



Microsoft GH-600

GitHub Agentic AI Developer

Thank You for Downloading GH-600 Updated Exam Questions

<https://certs4sale.com/microsoft/gh-600-pdf-exam-dumps>

MICROSOFT CERTIFICATION

Microsoft GH-600

Developing in Agentic AI Systems (Agentic AI Developer)

Vendor:	Microsoft
Exam Code:	GH-600
Document Type:	Practice Questions & Answers
Content Standard:	Verified Guide

Microsoft GH-600

PRACTICE EXAMINATION QUESTIONS

Question 1

SINGLE CHOICE

Your team is building an agentic AI system to automate code reviews in a CI/CD pipeline. The agent must identify potential bugs, suggest refactoring, and create pull request comments—but you need to prevent it from merging code without explicit human approval. You also want the agent to operate efficiently without requiring review for every comment it posts.

Which combination of configurations best balances autonomy with the compliance requirement?

A. Set autonomy level to full autonomous for all actions; implement post-hoc audit logging to track all decisions

B. Classify comment posting as low-risk (high autonomy), classify merging as high-risk (require explicit approval); configure the agent to block merge actions until human authorization is provided

C. Require human approval for every agent action before execution to ensure no code is modified without review

D. Configure the agent to suggest comments but disable PR creation entirely; require manual human intervention for all output review

Explanation:

The correct answer classifies agent actions by risk level and applies appropriate guardrails. Posting comments is low-risk and non-irreversible, so it can proceed autonomously to maximize delivery speed. Merging is high-risk and irreversible, so it requires explicit human authorization. This approach maximizes velocity while maintaining compliance—preserving execution velocity by minimizing approvals that do not materially reduce risk.

Why the others are wrong:

- **Option A:** Full autonomy for merging violates the stated compliance requirement and creates audit and rollback risk.
- **Option C:** Requiring approval for every action (including low-risk comments) slows delivery unnecessarily and contradicts the principle of risk-proportionate autonomy.
- **Option D:** Disabling PR creation entirely and requiring manual review for all output defeats the purpose of agentic automation and is overly restrictive.

Microsoft GH-600

PRACTICE EXAMINATION QUESTIONS

Question 2

MULTIPLE CHOICE

You are designing a multi-agent workflow where Agent A (code generation), Agent B (testing), and Agent C (documentation) work in parallel to complete a feature request. Both Agent A and Agent C have attempted to modify the same file, creating conflicting changes. Additionally, Agent B's test suite is stalled waiting for Agent A's output.

Which two of the following actions should you take to detect and resolve these multi-agent failures?
(Choose 2)

A. Implement agent isolation for parallel execution with scoped repository access so each agent operates on distinct file paths or branches

B. Increase the autonomy level of Agent A to override conflicting changes and force all agents to accept its output as canonical

C. Configure observability to detect stalled agent executions through logs and operational signals; implement a multi-agent recovery pattern such as rollback or requeue

D. Disable parallel execution entirely and serialize all agents into a linear pipeline to guarantee no conflicts occur

E. Assign all three agents identical tool permissions and memory scope so they can freely synchronize state without coordination overhead

Explanation:

The two correct answers address the root causes of multi-agent failures:

- **Option A (Correct):** Agent isolation with scoped access (e.g., branch-based scope, distinct file ownership) prevents overlapping code changes and conflicting modifications at the root. This is a proactive control matching the exam objective to detect and resolve agent conflicts.
- **Option C (Correct):** Observability (logs, traces, operational signals) detects stalled or failed executions; recovery patterns (rollback, human-in-the-loop, requeue) resolve them.

Why the others are wrong:

- **Option B:** Increasing autonomy to override conflicts masks symptoms without resolving underlying state conflicts.
- **Option D:** Serialization defeats the primary goal of parallel execution efficiency; isolation is the preferred architectural pattern.
- **Option E:** Identical permissions and shared memory increase the likelihood of state corruption and race conditions.

Microsoft GH-600

PRACTICE EXAMINATION QUESTIONS

Question 3

SINGLE CHOICE

Your team has deployed an agent that performs long-running code generation tasks across a codebase. After 6 hours of execution, the agent begins making decisions that contradict its earlier work, creating cascading errors. You suspect context drift—the agent's understanding of the task and prior decisions has diverged from the actual state of the codebase and its own earlier outputs.

Which combination of strategies should you implement to detect and prevent context drift in extended agent execution?

