THE COMPLETE STORY · 2026

The Complete Story

Every chapter from the GSDC **Certified Agentic AI Professional** learner journey, in one downloadable PDF: the 5-chapter story, the full 11-module syllabus, all 50+ Learn-by-Doing labs, the capstone outline and a sample exam.

Inside the brochure

✓ **5-chapter learner journey · expanded**	✓ **11 module syllabi · verbatim**
✓ **50+ Learn-by-Doing labs · full list**	✓ **Sample exam · printable checklist**
✓ **USA salary cards · 2026 figures**	✓ **Capstone outline · end-to-end**

5	11	250k+
Chapters	**Modules**	**Certified**
Learner journey	Verbatim syllabi	Already inside

Credential: **AAIPC · Certified Agentic AI Professional** · Issued by **Global Skill Development Council** · **Vendor-neutral** · Recognized in 100+ countries.

Page one carries no calls-to-action by design. The 5-chapter journey begins on page two.

JOURNEY HERO

The 5-chapter learner journey

Foundations → Architecture → Build → Govern → Ship. Five chapters move you from agentic AI literacy through your first multi-agent system to a deployed, governed agent — and the exam at the end.

THE 5-CHAPTER LEARNER JOURNEY · USED BY 250,000+ CERTIFIED

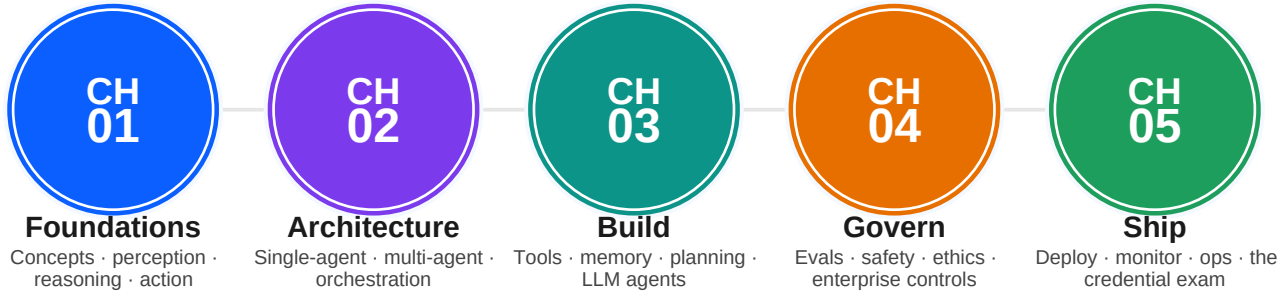

The story arc: the 5 chapters map the same arc every credential-holder describes in interviews. Pages 5–9 carry one page per chapter; the 11-module syllabus on pages 11–12 sits inside the same arc.

CONTEXT

Why agentic AI, why this credential

Agentic AI moved from research demos to production stacks faster than any AI shift since the original transformer paper. By 2026, agents are running in customer support, research, sales operations, internal IT, finance ops and software engineering — and the skills gap on building them defensibly has widened sharply.

Three forces converging

- **From chat to autonomy.** Conversational AI is a feature; autonomous agents are an architectural shift.
- **Multi-agent orchestration is going mainstream.** One agent rarely cuts it — coordination is the real skill.
- **Governance is racing to keep up.** Boards want defensible AI; agents need evals, controls and kill-switches.

What this credential does

- **Vendor-neutral framework coverage.** Skills transfer across LangGraph · CrewAI · AutoGen · OpenAI Assistants.
- **Hands-on through the GSDC Live Studio.** 50+ Learn-by-Doing labs with lifetime access.
- **Daily live sessions.** Industry experts and SMEs guide you through the stack.
- **Globally recognized credential.** Used by 250,000+ certified professionals across 100+ countries.

What you can do by the end: design, build, govern and deploy an agentic AI system end-to-end — with evals, safety controls, and the vocabulary your enterprise stakeholders actually use.

ADOPTION

Used by 250,000+ certified professionals worldwide

AAIPC sits inside the same GSDC credential family trusted by a quarter of a million certified professionals globally. The community spans the firms most actively shipping agentic AI in 2026 — across software, services, finance, healthcare and the public sector.

