

AZ-204^{Q&As}

Developing Solutions for Microsoft Azure

Pass Microsoft AZ-204 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/az-204.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

DRAG DROP

You are preparing to deploy an Azure virtual machine (VM)-based application.

The VMs that run the application have the following requirements:

1.
When a VM is provisioned the firewall must be automatically configured before it can access Azure resources.
2.
Supporting services must be installed by using an Azure PowerShell script that is stored in Azure Storage.

You need to ensure that the requirements are met.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Features	Answer Area	
Run Command	Requirement	Feature
Serial console	Firewall configuration	
Hybrid Runbook Worker	Supporting services script	
Custom Script Extension		

Correct Answer:

Features	Requirement	Feature
	Firewall configuration	Run Command
Serial console		
	Supporting services script	Hybrid Runbook Worker
Custom Script Extension		

Reference: <https://docs.microsoft.com/en-us/azure/automation/automation-hybrid-runbook-worker>
<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-command>

QUESTION 2

You need to ensure the security policies are met.

What code do you add at line CS07 of ConfigureSSE.ps1?

- A. -PermissionsToKeys create, encrypt, decrypt
- B. -PermissionsToCertificates create, encrypt, decrypt
- C. -PermissionsToCertificates wrapkey, unwrapkey, get
- D. -PermissionsToKeys wrapkey, unwrapkey, get

Correct Answer: B

Scenario: All certificates and secrets used to secure data must be stored in Azure Key Vault.

You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.

The Set-AzureRmKeyVaultAccessPolicy parameter -PermissionsToKeys specifies an array of key operation permissions to grant to a user or service principal. The acceptable values for this parameter: decrypt, encrypt, unwrapKey, wrapKey,

verify, sign, get, list, update, create, import, delete, backup, restore, recover, purge

Incorrect Answers:

A, C: The Set-AzureRmKeyVaultAccessPolicy parameter -PermissionsToCertificates specifies an array of certificate permissions to grant to a user or service principal. The acceptable values for this parameter: get, list, delete, create, import,

update, managecontacts, getissuers, listissuers, setissuers, deleteissuers, manageissuers, recover, purge, backup, restore

Reference:

<https://docs.microsoft.com/en-us/powershell/module/azurermskeyvault/set-azurermskeyvaultaccesspolicy>

QUESTION 3

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

You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is General-purpose V2.

When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.

You need to design the process that starts the photo processing.

Solution: Use the Azure Blob Storage change feed to trigger photo processing.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

The change feed is a log of changes that are organized into hourly segments but appended to and updated every few minutes. These segments are created only when there are blob change events that occur in that hour.

Instead catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload.

Reference: <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed>
<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

QUESTION 4

You are developing a web application that uses the Microsoft identity platform for user and resource authentication. The web application calls several REST APIs.

A REST API call must read the user's calendar. The web application requires permission to send an email as the user.

You need to authorize the web application and the API.

Which parameter should you use?

- A. tenant
- B. code_challenge
- C. state
- D. client_id
- E. scope

Correct Answer: E

Microsoft identity platform and OAuth 2.0 authorization code flow, Request an authorization code

<https://login.microsoftonline.com/{tenant}/oauth2/v2.0/authorize?>

The authorization code flow begins with the client directing the user to the /authorize endpoint. In this request, the client requests the openid, offline_access, and <https://graph.microsoft.com/mail.read> permissions from the user.

Parameters include:

*

scope required

A space-separated list of scopes that you want the user to consent to. For the /authorize leg of the request, this parameter can cover multiple resources. This value allows your app to get consent for multiple web APIs you want to call.

Incorrect:

*

tenant required

The {tenant} value in the path of the request can be used to control who can sign into the application. Valid values are common, organizations, consumers, and tenant identifiers. For guest scenarios where you sign a user from one tenant into

another tenant, you must provide the tenant identifier to sign them into the resource tenant.

*

code_challenge recommended / required

Used to secure authorization code grants by using Proof Key for Code Exchange (PKCE). Required if code_challenge_method is included. This parameter is now recommended for all application types, both public and confidential clients, and

required by the Microsoft identity platform for single page apps using the authorization code flow.

*

client_id

The Application (client) ID that the Azure portal – App registrations experience assigned to your app.

Reference: <https://learn.microsoft.com/en-us/azure/active-directory/develop/v2-oauth2-auth-code-flow>

QUESTION 5

Your company has an Azure Kubernetes Service (AKS) cluster that you manage from an Azure AD-joined device. The cluster is located in a resource group.

Developers have created an application named MyApp. MyApp was packaged into a container image.

You need to deploy the YAML manifest file for the application.

Solution: You install the docker client on the device and run the docker run -it microsoft/azure-cli:0.10.17 command.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

QUESTION 6

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 a medical records document management website. The website is used to store scanned copies of patient intake forms.

If the stored intake forms are downloaded from storage by a third party, the contents of the forms must not be compromised.

You need to store the intake forms according to the requirements.

Solution: Store the intake forms as Azure Key Vault secrets.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

QUESTION 7

HOTSPOT

You are developing a ticket reservation system for an airline.

