GrowCFO

Tech Innovation Report

AI Native and Agentic Solutions

A CFO's guide to evaluating, deploying, and governing AI-native finance technology

Table of Contents

Executive Summary

AI is moving fast inside finance technology. The terminology, AI-native, agentic, copilot, is being applied to almost everything, and CFOs are being asked to make real investment and deployment decisions in the middle of significant hype.

This report cuts through that noise with a practical framework for finance leaders who need to evaluate, deploy, and govern AI-native solutions responsibly.

We start with definitions that matter. In this report, AI-native describes products where AI is core to how the product works, not a bolt-on feature. It understands finance objects, operates within governed workflows, and produces outputs that can be evidenced. Agentic describes software that can progress work, not just advise, gathering information, executing steps across systems, and escalating exceptions within defined approval guardrails. If a tool cannot meet both tests, it is AI-enabled at best, and experimentation at worst.

The report then maps where these tools genuinely add value today across six core finance workflows: Record to Report, Planning and Forecasting, Billing and Accounts Receivable, Purchase to Pay, Treasury and Cash, and Risk, Fraud and Compliance. For each, we identify the AI patterns that work, the agentic workflows that are ready to deploy, the data and integration prerequisites, and the KPIs worth tracking.

A recurring theme is the build vs buy vs augment decision. For most finance teams, the right starting point is augmentation, adding AI capability on top of existing systems to reduce manual effort in specific workflows, without redesigning the underlying architecture. Buying a specialist AI-native solution makes sense where volume and friction justify deeper workflow ownership. Building is rarely the right answer unless the need is genuinely differentiating.

On governance, the report is unambiguous: AI can draft, classify, match, prioritize, and progress tasks. It cannot own accountability for financial statements, payments, or policy compliance. That accountability remains with the CFO. Finance-ready AI must be designed with clear ownership, explicit approvals, and traceable evidence from day one, not retrofitted after deployment.

The report closes with a 90-day deployment blueprint, a practical upskilling framework for finance teams, a vendor landscape overview, and a CFO checklist for evaluating whether any AI solution is genuinely finance-ready.

The objective is not more AI. It is measurable improvement in finance outcomes, delivered with finance-grade controls and the confidence that comes from having designed governance into the operating model from the start.

What do we mean by AI-native

What we mean by AI-native and agentic

The terms “AI-native” and “agentic” are being applied to almost everything in finance technology. In this report we use them in a **specific, practical** way that matters to CFOs: **does the technology reduce cycle time and touch time while strengthening control, governance, and auditability?**

AI-native (our definition)

AI-native describes products where AI is **core to how the product works**, not a bolt-on feature. In an AI-native finance workflow, AI does more than generate text: it understands finance entities (customers, suppliers, invoices, journals, dimensions), operates within the process, and produces outputs that can be evidenced.

An AI-native finance product isn’t simply a chatbot sitting on top of your ERP, nor is it a bundle of pre-written prompts. Instead, it is designed so that AI is embedded into the workflow: it understands finance objects (customers, suppliers, invoices, journals, dimensions), it can take action inside governed processes, and it improves the speed and quality of decision-making without weakening control.

Agentic (our definition)

Agentic describes software that can **progress work**, not just advise. A copilot helps a person complete tasks faster. An agent helps the organisation move work forward by gathering information, proposing a plan, executing steps across systems, and escalating exceptions—**within guardrails and approvals**.

Agentic is most valuable in finance where work is slowed by coordination: chasing evidence, triaging exceptions, routing approvals, reconciling mismatches, and progressing tasks across systems (ERP, billing, CRM, banking, procurement).

AI-native vs AI enabled

AI-native vs "AI-enabled" (the distinction that matters)

Most finance teams will encounter three broad patterns:

- **AI-enabled**: a traditional product that has added AI features (for example: summarisation, search, or suggested text). These can be valuable, but they rarely change the operating model.
- **AI-assisted**: AI is used to recommend actions (for example: "code this invoice to…", "these items look anomalous…", "this variance is driven by…"). A human still executes.
- **AI-native**: the product is built around AI-driven understanding and execution inside workflows (for example: classifying transactions, routing exceptions, drafting and evidencing narratives, assembling audit packs, and progressing tasks across systems with approvals).