250k+	100+	50+
Certified	**Countries**	**LBD labs**
Across all GSDC programs	Recognition	Hands-on

Where the community sits

Sector	Common roles in the community
Software & SaaS	Agentic AI engineer · LLM agent developer · AI platform engineer
Consulting & services	AI delivery consultant · agentic AI architect · solutions engineer
Financial services	AI risk analyst · agentic ops engineer · model validator (agents)
Healthcare & life sciences	Clinical AI engineer · agentic ops engineer · safety engineer
Customer support & ops	Support automation lead · agentic ops engineer · AI product manager
Public sector & policy	Supervisory analyst · policy advisor · enforcement analyst

RELATED · NEXT STEP

[50% OFF]

Live the story — start Chapter 1 this week

Every chapter in this brochure is taught hands-on inside the AAIPC curriculum. Enroll while the half-price window is open.

Enroll & Save 50% →

Tap anywhere on this card →

CH 01
CHAPTER

Foundations

Agentic AI literacy · perception · reasoning · action

The story of Chapter 1

Most readers come to AAIPC with comfort in LLMs and prompts, and curiosity about agents. Chapter 1 builds the shared vocabulary — perception, reasoning, action, memory, tool use, planning — and grounds it in the difference between an LLM call, a single-step agent, and a multi-step autonomous system. By the end of the chapter, you can explain agentic AI to a senior stakeholder in plain language.

What you learn

- The agentic AI vocabulary: perception · reasoning · action · memory · tool use · planning.
- How agents differ from chat assistants — and where the differences matter operationally.
- Vendor-neutral mental models that transfer across LangGraph · CrewAI · AutoGen · OpenAI Assistants.
- Where agents shine — and where they don't (the “not yet” list every engineer should keep).

What you ship (LBD highlights)

- **LBD 01:** Glossary of 40 agentic-AI terms in plain language.
- **LBD 02:** Comparison sheet — chat vs single-step agent vs multi-step autonomous system.
- **LBD 03:** A one-page “where agents fit / where they don't” rubric for your domain.
- **LBD 04:** A perception-reasoning-action diagram of a real use case from your team.

CH 02
CHAPTER

Architecture

Single-agent · multi-agent · orchestration patterns

The story of Chapter 2

Once you can describe an agent, the next step is choosing the right architecture. Chapter 2 covers single-agent loops, multi-agent coordination, graph-based orchestration, and the trade-offs each pattern carries — cost, latency, debuggability, governance. The chapter is deliberately vendor-neutral; patterns transfer across stacks.

What you learn

- The four canonical agent patterns: single-agent loop · supervisor-worker · debate · graph orchestration.
- How to choose between patterns by task complexity, tool count and reliability requirements.
- Coordination primitives: handoffs · shared memory · message buses · checkpointing.
- Trade-offs by pattern: cost · latency · debuggability · safety boundary.

What you ship (LBD highlights)

- **LBD 09:** Pattern decision matrix for one real use case (with picked pattern + rationale).
- **LBD 10:** Architecture diagram for a multi-agent system you might build at your firm.
- **LBD 11:** “Handoff contract” spec between two agents — inputs, outputs, failure modes.
- **LBD 13:** Cost & latency budget worksheet across your chosen pattern.

The architectural rule of thumb: reach for the simplest pattern that meets your requirements. Multi-agent is not better than single-agent — it's *different*, with real cost and debuggability trade-offs.

CH 03
CHAPTER

Build

Tools · memory · planning · LLM agents in code

The story of Chapter 3

Chapter 3 is where you build. Tool surfaces with explicit schemas, structured outputs, memory strategies, planning loops, and the LLM-agent code that ties them together. Vendor-neutral patterns first; vendor-specific implementations second. By the end of the chapter you have a working single-step agent *and* a working multi-step agent on a real use case.

What you learn

- Tool surfaces: schema design · read-only defaults · approval gates · safe error handling.
- Memory strategies: scratchpad · episodic · long-term semantic memory in agent loops.
- Planning patterns: ReAct · plan-then-execute · reflection · self-correction.
- LLM agent code: vendor-neutral patterns first; LangGraph, CrewAI, AutoGen, OpenAI Assistants second.

