

BROCHURE · 2026 EDITION

The Full 90-Day Pathway & Certification Brochure

Generative AI in Finance & Banking — the complete certification pathway as a downloadable PDF: 14 modules, 4 phases, role-mapped career roadmap, and sample exam questions.

Inside the toolkit

✓ 90-day weekly study plan	✓ 4 role-mapped pathways
✓ 14 official module syllabi	✓ Sample exam questions
✓ Salary band reference	✓ 28 pages, brochure-only

28
Pages
Brochure-only

14
Modules
4 phases

90
Days
Weekly plan

Credential code: **CGAIFB** · Issued by **Global Skill Development Council** · Recognized in 100+ countries · Aligned with employer expectations at JPMorgan, HSBC, Goldman Sachs, Deloitte and peers.

This is a brochure. Page 1 contains no calls-to-action by design — keep reading; the next pages walk you through the entire 90-day pathway, the four role tracks, and the official module map.

WELCOME

Why this brochure exists

You are reading the long-form brochure of the GSDC **Generative AI in Finance & Banking (CGAIFB)** certification — the credential employers screen for in AI finance jobs and AI banking jobs in 2026.

Hiring managers at global banks, fintech firms and consulting practices increasingly filter shortlists by AI fluency, not by traditional finance background alone. The CGAIFB pathway compresses that fluency into a structured 90-day journey: 14 modules, 4 phases, role-mapped capstones, daily live sessions, and a globally recognized digital badge.

What you will find inside these 28 pages

- **Pages 3–6** — Program overview, who it's for, outcomes, and the four phases.
- **Pages 7–10** — The 90-day weekly study plan with daily anchors and time budgets.
- **Pages 11–14** — The four role-mapped career pathways with capstone briefs.
- **Pages 15–22** — The 14 official module syllabi, learning outcomes and labs.
- **Pages 23–25** — Sample exam questions with rationales and study tips.
- **Pages 26–27** — Salary band reference, employer signals and ROI math.
- **Page 28** — Enrollment, contact and how to lock in your seat.

How to use this brochure: read top-to-bottom on first pass to understand the shape of the pathway, then revisit individual modules and pathways when you start studying. Every page links back to the live program page so you can enroll the moment you're ready.

OVERVIEW

Program at a glance

Credential	Generative AI in Finance & Banking (CGAIFB)
Issuing body	Global Skill Development Council (GSDC)
Pathway length	90 days, weekly study cadence
Structure	4 phases · 14 modules · role-mapped capstone
Total effort	~36+ hours of structured learning + labs + exam prep
Live sessions	Daily mentor-led live sessions (recordings retained)
Delivery	Self-paced LMS + cohort + live mentor office hours
Exam format	Multiple-choice + scenario items + applied case
Validity	Lifetime credential · refresh recommended every 2 years
Recognition	Recognized in 100+ countries
Audience	Finance, banking, fintech, risk, compliance, quant, audit
Pre-requisites	Working knowledge of finance/banking basics; coding is optional

Key outcomes after 90 days

- You can scope, design and govern generative AI use cases inside a regulated finance/banking environment.
- You can read, critique and contribute to model risk, validation and audit conversations involving LLMs.
- You can prototype assistants for KYC, fraud, credit, treasury, compliance, research and customer ops.
- You can pass employer screens that explicitly ask for AI fluency in finance and banking roles.
- You hold a globally recognized credential and a role-mapped capstone you can show in interviews.

AUDIENCE

Who this certification is for

Role today	Why CGAIFB fits	Where it takes you
Financial / business analyst	Layer generative AI on top of FP&A, reporting and research workflows.	AI Financial Analyst · AI Research Analyst
Banking operations / KYC analyst	Automate document review, KYC enrichment, sanctions and exception handling.	AI Ops Specialist · KYC Automation Lead
Risk / model risk professional	Govern, validate and audit LLM-driven models alongside traditional ones.	AI Risk Analyst · Model Validation Lead
Compliance / audit officer	Build defensible controls for genAI use cases and review AI-generated evidence.	AI Compliance Officer · Audit Lead — AI
Quant / data scientist	Combine classical quant techniques with retrieval, agents and structured outputs.	Quant Analyst (AI) · ML Engineer — Finance
Product / fintech PM	Ship customer-facing AI features safely inside regulated stacks.	Fintech AI PM · AI Product Lead — Banking
Treasury / corporate finance	Apply generative AI to forecasting, cash, scenario modelling and policy drafting.	AI Treasury Analyst · FP&A; AI Lead
Career switcher / fresh graduate	Pair a domain primer with AI fluency to enter regulated finance/banking tech.	AI Banking Associate · Entry-Level Fintech AI

Not sure if you fit?

