

Microsoft

Exam Questions AI-200

Developing AI Cloud Solutions on Azure



NEW QUESTION 1

You are investigating high latency in an AI search application that processes millions of requests daily. Telemetry is stored in Azure Monitor Logs.

You must create a KQL query that correlates information from the AppRequests table and the AppDependencies table. The query must meet the following requirements:

Include only data from the last 24 hours.

Filter for failed requests only.

Calculate the average duration of dependencies, grouped by operation.

The query must follow best practice to optimize the performance by minimizing the initial data scan.

You need to create the query.

In which order should 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

Answer Area

Summarize average dependency duration by operation.

Join the dependencies table.

Select the requests table.

Filter the failed requests.

Apply a time filter.

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Verified Answer:1) Select the requests table; 2) apply the 24-hour time filter; 3) filter failed requests; 4) join the dependencies table; 5) summarize average dependency duration by operation.

Detailed Explanation:KQL performs best when high-selectivity filters are applied as early as possible, especially the time predicate that limits the amount of data scanned. Starting from AppRequests, restricting to the last 24 hours, and then filtering failures reduces the left-side dataset before the join. The dependency table is then correlated with those requests, and aggregation is performed last to calculate average dependency duration by operation. Joining or summarizing before the time and failure filters would process more data than necessary.

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 | Optimize log queries in Azure Monitor

NEW QUESTION 2

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.

Ingress

Ingress

Deployment

ClusterIP service

VerticalPodAutoscaler

HorizontalPodAutoscaler

Map Kubernetes resources to scaling and internal networking requirements

Requirement

Scale out the application.

Expose the application internally only.

Resource

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 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 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 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 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.

`-n MyContainerApp -g MyResourceGroup --location eastus2`

env

create

add-on

env

create

add-on

`-n MyContainerApp -g MyResourceGroup --location eastus2`

--enable-mtls

--internal-only

--enable-workload-profiles

- 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 5

You develop an ASP.NET Core app that uses Azure App Configuration. You also create an App Configuration containing 100 settings. The app must meet the following requirements:

- Ensure the consistency of all configuration data when changes to individual settings occur.
- Handle configuration data changes dynamically without causing the application to restart.
- Reduce the overall number of requests made to App Configuration APIs

You must implement dynamic configuration updates in the app.

What are two ways to achieve this goal? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Create and configure Azure Key Vault
- B. Implement the Azure Key Vault configuration provider
- C. Register all keys in the App Configuration store. Set the `refreshAll` parameter of the `Register` method to `false`.
- D. Create and register a sentinel key in the App Configuration store
- E. Set the `refreshAll` parameter of the `Register` method to `true`
- F. Decrease the App Configuration cache expiration from the default value.
- G. Increase the App Configuration cache expiration from the default value.
- H. Create and implement environment variables for each App Configuration store setting

Answer: CE

Explanation:

Detailed Explanation:A sentinel key allows the application to check one value and refresh all related configuration only when that sentinel changes, preserving a consistent set of settings instead of refreshing 100 keys independently. Increasing the refresh/cache interval reduces how frequently the application contacts App Configuration APIs. Together these measures support dynamic refresh without application restart while controlling request volume. Registering every key separately or decreasing the cache interval would increase API traffic, and Key Vault/environment variables do not solve coordinated App Configuration refresh. 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

NEW QUESTION 6

You are developing a .NET application that uses Azure Cosmos DB for NoSQL to store application data The application uses the Azure Cosmos DB for NoSQL SDK to interact with the database account. The application must perform the following tasks

- Initialize the connection by using the account endpoint and key.
- Define shared throughput.
- Perform create, read, update, and delete (CRUD) operations on items stored in a container

Components

Database

CosmosClient

Container

Indexing policy

Answer Area

Azure Cosmos DB for NoSQL SDK object hierarchy

Requirement

Initialize the connection by using the account endpoint and key.

Define shared throughput.

Perform item CRUD operations.

Component

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Verified Answer:Initialize connection: CosmosClient. Define shared throughput: Database. Perform item CRUD: Container. Detailed Explanation:The Cosmos DB SDK follows the service resource hierarchy. CosmosClient represents the account connection and is initialized with the account endpoint and credential. Shared throughput is configured at the database level when multiple containers share database throughput. Item create, read, update, and delete operations are performed through a container object because containers hold items and define partitioning/indexing behavior. An indexing policy is configuration attached to a container, not the SDK object used for CRUD. 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 | Cosmos DB databases, containers, and items

NEW QUESTION 7

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.

Configurations

Use a parameterized query.

Pass the ID as an argument.

Escape quotation marks in the ID.

Supply the parameter tuple to the SDK method.