What you ship (LBD highlights)

- **LBD 18:** Tool surface spec with schemas, defaults and approval gates.
- **LBD 22:** Working single-step agent on a real read-only task.
- **LBD 25:** Multi-step agent with planning loop and reflection.
- **LBD 28:** A memory strategy doc — what you keep, where, and for how long.

CH 04
CHAPTER

Govern

Evals · safety · ethics · enterprise controls

The story of Chapter 4

Chapter 4 is where your agent becomes defensible. Evaluation harnesses for agentic behaviour (not just text), safety policies that ride along with the agent, ethics review frameworks, and the enterprise controls that auditors and second-line risk teams actually ask for. This chapter is what separates a demo from a system.

What you learn

- Agentic eval harnesses: trajectory evals · tool-use evals · refusal correctness · drift.
- Safety policies: tool-level capability boundaries · approval gates · kill-switches.
- Ethics review: fairness · accountability · transparency · explainability for agents.
- Enterprise controls: model cards · system cards · audit trails · vendor due-diligence.

What you ship (LBD highlights)

- **LBD 33:** Trajectory eval harness with a gold set of 20+ scenarios.
- **LBD 36:** A safety policy document for your agent (capability boundaries + escalation).
- **LBD 40:** Agentic system card: intended use · limitations · evals · controls.
- **LBD 42:** Mock regulator interview transcript on your agent — written + reviewed.

CH 05
CHAPTER

Ship

Deploy · monitor · ops · the exam at the end

The story of Chapter 5

Chapter 5 is the operational chapter — and the credential chapter. Deployment patterns, production monitoring, on-call playbooks for agentic systems, cost control, capstone build, sample exams and the AAIPC exam itself. By the end of this chapter you don't just hold the credential; you have a real artifact recruiters can open.

What you learn

- Deployment: containerization · staged rollouts · feature flags for agents · rollback patterns.
- Monitoring: per-trajectory observability · cost dashboards · latency budgets · drift alerts.
- Ops: on-call playbooks · agentic incident response · post-mortems · regression test gates.
- Capstone & exam: capstone build · two sample exams · mock viva · the AAIPC exam itself.

What you ship (LBD highlights)

- **LBD 46:** Deployment plan with staged rollout and rollback criteria.
- **LBD 48:** Monitoring spec: trajectories logged · cost dashboard · latency budget.
- **LBD 49:** Capstone — deployed agent with evals, controls and exec summary.
- **LBD 50+:** Two sample exams + mock viva + the AAIPC exam booked.

JOURNEY × SYLLABUS

Chapters × modules · the lookup matrix

Every module in the AAIPC syllabus sits inside one of the 5 chapters. The matrix below shows the canonical mapping. Use it to plan which modules to study when — and to describe the program clearly to a hiring manager.

Module	Ch 1 Found.	Ch 2 Arch.	Ch 3 Build	Ch 4 Govern	Ch 5 Ship
M1 · Agentic AI literacy	●	▪			
M2 · LLM mechanics for agents	●	▪	▪		
M3 · Architecture patterns	▪	●	▪	▪	
M4 · Tool use & tool surfaces	▪	▪	●	▪	▪
M5 · Memory & planning		▪	●	▪	
M6 · Multi-agent & orchestration		●	●	▪	▪
M7 · Evals & observability		▪	▪	●	●
M8 · Safety & controls		▪	▪	●	▪
M9 · Ethics & responsible AI	▪			●	▪
M10 · Deployment & ops			▪	▪	●
M11 · Capstone & exam	▪	▪	▪	▪	●

Legend: ● primary mapping · ▪ strong secondary. Use the matrix when scheduling your study — start with modules whose primary chapter is the one you most need to build.

SYLLABUS

11 module syllabi · verbatim (1 of 2)

#	Module · verbatim	LBD count
1	Agentic AI literacy & the autonomy spectrum	4
2	LLM mechanics for agents · context, tokens, reasoning	5
3	Architecture patterns · single-agent and multi-agent	5
4	Tool use & tool surface design	5
5	Memory, planning & reflection	5
6	Multi-agent orchestration · graphs & coordination	5

Learning outcomes per module · in one line

#	You can…
M1	Define agentic AI in plain language and place a use case on the autonomy spectrum.
M2	Reason about context windows, token costs and structured outputs in agentic loops.
M3	Pick the right architecture pattern for a given task with explicit trade-offs.
M4	Design a tool surface with safe defaults, schemas and approval gates.
M5	Implement memory and planning strategies fit for the agent's task complexity.
M6	Orchestrate multi-agent systems with clear handoffs, contracts and checkpoints.