A helpful test is: **would the product still deliver most of its value if you removed the AI?**

- If *yes*, it is likely "AI-enabled" (AI features added to a traditional product).
- If *no*, it is likely **AI-native** (AI is essential to its core operation).

AI-native systems typically add value by doing four things well, and that older systems struggle to do:

1. **Understand context:** policies, approval rules, chart of accounts, dimensions, master data, and materiality
2. **Work inside controls:** routing, approvals, sequencing, and ownership are built in
3. **Handle exceptions:** identifying anomalies and edge cases early and sending them to the right owner with context
4. **Create evidence:** clear traceability of what data was used, what changed, who approved, and what outcome resulted

AI-native and agentic systems must be finance ready

To be “finance-ready,” an agentic solution must be able to:

Sense: pull the right context from connected systems
Propose: recommend next steps with rationale and confidence
Act safely: execute only within defined permissions, thresholds, and approval workflows
Evidence: leave an auditable trail of actions and decisions

The CFO “agent stack” (simple view)

We view AI-native and agentic tools as operating across five layers:

Data: ERP, CRM, billing, procurement, banking, payroll
Context: policies, COA/dimensions, master data, thresholds

Workflow: tasks, approvals, controls, SLAs, ownership
AI: classification, matching, prediction, anomaly detection, narrative
Evidence: immutable logs, versions, approvals, audit outputs

The more fully a vendor covers **workflow + evidence**, the more “finance-ready” the solution is.

What we are not calling AI-native

We are not using these terms to describe:

- a chat interface over documents with no integration into finance systems
- prompt libraries that generate text without traceability
- tools that require manual export/import to take action
- autonomous posting or payment execution without approvals and auditability

In the sections that follow, we apply this lens consistently: **measurable outcomes**, **integration realism**, **control design**, and **quality of the evidence trail**.

Where can AI-native and Agentic workflows help?

