THE COMPLETE FIELD GUIDE · 2026

The Complete Field Guide

The full story of generative AI in 2026 — six chapters from “what is it” to “what does responsible deployment look like” — delivered through the GSDC **Certified Generative AI Professional** credential. Plus the 14-module syllabus, all 76 Learn-by-Doing labs, exam details and a sample certificate.

Inside the field guide

✓ **6-chapter story arc**	✓ **14-module syllabus · verbatim**
✓ **All 76 LBDs · full list**	✓ **Exam · validity · career outcomes**
✓ **Sample certificate**	✓ **Printable interview checklist**

6	14	76
Chapters	**Modules**	**LBDs**
Story arc	Verbatim syllabus	With deliverables

Credential: **CGAIP · Certified Generative AI Professional** · Issued by **Global Skill Development Council** · Recognized in 100+ countries.

Page one carries no calls-to-action by design. The story arc begins on page two.

STORY ARC

The 6-chapter journey

Six chapters take you from definitions to deployment. Most readers spend the most time in chapters 2 (how it works) and 4 (business use), and lean hardest on chapter 6 (responsible AI) in interviews.

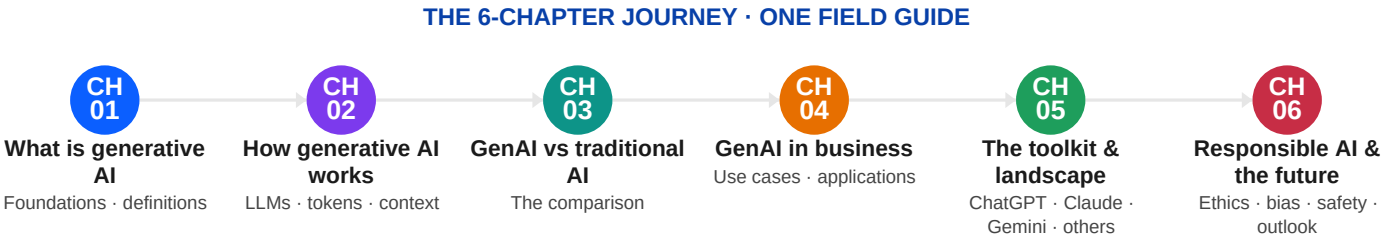

How chapters map to modules: chapters 1–2 = M01–M02 · chapter 3 = M01 cross-cut · chapter 4 = M11 + M13 · chapter 5 = M03–M07 · chapter 6 = M08–M10. The exam draws across all six chapters.

How readers actually use this field guide

- **Beginners** read chapters 1–3 first (pages 5–7) — the foundations.
- **Practitioners with one tool of experience** jump to chapters 4–5 (pages 8–9) — applications and the wider toolkit.
- **Senior or compliance-leaning readers** read chapter 6 first (page 10) — responsible AI is increasingly the differentiator at hiring.
- **Manager / sponsor readers** read the exam, career outcomes and sample certificate (pages 19–21) before deciding to fund a candidate.

READING GUIDE

How to read this field guide

Four blocks. Read the story arc first. Then trace your modules. Then exam, certificate and career outcomes.

Block	Pages	What you'll find
1 · The 6-chapter story arc	5–10	One page each: definition · how it works · vs traditional AI · in business · toolkit · responsible AI.
2 · 14-module syllabus · verbatim	11–12	All 14 modules with sub-topics and LBD counts.
3 · 76 LBDs · full index	13–15	Every Learn-by-Doing activity with named deliverables.
4 · Exam, certificate & career outcomes	19–21	Exam details, validity, sample certificate, employer view.
5 · FAQ, pitfalls & closer	22–28	Common pitfalls, FAQ, getting started, enrollment.

What this field guide does NOT cover

- **Career-by-career mapping.** See the Career Pathway brochure for the 5-career role-to-module matrix.
- **Agentic AI specifically.** See the AAIPC family for multi-step agents, tool surfaces and agent ops.
- **Generative AI in project management.** See the CGAIPM family for PM-specific workflows.

This field guide is the general-purpose CGAIP read. It walks the story arc, anchors the credential, and leaves career and specialism brochures for deeper drilling.

OVERVIEW

Headline numbers · what the credential moves

25–60%	8–12	100+
Reported uplift	**Weeks**	**Countries**
After credentialling	Typical time to certified	Recognition

Three signals that show up in 2026 hiring

- **Capstone artifact.** A working LLM-backed feature with evals and controls evidence — the single largest signal at screen.
- **Responsible-AI fluency.** Disclosure, bias awareness, escalation paths — the recruiter screens these in 2026 regardless of role.
- **Globally recognized credential.** Reduces the “unknown candidate” discount on shortlists across regions.

