

The 22-page toolkit **preview.**

Every artifact above, indexed. Plus the 11-module syllabus, sample exam, the framework decision matrix, and resume-ready bullet templates.

Used by 2,50,000+ certified professionals.

12

Toolkit artifacts previewed

11

Module syllabi, verbatim

6

Framework decision matrix,
scored

INSIDE THE GUIDE

What's indexed in this toolkit.

A complete map of the certification's working artifacts — what you'll receive, build, and ship. Each item below appears in full, paginated form across the next 21 pages.

- ✓ **12 toolkit artifact previews** — agent blueprints, prompt scaffolds, eval rubrics, deployment checklists, observability dashboards, and more.
- ✓ **6 framework decision matrix · scored** — LangGraph, CrewAI, AutoGen, OpenAI Assistants, Semantic Kernel, and Smol Agents — benchmarked across 9 dimensions.
- ✓ **11 module syllabi · verbatim** — exact learning objectives, hours, capstone deliverables, and assessment weighting from the live curriculum.
- ✓ **Sample exam · printable checklist** — five representative multiple-choice items plus the full 47-point readiness checklist used by passing candidates.
- ✓ **Resume bullet templates · ATS-tested** — 8 fill-in-the-blank bullets calibrated for ATS keyword density and senior-recruiter scanning patterns.

"The matrix on page 11 alone saved our team three weeks of framework evaluation. We picked, shipped, and iterated."

— Cohort 47, Lead AI Engineer

12 toolkit artifact previews.

The first six. Each artifact ships as an editable template plus a worked example from a live production deployment.

ARTIFACT 01

Agent Blueprint Canvas

One-page spec: goal, tools, memory, guardrails, escalation paths. Used pre-build to align engineering and product.

ARTIFACT 02

Prompt Scaffold Library

14 reusable scaffolds: ReAct, Plan-Execute, Reflexion, Tree-of-Thought. Variable slots and worked outputs.

ARTIFACT 03

Eval Rubric Generator

Weighted scorecards for accuracy, grounding, latency, cost. Includes auto-grader prompt patterns.

ARTIFACT 04

Tool Registry Schema

JSON spec for declaring tools, args, idempotency, and side-effects. Framework-agnostic.

ARTIFACT 05

Memory Layer Decision Tree

Episodic vs. semantic vs. procedural — when to use which, with vector DB and key-value reference setups.

ARTIFACT 06

Guardrail Pattern Library

12 patterns covering input validation, output filtering, PII redaction, and runaway-loop prevention.

TOOLKIT · SECTION 2 OF 2

Artifacts 7 through 12.

ARTIFACT 07

Multi-Agent Orchestration Map

Supervisor, hierarchical, and peer-to-peer topologies. Trade-offs, message flow diagrams, failure modes.

ARTIFACT 08

Cost & Latency Budget Sheet

Per-call estimates, per-session caps, monthly forecasts. Pre-filled for GPT-4o, Claude, Gemini, Llama tiers.

ARTIFACT 09

Observability Stack Recipe

Trace, log, replay. LangSmith, Langfuse, and OpenTelemetry-based setups with sample dashboards.

ARTIFACT 10

Deployment Readiness Checklist

47 line-items spanning security, rollback, monitoring, rate-limits, and human-in-the-loop fallbacks.

ARTIFACT 11

RAG-for-Agents Pattern Sheet

When agents need retrieval, how to chunk for tool-use, and how to cite sources in final responses.

ARTIFACT 12

Incident Response Runbook

Hallucination, prompt injection, infinite loop, tool failure. Triage matrix and recovery scripts.

★ RELATED OFFER

You've seen the toolkit. Now certify on it.

The same 12 artifacts you just previewed form the spine of the GSDC Agentic AI Certification — work them in full, defend them in the exam, ship them in your portfolio.

[Claim Your Certification Seat →](#)

Linked to the toolkit you just reviewed · Verified credential

FRAMEWORK DECISION MATRIX

6 frameworks. 9 dimensions. One scored verdict.

The agentic AI landscape moves weekly. The matrix on the next page is rebuilt every quarter by the GSDC technical board and the engineering leads of the cohort employer panel.