If you work in or around money — credit, payments, trading, insurance, audit, treasury, wealth, compliance — and your team is being asked the question *“how are we using AI?”*, this credential equips you to be the answer.

RELATED · NEXT STEP

[50% OFF]

Lock in your CGAIFB seat at half the price

Take the next logical step after this section. Enroll now and start the 90-day pathway with mentor access from day one.

Enroll & Save 50% →

Tap anywhere on this card →

STRUCTURE

The four phases of the 90-day pathway

The pathway is split into four progressive phases. Each phase has a clear theme, a set of modules, and a measurable outcome you can show in interviews.

Phase	Days	Theme	Outcome
Phase 1 — Foundations	Day 1–20	Generative AI literacy + finance & banking context.	You can explain LLMs, prompts, embeddings and risk vocabulary to a CFO.
Phase 2 — Applied building	Day 21–50	Prompting, retrieval, agents, structured outputs in regulated workflows.	You ship two prototypes that solve real finance/banking problems.
Phase 3 — Governance & risk	Day 51–75	Model risk, controls, audit, ethics, regulation, data protection.	You can lead an AI governance and model-risk conversation.
Phase 4 — Capstone & exam	Day 76–90	Role-mapped capstone + exam prep + sample exams + viva.	You sit and pass the CGAIFB exam, with a portfolio capstone.

How the phases reinforce each other

- **Foundations** gives you shared vocabulary — so building doesn't feel like guessing.
- **Applied building** grounds vocabulary in artifacts — prompts, agents, retrieval flows.
- **Governance & risk** wraps your artifacts in defensible controls — what banks demand.
- **Capstone & exam** compresses everything into a portfolio and a recognized credential.

STRUCTURE · DETAIL

Phase deep-dive

PHASE 1 · FOUNDATIONS · Days 1–20

- Modules 1–3: AI/LLM literacy, finance & banking context, data fundamentals.
- Daily 30–45 min self-study + 2 live sessions per week.
- Mini-quizzes after each topic to lock in vocabulary.

PHASE 2 · APPLIED BUILDING · Days 21–50

- Modules 4–7: prompting patterns, retrieval, agents, structured outputs.
- Two graded prototypes: one in lending/credit, one in operations/KYC.
- Peer review rubric: usefulness, safety, traceability, cost.

PHASE 3 · GOVERNANCE & RISK · Days 51–75

- Modules 8–11: model risk, audit, ethics, regulation, data protection.
- Workshop: write a model card and a control narrative for an LLM use case.
- Mock regulator interview with a mentor.

PHASE 4 · CAPSTONE & EXAM · Days 76–90

- Modules 12–14: capstone, integration, exam prep.
- Role-mapped capstone with a written brief and 10-minute viva.
- Two sample exams + answer-key walkthroughs + final exam booking.

WEEKLY PLAN

The 90-day weekly study plan — Month 1

Five to seven focused hours per week is enough. Plan one weekday anchor session, two shorter weekday touches, and a weekend consolidation. The plan below assumes you start on a Monday.

Week	Focus	Daily anchors	Deliverable
W1 · Days 1–7	AI & LLM literacy for finance	Read 1 concept → 1 prompt experiment → 1 reflection note	Glossary of 30 AI terms in your own words
W2 · Days 8–14	Banking & finance data fundamentals	Map one workflow you own to data inputs & outputs	Workflow map (1 page) — current vs AI-assisted
W3 · Days 15–21	Prompting patterns & chains	Practice 2 prompt patterns/day with finance prompts	Prompt library (20 prompts, finance themed)
W4 · Days 22–28	Retrieval & grounded answers	Build a tiny RAG demo on policy docs you have	RAG mini-prototype + 1-page note on failure modes

Time budget tips for Month 1

- Treat the live session as your fixed anchor; everything else is flexible.
- Don't aim for perfection in Week 1 — aim for vocabulary fluency.
- Keep a running 'questions' doc; bring them to weekly office hours.