Three signals readers tell us moved hardest after credentialling

- **Interviews scheduled within 2 weeks** of LinkedIn-profile refresh with the CGAIP badge.
- **Salary review conversation initiated** with a manager within 4 weeks of credential.
- **External offers received** within 8 weeks of the capstone being live on a portfolio link.

CH 01
CHAPTER

What is generative AI

Foundations · definitions · landscape

The story of Chapter 1

Generative AI is the family of models that produce new content — text, images, code, audio, video — instead of just classifying or scoring existing content. The shift from 2022 onwards came from large language models that scaled up well enough to be useful across general tasks, not just narrow ones. This chapter builds the shared vocabulary you'll use across the next 5 chapters.

What you'll learn

- The generative AI vocabulary: model · prompt · completion · context window · token.
- Where generative AI sits in the broader AI / ML landscape.
- Where generative AI shines — and where it doesn't (the “not yet” list).
- The 2022–2026 timeline of capability jumps every PM and engineer should be able to name.

Anchored in modules

M01 · Foundations of generative AI · **M02** (early sections) · LLM capability framing.

What you ship in this chapter (LBD highlights)

- **LBD 01:** Map 30 generative-AI terms in plain language.
- **LBD 02:** Compare 3 generative models with side-by-side outputs on the same prompt.
- **LBD 03:** A one-page rubric: where generative AI fits / where it doesn't.
- **LBD 04:** A landscape note covering the 2022–2026 capability timeline.

CH 02
CHAPTER

How generative AI works

LLMs · tokens · context · structured outputs

The story of Chapter 2

Most readers come to CGAIP with comfort using ChatGPT but not yet a working mental model of *why* it answers the way it does. Chapter 2 closes that gap — tokens, context windows, structured outputs, trace analysis. By the end of the chapter you can debug a bad output rather than just regenerating it.

What you'll learn

- How LLMs tokenize input and why tokenization matters across languages.
- What a context window is, and why your prompts get expensive when you fill it.
- How structured outputs (JSON schema, function calling) tame the “loose text” problem.
- How to read a trace — and how to use the trace to fix prompts faster.

Anchored in modules

M02 · LLMs · capabilities & limits · cross-cuts into **M04** (tool use / function calling).

What you ship in this chapter (LBD highlights)

- **LBD 05:** A context-window planner across 3 model sizes.
- **LBD 06:** Tokenization side-by-side across 3 languages.
- **LBD 07:** A structured-output JSON schema for one of your use cases.
- **LBD 08:** A trace-analysis memo on one real generation.
- **LBD 09:** A “what LLMs won't do well” one-page limits document.

CH 03
CHAPTER

GenAI vs traditional AI

The comparison · when each one wins

The story of Chapter 3

Generative AI didn't replace traditional ML — it joined it. Chapter 3 is the chapter that lets you have a confident hiring-manager conversation about *which one to reach for* on a given problem. Many failed AI projects in 2026 came from picking the wrong family for the problem, not from picking the wrong model inside a family.

Dimension	Traditional AI / ML	Generative AI
Output	Score · class · number	Free-form content
Training	Task-specific · supervised	General · self-supervised + RLHF
Best for	Predicting from clean labeled data	Producing varied content from prompts
Data needs	Labeled training set per task	Often zero-shot or few-shot
Latency	Milliseconds, usually	Hundreds of ms to seconds
Determinism	Deterministic per input	Stochastic by default
Cost shape	Mostly training-time	Mostly inference-time
Evaluation	Accuracy · F1 · AUC	Gold sets · LLM-as-judge · human review

The rule of thumb: if the answer is a number or a class, reach for traditional ML. If the answer is content, a summary, a response or a transformation, reach for generative AI. The hybrid case (predict + explain) is where both meet.

CH 04
CHAPTER

Generative AI in business

Use cases · applications · ROI

The story of Chapter 4

Chapter 4 is the most concrete chapter. Eight use-case patterns dominate generative AI in business in 2026 — and most enterprise programs use a handful of these in combination. This chapter helps you talk through them at the level a hiring manager expects.

