GrowCFO × AccountsIQ

# SECURE AI ADOPTION IN FINANCE

A practical CFO playbook for protecting company data without killing AI adoption

**01**
**Secure tools**

**02**
**Understand data flows**

**03**
**Draw the automation line**

ABOUT THIS PLAYBOOK

# Built by finance and AI experts, for finance leaders

This playbook was co-developed by GrowCFO and AccountsIQ as a companion to the roundtable “Can Finance Teams Trust AI?”. It turns the questions finance leaders are already asking into practical exercises, checklists and templates you can use the same week.

**The promise:** by the end of this guide, a CFO should know exactly what to ask the team, what to ask vendors, what to stop, and what to approve safely.

## CONTRIBUTORS




**James Hunter**
CFO, AccountsIQ



**Kevin Appleby**
COO & Podcast Host, GrowCFO



**Pascal Virmaud**
Director of Engineering, AccountsIQ



**Guy Weaver**
AI Expert, GrowCFO

PRACTICAL SPINE

# Start here: five next-day exercises

1 NEXT DAY

### Find actual AI use

Ask the finance team which AI tools they used last week and what data they entered.

2 NEXT DAY

### Create one safe lane

Approve one low-risk AI use case so governance feels enabling, not restrictive.

3 NEXT DAY

### Set three red lines

Name three data types that must not go into unapproved AI tools.

4 NEXT DAY

### Ask one vendor

Ask whether prompts, uploads, transaction data or outputs are used for model training.

5 NEXT DAY

### Place one workflow on the ladder

Decide whether AI is drafting, analyzing, recommending, triggering or executing.

*These exercises are repeated and expanded throughout the playbook so the reader can act as they read.*

DO THIS BEFORE WRITING A POLICY

# Exercise 1: find the AI use you already have

| Capture | Examples |
|---|---|
| **Tool used** | ChatGPT, Copilot, Gemini, Claude, finance system AI, transcription tool, spreadsheet add-in |
| **Task** | Drafting, summarizing, analyzing, forecasting, coding, reconciling, emailing |
| **Data entered** | Public, internal, confidential, personal, financial, customer, supplier |
| **Output used?** | Draft only, reviewed, copied into report, sent externally, used to make decision |

ASK THE TEAM — NO JUDGMENT

Send a short message: “Where did you use AI last week? Which tool did you use? What data did you put in? Did the output influence a report, decision or communication?”

GrowCFO × AccountsIQ

QUESTION 1

# How do you actually secure AI tools without blocking adoption?

**Practical answer:** create safe lanes, register risky use, and make red lines impossible to miss.

QUESTION 1

# Create safe lanes before you create roadblocks

## Safe lane

Generic drafting, checklists, public information summaries, spreadsheet help and training examples.

## Controlled lane

Internal reports, forecasts, aged debt, board commentary, contract summaries and customer/supplier analysis.

## Stop lane

Payroll, bank details, passwords, unreleased results, M&A, legal advice and confidential customer data.

DO TOMORROW

**Tomorrow morning action:** Pick one low-risk use case and explicitly approve it. Example: "You may use the approved AI assistant to draft generic credit-control email templates, provided no customer-specific data is entered."

QUESTION 1

# Exercise 2: sort today's AI use into three lanes

| Lane | Example finance use | Minimum control |
|---|---|---|
| **Encouraged** | Draft a generic policy checklist | Approved tool; human checks output |
| **Controlled** | Summarize a board pack draft | Approved tool only; register; CFO review |
| **Controlled** | Analyze anonymized aged debt trends | Approved tool; source checking; reviewer signs off |
| **Prohibited without approval** | Paste payroll data into public AI | Stop; escalate; approved secure workflow required |
| **Prohibited without approval** | Use AI to send customer collection emails automatically | Control review before any automation |

**TEAM TASK**

Take five live examples from your finance team and place them in one of these three lanes.

QUESTION 1

# The “before you paste” test

- ✓ Is this an approved company AI tool?
- ✓ Does the prompt include confidential, personal, commercial or regulated data?
- ✓ Do I know whether prompts, files or outputs may be stored, reviewed or used for training?
- ✓ Have I removed names, identifiers and details the AI does not need?
- ✓ Could the output affect reporting, customers, suppliers, employees, investors or lenders?
- ✓ Is there a competent human review step before the output is used?

**DO TOMORROW**

Print this checklist or add it to your AI usage guidance. Ask the team to use it before uploading finance data.

**5 MINUTES**

Micro-exercise: take one recent AI prompt and remove every data point that is not strictly needed.

QUESTION 1

# Build a one-page AI use case register

| Field | Example entry |
|---|---|
| **Use case** | Draft monthly board commentary |
| **Owner** | FP&A lead |
| **Tool** | Approved enterprise AI assistant |
| **Data** | Amber: management accounts and variance notes |
| **Automation level** | Drafting / analysis |
| **Review** | CFO review before publication |
| **Status** | Approved with controls |

**DO TOMORROW**

Create the register in a spreadsheet. Do not wait for the perfect governance system.

