# Microsoft

## Exam Questions AI-200

Developing AI Cloud Solutions on Azure

# About Exambible

***Your Partner of IT Exam***

# Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

# Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

**NEW QUESTION 1**
You are developing a microservices-based application that uses Azure Container Apps The application consists of several containerized services that handle tasks, such as processing orders, managing inventory, and generating reports.
You must secure the container apps. All apps must reside in the same virtual network, share the same Dapr configuration, and share the same logging location. Apps must support the configuration of the amount of memory and compute resources available to containers.
You need to configure the Azure Container App
How should you complete the CLI command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

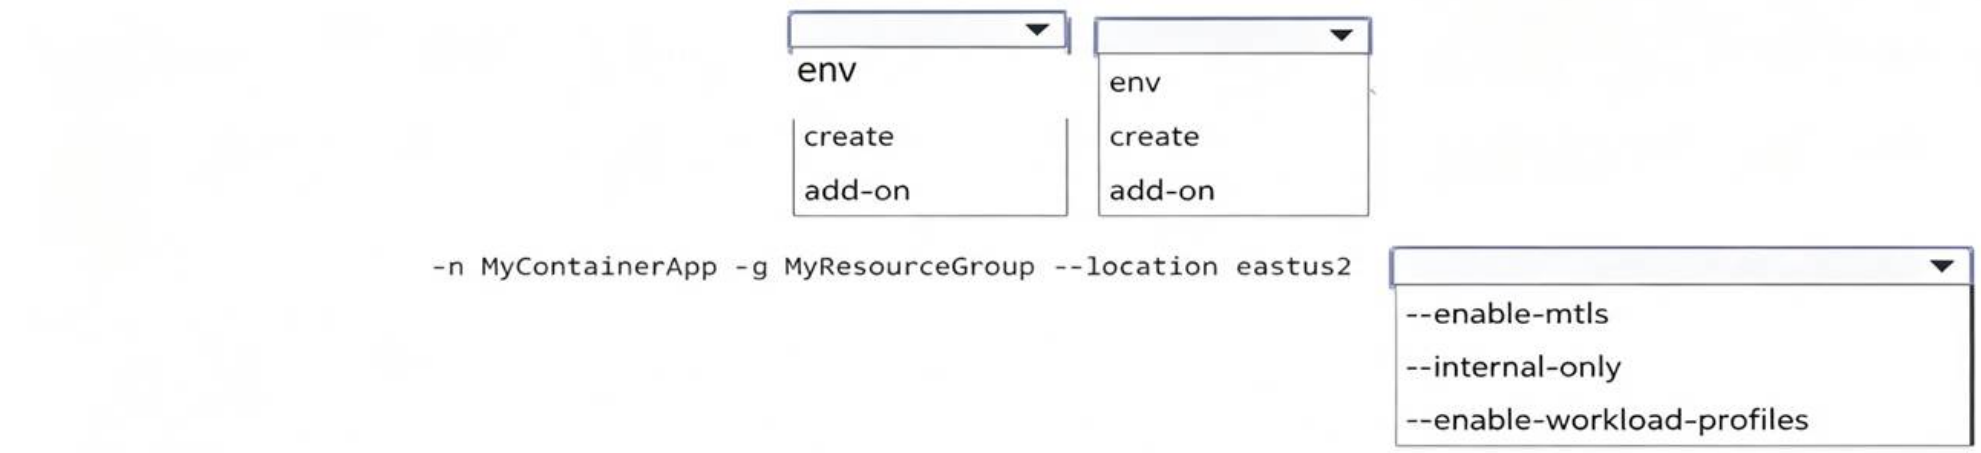


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Verified Answer:`az containerapp env create ... --enable-workload-profiles`.
Detailed Explanation:The shared boundary for Container Apps is the managed environment. Apps in one environment can share the environment??s virtual network integration, observability destination, and Dapr-related environment configuration. Enabling workload profiles provides selectable compute profiles and resource sizing options for workloads in that environment. The CLI must therefore create a Container Apps environment and enable workload profiles. Creating an individual container app alone would not establish the common environment-level network, logging, and resource-profile boundary required by the scenario.
Study Guide Alignment:Containerized Azure workloads: registry builds, App Service containers, Container Apps revision/scaling behavior, and AKS deployment choices.
Official Microsoft Learn References:AI-200 Study Guide | Azure Container Apps environments

**NEW QUESTION 2**
A Python API retrieves a document from Azure Database for PostgreSQL by using a SQL statement. The API accepts the document ID from user input. The current implementation inserts the document ID directly into the SQL statement.
You need to secure the SQL statement execution by using a parameterized query. You must minimize the possibility of SQL injection.
How should you update the implementation to execute the SQL statement safely? To answer, move the appropriate configurations to the correct requirements.
You may use each configuration once, more than once, or not at all. you may need to move the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

