AI-200 Dumps

Developing AI Cloud Solutions on Azure

https://www.certleader.com/AI-200-dumps.html

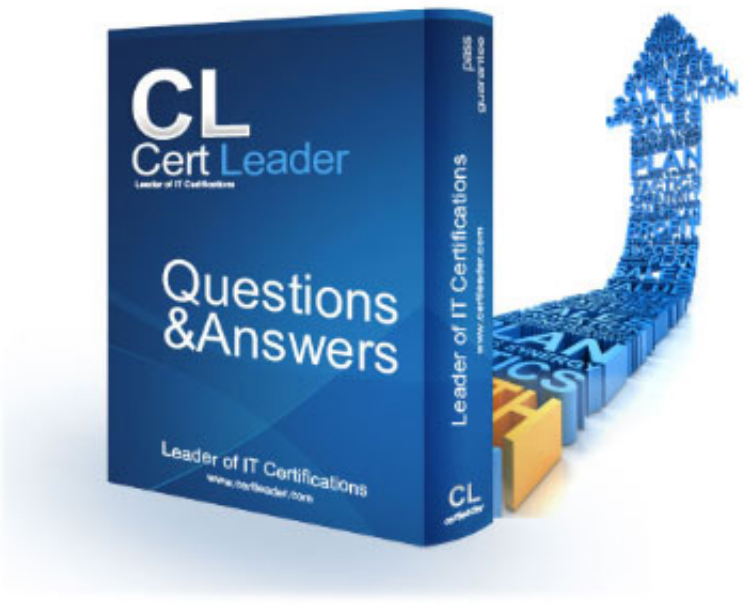

NEW QUESTION 1
You plan to deploy a web application to AKS.
The solution must:
• Scale out the application by adding more pods during peak CPU usage
• Expose the application internally within the cluster only
You need to configure a Kubernetes resource for each requirement
Which resources should you configure? To answer, move the appropriate resources to the correct requirements You may use each resource once, more than once, or not at all. You may need to move split bar between .. or scroll to view content.