SYLLABUS

11 module syllabi · verbatim (2 of 2)

#	Module · verbatim	LBD count
7	Evals & observability for agentic systems	5
8	Safety, capability boundaries & controls	5
9	Ethics & responsible AI for agents	4
10	Deployment, monitoring & ops	5
11	Capstone, exam prep & viva	2

Learning outcomes per module · in one line

#	You can…
M7	Build a trajectory eval harness with a gold set and observability dashboards.
M8	Define capability boundaries and approval gates that survive a security review.
M9	Frame an ethical review for an agent across fairness, accountability and transparency.
M10	Plan a staged rollout with monitoring, kill-switches and rollback criteria.
M11	Sit the AAIPC exam under timed conditions and ship a capstone.

Total LBDs across the program: 50. The full list is indexed on pages 13–14, with each LBD's name, the module it sits in, and the recruiter-readable deliverable it ships.

LBD INDEX

The 50+ Learn-by-Doing labs (1 of 2)

Every Learn-by-Doing activity, with the recruiter-readable deliverable. Print pages 13–14 as your portfolio checklist.

LBD	M	Activity	Deliverable
01	1	Glossary of 40 agentic-AI terms	Glossary doc
02	1	Chat vs single-step vs multi-step comparison	Comparison sheet
03	1	Where agents fit / don't fit rubric	1-page rubric
04	1	Perception-reasoning-action diagram of a real case	Architecture sketch
05	2	Context-window planner for a long-running agent	Context budget sheet
06	2	Structured output schema for an agent's outputs	JSON schema
07	2	Token cost simulator for an agent loop	Cost simulator
08	2	Reasoning trace analysis on three model sizes	Trace analysis memo
09	3	Pattern decision matrix for one real use case	Decision matrix
10	3	Architecture diagram for a multi-agent system	Architecture diagram
11	3	Handoff contract between two agents	Contract spec
12	3	Failure-mode catalogue for your chosen pattern	Failure catalogue
13	3	Cost & latency budget worksheet	Budget sheet
14	4	Read-only tool surface design	Surface spec
15	4	Approval-gate spec for write tools	Approval spec
16	4	Safe error-handling rules for tool calls	Error rules doc
17	4	Tool surface unit-test set	Test set
18	4	Tool surface spec with schemas, defaults, gates	Surface spec
19	5	Scratchpad design for a multi-step agent	Scratchpad spec
20	5	Episodic memory strategy with retention rules	Memory strategy doc
21	5	Long-term semantic memory design	Memory schema
22	5	Working single-step agent on a real read-only task	Single-step agent
23	5	Reflection loop with self-correction	Reflection agent
24	6	Supervisor-worker multi-agent prototype	Multi-agent prototype
25	6	Multi-step agent with planning loop & reflection	Multi-step agent

LBD INDEX

The 50+ Learn-by-Doing labs (2 of 2)

LBD	M	Activity	Deliverable
26	6	Graph-orchestrated agent flow	Graph flow
27	6	Handoff failure-injection test	Test report
28	6	Memory strategy doc for your system	Strategy doc
29	7	Trajectory eval design across 3 personas	Eval design
30	7	Tool-use eval set with a gold set	Gold set
31	7	Refusal-correctness eval on edge cases	Refusal evals
32	7	Drift monitor design for your agent	Drift spec
33	7	Trajectory eval harness with 20+ scenarios	Eval harness
34	8	Capability boundary spec for your agent	Boundary spec
35	8	Kill-switch design with triggers and approvers	Kill-switch spec
36	8	Safety policy document	Safety policy
37	8	Prompt-injection defence test set	Injection test set
38	8	Permission-aware retrieval design	Retrieval spec
39	9	Fairness audit on one consequential use case	Audit note
40	9	Agentic system card: use · limits · evals · controls	System card
41	9	AI-use disclosure for an enterprise rollout	Disclosure
42	9	Mock regulator interview transcript	Transcript
43	10	Staged rollout plan with feature flags	Rollout plan
44	10	On-call playbook for agentic incidents	Playbook
45	10	Post-mortem template for agent failures	Template
46	10	Deployment plan with rollback criteria	Deployment plan
47	10	Cost dashboard spec for your agent	Dashboard spec
48	10	Monitoring spec: trajectories · cost · latency	Monitoring spec
49	11	Capstone: deployed agent with evals, controls, summary	Capstone artifact
50	11	Two sample exams + mock viva	Exam evidence pack

50 named, recruiter-readable deliverables. Every LBD ships an artifact a hiring manager understands without translation — a system card, an eval harness, a monitoring spec, not just “a course completion certificate.”