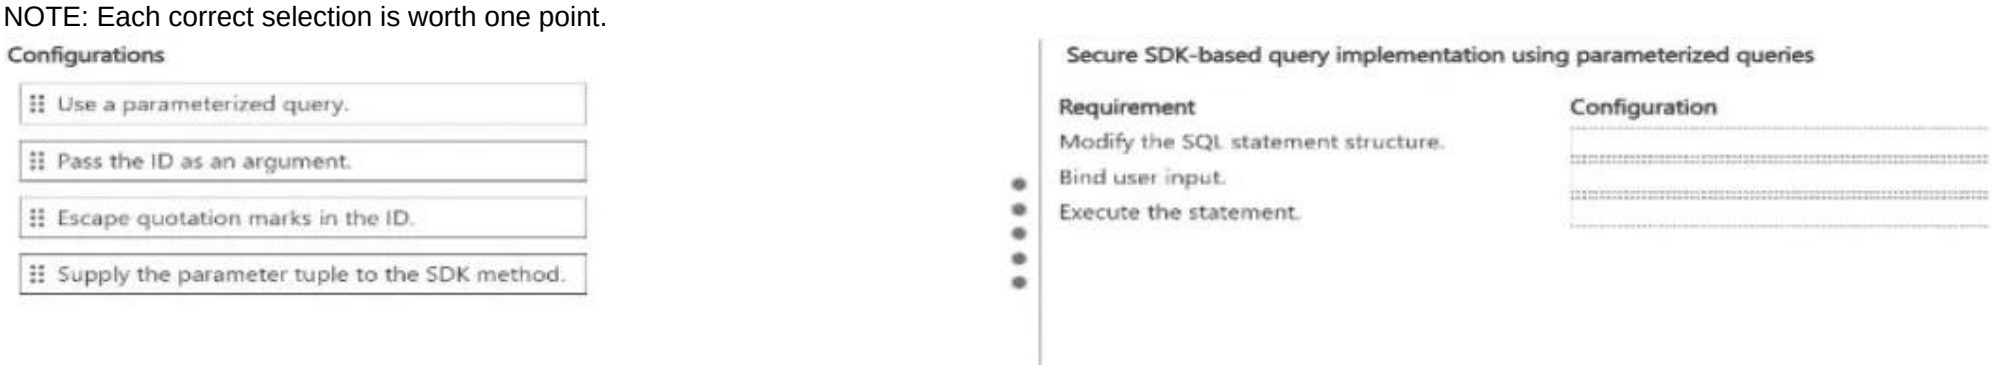


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

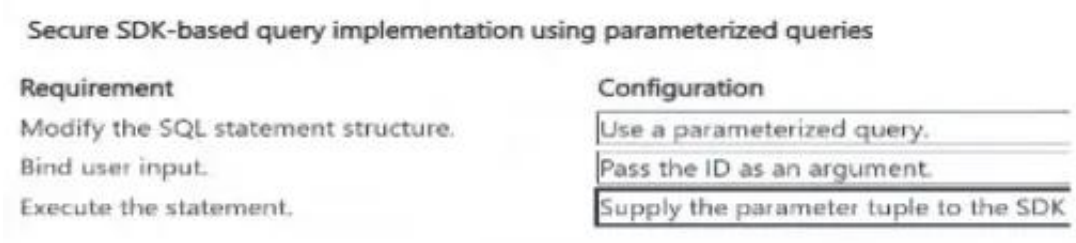


Verified Answer:Modify SQL: use a parameterized query. Bind input: pass the ID as an argument. Execute: supply the parameter tuple to the SDK/driver method.
Detailed Explanation:The security boundary is created by separating SQL syntax from user-supplied values. The statement should contain a parameter placeholder rather than concatenated input, and the document ID should be supplied separately through the database driver??s parameter mechanism. The driver then performs the correct protocol-level binding and escaping. Manually escaping quotation marks is error-prone and does not provide the same injection resistance as true parameterization.

Study Guide Alignment:AI data-management workloads: Cosmos DB, PostgreSQL, caching, vector storage, vector retrieval, consistency, and connection optimization.
Official Microsoft Learn References:AI-200 Study Guide | Vector similarity search with Azure PostgreSQL

**NEW QUESTION 3**
You need to configure vector embedding updates according to the business and technical requirements. Which information should you use? To answer, select the appropriate option in the answer area. NOTE: Each correct selection is worth one point.

**Update vector embeddings**

| Requirement | Configuration |
|---|---|
| Identify document changes that should trigger vectorization. | Change feed processor |
| Scale out the vectorization processing. | Periodic full container scan |
| | Cross-partition SPF CT query |
| Scale out the vectorization processing. | Lease container |
| | Strong consistency level |
| | RU throughput on the container |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
The change feed processor is designed to react to inserts and updates in a monitored Cosmos DB container, which directly matches the requirement to generate embeddings for new or changed documents. Its lease container stores processing state and coordinates work across multiple workers, allowing the processing workload to scale out without duplicating ownership of change-feed ranges. A periodic full scan would consume unnecessary RUs, and neither strong consistency nor container RU throughput is the coordination mechanism for distributed change-feed workers.
Study Guide Alignment: AI data-management workloads: Cosmos DB, PostgreSQL, caching, vector storage, vector retrieval, consistency, and connection optimization.
Official Microsoft Learn References: AI-200 Study Guide | Azure Cosmos DB change feed processor

**NEW QUESTION 4**
You are building a semantic search feature for a chatbot. You store document embeddings in Redis. You review the following Python code that connects to Redis and stores an embedding value:

```
01  import numpy as np
02  from redis import Redis
03  embedding = np.array([0.12, 0.98, 0.44], dtype=np.float32).tobytes()
04  r = Redis(
05      host="myredis.redis.cache.windows.net",
06      port=6380,
07      password="accesskey",
08      ssl=True
09  )
10  r.hset("doc:1", mapping={"embedding": embedding})
11  r.expire("doc:1", 600)
```