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

1.

Ensure at least 99.99% availability and provide low latency.

2.

Accept reservations event when localized network outages or other unforeseen failures occur.

3.

Process reservations in the exact sequence as reservations are submitted to minimize overbooking or selling the same seat to multiple travelers.

4.

Allow simultaneous and out-of-order reservations with a maximum five-second tolerance window.

You provision a resource group named `airlineResourceGroup` in the Azure South-Central US region.

You need to provision a SQL SPI Cosmos DB account to support the app.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

	▼
Strong	
Eventual	
ConsistentPrefix	
BoundedStaleness	

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

	▼
--enable-virtual-network true\	
--enable-automatic-failover true\	
--kind 'GlobalDocumentDB' \	
--kind 'MongoDB'\	

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

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

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

Correct Answer:

Answer Area

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

▼
Strong
Eventual
ConsistentPrefix
BoundedStaleness

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

▼
--enable-virtual-network true\
--enable-automatic-failover true\
--kind 'GlobalDocumentDB' \
--kind 'MongoDB'\

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

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

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

Box 1: BoundedStaleness

Bounded staleness: The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions (that is, "updates") of an item or by "T" time interval. In other words, when you choose

bounded staleness, the "staleness" can be configured in two ways:

The number of versions (K) of the item

The time interval (T) by which the reads might lag behind the writes

Incorrect Answers:

Strong

Strong consistency offers a linearizability guarantee. Linearizability refers to serving requests concurrently. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write.

Users are always guaranteed to read the latest committed write.

Box 2: --enable-automatic-failover true\

For multi-region Cosmos accounts that are configured with a single-write region, enable automatic-failover by using Azure CLI or Azure portal. After you enable automatic failover, whenever there is a regional disaster, Cosmos DB will automatically failover your account.

QUESTION 8

HOTSPOT

You are developing an application to collect the following telemetry data for delivery drivers: first name, last name, package count, item id, and current location coordinates. The app will store the data in Azure Cosmos DB.

You need to configure Azure Cosmos DB to query the data.

Which values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Configuration Parameter

Value

Azure Cosmos DB API

	▼
Gremlin	
Table API	
Core (SQL)	

Azure Cosmos DB partition key

	▼
first name	
last name	
package count	
item id	

Correct Answer:

Answer Area

Configuration Parameter

Value

Azure Cosmos DB API

	▼
Gremlin	
Table API	
Core (SQL)	

Azure Cosmos DB partition key

	▼
first name	
last name	
package count	
item id	

Box 1: Core (SQL)

Core(SQL) API stores data in document format. It offers the best end-to-end experience as we have full control over the interface, service, and the SDK client libraries. SQL API supports analytics and offers performance isolation between operational and analytical workloads.

Box 2: item id

item id is a unique identifier and is suitable for the partition key.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>

<https://docs.microsoft.com/en-us/azure/cosmos-db/partitioning-overview>

QUESTION 9

DRAG DROP

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?

Put the actions in the correct order.

Select and Place:

Select these

In the Azure portal, add a Job

In the Azure portal, create a Batch account.

In the Azure portal, create tasks

In the Azure portal, create a pool of compute nodes

Place here

Correct Answer:

Select these

Place here

In the Azure portal, create a Batch account.

In the Azure portal, create a pool of compute nodes

In the Azure portal, add a Job

In the Azure portal, create tasks

With the Azure Portal:

Step 1: In the Azure portal, create a Batch account.

First we create a batch account.

Step 2: In the Azure portal, create a pool of compute nodes

Now that you have a Batch account, create a sample pool of Windows compute nodes for test purposes.

Step 3: In the Azure portal, add a Job.

Now that you have a pool, create a job to run on it. A Batch job is a logical group for one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. Initially the job has no tasks.

Step 4: In the Azure portal, create tasks

Now create sample tasks to run in the job. Typically you create multiple tasks that Batch queues and distributes to run on the compute nodes.

References:

<https://docs.microsoft.com/en-us/azure/batch/quick-create-portal>

QUESTION 10

You are building an application to track cell towers that are available to phones in near real time. A phone will send information to the application by using the Azure Web PubSub service. The data will be processed by using an Azure

Functions app. Traffic will be transmitted by using a content delivery network (CDN).

The Azure function must be protected against misconfigured or unauthorized invocations.

You need to ensure that the CDN allows for the Azure function protection.

Which HTTP header should be on the allowed list?

- A. Authorization
- B. WebHook-Request-Callback
- C. Resource
- D. WebHook-Request-Origin

Correct Answer: D

CloudEvents extension for Azure Web PubSub event handler with HTTP protocol

The Web PubSub service delivers client events to the upstream webhook using the CloudEvents HTTP protocol binding.

Webhook validation

The Webhook validation follows CloudEvents. The request always contains WebHook-Request-Origin: xxx.webpubsub.azure.com in the header.

If and only if the delivery target does allow delivery of the events, it MUST reply to the request by including WebHook-Allowed-Origin header, for example:

WebHook-Allowed-Origin: *

Or:

WebHook-Allowed-Origin: xxx.webpubsub.azure.com

For now, WebHook-Request-Rate and WebHook-Request-Callback are not supported.

