Microsoft

Exam Questions AI-200

Developing AI Cloud Solutions on Azure

NEW QUESTION 1
You are configuring sampling for a distributed application that sends traces to Azure Monitor. The solution must:
• Preserve upstream sampling decisions across distributed traces
• Capture all spans during local testing.
• Sample 10 percent of traces in production.
You need to apply the appropriate sampling configuration for each requirement.
What should you do? 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.

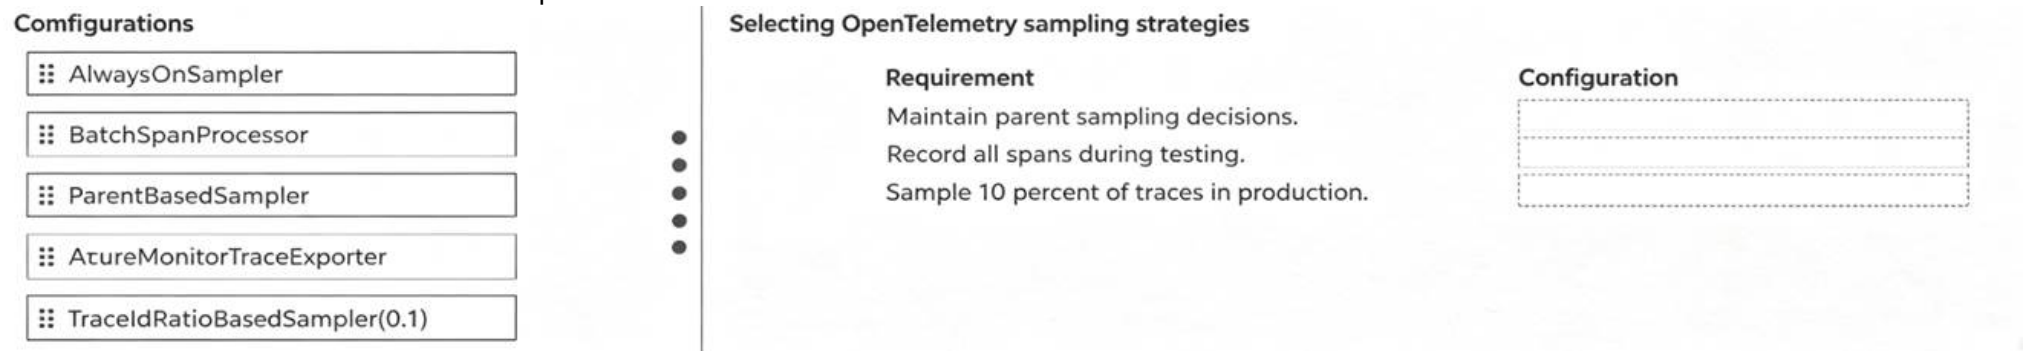

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Verified Answer:Preserve parent sampling decisions: ParentBasedSampler. Capture all spans locally: AlwaysOnSampler. Sample 10% in production: TraceIdRatioBasedSampler(0.1).
Detailed Explanation:Parent-based sampling propagates the sampling decision from the upstream parent, preventing broken sampling decisions across a distributed trace. AlwaysOnSampler records every span, which is suitable for local test environments where complete trace visibility is more important than telemetry volume. TraceIdRatioBasedSampler with 0.1 selects approximately ten percent of traces based on trace identifiers, providing deterministic fixed-ratio sampling for production. A batch span processor and an exporter control processing/export, not the sampling decision itself.
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 | OpenTelemetry sampling in Azure Monitor | Enable Azure Monitor OpenTelemetry

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.

