

[DOWNLOAD THE TOOLKIT BROCHURE](#)

Master the 30+ AI Tools Employers Actually Use in Production.

The full catalog of generative AI tools, all 76 Learn-by-Doing activities, and the complete 14-module GSDC syllabus — in one downloadable PDF.

TOOL FAMILIES COVERED

LLM APIs

Orchestration

Vector DBs

RAG stacks

Image gen

Fine-tuning

Serving

Eval

30+

AI Tools

14

Modules

76

LBDs

12

Weeks

Inside this toolkit

- Complete 30+ generative AI tools catalog
- All 76 LBDs and portfolio deliverables
- 14-module syllabus from GSDC
- Sample certificate and exam details

WELCOME

Tools, not theory

Employers in 2026 do not screen candidates on conceptual knowledge of generative AI. They screen on hands-on fluency with the exact tools production teams use every day — LLM APIs, orchestration libraries, vector stores, evaluation harnesses, image-gen pipelines, and the deployment stack around them. This brochure is the complete map of those tools, the order in which to learn them, and the 76 build-along activities that turn each one into a portfolio artifact.

Use it as a study plan, a hiring-readiness checklist, and a reference when comparing the GSDC certification against generic AI courses. Every tool listed here is one we have observed in current production environments at companies actively hiring.

Built on real listings

Tool selection comes from a rolling scan of generative-AI job posts, weighted by appearance frequency in technical interviews.

Hands-on by design

Every tool ships with at least one Learn-by-Doing activity — a small, scoped build that becomes a public portfolio artifact.

The tool families this brochure covers

LLM APIs

Orchestration

Vector DBs

RAG stacks

Image gen

Fine-tuning

Serving

Evaluation

WHAT YOU'LL GET

Inside the 28-page toolkit

This is a working document. Read it cover to cover before enrolling, then keep it open as your reference during the 12-week programme.

- **30+ tool catalog**
Every tool, what it does, where it fits, and which module covers it.
- **All 76 LBDs**
Learn-by-Doing activities with brief, success criteria, and portfolio artifact.
- **14-module syllabus**
Topics, outcomes, recommended tools, and reading list per module.
- **Sample certificate + exam details**
Format, scoring band, retake policy, digital badge specification.
- **Job-ready outcomes**
What graduates ship, how their portfolios are reviewed, where they end up.
- **Enrolment + FAQs**
Four-step enrolment, refund policy, retake policy, mentor pod structure.

TOOL LANDSCAPE

How 30+ tools split into 8 families

Production generative-AI work pulls from eight tool families. You don't need to master every tool in every family — but you do need fluency in at least one production-grade option from each. The catalog ahead lists representative tools per family, with notes on when each is the right pick.

LLM APIs

[OpenAI](#), [Anthropic](#), [Google](#), [Mistral](#), [Cohere](#)

Hosted models exposed over HTTP. The default starting point for any GenAI feature.

Orchestration

[LangChain](#), [LlamaIndex](#), [Haystack](#), [Semantic Kernel](#)

Framework layer for chaining LLM calls, tools, retrieval, and memory.

Vector databases

[Pinecone](#), [Weaviate](#), [Qdrant](#), [Chroma](#), [pgvector](#)

Storage and similarity search over embedding vectors for RAG and search.

RAG stacks

[Verba](#), [RAGFlow](#), [Llama-index RAG](#), [custom hybrid pipelines](#)

Opinionated retrieval-augmented-generation pipelines built on top of the above.

Image generation

[Stable Diffusion](#), [FLUX](#), [Midjourney API](#), [DALL-E API](#)

Image and visual-asset generation, including fine-tuning via LoRA / DreamBooth.

Fine-tuning

[Hugging Face TRL](#), [Unsloth](#), [Axolotl](#), [OpenAI fine-tune API](#)

SFT, LoRA/QLoRA, DPO, and instruction-tuning workflows for open-weights models.

Serving + ops

[vLLM](#), [TGI](#), [Ollama](#), [Modal](#), [Replicate](#), [BentoML](#)

Inference servers, autoscaling, and observability for production deployments.

Evaluation

[Ragas](#), [DeepEval](#), [Promptfoo](#), [Braintrust](#), [LangSmith](#)

Offline + online eval harnesses, LLM-as-judge, and human review pipelines.

