Microsoft

Exam Questions AI-200

Developing AI Cloud Solutions on Azure

NEW QUESTION 1
You need to configure a connection string for the partner-facing service according to the technical requirements.
What should you use?

A. Azure Key Vault references in App Service settings
B. GitHub secrets
C. Azure Container Registry Helm chart package
D. Dockerfile ENV instructions

Answer: A

Explanation:
Azure Key Vault references in App Service settings satisfy the requirement to keep secrets out of container images, source control, and directly stored application configuration. App Service resolves the referenced secret at runtime by using the app identity, so the application can consume the value as a normal setting without embedding credentials in the image. GitHub secrets are build/deployment secrets rather than a runtime App Service secret-delivery mechanism. Dockerfile ENV instructions would place secret material in the image configuration and violate the case requirements.
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 | Managed identities for Azure resources

NEW QUESTION 2
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 deploy a new revision of the processing orders app.
Processing orders must be triggered by a web request and must always be available based on incoming web requests.
You need to validate that the replica is ready to handle incoming requests.
What should you implement?

A. HTTP liveness probe
B. HTTP startup probe
C. TCP readiness probe
D. HTTP readiness probe
E. TCP liveness probe

Answer: D

Explanation:
Detailed Explanation:A readiness probe answers whether a running replica is currently ready to receive traffic. For a web-triggered order-processing service, an HTTP readiness probe can test the application endpoint and keep the replica out of the routing pool until it is able to handle requests. Liveness probes detect whether the process should be restarted, and startup probes protect slow-starting containers during initialization. A TCP readiness probe can test socket acceptance, but the HTTP readiness probe is the more application-aware choice for the HTTP service described.
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 | Health probes in Azure Container Apps

NEW QUESTION 3
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.

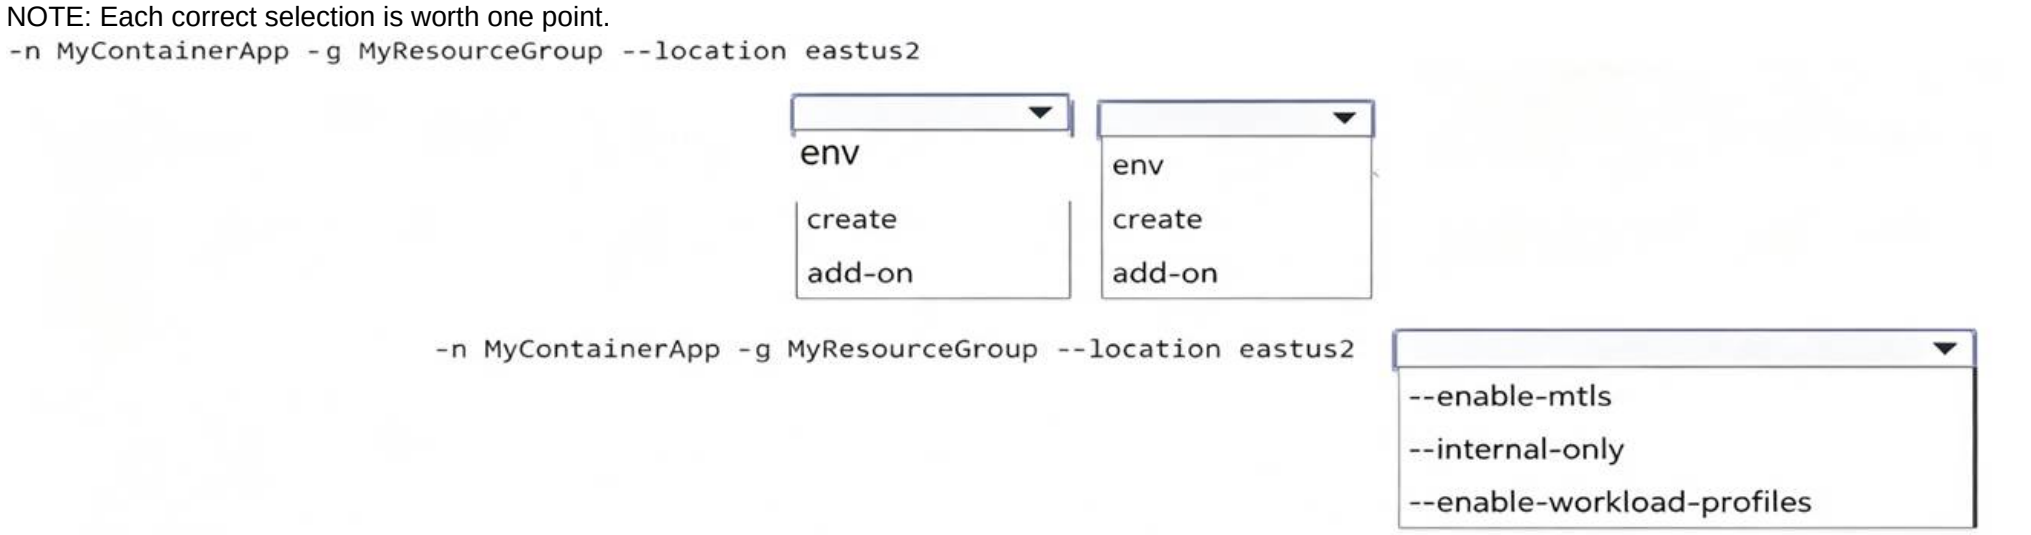

A. Mastered
B. Not Mastered

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 4
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.

