

The pathway plan brochure.

The 90-day roadmap, the 4-persona pathway map, the 11-module syllabus, sample exam, and the salary band table per pathway.

Used by 2,50,000+ certified professionals.

4

Persona pathway maps

90Day shared roadmap ·
printable**11**

Module syllabi · verbatim

INSIDE THE GUIDE

What you'll find in this brochure.

A single, opinionated plan that turns "I want to work on agentic AI" into a defended credential and a paying role. Across the next 21 pages, everything below is fully unfolded.

- ✓ **4 persona pathway maps** — Engineer, Architect, Product Lead, Researcher. Each with entry assumptions, exit outcomes, and per-module emphasis.
- ✓ **90-day shared roadmap · printable** — week-by-week milestones common to all four pathways, with self-check gates at day 30, 60, and 90.
- ✓ **11 module syllabi · verbatim** — pulled directly from the live certification handbook, including weighting and capstone deliverables.
- ✓ **Sample exam · printable checklist** — five representative items plus the 12-point go/no-go readiness checklist.
- ✓ **Per-pathway salary band table · USA 2026** — entry, mid, senior, and lead band ranges for each of the four pathways, sourced from cohort exit surveys.

"I followed the Engineer pathway end-to-end. By day 75 I had three live agents in staging and a credential to defend them with."

— Cohort 49, Senior AI Engineer

PATHWAY MAPS · OVERVIEW

4 persona pathway maps.

One certification, four routes through it. Pick the persona closest to your current role; the pathway adjusts emphasis, capstone scope, and assessment weighting accordingly.

PATHWAY 01

The Engineer

You ship code. You want to graduate from chat-completions to autonomous, tool-using agents in production.

Entry assumption: 2+ yrs Python, one shipped LLM feature.

Capstone focus: Build & deploy a tool-using agent.

Exit outcome: Senior AI Engineer / Agent Engineer.

PATHWAY 02

The Architect

You design systems. You need to compare frameworks, defend a stack, and own end-to-end observability.

Entry assumption: 5+ yrs systems design.

Capstone focus: Architecture review of a multi-agent platform.

Exit outcome: Staff AI Architect / Platform Lead.

PATHWAY 03

The Product Lead

You scope and prioritise. You translate "what should an agent do?" into specs engineers can build against.

Entry assumption: 3+ yrs technical PM, AI exposure.

Capstone focus: Agent PRD + eval scorecard.

Exit outcome: Senior PM, AI / Group Product Manager.

PATHWAY 04

The Researcher

You evaluate and publish. You care about reasoning patterns, eval rigour, and reproducibility above all.

Entry assumption: MSc/PhD or equivalent, ML background.

Capstone focus: Comparative eval across reasoning patterns.

Exit outcome: Applied Researcher / Research Engineer.

PATHWAY 01 · THE ENGINEER

The Engineer route.

Module emphasis is tilted toward tools, memory, deployment, and observability. Lighter on planning theory; heavier on shipping.

Module	Standard weight	Engineer weight	Why
M03 · Tool use & function calling	10%	14%	Tools are the engineer's daily surface.
M04 · Memory architectures	9%	12%	Where bugs hide; where cost lives.
M09 · Observability & debugging	8%	12%	Production correctness depends on it.
M10 · Production deployment	9%	13%	The point of the pathway.
M11 · Capstone — ship an agent	10%	12%	Shipped artifact, not a paper.

Engineer capstone: ship a tool-using agent with at least three downstream integrations, a trajectory-level eval harness, deployed observability, and a written rollback plan. Defended live for 30 minutes.

★ RELATED OFFER

Engineer pathway open — earn the credential.

If the route above maps to your role, enrolment is open for the GSDC Agentic AI Certification on the Engineer track. Same exam, persona-weighted capstone.

[Enrol on the Engineer Path →](#)

Persona-matched intake · Engineer track

PATHWAY 02 · THE ARCHITECT

The Architect route.

Module emphasis is tilted toward multi-agent orchestration, safety, and platform-grade observability. Lighter on hands-on lab churn; heavier on trade-off literacy.

Module	Standard weight	Architect weight	Why
M05 · Multi-agent orchestration	10%	14%	Architects own the topology.
M06 · Planning & task decomposition	9%	12%	System-level reasoning trade-offs.
M08 · Safety, guardrails & alignment	10%	13%	Platform liability sits here.
M09 · Observability & debugging	8%	12%	Platform-wide tracing standards.
M11 · Capstone — architecture review	10%	12%	Reviewable design document.

Architect capstone: deliver a 20-page architecture review of a multi-agent platform (real or proposed), including topology, message routing, failure modes, observability strategy, and a framework decision defended against three alternatives.

Where Architect-track graduates land

- ✓ Staff or Principal AI Architect roles owning the agentic platform.
- ✓ Platform Engineering Lead positions where agentic workloads are joining the platform.
- ✓ Internal AI Centre-of-Excellence leads inside larger enterprises.

PATHWAY 03 · THE PRODUCT LEAD

The Product Lead route.

Module emphasis is tilted toward scoping, evaluation, and safety policy. Lighter on framework internals; heavier on what to ask for, when, and how to measure it.