[TAKE THE NEXT STEP](#)

From toolkit to credential

You've just seen the eight tool families the brochure covers. The credential is what ties them together — a 12-week programme that walks you through every tool listed here, with mentor review, LBDs, and a portfolio you can show recruiters at the end.

[OFFER] · RELATED

Enrol in the Master GenAI Tools Certification

Same 30+ tool catalog you're reading about — now with mentor support, LBDs, and the GSDC credential.

[Start My Certification →](#)

Tap anywhere →

30+

Tools

14

Modules

76

LBDs

12

Weeks



Job ready

Tap the card above, or visit gsdcouncil.org · Master Generative AI Tools

[TOOL CATALOG · 1 OF 4](#)

LLM APIs & orchestration

OpenAI API

M2 · M4

GPT-4 / GPT-4o / o-series. The reference implementation for chat, function calls, structured output, vision, and the Responses API.

Anthropic Claude API

M2 · M5

Claude models with long context windows and strong tool-use. Reliable for agentic workflows and document-heavy RAG.

Google Gemini API

M2

Gemini Pro / Flash. Native multimodal input and built-in grounding via Google Search for fact-current queries.

LangChain

M3 · M5

Python / JS framework for chaining LLM calls, prompts, tools, memory, and agents. Default orchestration layer for most teams.

LlamaIndex

M3 · M4

RAG-first orchestration: data connectors, index abstractions, query engines, and evaluation hooks for retrieval pipelines.

Semantic Kernel

M5

Microsoft's orchestration SDK (C# / Python / Java) for plugging LLMs into existing enterprise applications.

[TOOL CATALOG · 2 OF 4](#)

Vector databases & RAG stacks

Pinecone

M4

Managed vector database. Easy ramp-up, generous free tier, strong for early production RAG before self-hosting becomes worth it.

Weaviate

M4

Open-source vector DB with built-in hybrid search (BM25 + vectors) and a strong module ecosystem for re-ranking.

Qdrant

M4

Rust-based vector DB known for performance and a clean Python client. Excellent for self-hosted RAG at scale.

Chroma

M3 · M4

Lightweight, developer-friendly vector store. The go-to for local prototyping and first-pass RAG before scaling out.

pgvector (Postgres)

M4

Vector extension for Postgres. The right call when your RAG data already lives in Postgres and you want one less moving part.

LlamaIndex RAG

M4 · M6

Opinionated end-to-end RAG pipelines: ingestion, indexing, retrieval, and evaluation, all wired together with sensible defaults.

[TOOL CATALOG · 3 OF 4](#)

Image generation & fine-tuning

Stable Diffusion (SDXL / SD3)

M7

Open-weights image generation. The reference stack for self-hosted image work, LoRA training, and ControlNet pipelines.

FLUX

M7

Newer open-weights family with strong prompt adherence and clean text rendering. Often paired with SD for variation work.

DALL-E / Midjourney APIs

M7

Hosted image generation. Lower friction than self-hosted, useful when brand consistency or compute cost matters.

Hugging Face TRL

M8

Reference library for SFT, DPO, PPO, and reward modelling on open-weights LLMs. The most-used training stack in 2026.

Unsloth

M8

Memory-efficient fine-tuning kernels. The right pick for LoRA / QLoRA work on a single GPU without compromising quality.

Axolotl

M8

Config-driven fine-tuning wrapper. Best when you need reproducible runs across a team rather than one-off notebooks.

[DON'T LET THE COHORT CLOSE](#)

Mentor-pod seats are filling

The Master Generative AI Tools cohorts run on a rolling basis, but mentor-pod slots are capped per intake. Reserving your seat now preserves access to the LBD review queue and the post-certification employer connect — both are cohort-bounded resources.

[OFFER] · LIMITED TIME

Lock your seat in the current intake

Mentor pods and employer-connect slots are capped per cohort. Reserve yours before this window closes.

[Secure My Seat Now →](#)

Tap anywhere →

30+

Tools

14

Modules

76

LBDs

12

Weeks



Job ready

Tap the card above, or visit gsdcouncil.org · Master Generative AI Tools

[TOOL CATALOG · 4 OF 4](#)

Serving, ops & evaluation

vLLM

M9

High-throughput inference server with paged attention and continuous batching. Default for self-hosted open-weights serving.

Text Generation Inference (TGI)

M9

Hugging Face's production inference server, used heavily on managed and self-hosted deployments alike.

Ollama

M9

Local LLM runner. Indispensable for laptop-grade development and offline prototyping before pushing to a real server.

