



CompTIA AgentOps+ Certification Exam Objectives

EXAM NUMBER: AG0-001 V1

About the Exam

The CompTIA AgentOps+ AG0-001 V1 certification exam will certify the successful candidate has the knowledge and skills required to:

- Architect, deploy, configure, and orchestrate AI agent systems, including harnesses and agentic workflows across enterprise environments.
- Integrate AI agents with enterprise tools.
- Implement AI agents with retrieval-augmented generation (RAG) pipelines, knowledge graphs, and data structures.
- Manage prompts, memory, context, and agent configuration.
- Implement governance, security, privacy, and compliance controls to ensure AI agents operate safely, responsibly, and within organizational policies and guardrails.
- Implement observability, monitoring, cost measurement, and performance measurement capabilities to validate agent behavior, reliability, and operational effectiveness.
- Manage and optimize the operational and financial performance of AI agent systems, including resource utilization, scalability, and cost.
- Troubleshoot deployment, integration, behavioral, and performance issues affecting AI agents and the systems with which they interact.
- Operate AI agent platforms throughout the software life cycle by implementing continuous integration/continuous delivery/deployment (CI/CD) processes, continuous evaluation, incident response procedures, and ongoing operational management practices.

The target audience consists of DevOps Engineer, Cloud Engineer (or Cloud+ certified), Systems Administrator (or Linux+ certified), Network Engineer, or General IT Administration professionals who have at least 3–4 years of experience, some programming/scripting background experience and 6–12 months of exposure to artificial intelligence/machine learning (AI/ML) concepts.

These content examples are meant to clarify the exam objectives and should not be construed as a comprehensive listing of all the content of this examination.

EXAM ACCREDITATION

TBD.

EXAM DEVELOPMENT

CompTIA exams result from subject matter expert workshops and industry-wide survey results regarding the skills and knowledge required of an IT professional.

CompTIA AUTHORIZED MATERIALS USE POLICY

CompTIA Certifications, LLC is not affiliated with and does not authorize, endorse, or condone utilizing any content provided by unauthorized third-party training sites (aka “brain dumps”). Individuals who utilize such materials in preparation for any CompTIA examination will have their certifications revoked and be suspended from future testing in accordance with the CompTIA Candidate Agreement. In an effort to more clearly communicate CompTIA’s exam policies on use of unauthorized study materials, CompTIA directs all certification candidates to the [CompTIA Certification Exam Policies](#). Please review all CompTIA policies before beginning the study process for any CompTIA exam. Candidates will be required to abide by the [CompTIA Candidate Agreement](#). If a candidate has a question as to whether study materials are considered unauthorized (aka “brain dumps”), they should contact CompTIA at examsecurity@comp-tia.org to confirm.

PLEASE NOTE

The lists of examples provided in bulleted format are not exhaustive lists. Other examples of technologies, processes, or tasks pertaining to each objective may also be included on the exam, although not listed or covered in this objectives document. CompTIA is constantly reviewing the content of our exams and updating test questions to be sure our exams are current, and the security of the questions is protected. When necessary, we will publish updated exams based on existing exam objectives. Please know that all related exam preparation materials will still be valid.

DRAFT

TEST DETAILS

Required exam	AgentOps+ AG0-001 V1
Number of questions	TBD
Length of test	TBD
Passing score	TBD
Types of questions	Multiple-choice and performance-based
Recommended experience	DevOps Engineer, Cloud Engineer (or Cloud+ certified), Systems Administrator (or Linux+ certified), Network Engineer, or General IT Administration professionals who have at least 3–4years of experience, some programming/scripting background experience, and 6–12 months of exposure to artificial intelligence/machine learning (AI/ML) concepts.

EXAM OBJECTIVES (DOMAINS)

The table below lists the domains measured by this examination and the extent to which they are represented.