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

**Vector search fundamentals**

| Statement | Yes | No |
|---|---|---|
| The embedding is stored as a hash field. | ○ | ○ |
| The code enables similarity search on the embedding field. | ○ | ○ |
| The key will expire after 10 minutes. | ○ | ○ |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
`HSET` stores the binary embedding as a field in the Redis hash keyed by `doc:1`, so the first statement is true. The code does not create a RediSearch/Redis Query Engine vector index or define a vector field schema; merely storing bytes does not enable similarity search, so the second statement is false.
`EXPIRE doc:1 600` sets a 600-second lifetime, which is ten minutes, making the third statement true.
Study Guide Alignment: AI data-management workloads: Cosmos DB, PostgreSQL, caching, vector storage, vector retrieval, consistency, and connection

optimization.
Official Microsoft Learn References: AI-200 Study Guide | Vector search in Azure Managed Redis

**NEW QUESTION 5**
You are implementing an application by using Azure Event Grid to push near-real-time information to customers.
You have the following requirements:
• You must send events to thousands of customers that include hundreds of various event types.
• The events must be filtered by event type before processing.
• Authentication and authorization must be handled by using Microsoft Entra ID.
• The events must be published to a single endpoint
You need to implement Azure Event Grid
Solution: Publish events to a system topic. Create an event subscription for each customer.
Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**Explanation:**
Explanation
System topics represent events emitted by Azure services and are not the correct resource for an application publishing its own hundreds of event types to thousands of customers through one custom endpoint. Event Grid domains are intended for exactly this type of application-level multi-tenant topic organization and expose one publishing endpoint for many topics. Therefore a system topic with a subscription per customer does not satisfy the stated publishing model, even though Event Grid supports filtering and secure delivery in other contexts.
Study Guide Alignment: Azure service integration: Service Bus, Event Grid, Azure Functions triggers/bindings, and event-driven processing.
Official Microsoft Learn References: AI-200 Study Guide | Azure Event Grid event domains
Verification completed using official Microsoft Learn sources only. Original exhibits have been retained unchanged.

**NEW QUESTION 6**
You are reviewing the message processing code in a backend worker service.
You review the following Python code that initializes and starts a Service Bus processor

```
from azure.servicebus import ServiceBusClient, ServiceBusReceiveMode

01 import json
02 from azure.servicebus import ServiceBusClient, ServiceBusReceiveMode
03 from azure.identity import DefaultAzureCredential
04 credential = DefaultAzureCredential()
05 with client.get_queue_receiver(
06     queue_name="inference-requests",
```

**Service Bus processor behavior and SDK defaults**

| Statement | Yes | No |
|---|---|---|
| Messages are removed from the queue as soon as they are received by the processor. | ○ | ○ |
| If message processing fails due to a transient service error, the message can be retried by another consumer. | ○ | ○ |
| Messages with invalid JSON payloads are retried until the maximum delivery count is reached. | ○ | ○ |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

The default Service Bus receive behavior is PeekLock rather than receive-and-delete, so merely receiving a message does not remove it from the queue; it must be settled successfully. If processing fails and the message remains unsettled or the lock is lost, it can become available for another delivery. Invalid JSON, however, is an application-level content problem: Service Bus does not inspect the payload as JSON and automatically retry it to MaxDeliveryCount solely because deserialization fails. The application??s settlement/error path determines what happens.
Study Guide Alignment: Azure service integration: Service Bus, Event Grid, Azure Functions triggers/bindings, and event-driven processing.
Official Microsoft Learn References: AI-200 Study Guide | Service Bus message transfers, locks, and settlement | Service Bus messages, routing, and correlation

**NEW QUESTION 7**
You are developing several microservices to run on Azure Container Apps.
The microservices must allow HTTPS access by using a custom domain.
You need to configure the custom domain in Azure Container Apps.
In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

