CertNexus® Agentic AI Builder™ Exam AGB-110 Blueprint

Date Issued: 06/01/2026
Version: 1.0

Agentic AI BUILDER™

Introduction to CertNexus

CertNexus is a vendor-neutral certification body, providing emerging technology certifications and micro-credentials for business, data, developer, IT, and security professionals. CertNexus' mission is to assist closing the emerging tech global skills gap while providing individuals with a path towards rewarding careers in Cybersecurity, Data Science, Data Ethics, Internet of Things, and Artificial Intelligence (AI)/ Machine Learning (ML).

We rely on our Subject Matter Experts (SMEs) to provide their industry expertise and help us develop these credentials by participating in a Job Task Analysis, Exam Item Development, and determining the Cut Score. We also depend upon practitioners in the field to participate in a survey of the Job Task Analysis and beta testing to ensure that our certifications validate knowledge and skills relevant to the industry.

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved. CertNexus® Agentic AI Builder™ (AGB-110) Exam Blueprint
CertNexus® Agentic AI Builder™ (AGB-110) Exam objectives are subject to change without notice.

CertNexus® Agentic AI Builder™ Exam AGB-110

Exam Information

Candidate Eligibility

The Agentic AI Builder™ assessment requires no application fee, supporting documentation, or other eligibility verification measures for you to be eligible to take it. Simply purchase an access key for the *Agentic AI Builder™ (Exam AGB-110): Setting Up a Functional AI Agent* course from the CertNexus Store here. This course includes access to the credential process directly through the CHOICE platform.

Exam Prerequisites

Successful candidates should have the skills and knowledge to build an agentic AI system through code, including implementing a reasoning loop within a manageable state, creating tools for the agent to use, evaluating agent behavior, and more. It is recommended that candidates acquire these skills and knowledge by attending the CertNexus® *Agentic AI Builder™ (Exam AGB-110): Setting Up a Functional AI Agent* course prior to taking the assessment.

Exam Specifications

Number of Items: 25

Passing Score: 80% or 20/25 items

Duration: Estimated 30-60 minutes; candidates receive 1 complimentary retake.

Exam Delivery: Online through the CHOICE platform.

Item Formats: Multiple Choice/Multiple Response

Upon successful completion, candidates will earn the CertNexus Agentic AI Builder credential.

Exam Description

Target Audience:

The Agentic AI Builder assessment is primarily designed for programmers, IT infrastructure personnel, DevOps/MLOps personnel, and any other technically minded professional who is responsible or may soon be responsible for implementing agentic solutions in their organization.

Exam Objective:

Upon successful completion of the AGB-110 assessment, technical professionals will demonstrate they can build an autonomous, tool-using AI system from scratch that brings value to the organization while minimizing risks.

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved.

To ensure that candidates possess the aforementioned knowledge, skills, and abilities, the AGB-110 assessment will test them on the following domains with the following weightings:

Domain	% of Examination
1.0 Set Up the Agent Development Environment	12%
2.0 Assess LLM Behavior in Agent Contexts	12%
3.0 Implement the Agent Reasoning Loop	16%
4.0 Use Agent Tools	16%
5.0 Add Knowledge with RAG	12%
6.0 Enforce Structure, Safety, and Reliability in an Agent	12%
7.0 Test an Agent	12%
8.0 Deploy a Single-Agent System	8%
Total	**100%**

The information that follows is meant to help you prepare for your credential assessment. This information does not represent an exhaustive list of all the concepts and skills that you may be tested on during your assessment. The domains, identified previously and included in the objectives listing, represent the large content areas covered in the assessment. The objectives within those domains represent the specific tasks associated with the job role(s) being tested. The information beyond the domains and objectives is meant to provide examples of the types of concepts, tools, skills, and abilities that relate to the corresponding domains and objectives. All this information represents the industry-expert analysis of the job role(s) related to the credential and does not necessarily correlate one-to-one with the content covered in your training program or on your assessment. We strongly recommend that you study independently to familiarize yourself with any concept identified here that was not explicitly covered in your training program or products.

Domains and Objectives

Domain 1.0 Set Up the Agent Development Environment [12%]

Objective 1.1 Configure a Python-based agent workspace.

- Python runtime requirements
  - Configuration
  - Recent versions of language and libraries
  - Foundation of agent
- Virtual environments
  - `venv`
  - VMs
  - Containers
- Dependency isolation
- Project structure
  - Entry point/execution starter
  - Agent logic
  - Tools
  - Data/storage
- Libraries, frameworks, and tools
  - `requests`, `httpx`

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved.