Module	Standard weight	PM weight	Why
M01 · Foundations of agentic AI	6%	10%	Vocabulary you negotiate with.
M07 · Evaluation & benchmarking	11%	15%	The PM owns the scorecard.
M08 · Safety, guardrails & alignment	10%	13%	Policy lives in product, not platform.
M10 · Production deployment	9%	11%	Release decisions are product decisions.
M11 · Capstone — agent PRD	10%	12%	Defensible PRD + eval scorecard.

Product Lead capstone: author a complete agent PRD for a real or hypothetical product, including job-to-be-done, autonomy tier, tool surface, guardrail policy, eval scorecard, rollout plan, and metric tree. Defended with engineering and design proxies.

"The Product Lead pathway is the first time I've seen a curriculum that takes PM-shaped AI work seriously. The eval scorecard module changed how I write specs."

— Cohort 46, Senior PM (AI)

PATHWAY 04 · THE RESEARCHER

The Researcher route.

Module emphasis is tilted toward reasoning patterns and evaluation rigour. Lighter on deployment plumbing; heavier on reproducibility, ablation, and method comparison.

Module	Standard weight	Researcher weight	Why
M02 · Reasoning patterns (ReAct, ToT, etc.)	8%	14%	The research surface itself.
M06 · Planning & task decomposition	9%	12%	Planning-as-research.
M07 · Evaluation & benchmarking	11%	15%	Reproducibility is non-negotiable.
M11 · Capstone — comparative eval	10%	13%	Publication-grade methods report.

Researcher capstone: publish a comparative evaluation report across at least three reasoning patterns on a non-trivial task, with golden dataset, ablation table, error analysis, and a reproducibility checklist. Defended in research-review format.

 LIMITED TIME OFFER

This intake window closes soon.

The current GSDC Agentic AI Certification enrolment is open for a limited window across all four pathways. Once filled, the next cohort rolls to the following quarter.

[Reserve a Pathway Seat →](#)

Time-bound enrolment · All four pathways available

90-DAY ROADMAP · DAYS 1–30

The 90-day shared roadmap.

Every pathway runs on the same calendar. Persona only changes the weighting; the dates do not move.

DAYS 1–7 · ONBOARD**Handbook, baseline assessment, peer intro**

Receive the full handbook, all artifact templates, and the 12-point readiness checklist. Sit a 30-minute baseline assessment so the instructor team can route you to the right peer pod.

DAYS 8–14 · FOUNDATIONS**Modules 01–02 — concepts and reasoning patterns**

Cover the perception-cognition-action loop, autonomy levels, and the three dominant reasoning patterns: ReAct, Reflexion, Tree-of-Thought. Submit Lab 1 (rebuild the same task three ways).

DAYS 15–21 · TOOLS**Module 03 — tool use and function calling**

Tool schemas, idempotency, retry semantics, schema-mismatch handling. Lab 2: ship a single-agent prototype with three tools wired and JSON-schema validation.

DAYS 22–30 · MEMORY · GATE**Module 04 + Day-30 self-check gate**

Memory architectures — episodic, semantic, procedural. Lab 3: layer two memory stores into the Lab-2 agent.

Day-30 gate: 20-question quick check. Score $\geq 75\%$ to proceed; below that, instructor 1:1 to triage.

90-DAY ROADMAP · DAYS 31–60

Month 2 — orchestration and rigour.

DAYS 31–38 · ORCHESTRATION

Module 05 — multi-agent orchestration

Supervisor, hierarchical, peer-to-peer topologies. Trade-offs and failure modes. Lab 4: build a 3-agent research pipeline with shared state and deadlock prevention.

DAYS 39–45 · PLANNING

Module 06 — planning and task decomposition

Hierarchical task networks, LLM-based planning, plan-monitor-replan loops. Workshop: choose reactive vs. planning explicitly for three target tasks.

DAYS 46–53 · EVALUATION

Module 07 — evaluation & benchmarking

Designing scorecards for systems where ground truth is fuzzy. LLM-as-judge patterns, golden datasets, trajectory-level eval. Lab 5: write a working eval harness for your Lab-4 pipeline.

DAYS 54–60 · SAFETY · GATE

Module 08 + Day-60 self-check gate

Safety, guardrails, alignment. Prompt-injection defences, PII redaction, runaway-loop circuit-breakers. **Day-60 gate:** instructor reviews your Lab-5 eval harness and your draft capstone scope. Green-light required to enter month 3.

90-DAY ROADMAP · DAYS 61-90

Month 3 — ship and defend.

DAYS 61-67 · OBSERVABILITY

Module 09 — observability & debugging

Tracing across agent runs, structured logging for non-deterministic flows, replay-based debugging. Wire a working stack (LangSmith / Langfuse / OTel) into your capstone build.

DAYS 68-74 · DEPLOYMENT

Module 10 — production deployment

Containerisation, autoscaling, rate-limits, human-in-the-loop fallbacks, rollback strategies, incident-response runbooks.

DAYS 75-84 · CAPSTONE BUILD

Module 11 — ship your capstone

Build, evaluate, deploy, monitor — end-to-end. Persona-specific deliverable. Submitted with a 6-slide defence deck and a written rollback plan.