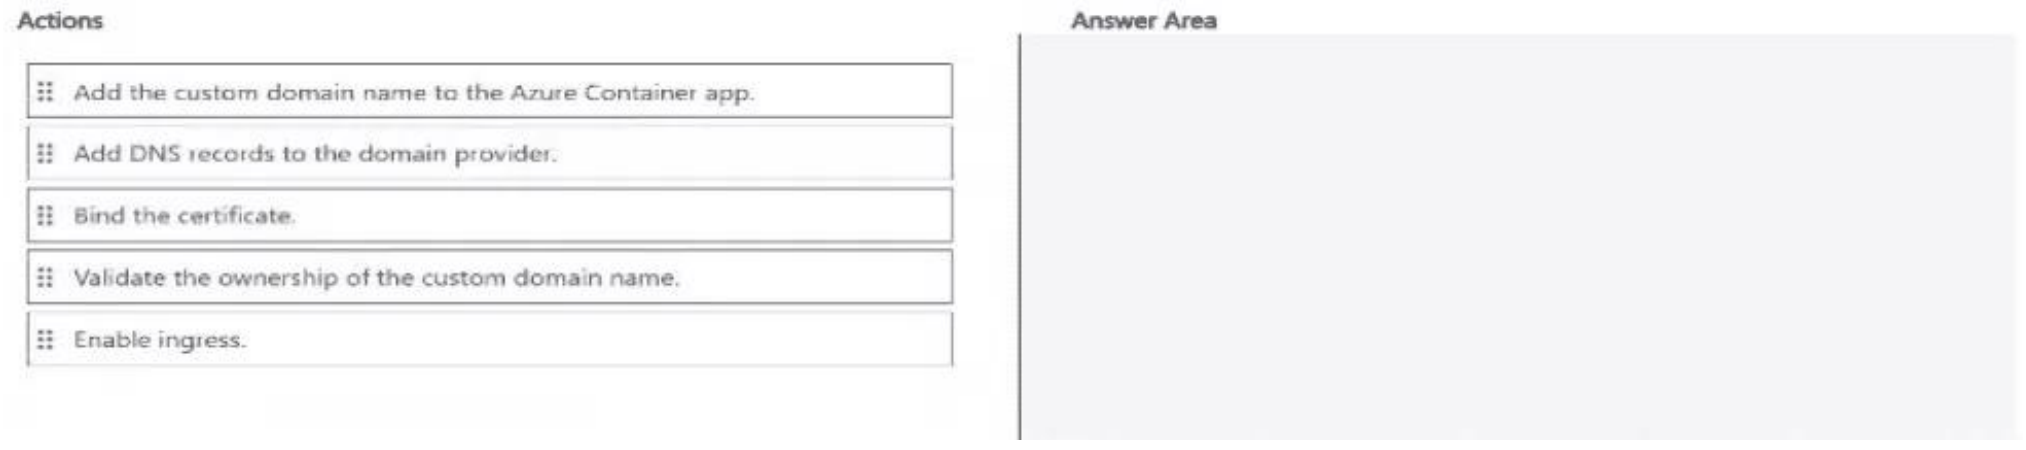


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

A custom HTTPS hostname requires an ingress endpoint, DNS proof that the requester controls the hostname, the hostname association itself, and a TLS certificate bound to that hostname. DNS records must exist before ownership validation can succeed. Once ownership is validated, the domain can be associated with the Container App and the certificate can be bound to provide HTTPS. Skipping ingress or DNS validation would prevent the hostname from being correctly routed and secured.
Study Guide Alignment: Containerized Azure workloads: registry builds, App Service containers, Container Apps revision/scaling behavior, and AKS deployment choices.
Official Microsoft Learn References: AI-200 Study Guide | Custom domains and certificates in Container Apps

**NEW QUESTION 8**
You are using the Python SDK to develop a containerized AI application that retrieves the value of a runtime setting stored in an Azure App Configuration resource. The application will run in Azure in the security context of a managed identity.
The application must work in a local development environment without any code changes.
You need to complete the code that implements the retrieval.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