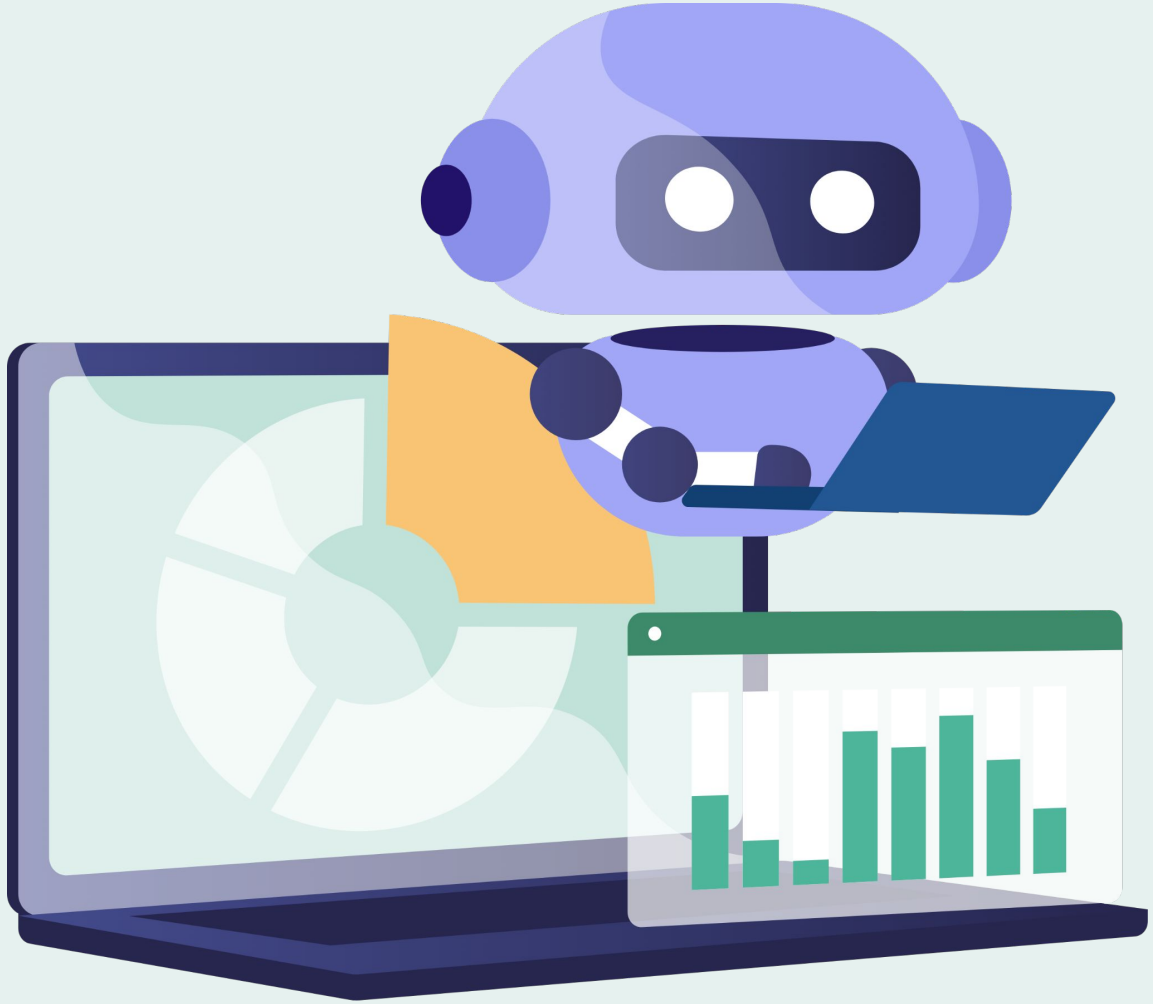

Record to Report

Where CFO teams feel the pain
- Close is slow because evidence is scattered, reconciliations are manual, and exceptions appear late
- Variance explanations are time-consuming and inconsistent
- Audit prep is reactive and repeats every period

AI patterns that work
- Reconciliation matching + exception detection (bank/GL/subledgers)
- Variance analysis + narrative drafting (with drill-through links)
- Anomaly detection (journals, unusual postings, duplicate entries)

“Agentic” workflows (with guardrails)
- Close orchestration agent: assigns tasks, chases evidence, escalates blockers
- Audit pack builder: compiles reconciliations/supporting docs and creates a traceable index
- Journal assistant: proposes journals + supporting rationale (human approval required)

Key integrations / data
- ERP/GL, subledgers, bank feeds, intercompany, consolidation tool, document store
- Must have: consistent chart/dimensions, period calendar, approval rules, evidence retention

KPIs to track
- Close days, late tasks, recon break rate, number of manual journals, audit queries

Planning and Forecasting

KPIs

- Forecast cycle time, forecast accuracy (by driver), scenario turnaround, stakeholder satisfaction

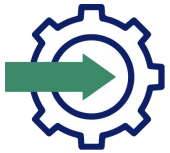

Key integrations / data

- ERP actuals, CRM pipeline, HR/headcount, billing/subscriptions, usage metrics, operational KPIs
- Must have: a single KPI definition set, agreed drivers, planning calendar, version control

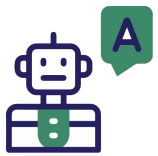

“Agentic” workflow

- Forecast refresh agent: pulls latest actuals, reconciles assumptions, proposes updated outlook
- KPI explanation agent: highlights root-cause contributors and risks/opportunities
- Executive pack assistant: drafts commentary + charts (human review)

