FRAMEWORK BROCHURE · 2026

The Complete Framework Brochure

Everything on the framework page, in one downloadable PDF: the 6-pillar framework, all 14 module syllabi verbatim, the full generative AI use-case map for finance and banking, a sample exam, and a printable 90-day plan.

Inside the brochure

✓ 6-pillar framework infographic	✓ 14 module syllabi · verbatim
✓ Use cases in finance · full map	✓ Use cases in banking · full map
✓ Sample exam · printable checklist	✓ 90-day plan · printable

6	14	250k+
Pillars	**Modules**	**Certified**
Framework	Verbatim syllabi	Already inside

Credential: **CGAIFB · Generative AI in Finance & Banking** · Issued by **Global Skill Development Council** · Recognized in 100+ countries.

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

INFOGRAPHIC

The 6-pillar framework

The CGAIFB framework organizes every concept, lab and use case into six pillars. Each pillar gets its own deep-dive page later in this brochure. Together they form the spine the certification curriculum is built on.

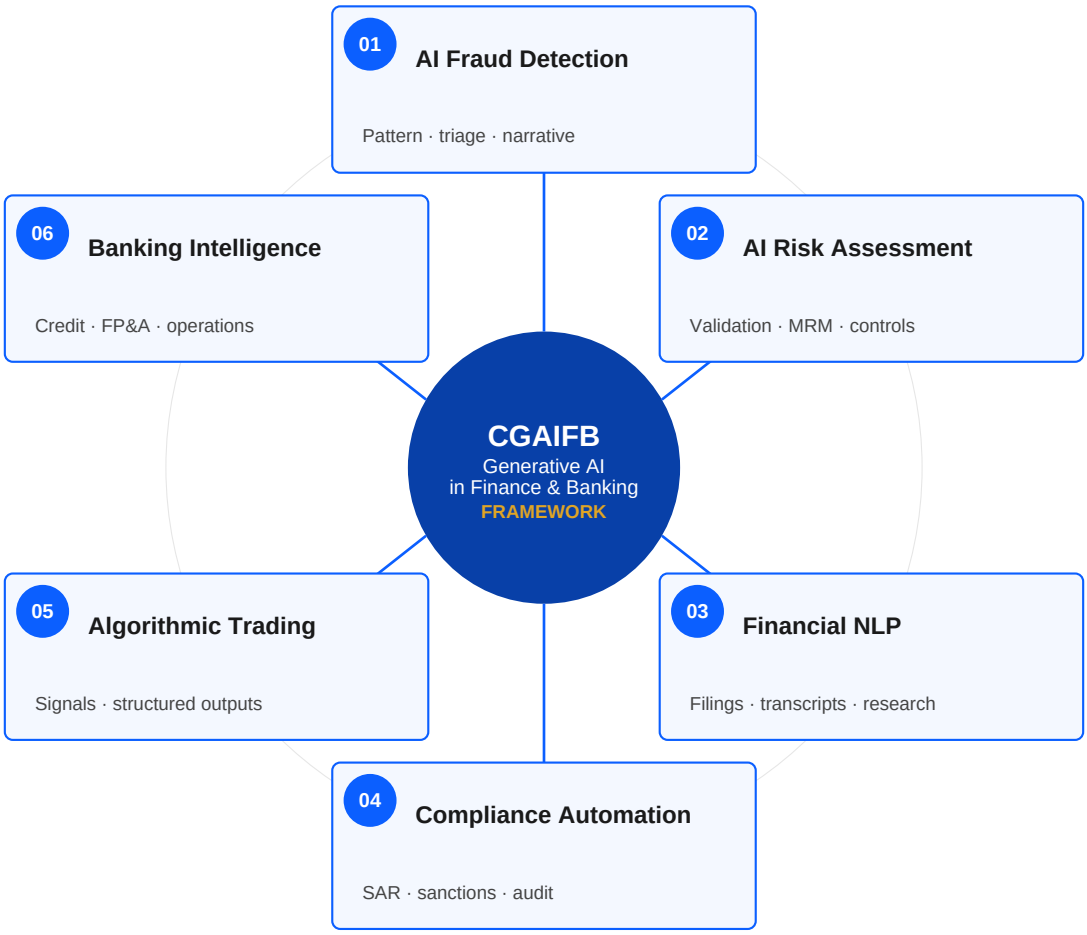

Pillars at a glance: 01 AI Fraud Detection · 02 AI Risk Assessment · 03 Financial NLP · 04 Compliance Automation · 05 Algorithmic Trading · 06 Banking Intelligence. Detail starts on page 5.