WEEKLY PLAN · MONTH 2

Month 2 — Applied building

Week	Focus	Daily anchors	Deliverable
W5 · Days 29–35	Agents & tool use	Read one agent pattern → implement on a finance task	Agent design doc for one role workflow
W6 · Days 36–42	Structured outputs & integrations	Practice JSON-mode prompts on real schemas	Validator + 3 schemas you ship internally
W7 · Days 43–49	Use cases in lending, fraud, KYC, trading ops	Pick one domain & build a vertical prototype	Prototype #1 with eval notes
W8 · Days 50–56	Evaluations, cost & latency in regulated stacks	Build a small eval set; measure quality vs cost	Eval scorecard for prototype #1

Why Month 2 is where careers actually move

By the end of Week 8 you will have **shipped something** — a vertical prototype with an evaluation scorecard. Hiring managers read this signal differently from a generic AI course completion. It says: *you have built inside a regulated context and you understand the trade-offs.*

LIMITED-TIME OFFER

[LIMITED TIME]

Cohort enrollment closing soon

Seats for the upcoming CGAIFB cohort are filling fast. Reserve your spot before the doors close on this enrollment window.

[Reserve My Seat →](#)

Tap anywhere on this card →

WEEKLY PLAN · MONTH 3

Month 3 — Governance, capstone, exam

Week	Focus	Daily anchors	Deliverable
W9 · Days 57–63	Model risk & validation for LLM-based systems	One MRM artifact/day: model card, control, test	Mini-MRM file on prototype #1
W10 · Days 64–70	Ethics, fairness, data protection & regulation	Map your prototype to 5 regulatory touchpoints	Regulatory mapping table (1 page)
W11 · Days 71–77	Capstone scoping & design	Lock pathway → scope → architecture	Capstone brief approved by mentor
W12 · Days 78–84	Capstone build & review	Build → measure → write up → peer review	Capstone artifact + 1-page exec summary
W13 · Day 85–90	Exam prep & final exam	Two sample exams + flashcard refresh + viva	Exam booked & cleared · credential issued

By Day 90 you have a recognized credential, a capstone artifact tied to a real role pathway, and a mock viva under your belt. That's the package employers screen for.

WEEKLY CADENCE

Inside a typical study week

Day	Activity	Time
Monday	Live mentor session — concept of the week	60 min
Tuesday	Self-study: read + 1 prompt experiment	30–45 min
Wednesday	Self-study: lab walkthrough on the LMS	45–60 min
Thursday	Office hours / Q&A drop-in (optional but recommended)	30 min
Friday	Apply: build / refactor a small artifact for your role	45–60 min
Saturday	Consolidation: notes, flashcards, peer review	45 min
Sunday	Light review + plan next week (or rest)	15–20 min

Design choices we made (and why)

- **Daily anchors, not marathon weekends.** Retention is higher when you touch the material daily.
- **Live + LMS hybrid.** Live sessions create accountability; LMS gives flexibility for shifts/timezones.
- **Artifacts over note-taking.** Every week ends with something you can show, not just something you read.
- **Peer review baked in.** Most banks work in peer-reviewed pipelines; the program mirrors that.

CAREER ROADMAP

Four role-mapped pathways

Phases and modules are the same for everyone. What changes is the **role pathway** — your labs, capstone and exam scenarios all bend toward the role you want to be hired into next. Pick one of the four pathways below at the end of Week 4.

Pathway	Best fit for	Capstone style
A · Fintech AI Specialist	Engineers, PMs, builders in fintech/banking tech	Production-style AI feature with eval harness
B · AI Risk & Compliance	Risk, MRM, compliance, audit, second-line teams	Model card + control narrative + regulator viva
C · Quant & AI Investment Analyst	Quants, research analysts, portfolio & trading desks	Research workflow with retrieval + structured outputs
D · AI Banking Operations	Ops, KYC, payments, customer service, back-office	End-to-end workflow automation with controls

All four pathways unlock the **same CGAIFB credential**. The pathway choice tailors your portfolio so that interviewers see exactly what they're hiring for.

