

CWNP CAIAE-101

CWNP AI Administrator and Engineer (CAIAE)

For More Information – Visit link below:

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

Product Version

- 1. Up to Date products, reliable and verified.**
- 2. Questions and Answers in PDF Format.**



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/caiae-101>

Latest Version: 6.0

Question: 1

An organization is standing up its AI risk program under the NIST AI Risk Management Framework. Leadership first wants to establish the policies, accountability structures, and risk culture that apply across every AI project.

Which NIST AI RMF function covers this work?

- A. Map
- B. Measure
- C. Govern
- D. Manage

Answer: C

Question: 2

A vector database has grown to hundreds of millions of embeddings, and exact brute-force similarity search can no longer meet the latency target.

Which indexing approach best restores fast retrieval?

- A. An approximate nearest neighbor index such as HNSW.
- B. A B-tree index built on the document IDs.
- C. Exact k-nearest-neighbor search across the full set.
- D. A keyword inverted index over the source text.

Answer: A

Question: 3

A developer wants to quickly run an open-weight LLM on a personal laptop for experimentation, using a single command and no cluster.

Which tool best fits this local-development need?

- A. vLLM
- B. Triton Inference Server
- C. KServe
- D. Ollama

Answer: D

Question: 4

A team is planning a multi-GPU cluster for distributed training of a large model, where the GPUs must exchange gradients frequently and rapidly during every training step. Which interconnect choice best fits this requirement?

- A. Standard 1 Gigabit Ethernet between the nodes, since it is the most widely available and lowest-cost option.
- B. High-bandwidth, low-latency interconnects such as NVLink or InfiniBand.
- C. A public-internet VPN between the GPUs so they can sit in different data centers for resilience.
- D. USB-attached external GPUs connected to a single workstation to keep the cabling simple.

Answer: B

Question: 5

Before any AI workload is deployed, an administrator must provision the underlying cloud infrastructure: a managed Kubernetes cluster, GPU node pools, and the surrounding networking. Which tool is designed for this infrastructure-provisioning step?

- A. Helm
- B. Ollama
- C. Terraform
- D. A model registry

Answer: C

Question: 6

An AI system is built to perform one specific task, such as filtering spam email, and cannot transfer its capability to unrelated problems. Which category of artificial intelligence does this describe?

- A. General AI
- B. Symbolic AI
- C. Superintelligence
- D. Narrow AI

Answer: D

Question: 7

A data subject exercises their right to be forgotten. A production model was already trained on a dataset that included that person's records, and deleting the source rows is not enough because the model has learned from them.

Which approach directly removes that individual's influence from the trained model?

- A. Data minimization
- B. Machine unlearning
- C. Differential privacy
- D. Encryption at rest

Answer: B

Question: 8

An engineer must load a model checkpoint downloaded from a public hub whose publisher is not fully trusted.

Which serialization format is the safer choice for loading that artifact?

- A. A Python pickle checkpoint
- B. Whichever format compresses to the smaller file
- C. The safetensors format
- D. Whichever format loads faster on the GPU

Answer: C

Question: 9

A healthcare provider requires that no patient data ever leaves its own controlled environment and that it retains full control over which model versions are used. It is comparing a closed hosted API model against a self-hosted open-weight model.

Which option best satisfies the stated constraints?

- A. Self-host an open-weight model inside the provider's own environment.
- B. Use the closed hosted API and rely on the vendor's contractual data-processing terms to keep the data private.
- C. Use the closed hosted API but enable client-side request logging so the provider can audit everything that was sent.
- D. Self-host the closed model by downloading its published weights from the vendor's public endpoint.

Answer: A

Question: 10

During risk planning for an AI project, a team wants a structured technique to systematically identify potential failure modes, rate their severity, likelihood, and detectability, and prioritize them for mitigation.

Which technique best fits this description?

- A. A confusion matrix, which summarizes a classifier's prediction errors on a test set.
- B. A risk register, which is the living log used to record and track risks once they are identified.
- C. A model card, which documents a model's intended use and evaluation results.
- D. Failure mode and effects analysis (FMEA).

Answer: D

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>