CONTEXT

Why a framework, why now

Generative AI in finance and banking has moved from slideware to production in less than two years. A framework matters because it gives professionals, hiring teams, regulators and educators a shared map of *what* the work actually is.

Three problems the framework solves

- **Fragmentation.** Use cases ship in many corners of the bank; a framework gives them a common vocabulary.
- **Mis-screening.** Hiring teams cannot screen for what they cannot name; pillars make the AI roles legible.
- **Mis-governance.** Risk and compliance partners need the same map as engineering to govern AI defensibly.

How the framework is used

By	How they use the framework
Professionals	To anchor their role and identify which pillars to invest in for the next 12 months.
Hiring teams	To write job descriptions and screen candidates against pillar-level capabilities.
Second-line risk & compliance	To structure controls and audit narratives by pillar instead of by ad-hoc use case.
Educators & trainers	To map curricula to a recognized, employer-screened framework.
Regulators & auditors	To structure inspections around the same pillars the bank organizes its AI work in.

ADOPTION

Used by 250,000+ certified professionals worldwide

The framework underpins one of the most-adopted skill credentials in finance and banking. The community sits across global banks, asset managers, insurers, fintechs and the consulting practices that serve them.

250k+	100+	6
Certified	**Countries**	**Pillars**
Across all GSDC programs	Recognition	Of the framework

Where the community sits

Sector	Common roles in the community
Global banks	AI risk analyst · AI compliance officer · fraud lead · credit / FP&A analyst
Asset managers & hedge funds	Quant analyst (AI) · research analyst · model validator
Insurance	AI claims lead · AI underwriting analyst · compliance officer
Fintech & payments	Fintech AI specialist · fraud product engineer · AI ops
Consulting & audit	Financial-services AI consultant · AI audit lead · advisory practitioner
Regulators & policy	Supervisory analyst · policy advisor · enforcement analyst

PILLAR 01

AI Fraud Detection

Generative AI applied to fraud and financial crime: pattern, triage, narrative.

What this pillar covers

Generative AI augments — but does not replace — fraud and financial-crime workflows. The pillar covers AI-assisted alert triage, pattern explanation, narrative generation for investigators, and the controls required to keep the system defensible to audit and regulators. Read-only data access and human-in-the-loop are non-negotiable here.

Use cases inside this pillar

- AI fraud-alert triage and narrative generation for investigators.
- Pattern explanation for card testing, account takeover, mule and scam typologies.
- Investigator copilot for evidence assembly and case writing.
- Drift detection on fraud patterns over weeks and quarters.

CGAIFB modules that anchor this pillar

M4 (prompting) · M6 (agents) · M8 (MRM) · M10 (privacy & security)

KPIs practitioners track

- Time saved per alert during triage (target: ~40–60% reduction).
- Narrative consistency across investigators (rated against a rubric).
- Per-alert log completeness (prompt · context · output · decision · timestamp).
- Drift signals on pattern distribution and model behaviour over time.

PILLAR 02

AI Risk Assessment

Governing and validating AI / LLM-driven systems inside risk and second line.

What this pillar covers

AI risk assessment extends classical model risk management to LLM-driven systems. The pillar covers model cards, control narratives, evaluation harnesses, drift monitoring and the cadence of validation. The pillar is what turns a prototype into a system that second line, audit and regulators will accept.

Use cases inside this pillar

- Model cards for LLM-driven systems with intended use, limitations and evaluations.
- Validation tests for grounded answers, citation faithfulness and refusal behavior.
- Drift monitoring against a gold set with periodic re-evaluation.
- MRM file structure: card, controls, tests, sign-offs, evidence retention.

CGAIFB modules that anchor this pillar

M5 (RAG) · M7 (structured outputs) · M8 (MRM) · M11 (regulation)

KPIs practitioners track

- Coverage of MRM template by AI use case (target: 100% sections addressed).
- Inter-rater agreement on quality eval (target: ≥ 0.7).
- Drift alert frequency and triage SLA.
- Audit-trail completeness on consequential outputs.

