

# **Generative AI Across the Trading Lifecycle: A Practical Guide**

## Introduction: Beyond Chatbots and Summaries

Financial institutions already use Generative AI for more than chatbots and document summaries. Banks, hedge funds and wealth managers are testing it for market research, portfolio analysis, risk assessment and trading strategy work. The common thread is volume: markets produce far more filings, news, research and data than any team can read.

Generative AI refers to AI systems that create new content, such as text, summaries, scenarios, code and structured insights, based on patterns learned from large datasets. In trading, that means it can read unstructured information (an earnings call transcript, a central bank statement) and turn it into something a person can act on.

**The core idea of this guide: Generative AI is an intelligent layer that supports existing trading systems. It does not decide when to buy or sell.**

Traditional algorithmic trading executes predefined rules or models quickly and consistently. Generative AI adds capabilities around those rules:

- Exploring new strategy ideas and drafting test code
- Reading and summarizing unstructured text at scale
- Generating hypothetical market conditions
- Explaining model outputs in plain language

**Where it fits in the lifecycle**

1. Research and strategy development: generate and refine ideas
2. Market monitoring: track sentiment from news and reports
3. Testing and risk review: stress-test under generated scenarios
4. Oversight and communication: explain signals to risk, compliance and management

## Stage 1: Strategy Development

Generative AI can speed up the early, creative part of strategy work by proposing ideas and writing the first draft of the code that tests them.

### What it can do

- Review historical patterns described by the analyst and suggest strategy hypotheses
- Draft or modify backtesting code (for example, in Python)
- Suggest parameter ranges and variations to test
- Document the logic of a strategy in plain language for later review

### Worked example: a momentum, volatility and sentiment strategy

A trader asks an AI assistant: *"Propose a rules-based equity strategy that combines 12-month momentum, a volatility filter and a news sentiment score. Describe entry and exit rules, then write backtesting code."*

The assistant might return:

- **Entry rule:** buy when 12-month momentum is positive, recent volatility is below its long-run median, and sentiment is neutral or better
- **Exit rule:** exit when momentum turns negative or volatility rises above a set threshold
- **Code:** a script that loads historical prices, calculates the three signals and records simulated trades

The analyst does not trade on this output. The next step is to run the code on historical data, review the logic line by line, and test it on data the strategy has not seen.

### **Pitfalls**

- **Overfitting:** an AI-suggested rule set may look excellent on past data simply because it was tuned to it
- **Unreliable logic or code:** generated code can contain errors, such as using future data in a past calculation (look-ahead bias)
- **Unsupported claims:** the assistant may describe a strategy as "well-known to work" without any evidence

### **Checkpoint: before any AI-generated strategy moves forward**

- ☐ Code reviewed by a developer or quant
- ☐ Backtested on historical data, including out-of-sample periods
- ☐ Assessed for risk, costs and liquidity
- ☐ Approved by a human reviewer, with the decision documented

### **A Step-by-Step Workflow for AI-Assisted Strategy Work**

1. **Define the objective and constraints.** State the market, asset class, holding period, risk limits and trading costs before asking for ideas. Example: "US large-cap equities, holding period of 1 to 3 months, maximum 10% position size, transaction costs of 10 basis points per trade."
2. **Generate hypotheses, not conclusions.** Ask the assistant for several competing ideas and the economic reasoning behind each. A strategy with no plausible reason to work is a warning sign.
3. **Draft the code.** Ask for readable code with comments, so a reviewer can follow every calculation.
4. **Review the code manually.** A developer or quant checks data handling, signal timing and cost assumptions.
5. **Backtest on in-sample data.** Use one historical period to build and tune the strategy.
6. **Validate on out-of-sample data.** Test on a later period the strategy has never seen. A large drop in results suggests overfitting.
7. **Stress and sensitivity test.** Change parameters slightly and add higher costs. A robust strategy should not collapse when a threshold moves from 20 days to 22 days.
8. **Document and decide.** Record the logic, data, results and limitations, then let a human committee approve, reject or request changes.

## Writing Better Prompts

The quality of the output depends heavily on the quality of the request. Specific prompts produce reviewable answers.

| Weak prompt                              | Stronger prompt                                                                                                                                                                                                 | Why it is better                                                |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|
| "Give me a profitable trading strategy." | "Propose three rules-based strategies for US large-cap equities using momentum and volatility, with a holding period of 1 to 3 months. For each, explain the economic rationale and the main failure scenario." | Sets scope, asks for reasoning and forces a discussion of risks |
| "Write backtest code."                   | "Write Python code that loads daily price data, computes a 12-month momentum signal, applies it with a one-day delay, subtracts 10 basis points per trade, and outputs annual return, volatility                | Defines timing and costs, which prevents common errors          |

|                          |                                                                                                           |                                                            |
|--------------------------|-----------------------------------------------------------------------------------------------------------|------------------------------------------------------------|
|                          | and maximum drawdown."                                                                                    |                                                            |
| "Is this strategy good?" | "List the ways this strategy could be overfitted or contain look-ahead bias, and suggest tests for each." | Turns the assistant into a critic instead of a cheerleader |

## Second Worked Example: Catching a Look-Ahead Error in AI-Generated Code

A trader asks for a 20-day moving-average strategy and receives code like this:

# Flawed version: uses today's close to decide today's position

```
df["signal"] = (df["close"] > df["close"].rolling(20).mean()).astype(int)
```

```
df["strategy_return"] = df["signal"] * df["close"].pct_change()
```

The strategy appears to earn excellent returns, but the signal uses today's closing price and is applied to today's return. In real trading, the signal is only known after the close, so the position cannot be held during that same day. A reviewer should catch this and request the fix:

# Corrected version: trade on the next day

```
df["strategy_return"] = df["signal"].shift(1) * df["close"].pct_change()
```

After the correction, the apparent performance usually drops sharply. This is exactly the kind of error a human review is meant to find. You can also ask the assistant itself:

*"Review this code for look-ahead bias and data leakage."* Use that as an extra check, never as a replacement for human review.