        - LangChain, Haystack
        - API-based data retrieval, database queries

Objective 1.2 Configure LLM access.

- Cloud LLM APIs
    - External access
    - RESTful methods
    - Provider-specific libraries
- Authentication and authorization
    - Token-based authentication
    - API keys
    - Secure storage of secret material
- Data requests
    - Payload data
    - Structure/content requirements
    - Common parameters
        - Model
        - Input message
        - Max tokens
- Data responses
    - JSON format
    - Common fields
        - Status
        - Timestamp
        - Output
        - Reasoning
    - Parsing routines
- Local models
    - Interfaces
    - Accessibility
- Environment configuration
    - Separate from code
    - Storage methods
    - Common configurations
        - Model name
        - API key
        - API endpoint

Objective 1.3 Configure runtime constraints.

- Cost awareness
- Cost limits
    - Implementations
        - Request
        - Session
        - Global
    - Outcomes when limit is reached
    - Influences on agent behavior
- Token limits
    - Implementations
        - Model input
        - Model output

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved. CertNexus® Agentic AI Builder™ (AGB-110) Exam Blueprint
CertNexus® Agentic AI Builder™ (AGB-110) Exam objectives are subject to change without notice.

        - Combined total
    - Differences in input/output cost
    - "Hidden" costs of reasoning
- Performance tradeoffs
    - Cost vs. quality
    - Latency vs. reasoning depth
    - Context size vs. relevance
    - Output length vs. efficiency

Domain 2.0 Assess LLM Behavior in Agent Contexts [12%]

Objective 2.1 Analyze LLM capabilities and limitations.

- Reasoning models
    - Attributes
    - Evaluation
- Context windows
    - Size constraints
    - Impacts
    - Requirements
- Information compression
    - Strategies
    - Tradeoffs
- Determinism vs. variability
    - Effect on output consistency/diversity
    - Temperature
- Confabulation patterns

Objective 2.2 Design prompts for agent reasoning.

- Prompting roles
    - Platform/root
    - System/developer
    - User
    - Tool
- Chain of command
- System/developer prompts
    - Purpose
    - Guidance
- User prompts
    - Purpose
    - Queries
- Prompt separation
- In-context prompting
    - Zero shot
    - One shot
    - Few shot
- Prompt templates
- Ad hoc prompts
- Prompt framing
    - Roles
    - Goals
    - Constraints

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved. CertNexus® Agentic AI Builder™ (AGB-110) Exam Blueprint
CertNexus® Agentic AI Builder™ (AGB-110) Exam objectives are subject to change without notice.

Domain 3.0 Implement the Agent Reasoning Loop [16%]

Objective 3.1 Implement the ReAct pattern.

- ReAct cycles
  - Adaptation
  - Phases
    - Reason
    - Act
    - Observe
  - Feedback loop
  - Practical artifacts
    - Thought
    - Action
    - Observation
  - Separation of artifacts
  - Iterative decision making
  - Compared to traditional LLM approach
- Stopping conditions
  - Goal completion
  - Max iterations
  - Sufficiency/confidence thresholds
  - No-progress detection
- Loop control
  - Iteration management
  - Conditional branching
  - Retry/recovery logic
  - Execution throttling

Objective 3.2 Manage agent state across iterations.

- Stateless agent design
  - Benefits
  - Drawbacks
  - Applicability
- Stateful agent design
  - Benefits
  - Drawbacks
  - Applicability
- Code representation of agent state
  - Dictionary based
  - Class based
  - Type of schema based
- Error states
  - Missing/invalid data
  - Failed tool execution
  - Unexpected/inconsistent outputs
  - Inability for agent to proceed
- Recovery paths
  - Retry same action
  - Modify action
  - Select alternative tool/path

Objective 3.3 Manage memory and persistence.

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved.

- Conversation memory
  - User inputs
  - Agent responses
  - Roll aggregation
  - Context-window implications
- Working memory
  - Internal, task-specific information
  - External to conversation thread
- Memory size limits
- Memory decay
  - Effects
  - Strategies
    - Summarize
    - Filter
    - Condense
  - Tradeoffs
- Memory as input, not truth
- Persistent memory
  - Purpose and application
  - Examples
    - User preferences/profiles
    - Historical interactions/summaries
    - Previously computed results
  - Storage
    - Lightweight options
    - File types
    - Embedded databases
- Versioning
  - Individual records
  - Entire datasets
- Schema evolution

Domain 4.0 Use Agent Tools [16%]

Objective 4.1 Design agent tools and interfaces.