#	Use-case pattern	Where you see it
1	Document summarisation & Q&A	Internal knowledge bases · legal · finance · clinical research.
2	Customer-service copilots	Tier-1 support · retail · banking · insurance.
3	Code assistants	Developer productivity · code review · test generation.
4	Content generation	Marketing copy · email · social · localized variants.
5	Data extraction from unstructured sources	Forms · contracts · receipts · medical notes.
6	RAG-powered enterprise search	Internal search · onboarding portals · sales enablement.
7	Translation & localization	Multilingual products · regulatory translations · subtitling.
8	Workflow automation with structured outputs	Ticketing · CRM enrichment · approvals routing.

CH 05
CHAPTER

The toolkit & landscape

ChatGPT · Claude · Gemini · others

The story of Chapter 5

Chapter 5 is the chapter you'll cite the most in interviews. The toolkit changes fast, but the categories are stable. CGAIP teaches vendor-neutral patterns so the tool list evolves while your skills don't have to be re-learned.

Category	Representative tools	What you use it for
General-purpose chat	ChatGPT · Claude · Gemini · Copilot	Drafting · summarisation · Q&A · brainstorming.
Open-weight LLMs	Llama · Mistral · Qwen · Gemma · DeepSeek	Self-hosted deployments · fine-tuning · cost-sensitive use.
Application frameworks	LangChain · LlamaIndex · Semantic Kernel	Building LLM-backed applications · RAG · agents.
Vector stores	Pinecone · Weaviate · Qdrant · pgvector · Chroma	Retrieval · embeddings · semantic search.
Eval platforms	Braintrust · LangSmith · Arize Phoenix · Helicone	Trajectories · gold sets · regression suites.
Image / video generation	Midjourney · DALL-E · Stable Diffusion · Sora · Veo	Creative · marketing · concept work.
Speech / voice	ElevenLabs · OpenAI Voice · Whisper	Voice cloning · TTS · transcription.
Productivity	Notion AI · Microsoft Copilot · Google AI features	Document & meeting AI inside existing tools.

What you'll be screened against: not which tool you know, but whether you can name the *category* and the *trade-off* that picks one tool over another for a given use case.

CH 06
CHAPTER

Responsible AI & the future

Ethics · bias · safety · 2026 outlook

The story of Chapter 6

Chapter 6 is the chapter that increasingly separates candidates at offer stage. Recruiters in 2026 ask three responsible-AI questions in nearly every screen: what would you do about bias? About disclosure? About escalation when the model gets it wrong? Strong answers to these distinguish the strong candidate from the average one.

The four responsible-AI domains you'll learn

Domain	What you learn
1 · Bias & fairness	How bias enters training data · how to audit outputs · how to mitigate.
2 · Privacy & data handling	PII handling · regulatory landscape · data-retention rules · consent & disclosure.
3 · Safety & abuse defence	Prompt-injection defence · output filters · red-teaming · escalation paths.
4 · Hallucination & grounding	Sources of hallucination · grounding via retrieval · citation strategies · audit.

Future of generative AI · what to watch in 2026

- **Multi-modal becomes default.** Text + image + voice in one model, not three.
- **Smaller models for the edge.** Open-weight 1–7B models running on-device.
- **Agentic workflows mainstreaming.** Single-prompt tools give way to multi-step systems (see AAIPC).
- **Regulation tightens unevenly.** EU AI Act enforcement · US state-level action · APAC sector-specific rules.

SYLLABUS

The 14-module syllabus · part 1 of 2

Modules 1–7. Phase, key topics, and the LBD count per module. The full LBD index with named deliverables is on pages 13–15.

#	Module	Phase	Key topics	LBDs
M01	Foundations of generative AI	Foundations	Generative models · LLMs vs traditional ML · history & landscape.	4
M02	LLMs · capabilities & limits	Foundations	Context windows · tokenization · structured outputs · trace analysis.	5
M03	Prompt design & patterns	Core skills	Zero/few-shot · CoT · system prompts · prompt templates · libraries.	6
M04	Tool use, function calling & APIs	Core skills	Function calling · JSON schemas · structured outputs · tool patterns.	5
M05	Retrieval-augmented generation (RAG)	Core skills	Embeddings · vector stores · re-ranking · permission-aware retrieval.	6
M06	Fine-tuning & adapters	Core skills	SFT · LoRA / PEFT · synthetic data · evaluation of fine-tuned models.	5
M07	Eval design · automatic + human	Core skills	Gold sets · eval harnesses · refusal correctness · LLM-as-judge.	6

How the modules thread the 6-chapter story arc

- **Chapter 1 · What is GenAI** — anchored in M01 + early M02.
- **Chapter 2 · How GenAI works** — anchored in M02 + M04 cross-cuts.
- **Chapter 3 · GenAI vs traditional AI** — anchored in M01 cross-cuts.
- **Chapter 4 · GenAI in business** — anchored in M11 (UX) and M13 (strategy).
- **Chapter 5 · The toolkit** — anchored in M03–M07 (the build-side modules).
- **Chapter 6 · Responsible AI** — anchored in M08–M10 (next page).