### Metrics to Review Before Approval

| Metric            | What it tells you           | What to watch for                                     |
|-------------------|-----------------------------|-------------------------------------------------------|
| Annualized return | Average yearly gain         | Compare after costs, not before                       |
| Volatility        | Size of return swings       | Higher than expected volatility may break risk limits |
| Maximum drawdown  | Largest peak-to-trough loss | Check whether the firm could tolerate it in practice  |
| Sharpe ratio      | Return per unit of risk     | Very high values often signal overfitting             |
| Turnover          | How often positions change  | High turnover raises costs and slippage               |

|                                 |                      |                                                         |
|---------------------------------|----------------------|---------------------------------------------------------|
| In-sample vs. out-of-sample gap | Stability of results | A large gap suggests the strategy was tuned to the past |
|---------------------------------|----------------------|---------------------------------------------------------|

## What to Document

- The prompt and the AI output that inspired the strategy (for traceability)
- The final rules as implemented, in plain language
- Data sources, time periods and cost assumptions
- Test results, including failures and rejected variants
- Named reviewers and the approval date

## Stage 2: Market Sentiment Analysis

Market sentiment reflects how participants feel about an asset or event, and it is shaped by news, social media, corporate announcements, economic data and investor expectations. Generative AI can read and summarize large volumes of this unstructured text far faster than a research team.

### What it can do

- Summarize the tone of news coverage, analyst notes and transcripts around a company or event
- Separate factual announcements from opinion and speculation
- Highlight what changed compared with the previous quarter or announcement

- Surface the specific phrases or topics driving the tone, so a human can check them

### Worked example: sentiment around an earnings release

A company reports quarterly earnings. Within hours there are the press release, the earnings call transcript, several analyst notes and thousands of social media posts. An analyst asks the assistant: *"Summarize the sentiment across these sources. List the three main positive themes, the three main concerns, and any statements where management's tone differs from the written release."*

A useful output could look like this:

- **Positive themes:** revenue growth in a core segment, improved margins, a raised full-year outlook
- **Concerns:** rising operating costs, slower growth in a second segment, cautious comments on customer demand
- **Tone gap:** the release is upbeat, while the call included hedged language about next-quarter orders

The analyst can then combine this with structured data (price moves, volume, volatility) and decide what deserves attention.

### Structured and unstructured inputs

| Input type | Examples | Role in the analysis |
|------------|----------|----------------------|
|            |          |                      |

|              |                                                  |                                          |
|--------------|--------------------------------------------------|------------------------------------------|
| Structured   | Prices, volume, volatility, financial ratios     | Measure what happened                    |
| Unstructured | News, transcripts, analyst reports, social posts | Explain how participants may be reacting |

### Limitations

- **Noise:** social media is full of repeated, low-quality or promotional posts
- **Nuance:** sarcasm, hedged language and industry jargon are easy to misread
- **Misinformation:** false or manipulated stories can distort a summary if sources are not checked
- **Lag and coverage:** the assistant only reflects the sources it is given, so missing sources mean missing signals

**Good practice:** always keep a link from each theme back to its source text so a reviewer can verify it.

### A Step-by-Step Sentiment Workflow

1. **Collect sources.** Choose a defined set: company filings, earnings transcripts, selected news outlets, analyst notes and, if relevant, social media. Record where each item came from.
2. **Clean and deduplicate.** Remove repeated headlines, syndicated copies and obvious spam, so one story does not count ten times.
3. **Tag each item.** Label the company, date, source type and topic (for example, guidance, costs, legal, product).

4. **Ask for structured output.** Request a consistent format, such as a score from -2 (very negative) to +2 (very positive), a one-line reason and a supporting quote reference.
5. **Aggregate.** Combine scores by company, topic and time window, weighting reliable sources more heavily than anonymous posts.
6. **Validate.** Have an analyst compare a sample of AI labels against their own judgment.
7. **Feed into the process.** Use the result as one input alongside price, volume and fundamentals, never as a standalone trigger.

### Example Prompt and Output Format

**Prompt:** *"For each of the following 20 articles about Company X, return: source, date, sentiment score from -2 to +2, the main topic (guidance, costs, legal, product, other), and a one-sentence reason. Do not infer anything that is not in the text. If the article is opinion rather than reporting, mark it as opinion."*

### Illustrative output:

| Source                | Date  | Score | Topic    | Reason                                       |
|-----------------------|-------|-------|----------|----------------------------------------------|
| Financial news site A | Oct 3 | +1    | Guidance | Company raised its full-year revenue outlook |
| Industry blog B       | Oct 3 | -1    | Costs    | Author expects higher input                  |

|                          |       |    |         |                                                                     |
|--------------------------|-------|----|---------|---------------------------------------------------------------------|
|                          |       |    |         | costs to<br>pressure<br>margins<br>(opinion)                        |
| NewsWire C               | Oct 4 | 0  | Product | Announces a<br>product update<br>with no<br>financial detail        |
| Financial news<br>site D | Oct 4 | -2 | Legal   | Reports a<br>regulatory<br>investigation<br>into a business<br>unit |

A consistent format like this lets the team calculate averages, spot sudden shifts and trace any score back to its source.

## Second Worked Example: Reading a Central Bank Statement

Central banks use carefully chosen wording, and small changes can move markets. An analyst provides the assistant with the previous and current policy statements and asks:

*"Compare these two statements. List every phrase that was added, removed or changed, and classify each change as more hawkish (favoring tighter policy), more dovish (favoring easier policy) or neutral. Explain your reasoning briefly."*

A useful answer would contain items such as:

- **Removed:** a phrase describing inflation as "transitory" or "easing" (classified as more hawkish, because it signals less confidence that inflation will fall)
- **Added:** a reference to "labor market remaining tight" (more hawkish)
- **Changed:** "policy will remain appropriate for some time" to "policy will be adjusted as data warrant" (more flexible, possibly more dovish)

The analyst reviews each classification, reads the original sentences, and decides what the overall message is. The assistant saves the time of line-by-line comparison, and the human remains responsible for judging the tone.

## Testing for Consistency

The same article can sometimes receive different scores on different runs. Before relying on sentiment output:

- Run the same 50 articles three times and measure how often the scores match
- Compare AI scores against a small set labeled by experienced analysts and calculate the agreement rate
- Check for systematic errors, such as treating every mention of "risk" as negative
- Re-test whenever the model, the prompt or the source mix changes

## Common Mistakes

- **Counting volume as sentiment.** A flood of posts is not the same as a strong opinion. Weight by source quality.

- **Ignoring timing.** A positive story from three days ago may already be reflected in prices.
- **Mixing facts and opinions.** Ask the assistant to separate them explicitly.
- **Over-trusting a single score.** Look at the distribution and the underlying reasons, not just the average.

## Stage 3: Scenario Generation

Generative models can create hypothetical market conditions to test how a strategy or portfolio might behave when things change. This helps teams look beyond the historical data they already have.