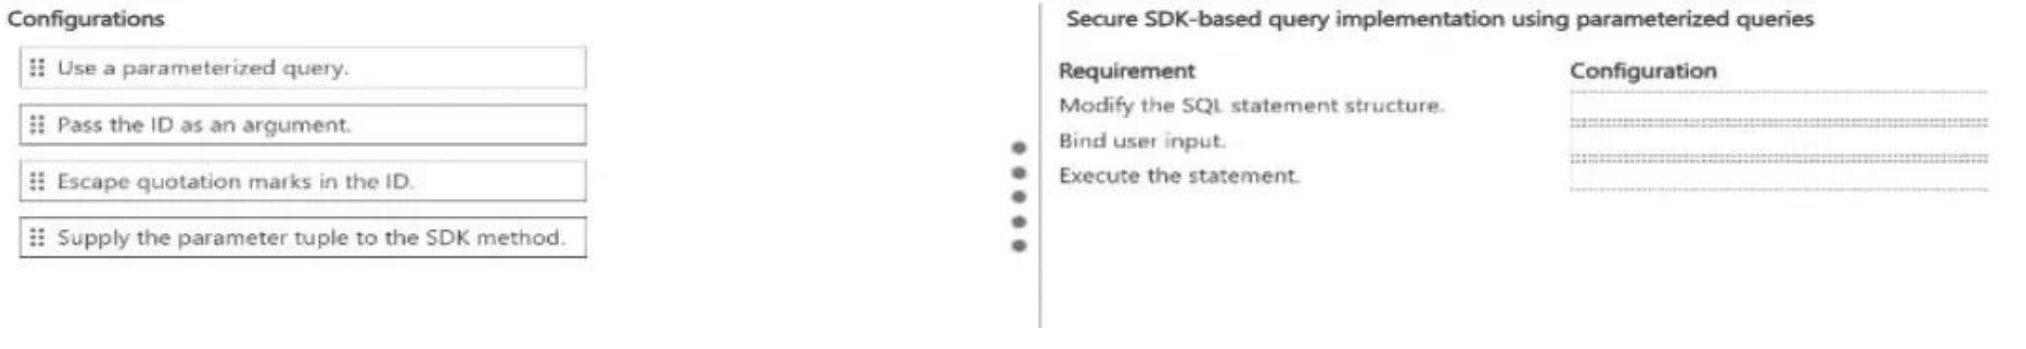

A. Mastered
B. Not Mastered

Answer: A

Explanation:

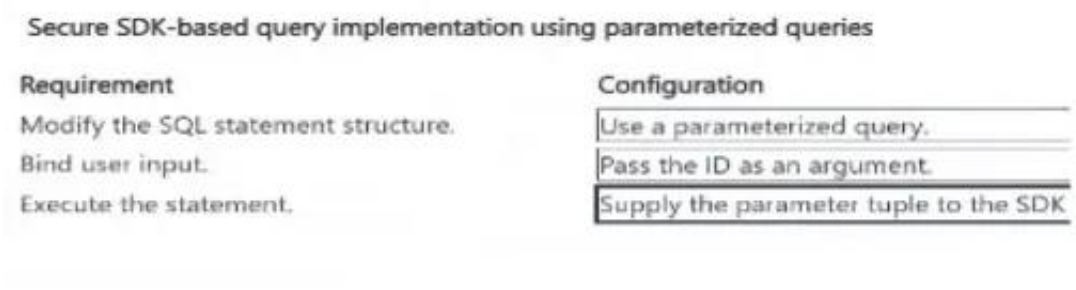

Verified Answer:Modify SQL: use a parameterized query. Bind input: pass the ID as an argument. Execute: supply the parameter tuple to the SDK/driver method.
Detailed Explanation:The security boundary is created by separating SQL syntax from user-supplied values. The statement should contain a parameter 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 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 4
You optimize an AI inference API that uses Redis caching.
You must reduce the risk of serving outdated data while minimizing cache management overhead. You need to implement the caching strategy that satisfies the requirements. What should you do?