PATHWAY A

Pathway A — Fintech AI Specialist

For people who want to **build** AI features inside banking and fintech stacks. You'll spend more lab time on prompting patterns, retrieval, agents, evaluations, and shipping safely.

- Capstone brief: an internal copilot for a banking workflow (credit, payments, treasury or KYC).
- Skills emphasized: prompt design, retrieval, structured outputs, evals, observability.
- Job titles unlocked: Fintech AI Specialist · AI Engineer (Finance) · AI Product Engineer.
- Salary band reference: see page 26.

Pathway B — AI Risk & Compliance

For people who govern, validate or audit models — including LLM-driven ones. This pathway leans into model risk management, regulation, ethics and defensible controls.

- Capstone brief: a full MRM file for a generative AI use case, with viva.
- Skills emphasized: model cards, control narratives, regulatory mapping, audit evidence.
- Job titles unlocked: AI Risk Analyst · AI Compliance Officer · Model Validation Lead · AI Auditor.
- Salary band reference: see page 26.

HALF-PRICE CERTIFICATION

[50% OFF]

Flat 50% off — full certification, full access

Same 14 modules, same exam, same globally recognized credential — at half the cost. Apply your discount at checkout.

[Claim My 50% Discount →](#)

Tap anywhere on this card →

PATHWAY C

Pathway C — Quant & AI Investment Analyst

For quants and research analysts who want to **fuse** classical methods with generative AI. More lab time on research retrieval, structured outputs against schemas, scenario modelling.

- Capstone brief: AI-assisted equity/credit/macro research workflow with grounded citations.
- Skills emphasized: retrieval over filings & research, structured analyst outputs, eval rigor.
- Job titles unlocked: Quant Analyst (AI) · AI Investment Analyst · Research Analyst (AI).
- Salary band reference: see page 26.

Pathway D — AI Banking Operations

For operations, KYC, fraud, payments, customer service and back-office teams who want to **automate** document-heavy and exception-heavy workflows responsibly.

- Capstone brief: an end-to-end workflow automation (e.g. KYC enrichment, fraud triage, dispute ops).
- Skills emphasized: document understanding, classification, agents, controls and audit trails.
- Job titles unlocked: AI Banking Ops Lead · KYC Automation Lead · Fraud Analyst (AI) · AI Ops Specialist.
- Salary band reference: see page 26.

DECISION GUIDE

Choosing your pathway

If you're between two pathways, use this lightweight matrix. There are no wrong choices — the credential is the same — but matching your pathway to your next-12-months goal sharpens interview signal.

If you say...	Pick	Why
"I want to build AI features in fintech / banking apps."	Pathway A	Production-style capstone
"I want to lead AI risk, validation or compliance."	Pathway B	MRM file + regulator viva
"I'm a quant or research analyst."	Pathway C	Research workflow capstone
"I run or automate banking operations."	Pathway D	End-to-end workflow capstone
"I'm a career switcher / new graduate."	A or D	Most teach-back, lowest pre-req lift
"I'm a regulator or policy person."	B (with C electives)	Frames work in defensible language

Lock-in: You officially choose your pathway at the end of Week 4. You can switch once with mentor approval before Week 8 — beyond that, your capstone is anchored to the pathway.

MODULES · OVERVIEW

The 14 official module syllabi

All 14 modules are required to sit the exam. Pathway choice shapes the labs and capstone, not the module list. Each module bundles concepts, labs, an exam-style quiz and a mentor checkpoint.

#	Module	Phase
1	Generative AI & LLM literacy for finance	Foundations
2	Finance & banking context for AI work	Foundations
3	Data foundations: structured, unstructured, market data	Foundations
4	Prompting patterns & prompt engineering for finance	Applied
5	Retrieval-augmented generation in regulated stacks	Applied
6	Agents & tool use for banking workflows	Applied
7	Structured outputs, integrations & evaluations	Applied
8	Model risk & validation for LLM systems	Governance
9	Ethics, fairness & responsible AI in finance	Governance
10	Data protection, security & privacy for genAI	Governance
11	Regulation & supervisory expectations	Governance
12	Role-mapped capstone design & build	Capstone
13	Integration, cost, latency & operations	Capstone
14	Exam preparation, sample exams & viva	Capstone

MODULE SYLLABI

Modules 1–3 - Foundations

MODULE 1 | Generative AI & LLM literacy for finance

Learning outcomes

- Explain LLMs, tokens, context windows and embeddings in finance terms.
- Distinguish chat, completion, function-calling and agentic patterns.