### Common scenario types

- Sudden interest-rate changes
- Spikes in market volatility
- Economic slowdowns or recessions
- Unexpected corporate events, such as a profit warning or leadership change
- Combined events, such as a rate shock during a liquidity squeeze

### Worked example: stress-testing a strategy under a rate shock

A team runs a bond-heavy strategy. They ask: *"Generate three scenarios in which central banks raise rates faster than the market expects over the next six months. For each, describe the likely moves in short- and long-term yields, equity volatility and credit spreads, and state your assumptions."*

The assistant produces three scenarios (for example, mild, moderate and severe), each with a narrative and a set of assumed market moves. The team then:

1. Converts each scenario's assumptions into parameters in their risk or simulation system
2. Runs the strategy through each scenario
3. Compares losses, drawdowns and the points at which risk limits are breached
4. Documents which assumptions drove the results

### Interpreting the results

- Treat scenarios as questions ("what would hurt us?"), not forecasts
- Check that the assumed moves are plausible and consistent with each other
- Look for patterns across scenarios, such as one position causing most of the loss
- Record who reviewed the assumptions and what limits or hedges changed as a result

### Anatomy of a Good Scenario

A scenario is more useful when it is specific. Ask the assistant to describe each of these elements, then have the risk team confirm or adjust them.

| Element | Question it answers | Example |
|---------|---------------------|---------|
|         |                     |         |

|                      |                                  |                                                                                      |
|----------------------|----------------------------------|--------------------------------------------------------------------------------------|
| Trigger              | What starts the event?           | An unexpected inflation report leads to a surprise rate hike                         |
| Magnitude            | How large is the shock?          | Policy rate up 100 basis points in three months                                      |
| Timing               | How fast does it unfold?         | Most of the move happens in the first two weeks                                      |
| Duration             | How long does it last?           | Elevated volatility for six months                                                   |
| Transmission         | How does it reach the portfolio? | Higher yields reduce bond prices, strengthen the currency and pressure growth stocks |
| Second-round effects | What follows?                    | Credit spreads widen as refinancing costs rise                                       |

### Second Worked Example: A Supply-Chain Shock for an Equity Portfolio

A portfolio manager holds stocks in manufacturing, retail and logistics. They ask:

*"Create three scenarios in which a major shipping route is disrupted for one, three and six months. For each, describe likely effects on input costs, delivery times, consumer demand and the sectors in my portfolio. State every assumption and rate how confident you are in each."*

A structured answer might include:

- **One-month disruption:** short-term delays, a modest rise in freight costs, limited earnings impact
- **Three-month disruption:** inventory shortages in manufacturing, higher costs passed to retailers, falling margins in two holdings
- **Six-month disruption:** sustained price increases, weaker consumer demand, and logistics firms benefiting from higher rates while retail and manufacturing suffer

The risk team then converts the narratives into numbers, such as a 15% fall in manufacturing earnings estimates in the severe case, runs the portfolio through them and examines which holdings drive the loss. The assistant's contribution is breadth and speed in imagining the chain of events. The quantitative step stays in the firm's validated risk models.

## Reverse Stress Testing

Instead of asking "what happens if X occurs," reverse the question: *"What combination of market events could cause this portfolio to lose more than 15% within a month?"*

The assistant can propose several plausible chains, such as a sharp equity sell-off combined with widening credit spreads and a liquidity squeeze. The team can then check whether the portfolio's actual exposures match these chains and whether hedges or limits are needed.

## Common Errors in Scenario Work

- **Treating narratives as forecasts.** A scenario shows vulnerability, not probability.
- **Inconsistent assumptions.** For example, assuming sharply higher interest rates alongside rising bond prices without explaining why.
- **Only testing familiar events.** Ask for scenarios unlike anything in recent history, then judge their plausibility.
- **Skipping the numbers.** A story without quantified parameters cannot be run through a risk system.
- **No follow-up.** A scenario exercise should end with decisions, such as new limits, hedges or monitoring triggers.

### Scenario Documentation Template

- **Scenario name and date:**
- **Description and trigger:**
- **Key assumptions (with source or rationale):**
- **Parameters used in the model:**
- **Results (losses, drawdown, limit breaches):**
- **Reviewer and approval:**
- **Actions taken or recommended:**

### Stage 4: Trade Explanation

Algorithmic systems can be difficult for non-technical colleagues to follow. Generative AI can translate technical signals and model outputs into plain-language explanations for risk teams, managers and compliance professionals.

### **Worked example: from signal to explanation**

#### **Technical output:**

*Signal: SELL. Momentum score -1.4; 20-day realized volatility in 92nd percentile; sentiment score -0.6; position limit utilization 78%.*

#### **Plain-language explanation generated by the assistant:**

*"The strategy flagged a sell because recent price momentum has turned clearly negative, volatility is higher than in nearly all of the past year, and news sentiment has worsened. The position is also close to its size limit, so reducing it lowers risk."*

A risk officer can read the second version, ask follow-up questions and decide whether the reasoning matches the strategy's documented rules.

#### **Why this matters for explainability and audit**

- Regulators and internal auditors increasingly expect firms to explain how AI-supported decisions are made
- Plain-language records make reviews and incident investigations faster
- Explanations help non-technical stakeholders challenge a model, not just accept it

**Caution:** the explanation must reflect the real model inputs. If the assistant guesses at reasons the model did not use, the explanation is misleading. Generate explanations

from the model's actual outputs, and have a human spot-check them against the system's logs.

## Tailoring Explanations to the Audience

The same signal needs different explanations depending on who reads it.

| Audience                    | What they need                              | Style of explanation                                                    |
|-----------------------------|---------------------------------------------|-------------------------------------------------------------------------|
| Trader or portfolio manager | Why the signal fired and what to watch next | Concise, with key drivers and thresholds                                |
| Risk officer                | Exposure, limits and what could go wrong    | Focus on position size, volatility, concentration and limit usage       |
| Compliance professional     | Whether the action follows approved rules   | Reference the documented strategy, approvals and any exceptions         |
| Senior management           | Business impact and decision required       | Short summary, no jargon, clear recommendation                          |
| Client (through an advisor) | What happened to their money and why        | Simple language, no technical terms, reviewed by a human before sending |

## A Reusable Explanation Template

Ask the assistant to follow a fixed structure so explanations are consistent and easy to audit.

1. **What happened:** the action or signal in one sentence
2. **Why:** the two or three main drivers, using the model's actual inputs
3. **Risk context:** position size, limit usage, volatility
4. **What is not known:** limitations or data gaps
5. **Next review:** when and by whom the position will be reviewed

## Second Worked Example: Explaining a Blocked Trade

An automated system blocks a purchase. The log reads:

*Order rejected. Rule 14: sector exposure limit exceeded (Technology 31% vs. limit 30%).*

*Pre-trade risk check failed.*

The assistant generates an explanation for the portfolio manager:

*"The system did not place this purchase because it would have pushed the portfolio's technology exposure to 31%, above the 30% limit. To proceed, you could reduce another technology holding or request an approved exception."*

A second version for the compliance file might read:

*"Order rejected under pre-trade control Rule 14 (sector concentration). Technology exposure would have been 31% against a 30% limit. No exception was requested. No further action required."*

Both versions come from the same log entry, which is why the explanation must always be generated from, and checked against, the real system output.