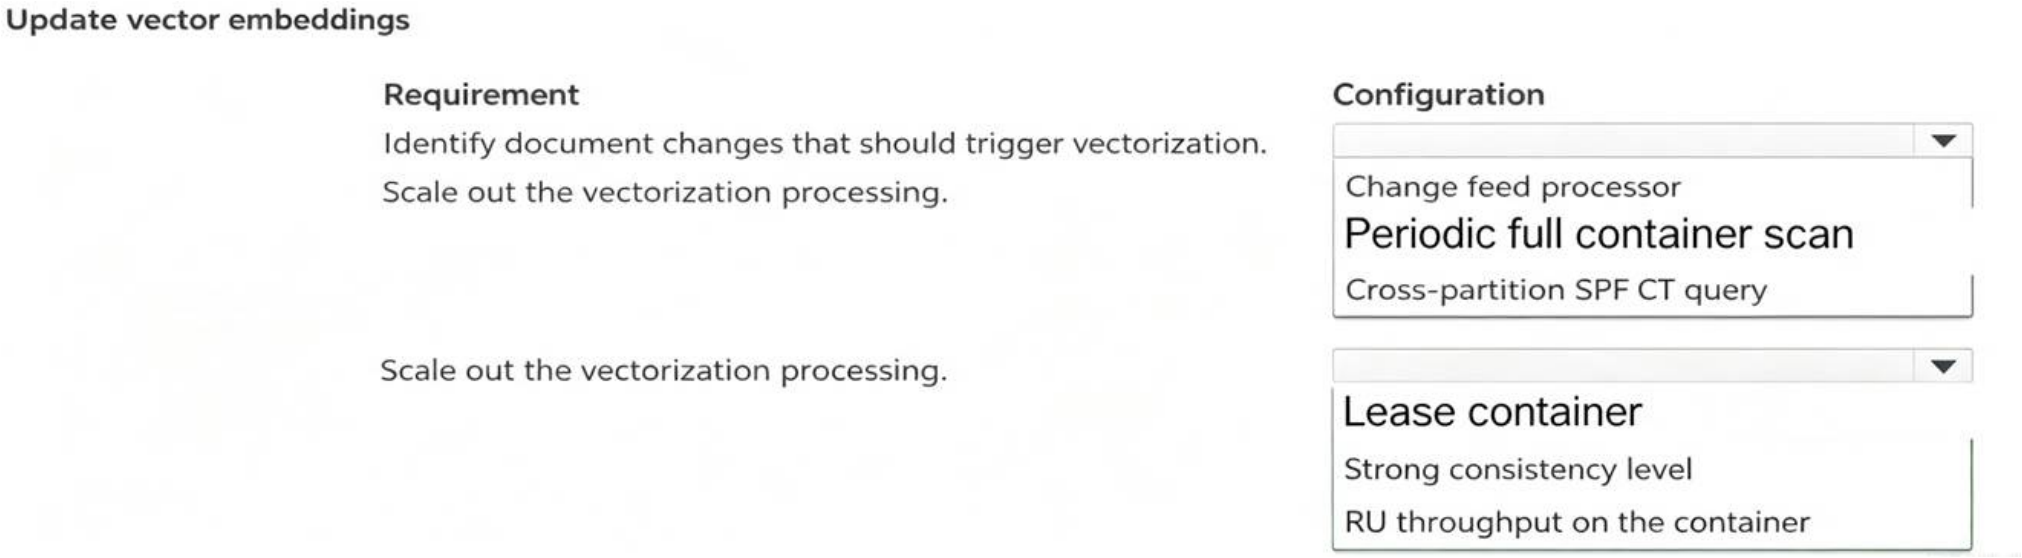

A. Mastered
B. Not Mastered

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 5
You need to implement trace correlation according to the business requirements.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.
NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

