USE-CASE PLAYBOOK · 2026

Download the Full Use-Case Playbook

A 28-page PDF: every generative AI use case in finance and generative AI use case in banking — workflow, architecture diagram, prompt template and tool stack. Plus the complete generative AI in finance and banking certification framework.

Inside the toolkit

✓ **9 production-ready use cases · annotated**	✓ **Architecture & stack for each**
✓ **Sample prompts & templates**	✓ **All 14 official module syllabi**
✓ **CGAIFB certification career roadmap**	✓ **Field-tested controls catalog**

9	14	28
Use cases	**Modules**	**Pages**
Production-ready	Official syllabi	Playbook PDF

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

Page one is content-only by design. The first use case begins on page four — read straight through.

ORIENTATION

How to read this playbook

This playbook is intentionally hands-on. Each of the nine use cases is presented on a single page, in the same compact format, so you can compare them at a glance and pull patterns into your own program of work.

Anatomy of a use-case page

Section	What you'll find
Problem	The business problem in one paragraph — what hurts today, and for whom.
Architecture diagram	A five-stage pipeline of the system: input → grounding → model → output → integration.
Sample prompt / template	The structural skeleton of the prompt or schema — generalized for use in your own bank.
Tool stack	The layers of the stack and the tools / patterns commonly seen in production.
Controls & risk hooks	The minimum controls that make this use case defensible to second line and audit.
Observed lift	The directional outcome practitioners report — not a guarantee, an order-of-magnitude.
Mapped to module(s)	Which CGAIFB module(s) the use case maps to in the official curriculum.

Every diagram, prompt and stack in this playbook is generalized. Names are removed; patterns and controls are not. Specific tools are illustrative — substitute your bank's approved equivalents.

INDEX

The nine use cases at a glance

All nine use cases below are **production-ready** — meaning teams across global banks, asset managers and fintechs are running them today inside their existing risk-and-controls fabric. Each is built on the same shape: ground, draft with constraints, log, keep a human on consequential calls.

#	Use case	Domain	Page
01	AI fraud detection software	Fraud & financial crime	4
02	AI credit scoring & memo drafting	Lending & credit	6
03	AI document processing in finance	Banking ops / KYC	8
04	AI compliance tools for SAR & sanctions	Compliance & AML	10
05	Generative AI for financial reporting	Finance / FP&A	12
06	Algorithmic trading AI tools	Markets / trading	14
07	AI risk assessment tools (banking)	Model risk / second line	16
08	AI portfolio management tools	Wealth & asset mgmt.	18
09	AI regulatory reporting tools	Regulatory reporting	20

How each maps to the CGAIFB curriculum

Every use case in this playbook is built and validated as part of the official CGAIFB lab curriculum. The mapping appears at the foot of each use-case page so you can move from the playbook to the relevant module in one step.

USE CASE 01

AI fraud detection software

Fraud & financial crime

Problem

Fraud investigators face thousands of low- and medium-risk alerts per day. Most are noise; the costly few hide in the volume. The task is triage, prioritization and auditable narrative — not classification alone.

Architecture

Controls: Read-only data · Human-in-the-loop · Per-alert log · Kill-switch

Sample prompt template

```
SYSTEM: You are an AML/fraud triage assistant for a retail bank. You never make
block or freeze decisions. You only classify the likely pattern and draft an
investigator-ready narrative grounded in the retrieved context.

USER: alert_id={id}
context={recent_txns, device_signals, kyc_flags, prior_alerts}

OUTPUT JSON: {
  "likely_pattern": one of [card_testing, mule_account, ato, scam, other],
  "confidence": low|medium|high,
  "narrative": <= 6 sentences, every fact must cite context,
  "recommended_action": one of [close, watch, escalate]
}
```

Tool stack

Layer	Tools / patterns commonly seen
Data & retrieval	Transaction store · feature store · vector store for narratives
Model layer	Frontier LLM with JSON-mode / function calling, banking-tuned guardrails
Integration	Case management system · SIEM · investigator UI
Observability	Per-alert prompt + context + output + decision log, retention per policy

Observed lift: **~40–60% time saved on triage; higher narrative consistency.** · **Maps to CGAIFB modules:** M4 (prompting), M5 (RAG), M6 (agents), M8 (MRM).

CONTROLS

Controls catalog used across the nine use cases

Every use case in this playbook leans on a small set of recurring controls. Master these and you have most of the defensible baseline a second line will accept.