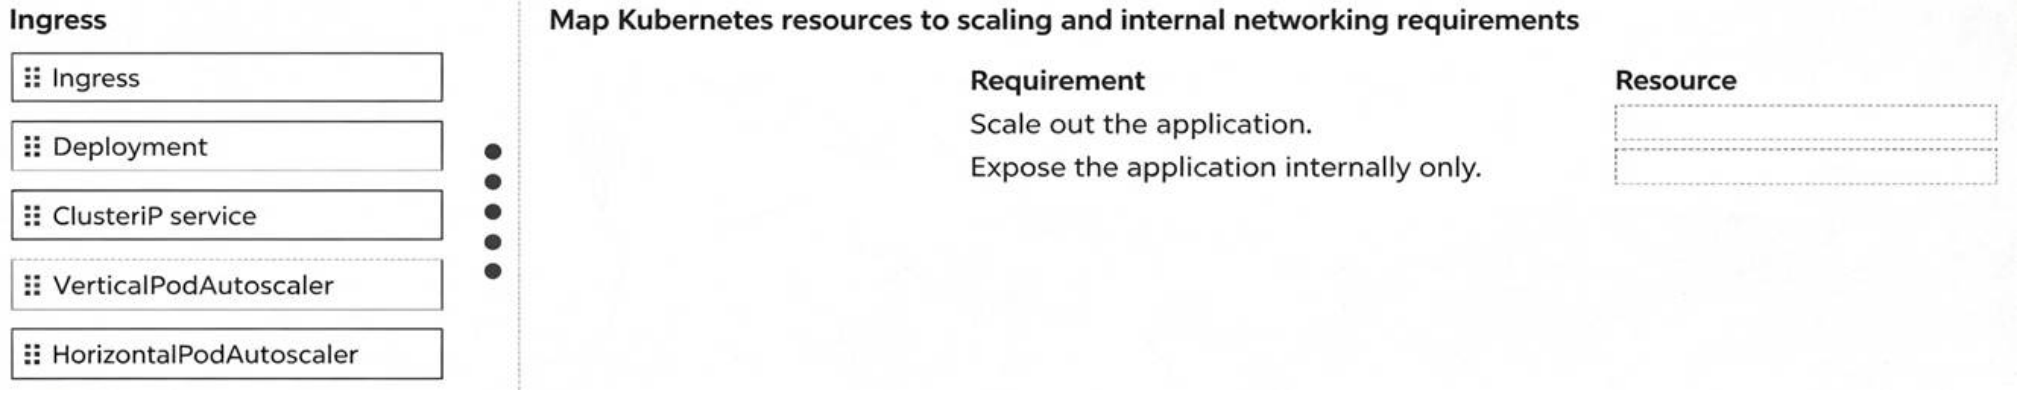

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Verified Answer:Scale out the application: HorizontalPodAutoscaler. Expose internally only: ClusterIP service.
Detailed Explanation:The Kubernetes HorizontalPodAutoscaler increases or decreases pod replicas in response to CPU, memory, or other metrics, so it is the resource that implements CPU-driven pod scale-out. A ClusterIP service creates an internal service IP that is reachable within the Kubernetes cluster and is the normal choice for an internal-only service. A Deployment defines the workload but does not itself implement metric-driven autoscaling, while Ingress is intended to route inbound network traffic.
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 | AKS scaling overview | Kubernetes services in AKS

NEW QUESTION 2
You need to deploy Azure function resources and apps by using an automated, version-controlled CI/CD pipeline that supports declarative infrastructure deployment. What should you use?

A. Local Git deployment
B. GitHub Actions
C. Azure Functions Core Tools
D. Azure CLI

Answer: B

Explanation:
Detailed Explanation:GitHub Actions is the correct orchestration mechanism among the options because the case requires an automated, version-controlled pipeline and Bicep-based declarative deployment. Microsoft documents GitHub Actions as a supported way to deploy Bicep and automate Azure resource deployment. Local Git, Azure Functions Core Tools, and Azure CLI can participate in development or scripted deployment, but they do not by themselves satisfy the stated requirement for a repository-driven, repeatable CI/CD pipeline.
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 | Deploy Bicep with GitHub Actions | Automate Function app resource deployment

NEW QUESTION 3
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 4
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 5
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:

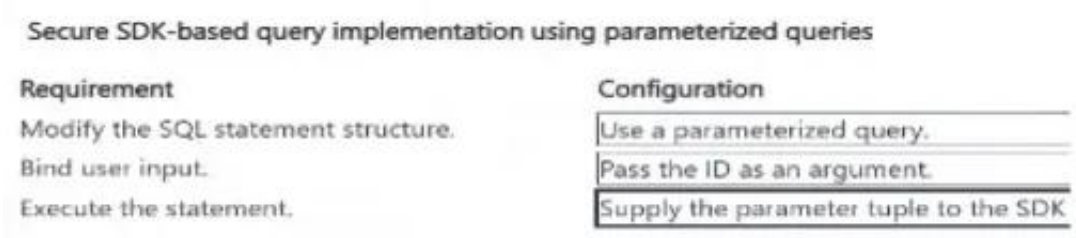

Verified Answer:Modify SQL: use a parameterized query. Bind input: pass the ID as an argument. Execute: supply the parameter tuple to the SDK/driver method.
Detailed Explanation:The security boundary is created by separating SQL syntax from user-supplied values. The statement should contain a parameter 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 6
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.