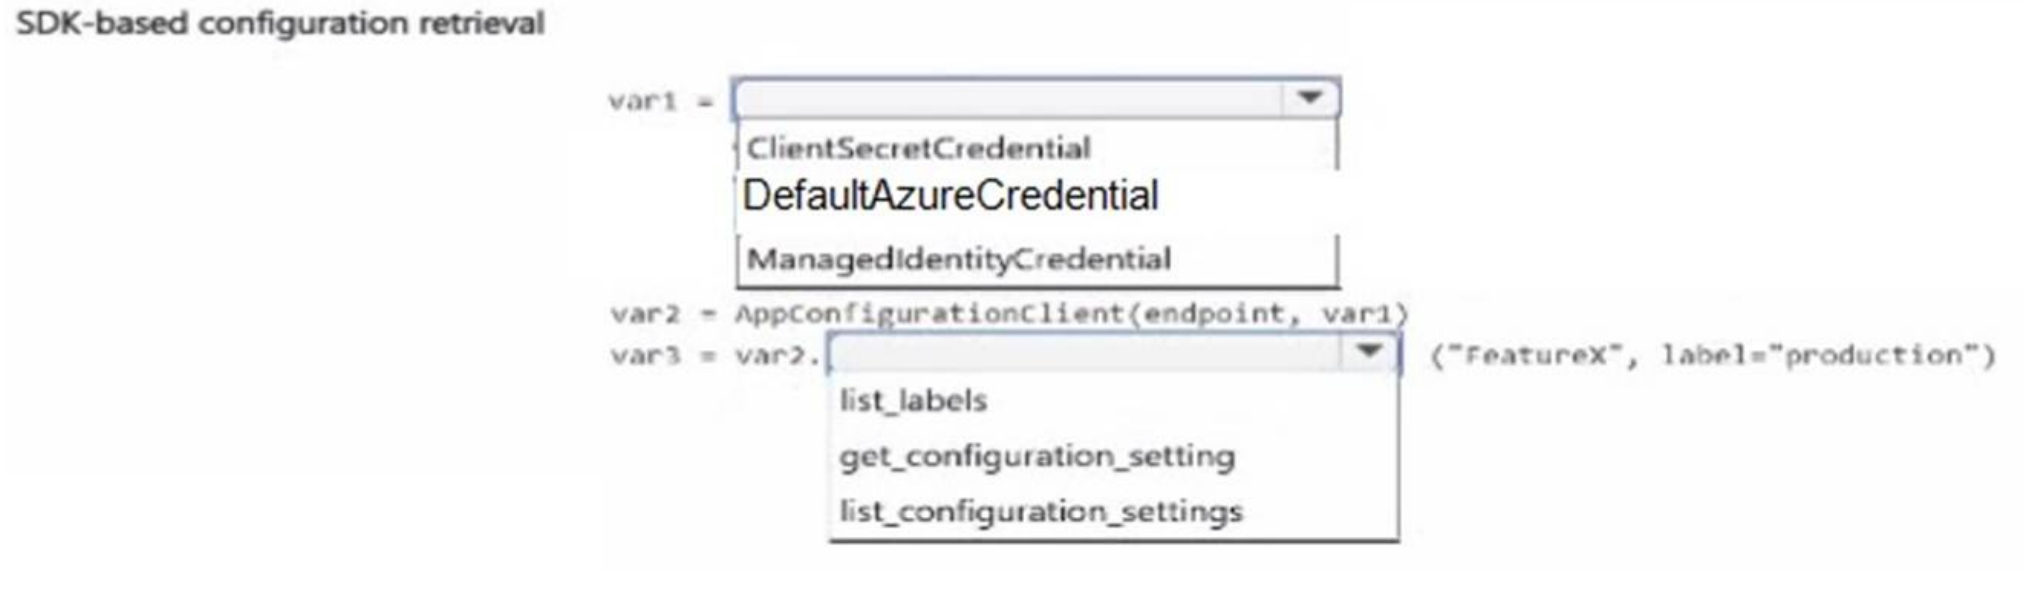


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

DefaultAzureCredential is designed to support the same application code across local development and Azure-hosted environments. Locally it can use developer credentials, while in Azure it can use the workload??s managed identity. The App Configuration client??s get_configuration_setting method retrieves a specific key/label pair, which matches the code shown. Using ManagedIdentityCredential directly would not provide the same local-development credential chain without code or environment-specific handling.
Study Guide Alignment: Security and operations: Key Vault, App Configuration, managed identity, OpenTelemetry, Azure Monitor, and KQL-based troubleshooting.
Official Microsoft Learn References: AI-200 Study Guide | Azure App Configuration Python client | Managed identities for Azure resources

**NEW QUESTION 9**
You are developing a Java application to be deployed in Azure. The application stores sensitive data in Azure Cosmos DB.
You need to configure Always Encrypted to encrypt the sensitive data inside the application. What should you do first?

A. Create a Microsoft Entra ID managed identity and assign the identity to a new Azure Key Vault instance.
B. Create a new container to include an encryption policy with the JSON properties to be encrypted.
C. Create a customer-managed key (CMK) and store the key in a new Azure Key Vault instance.
D. Create a data encryption key (DEK) by using the Azure Cosmos DB SDK and store the key in Azure Cosmos DB.

**Answer:** C

**Explanation:**
For Cosmos DB Always Encrypted, the key hierarchy starts with a customer-managed key held in Azure Key Vault. The data encryption key used by the client is protected by that customer- managed key. Microsoft??s setup sequence therefore begins by preparing Key Vault and the CMK before creating DEKs or an encryption policy. Creating a managed identity, a container policy, or a DEK can be necessary later, but none precedes establishing the CMK that protects the client-side encryption key material.
Study Guide Alignment: AI data-management workloads: Cosmos DB, PostgreSQL, caching, vector storage, vector retrieval, consistency, and connection optimization.
Official Microsoft Learn References: AI-200 Study Guide | Always Encrypted for Azure Cosmos DB

**NEW QUESTION 10**
You are designing an Azure Function app that exposes a public API. The solution must:
Validate incoming request data and return results immediately to the caller. Support Microsoft Entra ID authentication.
Scale automatically under variable load. You need to implement a trigger. Which trigger should you implement?

A. HTTP
B. Azure Queue storage
C. Azure Event Grid
D. Service Bus topic

**Answer:** A

**Explanation:**
An HTTP trigger is the only option that directly accepts a synchronous external request and can return a response immediately to the caller. Azure Functions uses HTTP triggers to implement serverless APIs and webhooks, while Microsoft Entra authentication can be applied to the Function App or API surface. Queue, Event Grid, and Service Bus triggers are asynchronous event-processing mechanisms and do not provide the request/response semantics required by this public API scenario.
Study Guide Alignment: Azure service integration: Service Bus, Event Grid, Azure Functions triggers /bindings, and event-driven processing.
Official Microsoft Learn References: AI-200 Study Guide | Azure Functions HTTP trigger

**NEW QUESTION 10**
You deploy a new revision of an app in Azure Container Apps.
You need to gradually shift production traffic to the revision while monitoring performance. In addition, you need to be able to quickly roll back. Which two actions should you perform?
Each correct answer presents part of the solution. Choose two. NOTE: Each correct selection is worth one point.

A. Use traffic splitting
B. Increase replica count.
C. Restart the revision.
D. Use single revision mode
E. Enable multiple revision mode.