Control	What it does	Where it shows up
Read-only data access	Restricts the model and agent to read-only APIs against systems of record.	Fraud · KYC · Credit · Reporting
Human-in-the-loop	Mandatory human approval on consequential outputs before they leave the system.	All customer-facing & high-impact use cases
Grounded answers only	Refuse if the retrieved context does not support the answer; never improvise.	RAG-based use cases (SAR, KYC, research)
Inline citations	Every fact in the output cites the exact source paragraph or row used.	SAR · sanctions · reporting · research
Numbers from data store	Numeric values come from structured stores, not the LLM's generation.	Credit · FP&A · regulatory reporting
Per-request audit log	Prompt, retrieved context, output, decision and timestamp logged immutably.	All use cases without exception
Model card	Intended use, limitations, evaluations, drift signals, owners.	MRM artifact for every use case
Drift & eval monitors	Periodic re-eval against a gold set; alerts if quality drops or distribution shifts.	Long-running production use cases
Kill-switch	One-click pause of the use case at any layer (retrieval, model, integration).	All production use cases
Schema-bound outputs	JSON-mode or function calling against strict schemas; reject malformed outputs.	Use cases integrating into downstream systems

USE CASE 02 | **AI credit scoring & memo drafting** | Lending & credit

Problem

Credit analysts spend hours pulling figures from statements and rewriting boilerplate into the bank's memo template. Decisioning remains a human call — but most of the work around it is high-volume document and data shuffling.

Architecture

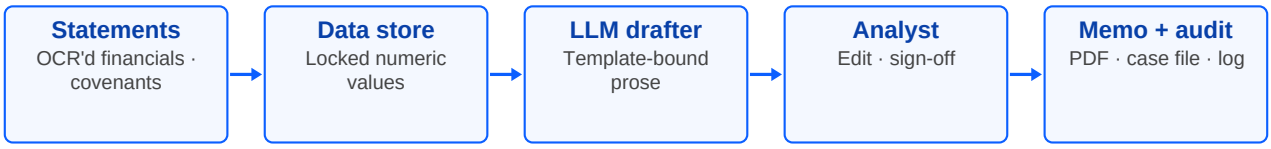

Controls: Numbers from store · Inline citations · Mandatory sign-off

Sample prompt template (excerpt)

```
SYSTEM: You draft a credit memo using the bank's official template. You must
never invent numbers. All numeric values come from {financials_json}.
If a value is missing, write <missing> and do not estimate.

USER: borrower={name}, sector={sector}, request={amount, tenor}
context={financials_json, covenants_json, prior_memos}

Sections to draft, in order:
1. Executive summary (3–5 sentences, no numbers invented)
2. Financial highlights (table-bound to financials_json)
3. Covenants & conditions (cite covenants_json clause IDs)
4. Risks & mitigants (each bullet must cite a source)
```

Tool stack

Layer	Tools / patterns commonly seen
Data & retrieval	Document AI for statements · structured covenants store · prior-memo corpus
Model layer	LLM with JSON-mode for tables · template engine for memo layout
Integration	Loan origination / decisioning system · DMS · analyst review UI
Observability	Per-memo log: prompt, retrieved context, output, analyst edits, sign-off

Observed lift: **~30–50% faster first drafts; analysts shift to judgement, not retyping.** · **Maps to modules:** M3 (data), M5 (RAG), M7 (structured outputs), M11 (regulation).

PATTERNS

The four patterns that repeat across all nine use cases

Across the nine use cases that follow, four architectural patterns repeat. Most production systems in finance and banking are a mix of these — almost never one alone.

Pattern 1 · Grounded RAG with strict refusal

Retrieve relevant context; answer only from that context; refuse cleanly when context does not support an answer. Used in: SAR drafting, KYC, policy Q&A, regulatory reporting, research assistants.

Pattern 2 · Template-bound drafting with locked numbers

The model fills in prose around values pulled deterministically from a data store. Numbers and tables come from the store; the model frames. Used in: credit memos, FP&A commentary, regulatory narratives.

Pattern 3 · Constrained agent with read-only tools

An agent can call a small, read-only tool surface; consequential actions go to a human approver. Used in: fraud triage, dispute drafting, customer-service copilot.

Pattern 4 · Schema-bound classifier with rationales

The model returns a strict JSON object — label, confidence, rationale, citations — consumed by downstream systems. Used in: sanctions screening review, risk assessment, transcript Q&A, claims triage.

Why this matters: when you scope a new use case, the first design question is *which of these four patterns fits?* The CGAIFB curriculum drills exactly this decision until it becomes instinct.

USE CASE 03 | **AI document processing in finance** | Banking ops / KYC

Problem

Onboarding is delayed by manual document review, name and address normalization, and exception research. Document AI alone solves only the extraction — humans still drown in exception narratives.