CAPSTONE

The capstone outline · end-to-end

The capstone is the artifact recruiters open first. Below is the canonical capstone shape — six sections, each with a measurable output, all built across the 5-chapter journey.

Section	What it contains	Built in
1 · The brief	Use case · users · success metrics · risks · controls · pattern choice.	Chapter 1–2
2 · The architecture	Pattern · handoff contracts · memory · tool surfaces · cost / latency budget.	Chapter 2
3 · The build	Working agent: single-step or multi-step. Code · tools · memory · planning.	Chapter 3
4 · The evals	Trajectory evals · tool-use evals · refusal correctness · drift monitor.	Chapter 4
5 · The governance	Safety policy · capability boundaries · kill-switch · system card.	Chapter 4
6 · The ship	Deployment plan · monitoring · on-call playbook · exec summary · viva.	Chapter 5

Three capstone examples from the community

- **Customer-support triage agent** · multi-agent · ~32% reduction in first-response time · adopted by team of 12.
- **Internal research assistant** · single-step + tools · gold-set eval at 0.91 citation faithfulness.
- **Sales-ops follow-up agent** · supervisor-worker · ~3 hrs saved / rep / week · governance pack approved by second-line risk.

EXAM PREP

Sample exam · structure & sample items

The AAIPC exam tests application of the 5-chapter journey, not recitation. The structure below mirrors real practitioner decisions across the 11 modules and 50+ LBDs.

Item	Detail
Credential code	AAIPC · Certified Agentic AI Professional
Issuing body	Global Skill Development Council (GSDC)
Vendor-neutral	Yes · framework-neutral coverage across LangGraph · CrewAI · AutoGen · OpenAI Assistants
Format	Multiple-choice + scenario items + applied mini-case
Items	Approx. 50–70 items
Duration	Approx. 75–100 minutes (varies by version)
Open book	No · closed-book, online proctored
Passing line	Set by GSDC psychometric review · communicated at booking
Retake policy	Retakes allowed per the GSDC retake policy
Validity	Lifetime credential · refresh recommended every 2 years
Recognition	Recognized in 100+ countries · used by 250,000+ certified
Money-back	7-day money-back guarantee per the GSDC refund policy

EXAM PREP

Sample exam questions (1 of 2)

Q1. An engineer ships a customer-support agent with write access to the order-management system. Which control is **most** important before go-live?

A. Use the largest available model.

B. Restrict the agent's tool surface to read-only by default; require human approval for any write action.

C. Disable evals to speed up release.

D. Run the agent only on weekends to limit blast radius.

Answer: B · *Capability boundaries live at the tool surface, not the prompt. Read-only defaults + human approval is the canonical pattern.*

Q2. Which of the following is the **best** reason to reach for a multi-agent pattern instead of a single-agent loop?

A. Multi-agent always has lower cost.

B. The task has clearly separable sub-tasks that benefit from specialised tool surfaces and contexts.

C. The team wants to use the latest framework.

D. Multi-agent is easier to debug than single-agent.

Answer: B · *Multi-agent shines when sub-tasks are separable. It is not cheaper or easier to debug; it's a different tool with real trade-offs.*

Q3. What is the **most** appropriate artifact to bring to a second-line risk review of a deployed agentic system?

A. A screenshot of the chat UI.

B. An agentic system card: intended use · limitations · evaluations · controls.

C. A list of all prompts ever sent in production.

D. Vendor marketing collateral.

Answer: B · *Agentic system cards are the standard artifact for risk review. They communicate intent, limits and controls in one document.*

Q4. Your agent occasionally invents a numeric value when summarising data. Which mitigation is **most** appropriate?