PILLAR 03

Financial NLP

Generative AI applied to filings, transcripts, research, policy and customer text.

What this pillar covers

Financial NLP is the pillar where retrieval and structured extraction live. The pillar covers chunking and embedding strategies for finance documents, retrieval design with citations, refusal behavior on ungrounded answers, and structured output extraction for downstream systems. It is the most-used pillar across the framework.

Use cases inside this pillar

- Earnings transcript Q&A with cited answers for research analysts.
- Filings extraction into structured analyst schemas.
- Policy and SOP Q&A for internal teams (credit, ops, KYC).
- Customer text classification and routing (complaints, disputes, queries).

CGAIFB modules that anchor this pillar

M3 (data) · M4 (prompting) · M5 (RAG) · M7 (structured outputs)

KPIs practitioners track

- Retrieval precision and recall against a labelled gold set.
- Citation faithfulness rate (target: ≥ 95% on a sampled audit).
- Refusal correctness on ungrounded questions.
- Per-question cost and p95 latency for interactive use.

PILLAR 04

Compliance Automation

Defensible AI controls and regulator-ready evidence for compliance and audit.

What this pillar covers

Compliance automation is the pillar where the bank's defensibility lives. The pillar covers SAR drafting, sanctions screening review, control narratives for AI use cases, regulatory mapping and audit-ready evidence packs. Practitioners here own the conversation with regulators on the bank's AI footprint.

Use cases inside this pillar

- SAR drafting with grounded retrieval and inline citations.
- Sanctions screening review assistant for hit triage.
- Control narratives for new AI use cases, written for second line.
- Regulatory mapping per use case across data, model and consumer-outcome regimes.

CGAIFB modules that anchor this pillar

M8 (MRM) · M9 (ethics) · M10 (privacy & security) · M11 (regulation)

KPIs practitioners track

- Per-section citation completeness on SAR drafts.
- Mean review time per sanctions hit.
- Inspection-readiness checklist completion across use cases.
- Regulator response cycle time and re-work rate.

PILLAR 05

Algorithmic Trading

Generative AI applied to research, signals and trading-adjacent workflows.

What this pillar covers

Algorithmic trading is the highest-paying pillar at senior levels. The pillar covers structured signal extraction from filings, transcripts and news, retrieval over research libraries, and the explicit non-decision boundary that keeps generative AI out of autonomous trading. Practitioners here often sit at the boundary of quant and product.

Use cases inside this pillar

- Structured signal extraction from filings, transcripts and news.
- Research-library Q&A with citations for buy- and sell-side analysts.
- Time-stamped signal stores for surveillance and post-trade review.
- Explicit non-decision boundaries (no autonomous order generation).

CGAIFB modules that anchor this pillar

M4 (prompting) · M5 (RAG) · M6 (agents) · M13 (operations)

KPIs practitioners track

- Signal extraction accuracy against a labelled gold set.
- Time saved per analyst per earnings season.
- Coverage of names per analyst per quarter.
- Surveillance-flag rate and per-signal provenance retention.

PILLAR 06

Banking Intelligence

Generative AI applied to credit, FP&A;, treasury and banking operations.

What this pillar covers

Banking intelligence is the pillar where most career-switchers start. The pillar covers credit memo drafting, FP&A commentary generation, treasury analytics, customer-service copilots and KYC document processing. The signature pattern is template-bound drafting with locked numbers — the model frames, the data store states.

Use cases inside this pillar

- Credit memo drafting from financials and covenant data.
- FP&A variance commentary with locked numbers.
- Customer-service copilot for contact centres.
- KYC document processing with field-level confidence scores.

CGAIFB modules that anchor this pillar

M3 (data) · M5 (RAG) · M7 (structured outputs) · M13 (operations)

KPIs practitioners track

- Draft acceptance rate by analysts (target: > 80%).
- Commentary numeric accuracy (target: 100% — numbers come from the store).
- Average handle time reduction on customer-service calls.
- Per-case onboarding cycle time reduction.

USE-CASE MAP

Generative AI use cases in finance · full map

The finance side of the framework. Each use case below maps to one primary pillar and often a secondary one. Use this page as the lookup table when you scope a new use case in your firm.