Architecture

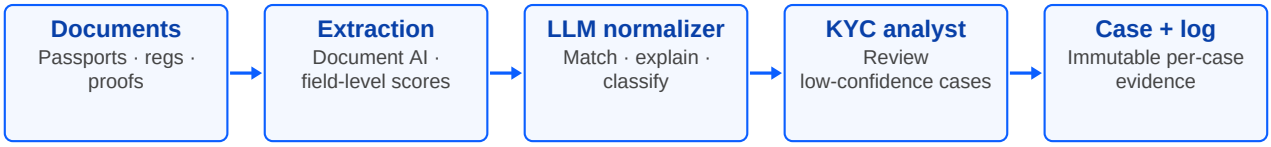

Controls: Field-level confidence · Human review on low-conf · Per-case log

Sample prompt template (excerpt)

```
SYSTEM: You normalize and reconcile extracted KYC fields against bank records.
You must NOT decide onboarding outcomes; only explain matches and exceptions.

USER: extracted_fields={dict with per-field confidence}
bank_record={canonical customer record}

OUTPUT JSON: {
  "matches": [{field, status: match|near|mismatch, reason}],
  "exception_narrative": <= 4 sentences for analyst review,
  "requires_human_review": boolean
}
```

Tool stack

Layer	Tools / patterns commonly seen
Data & retrieval	Document AI for extraction · golden-record store · sanctions / PEP feeds
Model layer	LLM in JSON-mode for normalization; classification head for exceptions
Integration	KYC case management · onboarding workflow · core banking customer master
Observability	Per-field confidence scores · per-case decision log retained per policy

Observed lift: **materially reduced onboarding cycle time; cleaner exception narratives for second line.** · **Maps to modules:** M3 (data), M5 (RAG), M7 (structured outputs), M10 (privacy).

EVAL

How practitioners evaluate these systems

Use cases that scale always have an **evaluation harness** sitting next to the live system. It is what catches drift, regressions and silent failures before users do. Below is the harness shape practitioners use across all nine use cases in this playbook.

The four-layer eval

Layer	What it measures	Tools / patterns
1 · Unit eval	Does a single prompt produce the right structure / schema?	Schema validators · regex / JSON-schema · unit-test harness
2 · Quality eval	Is the output factually grounded, well-cited and useful?	Gold set (30–200 items) · judge model or human raters · rubric
3 · Cost & latency	Does it run within budget at expected load?	Token / latency dashboards · cost-per-task · p50/p95 latency
4 · Drift	Has behaviour shifted vs last quarter?	Periodic re-eval against the gold set · distribution / output monitors

Minimum eval before go-live (checklist)

- A gold set of **at least 30 representative items** with expected outputs.
- A **rubric** that any rater on your team can apply consistently (≥ 0.7 inter-rater agreement).
- A **cost & latency budget** — and an alert when you breach it.
- A **drift plan**: how often you re-eval, and who owns it.
- A **kill-switch runbook** — what triggers it, what happens when it fires.

USE CASE 04

AI compliance tools for SAR & sanctions

Compliance & AML

Problem

Suspicious activity reports require dense, defensible narratives. Sanctions hits generate large false-positive volumes that still need rationale. Both are slow, error-prone, and audit-sensitive.

Architecture

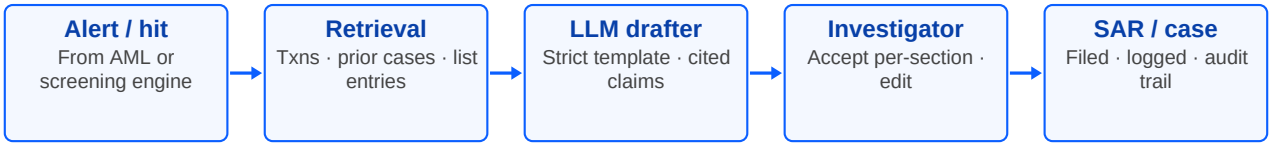

Controls: Inline citations · Tenant-isolated · Per-section sign-off

Sample prompt template (excerpt)

```
SYSTEM: You draft SAR sections grounded strictly in retrieved context. Every
factual claim must cite a source id. If the context does not support a section,
respond <insufficient evidence> for that section.

USER: case_id={id}
context={txns, prior_alerts, kyc_record, prior_sar_excerpts}

Draft per the regulator template:
- Subject info (cite kyc_record)
- Activity summary (cite txns)
- Relationship to prior cases (cite prior_alerts / prior_sar_excerpts)
- Reason for suspicion (cited bullets)
```

Tool stack

Layer	Tools / patterns commonly seen
Data & retrieval	Transaction store · prior-case archive · sanctions list · PEP feeds
Model layer	LLM with tenant-isolated deployment, strict refusal behavior, citation enforcement
Integration	AML case management · regulator filing portal · second-line review queue
Observability	Per-section drafts logged · investigator edits captured · evidence retention

Observed lift: faster SAR turnaround without losing investigator judgement or audit defensibility. · **Maps to modules:** M5 (RAG), M8 (MRM), M11 (regulation), M14 (capstone-grade).