A. Increase the model size.

B. Inject numeric values deterministically from the underlying data store; never let the model generate numbers.

C. Add a disclaimer to the output.

D. Reduce the temperature to zero.

Answer: B · *Numbers belong to the system of record. Template-bound injection — not model generation — is the standard mitigation.*

EXAM PREP

Sample exam questions (2 of 2)

Q5. A multi-agent system produces inconsistent outputs across runs. Which is the **most** useful first investigation?

A. Switch frameworks.

B. Inspect the handoff contracts and trajectory traces for non-determinism in the routing decisions.

C. Lower the model size.

D. Remove the reflection step.

Answer: B · *Inconsistency in multi-agent runs usually traces to routing variance and weak handoff contracts. Traces are the diagnostic surface.*

Q6. A regulator asks how your agent will behave under conditions it has not seen before. Which is the **best** response?

A. Our agent is deterministic, so it won't.

B. We have capability boundaries at the tool surface, drift monitors against a gold set, and a kill-switch with documented triggers.

C. We trust the vendor's safety claims.

D. We use the largest model, which doesn't have this problem.

Answer: B · *Bounded capability, monitored drift and kill-switches are the canonical regulator-facing answer for agentic systems.*

Q7. A “trajectory eval” in an agentic system is best described as:

A. An eval of the final answer only.

B. An eval of the agent's intermediate steps, tool calls and decisions across a scenario.

C. A latency benchmark.

D. A vendor-provided benchmark you cannot inspect.

Answer: B · *Trajectory evals measure the full sequence of decisions, not just the end output. They are the agentic equivalent of unit + integration tests.*

Q8. Which is the **least** appropriate use case for a fully autonomous agent today?

A. Triage of internal support tickets with a human approver before resolution.

B. Research summary with cited sources and an editorial review step.

C. Initiating high-value irreversible financial transactions without a human gate.

D. Drafting marketing copy with a brand reviewer step.

Answer: C · *High-value irreversible actions need explicit human gates by default. Other options have explicit human-in-the-loop reviews.*

CHECKLIST

Printable exam-prep checklist

Print this page. Tick each box as you complete it. The list below is the minimum evidence you should have on hand by the time you sit the AAIPC exam.

Chapter 1 · Foundations

- ☐ I can define agentic AI in plain language for a senior non-technical stakeholder.
- ☐ I can place a use case on the autonomy spectrum without hesitation.
- ☐ I can name where agents shine — and where they don't yet.
- ☐ I have a perception-reasoning-action diagram of a use case from my own work.

Chapter 2 · Architecture

- ☐ I can describe the four canonical patterns and pick between them with rationale.
- ☐ I have a pattern decision matrix for one real use case.
- ☐ I have written a handoff contract between two agents.
- ☐ I have a cost / latency budget for my chosen pattern.

Chapter 3 · Build

- ☐ I have shipped a working single-step agent on a real read-only task.
- ☐ I have shipped a multi-step agent with a planning loop and reflection.
- ☐ I have written a tool surface spec with schemas, defaults and approval gates.
- ☐ I have a memory strategy doc that explains what I keep, where, and for how long.

Chapter 4 · Govern

- ☐ I have a trajectory eval harness with a gold set of 20+ scenarios.
- ☐ I have a safety policy with capability boundaries and a kill-switch.
- ☐ I have written an agentic system card for my agent.
- ☐ I have rehearsed a mock regulator interview with a mentor.

Chapter 5 · Ship

- ☐ I have a deployment plan with staged rollout and rollback criteria.
- ☐ I have a monitoring spec covering trajectories, cost and latency.
- ☐ I have shipped a capstone with deployment, evals, controls and an exec summary.
- ☐ I have completed two sample exams under timed conditions; my exam slot is booked.

USA SALARY

USA salary cards · 2026 figures

Indicative 2026 annual base salary bands for the most common agentic AI roles in the United States. Bands triangulated from public salary data; exclude bonus, equity and sign-on. Use as directional benchmarks for negotiation anchoring.