DOMAIN		PERCENTAGE OF EXAMINATION
1.0	AI Agent Architecture, Deployment, and Configuration	TBD
2.0	Agent Governance, Security, and Compliance	TBD
3.0	Agent Observability, Performance, and Cost Management	TBD
4.0	Agent Operations	TBD
Total		100%

1.0 AI Agent Architecture, Deployment, and Configuration

1.1 Explain architectural components and design patterns for artificial intelligence (AI) agent systems.

- Control plane vs. data plane
 - Hosting and inference models
 - Retrieval architecture
 - Non-deterministic agents vs. traditional services
- Agent architecture types
 - Single-agent systems for task-specific automation
 - Multiagent systems
 - ◆ Collaborative
 - ◆ Hierarchical
 - ◆ Swarms
- Agentic workflows and reasoning patterns
 - ReAct
 - Chain-of-Thought (CoT)
 - Tree-of-Thought (ToT)
 - State machine
- Core components
 - Large language model (LLM) providers and model selection
 - Model routing logic and fallback strategies
 - Tool integration layers
- Memory management
 - Short-term
 - Mid-term
 - Long-term
- Application programming interface (API)-based vs. self-hosted
 - Air-gapped self-hosted systems
- Autonomy levels and promotion criteria
 - Assisted operations, in which a human confirms actions
 - Semi-autonomous operations, in which a human approves critical actions
 - Fully autonomous within predefined guardrails
 - Action-risk tiering framework
 - ◆ Read vs. write vs. destructive operations

1.2 Given a scenario, configure AI agent deployment and orchestration workflows.

- Workflow orchestration
 - Agent-to-agent (A2A) coordination and handoff logic
 - Multiagent-based context flow
 - Task delegation and workflow routing
- Cloud vs. on-premises deployment
 - Public
 - Private
- Deployment types
 - Containerized agent deployment
 - Serverless agent deployment
 - Container orchestration and resource allocation
- Autoscaling configurations

1.3 Given a scenario, implement tool integrations for AI agents.

- Tool integration
 - Model Context Protocol (MCP) server installation and configuration
 - OpenAPI-based tool definitions
 - Tool authentication and authorization
 - Tool versioning and deprecation strategies
- Enterprise integrations
 - Database and vector store connectors
 - API authentication methods
 - Filesystem and document access
 - MCP gateways
 - Knowledge graph
 - Feature flag
- Context and memory configuration
 - Session state persistence
 - Context window management
 - Knowledge base

1.4 Given a scenario, configure and manage agent prompt templates, instructions, and configuration systems.

- Prompt and instruction management
 - System instruction configuration
 - Versioning
 - Prompt and scenario routing
- Agent configuration management
 - Declarative agent configuration
 - Controlled configuration updates
 - Configuration service integration
 - Template-based agent creation workflows
 - Environment-specific configurations
- Response control
 - Output format specification and structured outputs
 - Task-specific instruction refinement
 - Response validation rules

1.5 Given a scenario, implement knowledge bases and retrieval-augmented generation (RAG) pipelines.

- Knowledge base operations
 - Content freshness monitoring and updates
 - Data ingestion and preprocessing pipelines
 - Metadata tagging and filtering
 - Permission-aware retrieval and tenant data isolation
- Vector index management
 - Vector index creation and maintenance
 - Bulk operations
 - Search quality tuning
- RAG pipeline configuration and retrieval architectures
 - RAG stages
 - ◆ Ingestion
 - ◆ Chunking
 - ◆ Embedding
 - ◆ Retrieval
 - ◆ Reranking
 - Vector similarity search and index operations
 - Grounding vs. hallucination as a quality concern
 - ◆ Positive fixture
 - ◆ Negative fixture
 - Retrieval scoring and ranking
 - Hybrid search strategies
 - ◆ Vectorless
 - ◆ Agentic RAG
 - ◆ Graph RAG
 - Context assembly and relevance filtering
 - Citation and evidence tracking

2.0 Agent Governance, Security, and Compliance

2.1 Given a scenario, implement governance policies for AI agent operations.

- Policy frameworks
 - Policy as code (PaC)
 - Pre-execution validation of tool calls
 - Policy hooks that trigger compliance checks
 - Policy violation handling and escalation
- Runtime access governance
 - Privileged access control for agent and operator accounts
 - Periodic access recertification
 - Network-based controls
 - Data residency-based controls
 - Time-based controls
 - ◆ Just-in-time access
 - Access audit procedures
 - Recording and log access controls
 - Segregation of duties enforcement
- Human oversight and accountability
 - Clear ownership assignment for each agent
 - ◆ Owner
 - ◆ Risk tier
 - Human-in-the-loop policy
 - Change management