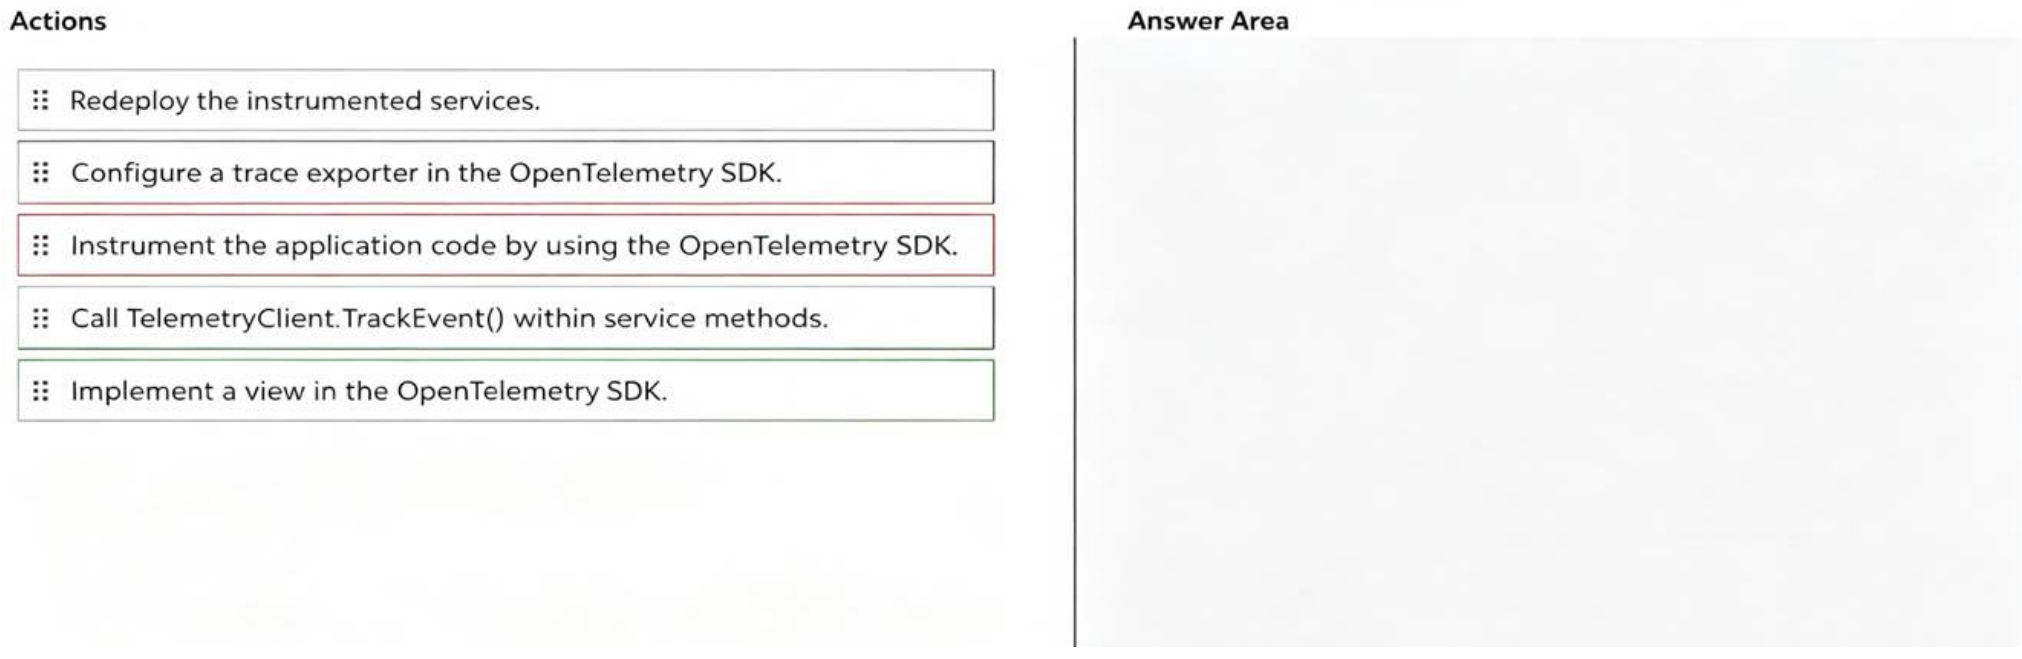

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Trace correlation requires the application to create OpenTelemetry spans and an exporter to send those spans to Azure Monitor. Instrumentation must therefore be present before the updated service is deployed. The previous Application Insights SDK instrumentation does not satisfy the case-study requirement that all tracing use OpenTelemetry. Configuring an OpenTelemetry view is unrelated to basic distributed-trace export, and calling TrackEvent is an Application Insights SDK pattern rather than the required OpenTelemetry approach.
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 developing an AI application that retrieves database credentials from Key Vault by using the Azure SDK for Python.
The application must use managed identity for authentication.
You review the following code segment that retrieves a secret from Key Vault:

```
01  from azure.identity import DefaultAzureCredential
02  from azure.keyvault.secrets import SecretClient
03
04  credential = DefaultAzureCredential()
05  client = SecretClient(
06      vault_url="https://vault.vault.azure.net/",
07      credential=credential
08  )
09
10  secret = client.get_secret("dbPassword", version="123")
```

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

Statement	Yes	No
The code retrieves a specific version of the dbPassword secret.	○	○
The code authenticates to Key Vault without storing client secrets in the application when deployed to an Azure hosted environment that has a managed identity enabled.	○	○
If dbPassword is rotated and a new version is created, subsequent calls to this code will retrieve the new version.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
The get_secret call supplies an explicit version value, so it retrieves that particular version of dbPassword. DefaultAzureCredential can authenticate to Key Vault by using the Azure-hosted application??s managed identity without embedding a client secret in the application. Because the code requests a fixed version, later secret rotation creates a newer version but subsequent executions of this exact code continue to request the specified old version. To follow the latest version automatically, the version parameter must be omitted.
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 | Managed identities for Azure resources | Use Key Vault references for App Service and Functions

NEW QUESTION 7
You are provisioning and configuring a Service Bus processor for AI batch jobs.
The processor must connect to an existing queue, register handlers for message and error processing, and then begin receiving messages.
You need to provision and configure the Service Bus processor for message and error handling.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Dead-letters failed messages.
Create the Service Bus client.
Create the Service Bus processor for the queue.
Register message and error handlers.
Start the message processor.