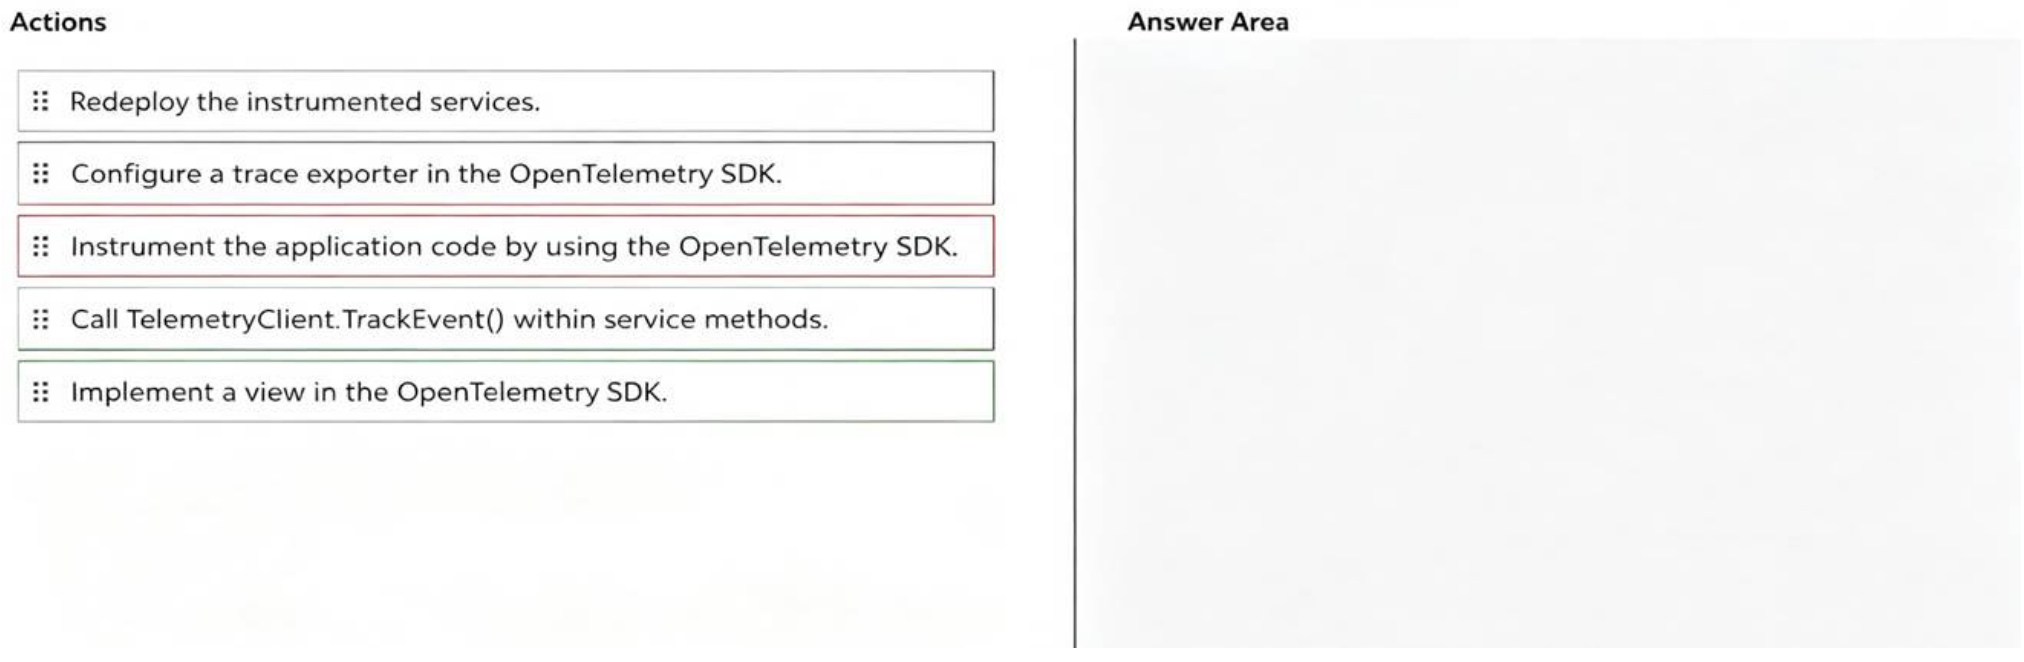

A. Mastered
B. Not Mastered

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 7
A company is implementing a publish-subscribe (Pub/Sub) messaging component by using Azure Service Bus. You are developing the first subscription application.
In the Azure portal you see that messages are being sent to the subscription for each topic. You create and initialize a subscription client object by supplying the correct details, but the subscription application is still not consuming the messages.
You need to ensure that the subscription client processes all messages. Which code segment should you use?

