*800-110: CISCO AI TECHNICAL PRACTITIONER (AITECH) EXAM QUESTIONS*

Q1: Which pair of risks is most likely to increase as agent autonomy grows without sufficient controls?

A. Unintended actions and loss of accountability

B. Reduced scalability and guaranteed accuracy

C. Simplified governance and guaranteed accuracy

D. Reduced scalability and simplified governance

ANSWER: A

Q2: A company wants an AI agent to pause before taking any action that changes production data. Which strategy best addresses this requirement?

A. Human-in-the-loop approval

B. Unlimited autonomous execution

C. Removal of escalation logic

D. Prompt shortening only

ANSWER: A

Q3: A team wants an agent to retrieve context from tools through a consistent interface rather than building separate custom logic for each application. Which concept best aligns with that goal?

A. Model Context Protocol primitives

B. Prompt token expansion

C. Diffusion sampling

D. Static report rendering

ANSWER: A

Q4: Which option correctly matches AI capabilities to the SDLC phases where they provide the most direct value?

A. Requirements clarification → Testing; Code generation → Deployment; Automated testing → Requirements; Deployment validation → Implementation

B. Requirements clarification → Requirements; Code generation → Implementation; Automated testing → Testing; Deployment validation → Deployment

C. Requirements clarification → Deployment; Code generation → Requirements; Automated testing → Implementation; Deployment validation → Testing

D. Requirements clarification → Implementation; Code generation → Testing; Automated testing → Deployment; Deployment validation → Requirements

ANSWER: B

Q5: An engineer wants to use AI to improve the readability of a complex function without changing its behavior. Which task is the best fit?

A. Refactoring suggestions and code explanation

B. Replacing all logic with placeholder comments

C. Deleting validation checks to shorten the file

D. Converting the source file into an image

ANSWER: A

Q6: Misinformation generated by AI systems poses reputational and operational risks. This risk increases when outputs are shared without verification. Which mitigation strategy best addresses this concern?

A. Autonomous deployment without review

B. Human-in-the-loop validation

C. Larger context windows

D. Reduced logging

ANSWER: B

Q7: Which prompt change is most likely to reduce hallucinated answers when the model is asked about a policy document?

A. Tell the model to answer confidently even when unsure

B. Instruct the model to rely only on provided source material and indicate uncertainty when evidence is missing

C. Remove all source text to encourage generalization

D. Ask for the fastest answer regardless of accuracy

ANSWER: B

Q8: A customer-support agent receives a request, retrieves account details, checks entitlement status, and then decides whether to escalate or resolve the issue. What does this most clearly illustrate?

A. Goal-directed decision flow in an agent

B. Pure image-generation behavior

C. Manual-only processing

D. Static file compression

ANSWER: A

Q9: Why are AI model repositories valuable to practitioners?

A. They eliminate prompt engineering

B. They centralize vetted models for reuse

C. They guarantee regulatory compliance

D. They reduce inference latency

ANSWER: B

Q10: Which option correctly matches prompting techniques to their typical use cases?

A. Zero-shot → multi-step orchestration; Few-shot → unrestricted creativity; Iterative → security hardening; Chained → no examples needed

B. Zero-shot → tasks where examples are unnecessary; Few-shot → tasks requiring demonstration of output format; Iterative → tasks improved through repeated refinement; Chained → tasks broken into sequential steps

C. Zero-shot → tasks requiring demonstration of output format; Few-shot → repeated refinement; Iterative → no examples needed; Chained → single-step summarization

D. Zero-shot → repeated refinement; Few-shot → no examples needed; Iterative → output-format demonstration; Chained → random exploration

ANSWER: B

Q11: Which factor most directly increases inference latency by increasing the amount of text the model must process?

A. Context window size

B. Model licensing type

C. Embedding dimensionality

D. Prompt temperature

ANSWER: A

Q12: A prototyping team notices that responses are degrading because the prompt includes repeated copies of the same code and prior outputs. What is the best corrective action?

A. Remove redundant context and keep only relevant information

B. Add more duplicate content for stability

C. Disable all output validation

D. Replace the prompt with unrelated logs

ANSWER: A

Q13: An AI agent plans a workflow, calls a pricing service, reformats the result, and then updates a sales quote. Which capability is most clearly being exercised across the workflow?

A. Orchestration across tools and task steps

B. Offline-only image rendering

C. Manual spreadsheet reconciliation only

D. Context removal before reasoning

ANSWER: A

Q14: A business deploys agents that can automatically perform actions impacting customers. Governance teams require review before execution in sensitive cases. Which strategy best satisfies this requirement?

A. Fully autonomous execution

B. Human-in-the-Loop approval gates

C. Increased temperature settings

D. Removal of logging

ANSWER: B

Q15: Why is standardized tool interaction important in agentic AI environments?