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:

The processor depends on a namespace client/connection, so the client is created first. Next, the queue-specific processor or receiver is created from that client. Message and error callbacks must be registered before processing starts so incoming messages and failures have defined handlers. Starting the processor is therefore the final step. Dead-lettering failed messages is an optional application settlement decision and is not a prerequisite for constructing and starting the processor.
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

NEW QUESTION 8
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 9
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 10
You need to redesign the configuration approach so that:
configuration supports dynamic refresh;
secrets are stored outside App Configuration;
the application can securely access both App Configuration and Key Vault.
You need to enable secure dynamic configuration management for the platform.
Which three actions should you perform? Each correct answer presents part of the solution. Choose three.
NOTE: Each correct selection is worth one point.

A. Store configuration values in environment variables.
B. Use a service principal secret for both App Configuration and Key Vault access.
C. Store API keys in Key Vault.
D. Allow configuration updates with a polling interval.
E. Use managed identity for both App Configuration and Key Vault access.

Answer: CDE

Explanation:
Explanation
The original answer omitted the secret-storage requirement. API keys and other secrets should be held in Key Vault, while App Configuration should retain non-secret configuration and Key Vault references. Dynamic refresh is implemented through the App Configuration refresh/polling mechanism rather than requiring an application restart. Managed identity provides credential-free access to both services in Azure. Environment variables do not provide centralized dynamic refresh, and a stored service-principal secret creates exactly the credential-management problem the design is intended to eliminate.
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 | Dynamic configuration in ASP.NET Core | Use Key Vault references in App Configuration | Managed

identities for Azure resources

NEW QUESTION 10
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 15
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.

Actions

- Add the custom domain name to the Azure Container app.
- Add DNS records to the domain provider.
- Bind the certificate.
- Validate the ownership of the custom domain name.
- Enable ingress.

Answer Area

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 16
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. Choose two
NOTE: Each correct selection is worth one point.

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

Answer: AD

Explanation:

The cache-aside design needs both bounded lifetime and explicit invalidation. A ten-minute TTL ensures cached search results do not persist indefinitely, while deleting related keys when source data changes prevents known-stale entries from being returned during that ten-minute window. Sliding expiration would extend lifetime based on access and can preserve stale data. Keyspace notifications report Redis key events; they do not replace application logic that invalidates entries when the authoritative source changes.
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 20
You are developing a solution that uses several Azure Service Bus queues. You create an Azure Event Grid subscription for the Azure Service Bus namespace.
You use Azure Functions as subscribers to process the messages.
You need to ensure that events can be sent to Azure Event Grid from the queues. The solution must use the principle of least privilege and minimize costs.
Which Azure Service Bus role and tier level should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