- What makes a good agent tool?
  - Clear/focused purpose
  - Predictable behavior
  - Well-defined I/O
  - Minimal ambiguity
- Tool inputs/outputs
  - Common types
  - Best practices
- Tool side effects
  - When they occur
    - Writing to DB
    - Sending email
    - Creating/updating records
  - Potential consequences
    - Cannot be undone

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved.

        - Unsafe
        - Triggered multiple times unintentionally
    - Methods for managing
- Idempotent vs. non-idempotent actions
    - Changes to environment
    - Duplication of effects
    - Mitigating effects of non-idempotent actions

Objective 4.2 Execute and validate tool calls.

- Function calling
    - Arguments
    - Avoid free-text requests
    - Structured invocation
- Asynchronous I/O
    - Applicability
        - External APIs
        - Network/storage operations
        - Multiple tool calls
    - Performance benefits
    - Complexity drawbacks
- Pre-execution validation
    - Consequences if not implemented
    - Approaches
        - Structural
        - Semantic
        - Authorization/scope
        - Safety/policy
    - Handling validation failures
- Guardrails
    - When to apply
        - Filtering/validating user input
        - During reasoning
        - During tool selection
    - Where to apply
        - System prompts
        - Programmatic design
- Dry-run vs. live-execution modes

Domain 5.0 Add Knowledge with RAG [12%]

Objective 5.1 Implement document ingestion and embeddings.

- RAG vs. standard tools
- Document sources
    - Internal documentation
    - External content
    - File-based sources
- Embeddings
    - Mathematical representation
    - Similarity
    - Distance in embedded space
    - Vectors

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved. CertNexus® Agentic AI Builder™ (AGB-110) Exam Blueprint
CertNexus® Agentic AI Builder™ (AGB-110) Exam objectives are subject to change without notice.

- Semantic search
- Transformer models
- Vector storage
- Chunking strategies
  - Fixed sized
  - Overlapping
  - Semantic
- Metadata strategies
  - Filtering
  - Ranking
  - Schema design

Objective 5.2 Retrieve and use context effectively.

- Context selection
  - Similarity score
  - Cosine similarity
  - Comparison between input vector and chunk vectors
  - Top-*k*
- Relevance filtering
  - Similarity thresholds
  - Metadata-based filtering
  - Heuristic filtering
- Query transformation
  - Expansion
  - Structuring
  - Simplification
- Over-contextualization
  - Causes
  - Mitigations

Domain 6.0 Enforce Structure, Safety, and Reliability in an Agent [12%]

Objective 6.1 Enforce structured outputs.

- Free-text outputs
  - Dangers
  - When/where to use natural language
- JSON schemas
  - Purpose
  - Relevance to agent output
- Type validation
  - Importance
  - Related issues
  - Consequences of poor validation
- Partial/invalid outputs
  - Missing required fields
  - Incomplete JSON
  - Incorrect data types

Objective 6.2 Handle uncertainty and failures.

- Confidence signaling
  - Certainty vs. uncertainty
  - Forms

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved. CertNexus® Agentic AI Builder™ (AGB-110) Exam Blueprint
CertNexus® Agentic AI Builder™ (AGB-110) Exam objectives are subject to change without notice.

    - Sources
- Abstention
    - Purpose
    - When to apply
    - Impacts
    - Effective design
- Retry logic
    - Issues
        - Invalid outputs
        - Validation failures
        - API issues
    - When to apply
        - Regenerate LLM response
        - Re-execute tool call
        - Reissue RAG query
    - Strategies
        - Retry as is
        - Modify retry
        - Guide retry
- Backoff strategies
    - Applications
        - Rate limiting
        - Temporary network failures
        - Service unavailability
    - Strategies
        - Fixed
        - Exponential
        - Jitter based
- Human-escalation triggers
    - Common conditions
    - Design principles

Domain 7.0 Test an Agent [12%]

Objective 7.1 Monitor agent behavior.

- Debugging difficulties
    - Probabilistic nature
    - Source of LLM-derived bugs
    - Propagation of mistakes
    - Reliance on external systems
- Execution traces
    - Purpose
    - Content
- Decision logs
    - Purpose
    - Content
- Token usage
    - Observable patterns
    - Repeated high usage
    - Sudden spikes

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved. CertNexus® Agentic AI Builder™ (AGB-110) Exam Blueprint
CertNexus® Agentic AI Builder™ (AGB-110) Exam objectives are subject to change without notice.