DATA & PRIVACY

Data flow and privacy across the playbook

Every use case in this playbook lives or dies on how cleanly its data flow is mapped — and how those flows are protected. The diagram below is the canonical shape practitioners trace before any go-live conversation with second line.

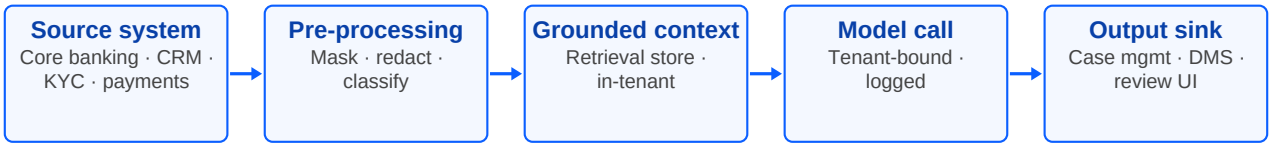

Controls: Masking · Tenancy · Retention · Right-to-delete

Privacy practices that show up everywhere

- **Minimization at ingestion.** Only fields needed for the use case enter the AI pipeline.
- **Masking before model call.** PII tokenized or redacted when the use case allows.
- **Tenant-bound model deployment.** No prompts or outputs leave the bank's tenant.
- **Retention policy mirrored.** Logs and outputs retained per the bank's retention rules.
- **Right-to-delete plumbed end-to-end.** Including vector stores, logs and caches.

USE CASE 05 | **Generative AI for financial reporting** | Finance / FP&A

Problem

Month-end reporting requires variance commentary across many cost centres and business lines. Most is mechanical; some needs real judgement. Speed and consistency suffer when analysts hand-write every line.

Architecture

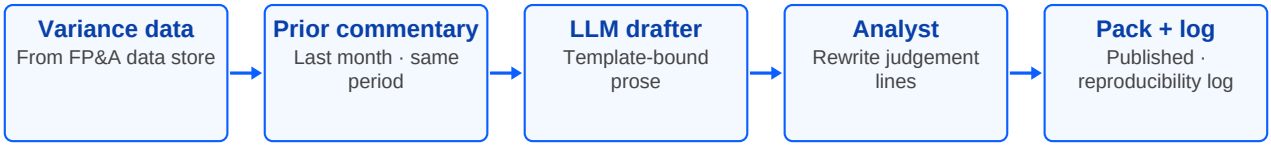

Controls: Numbers from store · Analyst edit gate · Reproducibility log

Sample prompt template (excerpt)

```
SYSTEM: You draft FP&A variance commentary. NEVER invent numbers. All values
are inserted from {variance_json}. Mark anything you cannot ground as <needs analyst>.

USER: period={month, year}, cost_centre={code, name}
context={variance_json, prior_commentary, narrative_style_guide}

Output sections:
1. Headline (one-line variance vs plan, cite variance_json keys)
2. Drivers (3 bullets, each tied to a variance line)
3. Outlook (1–2 sentences; mark as <needs analyst> if no signal)
```

Tool stack

Layer	Tools / patterns commonly seen
Data & retrieval	FP&A / GL data warehouse · prior-pack archive · narrative style guide
Model layer	LLM with strict numeric injection · template enforcement · reproducibility seed
Integration	Reporting pack generator · BI tools · finance reviewer UI
Observability	Per-line provenance: which data source, which template, which analyst edits

Observed lift: **reporting cycle shortened by days; analysts focus on commentary that matters.** · **Maps to modules:** M3 (data), M7 (structured outputs), M8 (MRM), M13 (operations).

OPERATIONS

Cost & latency budgets, in one page

Every use case has a cost and latency budget — explicit or implicit. Practitioners who ship treat these as first-class design constraints, not afterthoughts.