2.2 Given a scenario, implement guardrails for AI agent operations.

- Approved tool catalogs and allow lists
- Parameter validation and sanitization
- Scope limitations
- Guardrail enforcement points
- Input classification and output filtering
- Input and output guardrails
 - Human-in-the-loop checkpoint configuration in guardrail policies
- Emergency stop and override mechanisms
- Execution guardrails

2.3 Given a scenario, configure security controls for AI agent authentication, authorization, and secrets management.

- Agent identity and access management (IAM)
 - Service accounts
 - Role-based access control (RBAC)
 - Least privilege enforcement
 - Secure Production Identity Framework for Everyone (SPIFFE)/SPIFFE Runtime Environment (SPIRE)
- Secrets and credentials
 - Secrets management integration
 - API key management
 - Environment-based secrets isolation
- Agent-to-tool authorization
 - Open Authorization (OAuth) 2.0 token acquisition and refresh
 - Service principal authentication
 - Scoped permissions per tool
 - Certificate management
- MCP server security
 - Scoping permissions
 - Limiting exposed tools
 - Isolating credentials

2.4 Given a scenario, implement data governance and privacy controls for AI agents.

- Data protection control and compliance
 - Sensitive data detection and redaction in agent inputs/outputs
 - ◆ Personally identifiable information (PII), protected health information, and payment card information
 - Data minimization in prompts and memory
 - Data retention and purging policies
- Trust boundary enforcement
 - Controls preventing cross-boundary data leakage
 - Ingestion validation and filtering
 - Output sanitization
- Audit and compliance evidence
 - Recording and logging validation
 - Compliance evidence generation for regulatory review
 - Reference frameworks
 - ◆ Open Worldwide Application Security Project OWASP LLM Top 10
 - ◆ OWASP Agentic Top 10
 - ◆ MITRE Adversarial Threat Landscape for Artificial-Intelligence Systems (ATLAS)
- Transparency and auditing
 - Agent-level and model card documentation for production models
 - Decision and action audit logging (provenance)
 - Audit trail requirements

3.0 Agent Observability, Performance, and Cost Management

3.1 Given a scenario, configure observability instrumentation for AI agent systems.

- Agent-specific logging
 - Structured logging
 - Prompt and completion logging
 - Log hygiene
 - ◆ Metadata
 - ◆ Redacting sensitive data
 - ◆ Life cycle
 - ◆ Integrity
 - Tool call invocation logs
 - Correlation IDs for end-to-end tracing
- Tracing
 - Distributed tracing instrumentation
 - ◆ OpenTelemetry
 - Span creation for each reasoning step
 - Tool execution tracing
 - LLM API call tracing
 - Trace inspection and analysis
- Metrics collection
 - Token usage
 - Latency metrics
 - Cost per invocation
 - Success/failure rates
 - Concurrent agent sessions
 - Queue depth and backlog metrics
 - Anomalies
 - Resource utilization

3.2 Given a scenario, configure alerting and monitoring for agent execution, output validation, and behavioral anomalies.

- Execution monitoring
 - Agent workflow integration health
 - Tool invocation validation
 - Guardrail validation metrics
 - Failed session tracking
 - Workflow and routing status
- Output validation
 - Tool call performance
 - Structured output validation
 - Business rule validation
- Behavioral analysis
 - Tool usage patterns and anomalies
 - Context window utilization
 - Unexpected tool invocation detection
- Alert configuration
 - Threshold-based alerting
 - ◆ Task failure rate increases
 - ◆ Tool authorization failures
 - ◆ Policy violation detections
 - Anomaly detection using baselines
 - ◆ Token consumption anomalies
 - Rate-of-change alerts
 - Composite condition triggers
 - Notification and escalation
 - ◆ Incident management integration
 - ◆ Chat/webhook notifications
 - ◆ Escalation policies by severity
- Service level agreement (SLA)/service-level objective (SLO)/service-level indicator (SLI) and budgets
 - Availability targets
 - Latency targets
 - Cost threshold breaches
- Confidence scoring