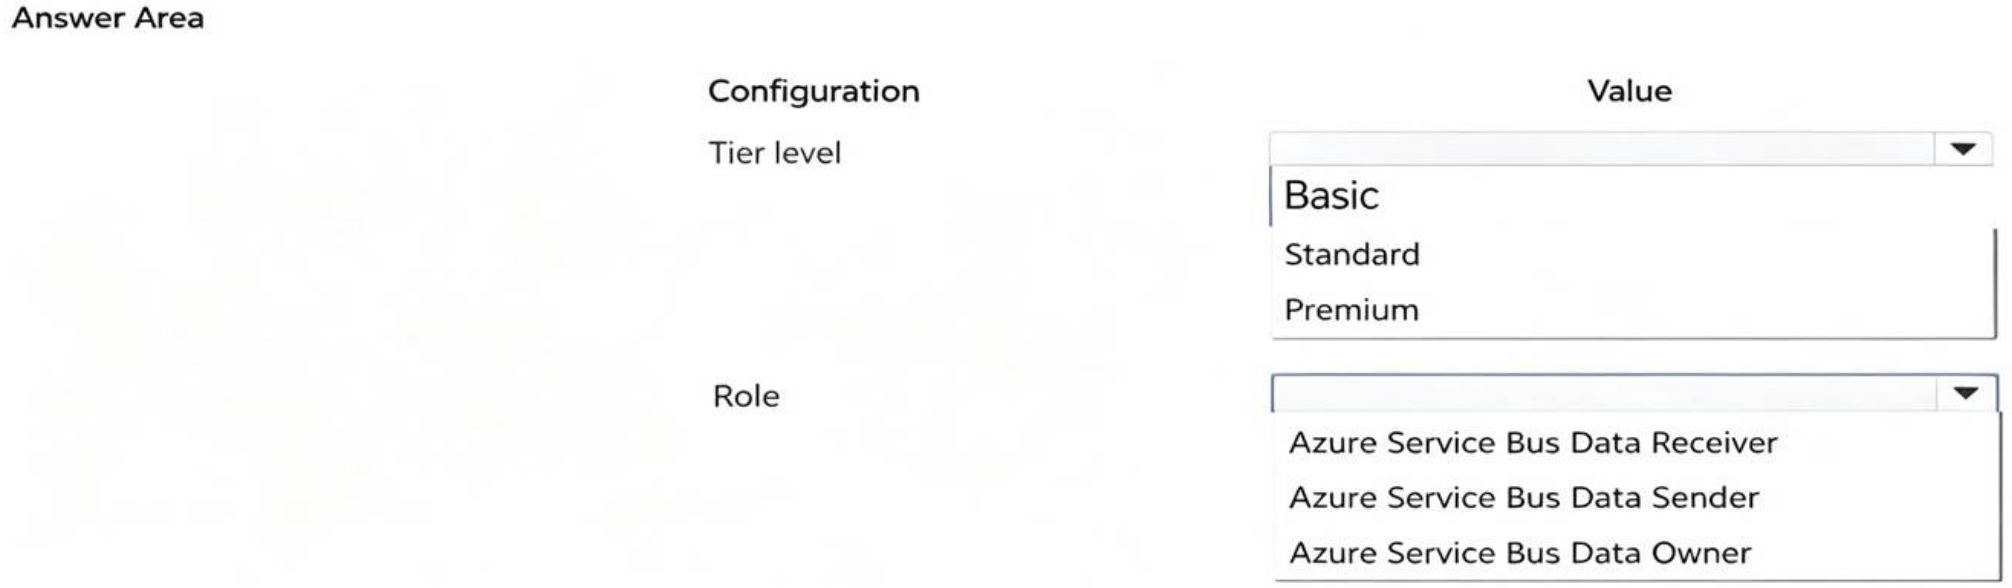

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Microsoft documents Event Grid integration for Service Bus queues and topics as requiring a Service Bus Premium namespace. The Azure Function subscriber must then receive messages from the queue; the least-privilege built-in data-plane role for that purpose is Azure Service Bus Data Receiver. Data Sender does not permit receiving, and Data Owner is broader than necessary. Premium therefore satisfies the feature prerequisite, while Data Receiver satisfies the stated least-privilege requirement.
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 to Event Grid integration | Authenticate Service Bus with Microsoft Entra ID

NEW QUESTION 24
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 28
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 32
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 33
You have an Azure web app that uses Azure Cosmos DB as a data store. You create a Cosmos DB container by running the following PowerShell script:

```
$resourceGroupName = "testResourceGroup"
$accountName = "testCosmosAccount"
$databaseName = "testDatabase"
$containerName = "testContainer"
$partitionKeyPath = "/EmployeeId"
$autoscaleMaxThroughput = 5000

New-AzCosmosDBSqlContainer `
   -ResourceGroupName $resourceGroupName `
   -AccountName $accountName `
   -DatabaseName $databaseName `
   -Name $containerName `
   -PartitionKeyKind Hash `
   -PartitionKeyPath $partitionKeyPath `
   -AutoscaleMaxThroughput $autoscaleMaxThroughput
```

You create the following queries that target the container:

```
SELECT * FROM c WHERE c.EmployeeId > '12345'
SELECT * FROM c WHERE c.UserId = '12345'
```

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 Ares

Statements	Yes	No
The minimum throughput for the container is 400 RU/s	○	○
The first query statement is an in-partition query.	○	○
The second query statement is a cross-partition query.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
The question maps directly to the AI-200 objective ??Develop AI solutions by using Azure Cosmos DB for NoSQL,?? which includes running queries and optimizing query performance and Request Unit (RU) consumption.
Statement 1: No — The minimum throughput is not 400 RU/s.
The PowerShell command provisions the container with:
-AutoscaleMaxThroughput 5000
Azure Cosmos DB autoscale operates between approximately 10% and 100% of the configured maximum throughput . Microsoft documentation specifically gives the example of an autoscale container provisioned with 5,000 RU/s scaling between 500 RU/s and 5,000 RU/s . Therefore, for this container, the minimum operating autoscale throughput is 500 RU/s , not 400 RU/s.