SYLLABUS

The 14-module syllabus · part 2 of 2

#	Module	Phase	Key topics	LBDs
M08	Hallucination, grounding & truthfulness	Core skills	Sources of hallucination · grounding patterns · citation strategies.	5
M09	Safety, abuse & jailbreak defence	Build & ship	Threat models · prompt-injection defence · output filters · red-teaming.	6
M10	Privacy, IP & data handling	Build & ship	PII handling · IP & copyright · regulatory landscape · retention rules.	5
M11	Product design with AI · UX patterns	Build & ship	AI UX patterns · disclosure design · streaming UX · error recovery.	5
M12	Cost, latency & production ops	Build & ship	Cost budgets · latency budgets · caching · model routing · monitoring.	5
M13	AI strategy, vendor selection & rollout	Build & ship	Vendor evaluation · build-vs-buy · pilot design · scaled rollout.	5
M14	Capstone, exam prep & viva	Capstone	Capstone project · sample exams · viva · interview practice.	4

Total LBDs across the 14 modules: 76. Every LBD ships a named, recruiter-readable artifact — the portfolio recruiters open during the screen.

LBD INDEX

The 76 Learn-by-Doing labs · part 1 of 3

Every LBD ships a recruiter-readable deliverable. Print pages 13–15 as your portfolio checklist.

LBD	M	Activity	Deliverable
01	1	Map 30 generative-AI terms	Glossary
02	1	Compare 3 generative models · output side-by-side	Comparison sheet
03	1	Where generative AI fits / doesn't fit rubric	1-page rubric
04	1	History & landscape note (one page)	Landscape note
05	2	Context-window planner across 3 model sizes	Budget sheet
06	2	Tokenization side-by-side · 3 languages	Token analysis
07	2	Structured output schema design	JSON schema
08	2	Trace analysis on a real generation	Trace analysis memo
09	2	Limits document · “what LLMs won't do well”	Limits doc
10	3	Zero-shot vs few-shot for same task	Comparison memo
11	3	Chain-of-thought prompt experiment	Experiment write-up
12	3	System prompt for a defined persona	System prompt
13	3	Prompt template library (10+ templates)	Template library
14	3	Prompt versioning + changelog	Versioning scheme
15	3	Prompt failure-mode catalogue	Failure catalogue
16	4	Function-calling specification for 3 tools	Tool spec
17	4	JSON-schema output enforcement	Schema + tests
18	4	Read-only tool surface vs write tool surface	Surface design
19	4	Tool-call retry & fallback rules	Rules doc
20	4	Test set for tool-call correctness	Test set
21	5	Embedding model choice with rationale	Selection memo
22	5	Vector store design with metadata schema	Store design
23	5	Hybrid retrieval (BM25 + dense) experiment	Hybrid retrieval
24	5	Re-ranking layer + eval	Re-rank + eval
25	5	Permission-aware retrieval design	Permission spec
26	5	RAG eval set with 30 scenarios	RAG eval set

LBD INDEX

The 76 Learn-by-Doing labs · part 2 of 3

LBD	M	Activity	Deliverable
27	6	Synthetic data set for a real task	Synthetic set
28	6	LoRA / PEFT fine-tune of a small model	Adapter weights
29	6	Fine-tuned vs base model · eval	Eval comparison
30	6	Fine-tuning failure-mode catalogue	Failure catalogue
31	6	Fine-tune cost vs prompt cost · break-even	Break-even memo
32	7	Gold set of 30 representative cases	Gold set
33	7	Automatic eval harness with 6+ metrics	Eval harness
34	7	Human eval rubric + double-blind protocol	Human rubric
35	7	LLM-as-judge eval with self-consistency check	Judge spec
36	7	Refusal-correctness eval on edge cases	Refusal evals
37	7	Grounding eval · citation correctness	Grounding evals
38	8	Catalogue 5 hallucination sources	Hallucination map
39	8	Grounding pattern selection per use case	Pattern rubric
40	8	Citation strategy with audit-friendly format	Citation spec
41	8	Truthfulness audit on one output	Audit note
42	8	Confidence calibration check	Calibration memo
43	9	Threat model for a real feature	Threat model
44	9	Prompt-injection defence test set	Injection test set
45	9	Output filter / moderation layer design	Filter spec
46	9	Red-team transcript on 10 attacks	Red-team report
47	9	Abuse-handling escalation path	Escalation playbook
48	9	Safety policy document	Safety policy
49	10	PII handling rules for your use case	PII rules
50	10	IP & copyright review for outputs	IP review
51	10	Regulatory landscape note (your region)	Regulatory note
52	10	Data-retention rule set	Retention rules
53	10	Consent & disclosure copy for users	Disclosure copy