Labs & artifacts

- Glossary of 30 AI terms in your own finance vocabulary, plus a workflow map.

MODULE 2 | Finance & banking context for AI work

Learning outcomes

- Map functions across retail, corporate, investment banking, insurance and fintech.
- Connect risk taxonomies (credit, market, ops, model) to AI use cases.

Labs & artifacts

- One-page workflow map and a list of 5 controls your AI use case must satisfy.

MODULE 3 | Data foundations: structured, unstructured, market data

Learning outcomes

- Classify financial data into structured, semi-structured, unstructured and market.
- Recognize data quality, lineage and access constraints in regulated stacks.

Labs & artifacts

- Data inventory for your role workflow + retrieval design for 2 document types.

OFFER VALID FOR 48 HOURS

[48 HOURS ONLY]

Your 50% offer expires in 48 hours

This bracket of the offer window is open for 48 hours from the moment you opened this brochure. Don't let the window close.

[Activate Offer Now →](#)

Tap anywhere on this card →

MODULE SYLLABI

Modules 4–5 - Applied building

MODULE 4 Prompting patterns & prompt engineering for finance

Learning outcomes

- Use core patterns: instruction, few-shot, chain-of-thought, self-critique, tool-use.
- Design prompts for risk-aware financial reasoning (with citations and limits).
- Manage hallucinations, jailbreaks and prompt injection in regulated contexts.

Labs & artifacts

- Build a finance prompt library (20+ prompts).
- Run a controlled experiment: same task, three patterns, score the outputs.
- Write a one-page anti-jailbreak checklist for your team.

MODULE 5 Retrieval-augmented generation in regulated stacks

Learning outcomes

- Architect a RAG pipeline (chunking, embedding, store, retriever, reranker).
- Choose retrieval depth, citation style and refusal behavior for finance.
- Evaluate retrieval quality with a small but rigorous eval set.

Labs & artifacts

- Mini-RAG prototype on a real policy / KYC / research corpus.
- Eval scorecard: top-k accuracy, citation faithfulness, refusal correctness.
- Write up: failure modes you found & how you would mitigate them.

MODULE SYLLABI

Modules 6–7 - Applied building

MODULE 6 Agents & tool use for banking workflows

Learning outcomes

- Distinguish single-step tool use from multi-step agents.
- Design safe action surfaces: read-only vs read/write tools in banking systems.
- Apply guardrails: budgets, timeouts, approval steps, audit logs.

Labs & artifacts

- Agent design doc for one role workflow (e.g. dispute triage, KYC enrichment).
- Implement a single-step tool-use agent with strict approvals.
- List 5 things this agent must **never** do, with justification.

MODULE 7 Structured outputs, integrations & evaluations

Learning outcomes

- Use JSON-mode / function calling to produce schema-bound outputs.
- Integrate AI outputs with downstream systems (RPA, BI, core banking).
- Build evaluation harnesses: gold sets, rubrics, judge models, cost/latency.

Labs & artifacts

- Validator + 3 schemas tied to real internal artifacts.
- Eval harness with a 30–50 item gold set and cost/latency scorecard.

MODULE SYLLABI

Modules 8–9 - Governance & risk

MODULE 8 Model risk & validation for LLM systems

Learning outcomes

- Extend traditional model risk management to LLM-driven systems.
- Author a model card and a control narrative your second line will accept.
- Design validation tests: replicability, robustness, fairness, drift.

Labs & artifacts

- Mini-MRM file on your prototype: card, controls, tests, sign-offs.
- Stress test plan: adversarial prompts + drift signals.

MODULE 9 Ethics, fairness & responsible AI in finance

Learning outcomes

- Apply fairness lenses to lending, fraud, KYC and pricing decisions.
- Detect and document bias in data, model and deployment layers.
- Translate ethical concerns into operational controls.