3.3 Given a scenario, implement cost monitoring and optimization.

- Cost tracking
 - Token consumption metrics by agent, team, and project*
 - LLM API spend monitoring
 - Compute cost allocation
 - Usage pattern analysis
- Cost governance
 - Token budget limits
 - LLM API cost ceilings and circuit breakers
 - Per-tenant/feature cost attribution
 - Budget alerting thresholds
 - Quota monitoring
- Cost optimization
 - Model selection
 - ◆ Agent function
 - ◆ Trade-offs
 - Semantic caching
 - Token usage optimization
 - Model right-sizing
- Resource allocation and tagging

**CompTIA is assuming that the target audience for the test already understands what tokens are, token basics, token usage, etc.*

3.4 Given a scenario, troubleshoot AI agent deployment and integration issues.

- Common deployment issues
 - Model API connectivity failures
 - Authentication and authorization errors
 - Resource constraints
 - Tool integration failures
 - State persistence errors
- Debugging techniques
 - Log analysis
 - Network connectivity testing
 - API endpoint validation
 - Configuration verification
 - Dependency health checks
- Resolution strategies
 - Rollback to previous stable version
 - Configuration drift remediation
 - Resource limit adjustments
 - Tool endpoint troubleshooting

3.5 Given a scenario, analyze and troubleshoot AI agent behavioral and performance issues.

- Common failure modes
 - Hallucinated tool calls
 - Infinite reasoning loops and timeouts
 - Incorrect tool selection
 - Parameter formatting errors
 - Stale or incorrect memory retrieval
 - Task-verification failures
 - Intent drift
 - Single-agent failure modes
 - ◆ Loops
 - ◆ Timeouts
 - ◆ Context overflow
 - Multiagent failure modes
 - ◆ Stalled subagents
 - ◆ Handoff failure
- Performance indicators
 - Task completion rate and success rate
 - Trajectory efficiency (tool calls vs. minimum required)
 - Response time and latency percentiles
 - Throughput
 - Error rates and retry frequency
 - Availability and uptime
 - Escalation to human frequency
- Diagnostic techniques
 - Trace analysis for decision paths
 - Prompt inspection and validation
 - Tool call replay in isolation
 - State inspection for stateful agents
 - End-to-end pipeline testing
- Remediation and mitigation strategies
 - Prompt refinement
 - Tool description and output schema improvements
 - Model selection adjustments
 - Context management
 - Retry, fallback, and error-handling implementation
 - Resource allocation adjustments
 - Policy changes

4.0 Agent Operations

4.1 Given a scenario, implement and operate continuous integration/continuous deployment/delivery (CI/CD) pipelines for agent systems.

- Agent artifact and versioning
 - Agent as a single versioned and deployable unit
 - ◆ Model configuration
 - ◆ Prompts
 - ◆ Tool and MCP definitions
 - ◆ Policies
 - ◆ Memory configuration
 - Source control for agent artifacts
 - Build and packaging procedures
- Automated pipeline gates
 - Tool unit and integration tests
 - Evaluation and regression gates
 - Adversarial and safety gates
 - Vulnerability scanning and dependency checks
- Deployment and promotion
 - Environment promotion
 - ◆ Development
 - ◆ Staging/sandbox
 - ◆ Production
 - Deployment automation
 - ◆ GitOps
 - ◆ Infrastructure as code (IaC)
 - Gradual rollout strategies
 - ◆ Canary
 - ◆ Blue-green
 - Post-deployment verification

4.2 Given a scenario, manage agent connectivity and rate controls.

- Network controls
 - LLM and tool API connectivity requirements
 - Agent-to-tool network policies
 - Private endpoints for sensitive integrations
 - Zero Trust principles
 - Agent identification and certificate management
- API gateway and resilience controls
 - Rate limiting and throttling per agent
 - Circuit breakers for fault tolerance
 - Retry policies

