PRACTICE QUESTIONS

EC-Council

Certified Responsible AI Governance & Ethics (CRAGE)

PUBLISHED BY aiexamcerts

EDITION	Free Set 1
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 practices reduce the risk of automation bias when staff rely on an AI recommendation? (Choose two.)

- **A** Treating benchmark accuracy above 90 percent as proof that individual outputs need no review
- **B** Logging each human override and reviewing how often recommendations are accepted without checks
- **C** Removing the human review step once the model has run in production for a full year
- **D** Displaying the model's confidence and main input factors so reviewers can question the output
- **E** Showing only the recommended action, with no alternatives or supporting detail

Answer: B, D

Logging overrides and watching for unchecked acceptance reveal when people have stopped questioning the model, and showing confidence and drivers gives reviewers something to challenge. Treating benchmark accuracy as proof is the most tempting wrong answer because it sounds like a sound reason to trust the model, but aggregate accuracy does not guarantee any single output. Removing review after a year and showing only the recommended action both make automation bias worse.

---

2. The business owner of an AI system is the right person to carry out the independent validation of that system's model.

- **A** True
- **B** False

Answer: B (False)

Independent validation is performed by people who did not build or own the system, so that they can challenge it objectively. The business owner has a stake in the system succeeding, which is why the owner should not perform the validation. So the statement is false.

---

3. What artifact best supports supply chain transparency for an AI system by listing its components, including libraries, pretrained models, and datasets, with their origins?

A A penetration test report that lists the vulnerabilities found during the last assessment

B A service level agreement summary showing uptime for the previous twelve months

C A software bill of materials extended to cover AI components and their origins

D A brand style guide that defines the approved logos, colors, and typefaces

Answer: C

A software bill of materials records components and their origins, and extending it to models and datasets gives visibility into an AI supply chain. The penetration test report is the most tempting wrong answer because reviewers ask for it routinely, but it reports findings from a test rather than listing components. The uptime summary and the brand guide do not describe what the system is made of.

---

4. Requiring a fraud model to produce the same true positive rate and the same false positive rate for every customer group describes which fairness criterion?

A Predictive parity, which requires equal precision for predicted positives in each group

B Demographic parity, which requires equal rates of positive predictions across groups

C Counterfactual fairness, which asks whether a decision would change if only a protected attribute changed

D Equalized odds, which requires equal true and false positive rates across groups

Answer: D

Equalized odds requires equal true positive and false positive rates across groups, which is exactly the condition in the scenario. Demographic parity is the most tempting wrong answer because it is the most familiar fairness metric, but it compares only the rate of positive predictions and ignores whether those predictions are correct. Predictive parity and counterfactual fairness test different properties.

---

5. The EU AI Act can apply to a provider established outside the European Union when its AI system is placed on the EU market or put into service there.

A True

B False

Answer: A (True)

The Act's reach follows where an AI system is placed on the market or put into service, not where the provider is established. A non-EU vendor selling into the EU therefore has to meet the Act's obligations for that system. So the statement is true.

---

6. An auditor finds that a bias review step is written into the approved release procedure, but it was skipped for 6 of the 30 releases sampled in the audit period. What is the auditor's conclusion?

A The control is designed appropriately but did not operate effectively during the period

B The control is ineffective by design, so the procedure itself must be rewritten

C The control was never documented, so its design cannot be assessed at all

D The control falls outside the audit scope because it concerns a model rather than a server

Answer: A

The procedure is well designed, but the sample shows it was not performed consistently, which is a failure of operating effectiveness. Calling the design ineffective is the most tempting wrong answer because the failures are visible and the procedure looks like the problem, but the design itself is adequate. The claim that the control was never documented contradicts the scenario, and scoping it out would hide the gap.

---

7. Faced with unsafe medical advice from a generative feature and a fix that would take longer than the business will wait, leadership removes the feature entirely. Which risk treatment option does this decision represent?

A Transfer, by moving the financial exposure to an insurer or contract partner

B Avoid, by removing the activity that creates the risk

C Accept, by documenting the risk and proceeding without further action

D Mitigate, by adding controls that reduce the likelihood or impact of the risk

Answer: B

Removing the feature eliminates the activity that creates the risk, which is avoidance. Mitigation is the most tempting wrong answer because the team could have added controls, but leadership chose to stop the activity instead of reducing the risk. Transfer shifts financial exposure without removing the risk, and acceptance keeps the risk in place.

---

8. Which two contract terms help an enterprise control how a vendor handles confidential data used with its AI service? (Choose two.)

A A license for the vendor to share outputs with other customers as examples of good results

B A clause that ends all data-handling duties on the day the contract is signed

C A ban on using the enterprise's data to train the vendor's shared models without written consent

D A promise that the vendor will publish its internal pricing model on a public website

E A right to audit the vendor or obtain independent assurance reports such as a SOC 2 report

Answer: C, E

A training restriction stops the vendor from reusing confidential data to improve models for others, and audit rights or SOC 2 reports give evidence that controls operate. Output sharing is the most tempting wrong answer because examples of good results sound like a harmless customer benefit, but it would release confidential material to other customers. The clause ending data duties at signature and the pricing promise have no role in protecting the data.

---

9. Needing a document summarization capability quickly, with no in-house machine learning team and no need for custom model logic, which strategy should a mid-size insurer choose?

A Build a foundation model from scratch on the company's own document archive

B Hire a research team to train a new in-house model over the next several years

C Buy a vendor-hosted service and govern it through third-party risk controls

D Postpone the capability until the company finishes a full data warehouse migration

Answer: C

A vendor-hosted service delivers the capability quickly without model-building skills, and the insurer manages the residual risk through due diligence and contract controls. Building a foundation model is the most tempting wrong answer because it promises full control, but it needs resources the insurer does not have and would take far longer. Hiring a research team and postponing for a warehouse migration are both unnecessary delays for a summarization need.

---

10. Before a public launch of a generative assistant, which practice most directly reduces the risk that harmful content reaches users?

A Adding a footer that says responses may be inaccurate and should be checked by the user

B Publishing the full system prompt so that users can review the assistant's instructions

C Reducing the maximum response length so that the assistant says less overall

D Harm-specific red teaming followed by a staged rollout with active monitoring

Answer: D

Harm-specific red teaming finds unsafe behavior before users do, and a staged rollout with monitoring catches what testing misses. A footer warning is the most tempting wrong answer because disclaimers are a common safety measure, but a warning does not stop harmful content from being generated. Publishing the system prompt and shortening responses do not reliably prevent harmful outputs.