LBD INDEX

The 76 Learn-by-Doing labs · part 3 of 3

LBD	M	Activity	Deliverable
54	11	Catalogue 5 AI UX patterns with examples	UX pattern library
55	11	Disclosure-design copy for one feature	Disclosure copy
56	11	Loading / streaming UX prototype	Prototype
57	11	Error-recovery UX for an LLM call	Recovery design
58	11	Confidence display experiment	Display variants
59	12	Cost budget spreadsheet for an LLM feature	Budget sheet
60	12	Latency budget with SLOs	Budget + SLOs
61	12	Caching strategy with hit-rate eval	Caching spec
62	12	Model-routing strategy across 3 models	Routing spec
63	12	Monitoring dashboard spec for an LLM feature	Dashboard spec
64	13	Build-vs-buy decision memo on one capability	Decision memo
65	13	Vendor due-diligence framework	DD framework
66	13	Pilot scope & success-metric document	Pilot doc
67	13	Staged rollout plan with feature flags	Rollout plan
68	13	Adoption-tracking dashboard mock	Dashboard mock
69	14	Capstone scope · 1-page brief	Capstone brief
70	14	Capstone build · working artifact	Capstone artifact
71	14	Capstone eval set + results	Eval results
72	14	Capstone exec summary	Exec summary
73	14	Two timed sample exams under exam conditions	Exam attempts
74	14	Mock viva with mentor	Viva transcript
75	14	Interview practice · technical + behavioural	Practice notes
76	14	Resume & LinkedIn refresh against the capstone	CV + LinkedIn

76 LBDs, 76 deliverables. Every artifact a hiring manager can open and verify.

CAPSTONE

Capstone preview · what you ship

The capstone is the artifact you point to in interviews. It's the answer to “show me one thing you actually built.” Below: the canonical capstone shape, with five example shapes readers ship.

Section	What it contains	Time investment
1 · Brief	Problem · users · data · success metrics · risks · controls.	Week 10
2 · Artifact	A built feature or document using ≥3 modules end-to-end.	Week 10–11
3 · Eval	Gold set · automatic evals · human eval rubric · results.	Week 11
4 · Controls	Disclosure · audit trail · escalation path · human gates.	Week 11
5 · Exec summary	One page: problem, what you built, results, controls, next steps.	Week 12
6 · Viva	10-minute mentor defence; rehearses before the real exam.	Week 12

Five capstone shapes readers have shipped

- **A RAG-backed internal search** with grounding evals · 3 tools · 60 scenarios.
- **A streaming customer-service chatbot** with citation and audit trail.
- **A 30-prompt library** with refusal-correctness and grounding eval sets.
- **A document-extraction pipeline** with JSON-schema enforcement and human review.
- **A use-case brief + autonomy rationale + disclosure** for a feature shipped in a real product.

CAREER OUTCOMES

Where this credential takes you

CGAIP is the general-purpose generative AI credential. It opens five distinct career doors. The Career Pathway brochure covers the role-to-module mapping in depth; the table below is the short version.

#	Career	Most-used module emphasis
C01	AI Engineer	M01–M10, M12, M14 — deep on build-side modules.
C02	Prompt Engineer	M01–M03, M05, M07–M09 — deep on prompts and evals.
C03	LLM Developer	M01–M10, M12 — deep on application-layer modules.
C04	AI Product Manager	M01–M03, M07–M14 — deep on product, controls and strategy.
C05	AI Consultant	M01–M03, M08–M11, M13 — deep on strategy, controls and rollout.

Typical career-outcome signals after the credential

- **Recruiter interest within 1–2 weeks** of the CGAIP badge going live on LinkedIn.
- **First interview slot within 3–4 weeks** for entry-level readers with a working capstone.
- **External offer within 8 weeks** is common for mid-career readers leveraging the capstone link.
- **Internal re-band within 3 months** reported by readers who stay in their current firm.

The credential alone doesn't land jobs. The credential *plus* the capstone *plus* the LinkedIn refresh *plus* targeted outreach is the combination readers consistently report works.

