PRACTICE QUESTIONS

EC-Council

Certified Responsible AI Governance & Ethics (CRAGE)

PUBLISHED BY aiexamcerts

EDITION	Free Set 2
QUESTIONS	10
ISSUED	November 8, 2026
EXAM	EC-Council Certified Responsible AI Governance & Ethics (CRAGE)

Shared by **aiexamcerts** on aiexamcerts

10 practice questions for the EC-Council Certified Responsible AI Governance & Ethics (CRAGE) exam, each followed by its answer and a short explanation. More free sets and the full practice exam: **aiexamcerts.com/ec-council/certified-responsible-ai-governance-ethics-crage**. Spotted a mistake? Email support@aiexamcerts.com with the exam name and the question number.

---

1. Which two descriptions of the EU AI Act's risk approach are accurate? (Choose two.)

- **A** High-risk AI systems must have a documented risk management system that runs across their lifecycle
- **B** Prohibited practices may be used for high-risk systems if the provider publishes a model card
- **C** A system's risk category depends mainly on the size of the company that built it
- **D** Limited-risk systems must tell users when they are interacting with AI, in the cases the Act specifies
- **E** The risk tiers apply only to AI systems that are developed inside the European Union

Answer: A, D

High-risk systems carry lifecycle duties, including a documented risk management system, and limited-risk systems carry transparency duties such as disclosure. Publishing a model card is the most tempting wrong answer because documentation is central to governance, but it never makes a prohibited practice lawful. The claims about company size and development location are wrong because the tiers depend on what a system is used for.

---

2. What deployment model keeps both patient records and the model inside a hospital's own infrastructure, with no patient data sent to outside services?

- **A** On-premises deployment on infrastructure that the hospital owns and controls
- **B** Public multi-tenant software-as-a-service hosted by the model vendor
- **C** Calls to a shared third-party API over the public internet
- **D** Community model hosting on a public marketplace shared with other customers

Answer: A

On-premises deployment keeps both the data and the model on infrastructure the hospital controls, which meets the no-outside-services requirement. Public multi-tenant software-as-a-service is the most tempting wrong answer because it is the easiest to adopt, but it places patient data on a vendor's shared platform. The shared API and the marketplace options both send patient data to outside parties.

---

3. At the bottom of the four EU AI Act risk tiers, below limited risk, sits _____ risk, which carries few obligations.

Answer: minimal (also accepted: minimal or no)

The lowest tier is minimal risk, which carries few obligations beyond voluntary codes of conduct. Limited risk is the category with transparency duties such as disclosure, so it describes a different level of obligation. Unacceptable risk covers the prohibited practices.

---

4. After a speech recognition product shows a much higher error rate for one regional dialect than for the others, which step most directly addresses the root cause?

A Raising the confidence threshold so that uncertain transcriptions are rejected for all users

B Auditing the training data and adding representative recordings of that dialect

C Adding a disclaimer on the product page that accuracy varies by accent

D Hiding the per-group error results from the model documentation

Answer: B

The root cause is that the training data under-represents the dialect, so auditing the data and adding representative recordings addresses it directly. Raising the confidence threshold is the most tempting wrong answer because it looks like a safety fix, but it changes behavior for every user and leaves the gap in place. A product-page disclaimer and hiding the per-group results only disclose or conceal the problem.

---

5. To learn whether a generative AI tool saves time in claims processing before committing budget, which approach gives the most reliable answer?

A A full enterprise rollout to every claims handler, with usage tracked but never compared to baseline

B A multi-year contract signed before any testing starts, so the vendor can tailor the product

C A bounded pilot with defined success criteria, a comparison group, and a scale-or-stop decision

D A survey asking executives how they feel about the tool after reading the vendor brochure

Answer: C

A bounded pilot with a comparison group and agreed success criteria produces evidence that can justify scaling or stopping. A full rollout without a baseline is the most tempting wrong answer because it looks like a faster route to results, but usage data alone cannot show that time was saved. A multi-year contract before testing and an executive survey both commit or decide without evidence.

---

6. A time-stamped, tamper-evident record of who accessed or changed a system, and what they did, is called an audit _____.

Answer: trail (also accepted: log)

An audit trail records actions in sequence so that auditors and investigators can reconstruct what happened and who did it. An audit log is accepted as a variant because it is the same kind of record in common usage. Tamper evidence matters because a record that can be edited without trace has little value as evidence.

---

7. What artifact lets a compliance lead show how one set of AI controls satisfies several overlapping regulations and frameworks at once?

A A press release announcing that the organization takes responsible AI seriously

B A monthly invoice register listing each AI vendor and the amount paid to it

C A completion report showing how many staff finished the annual AI awareness course

D A control crosswalk that maps each control to every requirement it helps satisfy

Answer: D

A crosswalk shows, for each control, which requirements across the regulations and frameworks it satisfies, so one control can be evidenced once and reused. A training completion report is the most tempting wrong answer because it looks like compliance evidence, but it shows attendance rather than whether requirements are met. The press release and the invoice register do not map controls to requirements at all.

---

8. When a production credit model starts denying loans at a much higher rate for one group after a data pipeline change, what is the most appropriate immediate containment step?

A Roll back to the last validated model or send decisions to manual review

B Wait a month for more complaints so that the pattern is clearly established

C Retrain the model during business hours with a higher learning rate to correct the drift

D Publish a blog post describing the model's recent accuracy gains for the public

Answer: A

Rolling back or routing decisions to manual review limits further harm while the root cause is investigated. Waiting for more complaints is the most tempting wrong answer because it looks like careful evidence gathering, but harmful decisions continue in the meantime. Retraining during business hours without analysis and publishing a blog post both add risk or ignore the incident.

---

9. The matrix that assigns each AI governance activity one Responsible, Accountable, Consulted, and Informed role is called a _____ matrix.

Answer: RACI

The matrix is named for its four roles: Responsible, Accountable, Consulted, and Informed. Keeping exactly one accountable party per activity prevents the diffuse ownership that leaves governance gaps. A risk register is easy to confuse with it, but a risk register records risks and their treatment rather than who owns each activity.

---

10. During a peak-load failure of a customer-facing AI triage tool, which business continuity design best keeps critical customer services running?

A Disabling all customer service until the tool has been repaired and retested

B A documented manual fallback process, with trained staff and a tested switchover

C Increasing the model size so that the tool can absorb peak load without failing

D Removing monitoring so that the tool responds faster and never reports overload

Answer: B

A documented manual fallback, with trained staff and a tested switchover, keeps critical services available when the tool fails. Increasing the model size is the most tempting wrong answer because it sounds like added capacity, but a larger model usually needs more compute and does not provide a fallback. Stopping all customer service and removing monitoring both leave customers without service or without visibility.