**Answer:** AE

**Explanation:**
Gradual production rollout requires more than one revision to remain active and a mechanism to divide traffic between them. Multiple revision mode permits concurrent active revisions, and traffic splitting lets a specified percentage of requests flow to the new revision while the remainder stays on the proven revision. This also enables rapid rollback by returning traffic to the prior revision. Increasing replicas or restarting a revision changes capacity or process state, not release traffic distribution.
Study Guide Alignment: Containerized Azure workloads: registry builds, App Service containers, Container Apps revision/scaling behavior, and AKS deployment choices.
Official Microsoft Learn References: AI-200 Study Guide | Azure Container Apps revisions

**NEW QUESTION 15**
You are developing an AI-powered API that retrieves connection strings and API keys from Azure Key Vault. You must configure a solution that provides the following security functionality:
• The API must authenticate to Key Vault without storing credentials in any application configuration files
• The identity used by the API must have only the minimum permissions necessary to read secrets.
• The configuration must minimize the blast radius if an identity or credential is compromised. You need to implement a secure access strategy for the API.
Which two actions should you perform? Each correct answer presents part of the solution. Choose two. NOTE: Each correct selection is worth one point.

A. Store a secret value in Azure App Configuration.
B. Assign the Key Vault Administrator role at subscription scope.
C. Grant the Key Vault Secrets User role at vault scope
D. Use system assigned managed identity.

**Answer:** CD

**Explanation:**
The API should use a system-assigned managed identity so no application credential is stored or rotated by the development team. That identity should receive only the permission needed to read secret values: Key Vault Secrets User at the vault scope is materially narrower than Key Vault Administrator at subscription scope. Storing a secret value in App Configuration would violate the requirement to protect secrets and increases exposure. The corrected pair therefore implements both credential-free authentication and least-privilege authorization with a limited blast radius.
Study Guide Alignment: Security and operations: Key Vault, App Configuration, managed identity, OpenTelemetry, Azure Monitor, and KQL-based troubleshooting.
Official Microsoft Learn References: AI-200 Study Guide | Azure Key Vault RBAC guide | Managed identities for Azure resources

**NEW QUESTION 19**
You deploy a private container image from Azure Container Registry (ACR) to App Service.
App Service must authenticate to ACR to pull the image.

The solution must NOT store static registry credentials.
You need to configure the secure image pull authentication.
Which configurations should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

**Configure secure container image pull**

| Requirement | Configuration |
|---|---|
| Authenticate App Service to ACR | Assign the Container Registry Repository Reader role to a managed identity.<br>Configure a webhook on ACR<br>Enable the admin user. |
| Avoid storing static credentials | Embed credentials in Dockerfile.<br>Store the registry password in application settings.<br>Use managed identity with role assignment. |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

App Service can pull a private image from Azure Container Registry by using a managed identity rather than registry username/password credentials. The identity must be granted an
appropriate pull-only registry role at the registry scope, preserving least privilege. A webhook does not
authenticate an image pull, and enabling the ACR admin account or storing a registry password would
introduce static credentials. The managed-identity approach satisfies both secure authentication and credentialelimination requirements.
Study Guide Alignment: Containerized Azure workloads: registry builds, App Service containers, Container
Apps revision/scaling behavior, and AKS deployment choices.
Official Microsoft Learn References: AI-200 Study Guide | Configure a custom container for App Service |
Managed identities for Azure resources

**NEW QUESTION 22**
You are reviewing an Azure Function app that processes incoming order requests for a company. The
function must:
• Accept order submissions from an external client application.
• Require controlled access for security
• Return a response containing the processed request payload.

You review the following code segment:

```
01  import azure.functions as func
02  app = func.FunctionApp(http_auth_level=func.AuthLevel.FUNCTION)
03    @app.route(route="processorder", methods=["POST"])
04    def processorder(req: func.HttpRequest) -> func.HttpResponse:
05        result = req.get_body().decode("utf-8")
06        return func.HttpResponse(result)
```

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

**Code-level analysis of HTTP trigger configuration**

| Statement | Yes | No |
|---|---|---|
| The function requires a function key for invocation. | ○ | ○ |
| The function supports HTTP GET requests. | ○ | ○ |
| The response returns the request body. | ○ | ○ |
| The function validates the request body before returning a response. | ○ | ○ |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

auth_level=FUNCTION` means invocation requires a function key unless a stronger platform authentication layer is configured, so the first statement is true. The route explicitly permits POST,
not GET. The function reads `req.get_body()` and returns that value in the HttpResponse, so the response

contains the request body. No schema, type, required-field, or other payload validation is performed before the response is constructed, making the final statement false.
Study Guide Alignment: Azure service integration: Service Bus, Event Grid, Azure Functions triggers /bindings, and event-driven processing.
Official Microsoft Learn References: AI-200 Study Guide | Azure Functions HTTP trigger

**NEW QUESTION 24**
Drag and Drop
You are developing a new page for a website that uses Azure Cosmos DB for data storage. The feature uses documents that have the following format:

```
{
  "name": "John",
  "city": "Seattle"
}
```

You must display data for the new page in a specific order. You create the following query for the page:

```
SELECT *
FROM People p
ORDER BY p.name, p.city DESC
```

You need to configure an Azure Cosmos DB policy to support the query.
How should you configure the policy? To answer, drag the appropriate JSON segments to the correct locations. Each JSON 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.