A. Reset expiration on each read.
B. Disable expiration for cache keys.
C. Set eviction policy to allkeys-lru.
D. Trigger invalidation when source data changes

Answer: D

Explanation:
If the priority is to minimize stale results, source-driven invalidation is the direct control: when the underlying record changes, delete the corresponding cache entry so the next request repopulates it from current data. Sliding expiration can keep a frequently accessed stale item alive, disabling expiration is worse, and an allkeys-LRU policy evicts according to memory pressure rather than data freshness. Invalidation therefore best addresses correctness without requiring constant cache-wide maintenance.
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 Managed Redis documentation

NEW QUESTION 5
A Python API running in ACA must send distributed traces to Azure Monitor. The API creates spans. However, no traces appear in Azure Monitor.
You need to configure the OpenTelemetry SDK pipeline to export traces to Azure Monitor.
What should you do? To answer, move the appropriate actions to the correct requirements. You may use each action 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.

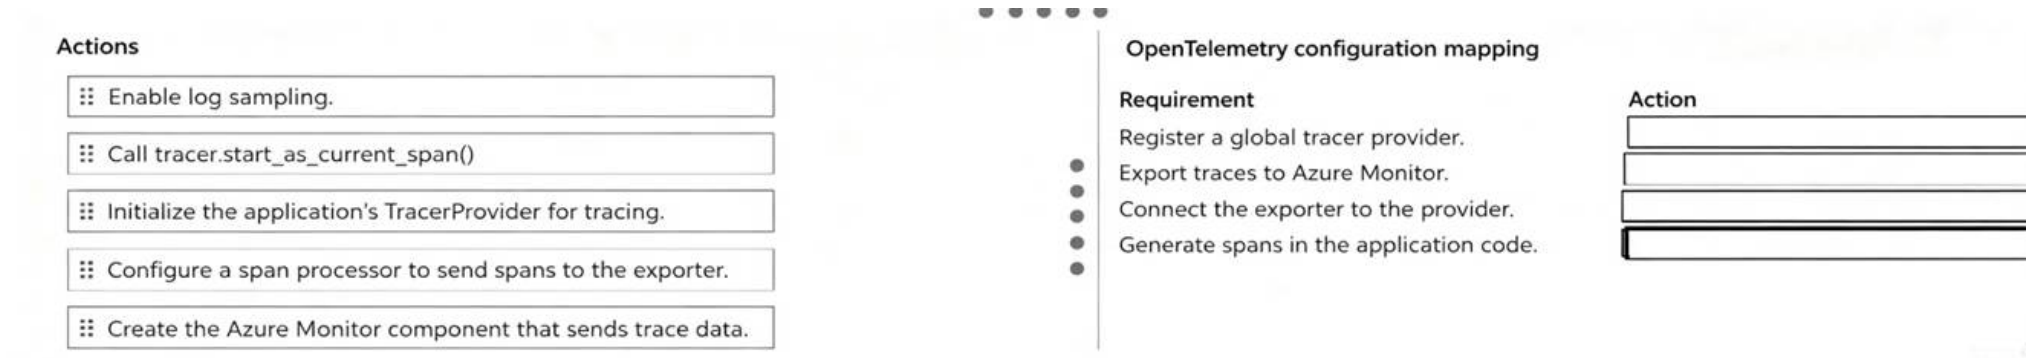

A. Mastered
B. Not Mastered

Answer: A

Explanation:
OpenTelemetry tracing is a pipeline: the application obtains a tracer from a configured provider, creates spans, passes ended spans through a span processor, and exports them using the Azure Monitor exporter. Creating spans alone is insufficient; without the provider/processor/exporter chain, no telemetry reaches Azure Monitor. Log sampling and legacy Application Insights TrackEvent calls are unrelated to the required OpenTelemetry trace-export path.
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 | Enable Azure Monitor OpenTelemetry

NEW QUESTION 6
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 7
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 8
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 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 developing several microservices to run on Azure Container Apps. External HTTP ingress traffic has been enabled for the Microservices.
A deployed microservice must be updated to allow users to test new features. You have the following requirements:
* Enable and maintain a single URL for the updated microservice to provide to test users.
* Update the microservice that corresponds to the current microservice version. You need to configure Azure Container Apps.
Which features should you configure? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Anwer Area