Agentic AI Engineer (Mid) **USD 140–200k** 3–5 yrs · major metro	**Senior Agentic AI Engineer** **USD 180–260k** 5–8 yrs · major metro	**Staff / Principal Agent Engineer** **USD 240–340k** 8+ yrs · NYC / SF Bay
Agentic AI Architect **USD 200–290k** Architecture · multi-agent	**Agentic Ops Engineer** **USD 150–215k** Monitoring · ops · evals	**AI Product Manager (Agents)** **USD 160–240k** Product · go-to-market

How to read these cards

- **Mid (3–5 yrs)** assumes major-metro base; NYC/SF Bay typically sits 20–30% higher on base.
- **Senior (5–8 yrs)** bands reward demonstrable production-agent experience and eval craft.
- **Staff / Principal (8+ yrs)** bands assume multi-agent system ownership and architectural leadership.
- **Architect** bands reward vendor-neutral pattern fluency — exactly what AAIPC teaches.
- **Ops & PM** roles are the fastest-growing agentic seats outside engineering.

PROGRAM

How the program works

The credential behind the journey. Same delivery format used across the GSDC AI certification family — Learn-by-Doing on the GSDC Live Studio with lifetime access.

Element	Detail
Credential	AAIPC · Certified Agentic AI Professional
Issuing body	Global Skill Development Council (GSDC)
Vendor-neutral	Yes · framework-neutral coverage by design
Recognition	100+ countries · 250,000+ certified across the GSDC family
Delivery	GSDC Live Studio + self-paced LMS + daily live sessions
Modules	11 official modules across the 5-chapter journey
Learn-by-Doing labs	50+ hands-on labs with lifetime access
Live sessions	Daily live sessions with industry experts; ask questions in real time
1:1 SME sessions	Personalised 1-on-1 Trainer / SME connect sessions
Capstone	Role-mapped capstone built across the 5-chapter journey
Exam preparation	Two sample exams + mock viva built into the program
Money-back	7-day money-back guarantee per the GSDC refund policy

CADENCE

Inside a typical study week

Most readers anchor a single live session per week and treat everything else as flexible. Five to seven focused hours per week through the 5-chapter journey.

Day	Activity	Time
Monday	Daily live session — concept of the day	45 min
Tuesday	Self-study: read + 1 prompt / tool experiment	30–45 min
Wednesday	Self-study: lab walkthrough on the GSDC Live Studio	45–60 min
Thursday	Office hours / Q&A with SMEs (optional)	30 min
Friday	Apply: ship one LBD artifact	45–60 min
Saturday	Peer review · consolidate notes · flashcards	45 min
Sunday	Light review + plan next week (or rest)	15–20 min

Why this cadence works

- **Daily anchors beat marathon weekends.** Retention is measurably higher with daily contact.
- **Live + LMS hybrid.** Live sessions create accountability; lifetime LMS access gives flexibility.
- **Artifacts over notes.** Each week ends with one LBD deliverable, not a stack of highlights.
- **Peer review is built in.** Most agentic teams ship in peer-reviewed pipelines; the program mirrors that.

VENDOR-NEUTRAL

Vendor-neutral coverage

AAIPC is vendor-neutral by design. The same patterns and skills transfer across the four major agentic frameworks. The matrix below shows where each framework is best suited — and where the program teaches each.

Pattern / capability	LangGraph	CrewAI	AutoGen	OpenAI Assistants
Single-agent loop	Excellent	Excellent	Excellent	Excellent
Supervisor-worker	Excellent	Excellent	Good	Good
Graph orchestration	Excellent	Good	Excellent	Limited
Built-in tool framework	Strong	Strong	Strong	Excellent
Long-running tasks	Excellent	Good	Strong	Strong
Observability hooks	Excellent	Good	Strong	Limited
Cost control primitives	Strong	Good	Strong	Limited

The vendor-neutral promise: learn the *patterns* first; the *framework* is a translation step that takes hours, not weeks. Your skills don't depreciate when the stack changes.

CERTIFICATE

Sample certificate

Below is a stylised preview of the digital certificate issued on completion. The live certificate is dated, individually numbered, and accompanied by a verifiable digital badge you can share on LinkedIn, your CV and email signature.