A. subscriptionClient = new SubscriptionClient(ServiceBusConnectionString, TopicName, SubscriptionName);
B. subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync, messageHandlerOptions);
C. await subscriptionClient.AddRuleAsync(new RuleDescription(RuleDescription.DefaultRuleName, new TrueFilter()));
D. await subscriptionClient.CloseAsync();

Answer: B

Explanation:
Creating a subscription client establishes the client object but does not start application-level message processing. The message handler must be registered so received subscription messages are delivered to the application callback. Adding a default TrueFilter is only necessary when subscription rules require adjustment and does not start the receiver; the default subscription rule already accepts messages in a normal subscription. Closing the client would stop processing rather than enable it.
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 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.

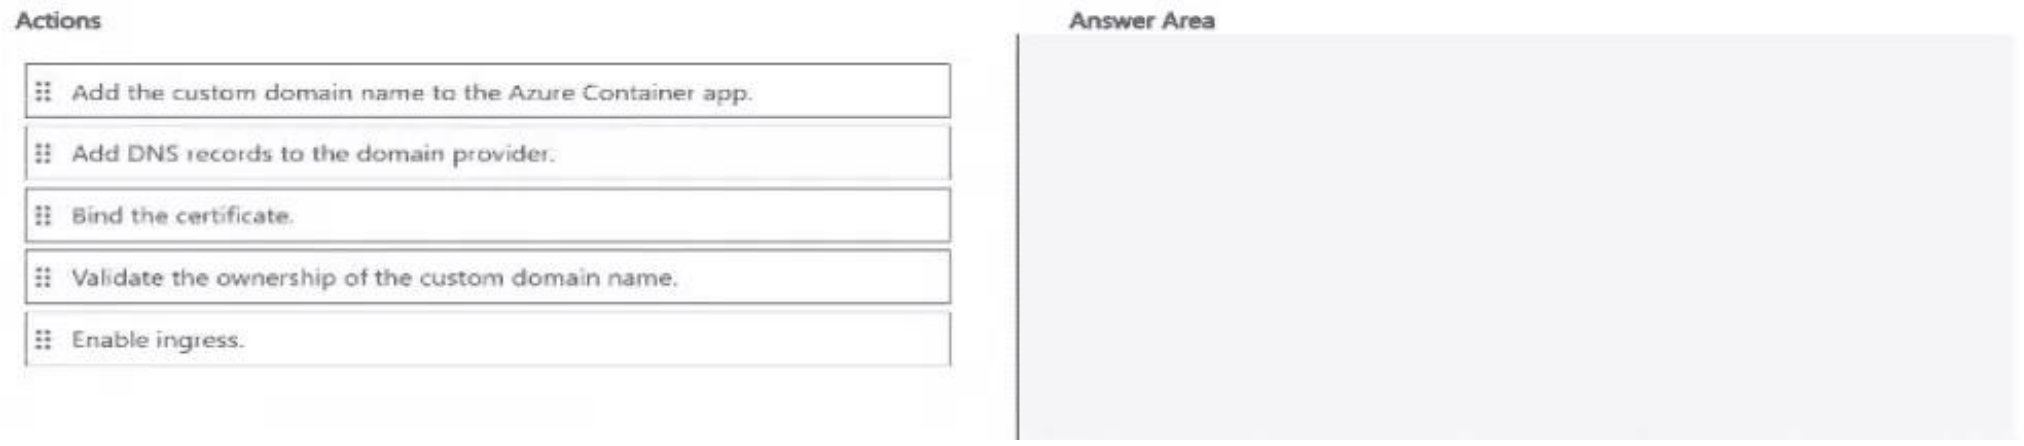

A. Mastered
B. Not Mastered

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 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.