Use case	Primary pillar	Secondary pillar
Earnings transcript Q&A	03 NLP	06 Banking Intelligence
Filings extraction → analyst schemas	03 NLP	05 Algo Trading
Research-library Q&A copilot	03 NLP	05 Algo Trading
Signal extraction from news / transcripts	05 Algo Trading	03 NLP
Portfolio Q&A copilot (RM-internal)	03 NLP	04 Compliance
FP&A variance commentary	06 Banking Intel.	02 Risk Assessment
Treasury policy & SOP Q&A	06 Banking Intel.	04 Compliance
Regulatory reporting narrative drafting	04 Compliance	06 Banking Intel.
Model documentation summarization	02 Risk Assessment	04 Compliance
AI risk control-narrative writer	02 Risk Assessment	04 Compliance

Reader hint: if a use case maps cleanly to two pillars, treat it as a learning milestone — your role probably crosses both.

USE-CASE MAP

Generative AI use cases in banking · full map

The banking side of the framework. The same lookup-table layout: primary pillar, secondary pillar. Together with the finance map on page 11, these cover the production footprint of generative AI across the industry in 2026.

Use case	Primary pillar	Secondary pillar
AI fraud-alert triage & narrative	01 Fraud	02 Risk Assessment
Card / payments fraud pattern explainer	01 Fraud	03 NLP
Customer service copilot (retail)	06 Banking Intel.	04 Compliance
Dispute & chargeback drafting assistant	06 Banking Intel.	04 Compliance
KYC document processing & exception narration	06 Banking Intel.	04 Compliance
AML SAR drafting with grounded retrieval	04 Compliance	03 NLP
Sanctions screening review assistant	04 Compliance	01 Fraud
Credit memo drafting from documents & data	06 Banking Intel.	03 NLP
Covenant tracking and exception narration	06 Banking Intel.	02 Risk Assessment
Onboarding / due-diligence file enrichment	06 Banking Intel.	04 Compliance

FRAMEWORK · MAPPING

Pillars × modules · the lookup matrix

Every CGAIFB module touches at least one pillar. The matrix below shows the canonical mapping — a full blue dot is the module's primary pillar, a small dot is a strong secondary. Use it to plan which modules to study first for your pillar(s) of interest.

Module	01 Fraud	02 Risk	03 NLP	04 Compl.	05 Algo	06 Banking
M1 · GenAI & LLM literacy	▪	▪	▪	▪	▪	▪
M2 · Finance & banking context	▪	▪	▪	▪	▪	▪
M3 · Data foundations	▪		●	▪	▪	●
M4 · Prompting patterns	●		●		●	▪
M5 · RAG in regulated stacks	▪	●	●	●	●	●
M6 · Agents & tool use	●				●	▪
M7 · Structured outputs & eval	▪	●	●	▪	●	●
M8 · Model risk & validation (MRM)	●	●		●		▪
M9 · Ethics & responsible AI	▪	▪		●		▪
M10 · Data protection & security	●	▪	▪	●		▪
M11 · Regulation & supervision	▪	●		●	▪	▪
M12 · Capstone design & build	▪	▪	▪	▪	▪	▪
M13 · Integration, cost & ops	▪	▪	▪	▪	●	●
M14 · Exam prep & viva	▪	▪	▪	▪	▪	▪

Legend: ● primary mapping · ▪ strong secondary mapping. Use the matrix to pick which module to study next for your target pillar(s).

SYLLABUS

14 module syllabi · verbatim (1 of 2)

#	Module · verbatim	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

Learning outcomes per module · in one line

#	You can…
M1	Explain LLMs, tokens, context windows and embeddings in finance terms.
M2	Map functions across retail / corporate / IB / insurance to AI use cases.
M3	Classify data types and pick the right retrieval strategy for each.
M4	Apply prompting patterns to finance tasks, including refusal and grounded prompts.
M5	Architect a RAG pipeline in a regulated stack and evaluate retrieval quality.
M6	Design safe agent tool surfaces with read-only defaults and approval gates.
M7	Build schema-bound outputs with validators and a gold-set evaluation harness.

SYLLABUS

14 module syllabi · verbatim (2 of 2)

#	Module · verbatim	Phase
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

Learning outcomes per module · in one line

#	You can…
M8	Author a model card and a control narrative your second line will accept.
M9	Apply fairness lenses and translate ethical concerns into operational controls.
M10	Map data flows for a genAI app and apply masking, tenancy and retention controls.
M11	Map an AI use case to relevant regulatory regimes and anticipate inspection questions.
M12	Scope, design, build and present a role-mapped capstone with a 10-minute viva.
M13	Plan integration, cost, latency and rollout with kill-switch criteria.
M14	Sit the CGAIFB exam under timed conditions and pass on the first attempt.