VALIDITY

Validity, recognition & recertification

What happens after you pass the exam. The credential is lifetime-valid, with an industry-recommended refresh cycle as generative AI evolves.

Item	Detail
Credential code	CGAIP · Certified Generative AI Professional
Issuing body	Global Skill Development Council (GSDC)
Validity	Lifetime credential · no expiry
Refresh recommendation	Every 2 years · short refresh module on what's changed
Recognition	100+ countries · 250,000+ certified community
Verification	gsdcouncil.org/verify · independent verification of credential ID
Digital badge	Shareable on LinkedIn, CV, email signature, personal site
Continuing community	Free GSDC Membership · access to certified community
Career boosters	LinkedIn Enhancer · Resume Builder · interview practice platform

What the recertification refresh covers

- **New capability jumps** in foundation models since your original credential.
- **New regulatory developments** in your region (e.g. EU AI Act updates).
- **New tools and patterns** that have become standard since your original credential.
- **New evaluation methods** as eval craft matures.

EXAM

Exam details

The CGAIP exam tests application across the 14 modules and 76 LBDs. The structure mirrors real practitioner decisions, not vocabulary recitation.

Item	Detail
Credential code	CGAIP · Certified Generative AI Professional
Format	Multiple-choice + scenario items + applied mini-case
Items	Approx. 60–80 items
Duration	Approx. 90–120 minutes
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
Money-back	7-day money-back guarantee per the GSDC refund policy

How the chapters weight the exam

- **Chapters 2 + 5 (how it works + the toolkit):** heaviest weighting at ~40% combined.
- **Chapter 6 (responsible AI):** ~25% weighting; growing in 2026.
- **Chapter 4 (business use cases):** ~15% weighting; tested via mini-case.
- **Chapters 1 + 3 (foundations + vs traditional AI):** ~20% weighting; mostly conceptual.

EXAM PREP

Sample exam questions

Q1. A team wants to reduce hallucinations in an internal Q&A system. Which is the **most** appropriate **first** intervention?

A. Switch to a larger model.

B. Add a retrieval layer that grounds answers in a controlled document store.

C. Lower the temperature to 0.

D. Add a system prompt instructing the model not to hallucinate.

Answer: B · *Grounding via retrieval is the standard first move on internal Q&A.*

Q2. When evaluating a prompt change, which control gives you the **most** defensible signal?

A. One-off side-by-side comparison of outputs.

B. A gold set with refusal-correctness, grounding and task-completion evals.

C. Reviewer's subjective “feels better”.

D. Letting the model self-rate.

Answer: B · *Multi-axis gold-set evals are the audit-grade control.*

Q3. Which is the **most** appropriate **first** control for prompt-injection risk in an LLM-backed customer-service tool?

A. Disable the tool entirely.

B. Validate and sanitize user input; treat retrieved content as untrusted; enforce output filters.

C. Add a single prompt instructing the model to ignore injection attempts.

D. Increase the temperature so outputs are varied.

Answer: B · *Defence-in-depth is the standard pattern; single instructions are bypassable.*

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.