Therefore:
??The minimum throughput for the container is 400 RU/s.?? # No Statement 2: No — The first query is not an in-partition query. The container uses:
/EmployeeId
as its partition key. The first query is:
SELECT * FROM c WHERE c.EmployeeId > ' 12345 '
Although the query references the partition key, it uses a range predicate (>) , not an equality predicate. Microsoft explicitly states that a range filter on a partition key is not scoped to a single physical partition . To qualify as an in-partition query, the filter must identify the applicable partition, typically through an equality predicate such as:
WHERE c.EmployeeId = ' 12345 '
Microsoft ' s documentation provides essentially the same example: a query using DeviceId > ... against a container partitioned by DeviceId is not an in-partition query .
Therefore:
??The first query statement is an in-partition query.?? # No Statement 3: Yes — The second query is a cross-partition query. The second query is:
SELECT * FROM c WHERE c.UserId = ' 12345 '
The container ' s partition key is /EmployeeId, not /UserId. Because the query contains no filter on the partition key , Azure Cosmos DB cannot route it to one logical partition. It must fan out the query across the applicable physical partitions and combine the results.
Microsoft describes this behavior directly: when a query does not contain a filter on the partition key, it must execute across the partitions.
Therefore:
??The second query statement is a cross-partition query.?? # Yes

NEW QUESTION 35
A semantic search application queries Azure Database for PostgreSQL and stores document embeddings and metadata in a table with the following columns:
• embedding (pgvector)
• department
• created_at
The application must return the top five most similar documents for a given query embedding only from the finance department. You need to implement semantic retrieval with metadata filtering.
Which query components should you select? To answer, move the appropriate query components to the correct requirements. You may use each query component once, more than once, or not at all. You may need to move the split bar between panes or scroll to view content.

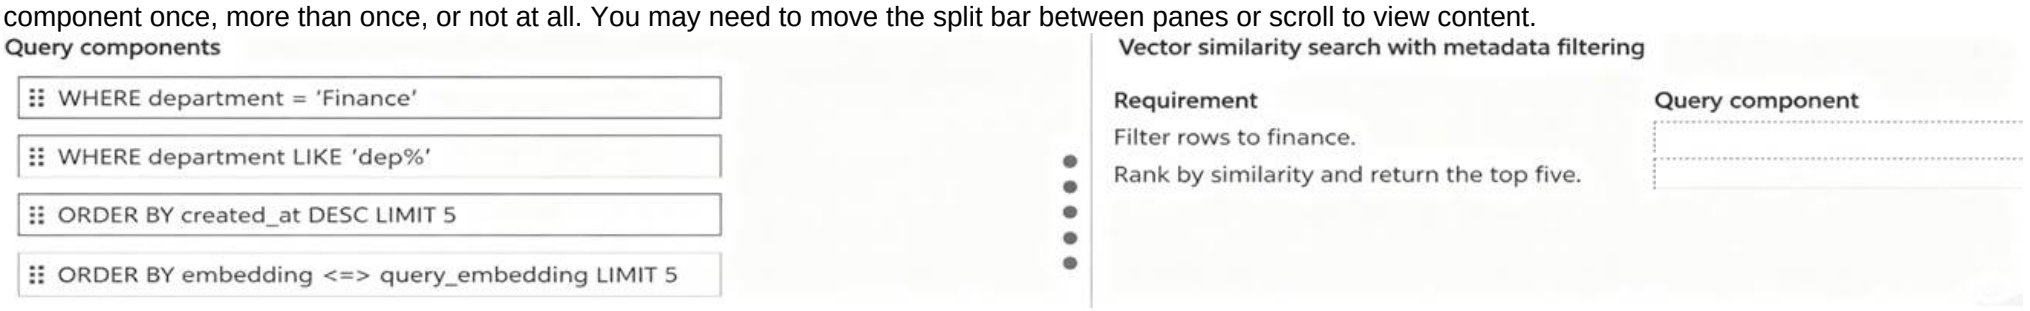

A. Mastered
B. Not Mastered

Answer: A

Explanation:
The metadata predicate must restrict the candidate rows to the Finance department, so the `WHERE department = ' finance ' ` component supplies the required filter. The pgvector cosine-distance operator ` < = > ` orders rows by vector distance to the supplied query embedding, and `LIMIT 5` keeps only the five nearest matches. Ordering by creation date would rank recency rather than semantic similarity, and a wildcard department filter would not satisfy the Finance-only requirement.
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 36
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.