Secure SDK-based query implementation using parameterized queries

Requirement

Modify the SQL statement structure.

Bind user input.

Execute the statement.

Configuration

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Secure SDK-based query implementation using parameterized queries

Requirement

Modify the SQL statement structure.

Bind user input.

Execute the statement.

Configuration

Use a parameterized query.

Pass the ID as an argument.

Supply the parameter tuple to the SDK

Passing Certification Exams Made Easy

visit - <https://www.surepassexam.com>

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 8

You need to address the known issue resulting from vector similarity queries.
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. Add a composite index on the vector fields and metadata properties of the container.
- B. Reduce the numeric precision of the vector embeddings where supported.
- C. Set the account consistency level to Strong.
- D. Change the vector index type from flat to quantizedFlat or diskANN.

Answer: BD

Explanation:

The objective is to reduce RU consumption from vector similarity queries. Microsoft guidance identifies vector numeric precision and vector-index selection as major cost and performance levers. Using a lower supported vector precision can reduce storage and processing cost, while quantizedFlat and DiskANN are designed to reduce latency and RU consumption compared with flat search for appropriate data sizes. Strong consistency would increase resource cost rather than solve vector-search efficiency. A regular composite index is not a substitute for the vector index. Option B should be read as reducing vector numeric precision, not changing a generic indexing precision setting.
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 Cosmos DB | Optimize Cosmos DB vector search performance

NEW QUESTION 9

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 10

You deploy an API to Azure Container Apps.
The solution must provide the following functionality:

- Support the concurrent activation of multiple application versions.
- Allocate a specific percentage of incoming requests to a secondary version

You need to configure revision behavior.

Which configurations should you use? 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.

Configurations

⋮ Traffic splitting	⋮ Traffic isolation
⋮ Single revision mode	⋮ Multiple revision mode

Revision management and traffic routing in Azure Container Apps

Requirement	Configuration
Test a new revision without affecting production traffic.	
Route 20 percent of traffic to a new revision.	

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Azure Container Apps supports side-by-side application versions through multiple revision mode. Once multiple revisions can be active, ingress traffic can be distributed by percentage across those revisions. This is the required combination for canary or controlled rollout scenarios. Single revision mode deactivates the previous revision as a new one becomes active, while traffic splitting is the specific feature that assigns a percentage of incoming requests to each active revision. 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 10

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 13

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 17

You process Azure Service Bus messages that require a dependent external API call.

If the API is temporarily unavailable, you must delay processing of the message without incrementing the delivery count. You need to find a way to process the message when the API is available while keeping the message accessible. Which message action should you perform?

- A. Defer
- B. Complete
- C. Dead-letter
- D. Abandon

Answer: A

Explanation:

Deferral is intended for a message that cannot be processed now because a particular dependency is temporarily unavailable, while the application wants to keep the message in the main queue for later retrieval by sequence number. Complete would permanently remove it, dead-letter would move it aside as a failed message, and abandon immediately releases the lock for redelivery. Deferral therefore best matches the requirement to postpone the individual message while keeping it available for deliberate later processing.

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 deferral | Service Bus message transfers, locks, and settlement

NEW QUESTION 18

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.

SDK-based configuration retrieval



- 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 23

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.

Answer Area

Requirement:

Single URL for test users

Feature

	▼
Revision Label	
Revision mode	
Container image	
Container registry	

	▼
Revision Label	
Revision mode	
Container image	
Container registry	

Current microservice activation

- 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 24

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 29

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 33

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.

Query components

WHERE department = 'Finance'

WHERE department LIKE 'dep%'

ORDER BY created_at DESC LIMIT 5

ORDER BY embedding <=> query_embedding LIMIT 5

Vector similarity search with metadata filtering

Requirement

Filter rows to finance.

Rank by similarity and return the top five.

Query component

- 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 37

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 40

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 44

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 48

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.

Actions

- Add the custom domain name.
- Bind the certificate.
- Enable ingress.
- Validate the custom domain name.
- Add DNS records to the domain provider.

Answer area

1.
2.
3.
4.
5.

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer area

1.
2.
3.
4.
5.

NEW QUESTION 49

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.

JSON segments

orderBy

sortOrder

ascending

descending

compositeIndexes

Answer Area

```
{
  "automatic": true,
  "ngMode": "Consistent",
  "includedPaths": [
    {
      "path": "/*"
    }
  ], "excludedPaths": [ ],
  "JSON segment": [
    {
      "path": "/name", "order": "descending"
    },
    {
      "path": "/city", "order": "JSON segment"
    }
  ]
}
```

- 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 52

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 - Questions 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 First Try
- * AI-200 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AI-200 Practice Test Here](#)