

The full field guide, on paper.

6-pillar architecture diagram. Framework comparison matrix. Six vendor write-ups. The 11-module syllabus. A sample exam.

Used by 2,50,000+ certified pros.

6

Architecture pillars

6

Vendor write-ups

11

Module syllabi

1

Sample exam

INSIDE THE GUIDE

What's in this field guide.

The framework you need to defend an agentic stack — to architects, to security review, to the steering committee. Printed, paginated, and quotable. Over the next 21 pages:

- ✓ **6-pillar architecture** — the reference model the GSDC technical board reviews every quarter. The vocabulary your design docs should be using.
- ✓ **Framework matrix · scored** — six frameworks across nine dimensions, with a verdict for each.
- ✓ **6 vendor write-ups** — LangGraph, CrewAI, AutoGen, OpenAI Assistants, Semantic Kernel, Smol Agents. What each is good at, where it cracks, who picks it.
- ✓ **11 module syllabi · verbatim** — pulled directly from the certification handbook, with hours and weighting.
- ✓ **Sample exam** — five representative items plus the 12-point readiness checklist.

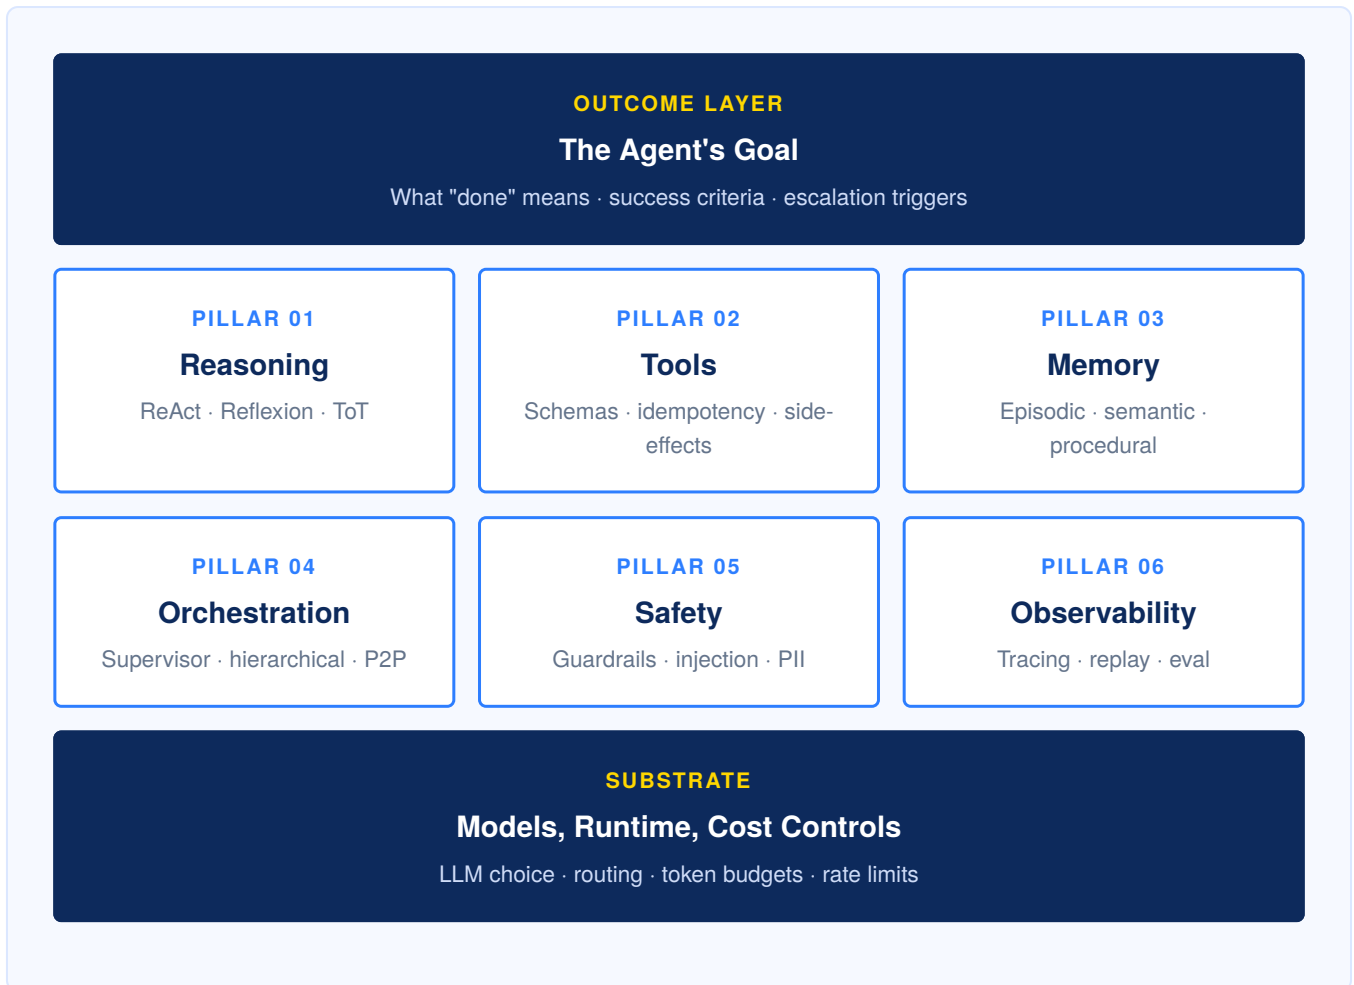
"I handed pages 4–6 to my CTO. We picked a framework that afternoon and shipped a pilot in eight weeks."