Use case	Latency target	Cost shape
Fraud triage assistant	≤ 2s per alert (interactive)	Per-alert cost; volume * cost-per-alert dashboards
Credit memo drafting	≤ 60s per memo (batch-friendly)	Per-memo cost; cap on retries; deterministic where possible
KYC document processing	≤ 30s per case	Per-case cost; document AI tier vs LLM tier separated
SAR drafting	Minutes per case (long-form)	Per-case cost; second-line review time accounted for
FP&A commentary	Minutes per cost-centre (batch overnight)	Per-pack cost; ran in cycle, not real-time
Algo-trading research assist	≤ 1s for retrieval, model TBD	Per-query cost; bounded concurrency on model layer
AI risk assessment	Minutes per model (review-time)	Per-model-doc cost; concentrated at validation cycles
Portfolio Q&A	≤ 3s per question (interactive)	Per-query cost; cache common questions
Regulatory reporting	Hours per filing (long-form)	Per-filing cost; bounded against the filing calendar

Rule of thumb: if you cannot state your use case's per-task cost and p95 latency in one sentence each, you are not ready for production review.

USE CASE 06 | **Algorithmic trading AI tools** | Markets / trading

Problem

Trading desks consume vast amounts of unstructured information — filings, transcripts, news, broker notes. Reading every source end-to-end is no longer feasible; structured extraction with grounded answers is.

Architecture

Controls: No autonomous trading · Citations required · Per-signal log

Sample prompt template (excerpt)

```
SYSTEM: You extract structured market signals from a single source. You do NOT
make trading recommendations. You report what the source says, with citations.

USER: source_type={filing|transcript|news}
source_text={raw_text}

OUTPUT JSON: {
  "entity": ticker / issuer,
  "period": reporting period if applicable,
  "signals": [{topic, claim, sentiment, quote_span}],
  "refusals": [reasons for non-extractable sections]
}
```

Tool stack

Layer	Tools / patterns commonly seen
Data & retrieval	Filings feed · transcript provider · news pipeline · vector index
Model layer	LLM with JSON-mode · structured-output enforcement · refusal on ungrounded claims
Integration	Signal store · research notebook · risk & surveillance hooks
Observability	Per-signal provenance · drift monitors · explicit non-decision boundary

Observed lift: **broader coverage per quant / analyst; less missed colour; more time for thesis work.** · **Maps to modules:** M5 (RAG), M6 (agents), M7 (structured outputs), M11 (regulation).

FAILURE MODES

The failure modes that bite — and how to catch them

Almost every production incident in finance and banking AI comes back to a small set of failure modes. Memorize this list. The CGAIFB curriculum drills exactly these.

Failure mode	What it looks like	How to catch / prevent
Hallucination	Confident, fluent claims that aren't supported by the retrieved context.	Citation enforcement · refusal behavior · grounded-only design.
Citation faithfulness gap	The model cites a doc that doesn't actually contain the claimed fact.	Quality eval with rater check on citations · automated source-match.
Number drift	The model writes a slightly different number than the source.	Numbers from the data store · template injection · never from the model.
Prompt injection	User or document content overrides the system instructions.	System prompt isolation · output validators · escape rules on retrieved text.
Schema breakage	Output doesn't conform to the expected JSON schema.	JSON-mode / function calling · validator · retries with reduced creativity.
Silent drift	Quality degrades quietly over weeks as model / data / users shift.	Periodic re-eval against gold set · alerts on dips · model card updated.
Cost runaway	Token consumption climbs as prompts / context grow.	Per-task cost dashboards · budget alerts · context-size budget enforced.
Permission leakage	User sees grounded content they shouldn't have permission to see.	Permission-aware retrieval · per-user filters at retrieval, not at the model.

USE CASE 07 | **AI risk assessment tools (banking)** | Model risk / second line

Problem

Model documentation is dense and inconsistently structured. Validators are slow because every model owner writes the file differently — and LLM-driven systems multiply the review burden.

Architecture

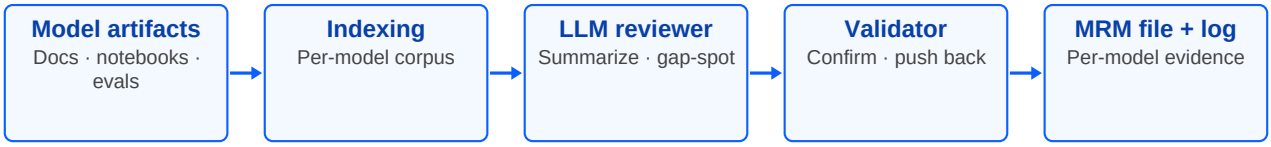

Controls: Validator approval · Ungrounded-claim refusal · Per-model log

Sample prompt template (excerpt)

```
SYSTEM: You draft structured summaries of model documentation against the bank's
MRM template. You may NOT add claims the source does not make. Mark gaps clearly.

USER: model_id={id}, template_sections={...}
context={model_doc_chunks, eval_artifacts, prior_validation_notes}

OUTPUT JSON per section: {
  "section": MRM template section name,
  "summary": <= 5 sentences, every claim cited,
  "open_questions": [list of validator questions],
  "gaps": [missing or unsupported areas]
}
```