## Checking That an Explanation Is Faithful

A fluent explanation is not necessarily a correct one. Before relying on explanations:

- Compare a sample against the system logs and model inputs
- Check that every number mentioned appears in the source data
- Make sure the explanation does not invent reasons, such as citing news the model never used
- Have compliance review the template wording once, then spot-check periodically
- Record which explanations were generated by AI and who reviewed them

## End-to-End Example: One Signal Across All Four Stages

The following illustrative case shows how the four stages connect. The company and numbers are fictional.

**Stage 1: Strategy development.** A quant team asks an AI assistant to propose a strategy for mid-cap industrial companies that combines price momentum, volatility and news sentiment. The assistant drafts the rules and the backtesting code. A developer finds and fixes a look-ahead error, the strategy is tested in-sample and out-of-sample, and an investment committee approves a small pilot allocation.

**Stage 2: Market sentiment analysis.** Each morning the system collects news and earnings-call transcripts for the pilot universe. The assistant scores each item from -2 to +2 with a one-line reason. Analysts review a daily sample for accuracy. One day, sentiment for "Company A" drops from +0.8 to -0.9 after a profit warning.

**Stage 3: Scenario generation.** Before increasing exposure, the risk team asks the assistant to generate scenarios for the industrial sector: a demand slowdown, a spike in raw material costs and a sudden rate rise. They convert the assumptions into model parameters and find that a combined cost spike and slowdown would breach the pilot's drawdown limit. The committee lowers the maximum position size.

**Stage 4: Trade explanation.** The strategy generates a sell signal on Company A. The assistant produces a plain-language note: momentum is negative, volatility is elevated, sentiment has deteriorated after the profit warning, and the position is near its reduced limit. The risk officer checks the note against the logs, approves the reduction and files the explanation.

### What the example shows

- The AI contributed speed and breadth at every stage
- Humans reviewed code, sampled sentiment scores, confirmed scenario assumptions and approved the trade
- Every step left a record that could be reviewed later

## Stage-by-Stage Summary

Across all four stages, the benefit is speed and breadth, and the control is always a human checkpoint.

| Stage | Generative AI | Key benefit | Main risk | Human      |
|-------|---------------|-------------|-----------|------------|
|       | use           |             |           | checkpoint |

|                           |                                              |                                     |                                          |                                              |
|---------------------------|----------------------------------------------|-------------------------------------|------------------------------------------|----------------------------------------------|
| Strategy development      | Generate strategy ideas and backtesting code | Faster idea exploration             | Overfitting, flawed code                 | Code review, out-of-sample testing, sign-off |
| Market sentiment analysis | Summarize news, transcripts and reports      | Faster reading of unstructured data | Noise, misread tone, misinformation      | Source verification by an analyst            |
| Scenario generation       | Create hypothetical market conditions        | Broader stress testing              | Implausible or inconsistent assumptions  | Risk team reviews assumptions                |
| Trade explanation         | Translate signals into plain language        | Clearer oversight and audit trail   | Explanations that do not match the model | Spot-check against system logs               |

## Risks and Guardrails

Financial institutions operate in heavily regulated environments, so the benefits of Generative AI have to be balanced against specific risks. Set clear rules in advance for where AI may be used and where human approval is mandatory.

## Key risks

- **Accuracy and hallucination:** models can state invented facts, figures or sources with confidence. Example: an assistant cites an earnings figure that does not appear in the filing.
- **Algorithmic bias:** models learn from historical data that may carry gaps or bias, which can produce unfair or unreliable outputs, especially in customer-facing work such as credit or advice.
- **Data privacy and security:** client and trading data must not be pasted into tools without approved access controls.
- **Regulatory compliance:** record-keeping, suitability and disclosure requirements still apply to AI-supported work.
- **Explainability and auditability:** firms need to show how AI outputs were used in important decisions.
- **Third-party AI risk:** external AI providers add dependency, data-handling and continuity questions.

## Guardrails that work

- Define approved and prohibited uses, such as "AI may draft research summaries, but may not place orders"
- Require human sign-off for anything that affects client outcomes or capital
- Test models regularly for accuracy and bias, and keep the results
- Log prompts, outputs and approvals so decisions can be reconstructed later
- Review vendors for security, data handling and regulatory fit before adoption

## Risk Register Example

| Risk                   | Example in practice                                                                                   | Mitigation                                                                                 | Suggested owner      |
|------------------------|-------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|----------------------|
| Hallucinated data      | The assistant quotes a revenue figure that is not in the filing                                       | Require source references; verify figures against the original document                    | Research lead        |
| Algorithmic bias       | A client-profiling tool consistently rates one group as higher risk because of skewed historical data | Test outcomes across groups; review training data; keep human review for adverse decisions | Model risk team      |
| Data leakage           | An employee pastes confidential client details into an unapproved tool                                | Approved tools only; access controls; staff training; monitoring                           | Information security |
| Prompt or output drift | A prompt change alters sentiment                                                                      | Version control for prompts; regular re-testing                                            | Model owner          |

|                   |                                                          |                                                        |                            |
|-------------------|----------------------------------------------------------|--------------------------------------------------------|----------------------------|
|                   | scores without anyone noticing                           |                                                        |                            |
| Over-reliance     | Analysts accept AI summaries without reading the sources | Spot-check requirements; require source links          | Team managers              |
| Vendor dependency | An AI provider changes a model version or has an outage  | Contract terms, fallback processes, version monitoring | Procurement and technology |

## Roles and Responsibilities

- **Business owner:** decides where AI is used and accepts responsibility for the outcomes
- **Model or AI owner:** maintains prompts, tests, versions and performance records
- **Risk and compliance:** sets rules, reviews high-impact uses, monitors for breaches
- **Information security:** controls data access and vendor security
- **Internal audit:** periodically checks that controls work in practice
- **Human reviewers:** validate outputs and sign off on decisions

## When Something Goes Wrong: A Short Example

An assistant summarizing earnings reports reports a company's profit as higher than actual because it confused two reporting periods. An analyst notices before publishing.