LangSmith

M6

Tracing, eval, and prompt management for LangChain-based applications. The easiest way to add observability to a RAG stack.

Ragas

M6

Open-source RAG evaluation library. Faithfulness, answer relevance, context precision/recall — all out-of-the-box.

Promptfoo

M6

CLI-driven prompt evaluation harness. Run hundreds of prompt variants against graded test sets locally or in CI.

TOOL-TO-MODULE MAP

Where each tool family is covered

Use this matrix to see which modules introduce each tool family, which modules deepen it, and which modules treat it as optional context. Core coverage is where the tool is taught hands-on; reinforced is where it shows up again in a different scenario.

- Core (hands-on)
- Reinforced
- Context only

Tool family	M1-2	M3-4	M5-6	M7-8	M9-10	M11-14
LLM APIs	C	R	R	R	R	R
Orchestration	O	C	C	R	R	R
Vector DBs	O	C	R	O	R	R
RAG stacks	O	C	C	R	R	C
Image gen	O	O	O	C	O	R
Fine-tuning	O	O	O	C	R	R
Serving + ops	O	O	O	R	C	C
Evaluation	R	R	C	R	R	C

14-MODULE SYLLABUS · 1 OF 2

Modules M1 – M7

M1**Generative AI Foundations**

Transformer basics, capabilities and limits of modern LLMs, environment setup.

M2**Working with LLM APIs**

OpenAI, Anthropic, Gemini APIs. Function calls, structured output, streaming.

M3**Embeddings & Vector Search**

Embedding models, similarity, hybrid search, Chroma / pgvector / Pinecone basics.

M4**Retrieval-Augmented Generation**

Chunking, hybrid retrieval, re-ranking, RAG eval with Ragas and LangSmith.

M5**Agents & Orchestration**

LangChain / LlamaIndex agents, tools, memory, multi-step planners, controllability.

M6**Evaluation in Practice**

Promptfoo + Ragas pipelines, LLM-as-judge, human review, eval in CI.

M7**Image Generation**

Stable Diffusion + FLUX workflows, LoRA training, ControlNet, hosted-API tradeoffs.

[HALF-PRICE ACCESS](#)

The certification at 50% off

GSDC opens a recurring promotional window where the Master Generative AI Tools Certification is available at half the standard fee. The promotional rate covers the full 12-week pathway, all 76 LBDs, mentor pods, and the post-certification employer connect — nothing is removed at the lower rate.

[OFFER] · 50% OFF

Get the full tool catalog at half the standard rate

Same 14 modules, same 76 LBDs, same mentor and employer-connect package — at 50% off.

[Claim 50% Off Now →](#)

Tap anywhere →

30+

Tools

14

Modules

76

LBDs

12

Weeks



Job ready

Tap the card above, or visit gsdcouncil.org · Master Generative AI Tools

14-MODULE SYLLABUS · 2 OF 2

Modules M8 – M14

M8**Fine-tuning Workflows**

TRL, Unsloth, Axolotl. SFT, LoRA / QLoRA, DPO, dataset curation, eval after tuning.

M9**Production Serving**

vLLM, TGI, Ollama. Streaming, batching, quantization, autoscaling, observability.

M10**Guardrails & Safety**

Input/output filtering, jailbreak resistance, sensitive-data masking, policy.

M11**AI Workflow Automation**

n8n + Zapier + custom orchestration. Building internal AI tooling and copilots.

M12**Cost & Unit Economics**

Per-request cost models, caching tactics, model-tier routing, hosted-vs-open math.

M13**AI Product Patterns**

Common UI/UX patterns for GenAI features: chat, search, copilots, agents, eval-in-UI.

M14**Capstone Project**

End-to-end shipped GenAI feature with full eval and write-up, presented to mentors.

LEARN-BY-DOING ACTIVITIES

All 76 LBDs across the toolkit

Each module ships with a set of Learn-by-Doing activities — small, scoped builds that produce a reviewable portfolio artifact. They are graded by a mentor and are the backbone of your post-certification employability story.



TOTAL · 76 Learn-by-Doing activities

[SAMPLE LBD](#)

What one Learn-by-Doing activity looks like

LBD M4-3 · BUILD A HYBRID RAG PIPELINE WITH RE-RANKING

Brief

Build a question-answering service over a 200-document corpus that combines Weaviate hybrid search (BM25 + vectors) with a cross-encoder re-ranker, exposes a streaming HTTP endpoint, and ships with a Ragas eval harness covering 30 gold-set questions.