### Minimum fields

Use case · Owner · Tool · Data type · Decision impact · Automation level · Human review · Status

GrowCFO × AccountsIQ

QUESTION 2

# What happens to your data once it enters an AI model?

**Practical answer:** map the data journey and ask for contractual evidence, not reassurance.

QUESTION 2

# Map the AI data journey

1 → 2 → 3 → 4 → 5 → 6

### 1 Prompt / upload

What did the user enter?

### 2 Application

Which product captured it?

### 3 Model provider

Which model processed it?

### 4 Storage / logs

What is retained?

### 5 Training rules

Can it improve models?

### 6 Output / audit

Can it be traced later?

**DO TOMORROW**

**Tomorrow morning action:** Choose one AI tool your team uses and map this data journey. Mark every unknown with a question mark.

**COLLECT**

Evidence to collect: admin settings screenshots, vendor terms, DPA clauses, subprocessor list, retention settings and audit-log availability.

QUESTION 2

# Exercise 3: ask the vendor the question that matters

### Send this to one AI-enabled vendor

*"Can you confirm whether our prompts, uploaded files, transaction data or outputs are used to train any AI model? Please also point us to the clause in the contract, DPA or product terms that confirms this, and identify any model providers or subprocessors involved."*

### Record the answer in your register

- Training use: Yes / No / Unclear
- Retention: Known / Unknown
- Subprocessors: Listed / Not listed
- Audit logs: Available / Not available
- Contract evidence: Clause captured / Not captured

QUESTION 2

# Ten vendor questions every CFO should ask

1. What AI features are live, optional, beta or experimental?
2. Can administrators disable or restrict AI features?
3. Is our data used to train any model?
4. Which model providers and subprocessors are used?
5. Where is data processed and stored?
6. How long are prompts, uploads and outputs retained?
7. Does the AI respect existing permissions?
8. Are prompts, outputs and AI actions logged?
9. Can outputs be traced back to source data?
10. Where are these commitments documented contractually?

**DO TOMORROW**

Pick your most important vendor. Ask about training, retention and permissions.

**DO NOT ACCEPT**

"Enterprise-grade", "secure" or "trusted" without evidence in settings, logs or contract terms.

QUESTION 2

# Exercise 4: run a 30-minute AI settings audit

| Check | Question to answer |
| --- | --- |
| **Training setting** | Is customer data excluded from model training? |
| **Retention setting** | Can prompt/file/output retention be limited or configured? |
| **Connectors** | Which email, drive, finance or workflow systems are connected? |
| **Permissions** | Can AI expose data a user could not otherwise access? |
| **Logging** | Are prompts, outputs and AI-triggered actions visible to admins? |
| **Subprocessors** | Are model providers and subprocessors documented? |

**30 MIN**

Output: one page showing green = confirmed, amber = unclear, red = unacceptable.

**OWNER**

Ask CTO/IT to confirm; CFO owns finance process risk.

GrowCFO × AccountsIQ

QUESTION 3

# Where is the line between useful automation and dangerous exposure?

**Practical answer:** the risk rises when AI moves from drafting to recommending, triggering or executing.

QUESTION 3

# Use the Automation Exposure Ladder

| | | |
|---|---|---|
| **1. Drafting** | Board commentary, policy wording, meeting questions | *Human review* |
| **2. Analysis** | Variance drivers, spend categories, aged debt trends | *Source checking* |
| **3. Recommendation** | Collect, accrue, reforecast, investigate | *Documented human decision* |
| **4. Workflow trigger** | Routes invoice, creates task, flags account | *Permissions, logs, thresholds* |
| **5. Execution** | Posts, approves, pays, updates, sends | *Formal control review* |

DO TOMORROW

**Tomorrow morning action:** Place three finance workflows on the ladder.

RULE OF THUMB

Level 3 needs documented human judgment. Levels 4–5 need formal control review.

QUESTION 3

# Exercise 5: classify the data before classifying the tool

## Green

Public info, blank templates, generic process descriptions, anonymized examples. Usually safe in approved AI tools.

## Amber

Management accounts, budgets, forecasts, board packs, aged debt, supplier data. Controlled use only.

## Red

Payroll, bank details, passwords, unreleased results, M&A, legal advice, personal data. Do not upload without explicit approval.

DISCUSS

**Tomorrow morning action:** List the five finance datasets most likely to be used with AI and mark them green, amber or red. Where do people disagree? That disagreement identifies where guidance is needed.