A good process then:

1. Records the error and the prompt that caused it
2. Corrects the output and checks whether other summaries used the same source
3. Adjusts the prompt, for example by requiring the reporting period to be stated for every figure
4. Re-tests on a sample of previous reports
5. Updates the team guidance and logs the incident for review

## Implementation Checklist

Use this list when planning a first pilot.

- ☐ Pick one narrow use case, such as earnings-call summaries or scenario drafts
- ☐ Define success measures, such as time saved or error rate, before starting
- ☐ Confirm which data the tool may and may not access
- ☐ Document what the AI may do and where human approval is required
- ☐ Assign an owner for accuracy and bias monitoring
- ☐ Set up logging of prompts, outputs and reviewer decisions
- ☐ Brief compliance and risk teams before launch

- ☐ Run the pilot in a controlled environment, with no live trading or client impact
- ☐ Review results, lessons and incidents, then decide whether to expand

### A Phased Rollout Plan

| Phase    | Timing         | Focus                                                            | Exit criteria                                       |
|----------|----------------|------------------------------------------------------------------|-----------------------------------------------------|
| Prepare  | Weeks 1 to 4   | Choose the use case, set rules, select tools, brief stakeholders | Approved scope and governance rules                 |
| Pilot    | Weeks 5 to 12  | Run in a controlled setting with human review of every output    | Accuracy, time savings and risk findings documented |
| Evaluate | Weeks 13 to 16 | Compare results against the success measures, fix issues         | Sign-off from business, risk and compliance         |
| Scale    | After approval | Expand to additional teams or use cases, with                    | Ongoing monitoring and periodic review in place     |

|  |  |                         |  |
|--|--|-------------------------|--|
|  |  | continued<br>monitoring |  |
|--|--|-------------------------|--|

## Next Steps

Generative AI works best when the people using it understand both financial concepts and the technology's limits. As adoption grows across trading, wealth management, banking and finance operations, professionals benefit from structured learning that covers use cases, risks and governance.

- GSDC's Generative AI certification builds structured knowledge of Generative AI concepts, applications, risks and governance
- The Generative AI in Finance and Banking Certification focuses on AI use cases in banking, investments, financial operations, risk management and decision-making

These learning paths suit finance professionals, technology teams, risk managers and business leaders involved in AI transformation.

## Glossary

| Term                | Meaning                                                            |
|---------------------|--------------------------------------------------------------------|
| Algorithmic trading | Executing trades using predefined rules or models run by computers |

|                       |                                                                                |
|-----------------------|--------------------------------------------------------------------------------|
| Backtesting           | Testing a strategy on historical data to see how it would have performed       |
| Out-of-sample testing | Testing on data not used to build or tune the strategy                         |
| Overfitting           | Tuning a strategy so closely to past data that it fails on new data            |
| Look-ahead bias       | Using information in a test that would not have been available at the time     |
| Drawdown              | The decline from a peak value to a later low                                   |
| Stress test           | Evaluating how a portfolio behaves under extreme or unusual conditions         |
| Hallucination         | A confident but false or invented statement from an AI model                   |
| Explainability        | The ability to describe how and why a model produced a result                  |
| Algorithmic bias      | Systematic unfairness or error in model outputs, often inherited from the data |

*This guide is for educational purposes only and does not constitute financial, investment, legal or compliance advice.*

| Item              | Detail                                                     |
|-------------------|------------------------------------------------------------|
| Version           | 1.0                                                        |
| Date              | [Insert date]                                              |
| Prepared by       | [Your blog / organization name]                            |
| Intended audience | AI, risk, compliance, security, data, and business leaders |

## How to Use This Kit

Artificial intelligence is moving from experimentation into business-critical work. As it does, governance can no longer sit with a single technology or compliance team. This kit gives you the core documents needed to set up a cross-functional AI Governance Committee and get it running in its first three months.

## What is inside

| Part | Document                             | Use it to                                                        |
|------|--------------------------------------|------------------------------------------------------------------|
| 1    | Committee Charter Template           | Define purpose, scope, membership, and decision rights           |
| 2    | RACI Matrix                          | Clarify who is Responsible, Accountable, Consulted, and Informed |
| 3    | Meeting Agenda and Minutes Templates | Run consistent meetings and record decisions                     |
| 4    | First 90 Days Checklist              | Move from kickoff to an operating governance process             |

## Suggested order of use

Work through the parts in sequence, because each one builds on the last:

- **Start with the Charter (Part 1).** It settles authority and membership. Without it, the RACI and agendas have nothing to anchor to.
- **Adapt the RACI (Part 2)** once you know who sits on the committee.
- **Adopt the meeting templates (Part 3)** before your first meeting so decisions are documented from day one.
- **Follow the 90-day checklist (Part 4)** to track progress and keep momentum.

## How to customize

Text in [square brackets] is a placeholder to replace. Example wording is provided throughout, but every organization is different. A financial institution will need deeper legal, risk, and compliance involvement than a ten-person software company, and a global enterprise may need regional or business-unit representatives. Treat this kit as a starting point, not a finished policy.

**Tip:** Keep the committee working alongside existing enterprise risk, cybersecurity, privacy, data governance, procurement, and compliance functions. The goal is to connect them for AI decisions, not to build a separate parallel bureaucracy.

# Part 1: AI Governance Committee Charter

## Template

The charter is the committee's founding document. It states why the committee exists, what it may decide, who sits on it, and how it reports. A clear charter prevents the two most common early failures: a committee with no real authority, and a committee that tries to review everything and becomes a bottleneck.

Each section below explains what to include, gives example wording, and points out common pitfalls.

# 1. Purpose and Mandate

State in two or three sentences why the committee exists. The purpose should go beyond approving or rejecting AI projects. It should describe the committee as the mechanism for responsible, accountable AI adoption.

**Example:** "The AI Governance Committee is a cross-functional body responsible for overseeing the organization's use and management of artificial intelligence. It establishes policies, accountability mechanisms, and review processes that enable responsible AI adoption while managing legal, security, privacy, ethical, and operational risk."

Also name the authority under which the committee operates, for example the executive leadership team or the risk committee of the board.

# 2. Scope

Define exactly what falls under the committee's oversight. Be explicit about the range of AI, since risk now comes from far more than in-house models.

- **Built AI:** models and applications developed internally
- **Bought AI:** third-party AI products and services, including AI features embedded in existing SaaS tools
- **Generative AI:** chatbots, coding assistants, and content generation tools
- **Agentic and automated workflows:** systems that take actions or make decisions with limited human input
- **Shadow AI:** unapproved AI tools employees may be using without authorization

**Example:** A customer support team enables an AI summarization feature that ships inside its ticketing platform. Nobody "bought an AI tool," yet the feature processes customer data through a third-party model. A scope statement that covers embedded AI features brings this under review.