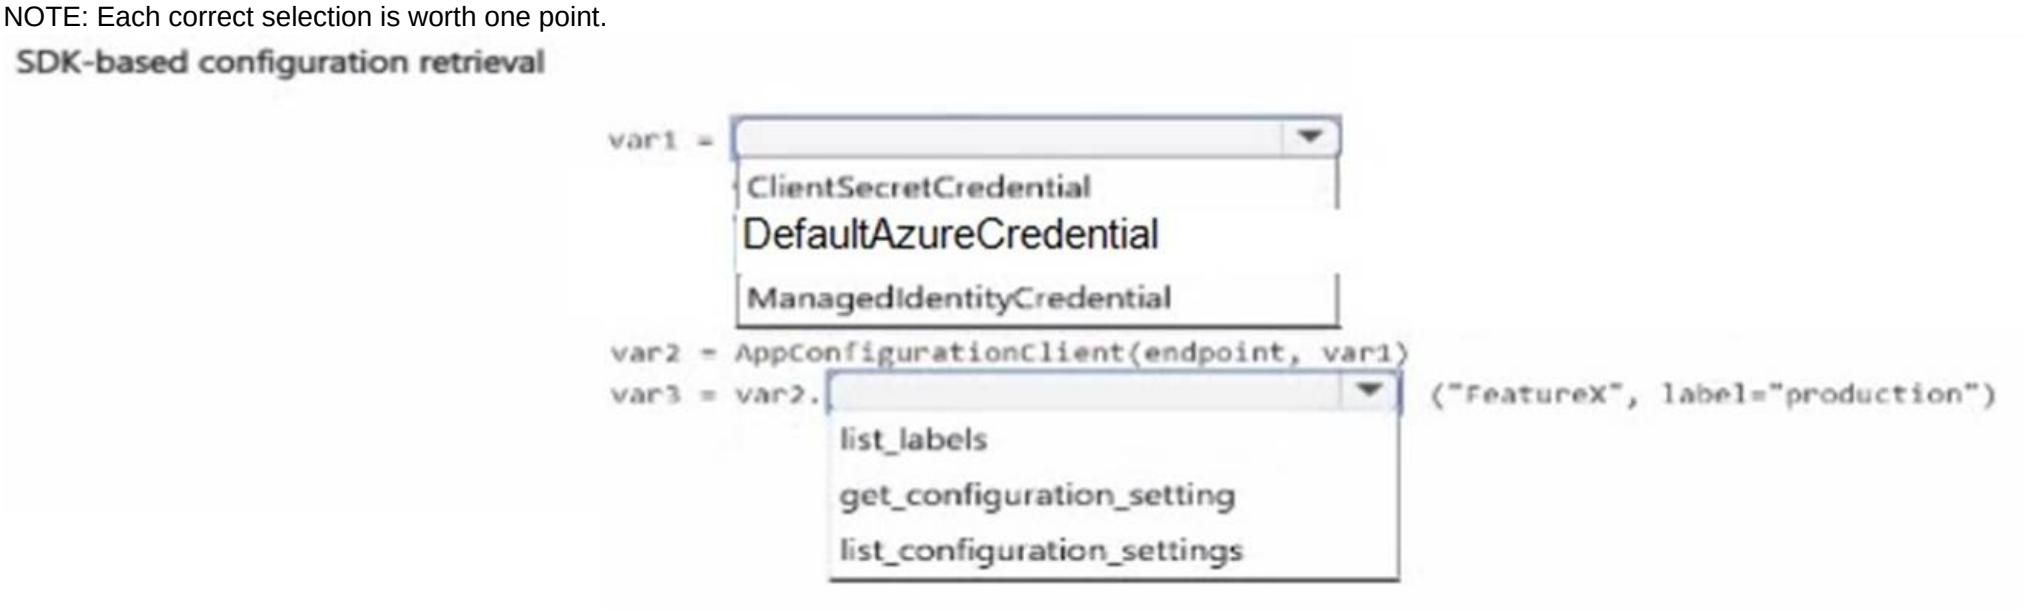

A. Mastered
B. Not Mastered

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 10
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.
Answer Area