Tool stack

Layer	Tools / patterns commonly seen
Data & retrieval	Per-model corpus (docs + notebooks + evals) · MRM template store
Model layer	LLM with strict citation enforcement · structured output per template section
Integration	MRM workflow tooling · validator review UI · evidence vault
Observability	Per-model file · validator overrides logged · drift in summary quality monitored

Observed lift: **faster model reviews · standardised first-pass narratives · sharper validator focus.** · **Maps to modules:** M5 (RAG), M7 (structured outputs), M8 (MRM), M11 (regulation).

DESIGN

Designing the human-in-the-loop properly

“Human-in-the-loop” is the most-claimed and least-engineered control in finance AI. Below is the design grid practitioners use to make HITL actually work.

Decision class	Recommended HITL design	Example use case
Customer-facing send	Block-by-default; human must explicitly send.	Dispute / chargeback drafting · customer service copilot
Money movement	AI cannot trigger; human approves through a separate write path.	Account operations · payment authorisation copilots
Risk / case decision	AI recommends + rationale; human accepts per item with audit trail.	Fraud triage · sanctions review · claims triage
External filing	Human signs each section; immutable evidence retained.	SAR · regulatory reporting · audit packs
Internal artifact	Edit-by-default; system tracks human edits as ground truth signal.	Credit memos · FP&A commentary · MRM summaries
Information retrieval	No HITL on read; permission-aware retrieval prevents over-exposure.	Research assistants · policy Q&A · portfolio Q&A

Anti-pattern: “a human can review if they want” is not human-in-the-loop. It is a wish. Real HITL has a forced gate, an explicit signal, and a fallback path when the human does not act.

USE CASE 08 | **AI portfolio management tools** | Wealth & asset management

Problem

Relationship managers and portfolio analysts answer the same product, tax, market and research questions repeatedly. Quality varies; audit trails are thin; advice boundaries are easy to cross by accident.

Architecture

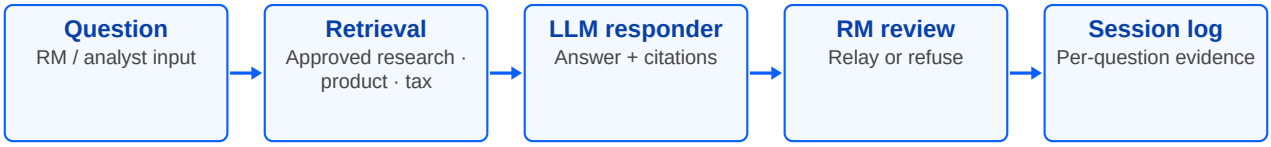

Controls: Internal-only · No advice generation · Strict guardrails

Sample prompt template (excerpt)

```
SYSTEM: You answer wealth-management questions for an internal RM audience.
You may NOT give personalised investment advice. You may summarise approved
research and product / tax content with citations. If a question crosses the
advice line, respond <cannot answer — please consult investment advisor>.

USER: question={text}, client_segment={segment}
context={approved_research_chunks, product_docs, tax_summaries}

OUTPUT JSON: {
  "answer": grounded answer with inline citations,
  "citations": [{source, chunk_id}],
  "refusal_reason": null or one of [advice, off-policy, ungrounded]
}
```

Tool stack

Layer	Tools / patterns commonly seen
Data & retrieval	Approved research repo · product docs · tax summaries · client segments store
Model layer	LLM with refusal behavior · permission-aware retrieval · advice-line classifier
Integration	RM workstation · CRM · advisory-line workflow on refusals
Observability	Per-session log · refusal-reason analytics · drift on advice-line classifier

Observed lift: **RMs handle more accounts; consistent answers across the bank; cleaner audit trail.** · **Maps to modules:** M5 (RAG), M6 (agents), M9 (ethics), M11 (regulation).

INTEGRATION

Integration patterns into existing banking stacks

AI use cases do not live in isolation. They plug into core banking, CRM, case management, DMS, reporting and BI. Four integration patterns dominate across the playbook.