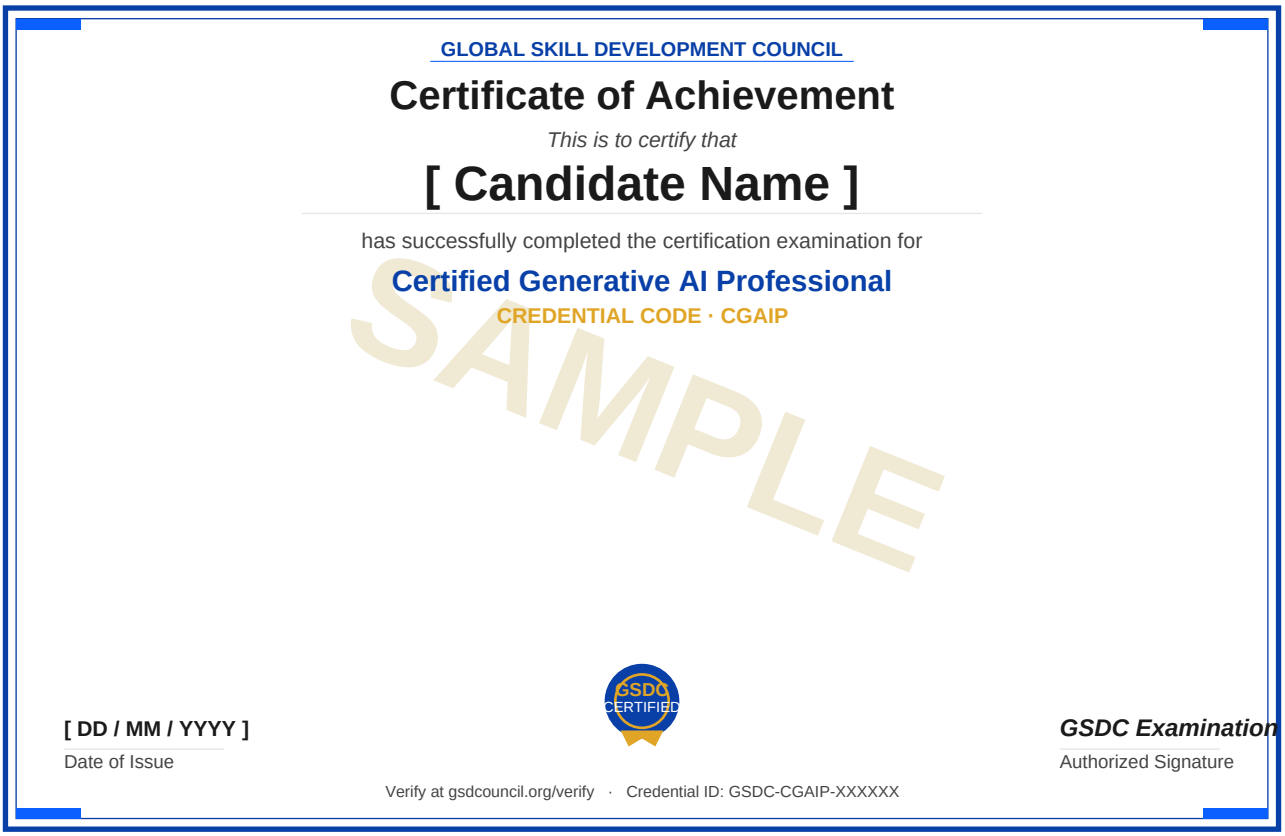

Verification: every certificate is independently verifiable at gsdcouncil.org/verify using the credential ID printed on the certificate.

CHECKLIST

Printable chapter checklist

Print this page. Tick the boxes as you complete each chapter and the LBD highlights. Six chapters; one set of milestones per chapter.

Chapter 1 · What is generative AI

- ☐ I can define generative AI in 60 seconds without notes.
- ☐ I can name 3 things generative AI does well and 3 it doesn't.
- ☐ I can sketch the 2022–2026 capability timeline by heart.

Chapter 2 · How it works

- ☐ I can explain tokenization and context windows to a non-technical stakeholder.
- ☐ I have shipped one structured-output JSON schema.
- ☐ I have read and explained one trace from a real generation.

Chapter 3 · vs traditional AI

- ☐ I can name 3 problems where traditional ML wins and 3 where generative AI wins.
- ☐ I have a worked hybrid example (predict + explain) I can describe.

Chapter 4 · in business

- ☐ I can name 5 of the 8 business use-case patterns from page 8.
- ☐ I have ROI math for at least one of them in my own domain.

Chapter 5 · the toolkit

- ☐ I can name 8 tool categories and the trade-off picking one over another.
- ☐ I have hands-on experience with at least 2 tools per category I'd use.

Chapter 6 · responsible AI

- ☐ I can describe a bias-mitigation strategy on a real use case.
- ☐ I have one-page disclosure copy ready for a real feature.
- ☐ I can describe a prompt-injection defence in 3 specific layers.

PITFALLS

Common pitfalls in a generative AI journey

Eight pitfalls account for the majority of “I read the field guide but didn't land anything.” Memorize them — knowing each one in advance is half the protection.

Pitfall	How to avoid it
Reading without shipping	Pair every chapter with one LBD. By the end of 6 chapters you have 12+ artifacts.
Tool partisanship	Pick patterns, not vendors. CGAIP teaches vendor-neutral patterns by design.
Skipping responsible AI	Chapter 6 is increasingly the interview differentiator. Don't treat it as optional.
No capstone	Ship one. Even a small one. “Course completed” loses to “I built X”.
No eval evidence	Always mention your gold set, your eval harness and your refusal-correctness work.
Treating AI as “autonomous”	The human gate is non-negotiable. The PM owns the output.
Letting AI generate numbers	Numbers come from the system of record. The model frames; the data store states.
Not naming the credential	“CGAIP · Certified Generative AI Professional” should appear in your LinkedIn headline.