— Cohort 50, Principal Engineer

6-PILLAR ARCHITECTURE · DIAGRAM

The 6-pillar architecture.

Every production-grade agentic system the GSDC board has reviewed maps to these six pillars. If your design doc skips one, the review will surface it.



Read top-down: the goal frames the pillars; the substrate constrains them. Each pillar is detailed on pages 4–6.

Reasoning & Tools.

PILLAR 01 Reasoning

The loop the agent runs to decide what to do next. Determines latency, cost, and recovery behaviour. Wrong reasoning pattern is the single most common root-cause of "the agent loops forever" tickets.

- ▶ **ReAct** — interleaved thought/action; the default for tool-using agents.
- ▶ **Reflexion** — self-critique each iteration; the cure for non-convergence.
- ▶ **Tree-of-Thought** — branch and search; the right call for combinatorial reasoning.
- ▶ **Plan-Execute** — plan once, execute deterministically; fewer surprises, less flexibility.

PILLAR 02 Tools

The agent's interface to the world. Every tool is a contract: name, args, side-effects, idempotency. Treat tools as APIs, not as prompts.

- ▶ Declarative schemas (JSON-schema or equivalent) — non-negotiable.
- ▶ Idempotency keys for any tool with write side-effects.
- ▶ Retry semantics defined per tool, not globally.
- ▶ Tool-poisoning is a real attack surface — validate inputs before execution.

★ RELATED OFFER

You've read the pillars. Now defend them.

The same 6-pillar architecture is the spine of the GSDC Agentic AI Certification. The exam tests pillar-by-pillar literacy; the capstone tests pillar-by-pillar design.

[Enrol & Defend the Pillars →](#)

[Linked to the field guide](#) · [Verified credential](#)

Memory & Orchestration.

PILLAR 03 Memory

What the agent remembers and how it retrieves it. The pillar where cost lives, where leakage happens, and where staleness compounds.

- ▶ **Short-term** — scratchpad / working context. Token-budget-sensitive.
- ▶ **Episodic** — what happened, when. Vector stores work; KV stores often suffice.
- ▶ **Semantic** — facts and concepts. RAG-adjacent. Watch for stale embeddings.
- ▶ **Procedural** — how to do something. Often best encoded as tools, not memory.