The 9 scoring dimensions

- ✓ **Production maturity** — release cadence, breaking-change frequency, enterprise adopters on record.
- ✓ **Observability** — built-in tracing, third-party integrations, replay fidelity.
- ✓ **Multi-agent support** — supervisor patterns, peer messaging, shared state primitives.
- ✓ **Tool ecosystem** — first-party integrations, community connectors, schema standards.
- ✓ **Memory primitives** — short-term, long-term, vector store hooks.
- ✓ **Cost controls** — token budgets, request caching, model-routing helpers.
- ✓ **Learning curve** — time-to-first-agent for an engineer with two years of Python.
- ✓ **Vendor lock-in** — portability of agent definitions to other runtimes.
- ✓ **Community velocity** — issues closed per week, PR throughput, discussion volume.

Scoring scale: 1 (weak) → 5 (best-in-class). Weighted total normalized to 100.

FRAMEWORK DECISION MATRIX · SCORED

The scored matrix.

Six frameworks across nine dimensions. Scores reflect the Q1 2026 cohort review.

| Framework | Prod. | Obs. | Multi-Agt. | Tools | Memory | Cost | Curve | Lock-in | Comm. | Total |
|-------------------|-------|------|------------|-------|--------|------|-------|---------|-------|-------|
| LangGraph | 5 | 5 | 5 | 4 | 4 | 4 | 3 | 4 | 5 | 87 |
| CrewAI | 4 | 3 | 5 | 4 | 3 | 3 | 5 | 3 | 4 | 76 |
| AutoGen | 4 | 4 | 5 | 4 | 3 | 3 | 3 | 4 | 4 | 76 |
| OpenAI Assistants | 5 | 3 | 3 | 5 | 4 | 3 | 5 | 2 | 4 | 76 |
| Semantic Kernel | 4 | 3 | 4 | 4 | 4 | 4 | 3 | 4 | 3 | 73 |
| Smol Agents | 3 | 2 | 3 | 3 | 3 | 4 | 5 | 5 | 3 | 69 |

Verdict — when to pick what

LangGraph for production multi-agent systems with strong observability needs. **CrewAI** when speed-to-prototype matters more than runtime control. **OpenAI Assistants** for single-tenant SaaS with low ops overhead. **Semantic Kernel** in Microsoft-aligned enterprise stacks. **AutoGen** for research-flavored conversational agents. **Smol Agents** for code-execution-first lightweight tasks.

FRAMEWORK DECISION FLOW

The 5-question decision flow.

Use this top-down. Stop at the first match.

1. Are you shipping to production within 90 days?

→ Yes: shortlist LangGraph, OpenAI Assistants, Semantic Kernel. → No: continue.

2. Do you need more than one agent collaborating?

→ Yes: LangGraph, CrewAI, or AutoGen. → No: OpenAI Assistants or Smol Agents.

3. Is observability a hard requirement (audited workloads)?

→ Yes: LangGraph (LangSmith-native). → No: continue.

4. Is your stack Microsoft/.NET-aligned?

→ Yes: Semantic Kernel. → No: continue.

5. Do you need rapid prototyping over runtime control?

→ Yes: CrewAI. → No: AutoGen research; Smol Agents code-heavy.

 LIMITED TIME OFFER

This intake closes soon — don't miss the next cohort.

Registrations for the current GSDC Agentic AI Certification window are capped. When the seats fill, the next opening rolls to the following quarter.

[Reserve a Seat Now →](#)

Time-bound enrollment · Limited cohort size

SYLLABUS · MODULES 1 OF 4

11 module syllabi · verbatim.

Lifted directly from the certification handbook. No paraphrase, no edits.

MODULE 01

8 hours · 6% weight

Foundations of Agentic AI

Distinguishes agents from chatbots, RPA, and traditional LLM apps. Covers the perception-cognition-action loop, autonomy levels, and the historical lineage from rule-based agents to LLM-driven reasoners. Capstone: produce a one-page autonomy-tier classification of three real-world systems.

MODULE 02

10 hours · 8% weight

Reasoning Patterns: ReAct, Reflexion, Tree-of-Thought

Implementations and trade-offs of the three dominant reasoning loops. When to chain, when to branch, when to reflect. Hands-on: build the same task using all three patterns and benchmark token cost, latency, and solution quality.

MODULE 03

12 hours · 10% weight

Tool Use and Function Calling

Declaring tools, handling schema mismatches, retry semantics, idempotency, and side-effect management. Includes a deep-dive on JSON-schema validation and a section on tool-poisoning attack vectors.