Incorrect:

* WebHook-Request-Callback. An optional field that provides the webhook with an alternative to grant permission asynchronously, by way of a HTTP callback.

Reference: <https://learn.microsoft.com/en-us/azure/azure-web-pubsub/reference-cloud-events>

QUESTION 11

You are developing several microservices to run on Azure Container Apps for a company. External TCP ingress traffic from the internet has been enabled for the microservices.

The company requires that the microservices must scale based on an Azure Event Hub trigger.

You need to scale the microservices by using a custom scaling rule.

Which two Kubernetes Event-driven Autoscaling (KEDA) trigger fields should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. metadata
- B. type
- C. authenticationRef
- D. name
- E. metricType

Correct Answer: AB

Explanation: Example This example shows how to convert an Azure Service Bus scaler to a Container Apps scale rule, but you use the same process for any other ScaledObject-based KEDA scaler specification. For authentication, KEDA scaler authentication parameters convert into Container Apps secrets. From the KEDA scaler specification, find the type value. triggers:

```
-type: azure-servicebus
```

```
metadata:
```

```
queueName: my-queue
```

```
namespace: service-bus-namespace
```

```
messageCount: "5"
```

Reference: <https://learn.microsoft.com/en-us/azure/container-apps/scale-app>

QUESTION 12

You need to investigate the Azure Function app error message in the development environment. What should you do?

- A. Connect Live Metrics Stream from Application Insights to the Azure Function app and filter the metrics.
- B. Create a new Azure Log Analytics workspace and instrument the Azure Function app with Application Insights.
- C. Update the Azure Function app with extension methods from Microsoft.Extensions.Logging to log events by using the log instance.
- D. Add a new diagnostic setting to the Azure Function app to send logs to Log Analytics.

Correct Answer: A

Azure Functions offers built-in integration with Azure Application Insights to monitor functions.

The following areas of Application Insights can be helpful when evaluating the behavior, performance, and errors in your functions:

Live Metrics: View metrics data as it's created in near real-time. Failures Performance Metrics

Reference: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-monitoring>

QUESTION 13

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:

1.
Share session state across all ASP.NET web applications
2.
Support controlled, concurrent access to the same session state data for multiple readers and a single writer
3.
Save full HTTP responses for concurrent requests

You need to store the information.

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

Does the solution meet the goal?

- A. Yes

B. No

Correct Answer: B

Instead use Azure Cache for Redis.

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

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

QUESTION 14

HOTSPOT

You need to retrieve the database connection string.

Which values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

REST API Endpoint:

https:// /

cpandlkeyvault
PostgreSQLConn
80df3e46ffcd4f1cb187f79905e9a1e8

Variable type to access Azure Key Vault secret values:

<input type="text"/>
Environment
Session
ViewState
QueryString

Correct Answer:

REST API Endpoint:

https:// /

cpandlkeyvault
PostgreSQLConn
80df3e46ffcd4f1cb187f79905e9a1e8

cpandlkeyvault
PostgreSQLConn
80df3e46ffcd4f1cb187f79905e9a1e8

Variable type to access Azure Key Vault secret values:

<input type="text"/>
Environment
Session
ViewState
Querystring

Azure database connection string retrieve REST API vault.azure.net/secrets/

Box 1: cpandlkeyvault

We specify the key vault, cpandlkeyvault.

Scenario: The database connection string is stored in Azure Key Vault with the following attributes:

Azure Key Vault name: cpandlkeyvault

Secret name: PostgreSQLConn

Id: 80df3e46ffcd4f1cb187f79905e9a1e8

Box 2: PostgreSQLConn

We specify the secret, PostgreSQLConn

Example, sample request:

<https://myvault.vault.azure.net/secrets/mysecretname/4387e9f3d6e14c459867679a90fd0f79?api-version=7.1>

Box 3: Querystring

Reference:

<https://docs.microsoft.com/en-us/rest/api/keyvault/getsecret/getsecret>

QUESTION 15

You are developing an internal website for employees to view sensitive data. The website uses Azure Active Directory (AAD) for authentication.

You need to implement multifactor authentication for the website.

What should you do? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Upgrade to Azure AD Premium.

B. In Azure AD conditional access, enable the baseline policy.

C. In Azure AD, create a new conditional access policy.

D. In Azure AD, enable application proxy.

E. Configure the website to use Azure AD B2C.

Correct Answer: AC

A: Multi-Factor Authentication comes as part of the following offerings:

Azure Active Directory Premium licenses - Full featured use of Azure Multi-Factor Authentication Service (Cloud) or Azure Multi-Factor Authentication Server (On-premises).

Multi-Factor Authentication for Office 365

Azure Active Directory Global Administrators

C: MFA Enabled by conditional access policy. It is the most flexible means to enable two-step verification for your users. Enabling using conditional access policy only works for Azure MFA in the cloud and is a premium feature of Azure AD.

References:

<https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfa-getstarted>

[AZ-204 VCE Dumps](#)

[AZ-204 Practice Test](#)

[AZ-204 Exam Questions](#)