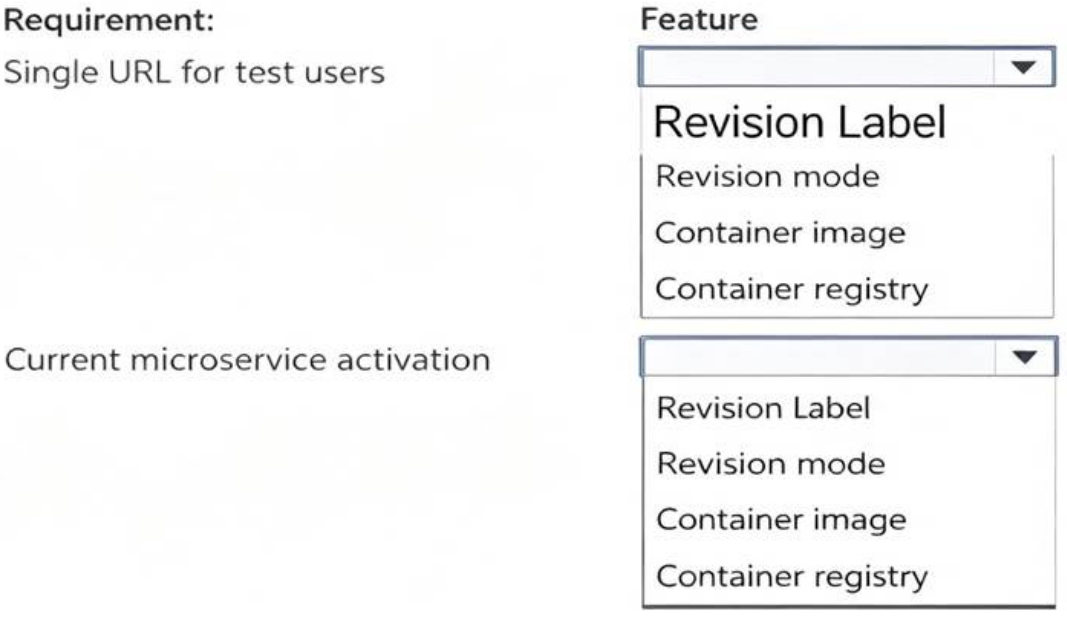

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Azure Container Apps revision labels provide a stable label-specific URL that can be reassigned to a different revision, which is appropriate for a consistent test endpoint. Revision mode controls
whether one or multiple revisions can be active at the same time and therefore governs how the current and updated application versions are activated. The container image and registry identify what is deployed, but they do not provide the stable test URL or the application-level revision activation behavior requested.
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 | Azure Container Apps revision labels

NEW QUESTION 10
You deploy a production Azure Function app that connects to an Azure SQL Database. The solution must provide the following functionality:
• Prevent secrets from being exposed in source control.
• Support secret rotation without redeploying the function app.
• Avoid downtime during credential updates.
You need to configure secure and maintainable secret management. What should you configure?

A. Application settings with Key Vault references
B. Environment variables in local.settings.json
C. Parameter file stored in source control
D. Hard coded connection string in the startup class

Answer: A

Explanation:
Key Vault references in Function/App Service application settings keep the connection secret outside source control while allowing the platform to resolve it at runtime by identity. When a referenced secret version is not fixed, rotation can be picked up without rebuilding or redeploying the function code.
`local.settings.json` is intended for local development and must not be used as a production secret store. A source-controlled parameter file or hard-coded connection string directly violates the security requirement.
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 | Use Key Vault references for App Service and Functions