Labs & artifacts

- Bias audit on a real or sample finance dataset/use case.
- Ethics statement + escalation playbook for your use case.

MODULE SYLLABI

Modules 10–11 · Governance & risk

MODULE 10 Data protection, security & privacy for genAI

Learning outcomes

- Map data flows in a genAI app: prompts, retrieval, logs, outputs.
- Apply privacy techniques: minimization, masking, tenancy, retention.
- Address security: prompt injection, exfiltration, supply chain, secrets.

Labs & artifacts

- Data flow diagram for your prototype with control points.
- Security checklist: 15+ items mapped to OWASP & banking practice.

MODULE 11 Regulation & supervisory expectations

Learning outcomes

- Read & reference major regulatory and supervisory expectations on AI.
- Map your use case to relevant regimes (model risk, data protection, sectoral).
- Anticipate questions a regulator or auditor will ask of your AI use case.

Labs & artifacts

- Regulatory mapping table for your capstone use case.
- Mock regulator interview with your mentor — 15 minutes.

UNLOCK FULL ACCESS

[50% OFF]

Unlock all 14 modules, labs & exam attempts

Self-paced LMS + live mentor cohorts + sample exams + recognized digital badge. Everything you need to certify.

[Unlock Full Access →](#)

Tap anywhere on this card →

MODULE SYLLABI

Modules 12–13 - Capstone & integration

MODULE 12 Role-mapped capstone design & build

Learning outcomes

- Scope a problem worth solving in your role with clear success metrics.
- Design the technical, control and change-management layers together.
- Build, measure, write up and present your capstone artifact.

Labs & artifacts

- Capstone brief approved by your mentor by end of Week 11.
- Capstone artifact: prototype + eval scorecard + 1-page exec summary.
- 10-minute viva with mentor and peer panel.

MODULE 13 Integration, cost, latency & operations

Learning outcomes

- Plan integration with existing finance/banking systems and identity stacks.
- Budget for cost, latency, observability and on-call from day one.
- Plan rollout: pilot, controlled expansion, kill-switch criteria.

Labs & artifacts

- Integration architecture sketch with control points.
- Cost & latency budget for your capstone use case.
- Rollout plan with kill-switch criteria.

MODULE SYLLABI

Module 14 · Exam preparation & exam mechanics

MODULE 14 Exam preparation, sample exams & viva

Learning outcomes

- Consolidate concepts from all 13 prior modules with structured revision.
- Practice exam timing & question-type strategy.
- Walk through the official sample exam with answer key & rationale.

Labs & artifacts

- Two sample exams completed under timed conditions.
- Flashcard set covering all 14 modules' must-know terms.
- Mock viva with a mentor.

How the exam works

Item	Detail
Format	Multiple-choice + scenario items + applied mini-case
Items	Approx. 60–80 items in total
Duration	Approx. 90 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 GSDC retake policy (see live page)
Credential issued	Digital certificate & verifiable badge on completion

EXAM PREP

Sample exam questions (1 of 3)

These items mirror the style of the real CGAIFB exam. They are illustrative only and are not scored. Each question is followed by the correct answer and a short rationale.

Q1. A retail bank wants to use an LLM to draft customer dispute responses. Which control is most important to add **before** turning the feature on?

- A. Switch to the largest possible model for higher quality.
- B. Require human-in-the-loop review of every drafted response prior to send.
- C. Cache all responses to reduce inference cost.
- D. Disable retrieval and let the model freely generate from memory.

Answer: B · Customer communications are high-impact, externally visible artifacts. A human-in-the-loop review is the minimum defensible control before any go-live; cost and model size are secondary considerations.

Q2. Which of the following is the **best** definition of retrieval-augmented generation (RAG) in a regulated banking context?

- A. A method to fine-tune a model on private banking data.
- B. A technique to compress prompts to reduce token usage.
- C. An architecture where the model answers using documents retrieved at query time, ideally with citations.
- D. A way to bypass content filters by retrieving alternative phrasings.

Answer: C · RAG grounds the model in retrieved documents at query time. In a regulated context, citations and refusal behavior are critical for traceability.