        - Overly verbose outputs
        - Per step vs. per session
    - Cost signals
        - As a behavioral indicator
        - Per step vs. per session
        - Obtaining/estimating cost
        - Contribution of external tools

Objective 7.2 Evaluate agent performance.

- Unit testing
    - Purpose
    - Application in agentic AI
        - Prompt-construction logic
        - Tool interfaces
        - Parsing/validation logic
- Scenario-based agent evaluation
    - When to use
    - Composition
        - User request/goal
        - Relevant context/constraints
        - Expected behaviors/outcomes
    - Definition of success/failure
- Regression risks
    - Signs
        - Scenario that once worked now fails
        - Worsening performance
    - Causes
        - Prompt modifications
        - Tool updates
        - Memory/context adjustments
    - Mitigations
        - Monitor drift
        - Compare against baselines
        - Test continuously

Objective 7.3 Analyze an agent for security flaws.

- OWASP ASI
    - Guidance
    - LLM Top Ten
- Threat modeling
    - Purpose
    - Process
    - OWASP threat model for AI agents
        - Threats
        - Mitigations
- Privilege management
    - Least privilege
    - Threats to agents
    - Mitigations
        - RBAC
        - Cryptographic authentication

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved. CertNexus® Agentic AI Builder™ (AGB-110) Exam Blueprint
CertNexus® Agentic AI Builder™ (AGB-110) Exam objectives are subject to change without notice.

        - Human validation
    - Secure tool execution
        - Misuse
        - Mitigations
    - Supply-chain concerns
        - Reliance on external resources
        - Mitigations
    - Prompt injection
        - Effects
        - Forms
            - Direct user input
            - Documentation used in RAG
            - Tool outputs
        - Tests

Domain 8.0 Deploy a Single-Agent System [8%]

Objective 8.1 Expose agent interfaces.

- CLI
    - Forms
    - Benefits and drawbacks
    - Suitability
        - Development and testing
        - Debugging workflows
- Notebook interfaces
    - Contents of cells
    - Sequential execution
    - Benefits and drawbacks
    - Suitability
        - Prototyping
        - Experimentation
- Web interfaces
    - User friendliness
    - Target audiences
    - Common elements
    - Client–server model
    - Interaction patterns
        - Chat-based interface
        - Form-based inputs
        - Event-driven actions
    - Benefits and drawbacks
    - Suitability
        - Production use for end users
        - Demonstration to business stakeholders
- UI as a debugging tool
- Separation of UI and agent logic
    - Improvements
    - Code modules for each concern
    - Point of integration

Objective 8.2 Review readiness and limitations.

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved. CertNexus® Agentic AI Builder™ (AGB-110) Exam Blueprint
CertNexus® Agentic AI Builder™ (AGB-110) Exam objectives are subject to change without notice.

- Production readiness
    - Reliability/stability
    - Operational readiness
    - Security/safety readiness
- Known limitations
    - Model
    - Tool/integration
    - Design
    - Operational
- Risk disclosure
    - What to disclose
    - Audience perspectives
- Multi-agent expansion
    - Identify need for expansion
    - Code modularity
    - Tools as reusable services

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved.

Agentic AI Builder Acronyms

Acronym	Expanded Form
AI	Artificial intelligence
ASI	Agentic Security Initiative
API	Application programming interface
JSON	JavaScript Object Notation
LLM	Large language model
OWASP	Open Worldwide Application Security Project
RAG	Retrieval-augmented generation
RBAC	Role-based access control
ReAct	Reason–act
REST	Representational State Transfer
SBOM	Software bill of materials
VM	Virtual machine

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved. CertNexus® Agentic AI Builder™ (AGB-110) Exam Blueprint
CertNexus® Agentic AI Builder™ (AGB-110) Exam objectives are subject to change without notice.

CertNexus offers personnel certifications and micro-credentials in a variety of emerging technology skills including Cybersecurity, Cyber Secure Coding, the Internet of Things (IoT), IoT Security, Data Science, Artificial Intelligence, and Data Ethics. For a complete list of our credentials visit https://certnexus.com/certification/.

CERTNEXUS®

1150 University Ave, Suite 20, Rochester, NY 14607
1-800-326-8724 | info@certnexus.com
certnexus.com

Copyright ©2026 by CertNexus Inc. (a division of Logical Operations). All rights reserved. CertNexus® Agentic AI Builder™ (AGB-110) Exam Blueprint
CertNexus® Agentic AI Builder™ (AGB-110) Exam objectives are subject to change without notice.