GLOBAL SKILL DEVELOPMENT COUNCIL

Certificate of Achievement

This is to certify that

[Candidate Name]

has successfully completed the certification examination for

Certified Agentic AI Professional

CREDENTIAL CODE · AAIPC

SAMPLE

GSDC CERTIFIED

[DD / MM / YYYY]
Date of Issue

GSDC Examination Board
Authorized Signature

Verify at gsdcouncil.org/verify · Credential ID: GSDC-AAIPC-XXXXXX

Verification: every certificate is independently verifiable at gsdcouncil.org/verify using the credential ID printed on the certificate.

EMPLOYER VIEW

The employer view in 2026

Hiring teams shipping agentic AI in 2026 don't ask “do you know agents?” They ask “show me a multi-agent system you built, evaluated, governed and deployed.” The 5-chapter journey is the answer.

What hiring teams actually look at

- **A built capstone.** Deployed agent · evals · controls · system card · exec summary.
- **Pattern fluency.** You can pick between single-agent and multi-agent with explicit trade-offs.
- **Eval craft.** You know what trajectory evals are; you've built one against a gold set.
- **Governance instinct.** Capability boundaries, kill-switches, system cards — vocabulary on demand.
- **Vendor-neutrality.** You're not locked into one framework; your skills transfer.

What gets rejected at the screen

- Vague “agent experience” with no system card or eval harness behind it.
- Framework partisanship — “X is the only framework that matters.”
- Treating agents as autonomous teammates without capability boundaries.
- Skipping evals because “the demo worked.”

FAQ

Frequently asked questions

Question	Short answer
Is AAIPC vendor-neutral?	Yes. The program covers patterns first; frameworks (LangGraph, CrewAI, AutoGen, OpenAI Assistants) are translation steps.
Do I need prior AI experience?	Recommended but not mandatory. The 5-chapter journey scaffolds from foundations through ship.
Is the credential globally recognized?	Yes — recognized in 100+ countries; used by 250,000+ certified professionals across the GSDC family.
How long does it realistically take?	Most readers finish in 10–14 weeks at 5–7 focused hours per week.
Are all 50+ LBDs included?	Yes. Every Learn-by-Doing lab is included; the full list is indexed on pages 13–14.
What if I fail the exam?	Retakes are allowed per the GSDC retake policy. Two sample exams and a mock viva are included.
Can my employer sponsor me?	Yes. Corporate / cohort sponsorship is available via the live program page.
What if the program isn't a fit?	GSDC offers a 7-day money-back guarantee per its refund policy.
Do I need to code to ship the capstone?	Light coding helps. The capstone can be built using a low-code agent framework if you prefer.
Does the credential expire?	Lifetime credential · refresh recommended every 2 years as the field moves fast.

GETTING STARTED

Start the journey this week

Four moves to start the 5-chapter journey before you enroll — so Module 1 begins on a running start.

Move	What you do	Time
1 · Pick one use case	Pick one real use case you'd build an agent for. Write it in one sentence.	5 min
2 · Place it on the autonomy spectrum	Where does the use case sit between chat and full autonomy?	10 min
3 · Sketch the perception-reasoning-action loop	On one page. Don't optimise yet — just sketch the loop.	20 min
4 · Tell one peer	Share the sketch with a teammate. Get one piece of feedback before Day 1.	10 min

By the time Module 1 starts, you already have a use case, a placement, a sketch and a peer's feedback. The program then formalises and accelerates what you've already started.

ENROLLMENT

Start the story

You've read the 5-chapter journey, the syllabus, the LBDs and the exam. The next step is starting Chapter 1. Enrol on the live program page; access opens immediately; your cohort is assigned within 24 hours; and Module 1 is waiting in the LMS.

Step	What happens
1 · Click any CTA in this brochure	You land on the official AAIPC program page.
2 · Apply your offer at checkout	Your 50% discount is auto-applied within the offer window.
3 · Complete enrolment	Your access details, cohort schedule and SME introduction arrive in your inbox.
4 · Start Chapter 1	Open the GSDC Live Studio and begin the 5-chapter journey.

Direct contact

Program page: gsdcouncil.org/certification-program/became-a-certified-agentic-ai-professional
Issuing body: Global Skill Development Council (GSDC)
Vendor-neutral: Yes · framework-neutral coverage
Recognition: 100+ countries · 250,000+ certified
Credential code: AAIPC · Certified Agentic AI Professional

Thank you for reading the complete story. We'll see you inside Chapter 1.