A. Capture task progress and prior decisions as durable artifacts (e.g., decision logs, state checkpoints); periodically validate agent understanding against these artifacts to detect divergence; implement a reset or rewind mechanism when drift is detected

B. Increase the agent's long-term memory size to retain all historical context indefinitely without pruning or expiration

C. Disable the agent's ability to revise prior decisions to ensure consistency, even if those decisions become obsolete due to changing requirements

D. Configure the agent to re-scan the entire codebase at every decision point to refresh its context dynamically

E. Resume the agent without resetting state and allow it to continue from where it left off, trusting that later decisions will correct earlier mistakes

Explanation:

The correct answer implements the three critical components for managing context drift:

- **Durable artifacts:** Decision logs and state checkpoints create an authoritative record to verify agent reasoning.
- **Validation:** Periodically comparing current understanding against artifacts identifies divergence early.
- **Recovery:** Reset or rewind mechanisms allow the agent to revert to a known-good state.

Microsoft GH-600

PRACTICE EXAMINATION QUESTIONS

Question 4

MULTIPLE CHOICE

You are configuring an agent to use the GitHub remote MCP server to automate pull request creation and branch management in your CI/CD pipeline. Before deploying to production, you need to ensure the agent has appropriate tool permissions and that the MCP server integration is secure and compliant.

Which two of the following configurations should you implement? (Choose 2)

A. Configure MCP allow lists to restrict the agent's access to only the GitHub operations required for your workflow (e.g., create branch, push commits, create PR), and review the MCP registries to verify the server version and integrity

B. Grant the agent full administrative access to all GitHub operations to maximize flexibility and avoid permission errors during execution

C. Configure the agent's scope to the target repository to prevent unintended actions against other repositories in your organization

D. Disable MCP server authentication to reduce latency and simplify configuration

E. Allow the agent to dynamically discover and load MCP servers at runtime without pre-approval or registry validation

Explanation:

The two correct options enforce the principle of least privilege and secure integration:

- **Option A (Correct):** MCP allow lists restrict available tools exclusively to required actions, and registry verification prevents compromised tool integration.
- **Option C (Correct):** Scoping access strictly to target repositories prevents cross-repository pollution or unauthorized operations.

Microsoft GH-600

PRACTICE EXAMINATION QUESTIONS

Question 5

SINGLE CHOICE

Your organization has deployed an agent that automatically refactors legacy code and creates pull requests. After analyzing execution logs and artifacts from 50 test runs, you identify that the agent succeeds 92% of the time but fails in a specific pattern: when the target file contains conditional compilation directives, the agent misunderstands the code structure and produces syntactically invalid refactorings.

Which of the following root cause classifications and tuning approaches is most appropriate?

A. The root cause is a tool misuse error; revise the agent's tool permissions to prevent access to files with conditional directives

B. The root cause is a reasoning error; revise the agent's instructions to explicitly address conditional compilation, include examples in the prompt, or refine the code analysis tool to detect and preserve these directives

C. The root cause is a context or environment issue; the CI environment is misconfigured and must be rebuilt

D. The root cause is an evaluation signal misalignment; increase the success threshold to 95% so fewer runs are flagged as failures

Explanation:

The correct answer identifies a reasoning error. The agent possesses the correct execution tools, but its cognitive parsing of conditional directives fails without explicit guidance or tuned prompt instructions. Adding target examples and refining instructions directly addresses reasoning deficits.

Microsoft GH-600

ANSWER KEY & DOMAIN ALIGNMENT

Quick Answer Key

Question #	Topic Domain	Correct Answer
Question 1	Autonomy & Guardrails	Option B
Question 2	Multi-Agent Orchestration & Recovery	Options A, C
Question 3	Context Drift & Long Execution Memory	Option A
Question 4	MCP Server Security & Permissions	Options A, C
Question 5	Root Cause Classification & Tuning	Option B

Exam Focus Summary

This practice set covers core Agentic AI Development topics including multi-agent failure resolution, context drift mitigation via state checkpoints, safe Model Context Protocol (MCP) integrations, and root cause classification for iterative prompt/tool tuning.

Microsoft GH-600

COMPLETION NOTICE

End of Practice Guide

You have completed the sample practice module for Microsoft GH-600 Developing in Agentic AI Systems. Use these key architectural insights to prepare for the full certification examination.

Verified Study Guide

THANK YOU FOR DOWNLOADING GH-600 UPDATED EXAM QUESTIONS

Note: Thanks For Trying The Demo Of Our GH-600 Exam Product

Visit Our Site to Purchase the Full Set of Actual GH-600 Exam Questions With Answers.

Money Back Guarantee



Click The Link Below

<https://certs4sale.com/microsoft/gh-600-pdf-exam-dumps>