STUDY PLAN

The 90-day plan · printable

Five to seven focused hours per week. Print this page; check off each week as you complete it. The plan below is the canonical pathway most cohort members follow.

Week	Focus	Deliverable
W1 · Days 1–7	M1 · AI & LLM literacy for finance	30-term finance-flavoured glossary
W2 · Days 8–14	M2 · Finance & banking context	Workflow map · 5 controls list
W3 · Days 15–21	M3 · Data foundations	Data inventory + retrieval sketch
W4 · Days 22–28	M4 · Prompting patterns	Finance prompt library (20+ prompts)
W5 · Days 29–35	M5 · RAG in regulated stacks	Mini-RAG prototype with eval scorecard
W6 · Days 36–42	M6 · Agents & tool use	Agent design doc + single-step agent
W7 · Days 43–49	M7 · Structured outputs & eval	Validator · 3 schemas · gold set
W8 · Days 50–56	M8 · Model risk & validation	Mini-MRM file on your prototype
W9 · Days 57–63	M9 · Ethics & responsible AI	Bias audit + ethics statement
W10 · Days 64–70	M10 · Data protection & security	Data flow + 15-item security checklist
W11 · Days 71–77	M11 · Regulation & supervision	Regulatory mapping table + mock viva
W12 · Days 78–84	M12 · Capstone build & review	Capstone artifact + 1-page summary
W13 · Days 85–90	M13/M14 · Integration · exam	Exam booked & cleared · credential issued

CADENCE

Inside a typical study week

The plan on page 16 assumes the weekly cadence below. Most readers anchor a single live session per week and treat everything else as flexible.

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 (optional)	30 min
Friday	Apply: build / refactor a small artifact	45–60 min
Saturday	Consolidation: notes, flashcards, peer review	45 min
Sunday	Light review + plan next week (or rest)	15–20 min

Why this cadence works

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

EXAM

Sample exam · structure & sample items

The CGAIFB exam tests application, not recitation. The structure mirrors real practitioner decisions; pages 19–20 carry sample items in the same style as the live exam.

Item	Detail
Format	Multiple-choice + scenario items + applied mini-case
Items	Approx. 60–80 items
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 the GSDC retake policy
Credential issued	Digital certificate & verifiable badge on completion
Validity	Lifetime credential · refresh recommended every 2 years

Mapping the exam to the framework

Exam items are sampled across all six pillars and all 14 modules. A typical paper covers every pillar at least once, with weighting reflecting the live curriculum's emphasis on applied building (modules 4–7) and governance (modules 8–11).

EXAM PREP

Sample exam questions (1 of 2)

Q1. A retail bank rolls out an LLM-driven fraud triage assistant. Which control is **most** important to enforce **before** go-live?

A. Switch to the largest available model.

B. Restrict the assistant to read-only data and require analyst confirmation for any block / freeze action.

C. Cache all responses to reduce cost.

D. Disable retrieval and let the model generate from memory.

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

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

A. A method to fine-tune a model on private banking data.

B. A compression technique to reduce token cost.

C. An architecture where the model answers using documents retrieved at query time, ideally with citations.

D. A way to bypass content filters via retrieved phrasings.

Answer: C · *RAG grounds the model in retrieved documents at query time. Citations and refusal behavior are critical for traceability.*

Q3. Which artifact is **most** useful for a second-line MRM team reviewing an LLM-driven credit memo assistant?

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.

Answer: B · *Model cards are the standard artifact for model risk review of AI systems.*

Q4. An LLM-based KYC assistant is asked a question its retrieved context does not answer. The **most** defensible response is to:

A. Answer from general knowledge to be helpful.

B. Refuse and explain that the answer is not in the supporting documents.

C. Generate the most likely answer with a confidence score.

D. Forward the query to the customer without any response.

Answer: B · *Ungrounded answers create audit and customer-harm risk. Refusing with a clear reason is the safe, defensible behavior.*

EXAM PREP

Sample exam questions (2 of 2)

Q5. Your RAG-based credit memo drafter occasionally writes a slightly different number than the underlying statement. Which fix is **most** appropriate?