NEW QUESTION 11
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

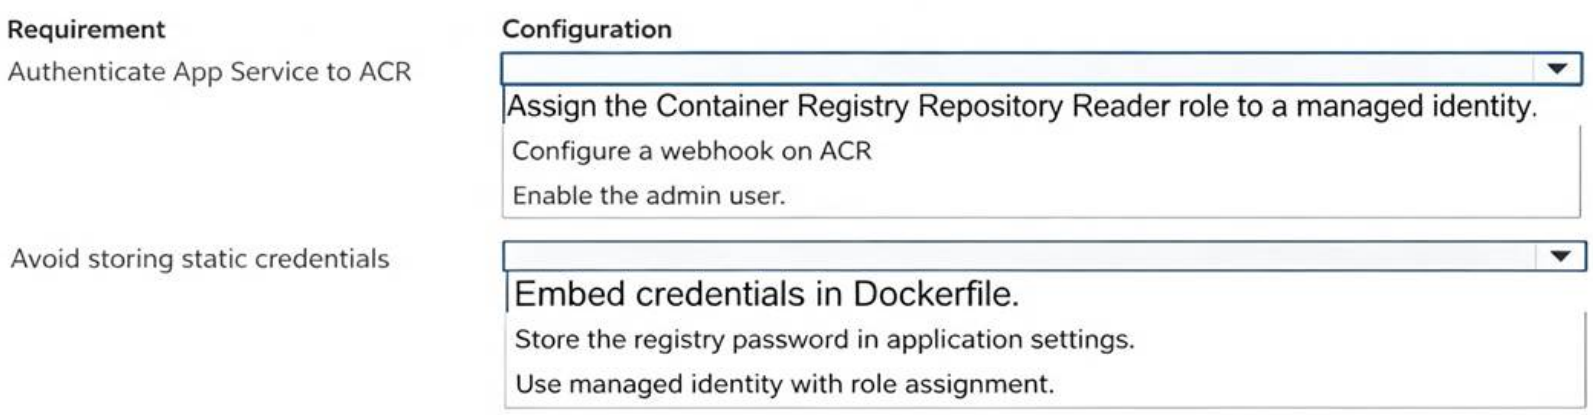

A. Mastered
B. Not Mastered

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 13
You develop a message-processing service deployed to Azure Container Apps. The service reads messages from an Azure Service Bus queue. The solution must minimize costs by ensuring NO compute resources are consumed when the queue is empty. You need to configure scaling for the service. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. A.Increase the scaling rule to allow for the maximum running replica count.B.Configure the scaling rule to allow for the termination of all active replicas.C.Configure a Kubernetes Event-driven Autoscaler rule that monitors queue length.D.Enable HTTP ingress concurrency scaling.

Answer: BC

NEW QUESTION 18
You configure ACR Tasks to automate image builds. Container images must rebuild when: Application updates occur. Base image updates occur, such as when the underlying OS image is updated. Regular scheduled rebuilds are required. You need to configure ACR Tasks to support automated image rebuilds. Which three triggers should you configure? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. A.Timer triggerB.Source code commit triggerC.Registry event triggerD.Base image update triggerE.Webhook notification trigger

Answer: ABD

Explanation:

To support automated container image rebuilds based on your requirements, you need to configure source triggers, base image triggers, and timer triggers in Azure Container Registry (ACR) Tasks. Timer Triggers: Automates regular scheduled rebuilds. This uses a cron schedule expression to run image builds at specific intervals (e.g., weekly or monthly) regardless of code changes. Source Triggers: Automates rebuilds when application updates occur. This monitors changes in your source code repository (like GitHub or Azure Repos) and fires a build when code is committed. Base Image Triggers: Automates rebuilds when the underlying OS or framework image updates. ACR tracks dependencies and automatically kicks off a new build when your defined FROM image changes in the public registry or your private registry. Reference: https://oneuptime.com/blog/post/2026-02-16-how-to-set-up-acr-tasks-for-automated-containerimage-builds-on-git-commit/view

NEW QUESTION 22
DRAG DROP - 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.