Pattern	Where it fits	Typical interface
Inline copilot	Inside an existing app, sitting next to the user's main workflow.	Browser extension · embedded panel · API-backed widget
Workflow injection	AI step inserted into an existing business workflow.	Workflow engine call · webhook · queue worker
Batch enrichment	Bulk AI processing of cases / documents off the live path.	Scheduled job · staging table · downstream sink
Conversational entry	User starts in a chat / search interface and routes into downstream actions.	Chat front-end · intent router · downstream system handover

Practitioner heuristics

- **Start where the user already is.** Inline copilots adopt fastest.
- **Use batch where latency is forgiving.** Reporting, MRM, regulatory work — almost always batch.
- **Separate read and write planes.** AI inhabits the read plane by default; writes are a separate path.
- **Don't replace the system of record.** The bank's existing systems remain the source of truth.

USE CASE 09

AI regulatory reporting tools

Regulatory reporting

Problem

Regulatory filings combine large structured data sets with long-form narrative. The data is unforgiving; the narrative is dense and dated. Both must reconcile to the structured store, and both must be defensible.

Architecture

Controls: Numbers from store · Per-section sign-off · Evidence retained

Sample prompt template (excerpt)

```
SYSTEM: You draft regulatory-filing narrative sections. All numeric values are
injected from {data_store_json}. You never generate numbers. Cite the report
section / line id for every claim.

USER: filing={code}, section={name}
context={data_store_json, prior_filings_archive, reg_template}

OUTPUT: a markdown section that conforms to the regulator's template, with
inline citations to data_store_json keys and prior_filings_archive ids.
```

Tool stack

Layer	Tools / patterns commonly seen
Data & retrieval	Regulatory data warehouse · prior filings archive · template store
Model layer	LLM with strict template enforcement · numeric injection · reviewer attention markers
Integration	Regulator filing portal · second-line review queue · evidence vault
Observability	Per-section provenance · reviewer overrides logged · regulator response logged

Observed lift: **shorter filing cycles · cleaner reconciliation · fewer late-stage rewrites.** · **Maps to modules:** M3 (data), M7 (structured outputs), M11 (regulation), M14 (capstone).

CERTIFICATION

The certification framework behind the playbook

Every use case in this playbook is built and validated inside the official CGAIFB curriculum. The 14 modules are organised in four phases — Foundations, Applied, Governance, Capstone. Pages 22–23 list every module.

#	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

SYLLABUS · DETAIL

Modules 1–7 · Foundations & Applied

#	Module	Signature lab
M1	GenAI & LLM literacy for finance	30-term finance-flavoured glossary + your own workflow map.
M2	Finance & banking context for AI work	One-page workflow map · 5 controls your AI use case must satisfy.
M3	Data foundations: structured, unstructured, market	Data inventory for your role + retrieval design for 2 doc types.
M4	Prompting patterns & engineering for finance	Finance prompt library (20+ prompts) + side-by-side pattern experiment.
M5	Retrieval-augmented generation in regulated stacks	Mini-RAG prototype on policy / KYC / research corpus with eval scorecard.
M6	Agents & tool use for banking workflows	Agent design doc + single-step agent with strict approval gates.
M7	Structured outputs, integrations & evaluations	Validator + 3 schemas + gold set + cost / latency scorecard.

Modules 1–7 give you the building blocks behind every use case in this playbook. By the end of M7, you have **shipped** a small but real system — not just read about one.

SYLLABUS · DETAIL

Modules 8–14 · Governance & Capstone

#	Module	Signature lab
M8	Model risk & validation for LLM systems	Mini-MRM file on your prototype: card · controls · tests · sign-offs.
M9	Ethics, fairness & responsible AI in finance	Bias audit on a use case + ethics statement + escalation playbook.
M10	Data protection, security & privacy for genAI	Data flow diagram + 15-item security checklist for your prototype.
M11	Regulation & supervisory expectations	Regulatory mapping table + mock regulator interview with your mentor.
M12	Role-mapped capstone design & build	Capstone brief + artifact + 1-page exec summary + 10-minute viva.
M13	Integration, cost, latency & operations	Integration sketch · cost & latency budget · rollout with kill-switch.
M14	Exam preparation, sample exams & viva	Two sample exams under timed conditions + flashcards + mock viva.

Modules 8–14 are the difference between a prototype and a credentialed practitioner. You finish with a defensible capstone, a recognized credential, and a portfolio you can show in interviews.

CAREER

CGAIFB career roadmap

The credential maps onto five career tracks. Most readers move along one of these for the next 12–24 months, then re-skill into the next as their team scales.