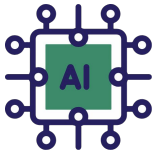

AI patterns that work

- Driver-based forecasting support (suggest drivers; flag drift/step-changes)
- Scenario generation + sensitivity analysis
- Narrative reporting: “what changed / why / what to do next” with evidence links

Where CFO teams feel the pain

- Forecasts are slow to refresh and rely on manual updates
- Driver definitions are unclear, scenario modelling is hard, and narrative takes forever
- Stakeholders want faster answers, not more spreadsheets

Billing & Accounts Receivable

Where CFO teams feel the pain
- Billing errors, revenue leakage, disputes, slow cash application
- Collections effort isn't prioritised; teams chase the wrong accounts
- Data split across ERP + CRM + billing + customer comms

AI patterns that work
- Cash application matching (remittance to invoices) + exception routing
- Collections prioritisation (propensity-to-pay, dispute likelihood, next best action)
- Revenue leakage detection (contract vs invoice vs cash vs usage)

"Agentic" workflow
- Collections assistant: drafts emails, schedules follow-ups, recommends settlement options (approval gates)
- Dispute triage agent: classifies disputes, routes to owner, pulls supporting documents
- Billing QA agent: flags anomalies pre-send (pricing, terms, usage mismatch)

Key integrations / data
- ERP AR, billing platform, CRM, ticketing/helpdesk, email/communications, payments gateways
- Must have: customer master hygiene, dispute reason codes, clear dunning policies

KPIs
- DSO, overdue %, cash application touch-time, dispute cycle time, write-offs, leakage %

Purchase to Pay

Where CFO teams feel the pain

- Invoice processing still has human touch points (coding, approvals, exceptions)
- “Where is my invoice?” supplier queries eat time
- Policy compliance is inconsistent; approvals are bottlenecks

AI patterns that work

- Invoice capture + coding suggestions with confidence scoring
- Policy checks (spend limits, vendor rules, category rules)
- Duplicate invoice detection and three-way-match exception explanation

“Agentic” workflow

- AP intake agent: reads inbox, extracts invoice, proposes coding/approver, routes into workflow
- Supplier query agent: responds using status lookups + approved templates (no free-text hallucination)
- Exception resolver: groups exceptions, suggests actions, escalates only true outliers

Key integrations / data

- ERP AP, OCR/intake, procurement/PO system, expense tool, vendor portal, document store
- Must have: vendor master controls, approval matrix, PO discipline (where relevant)

KPIs

- Cost per invoice, cycle time to approve/pay, exception rate, duplicate rate, supplier query volume

Treasury & Cash (Liquidity, Forecasting, Banking & Risk)

Where CFO teams feel the pain

- Cash visibility is fragmented; forecasts don't reconcile to reality
- Manual bank reconciliations and reporting
- Covenant monitoring and funding decisions are slower than needed

AI patterns that work

- Cash positioning and short-term forecasting support
- Automated reconciliation + anomaly detection
- Alerting on covenant triggers / unusual movements

"Agentic" workflow

- Daily cash brief agent: pulls balances, expected inflows/outflows, highlights risk
- Exception monitor: flags unusual payments/receipts, reconciles to underlying transactions
- Funding readiness assistant: assembles lender pack inputs (human review)

Key integrations / data

- Banks/TMS, ERP, AR/AP, payroll, capex schedule, debt schedules
- Must have: bank account structure, reliable payment calendars, approval rules for payments

KPIs

- Cash forecast accuracy, time to produce cash position, number of reconciliation breaks, risk alerts resolved

Risk, Fraud and Compliance

Where CFO teams feel the pain
- Controls are documented but not continuously monitored
- Fraud/duplicate payments are detected late
- Segregation of duties and access risk isn't visible day-to-day

AI patterns that work
- Continuous controls monitoring (CCM) on transactions and access
- Duplicate and anomaly detection for payments/journals/expenses
- Policy compliance monitoring and evidence logging

"Agentic" workflow
- Control monitor: runs tests, creates exceptions, routes to owners, tracks resolution
- Evidence agent: collects proof of control operation for audit trails
- Access risk watcher: flags role changes and SoD conflicts for review

Key integrations / data
- ERP audit logs, IAM/SSO, approvals workflow, payments, expense data
- Must have: clear control definitions, owners, thresholds, and remediation workflows

KPIs
- Exceptions per control, time to resolve, audit findings, duplicate payment rate, SoD breaches

Top 10 CFO use cases worth piloting (vendor-agnostic)

Designed for 30-90 day pilots: improve cycle time, touch time, and control quality with human-in-the-loop guardrails.