Q3. A model risk team is reviewing an LLM-driven credit memo assistant. Which artifact is **most** useful for them to assess the system?

- A. A screenshot of the chat UI.
- B. A model card with intended use, limitations, evaluations and controls.
- C. A list of all prompts ever sent in production.
- D. Vendor marketing collateral about the model provider.

Answer: B · Model cards (with intended use, limitations, eval results and controls) are the standard artifact for model risk review.

CAREER FAST-TRACK

[50% OFF]

Fast-track AI finance & banking roles

Built for fintech AI, quant, risk, compliance and investment analyst tracks at JPMorgan, HSBC, Goldman, Deloitte and peers.

[Fast-Track My Career →](#)

Tap anywhere on this card →

EXAM PREP

Sample exam questions (2 of 3)

Q4. Which is the **most** defensible response when an LLM-based KYC assistant is asked a question whose answer it cannot ground in retrieved documents?

- A. Answer based on general knowledge to be helpful.
- B. Refuse and explain that the requested information is not in the supporting documents.
- C. Generate the most likely answer and add a confidence score.
- D. Forward the query to the customer without any response.

Answer: B · *In regulated finance, ungrounded answers create audit and customer-harm risk. Refusing with a clear reason is the safer, defensible behavior.*

Q5. You are designing an agent that can read a customer's account but should never **move** money. Which design choice is essential?

- A. Use a more capable model.
- B. Allow the agent direct write access for speed.
- C. Restrict the tool surface to read-only APIs and require human approval for any write.
- D. Add a longer system prompt warning the agent not to move money.

Answer: C · *Capabilities should be enforced at the **tool surface**, not just in prompts. Read-only tools + human approval for writes is the standard pattern.*

Q6. Which of the following is the **best** example of a structured output use case in investment research?

- A. A free-form essay summarizing an earnings call.
- B. A JSON object with fields: ticker, period, key_metrics, sentiment, citations.
- C. A chat history of analyst conversations.
- D. A scanned PDF of a research note.

Answer: B · *Structured outputs (typed fields) integrate cleanly with downstream systems and are easier to evaluate and audit.*

Q7. Which is the **strongest** signal that your RAG pipeline has a citation faithfulness problem?

- A. Users report the answers feel helpful.
- B. The model cites documents that, on inspection, do not contain the claimed facts.
- C. Retrieval latency is higher than 1 second.
- D. Cost per query is rising.

Answer: B · *Citation faithfulness fails when cited documents do not actually support the claim. This is a top-priority defect in regulated settings.*

EXAM PREP

Sample exam questions (3 of 3)

Q8. A second-line risk reviewer asks: “how do you know this LLM agent will behave the same next quarter as it does today?” Which response is **most** appropriate?

- A. “LLMs are deterministic, so behavior won't change.”
- B. “We have drift monitors, periodic re-evaluation against a gold set, and a kill-switch.”
- C. “We trust the vendor's marketing.”
- D. “We only use the latest model so this is not a concern.”

Answer: B · Drift is a real concern for LLM systems. Monitors, periodic re-eval against a gold set, and a kill-switch are the standard mitigations.

Q9. Which of the following is the **most** appropriate first artifact to produce when you scope a new genAI use case in a bank?

- A. Production-ready code.
- B. A press release.
- C. A one-page use-case brief: problem, users, data, success metrics, risks, controls.
- D. A blog post about LLMs.

Answer: C · A concise use-case brief aligns engineering, risk, compliance and business early. It is the standard first artifact.

Q10. An auditor asks for evidence that your AI assistant did not give a harmful answer last Tuesday. Which artifact is **most** directly useful?

- A. The latest training data dump.
- B. Application logs with prompt, retrieved context, output, decision and timestamp.
- C. A model architecture diagram.
- D. A summary of all sessions for that user across all systems.

Answer: B · Per-request logs with prompt, context, output and decision/timestamp are the core audit evidence for genAI systems.

Study tips for the exam

- Re-do the sample exams twice — once for understanding, once under strict timing.
- For each missed question, write the principle behind it in your own words.
- Stress-test your understanding by explaining a topic out loud as if to a CFO.
- Don't memorize regulation numbers; understand the *shape* of expectations.