Success criteria

Retrieval quality	Recall@5 \geq 0.80 on the gold set
Re-rank lift	\geq 12% NDCG improvement vs. BM25-only baseline
Latency	P95 end-to-end \leq 1.8s on a 4-CPU host
Faithfulness	Ragas faithfulness score \geq 0.85
Code quality	Tests pass; clean retrieve / rerank / generate separation
Write-up	1-page architecture decision record (ADR)

Portfolio artifact

GitHub repository with reproducible run instructions, a short Loom-style demo (3 minutes), and the ADR PDF. Reviewable by a hiring manager in under 10 minutes.

TIME-BOUND WINDOW

Your fast-track offer is valid for 48 hours

If you are reading this brochure as part of an outreach campaign, your enrolment window is time-bound — typically 48 hours from when the document was shared with you. Within that window, your mentor-pod slot is held without cost and the promotional rate is applied automatically at checkout.

[OFFER] · VALID 48 HOURS

Activate your reserved seat in the next 48 hours

Your seat hold and promotional rate are auto-applied at checkout. After the window closes, both reset.

Activate My Seat →

Tap anywhere →

30+

Tools

14

Modules

76

LBDs

12

Weeks



Job ready

Tap the card above, or visit gsdcouncil.org · Master Generative AI Tools

PORTFOLIO DELIVERABLES

What you actually ship by the end

Beyond passing the exam, every graduate exits with a small but high-signal portfolio. These are the six artifacts hiring managers consistently spend time on:

Multi-provider LLM playground

Single UI that hits OpenAI, Anthropic, and Gemini side-by-side with shared eval grid. Tools used: OpenAI API · Anthropic API · Gemini API · LangChain

Production-grade RAG service

Streaming HTTP API with hybrid search, re-ranking, Ragas eval, and observability. Tools used: Weaviate / Qdrant · LangChain · Ragas · LangSmith

Agent with real tools

Multi-step agent that uses three external tools (search, calculator, internal API) with retry, memory, and trace UI. Tools used: LangChain / LlamaIndex · LangSmith

Fine-tuned image LoRA

Custom-style LoRA trained on a small dataset with reproducible config and gallery. Tools used: Stable Diffusion / FLUX · Diffusers · Unsloth

Fine-tuned LLM adapter

QLoRA-tuned 7B model with eval-before / eval-after report. Tools used: TRL · Unsloth · Axolotl · Hugging Face Hub

Capstone end-to-end app

A shipped GenAI feature: chosen problem, chosen tools, evaluation in CI, deployed and demo-able. Tools: any from the catalog.

EXAM DETAILS

How the certification exam works

The exam is the final gate, in addition to LBD review. It is proctored, online, and designed to be completed in a single sitting.