A. Increase the model size.

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

C. Add a disclaimer to the memo.

D. Reduce the temperature to zero.

Answer: B · *Numbers should come from the data store, not the model. Template injection is the standard mitigation.*

Q6. A second-line MRM reviewer asks: *how do you know this LLM agent will behave the same next quarter as it does today?* The **best** response is:

A. LLMs are deterministic, so behaviour 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 claims.

D. We always use the latest model so this is not a concern.

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

CHECKLIST

Printable exam-prep checklist

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

Foundations & vocabulary

- ☐ I can explain LLMs, tokens, context windows and embeddings in finance terms.
- ☐ I can name the four architectural patterns (RAG · template-bound · constrained agent · schema-bound classifier).
- ☐ I can name three controls every AI use case in a bank needs.
- ☐ I can read a model card and spot what's missing.

Applied building

- ☐ I have built a finance prompt library (20+ prompts) and tested them side by side.
- ☐ I have built a mini-RAG prototype with a gold-set eval scorecard.
- ☐ I have designed a single-step agent with read-only tools and an approval gate.
- ☐ I have built a schema-bound JSON output with a validator and retries.

Governance & risk

- ☐ I have authored a model card for my prototype, end to end.
- ☐ I have written a control narrative for my prototype, end to end.
- ☐ I have produced a data flow diagram and a 15-item security checklist for my use case.
- ☐ I have run a mock regulator interview on my use case with a mentor.

Capstone & exam day

- ☐ My capstone artifact is built, reviewed and signed off by my mentor.
- ☐ I have completed two sample exams under timed conditions.
- ☐ I have written a flashcard set covering all 14 modules' must-know terms.
- ☐ My exam slot is booked.

ADOPTION GUIDE

How to use the framework inside your firm

The framework is most useful when it stops living in a brochure and starts living in your firm's vocabulary. Below are the four ways readers actually deploy it.

Where	How to use the framework
1 · In job descriptions	Frame each AI-adjacent role by its anchor pillar(s); name the pillar in the JD title.
2 · In capability mapping	Score your team's capability per pillar (1–5). Use it as the gap analysis for hiring & training.
3 · In controls libraries	Organize your bank's AI controls catalog by pillar instead of by ad-hoc use case.
4 · In regulator briefings	Present the bank's AI footprint by pillar; regulators can navigate the framework alongside you.

Pattern: pillar-first job descriptions

Most banks today still write generic “AI in finance” job descriptions. The pillar-first approach is sharper: name the pillar in the JD title (e.g. *AI Risk Assessment Lead* instead of *AI Lead in Risk*), then anchor the screening rubric to that pillar's capabilities. Candidates self-screen; recruiters screen consistently; the role's career path is legible from day one.

CAREER

Career roadmap, aligned to the six pillars

Five role families anchor the bulk of pillar-aligned hiring. Each role anchors to one primary pillar and frequently a secondary one — exactly the way the use-case maps on pages 11–12 are structured.

Role family	Primary pillar	Secondary pillar	Anchor modules
AI Risk Analyst	02 Risk	04 Compliance	M5, M7, M8, M11
Quant Analyst (AI)	05 Algo	03 NLP	M4, M5, M7, M13
AI Fraud Detection Specialist	01 Fraud	02 Risk	M4, M6, M8, M10
AI Compliance Officer	04 Compliance	02 Risk	M8, M9, M10, M11
AI Financial Modeller	06 Banking	03 NLP	M3, M5, M7, M13

How readers pick a role

- **Anchor your interest in a pillar first.** Roles flow from pillars, not the other way around.
- **Map your existing skills to the secondary pillar.** That's where domain experience pays off fastest.
- **Pick anchor modules with the matrix on page 13.** Study those first, regardless of phase order.
- **Build the capstone in your primary pillar.** It's the artifact recruiters actually read.

MEASUREMENT

KPIs every pillar tracks

Every pillar has its own KPIs (pages 5–10). But a small set repeats across the entire framework. Practitioners track these regardless of which pillar they sit in.