Consider also listing what is out of scope, such as basic rules-based automation, to avoid overreach.

# 3. Objectives

List four to six measurable goals. Linking each objective to an outcome makes it easier to report progress later.

- Establish and maintain AI policies and standards
- Maintain a complete, current AI inventory
- Apply risk-based review to AI use cases before deployment
- Define and monitor AI risk tolerance
- Provide visibility into unauthorized AI use and address it constructively
- Report significant AI risks to senior leadership

## 4. Membership and Roles

Membership should bring in the expertise needed to assess AI across departmental boundaries. No single department has all of it. A typical committee has seven roles:

| Role                        | Core responsibilities                                                                                     | Example of their input                                                            |
|-----------------------------|-----------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|
| <b>Executive Sponsor</b>    | Approves governance strategy, sets or approves risk tolerance, provides resources, escalates to the board | Decides the organization will not use AI for automated hiring rejections          |
| <b>AI Governance Chair</b>  | Sets the agenda, coordinates risk reviews, maintains documentation, tracks decisions and remediation      | Schedules an urgent review when a business unit finds an unregistered AI tool     |
| <b>Legal and Compliance</b> | Assesses regulatory, contractual, IP, and transparency requirements                                       | Flags that a vendor's terms allow training on customer data                       |
| <b>Cybersecurity</b>        | Evaluates access controls, third-party AI services, prompt injection, and incident response               | Requires single sign-on and logging before a chatbot connects to internal systems |
| <b>Data Governance</b>      | Sets data classification, quality, provenance, retention, and privacy requirements                        | Prohibits uploading restricted data to external AI services                       |

|                            |                                                                              |                                                                                   |
|----------------------------|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|
| <b>AI/ML or Technology</b> | Model evaluation, testing, monitoring, documentation, and technical controls | Requires a bias and accuracy test before a model goes live                        |
| <b>Business or Product</b> | Defines objectives, users, acceptable performance, and human oversight needs | Explains that a tool will be used by frontline staff with no technical background |

Add each member's name, title, and term in a roster table at the end of the charter. Where possible, name an alternate for every seat so that meetings are not cancelled when one person is unavailable.

**Tip:** The chair can come from AI, risk, compliance, technology, or another function. Choose someone with the time, credibility, and cross-functional relationships to follow decisions through, rather than simply the most senior person available.

## 5. Responsibilities

Document the committee's standing duties. These are the ongoing work, distinct from the objectives above.

- **Policies and standards:** approved tools, prohibited use cases, sensitive-data restrictions, human oversight, third-party requirements, incident reporting, and employee responsibilities. Review them periodically as capabilities and regulations change.
- **AI inventory:** record each system's business owner, purpose, model or provider, data used, risk classification, deployment environment, user population, monitoring needs, and review date.
- **Risk management:** apply a consistent process to identify, assess, treat, and monitor risks such as privacy, security, bias, inaccurate outputs, lack of transparency, IP concerns, and third-party dependencies.
- **Shadow AI governance:** understand why employees turn to unapproved tools and offer safer alternatives rather than only blocking them.
- **Incidents and exceptions:** review AI incidents and approve or deny policy exceptions.
- **Post-deployment monitoring:** confirm that approved systems are monitored for performance, security, and compliance.

**Example:** Employees are pasting meeting notes into a public chatbot to summarize them. Rather than issuing a blanket ban, the committee approves an enterprise-licensed assistant with data protections, publishes guidance on what may be entered, and adds the tool to the approved registry.

## 6. Decision Rights and Authority

This is the section that gives the charter teeth. Spell out which decisions the committee can make on its own, which need executive approval, and which are delegated to business owners.

| Committee decides                       | Escalate to Executive Sponsor              | Delegated to business owner                        |
|-----------------------------------------|--------------------------------------------|----------------------------------------------------|
| Approve or reject medium-risk use cases | High-risk or prohibited-category use cases | Day-to-day operation of approved low-risk tools    |
| Approve standard AI policy updates      | Changes to organizational risk tolerance   | User training and adoption                         |
| Grant time-limited exceptions           | Exceptions affecting regulated data        | Routine performance checks per the monitoring plan |

Also state what happens if the committee cannot reach a decision, for example a simple majority with the chair holding a casting vote, or automatic escalation to the sponsor.

## 7. Meeting Cadence and Quorum

- **Frequency:** monthly is common at the start, moving to quarterly once processes mature, with ad hoc sessions for urgent reviews.
- **Quorum:** for example, the chair plus at least four members, which must include legal or compliance and cybersecurity.
- **Decision rule:** consensus first, then majority vote if needed.
- **Urgent reviews:** define a fast path, such as a written review by three members within five business days, for time-sensitive requests.

**Tip:** A slow committee drives people to shadow AI. Publish target turnaround times, for example five business days for low-risk requests, and track them.

## 8. Escalation Path

Describe when and how issues travel upward. Triggers might include a serious AI incident, a regulatory inquiry, a disagreement the committee cannot resolve, or a proposed use that exceeds the approved risk tolerance.

**Example:** An AI-assisted customer chatbot begins giving customers incorrect refund information. The chair convenes an urgent review, security and legal assess the exposure, and because customers were affected the chair escalates to the Executive Sponsor within one business day.

## 9. Reporting and Documentation

Set out what the committee records and what it reports. Good documentation is what turns governance from opinion into evidence.

- **Records to keep:** approvals, assessments, conditions, exceptions, incidents, and monitoring outcomes
- **Regular report to leadership:** number of AI systems by risk tier, use cases reviewed, open exceptions, incidents, and shadow AI findings
- **Retention:** how long records are kept and where they are stored

## 10. Charter Review and Amendment

AI capabilities and regulations change quickly. Commit to reviewing the charter at least annually, and after any major regulatory change or significant incident. State who may propose amendments and who approves them.

## 11. Approval Block

| Name and title        | Signature | Date |
|-----------------------|-----------|------|
| [Executive Sponsor]   |           |      |
| [AI Governance Chair] |           |      |

## Part 2: RACI Matrix

A RACI matrix removes ambiguity about who does what. When responsibilities are unclear, tasks are duplicated, skipped, or argued over at the worst possible moment, usually during an incident.