QUESTION 3

# Decide where human review is non-negotiable

### Reporting impact

Board pack, forecast, investor update, lender report

### People impact

Customer, supplier, employee or applicant affected

### Action impact

AI recommends, triggers, posts, approves or sends

### Sensitive data

Confidential, personal, regulated or commercially sensitive

### Hard to reverse

Error would be material, unfair, embarrassing or unlawful

DO TOMORROW

**Tomorrow morning action:** Pick one AI-assisted output currently used in finance. Decide who is competent to challenge it before it is used.

QUESTION 3

# Exercise 6: validate one AI output before using it

- ✓ Can I trace the output to source data?
- ✓ Do the numbers reconcile?
- ✓ Are dates, currencies and periods correct?
- ✓ Has the AI invented facts, citations or assumptions?
- ✓ Has it changed definitions or KPI logic?
- ✓ Would this stand up to board or auditor challenge?s

**DO TOMORROW**

Take one AI-generated paragraph, variance explanation or analysis and mark every claim as verified, unverified or wrong.

**CORE RULE**

AI output is a draft until checked.

QUESTION 3

# Set the CFO's red lines

- No unapproved AI tools for confidential finance data.
- No payroll, banking, legal, M&A or unreleased results in general-purpose AI tools.
- No AI-generated external communication without human review.
- No AI-generated financial reporting output without reconciliation to source data.
- No AI workflow that bypasses existing permissions or segregation of duties.
- No automated posting, approval or payment action without formal control review.

DO TOMORROW

**Tomorrow morning action:** Choose three red lines and send them to the finance team as interim guidance.

IMPLEMENTATION

# Turn next-day actions into a 30-day rhythm

| Timing | Focus | Practical output |
|---|---|---|
| **Day 1** | Find actual use | Ask team; capture tools, tasks, data and output use. |
| **Day 2** | Create safe lanes | Approve one low-risk use and define three red lines. |
| **Week 1** | Build the register | Record live use cases and basic risk status. |
| **Week 2** | Interrogate vendors | Training, retention, subprocessors, permissions and logs. |
| **Week 3** | Train behavior | Prompt hygiene, data minimization, validation and escalation. |
| **Week 4** | Review risky workflows | Automation ladder, human review, finance controls and approvals. |

TOOLKIT

# Template: AI use case register

| Field | Prompt |
|---|---|
| **Use case** | What is AI being used for? |
| **Owner** | Who is accountable? |
| **Tool / vendor** | Which product and model provider? |
| **Data involved** | Green, amber or red? Personal/confidential data? |
| **Decision impact** | Does it affect reporting, customers, suppliers, employees or external communications? |
| **Automation level** | Drafting, analysis, recommendation, workflow trigger or execution? |
| **Controls** | Human review, audit trail, approval, monitoring. |
| **Status** | Approved, controlled, under review or prohibited. |

TOOLKIT

# Template: simple AI risk score

### Score 1–5: data and impact

- Data sensitivity
- Personal data
- Commercial confidentiality
- Financial reporting impact
- Customer / supplier / employee impact
- External visibility

### Score 1–5: automation and control

- Automation level
- Reliance on AI output
- Legal / contractual exposure
- Difficulty of reversing error
- Vendor maturity
- Auditability

**Low risk: encourage**

**Medium risk: control**

**High risk: formal approval**

TOOLKIT

# Template: AI incident response mini-flow

### 1 Stop

Pause the output or workflow.

### 2 Capture

Tool, user, prompt, input, output, date/time.

### 3 Assess

What data was exposed? Was it personal or confidential?

### 4 Escalate

Notify IT/security, DPO/legal, CFO/CTO or process owner.

### 5 Remediate

Delete, revoke, disable, notify, retrain or update controls.

### 6 Learn

Review root cause and update governance.

DO TOMORROW

**Tomorrow morning action:** Agree who receives an AI incident: IT/security, DPO/legal, CFO, CTO and the finance process owner.

AccountsIQ

# Finance Function Health Check

**How much is inefficient finance costing your business?**

**Long month-end closes**

**Manual reporting**

**Disconnected systems**

These hidden inefficiencies don't just waste time, they can slow decision-making and limit growth.

**Scan the QR code to answer 5 quick questions** and find out how your finance function compares and where the biggest opportunities for improvement are.

GrowCFO × AccountsIQ

# Turn this playbook into lasting capability

**The GrowCFO Academy** is launching soon with 18 practical, CPE-accredited AI courses built specifically for finance.

**INCLUDED AI COURSES**

- Introduction to AI in Finance
- Prompt Engineering for Finance
- AI for FP&A
- AI for Board Reporting
- AI Writing for Finance
- AI for Cash Flow Forecasting
- Financial Modeling with AI
- AI Productivity Tools
- AI Risk & Governance

**...and 9 more courses.**

**Sept 2026: AI for Board Reporting**

A live, CPE/CPD-accredited workshop with Guy Weaver on using AI to sharpen board pack structure, narrative and clarity.