



Accelerate Your Certification with Microsoft AI-200 Practice Questions

Last chance to prepare smart! Get your hands on free Microsoft Developing AI Cloud Solutions on Azure exam PDF questions. Study real AI-200 dumps with verified answers and fast-track your certification success with [PrepBolt](https://prepbolt.com/AI-200.html) AI-200 exam pdf questions and answers.

Thank you for Downloading AI-200 exam PDF Demo

<https://prepbolt.com/AI-200.html>

QUESTIONS & ANSWERS
DEMO VERSION
(LIMITED CONTENT)

Question 1

Question Type: Hotspot

Case Study: Mix Questions

Mix Questions

AI-200 Mix Questions IN THIS CASE STUDY

You are building a semantic search feature for a chatbot. You store document embedding 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.		☐	☐

Answer:

See the Answer in the Premium Version!

Question 2

Question Type: Hotspot

Case Study: Mix Questions

Mix Questions

AI-200 Mix Questions IN THIS CASE STUDY

You are developing a Retrieval-Augmented Generation (RAG) solution for a company. AI responses and

embedding vectors are cached in Redis. The solution must meet the following requirements.

- * AI responses must expire exactly 24 hours after they are cached
- * Cached embeddings must always reflect the current source data.

You need to configure Redis to meet the requirements.

Cache expiration and invalidation strategy

Requirement	Action
AI responses must expire exactly 24 hours after they are cached.	Configure allkeys-lru eviction. Reset expiration on every read. Set a time to Live (TTL) on each key.
Cached embeddings must always reflect the current source data.	Configure volatile-lru eviction. Delete related keys when the document changes Increase maximum memory allocation.

Answer:

See the Answer in the Premium Version!

Question 3

Question Type: DragDrop

Case Study: Mix Questions

Mix Questions

AI-200 Mix Questions IN THIS CASE STUDY

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

Answer:

See the Answer in the Premium Version!

Question 4

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

AI-200 Mix Questions IN THIS CASE STUDY

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?

Options:

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

Answer:

D

Question 5

Question Type: Hotspot

Case Study: Mix Questions

Mix Questions

AI-200 Mix Questions IN THIS CASE STUDY

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

Answer:

See the Answer in the Premium Version!

Question 6

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

AI-200 Mix Questions IN THIS CASE STUDY

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.

Options:

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

Answer:

B, D

Question 7

Question Type: DragDrop

Case Study: Mix Questions

Mix Questions

AI-200 Mix Questions IN THIS CASE STUDY

You plan to deploy a web application to AKS.

The solution must:

- * Scale out the application oy add ng more pods during peak CPU usage
- * Expose the application internal y within the cluster only

You need to configure a Kubemetes resource for each requirements

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.

Resources

Ingress

Deployment

ClusterIP service

VerticalPodAutoscaler

HorizontalPodAutoscaler

Map Kubernetes resources to scaling and internal networking requirements

Requirement	Resource
Scale out the application.	
Expose the application internally only.	

Answer:

See the Answer in the Premium Version!

Question 8

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

AI-200 Mix Questions IN THIS CASE STUDY

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. Choose two

NOTE: Each correct selection is worth one point.

Options:

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

Answer:

A, B

Thank You for trying AI-200 PDF Demo

To try our AI-200 practice exam software visit
link below

<https://prepbolt.com/AI-200.html>

Start Your AI-200 Preparation

Use Coupon “**SAVE50**” for extra 50% discount on the purchase of Practice Test Software. Test your AI-200 preparation with actual exam questions.