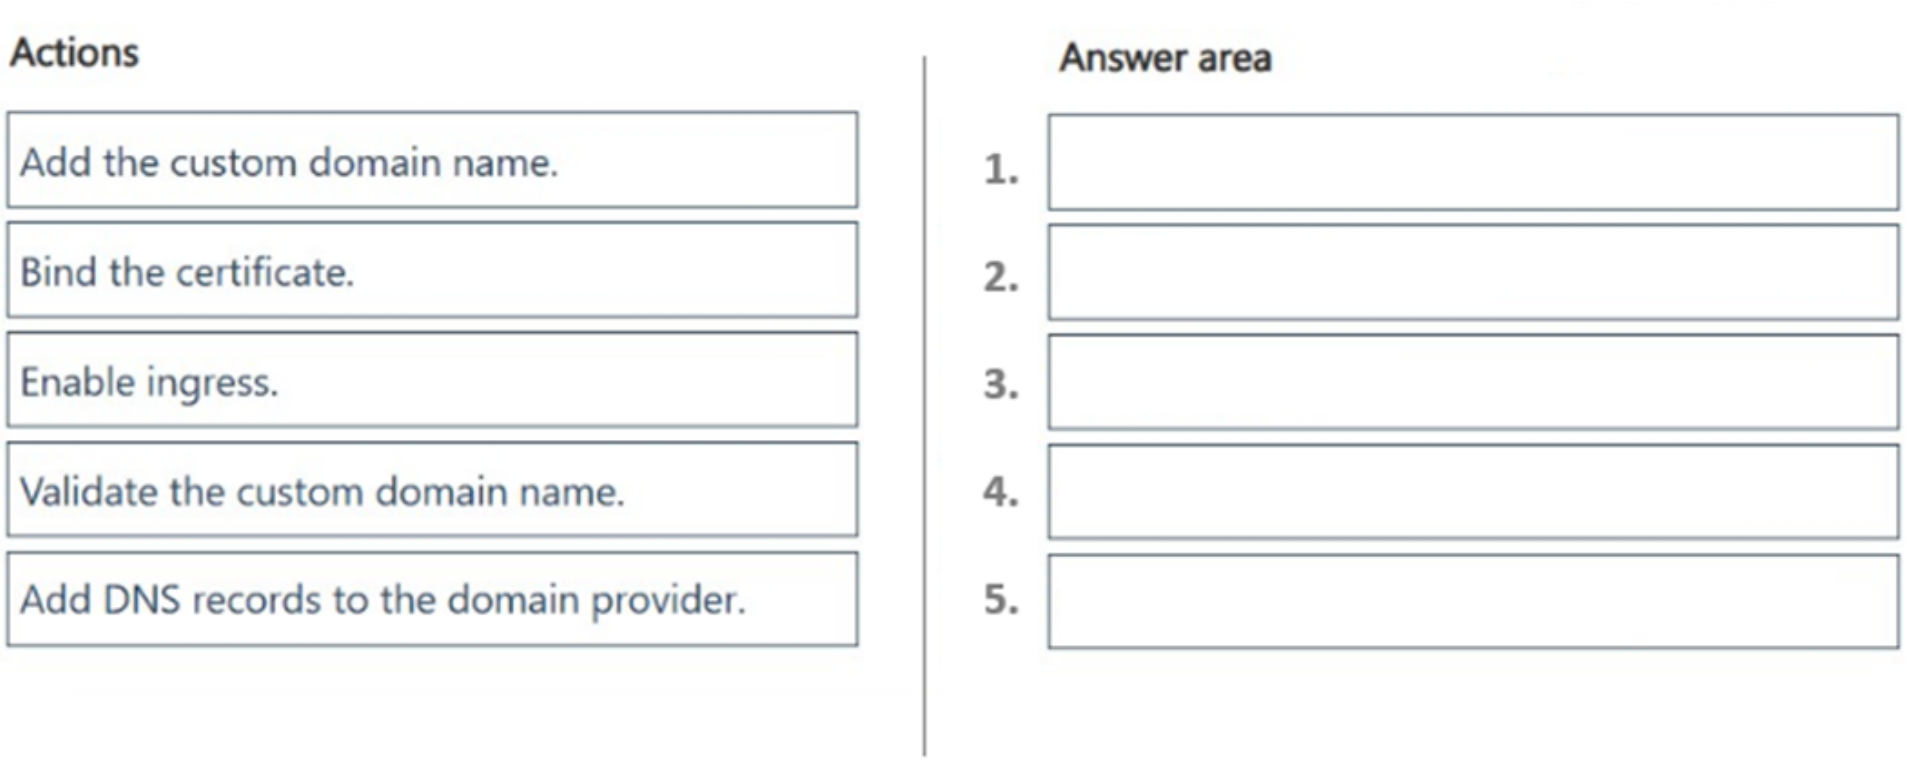

A. Mastered
B. Not Mastered

Answer: A

Explanation:

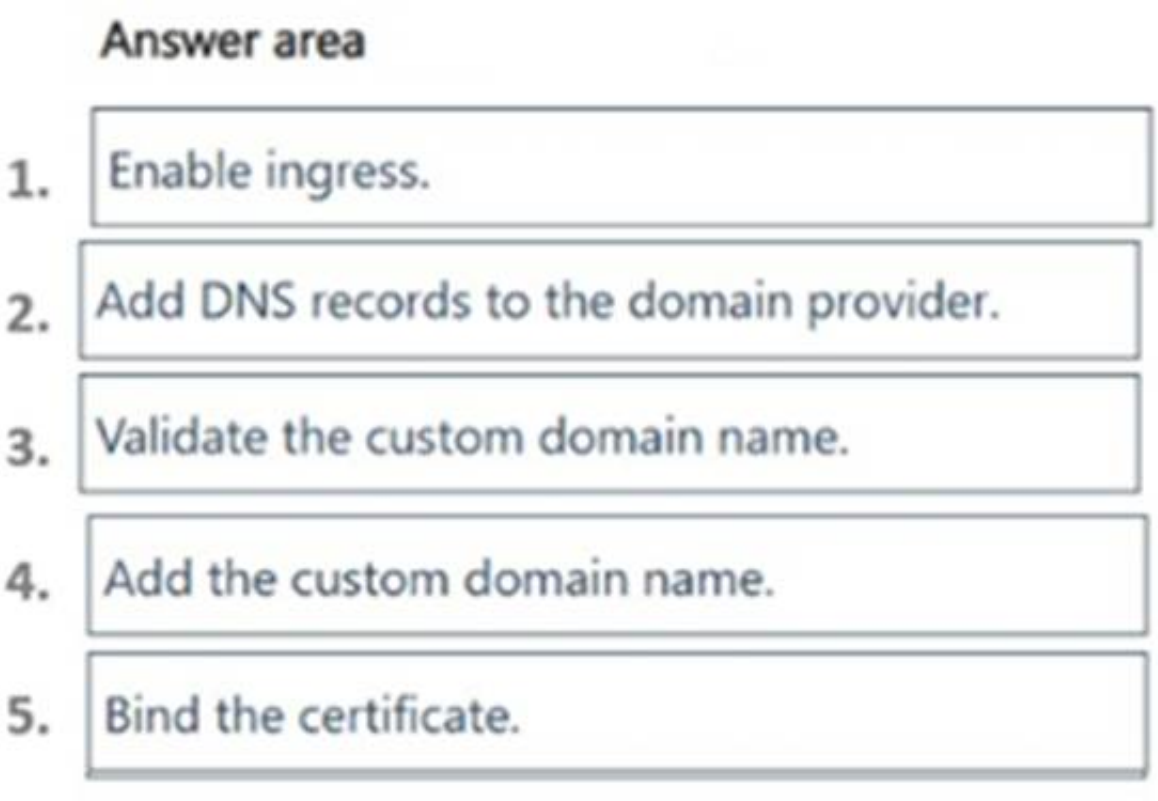

NEW QUESTION 26
You plan to deploy a container to an Azure App Service API app named api1. You host the source code for api1 in a GitHub repository. The container uses the API key at runtime to connect to a backend service.
The container must be able to retrieve the API key at runtime without exposing 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: Store the API key as a GitHub repository secret.
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 27
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 37
.....

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

AI-200 Practice Exam Features:

* AI-200 Questions and Answers Updated Frequently

* AI-200 Practice Questions Verified by Expert Senior Certified Staff

* AI-200 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AI-200 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year