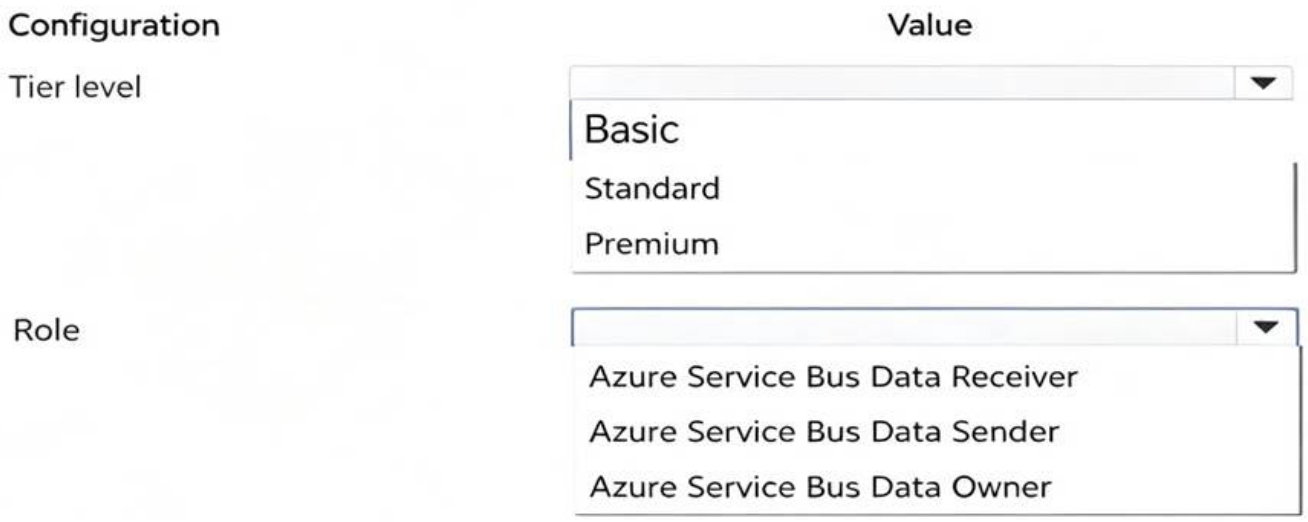

A. Mastered
B. Not Mastered

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 10
You are designing an Azure Database for PostgreSQL table for semantic search. Queries frequently filter on the created_at column. The schema must support vector similarity search and reliable date filtering You need to ensure that the table meets the requirements.
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 embeddings in a pgvector column.
B. Store embeddings in a varchar column.
C. Use a typed timestamp column for created_at.
D. Store created_at as a free-form string.

Answer: AC

Explanation:
A semantic-search schema should use the pgvector vector type for embeddings so PostgreSQL can apply vector distance operators and vector indexes. The creation timestamp should use an actual timestamp data type so comparisons, range predicates, ordering, and indexes use date semantics rather than string ordering. Storing vectors or dates as free-form character data would sacrifice type validation and make similarity or date filtering less efficient and less reliable.
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 14
An ACA app processes messages from an Azure Storage queue.
The app must scale automatically based on messages in a specific Azure Storage queue by using a Kubernetes Event-driven Autoscaler (KEDA) custom scale rule. You need to configure the required scale rule values.
Which two values should you configure? Each correct answer presents part of the solution. Choose two. NOTE: Each correct selection is worth one point.

A. Polling interval
B. Trigger type
C. Maximum replicas
D. Queue name

Answer: BD

Explanation:
A KEDA custom scale rule must identify what scaler to use and which resource within that scaler should be observed. For Azure Storage queues, the trigger type selects the Storage Queue scaler and the queue name identifies the specific queue whose message count drives scaling. Polling interval and maximum replicas are general scaling controls, but they do not define the required queue-specific trigger. Therefore trigger type and queue name are the essential rule values tested here.
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 | Set scaling rules in Azure Container Apps

NEW QUESTION 16
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 21
A container in an AKS cluster repeatedly restarts.
Pod events show probe failures, although node-level CPU and memory metrics are normal.
You need to diagnose the cause of the repeating restarts.
What should you do first?