SYLLABUS · MODULES 2 OF 4

Modules 4 through 6.

MODULE 04

10 hours · 9% weight

Memory Architectures

Short-term scratchpad, episodic memory, semantic memory, and procedural memory. Vector stores, key-value stores, and hybrid retrieval. Implementation walkthrough with Pinecone, Chroma, and pgvector.

MODULE 05

12 hours · 10% weight

Multi-Agent Orchestration

Supervisor patterns, hierarchical teams, peer-to-peer collaboration, and debate-based ensembles. Message-routing primitives, shared state management, and deadlock prevention. Hands-on lab: build a 3-agent research pipeline.

MODULE 06

10 hours · 9% weight

Planning and Task Decomposition

Hierarchical task networks, LLM-based planning, plan-monitor-replan loops. Covers when planning beats reactive execution and when it doesn't.

SYLLABUS · MODULES 3 OF 4

Modules 7 and 8.

MODULE 07

12 hours · 11% weight

Evaluation and Benchmarking

Designing scorecards for autonomous systems where ground truth is fuzzy. LLM-as-judge patterns, golden datasets, A/B harnesses, and trajectory-level evaluation. Includes the GSDC eval template used by Cohorts 38 onwards.

MODULE 08

10 hours · 10% weight

Safety, Guardrails & Alignment

Input/output filtering, jailbreak resistance, prompt-injection defenses, PII redaction, runaway-loop circuit-breakers. Maps to NIST AI RMF and the EU AI Act risk categories.

 50% OFF

Half off the Agentic AI Certification — this window only.

The GSDC Agentic AI Certification is currently available at a 50% reduction on the standard enrollment. Same syllabus, same exam, same credential — half the investment.

[Apply the 50% Discount →](#)

Promotional pricing · Apply at registration

SYLLABUS · MODULES 4 OF 4

Modules 9 through 11.

MODULE 09

8 hours · 8% weight

Observability & Debugging

Tracing across agent runs, structured logging for non-deterministic flows, replay-based debugging, and root-causing tool failures. Stack walkthroughs for LangSmith, Langfuse, and OpenTelemetry-based custom pipelines.

MODULE 10

10 hours · 9% weight

Production Deployment

Containerization, autoscaling, statefulness, rate limiting, and cost controls. Covers human-in-the-loop fallbacks, rollback strategies, and incident-response runbooks.

MODULE 11

14 hours · 10% weight

Capstone — Ship an Agent

End-to-end project: blueprint → build → evaluate → deploy → monitor. Defended in a live review with two GSDC technical board members and one industry assessor.

Total program: 116 instructional hours · 100% weighted assessment · capstone defense required for credentialing.

SAMPLE EXAM · FORMAT

Sample exam · printable checklist.

The certification exam is 90 minutes, 60 multiple-choice items, calibrated to map 1:1 with the 11 modules. A representative subset of items follows on the next two pages.

Exam blueprint

- ✓ **60 questions**, 90 minutes — single-best-answer multiple choice.
- ✓ **Passing threshold:** 70% — modular minimums apply on each section.
- ✓ **Item types:** conceptual, applied scenario, debugging trace, framework-trade-off.
- ✓ **References:** closed-book. No external lookups, no LLM assistance during the exam.
- ✓ **Retake policy:** two attempts included; 14-day cool-off between attempts.

Section weighting

| Section | Items | Weight |
|------------------------------------|-------|--------|
| Foundations & Reasoning Patterns | 11 | 18% |
| Tools, Memory & Orchestration | 17 | 28% |
| Evaluation, Safety & Observability | 17 | 28% |
| Planning & Production Deployment | 11 | 18% |
| Capstone Defense (oral) | 4 | 8% |

SAMPLE EXAM · ITEMS

Five representative items.

QUESTION 01 · MODULE 02

An agent loops repeatedly between "Thought" and "Action" without converging. Which reasoning pattern most directly addresses this?

- A. Plan-Execute with frozen plans
- B. Reflexion with self-critique on each iteration
- C. Tree-of-Thought with breadth-first search
- D. Single-shot Chain-of-Thought

QUESTION 02 · MODULE 05

In a supervisor-based multi-agent topology, which component is responsible for routing the next sub-task?