KPI family	What it measures	Why it matters
Quality	Grounded answer rate · citation faithfulness · refusal correctness.	Defends the system to second line and audit.
Cost	Per-task cost · token consumption per user · cost-per-business-outcome.	Keeps the use case economic at scale.
Latency	p50 / p95 latency per task; per-region SLAs.	Decides whether the system can be inline or only batch.
Drift	Quality dip vs gold set; distribution shifts on inputs / outputs.	Catches silent degradation before users do.
Coverage	% of use cases with model cards, controls, eval scorecards.	Tells leadership how mature the AI footprint actually is.
Adoption	% of target users using the assistant weekly; edit-vs-accept rates.	Distinguishes shipped tools from used tools.

Rule of thumb: if your team cannot produce these six KPI families for any AI use case in production, the use case is not yet ready for second-line review.

PITFALLS

Common pitfalls across the framework

Eight pitfalls account for the majority of AI use-case incidents inside banks. Memorize them — the framework's structure exists in large part to prevent each one.

Pitfall	How the framework prevents it
Hallucination on consequential outputs	Pillar 02 + Module 5: grounded-only design with refusal behavior.
Citation faithfulness gap	Pillar 03 + Module 5: citation enforcement and quality eval on a gold set.
Number drift in drafted text	Pillar 06 + Module 7: numbers from the data store; the model never generates them.
Prompt injection from retrieved content	Pillar 04 + Module 10: input isolation and output validators.
Schema breakage downstream	Module 7: JSON-mode / function calling with validators and retries.
Silent drift over weeks / quarters	Pillar 02 + Modules 8/13: periodic re-eval against gold set + alerts.
Cost runaway as context grows	Module 13: per-task cost dashboards and context-size budget enforcement.
Permission leakage in retrieval	Pillar 04 + Module 10: permission-aware retrieval; filter at retrieval, not the model.

30-DAY APPLICATION

Apply the framework in 30 days

If you have not yet started the credential, you can still apply the framework to your current work this month. Below is the 30-day pattern most early readers follow.

Week	Action	Output
W1 · Map yourself	Score your current AI fluency 1–5 against each pillar.	A radar of strengths and gaps.
W2 · Pick one pillar	Pick the pillar that fits your role today; identify two use cases inside it.	Two named use cases and their primary / secondary pillars.
W3 · Brief one use case	Write a one-page use-case brief: problem · users · data · metrics · risks · controls.	An aligned, internal-quality brief you can share.
W4 · Socialize the framework	Share the pillar map with one peer in risk, compliance or product.	A common vocabulary inside your team.

Most readers tell us the 30-day pattern above pays back on its own — and that the credential simply formalises the playbook they've already begun.

EMPLOYER VIEW

Employer view & common questions

What hiring teams in finance and banking look for

- **Framework literacy.** You can name the six pillars and where your role sits in them.
- **Built artifacts.** A capstone with an eval scorecard reads at one level above a course completion.
- **Controls instinct.** You can name three controls a use case needs without thinking.
- **Globally recognized credential.** Reduces the “unknown candidate” discount on shortlists.

Common questions

Question	Short answer
Is the framework GSDC-proprietary?	The framework structure is published by GSDC and used in the CGAIFB curriculum.
Are all six pillars covered in the exam?	Yes. The exam samples across all six pillars and all 14 modules.
Do I need to code to do the credential?	No. Engineering tracks lean technical; risk, ops, compliance and credit tracks do not.
Is the credential globally recognized?	Yes — recognized in 100+ countries; used by 250,000+ certified professionals across GSDC.
How long does it take, realistically?	Most readers finish in 90 days at 5–7 focused hours per week.
Can my employer sponsor me?	Yes. Corporate / cohort sponsorship is available via the live program page.

ENROLLMENT

Step into the framework

You've read the framework end to end. The next step is being inside it. Enrol on the live program page; your access opens immediately; your cohort is assigned within 24 hours; and your 90-day pathway starts the moment you open Module 1.

Step	What happens
1 · Click any CTA in this brochure	You land on the official framework 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 Module 1	Open the LMS and begin the 90-day pathway.

Direct contact

Program page: gsdcouncil.org/certification-program/generative-ai-in-finance-and-banking-framework
Issuing body: Global Skill Development Council (GSDC)
Recognition: 100+ countries · 250,000+ certified professionals
Credential code: CGAIFB · Generative AI in Finance & Banking

Thank you for reading the complete framework brochure. We'll see you inside the framework.