PILLAR 04 Orchestration

How multiple agents (or sub-agents) collaborate. Picking the wrong topology costs more than picking the wrong framework — it forces a rewrite.

- ▶ **Supervisor** — a routing agent dispatches to specialised workers. Most common.
- ▶ **Hierarchical** — supervisors-of-supervisors. Powerful, observable, slower.
- ▶ **Peer-to-peer** — agents negotiate directly. Research-flavoured; rare in production.
- ▶ **Debate / ensemble** — same task, multiple agents, vote or judge. Expensive.

Safety & Observability.

PILLAR 05 Safety

The pillar that decides whether the system passes legal review. Maps to NIST AI RMF and the EU AI Act risk categories — both expect documented controls, not vibes.

- ▶ Input validation and PII redaction before the model sees the data.
- ▶ Prompt-injection defences — distinguish instructions from data.
- ▶ Output filtering — toxicity, secrets, structured-output schema enforcement.
- ▶ Runaway-loop circuit-breakers — token caps, step caps, cost caps.
- ▶ Human-in-the-loop gates on any tool with irreversible side-effects.

PILLAR 06 Observability

If you can't replay it, you can't debug it. If you can't measure it, you can't ship it. The pillar that separates side-projects from production systems.

- ▶ **Tracing** — every reasoning step, every tool call, every model response.
- ▶ **Structured logging** — for non-deterministic flows, structure beats verbosity.
- ▶ **Replay** — re-run a session end-to-end with the same inputs and intermediates.
- ▶ **Eval** — trajectory-level scorecards in CI, not just unit tests.

The six pillars run on the same substrate: model choice, runtime, cost controls. The substrate is not a pillar — it's the floor every pillar stands on.

FRAMEWORK MATRIX · SCORED

Six frameworks. Nine dimensions.

Rebuilt quarterly by the GSDC technical board and the engineering leads of the employer panel. Scores reflect the Q1 2026 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 |

Reading the table: 1 (weak) → 5 (best-in-class) per dimension; total weighted and normalized to 100. Individual vendor write-ups follow on pages 8–10.

 LIMITED TIME OFFER

This intake window closes soon.

The current GSDC Agentic AI Certification enrolment is open for a limited window. Once filled, the next opening rolls to the following quarter.

[Reserve a Seat Now →](#)

Time-bound enrolment · Limited cohort size

VENDOR WRITE-UPS · 1 OF 3

LangGraph & CrewAI.

LangGraph

Score 87 / 100

"The default for production multi-agent systems that need real observability."

A graph-based runtime built on top of LangChain. State is first-class; nodes are functions; edges encode the topology. The strongest production story in the matrix — durable execution, time-travel debugging, LangSmith-native tracing. The cost is a real learning curve and a deeper LangChain dependency than some teams want.

★ Best observability

★ Production-grade

★ Multi-agent first-class

⚠ Learning curve

⚠ LangChain coupling

CrewAI

Score 76 / 100

"Speed-to-prototype champion. The first agent you ship in a weekend."

Role-based agents with a deliberately small surface. Define a crew, give each member a role and tools, run the task. Excellent for demos and internal pilots; observability and runtime control are weaker, so production-grade deployments often graduate to LangGraph after a quarter.

★ Fastest to prototype

★ Role-based clarity

★ Strong community

⚠ Lighter observability

⚠ Less runtime control

AutoGen & OpenAI Assistants.

AutoGen

Score 76 / 100

"Microsoft Research's conversational multi-agent stack."

Multi-agent conversations are the default abstraction. Excellent for research-flavoured workflows where agents negotiate, debate, and review each other's output. The conversational metaphor cuts both ways — easy to reason about, sometimes harder to bound. The 0.4 rewrite improved production posture meaningfully.

★ Best for research workflows

★ Conversational multi-agent

★ Research-backed

⚠ Conversational abstraction can leak

OpenAI Assistants API

Score 76 / 100

"Lowest-ops single-tenant SaaS path. Vendor lock-in is the price."