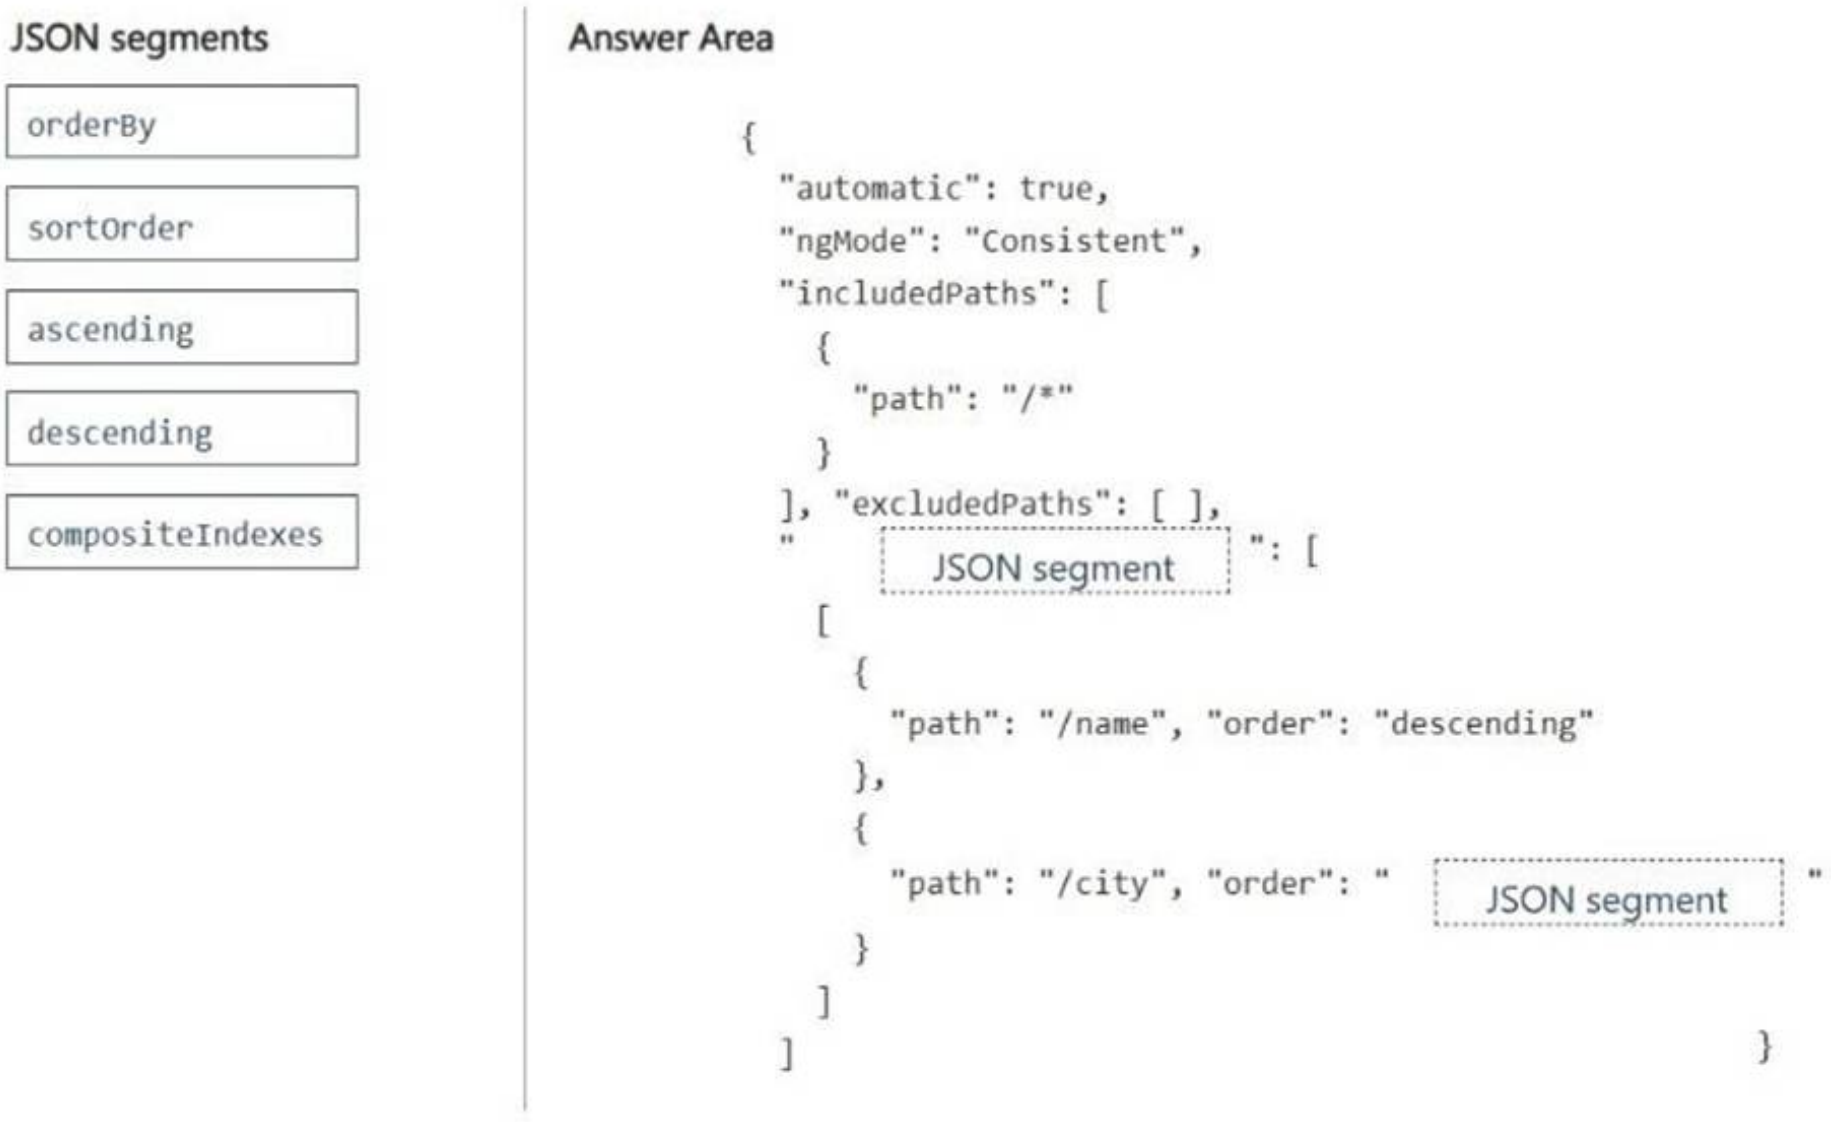


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: compositeIndexes
Azure Cosmos DB requires a defined composite index for any query utilizing an ORDER BY clause with multiple properties.
Box 2: descending
In an ORDER BY clause, properties without an explicit direction default to ascending order. Your query ORDER BY p.name, p.city DESC translates to an ascending sort on /name and a descending sort on /city.
Reference:
https://docs.azure.cn/en-us/cosmos-db/index-policy

**NEW QUESTION 26**
You are developing an AI API deployed to ACA. The API requires database credentials that are stored in Key Vault. Key Vault is configured to use Azure RBAC for access control.
The database credentials are rotated periodically by the security team. The application must always use the latest version of each credential without being redeployed and without exposing secrets in code or configurations.
You need to implement a secure secret access strategy that prevents credential exposure and fetches the latest version of each secret at runtime without redeploying the container.
Which three actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Configure a Key Vault RBAC role assignmen
B. Configure a Key Vault access polic
C. Export the secret during deploymen
D. Assign a system-assigned managed identit
E. Retrieve the secret at runtime by using the SDK.

**Answer:** ADE

**Explanation:**
Step-by-Step Implementation Guide
* 1. Assign a System-Assigned Managed Identity
* 2. Configure Key Vault RBAC Role Assignment
Role Selection: Assign the Key Vault Secrets User role to the container app's managed identity.
Scope Limitation: Limit the scope of this assignment to the specific Key Vault or individual secrets rather than the entire resource group.
* 3. Retrieve the Secret at Runtime Using the SDK
Reference:
https://oneuptime.com/blog/post/2026-02-16-how-to-use-managed-identity-with-azure-container-apps-to-access-azure-services/view

**NEW QUESTION 29**
You plan to deploy a container to an Azure App Service API app named api1. You host the source code for api1 in a GitHub repository. The container uses the API key at runtime to connect to a backend service.
The container must be able to retrieve the API key at runtime without exposing it in the source repository or Git commit history.