EMPLOYER VIEW

The employer view in 2026

Hiring teams across enterprise, consulting, startups, banks and fintech screen against a consistent shape, even when their JDs read very differently. CGAIP moves the needle *before* the interview — at the stages where most candidates get filtered out.

Where CGAIP moves the needle first

Stage	How CGAIP moves the needle
Recruiter screen	A globally recognized credential reduces the “unknown candidate” discount.
ATS filter	CGAIP + canonical keywords lifts you past automated scans.
First call	“Tell me about your generative AI work” has a credible answer: the capstone.
Technical interview	Eval mindset and responsible-AI fluency shine in the hiring-manager rubric.
Late-stage / panel	The 76-LBD portfolio gives you 5+ named artifacts to defend at panel.
Offer / band placement	Salary conversation anchored against demonstrated capability, not gut figures.

RESPONSIBLE AI INCLUDED

[50% OFF]

Ethics, bias and safety are core, not optional

Chapter 6 covers responsible-AI practices, including bias, safety, disclosure and audit. The exam tests it. Recruiters look for it.

Get Responsible-AI Ready →

Tap anywhere on this card →

FAQ

Frequently asked questions

Question	Short answer
Do I need prior production-AI experience?	No. CGAIP is designed for both upskillers and beginners. The capstone is the substitute.
Do I need to code to do this credential?	Some roles (engineer / developer) require Python. Product / consulting roles don't.
How long does CGAIP realistically take?	Typically 8–12 weeks at 5–7 focused hours per week.
What about the toolkit — will it be out of date in 6 months?	Tools change; patterns don't. CGAIP teaches vendor-neutral patterns that transfer.
Is the credential globally recognized?	Yes — recognized in 100+ countries; 250,000+ certified community.
What if I'm just curious — not job-hunting?	The field guide is the best read for that. The credential is optional but recognized everywhere.
How does this compare to AAIPC (agentic) or CGAIPM (PM)?	CGAIP is the foundation. AAIPC and CGAIPM are specialism credentials. Many readers do CGAIP first.
What if I fail the exam?	Retakes are allowed per the GSDC retake policy. Two sample exams + mock viva 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.

FAMILY

Sister credentials in the GSDC AI family

CGAIP is the foundation credential. Several specialism credentials sit alongside it — many readers stack them over time as their career evolves.

Credential	What it adds	Typical reader
CGAIP · Certified Generative AI Professional	The foundation: 14 modules, 76 LBDs, 6-chapter story arc (this brochure).	Anyone working with or around generative AI.
AAIPC · Certified Agentic AI Professional	Specialism: multi-step agents, tool surfaces, agent ops, agent safety.	Engineers / PMs working on autonomous agentic systems.
CGAIPM · Certified Generative AI in Project Management	Specialism: how PMs run projects with generative AI in the loop.	Project managers in any function.
CGAIFB · Generative AI in Finance & Banking	Specialism: agentic finance, RAG for compliance, banking-specific use cases.	Finance / banking professionals.

The typical stacking pattern: CGAIP first, then one specialism within 6–12 months. The credentials are independently valuable but compound when stacked.

GETTING STARTED

How to start this week

Four moves to take this week — before you enroll — to line up your CGAIP pathway.

Move	What you do	Time
1 · Pick your chapter focus	Read all 6 chapter pages (5–10). Mark the chapter you feel weakest on.	30 min
2 · Sketch your capstone	Pick one capstone shape from page 16. Adapt to your domain in 5 sentences.	15 min
3 · Update LinkedIn headline	Add “CGAIP candidate · [your career focus]” to your headline.	10 min
4 · Block your weekly study slot	Block 5–7 hours per week for the CGAIP pathway over 8–12 weeks.	5 min

By the time you start Module 1, you already know your weak chapter, your capstone shape and your weekly anchor. The program accelerates everything else.

ENROLLMENT

Walk the field guide for real

You've read the six chapters, the syllabus, the LBDs and the exam. The next step is starting Module 1. Access opens immediately; your cohort is assigned within 24 hours; Module 1 is waiting in the LMS.

Step	What happens
1 · Click any CTA in this field guide	You land on the official story / CGAIP 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 mentor introduction arrive in your inbox.
4 · Start Chapter 1 / Module 1	Open the GSDC Live Studio and begin the 8–12 week pathway.

Direct contact

Program page: gsdcouncil.org/certification-program/the-story-of-generative-ai
Issuing body: Global Skill Development Council (GSDC)
Recognition: 100+ countries · 250,000+ certified
Credential code: CGAIP · Certified Generative AI Professional

Thank you for reading the field guide. We'll see you inside Chapter 1.