- A. The orchestrator agent
- B. The first available worker
- C. A round-robin scheduler
- D. The user, in every loop

QUESTION 03 · MODULE 08

A prompt-injection attack succeeds when the model treats untrusted input as ____ ?

- A. A latency event
- B. An instruction rather than data
- C. A function-call argument
- D. A retrieval result

 VALID 48 HOURS

This offer expires in 48 hours.

The current enrollment incentive for the GSDC Agentic AI Certification is active for the next 48 hours only. After it closes, standard terms resume for the rest of the quarter.

[Lock In Within 48 Hours →](#)

48-hour window · Standard terms resume after

SAMPLE EXAM · ITEMS (CONT.)

Questions 4 and 5.

QUESTION 04 · MODULE 07

An agent passes unit tests but degrades in production. Which eval pattern is best suited to catch this in CI?

- A. Single-shot accuracy on a held-out set
- B. Trajectory-level evaluation with shadow traffic
- C. A static linter run
- D. Latency benchmarks only

QUESTION 05 · MODULE 10

Which is the most important deployment safeguard for an agent with write-access to production databases?

- A. Hot reload on config change
- B. A human-in-the-loop approval gate on write tools
- C. A faster underlying model
- D. A larger context window

Printable readiness checklist · 8 of 47

- ✓ I can articulate the ReAct, Reflexion, and ToT differences in one breath.

- ✓ I have shipped at least one tool-using agent end-to-end.

- ✓ I have implemented a memory layer with at least two stores.

- ✓ I have a rollback plan for an in-production agent.

The full 47-point checklist ships with the certification handbook.

RESUME BULLET TEMPLATES · ATS-TESTED

Resume bullets, calibrated.

Eight templates. Plug your numbers in. They've been keyword-tested against the top three ATS engines used by Fortune-500 recruiters and the senior-recruiter scan-pattern study published by GSDC's careers council.

Template set 1 of 2 — Build & Ship

- Designed and deployed a production [multi-agent / single-agent] system using [LangGraph / CrewAI / AutoGen] that automated [task], reducing handle time by [X%] across [N] daily sessions.

- Architected a [tool-using / RAG-enabled] agent integrating [N] downstream APIs, achieving [Y%] task-completion rate on a [Z-item] internal benchmark.

- Built and maintained an evaluation harness covering [N] trajectory-level test cases, catching [X] regressions before production release across [Q] quarters.

- Reduced agent inference cost by [X%] through model-routing, prompt compression, and tool-call caching, saving [\$Y] monthly at [N]-session scale.

RESUME BULLET TEMPLATES · SET 2 OF 2

Lead & measure.

- Led a cross-functional team of [N] engineers and [M] product partners to deliver an autonomous [domain] agent serving [X] internal users.

- Authored and rolled out a [N-page] guardrails policy covering input validation, prompt-injection defenses, and PII redaction; cleared [audit-type] review on first pass.

- Implemented end-to-end observability with [LangSmith / Langfuse / OpenTelemetry], reducing mean-time-to-resolution for agent incidents from [X hours] to [Y minutes].

- Earned GSDC Agentic AI Certification (Cohort [N]); capstone defended on a [domain] agent shipped to [X] live users.

 LAST FEW SEATS

Cohort closing — final seats available.

Each GSDC cohort caps capstone-defense slots to preserve assessor-to-candidate ratios. The current intake is almost full. After full, the next opening rolls forward.

[Grab a Remaining Seat →](#)

Capacity-capped cohort · Assessor ratios protected

OUTCOMES

What certified cohorts report.

Self-reported outcomes from the most recent four cohorts (n=1,847 respondents, 60-day post-credential survey).

2,50,000+

Professionals certified

73%

Reported a role expansion within 6 months

4.7 / 5

Cohort-average rating, last 4 intakes

What changes after credentialing

- ✓ Recruiter inbound increases meaningfully when the credential appears on LinkedIn and resume header.
- ✓ Internal role-scoping shifts: certified individuals are routed to agentic AI workstreams over generic ML tickets.
- ✓ Cross-team credibility improves — the credential signals "ready to defend a system in front of an assessor."

"I went from 'the LLM person' on my team to leading the agentic platform initiative within the quarter."

— Cohort 51, Staff Engineer

AUDIENCE FIT

Who this certifies.

