caching-type

This attribute must be set to one of the following values.

none - downstream caching is not allowed.

private - downstream private caching is allowed.

public - private and shared downstream caching is allowed.

Box 3: Authorization

<vary-by-header>Authorization</vary-by-header>

<!-- should be present when allow-private-response-caching is "true"-->

Note: Start caching responses per value of specified header, such as Accept, Accept-Charset, Accept- Encoding, Accept-Language, Authorization, Expect, From, Host, If-Match

Reference:

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

QUESTION 106

Hotspot Question

A company is developing a gaming platform. Users can join teams to play online and see leaderboards that include player statistics. The solution includes an entity named Team.

You plan to implement an Azure Redis Cache instance to improve the efficiency of data operations for entities that rarely change.

You need to invalidate the cache when team data is changed.

How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
void ClearCachedTeams()
{
    IDatabase cache = Connection.GetDatabase();
    ICache cache = Connection.GetDatabase();

    cache.KeyDelete("Team");
    cache.StringSet("Team", "");
    cache.ValueDelete("Team");
    cache.StringGet("Team", "");

    ViewBag.msg += "Team data removed from cache.";
}
```

Answer:

Answer Area

```
void ClearCachedTeams()
{
    IDatabase cache = Connection.GetDatabase();
    ICache cache = Connection.GetDatabase();

    cache.KeyDelete("Team");
    cache.StringSet("Team", "");
    cache.ValueDelete("Team");
    cache.StringGet("Team", "");