A. Scale the deployment to more replicas.
B. Decrease the initialDelaySeconds for the container liveness probe.
C. Drain and reboot the node hosting the pod.
D. Inspect the pod events and container logs.

Answer: D

NEW QUESTION 25
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.

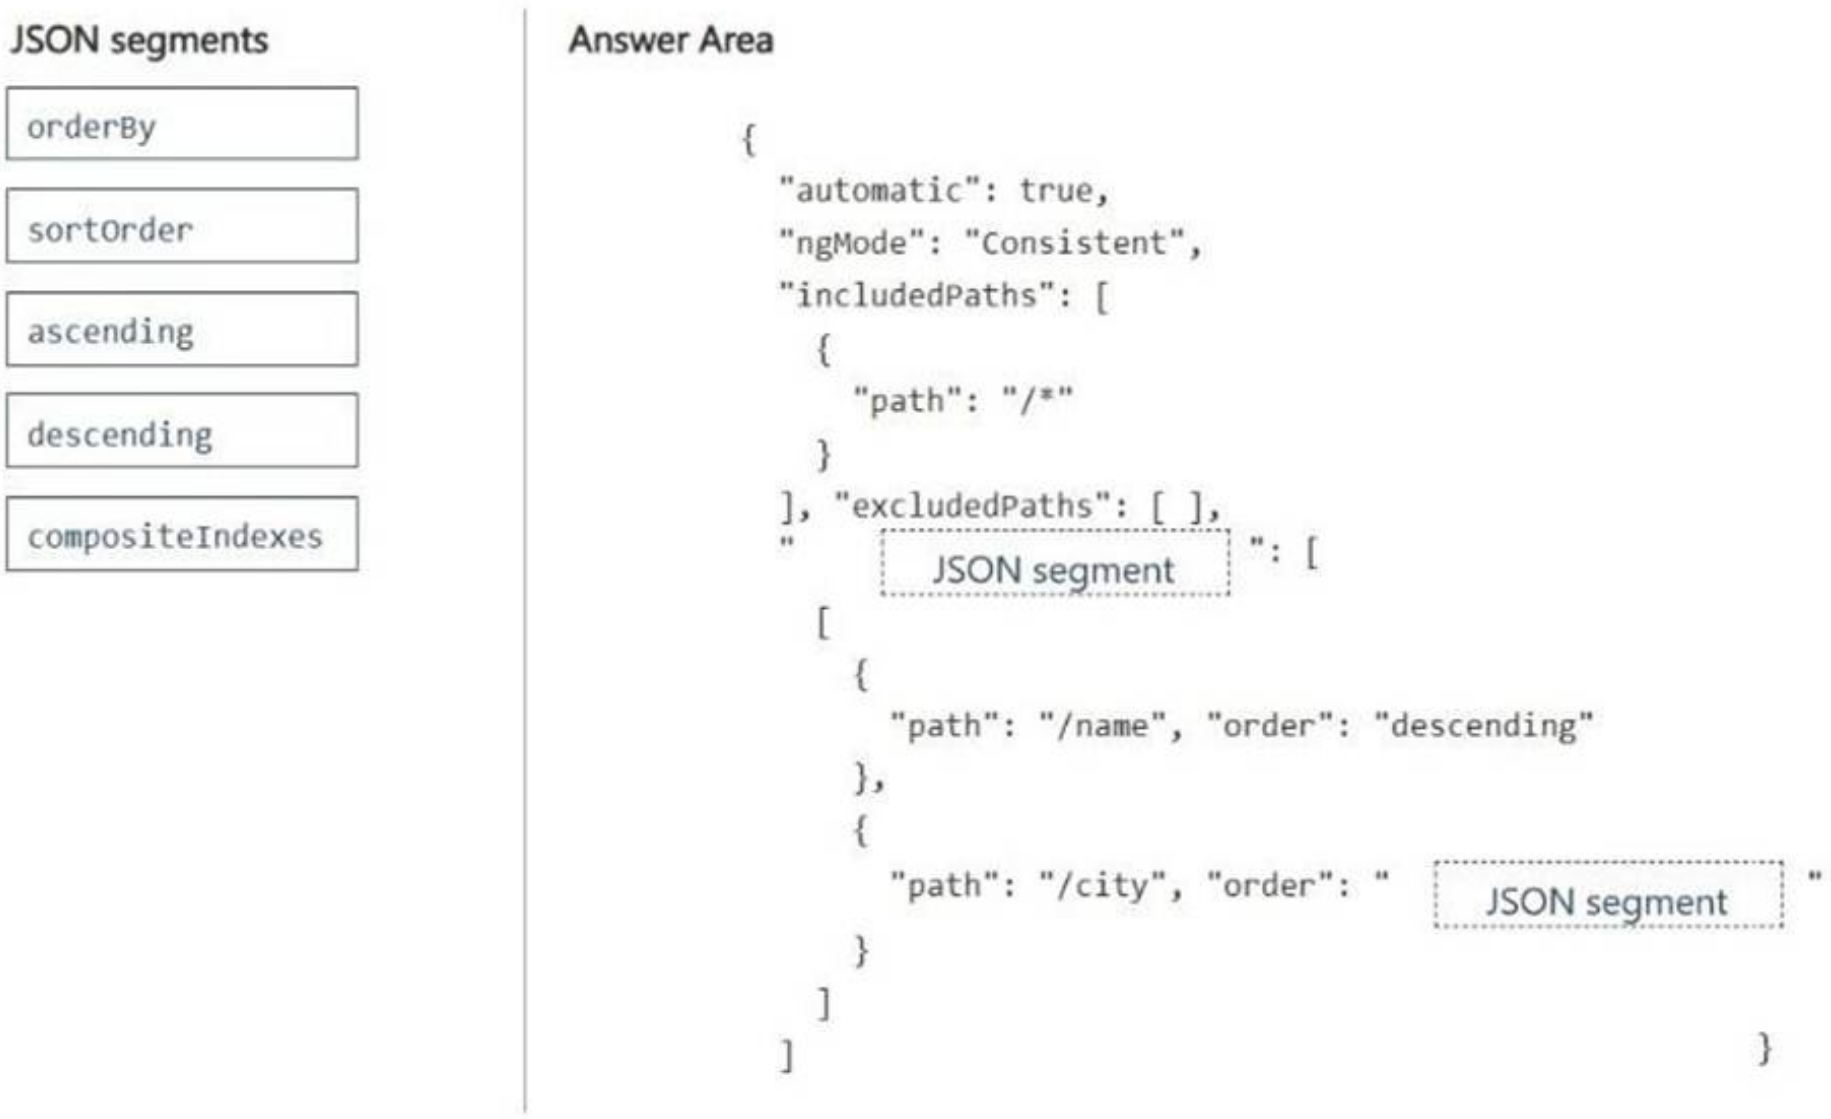

A. Mastered
B. Not Mastered

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 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 deploy multiple instances of a change feed processor to handle a high ingestion rate within Azure Cosmos DB for NoSQL.
Each processor instance must process a different subset of partitions.
You need to ensure the workload is load-balanced across all processor instances.
What should you configure?

A. indexing precision
B. strong consistency
C. lease container
D. autoscale throughput

Answer: C

Explanation:
To handle a high ingestion rate in Azure Cosmos DB for NoSQL, you must deploy multiple instances of the change feed processor using the same workflow name and same lease container.
Reference:
https://github.com/Azure/azure-cosmos-dotnet-v3/wiki/Comprehensive-Test-Suite-for-Azure-Cosmos-DB-Change-Feed-Processor:-All-Versions-and-Deletes

NEW QUESTION 33
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AI-200 Exam with Our Prep Materials Via below:

https://www.certleader.com/AI-200-dumps.html