Track	Typical entry role	12-month role	24-month role
Build & ship	Fintech AI Specialist · AI Engineer (Finance)	AI Product Engineer	AI Product Lead — Banking
Risk & control	AI Risk Analyst	Model Validation Lead — AI	Head of AI Risk / MRM
Compliance & audit	AI Compliance Officer	AI Audit Lead	Head of AI Compliance
Quant & research	AI Investment Analyst · Quant Analyst (AI)	Senior AI Research Analyst	AI Research Lead
Operations & KYC	AI Banking Ops Specialist	KYC Automation Lead	Head of AI Banking Operations

How playbook readers position themselves

- **Lead with a built artifact.** The capstone, not the badge, is the strongest signal.
- **Pick one use case as a depth area.** Pair the credential with one of the nine workflows here.
- **Speak both AI and risk.** Use the controls catalog on p5 as your interview vocabulary.
- **Show the eval mindset.** Quote your gold-set size and your p95 latency without flinching.

EARNINGS

How use cases map to roles and salary tiers

Indicative tier mapping: which of the nine use cases tends to anchor which role family, and where on the salary band that role sits. Bands vary by region and employer — use as directional benchmarks, not guarantees.

Use case	Anchor role family	Typical tier
01 · Fraud detection software	Fraud / financial crime · AI ops	$$ → $$$$
02 · Credit scoring & memo drafting	Credit analyst · AI lending	$$ → $$$$
03 · AI document processing	Banking ops · KYC automation	$$ → $$$$
04 · SAR & sanctions compliance	AML · compliance · audit	$$ → $$$$$
05 · Generative AI for reporting	FP&A · finance · controlling	$$ → $$$$
06 · Algorithmic trading AI	Quant · markets · research	$$$ → $$$$$
07 · AI risk assessment	Model risk · second line	$$$ → $$$$$
08 · AI portfolio management	Wealth · asset management	$$ → $$$$$
09 · AI regulatory reporting	Reg reporting · finance · compliance	$$ → $$$$

Legend: $ = entry tier · $$ = early-career · $$$ = mid-senior · $$$$ = senior · $$$$$ = principal / lead. The live program page maintains the current regional ranges.

Specialist premium: Practitioners who own a single use case end-to-end — design, build, controls, MRM, audit — consistently sit one tier above generalists of the same seniority.

SHIPPING GUIDE

How readers ship their first use case

Most readers move from playbook to production along the same arc. Whether you're inside a global bank or a fintech, this is the shape the credential walks you through.

Week	What you do	Artifact
W1–2	Pick one of the nine use cases and write a one-page brief.	Use-case brief (problem · users · data · metrics · risks · controls).
W3–5	Build the grounded retrieval layer for your chosen use case.	Mini-RAG prototype + eval scorecard (30+ gold-set items).
W6–7	Add structured outputs, integration sketch and HITL design.	Schema · validator · integration architecture · HITL grid.
W8–9	Wrap with MRM artifacts: model card, controls, drift plan.	Mini-MRM file ready for second-line conversation.
W10–12	Capstone consolidation, viva and exam booking.	Capstone + viva + booked exam slot.
W13	Sit the exam · issue digital credential and badge.	CGAIFB credential issued · badge live on profile.

Reading the playbook teaches you the *shape*. The credential teaches you the *discipline* — and gives you the credential that recruiters and second-line reviewers already screen for.

EMPLOYER VIEW

Employer view & common questions

What hiring teams in finance and banking screen for

- **Built artifacts.** A capstone with an eval scorecard reads differently from a generic AI badge.
- **Controls instinct.** You can name three controls a use case needs without thinking.
- **Pattern fluency.** You can name which of the four patterns on page 7 fits a problem.
- **Globally recognized credential.** Reduces the “unknown candidate” discount on shortlists.

Common questions

Question	Short answer
Do I need to code to do this credential?	No. Engineering tracks lean technical; risk, ops, compliance and credit tracks do not.
Are all nine use cases taught hands-on?	Yes. Every use case in this playbook has a corresponding lab in the curriculum.
Is the credential globally recognized?	Yes — recognized in 100+ countries. The live page maintains current specifics.
How long does it take, realistically?	Most readers finish in 90 days at 5–7 focused hours per week.
What happens after I pass?	Digital certificate, verifiable badge and a lifetime credential record.
Can my employer sponsor me?	Yes. Corporate / cohort sponsorship is available via the live program page.

ENROLLMENT

Download the playbook into your career

You've reached the last page of the playbook. The next step is the easiest one: enrol on the live program page. Your access opens immediately, your cohort is assigned within 24 hours, and you start building one of these nine use cases under mentor guidance.

Step	What happens
1 · Click any CTA in this playbook	You land on the official use-cases 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 · Pick your use case & start Module 1	Choose one of the nine use cases as your capstone target and start building.

Direct contact

Program page: gsdcouncil.org/certification-program/generative-ai-use-cases-in-banking-and-finance
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 use-case playbook. We'll see you inside the cohort, building one of these nine workflows for real.