### Legend

| Code     | Meaning     | What it looks like in practice                                  |
|----------|-------------|-----------------------------------------------------------------|
| <b>R</b> | Responsible | Does the work. There may be several.                            |
| <b>A</b> | Accountable | Owns the outcome and has final sign-off. Only one per activity. |
| <b>C</b> | Consulted   | Provides input before a decision. Two-way communication.        |
| <b>I</b> | Informed    | Told the outcome after the fact. One-way communication.         |

### Sample matrix

The matrix below maps common governance activities to the seven committee roles. Abbreviations: **ES** Executive Sponsor, **Chair** AI Governance Chair, **L&C** Legal and Compliance, **Sec** Cybersecurity, **Data** Data Governance, **Tech** AI/ML or Technology, **Biz** Business or Product.

| Activity                     | ES | Chair | L&C | Sec | Data | Tech | Biz |
|------------------------------|----|-------|-----|-----|------|------|-----|
| Approve AI governance policy | A  | R     | C   | C   | C    | C    | I   |
| Set AI risk tolerance        | A  | R     | C   | C   | C    | C    | C   |

|                                |   |     |   |   |   |     |     |
|--------------------------------|---|-----|---|---|---|-----|-----|
| Register a new AI use case     | I | I   | I | I | I | C   | A/R |
| Conduct risk assessment        | I | A   | R | R | R | R   | C   |
| Classify risk level            | I | A   | C | C | C | R   | C   |
| Approve or reject use case     | I | A/R | C | C | C | C   | I   |
| Implement technical controls   | I | I   | I | C | C | A/R | C   |
| Monitor deployed systems       | I | A   | C | R | R | R   | R   |
| Manage AI incidents            | I | A   | R | R | C | R   | C   |
| Approve exceptions             | A | R   | C | C | C | C   | I   |
| Review shadow AI findings      | I | A   | C | R | C | C   | C   |
| Report to executive leadership | A | R   | C | C | C | C   | I   |

## Reading the matrix: two examples

**Example: Manage AI incidents.** The chair is Accountable, meaning they make sure the response happens and is closed out. Legal, security, and technology are Responsible for the hands-on work: assessing legal exposure, containing the issue, and fixing the system. Data and business are Consulted for context. The sponsor is Informed.

**Example: Register a new AI use case.** The business owner is both Accountable and Responsible, because the team proposing the AI is best placed to describe what it will do, who will use it, and what data it will touch. Technology is Consulted to help complete the technical details.

## Guidance notes

- **One Accountable per activity.** If two people are accountable, nobody is. A/R means one role is both Accountable and Responsible.
- **Avoid over-consulting.** If everyone is Consulted on everything, decisions stall. Consult only those whose expertise changes the outcome.
- **Make sure every row has at least one R.** Someone must actually do the work.
- **Adapt for smaller organizations.** One person may hold several seats, for instance a head of IT covering both security and technology. Combine columns rather than leaving roles unstaffed.
- **Adapt for larger organizations.** Add columns for regional leads, procurement, privacy, or internal audit, and add rows for activities such as third-party AI vendor onboarding.
- **Match RACI to decision rights.** The Accountable role for approvals should agree with the decision-rights table in your charter.

## Blank template

| Activity | ES | Chair | L&C | Sec | Data | Tech | Biz |
|----------|----|-------|-----|-----|------|------|-----|
|          |    |       |     |     |      |      |     |
|          |    |       |     |     |      |      |     |
|          |    |       |     |     |      |      |     |
|          |    |       |     |     |      |      |     |
|          |    |       |     |     |      |      |     |
|          |    |       |     |     |      |      |     |

## Part 3: Meeting Agenda and Minutes

### Templates

Consistent meetings produce consistent decisions. These templates keep the committee focused on governance work and leave a clear record of what was decided and why.

## 1. Standard Meeting Agenda (60 minutes)

| Time  | Item                            | Lead            | Purpose                                                              |
|-------|---------------------------------|-----------------|----------------------------------------------------------------------|
| 0-5   | Opening and attendance          | Chair           | Confirm quorum and note conflicts of interest                        |
| 5-10  | Review of previous actions      | Chair           | Confirm what is complete, overdue, or blocked                        |
| 10-30 | New AI use case reviews         | Business owners | Decide on submitted use cases using the review summary sheet         |
| 30-38 | High-risk system updates        | Tech / Sec      | Report on monitoring results and changes to high-risk systems        |
| 38-45 | Incidents and exceptions        | Sec / L&C       | Review incidents since the last meeting and any requested exceptions |
| 45-50 | Shadow AI and inventory updates | Chair / Data    | Share discovery findings and inventory gaps                          |
| 50-55 | Regulatory and policy updates   | L&C             | Flag upcoming changes that affect policy or use cases                |
| 55-60 | Decisions and action items      | Chair           | Read back decisions, owners, and due dates                           |

**Tip:** Circulate papers at least three business days before the meeting so members arrive prepared. Reserve meeting time for discussion and decisions, not for reading.

## 2. Use Case Review Summary Sheet

Complete one sheet per AI use case. It gives members the same information in the same format every time, which speeds up review and makes comparisons easier.

| Field                               | Entry                                                                        |
|-------------------------------------|------------------------------------------------------------------------------|
| <b>Use case name</b>                | [e.g., Customer support reply drafting assistant]                            |
| <b>Business owner</b>               | [Name, department]                                                           |
| <b>Purpose and expected outcome</b> | [What the system will do and the business result expected]                   |
| <b>Users</b>                        | [Who will use it and how many]                                               |
| <b>Model or provider</b>            | [Internal model, or vendor and product name]                                 |
| <b>Data processed</b>               | [Data types and classification, including any personal or confidential data] |
| <b>Level of automation</b>          | [Advisory only / human review required / fully automated]                    |
| <b>Regulatory exposure</b>          | [Applicable laws, regulations, contractual obligations]                      |
| <b>Risk level</b>                   | [Low / Medium / High, with rationale]                                        |
| <b>Key risks identified</b>         | [Privacy, security, accuracy, bias, third-party dependency, etc.]            |
| <b>Required controls</b>            | [Access restrictions, human review, logging, testing, training]              |
| <b>Decision</b>                     | [Approved / Approved with conditions / Deferred / Rejected]                  |
| <b>Conditions and review date</b>   | [Conditions attached and when the system will be reviewed again]             |

### Completed example

| Field | Entry |
|-------|-------|
|-------|-------|

|                                     |                                                                                                                 |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Use case name</b>                | Customer support reply drafting assistant                                                                       |
| <b>Business owner</b>               | Priya Nair, Head of Customer Support                                                                            |
| <b>Purpose and expected outcome</b> | Draft first-pass replies to customer tickets, reducing average response time by 30%                             |
| <b>Data processed</b>               | Customer names, order details, and free-text ticket content (personal data)                                     |
| <b>Model or provider</b>            | Third-party large language model accessed through the ticketing platform                                        |
| <b>Level of automation</b>          | Human review required before any reply is sent                                                                  |
| <b>Risk level</b>                   | Medium: personal data and a third-party provider, but a human reviews every output                              |
| <b>Decision</b>                     | Approved with conditions                                                                                        |
| <b>Conditions and review date</b>   | Vendor confirms customer data is not used for training; reply logs retained for 12 months; review in six months |

### 3. Minutes and Decision Log Template

Minutes should capture decisions and rationale, not a transcript. The decision log matters most, because it is the evidence you will point to in an audit or after an incident.