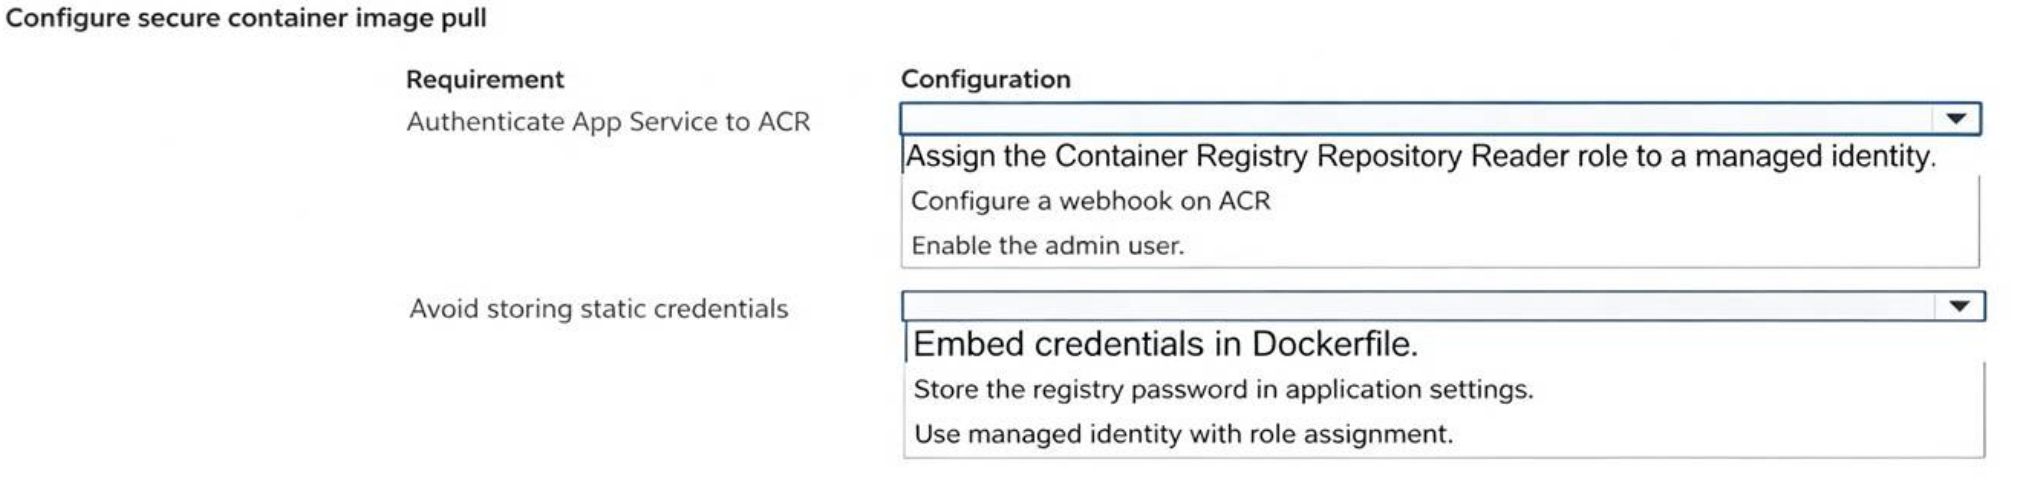

A. Mastered
B. Not Mastered

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 40
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 44
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.