FORMAT

60 multiple-choice + 4 scenario-based questions

DURATION

90 minutes, single sitting

PASS MARK

65% overall, with a 50% floor on scenario items

PROCTORING

Browser-based with live invigilation

RETAKE POLICY

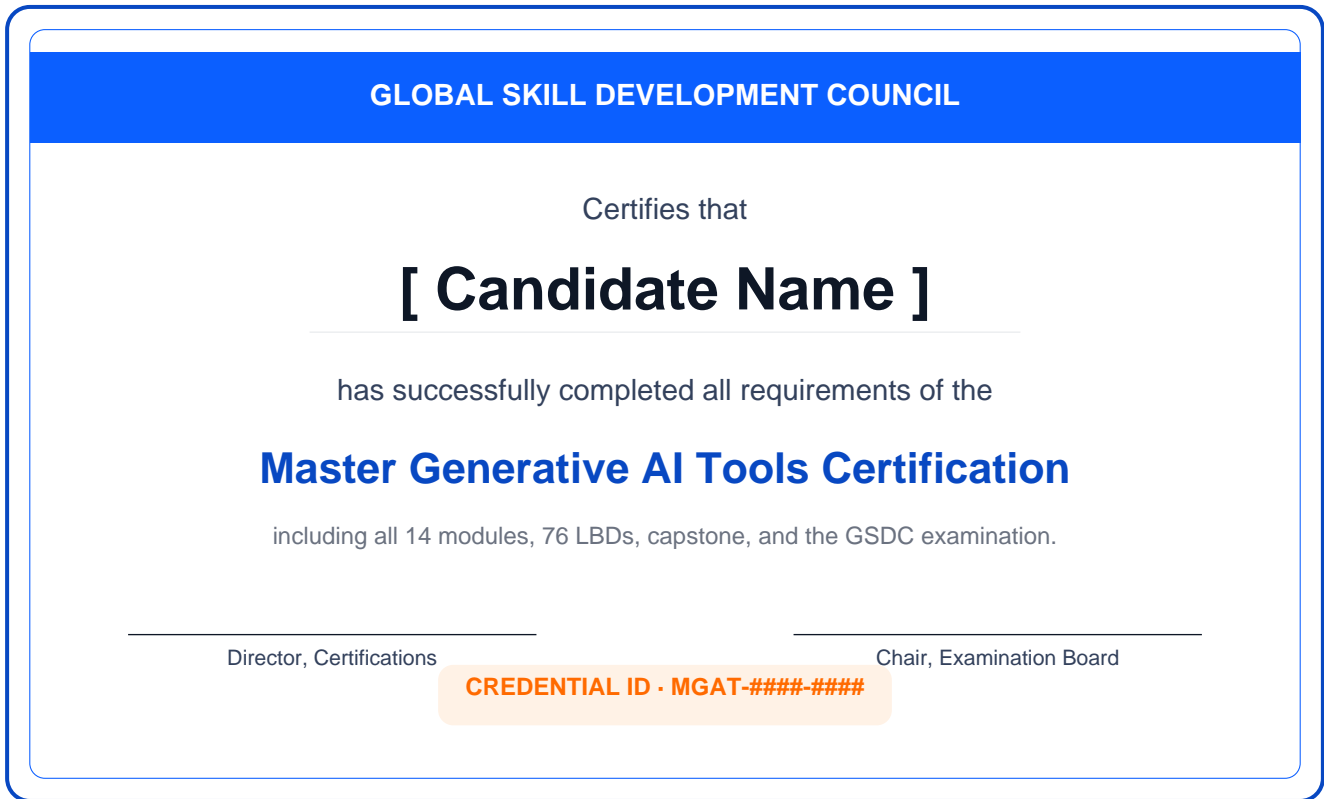
Two free retakes within 12 months

SCORING BAND

Pass / Merit / Distinction ($\geq 75\%$ / $\geq 85\%$)

SAMPLE CERTIFICATE

What you receive on passing



Each credential is paired with a verifiable digital badge and a public verification URL hosted on the GSDC credentialing system.

HALF OFF, FULL TOOLKIT

Half-price access · second window

If you missed the earlier promotional rate, a second 50%-off window is currently open for candidates who reach this page of the brochure. It's the lowest publicly available pricing on the Master Generative AI Tools credential and includes everything described in this document.

[OFFER] · 50% OFF

Half-price certification — full tool catalog, no add-ons removed

Includes the full syllabus, all 76 LBDs, mentor pods, capstone review, and the employer-connect package.

[Apply 50% Discount →](#)

Tap anywhere →

30+

Tools

14

Modules

76

LBDs

12

Weeks



Job ready

Tap the card above, or visit gsdcouncil.org · Master Generative AI Tools

12-WEEK PATHWAY

How your 12 weeks are structured

The pathway is structured into four phases of three weeks each. Each phase ends in a checkpoint — a graded LBD bundle — so progress is visible long before the final exam.

PHASE 1 Weeks 1-3	APIs + Embeddings MODULES IN THIS PHASE M1 Foundations M2 LLM APIs M3 Embeddings
PHASE 2 Weeks 4-6	RAG + Agents + Eval MODULES IN THIS PHASE M4 RAG M5 Agents M6 Evaluation
PHASE 3 Weeks 7-9	Image + Fine-tune + Serve MODULES IN THIS PHASE M7 Image gen M8 Fine-tuning M9 Serving
PHASE 4 Weeks 10-12	Guardrails, Product, Capstone MODULES IN THIS PHASE M10 Guardrails M11-13 Product / Cost / Automation M14 Capstone

OUTCOMES

What graduates do after the certification

Outcomes are tracked at 30, 60, and 90 days post-certification. The figures below come from the 2025 graduating cohorts that opted into outcome tracking.