You need to ensure that the API key remains outside of Git commit history and is available to the container at runtime.
Solution: Embed the API key as a hardcoded environment variable in the Dockerfile.
Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**Explanation:**
Correct:
* Store the API key in Azure Key Vault and reference it from an App Service application setting.
Storing the API key in Azure Key Vault and referencing it via App Service application settings is the recommended, secure approach. This strategy completely removes sensitive credentials from your GitHub repository and Git commit history while injecting them safely into your container environment at runtime.
Incorrect:
* Embed the API key as a hardcoded environment variable in the Dockerfile.
* Store the API key as a GitHub repository secret.
Reference:
https://www.qservicesit.com/full-stack-applications-on-azure

**NEW QUESTION 32**
You are developing an AI search API that caches semantic search results in Redis.
Search results must remain cached for 10 minutes. If the underlying data changes, cached entries must NOT be returned.
You need to implement a cache-aside strategy to ensure data consistency.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Configure a cache notification for key space event
B. Delete related cache keys when the source data change
C. Implement sliding expiration based on key acces
D. Configure a 10-minute Time to Live on each key.

**Answer:** BD

**Explanation:**
Delete keys on mutation. Implement an event-driven or database-hook system to purge related Redis keys the moment the source data updates.
Use absolute expiration. Set a strict, un-extendable Time-To-Live (TTL) of 10 minutes when writing to the cache, ensuring the data naturally expires even if a deletion event is missed.
Incorrect:
[Not C]
Implementing a sliding expiration is not a good step for this specific architecture because it violates the requirement that entries must not be returned if the underlying data changes.
Reference:
https://medium.com/real-world-net/net-caching-strategies-compared-redis-vs-memory-vs-hybrid-b8b8be1e708d

**NEW QUESTION 33**
.....

# Relate Links

**100% Pass Your AI-200 Exam with Exambible Prep Materials**

https://www.exambible.com/AI-200-exam/

# Contact us

**We are proud of our high-quality customer service, which serves you around the clock 24/7.**

**Viste -** https://www.exambible.com/