Strong fit

- ✓ ML / AI engineers shipping LLM-backed systems and wanting to graduate from chat-completions to agentic systems.
- ✓ Senior software engineers leading an agentic initiative at a 50+ person company.
- ✓ Solution architects who need to compare frameworks and defend a choice to a steering committee.
- ✓ Tech leads whose teams have built one or two agents and are now hitting evaluation, observability, and safety walls.
- ✓ Product engineers transitioning from RAG-heavy to action-heavy systems.

Less ideal fit

- ✓ Pre-engineering students with no production exposure — start with foundations programs first.
- ✓ Engineers seeking a pure-research credential — this is a build-and-ship certification.

BONUS INCLUDED

Free bonus: 1-on-1 capstone review on enrollment.

For the current intake, every new GSDC Agentic AI Certification registration includes a complimentary 30-minute capstone scoping session with a senior assessor — normally an add-on.

[Enroll & Claim the Bonus →](#)

Bonus session · Auto-applied to current intake

HOW IT WORKS

From enrollment to credential.

The end-to-end flow. Self-paced where it works, instructor-led where it matters.

STEP 01 · DAY 0

Enroll & receive the handbook

Access opens within minutes. Receive the full handbook, all 12 artifact templates, the matrix, the syllabus, and the readiness checklist.

STEP 02 · WEEKS 1–8

Work through the 11 modules

Self-paced video + reading + lab. Weekly office hours with a GSDC instructor. Slack-based cohort discussion.

STEP 03 · WEEKS 9–10

Build the capstone

Ship an agent end-to-end. Blueprint, build, evaluate, deploy, monitor. Submitted with a 6-slide defense deck.

STEP 04 · WEEK 11

Sit the exam & defend the capstone

90-minute written exam + 30-minute live capstone defense with two assessors. Both required.

STEP 05 · WEEK 12

Credential issued

Verifiable digital badge + serialised certificate + LinkedIn integration. Independently verifiable from the GSDC registry.

FAQ

Quick answers.

Is this self-paced or cohort-based?

Both. The 11 modules are self-paced; cohort touchpoints (office hours, capstone defense, exam sittings) anchor the timeline.

How long is the credential valid?

Three years from the issuance date. Renewal involves a delta-exam on syllabus updates since your cohort.

Is the credential independently verifiable?

Yes. Every certificate carries a serial that resolves on the public GSDC verification registry.

What if I fail the exam?

You get a second attempt within 14 days at no additional charge. Capstone re-submission is included once.

17 NEW COHORT FORMING

The next cohort starts shortly — secure your spot.

The upcoming GSDC Agentic AI Certification cohort is forming now. Enrol before kickoff to receive the full handbook, lab access, and pre-work materials ahead of week 1.

[Join the Next Cohort →](#)

Cohort-based start · Pre-work unlocks at enrolment

GET THE FULL TOOLKIT

Download the brochure.

This 22-page preview is the trailer. The full toolkit ships with editable templates, the worked-example archive, and the full 47-point readiness checklist.

Download Brochure

FULL NAME

e.g. Priya Sharma

WORK EMAIL

name@company.com

CODE

+91

PHONE

WHAT MOTIVATED YOU TO REGISTER?

Select Input ▼

COMPANY NAME

e.g. Acme Corp

Send Me the Free Guide →

Quick Primer · 2 fields

If forms slow you down, this is the short version.

FULL NAME

e.g. Priya Sharma

WORK EMAIL

name@company.com

CODE

+91

PHONE

Send Me the Free Primer



Instant PDF · We respect your inbox

THE WRAP

You've seen the index. Now run it.

Twelve artifacts. Eleven modules. Six frameworks compared. One exam. One capstone defense. The shortest distance between "I work with LLMs" and "I ship and defend agentic systems."

What to do next

- ✓ Skim the matrix on page 6 one more time and pick a framework to commit to.
- ✓ Take the 12-item readiness checklist on page 14. Note your gaps.
- ✓ Tap any CTA in this guide — they all land on the certification resource page.

LIFETIME ALUMNI ACCESS

Once certified, you stay in the room.

The GSDC Agentic AI Certification includes lifetime alumni access — every syllabus refresh, every new artifact, every cohort discussion, indefinitely. Earn it once, benefit from it as the field moves.

[Earn the Credential →](#)

Lifetime alumni access · Syllabus updates included

© Global Skill Development Council · All links resolve to www.gsdccouncil.org/certification-program/agentic-ai-resources