4.3 Given a scenario, operate continuous evaluation, testing, and validation.

- Evalsuites and gates
 - Smoke and regression evaluation suites
 - Golden dataset validation
 - ◆ Regression testing
 - Sanitization testing
 - Regression gates that block release on quality regression
 - Evaluation frameworks (eval frameworks)
 - Human-in-the-loop evaluation sampling
- Drift signals to monitor
 - Input distribution drift
 - ◆ User query patterns
 - Output pattern changes
 - ◆ Response structure
 - ◆ Tone
 - Tool-usage pattern shifts
 - Guardrail intervention-rate trends and grounding-score trends
 - Performance and quality degradation
- Review cadence
 - Automated evaluation using drift dashboards
 - Periodic (weekly/monthly) performance reviews
 - User feedback aggregation
- Security testing
 - Prompt injection attack simulation
 - Jailbreak attempt testing
 - Tool manipulation testing
 - Authorization bypass attempts
 - Data exfiltration scenarios
 - Memory/context poisoning that persists across sessions
- Edge case testing
 - Unexpected input handling
 - Boundary condition validation
 - Malformed request processing
 - Resource exhaustion scenarios
 - Fallback behavior verification
 - Sandbox isolation testing
 - A/B testing

4.4 Given a scenario, perform incident response for agent misbehavior and outages.

- Incident detection and triage
 - Agent misbehavior identification
 - Harmful output detection
 - Unauthorized action detection
 - Data leakage incidents
 - Service degradation or outage detection
- Response procedures
 - Activate kill switch during live incident
 - Agent suspension or isolation
 - Credential revocation
 - Workflow bypass and workaround management
- Root cause analysis
 - Failed interaction/workflow review
 - Broken automation diagnosis
 - Integration failure analysis
 - Post-incident reporting
- Recovery and continuity
 - Service restoration procedures
 - Failover testing and alternate routing
 - Backup routing activation
 - Disaster recovery playbook
 - Recovery time objective/recovery point objective (RTO/RPO)
- Multiagent failure modes
 - Cascading failures across agents
 - Interagent trust escalation (confused deputy problem)
 - Infinite delegation loops
 - Session-context contamination

CompTIA AgentOps+ Acronym List

The following is a list of acronyms that appear on the CompTIA AgentOps+ AG0-001 V1 certification exam. Candidates are encouraged to review the complete list and attain a working knowledge of all listed acronyms as part of a comprehensive exam preparation program.

ACRONYM	DEFINITION
A2A	Agent-to-agent
AI	artificial intelligence
API	application programming interface
ATLAS	Adversarial Threat Landscape for Artificial-Intelligence Systems
CI/CD	Continuous integration/continuous deployment/delivery
CoT	Chain-of-Thought
GPU	graphics processing unit
IaC	infrastructure as code
IAM	identity and access management
IDE	integrated developer environment
LLM	large language model
MCP	Model Context Protocol
OAuth	Open Authorization
OWASP	Open Worldwide Application Security Project
PaC	policy as code
RAG	retrieval-augmented generation
RBAC	role-based access control
REST	Representational State Transfer
RPO	recovery point objective
RTO	recovery time objective
SLA	service-level agreement
SLI	service-level indicator
SLO	service-level objective
SPIFFE	Secure Production Identity Framework for Everyone
SPIRE	SPIFFE [Secure Production Identity Framework for Everyone] Runtime Environment
ToT	Tree-of-Thought

CompTIA AgentOps+ Hardware and Software List

CompTIA has included this sample list of hardware and software to assist candidates as they prepare for the AgentOps+ AGO-001 V1 certification exam. This list may also be helpful for training companies that wish to create a lab component for their training offering. The bulleted lists below each topic are sample lists and are not exhaustive.

EQUIPMENT

- Laptop
- Internet access
- Graphics processing unit (GPU) (optional)

SOFTWARE

- Python environment
- Container orchestration platform
- Source control/git
- Integrated developer environment (IDE)
- Model server
- Agent orchestrator/framework
- Representational State Transfer (REST) tools
- Secrets manager
- Monitoring or observability stack
- Databases
- IAC

OTHER

- Cloud account
- Sandbox account
- LLM provider