A managed runtime. Threads, tools, file search, code execution — all hosted. The fastest path to a credible single-agent product if your stack and audience already live in the OpenAI ecosystem. Multi-agent is weaker, portability is the worst in the matrix, and pricing surprises happen at scale.

★ Lowest ops overhead

★ Built-in tools

★ Gentlest learning curve

⚠ High vendor lock-in

⚠ Weaker multi-agent

VENDOR WRITE-UPS · 3 OF 3

Semantic Kernel & Smol Agents.

Semantic Kernel

Score 73 / 100

"The right answer inside Microsoft-aligned enterprise stacks."

Skills, planners, memory — Microsoft's first-party SDK with strong .NET and Azure integration. Picks up enterprise patterns (auth, identity, governance) more cleanly than alternatives in the same stack. Documentation has stabilised; ecosystem is more focused than viral.

★ .NET / Azure native

★ Enterprise-aligned

★ Strong cost controls

⚠ Smaller community velocity

Smol Agents

Score 69 / 100

"Lightweight, code-execution-first, almost no lock-in."

Hugging Face's deliberately small framework. Agents write and execute code as their primary action. Excellent for code-heavy lightweight tasks; less appropriate for long-running multi-agent systems. The lowest-friction option for prototyping.

★ Lowest lock-in

★ Code-execution-first

★ Tiny footprint

⚠ Limited multi-agent

⚠ Observability is DIY

 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 enrolment. Same syllabus, same exam, same credential — half the investment.

[Apply the 50% Discount →](#)

Promotional pricing · Applied at registration

VENDOR PICK · DECISION GUIDE

One-line verdicts.

Use this as a shortlist generator. The matrix on page 7 and the write-ups on pages 8–10 give you the full defence.

| Pick if... | Framework |
|---|--------------------------|
| You're shipping a production multi-agent platform with audited observability. | LangGraph |
| Speed to first agent matters more than runtime control. | CrewAI |
| You're running research-flavoured conversational workflows. | AutoGen |
| You're inside the OpenAI ecosystem and need the lowest-ops single-agent path. | OpenAI Assistants |
| Your stack is .NET / Azure-aligned. | Semantic Kernel |
| Code-execution-first lightweight tasks with minimal lock-in. | Smol Agents |

Three traps to avoid

- ✓ **Picking on score alone.** Total score is a useful tiebreaker, not a verdict. The dimension that matters most to your stack should drive the call.
- ✓ **Picking on demos.** Every framework looks good in a five-minute video. Build the same toy task in two finalists before committing.
- ✓ **Picking once, forever.** Frameworks move fast. Re-audit your choice every two quarters; the matrix is published quarterly for exactly this reason.

SYLLABUS · 1 OF 4

11 module syllabi · **verbatim.**

Lifted directly from the certification handbook. No paraphrase.

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.

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 cost, latency, and quality.

MODULE 03

12 hours · 10% weight

Tool Use and Function Calling

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

SYLLABUS · 2 OF 4

Modules 4 through 6.

MODULE 04

10 hours · 9% weight

Memory Architectures

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

MODULE 05

12 hours · 10% weight

Multi-Agent Orchestration

Supervisor patterns, hierarchical teams, peer-to-peer collaboration, debate-based ensembles. Message routing, shared state, deadlock prevention. Hands-on: 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. When planning beats reactive execution and when it doesn't.

 **VALID 48 HOURS**

This offer expires in 48 hours.

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

Lock In Within 48 Hours →

48-hour window · Standard terms resume after

SYLLABUS · 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, 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 defences, PII redaction, runaway-loop circuit-breakers. Maps to NIST AI RMF and the EU AI Act risk categories.

"Module 07 is where the field guide becomes a job. Walking out with a working eval harness changed how my team scoped every shipping decision."