EARNINGS & ROI

Salary band reference

Indicative annual compensation ranges for AI-credentialed finance and banking roles in 2026. Bands vary by geography, employer tier and seniority — use these as directional benchmarks to anchor your negotiations, not as guarantees.

Role family	Entry	Mid	Senior
AI Banking Operations / KYC Automation	\$ low	\$\$ mid	\$\$\$ senior
Fintech AI Specialist / AI Engineer (Finance)	\$\$	\$\$\$	\$\$\$\$
AI Risk Analyst / Model Validation	\$\$	\$\$\$	\$\$\$\$
AI Compliance Officer / Audit Lead — AI	\$\$	\$\$\$	\$\$\$\$
AI Investment Analyst / Quant Analyst (AI)	\$\$\$	\$\$\$\$	\$\$\$\$\$
AI Product Manager — Banking / Fintech	\$\$\$	\$\$\$\$	\$\$\$\$\$

Legend: \$ = entry tier · \$\$ = early-career tier · \$\$\$ = mid-senior tier · \$\$\$\$ = senior tier · \$\$\$\$\$ = principal/lead tier. Exact ranges depend on country, employer and total-comp mix. The GSDC live page maintains the most current ranges.

Where CGAIFB shows up on a CV

- **Headline:** “GSDC Certified — Generative AI in Finance & Banking (CGAIFB)”.
- **Experience:** reference your capstone artifact as a concrete deliverable, with metrics.
- **Skills section:** list the four pathway-aligned competencies (e.g. RAG in regulated stacks).
- **Verification:** include a link to your verifiable digital badge.

EMPLOYER VIEW

Employer signals & common questions

What hiring managers screen for in 2026

- **AI fluency, grounded in finance.** Can you talk about LLMs *and* credit risk in one breath?
- **Built artifacts, not just opinions.** A capstone with eval scorecards is a real signal.
- **Risk & controls instincts.** You know what a model card is and why it matters.
- **Globally recognized credential.** Reduces the “unknown candidate” discount on shortlists.

Common questions, briefly answered

Question	Short answer
Do I need to code?	No. Pathway A leans technical; the other three are role-coded, not engineering-coded.
Is the credential globally recognized?	Yes — recognized in 100+ countries. See the live program page for current details.
How is it different from a generic AI course?	Every module, lab, capstone and sample question is built for finance & banking.
What happens after I pass?	You receive a digital certificate, verifiable badge, and lifetime credential record.
Can my employer sponsor me?	Yes. Corporate / cohort sponsorship is available via the live program page.
What if I miss live sessions?	Sessions are recorded; office hours run weekly to compensate for missed cohorts.

NEXT COHORT INTAKE

[50% OFF]

Join the next CGAIFB cohort intake

Daily live sessions, peer cohort, role-mapped capstones. Move from brochure to enrolled candidate in under two minutes.

[Join the Next Cohort →](#)

Tap anywhere on this card →

ENROLLMENT

Lock in your CGAIFB seat

You've reached the last page of the brochure. The next step is the easiest one: enroll on the live program page. Your access opens immediately, your cohort is assigned within 24 hours, and your 90-day clock starts the moment you finish Module 1.

Step	What happens
1 · Click any CTA in this brochure	You land on the official CGAIFB program page.
2 · Apply your offer at checkout	Your 50% discount is auto-applied within the offer window.
3 · Complete enrollment	Receive your LMS access, cohort schedule and mentor introduction.
4 · Start Phase 1	Open Module 1 and begin your 90-day pathway.

Direct contact

Program URL: gsdcouncil.org/certification-program/credential-for-ai-finance-jobs-and-ai-banking-jobs

Issuing body: Global Skill Development Council (GSDC)

Recognition: 100+ countries

Credential code: CGAIFB · Generative AI in Finance & Banking

globally recognized

[50% OFF]

A credential recognized in 100+ countries

Add a globally recognized GSDC credential to your LinkedIn, CV and recruiter screens — at half the standard certification fee.

[Get Certified Globally →](#)

Tap anywhere on this card →

Thank you for reading the full 90-day pathway & certification brochure. We'll see you inside the cohort.