

[Your Company Name] — AI Fundamentals & Safety — Quiz Bank (12 Questions)

Version: 1.0 | Date: 2025-10-14 | Passing Score: 85% | Attempts: 2

Multiple Choice — Select the best answer

1) Which tools should staff use for company work?

- A. Any chatbot if content seems harmless
- B. Only allow-listed enterprise AI (e.g., M365 Copilot Enterprise, Azure OpenAI, Amazon Bedrock)
- C. Personal AI accounts are fine
- D. Whatever the team prefers

Answer: B

Rationale: Use only approved enterprise tools; consumer endpoints should be disabled for work accounts.

2) Which input is prohibited in AI prompts?

- A. Public product brochure
- B. Generic, non-sensitive meeting notes
- C. Secrets/keys or device logs with identifiers
- D. A link to a public RFP

Answer: C

Rationale: Secrets, credentials, and identifiable logs must never be pasted into AI tools.

3) Human-in-the-loop means:

- A. AI results can be sent externally without review
- B. A qualified person reviews and approves consequential outputs before use
- C. Two AI systems review each other
- D. No review is needed

Answer: B

Rationale: Consequential actions require documented human approval.

4) Facial recognition/biometrics are:

- A. Allowed by default
- B. Allowed with manager approval only
- C. Disallowed unless a formal exception is approved with evidence and legal sign-off
- D. Irrelevant to MSPs

Answer: C

Rationale: Risky features require formal exceptions and legal approval.

5) If you paste client PI/PHI into a non-approved tool, first you should:

- A. Close the browser and ignore it
- B. Try again in another tool
- C. Stop and notify Security immediately with details
- D. Email the client to apologize

Answer: C

Rationale: Treat as a security incident; escalate immediately.

6) Which statement about data training is most accurate?

- A. Enterprise AI always trains on your prompts
- B. Enterprise AI can be configured so prompts/responses are not used to train foundation models
- C. Training status doesn't matter for internal drafts
- D. None of the above

Answer: B

Rationale: TAG members should configure enterprise AI so customer content is not used to train base models.

7) Where should AI-assisted work be stored?

- A. Personal cloud
- B. Local desktop
- C. Approved repositories with retention policies
- D. Email drafts forever

Answer: C

Rationale: Keep work products only in approved systems with retention controls.

8) Which scenario is acceptable?

- A. Using a personal chatbot to summarize a client incident report
- B. Uploading unreleased firmware to a public AI to get help
- C. Using M365 Copilot Enterprise to re-format non-sensitive internal notes, then getting human review
- D. Posting a BWC clip with faces to a consumer video AI

Answer: C

Rationale: Internal non-sensitive text in approved tools is acceptable with review.

9) FOIA/open-records requests from public-sector clients should be:

- A. Answered directly by staff
- B. Ignored unless a subpoena arrives
- C. Routed to Legal and handled via approved records systems
- D. Posted to the company wiki

Answer: C

Rationale: Public-sector clients are custodians; route through Legal.

10) For external-facing AI features, engineering must:

- A. Launch first, test later
- B. Complete a risk review, prepare a model factsheet, run red-team/bias tests, and set monitoring
- C. Let sales decide
- D. Fine-tune by default

Answer: B

Rationale: Governance and testing are required before launch.

11) Prompt hygiene means:

- A. Using the longest prompt possible
- B. Minimizing data, redacting sensitive elements, and avoiding secrets
- C. Asking the AI to ignore rules
- D. Pasting screenshots of everything

Answer: B

Rationale: Use minimum-necessary, redacted inputs and avoid secrets.

12) Who must complete this course and acknowledge the policy?

- A. Only engineers
- B. Only managers
- C. All employees, contractors, and vendors with system access
- D. New hires only

Answer: C

Rationale: Everyone with access must complete training and acknowledge the policy.