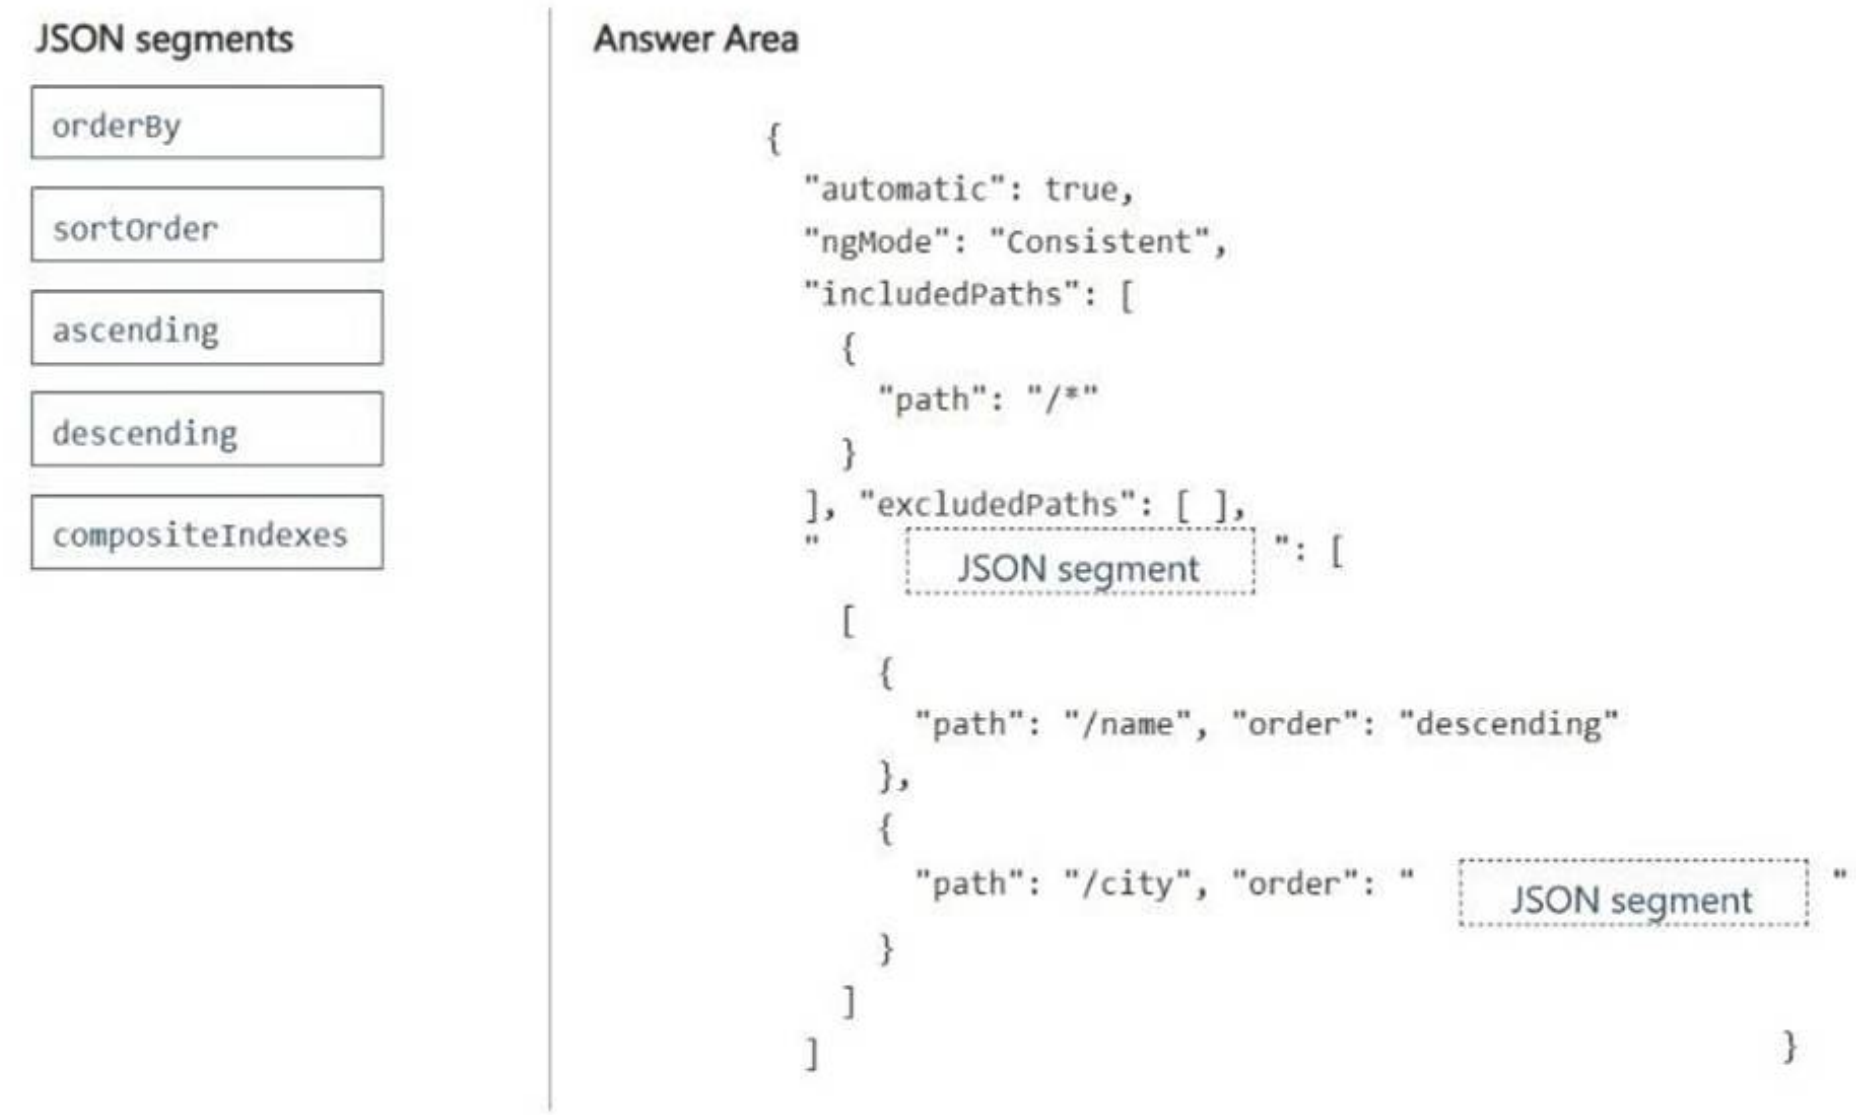

A. Mastered
B. Not Mastered

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 45
You have an Event Grid subscription that triggers an Azure Function.
You need to prevent loss of events in case the endpoint returns an HTTP 400 response.
Which action should you perform?

A. Implement optimistic batchin
B. Implement asynchronous handshake validatio
C. Configure a dead-letter destinatio
D. Configure a retry policy.

Answer: C

Explanation:
Configuring a dead-letter destination is exactly the right step to take.
By default, when an endpoint returns an HTTP 400 (Bad Request) or HTTP 413 (Payload Too Large) response, Azure Event Grid immediately stops delivery attempts and drops the message. It treats these specific errors as non-transient, client-side issues, meaning it skips its standard 24-hour retry policy.
Configuring a dead-letter destination ensures that these dropped events are safely preserved for future troubleshooting and reprocessing.
Reference:
https://turbo360.com/blog/azure-event-grid-dead-letter-monitoring

NEW QUESTION 49
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 53
.....

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