    ViewBag.msg += "Team data removed from cache.";
}
```

Explanation:

Box 1: IDatabase cache = connection.GetDatabase();

Connection refers to a previously configured ConnectionMultiplexer.

Box 2: `cache.StringSet("teams",")`

To specify the expiration of an item in the cache, use the TimeSpan parameter of StringSet. `cache.StringSet("key1", "value1", TimeSpan.FromMinutes(90));`

Reference:

<https://azure.microsoft.com/sv-se/blog/lap-around-azure-redis-cache-preview/>

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

QUESTION 107

Drag and Drop Question

You have an application that provides weather forecasting data to external partners. You use Azure API Management to publish APIs.

You must change the behavior of the API to meet the following requirements:

- Support alternative input parameters
- Remove formatting text from responses
- Provide additional context to back-end services

Which types of policies should you implement? To answer, drag the policy types to the correct scenarios. Each policy type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

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

Answer:

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

QUESTION 108

Drag and Drop Question

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 blobs to Archive tier after they have not been modified within 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.

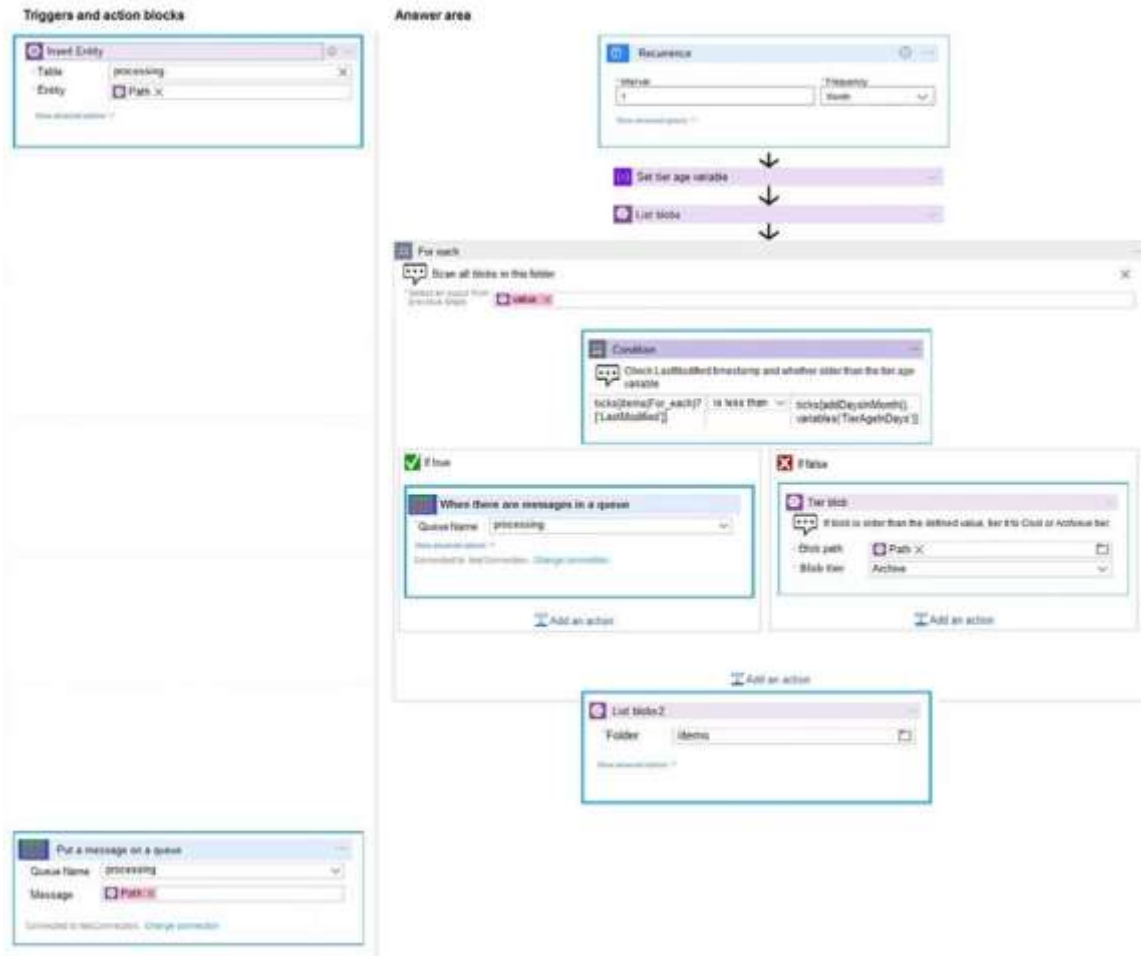
NOTE: Each correct selection is worth one point.

Triggers and action blocks

- Typed Entry**
Table: processing
Entry: Path X
- Tim block**
If time is older than the defined value, Set its Cost or Archive Set
Block path: Path X
Block size: Archive
- When there are messages in a queue**
Queue Name: processing
- Recursive**
Name: i
Frequency: Hour
- List files 2**
Folder: ibetis
- Condition**
Check LastModified timestamps and whether older than the file age variable
IsOlderThanFor_each? Is less than? (to)(getDaysOfMonth) variables/TimeAgeDays
- Put a message on a queue**
Queue Name: processing
Message: Path X



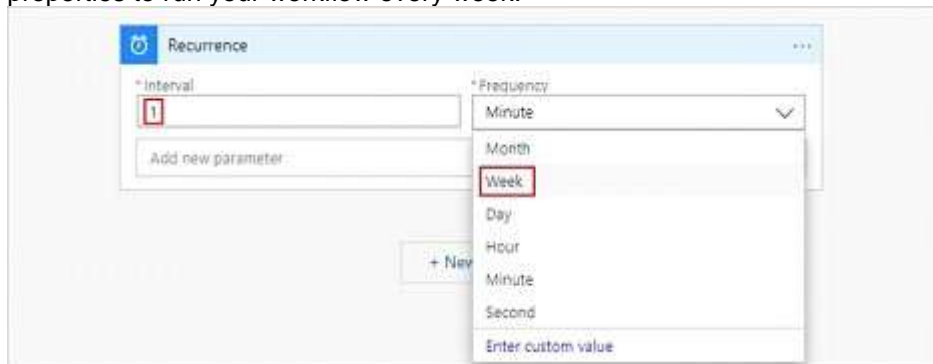
Answer:



Explanation:

Box 1: Reoccurrence..

To regularly run tasks, processes, or jobs on specific schedule, you can start your logic app workflow with the built-in Recurrence - Schedule trigger. You can set a date and time as well as a time zone for starting the workflow and a recurrence for repeating that workflow. Set the interval and frequency for the recurrence. In this example, set these properties to run your workflow every week.



Box 2: Condition..

To run specific actions in your logic app only after passing a specified condition, add a conditional statement. This control structure compares the data in your workflow against specific values or fields. You can then specify different actions that run based on whether or not the data meets the condition.

Box 3: Put a message on a queue

The path for any item that is not archived must be placed in an existing queue.

Note: Under If true and If false, add the steps to perform based on whether the condition is met.

Box 4: ..tier it to Cool or Archive tier.

Archive item.

Box 5: List blobs 2

Reference:

<https://docs.microsoft.com/en-us/azure/connectors/connectors-native-recurrence>

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-control-flow-loops>

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-control-flow-conditional-statement>