82%

report a measurable role or scope upgrade within 90 days

3.4x

more interview invitations vs. their pre-cert baseline

68%

complete at least one warm employer introduction

94%

complete the capstone and earn the credential

Where graduates typically end up

- **Internal AI-team pivot** — Move into an AI-focused team at the same employer.
- **New-employer offer** — External offer for AI engineering, prompt, or LLM-developer roles.
- **Founding / early start-up** — Founding-engineer at an early-stage AI start-up or independent consulting.
- **Side-project to product** — Take a capstone artifact into a real product or open-source release.

[PAIR THE BROCHURE WITH THE CREDENTIAL](#)

The natural companion to this document

This brochure is the planning artifact. The credential is the execution artifact. Candidates who succeed almost always pair the two — the brochure tells them which tools to learn and in what order; the credential gives them the structure, mentors, and evidence to actually master them.

[OFFER] · RELATED

Pair this brochure with the GSDC credential

Same 30+ tools, same mentor pods, same LBDs and employer connect — packaged into one structured 12-week credential.

[Get the Credential →](#)

Tap anywhere →

30+

Tools

14

Modules

76

LBDs

12

Weeks



Job ready

Tap the card above, or visit gsdcouncil.org · Master Generative AI Tools

MENTORS & COMMUNITY

Who you'll learn from, and with

Every mentor pod is led by a practising senior GenAI engineer or product leader, vetted by GSDC's accreditation board. The community layer — Slack, weekly office hours, and a private LinkedIn network — stays active long after the credential is awarded.

Mentor pods

Cohorts are organised into pods of 8–10 candidates plus one lead mentor. Pods meet weekly for two live hours: one for code/LBD review, one for the current module's tools and discussion. Pods stay together for the full 12 weeks.

Community access

Lifetime access to the GSDC Alumni Slack and quarterly virtual meetups. Alumni regularly post role openings from their own teams — about 40% of post-cert offers in 2025 came through the alumni channel.

Mentor bench at a glance

120+ Active mentors

9 Avg years senior

18 Industries

4.8/5 Mentor rating

FINAL 48 HOURS

The fast-track window closes soon

Candidates who reach this section of the brochure typically convert into enrolled learners within the next 48 hours, because that is the window in which their reserved mentor-pod slot is held. After 48 hours the slot is released back into the general intake queue and the discount auto-applied to your account expires.

[OFFER] · VALID 48 HOURS

Convert your reservation into an enrolment

Your held pod slot and auto-applied discount expire after 48 hours. Confirm to keep both.

Confirm My Enrolment →

Tap anywhere →

30+

Tools

14

Modules

76

LBDs

12

Weeks



Job ready

Tap the card above, or visit gsdcouncil.org · Master Generative AI Tools

FINAL STEPS

How to enrol - quick FAQs

Enrolment in 4 steps

- 1 Review the tool-to-module map (page 11) to confirm fit with your goals.
- 2 Reserve your seat — tap any CTA in this brochure to open the form.
- 3 Onboarding call within 24 hours of payment to confirm pod and schedule.
- 4 Week 1 unlocks on the cohort start date with M1 lectures and LBDs.

Quick FAQs

Q. Do I need a CS degree?

A. No. Comfort with one programming language is enough; M1 covers the rest.

Q. Will I get API access during the programme?

A. We provide credits for OpenAI, Anthropic, and Hugging Face for the duration of the cohort.

Q. Is the credential industry-recognised?

A. Yes — listed on the GSDC public registry with a verifiable digital badge.

Q. What if I miss the 12-week window?

A. The dashboard stays open 12 months. Only mentor-pod access is time-bounded.

Q. Are retakes allowed?

A. Two free retakes within 12 months — no effect on the final credential.

Q. Is there a refund policy?

A. Full refund within 7 days; partial after that per the published schedule.

[END OF THE BROCHURE](#)

From toolkit to credential — start today

You have now seen the full 30+ tool catalog, the tool-to-module map, the 14-module syllabus, the 76 LBDs, the portfolio you'll ship, the exam and certificate, the 12-week pathway, the mentor structure, and the outcomes data. The final step is yours — take it while the cohort, the reserved slot, and the promotional rate are still in play in the same window.

[OFFER] · RELATED

Begin your Master Generative AI Tools credential now

Everything in this 28-page brochure ships under one credential. Start where you finished reading.

[Start My Certification →](#)

Tap anywhere →

30+

Tools

14

Modules

76

LBDs

12

Weeks



Job ready

Tap the card above, or visit gsdcouncil.org · Master Generative AI Tools

Thank you for reading. We'll see you in the cohort.