#### Meeting record

| Field                        | Entry                           |
|------------------------------|---------------------------------|
| <b>Meeting date and time</b> | [Date, time]                    |
| <b>Attendees / apologies</b> | [Names, roles; note alternates] |
| <b>Quorum confirmed</b>      | [Yes / No]                      |

|                                       |                  |
|---------------------------------------|------------------|
| <b>Conflicts of interest declared</b> | [None / details] |
|---------------------------------------|------------------|

## Decision log

| Ref   | Decision                                           | Owner   | Rationale and conditions                                     | Due    | Status |
|-------|----------------------------------------------------|---------|--------------------------------------------------------------|--------|--------|
| D-001 | Approved support reply assistant with conditions   | P. Nair | Medium risk; vendor no-training clause required              | 15 Nov | Open   |
| D-002 | Rejected use of public chatbot for contract review | Legal   | Confidential data exposure; enterprise tool proposed instead | N/A    | Closed |
|       |                                                    |         |                                                              |        |        |

## Action items

| Action                                                | Owner   | Due date | Status |
|-------------------------------------------------------|---------|----------|--------|
| [e.g., Confirm vendor data-training terms in writing] | [Legal] | [Date]   | [Open] |
|                                                       |         |          |        |

# Part 4: First 90 Days Checklist

A committee that meets but never operationalizes its processes quickly loses credibility. This checklist follows the lifecycle logic used throughout AI governance (identify, assess, classify, approve, deploy, monitor, and review) and builds the committee's capability in three 30-day phases.

## Days 1-30: Foundation

**Goal:** establish authority, people, and a first view of where AI is being used.

- ☐ Secure an executive sponsor and confirm their time commitment
- ☐ Identify and invite committee members from each function, plus alternates
- ☐ Adapt and adopt the charter, with sponsor sign-off
- ☐ Draft an initial AI inventory from IT records, procurement data, and department surveys
- ☐ Hold the first meeting and agree on meeting cadence

**Example:** The initial inventory will be incomplete, and that is fine. A common first finding is that teams are using AI features embedded in existing software that nobody thought of as "AI". Record what you find and refine over time.

## Days 31-60: Structure

**Goal:** build the working tools of governance.

- ☐ Finalize the RACI matrix and share it with all stakeholders
- ☐ Define risk classification criteria (impact, data sensitivity, user population, automation, regulatory exposure, business criticality)
- ☐ Launch a use case intake process with a simple registration form
- ☐ Draft or update the AI policy: approved tools, prohibited uses, sensitive-data rules, human oversight
- ☐ Begin shadow AI discovery using SaaS discovery tools, network monitoring, and an employee reporting channel

**Example:** A three-tier classification works well to start. **Low:** internal productivity tools with no sensitive data. **Medium:** tools processing personal or confidential data with human review. **High:** systems that influence decisions about people, operate with significant automation, or face heavy regulation.

## Days 61-90: Operationalize

**Goal:** run real use cases through the full process and start reporting.

- ☐ Review the first use cases end to end, from registration to documented decision
- ☐ Set the monitoring approach and reporting cadence for approved systems
- ☐ Establish incident and exception handling procedures, including the escalation path
- ☐ Deliver the first report to executive leadership
- ☐ Schedule a 90-day retrospective to review what worked and what to change

## Success indicators

Track a few simple measures so you can show progress and spot problems early:

| Indicator                  | What good looks like at 90 days                                                  |
|----------------------------|----------------------------------------------------------------------------------|
| Inventory coverage         | Known AI systems recorded, with owners and risk tiers assigned                   |
| Use case reviews completed | First use cases taken through the full workflow                                  |
| Review turnaround time     | Within the published target (for example, five business days for low risk)       |
| Shadow AI visibility       | Discovery under way, with approved alternatives identified for the top use cases |
| Meeting attendance         | Quorum met at every meeting                                                      |
| Leadership reporting       | First report delivered and acted on                                              |

## Next Steps and References

### Where to go from here

- Build out your AI inventory and risk register so the committee has a reliable basis for decisions. Our AI Inventory and Risk Register Template is a natural companion to this kit.
- Run a pilot review on one real use case to test the charter, RACI, and templates before rolling out more widely.
- Revisit the documents after your 90-day retrospective and adjust them based on what you learned.

## Framework references

- **NIST AI Risk Management Framework (AI RMF)** describes governance as a cross-cutting function informing AI risk management throughout the lifecycle, emphasizing policies, accountability structures, defined roles, and AI inventories.
- **NIST Generative AI Profile** identifies risks specific to or intensified by generative AI and recommends oversight, testing, and documentation.
- **ISO/IEC 42001:2023** sets requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (AIMS).

## Build deeper skills

If you want structured, practical training in AI governance, the GSDC Certified AI Governance Professional certification covers frameworks such as ISO/IEC 42001 and NIST AI RMF, along with AI inventories, risk registers, governance committees, RACI structures, impact assessments, and generative AI governance.

**Disclaimer:** The templates in this kit are general starting points and do not constitute legal advice. Adapt them to your organization's size, industry, and legal and regulatory context, and consult qualified legal counsel where appropriate.

# CERTIFICATION IN GENERATIVE AI IN FINANCE AND BANKING

LEARN DIRECTLY FROM GLOBAL PRACTITIONERS, AI AND FINTECH EXPERTS, AND INDUSTRY LEADERS WHO ARE SHAPING THE FUTURE OF AI-DRIVEN FINANCIAL SERVICES, DIGITAL BANKING INNOVATION, AND DATA-POWERED FINANCIAL DECISION-MAKING.



## ABOUT GSDC CERTIFICATION



### EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.



### EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.



### CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.



### LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

## LEARNING OBJECTIVE

- Build real-world expertise through 36+ hours of expert-led sessions
- Earn a globally recognized AI professional certificate that employers trust and prove your skills.
- Understand governance in AI banking adoption

Enroll now with the  
code **LEARN20** To  
avail **20%** discount

**Enroll Now**



[www.gsdccouncil.org](http://www.gsdccouncil.org)