DAYS 85-90 · EXAM + DEFENCE

Final exam + live capstone defence

90-minute written exam followed by a 30-minute live capstone defence with two GSDC assessors. Credential issued within 7 working days of passing both.

 50% OFF

Half off across all four pathways — this window only.

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

[Apply the 50% Discount →](#)

Promotional pricing · Applied at registration

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. Capstone: produce a one-page autonomy-tier classification of three real 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 · 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 · 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 defences, PII redaction, runaway-loop circuit-breakers. Maps to NIST AI RMF and the EU AI Act risk categories.

 **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 · 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

Containerisation, 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. Persona-specific scope. Defended in a live review 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 · FORMAT

Sample exam · printable checklist.

90 minutes, 60 multiple-choice items, calibrated to map 1:1 with the 11 modules. A representative subset of items appears on the next page.

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 Defence (oral)	4	8%

SAMPLE EXAM · ITEMS

Three 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 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 03 · 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

📢 LAST FEW SEATS

Cohort closing — final pathway seats available.

Each GSDC cohort caps capstone-defence slots to protect assessor-to-candidate ratios. The current intake is nearly full across all four pathways. 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.

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

- ✓ I have implemented a memory layer with at least two stores (e.g. vector + KV).

- ✓ 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 design choice against at least two alternatives in writing.

- ✓ I have a 6-slide deck ready to defend my capstone.

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

SALARY BAND TABLE · USA 2026

Per-pathway salary bands.

Self-reported base + target bonus from cohort exit surveys, calibrated against three public compensation datasets. Bands are 25th–75th percentile, USD, USA, 2026.

Pathway	Entry	Mid	Senior	Lead / Staff
Engineer	\$120k – \$155k	\$165k – \$210k	\$220k – \$290k	\$310k – \$410k
Architect	\$155k – \$190k	\$200k – \$250k	\$265k – \$340k	\$360k – \$480k
Product Lead	\$140k – \$175k	\$185k – \$235k	\$245k – \$315k	\$330k – \$440k
Researcher	\$135k – \$170k	\$180k – \$230k	\$240k – \$315k	\$325k – \$450k

Methodology: 60-day post-credential survey (n = 1,847 across last 4 cohorts), reconciled against Levels.fyi, Built-In, and Lightcast 2026 medians. Bay Area / NYC / Seattle bands trend toward the upper quartile; remote-friendly roles cluster mid-band.

 **BONUS INCLUDED**

Free bonus: 1-on-1 pathway-selection call.

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

Enrol & Claim the Bonus →

Bonus session · Auto-applied to current intake

HOW THE PATHWAY WORKS

From enrolment to credential.

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

STEP 01 · DAY 0

Pick your pathway, enrol, receive the handbook

Choose Engineer, Architect, Product Lead, or Researcher. Access opens within minutes. Handbook, all templates, the roadmap, and the readiness checklist are delivered to your dashboard.

STEP 02 · DAYS 1–60

Months 1 & 2 — modules 1–8, with Day-30 & Day-60 gates

Self-paced video + reading + labs. Weekly instructor office hours. Two self-check gates ensure you don't carry forward gaps.

STEP 03 · DAYS 61–84

Month 3 — modules 9–11, capstone build

Observability, deployment, then capstone. Pathway-specific deliverable. 6-slide defence deck + written rollback plan.

STEP 04 · DAYS 85–90

Final exam + live defence

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

STEP 05 · WITHIN 7 DAYS

Credential issued

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

FAQ

Quick answers.

Can I switch pathways mid-course?

Once. Up to Day 30. After that, the persona weighting is locked to preserve the integrity of the capstone defence.

Is the credential the same across all four pathways?

Yes — single credential, single registry entry. The pathway is noted on the badge metadata and on your transcript.

How long is the credential valid?

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

What if I fail the exam or the defence?

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 pathway.

The upcoming GSDC Agentic AI Certification cohort is forming now across all four pathways. Enrol before kickoff to receive the handbook and pre-work materials ahead of week 1.

[Join the Next Cohort →](#)

Cohort-based start · Pre-work unlocks at enrolment

GET THE FULL BROCHURE

Download the **pathway plan**.

This 22-page preview is the trailer. The full pathway plan brochure ships with editable templates, the per-week reading list, and the pathway-selection 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

Pathway picked. Now walk it.

Four pathways. One 90-day roadmap. Eleven modules. One exam. One capstone defence. The shortest distance between "I'm interested in agentic AI" and "I've shipped, defended, and earned the credential."

What to do next

- ✓ Re-read pages 3–7 and pick your pathway honestly.
- ✓ Print page 17 and take the 12-point readiness checklist.
- ✓ Tap any CTA in this brochure — every link lands on the pathway page.

 LIFETIME ALUMNI ACCESS

Once certified, you stay in the room.

The GSDC Agentic AI Certification includes lifetime alumni access on every pathway — every syllabus refresh, every new artifact, every cohort discussion, indefinitely.

[Earn the Credential →](#)

Lifetime alumni access · Syllabus updates included

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