

Microsoft AI-200

Developing AI Cloud Solutions on Azure



For More Information – Visit link below:

<https://www.examsboost.com/>

Product Version

- ✓ Up to Date products, reliable and verified.
- ✓ Questions and Answers in PDF Format.

Latest Version: 6.0

Topic 1, Proseware Inc.

Background

Proseware Inc. develops AI-powered knowledge management solutions for enterprise customers. The company is modernizing its platform to support semantic search, intelligent document retrieval, and real-time partner integrations.

The engineering team uses Python and Azure SDKs. The architecture is being redesigned to support containerized microservices, vector search workloads, and serverless backend processing.

Planned Application Architecture

- ☑ Microservices are containerized by using Docker.
- ☑ Code for containerized microservices and Azure Function apps is developed locally but stored in a GitHub repository.
- ☑ Custom images for containerized microservices are stored in Azure Container Registry (ACR).
- ☑ Base images are stored in Docker Hub. Custom images must be rebuilt automatically whenever their base images are updated.
- ☑ Azure Cosmos DB for NoSQL stores documents, metadata, and vector embeddings.
- ☑ Azure Functions generate vector embeddings of Azure Cosmos DB for NoSQL-hosted documents and send messages to Service Bus to trigger search index updates.
- ☑ Azure Container Apps (ACA) apps host backend API services that provide semantic search across Azure Cosmos DB for NoSQL documents. API services process Service Bus messages and update search indexes.
- ☑ Azure Kubernetes Service (AKS) processes batch vector embedding regeneration for existing Azure Cosmos DB for NoSQL documents (whenever the embedding model is changed).
- ☑ An extranet-facing containerized webhook allows business partners to submit documents to be processed by internal AI workflows for semantic search and retrieval.

Monitoring

Telemetry generated by Azure resources is sent to Azure Monitor.

A Log Analytics workspace is used to collect ACA apps logs, AKS container logs, and Azure Functions apps logs.

Monitoring of Azure Functions is currently implemented by using Azure Application Insights SDK instrumentation.

Business Requirements

Embeddings for new or updated Azure Cosmos DB for NoSQL-hosted documents must be automatically generated.

Backend API services must scale automatically during business hours.

Cold start delay of backend APIs must be minimized.

Secrets must be stored outside of container images.

Developers must be able to correlate telemetry across Azure Functions hosts and apps.

All tracing must be implemented by using OpenTelemetry SDK instrumentation.

Development efforts must be minimized.

Technical Requirements

- ☑ Container images must be built automatically and validated before code updates are merged into the main branch.
- ☑ Image build automation must run inside the Azure Container Registry, eliminating dependency on local developer machines and external build services.
- ☑ Dependency of image builds on local developer machines must be eliminated.
- ☑ Event-driven scaling in ACA must occur based on the number of pending messages in the Azure Service Bus queue.
- ☑ Azure Cosmos DB for NoSQL RU consumption must be minimized.
- ☑ Vector similarity search must use embeddings stored in Azure Cosmos DB for NoSQL.
- ☑ The partner-facing containerized webhook service must run on Azure App Service.
- ☑ Secrets must NOT be stored in container images, source control, or application configuration directly. They must be accessed securely at runtime.
- ☑ All secrets must be stored centrally in Azure Key Vault and accessed at runtime through a managed identity.
- ☑ Azure App Service must supply secrets at runtime without relying on external services.
- ☑ Resources and workloads must be deployed by using Bicep templates through an automated, version-controlled pipeline. Local and command-line deployments must be eliminated to ensure repeatable, auditable deployments.

Known issues

RU consumption spikes during vector similarity queries.

Question: 1

DRAGDROP

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.

Actions

- Redeploy the instrumented services.
- Configure a trace exporter in the OpenTelemetry SDK.
- Instrument the application code by using the OpenTelemetry SDK.
- Call `TelemetryClient.TrackEvent()` within service methods.
- Implement a view in the OpenTelemetry SDK.

Answer Area

Answer:

Actions

- Redeploy the instrumented services.
- Configure a trace exporter in the OpenTelemetry SDK.
- Instrument the application code by using the OpenTelemetry SDK.
- Call `TelemetryClient.TrackEvent()` within service methods.
- Implement a view in the OpenTelemetry SDK.

Answer Area

- Instrument the application code by using the OpenTelemetry SDK.
- Configure a trace exporter in the OpenTelemetry SDK.
- Redeploy the instrumented services.

Question: 2

You need to deploy Azure function resources and apps by using an automated, version-controlled failed 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

Question: 3

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. Modify the indexing precision of the vector fields.
- C. Set the account consistency level to Strong.
- D. Change the vector index type from flat to quantizedFlat or diskANN.

Answer: B, D

Question: 4

HOTSPOT

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 select is worth one point.

Update vector embeddings

Requirement

Identify document changes that should trigger vectorization.
Scale out the vectorization processing.

Scale out the vectorization processing.

Configuration

Change feed processor
Periodic full container scan
Cross-partition SELECT query

Lease container
Strong consistency level
RU throughput on the container

Answer:

Update vector embeddings

Requirement

Identify document changes that should trigger vectorization.
Scale out the vectorization processing.

Scale out the vectorization processing.

Configuration

Change feed processor
Periodic full container scan
Cross-partition SELECT query

Lease container
Strong consistency level
RU throughput on the container

Question: 5

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

Thank You for Trying Our Product

For More Information – **Visit link below:**

<https://www.examsboost.com/>

15 USD Discount Coupon Code:

G74JA8UF

FEATURES

- ✓ **90 Days Free Updates**
- ✓ **Money Back Pass Guarantee**
- ✓ **Instant Download or Email Attachment**
- ✓ **24/7 Live Chat Support**
- ✓ **PDF file could be used at any Platform**
- ✓ **50,000 Happy Customer**



Visit us at: <https://www.examsboost.com/test/ai-200>