— Cohort 44, Engineering Manager

SYLLABUS · 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, root-causing tool failures. Walkthroughs for LangSmith, Langfuse, and OpenTelemetry-based custom pipelines.

MODULE 10

10 hours · 9% weight

Production Deployment

Containerisation, autoscaling, statefulness, rate limiting, cost controls. Human-in-the-loop fallbacks, rollback strategies, incident-response runbooks.

MODULE 11

14 hours · 10% weight

Capstone — Ship an Agent

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

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

SAMPLE EXAM · ITEMS

Sample exam.

90 minutes total, 60 multiple-choice items, 70% passing threshold, closed-book. Two representative items below; the full handbook ships a longer reference set.

QUESTION 01 · MODULE 02 · PILLAR 01

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 10 · PILLAR 06

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

📢 LAST FEW SEATS

Cohort closing — final seats available.

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

[Grab a Remaining Seat →](#)

Capacity-capped cohort · Assessor ratios protected

PRINTABLE READINESS CHECKLIST

The 12-point go / no-go.

Print, tick, file. If you can honestly tick 10 of 12, you're ready to sit the exam — and to defend a framework choice to your steering committee.

- ✓ I can draw the 6-pillar architecture from memory and explain each pillar in one breath.

- ✓ I can articulate the ReAct, Reflexion, and Tree-of-Thought differences in one breath.

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

- ✓ I can name three failure modes of a supervisor pattern.

- ✓ I have implemented a memory layer with at least two stores (vector + KV, for instance).

- ✓ I can draw the GSDC eval scorecard from memory.

- ✓ I have a written guardrail policy for one of my agents.

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

- ✓ I can name 6 mainstream frameworks and pick between them in 60 seconds.

- ✓ I have wired one observability stack (LangSmith / Langfuse / OTel).

- ✓ I have written a trajectory-level eval case at least once.

- ✓ I have defended a framework choice against at least two alternatives in writing.

The full handbook ships a deeper 47-point checklist for the day-of-exam audit.

AUDIENCE FIT

Who this certifies.

Strong fit

- ✓ ML / AI engineers shipping LLM-backed systems and ready 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 hitting evaluation, observability, and safety walls.
- ✓ Platform engineers absorbing agentic workloads into existing platforms and runtimes.

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 scoping call.

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

[Enrol & Claim the Bonus →](#)

Bonus session · Auto-applied to current intake

HOW IT WORKS

From enrolment to **credential.**

Self-paced where it works, cohort-anchored where it matters.

STEP 01 · DAY 0

Enrol & receive the handbook

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

STEP 02 · WEEKS 1–8

Work through the 11 modules

Self-paced video + reading + labs. 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 defence deck.

STEP 04 · WEEK 11

Sit the exam & defend the capstone

90-minute written exam + 30-minute live capstone defence 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 the certification framework-specific?

No. The 6-pillar architecture and the framework matrix are deliberately framework-agnostic. The exam tests literacy across all six tracked vendors; the capstone lets you pick one to ship on.

How often is the framework matrix rebuilt?

Quarterly. The Q1 2026 scores are reflected on page 7. Active candidates receive the updated matrix with each release.

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?

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

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 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 FIELD GUIDE

Download the field guide.

This 22-page preview is the trailer. The full field guide ships with editable templates, the full vendor-comparison spreadsheet, and the pillar-by-pillar design worksheet.

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

The short version. Phone optional.

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 read the guide. Now own it.

Six pillars. Six vendors. Eleven modules. One exam. One capstone defence. The shortest distance between "I work with LLMs" and "I can defend an agentic stack on paper, in a room, in production."

What to do next

- ✓ Re-read pages 3–6. Can you draw the 6-pillar architecture from memory?
- ✓ Take the 12-point readiness checklist on page 17.
- ✓ Tap any CTA in this guide — every link lands on the framework certification 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 matrix release, every cohort discussion, indefinitely.

[Earn the Credential →](#)

Lifetime alumni access · Matrix updates included

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