1 Close orchestration + evidence chasing

Outcome: Faster, more predictable close

KPIs: Close days; % tasks on time; hours chasing

Control: Sign-off on critical tasks (recs, journals, etc,)

Data: Close checklist/owners; ERP period status; workpaper/evidence links

4-8 weeks

2 Reconciliations + matching assistance

Outcome: Less effort, fewer recon breaks

KPIs: Time per recon; #done/day unreconciled items;

Control: Reviewer approval of matches + final reconciliation

Data: Bank feeds; GL/sub-ledger; tolerance rules

4-8 weeks

3 Variance analysis + commentary drafts

Outcome: Faster pack with better narratives

KPIs: Pack cycle time; narrative completeness; rework cycles

Control: Finance review/edit before publication

Data: Actuals + budget/forecast; dimensions; key KPIs

2-6 weeks

4 AP coding suggestions + policy checks

Outcome: Lower invoice touch time; fewer errors

KPIs: Touch time/invoice; % straight-through; re-coding level

Control: Approvals for exceptions /high value; override review

Data: Invoice data; supplier master; COA/other dimensions; approval matrix; policy rules

4-8 weeks

5 AP exception triage (PO/GRN/Invoice)

Outcome: Faster resolution; reduced backlog

KPIs: Exception cycle time; invoice aging; late fees

Control: Approval for value/timing decisions

Data: Invoice queue status; PO/receipting; exception reason codes; comms channel

4-10 weeks

6 Supplier & internal finance inbox agent

Outcome: Fewer inbound queries; faster...

KPIs: Deflection rate; response time; CSAT

Control: Approved templates; review-before-send

Data: AP/ERP status; payment calendar; remittance info; template library

2-6 weeks

7 Cash application assistance

Outcome: Faster cash posting; less unapplied cash

KPIs: Unapplied cash; time to apply; dispute cycle time

Control: Review exceptions, deductions and...

Data: Bank txns; remittances; AR ledger; tolerance rules

4-8 weeks

8 Collection prioritisation + next best action

Outcome: Improved collections efficiency and cash

KPIs: Overdues; DSO; collector productivity

Control: Approval for credit holds, settlements, etc

Data: AR aging; customer History & master data; dispute status; CRM notes; terms

4-10 weeks

9 Revenue leakage + billing exception detection

Outcome: Reduced leakage and billing disputes

KPIs: Leakage recovered; credit note rate; dispute volume

Control: Approval for billing changes, credits, rev...

Data: Contracts/price lists; billing/invoices; order/CRM data; rev rec policies

6-12 weeks

10 Forecast refresh agent (drivers + assumptions)

Outcome: Faster forecast cycles; better...

KPIs: Time to refresh; forecast accuracy; scenarios

Control: Sign-off on forecast version + assumptions

Data: Actuals; model; driver KPIs; pipeline/headcount etc.

4-10 weeks

What does this mean for your tech stack?

Build vs buy vs augment the stack

AI-native finance solutions are evolving quickly, and CFO teams are increasingly faced with a practical choice: **build** capabilities internally, **buy** a specialist solution, or **augment** what they already have (ERP, FP&A, close, AP/AR) with new AI layers.

In this report we frame the decision around three finance realities:

TIME-TO-VALUE
(HOW QUICKLY YOU CAN DELIVER MEASURABLE OUTCOMES)

CONTROL-READINESS
(HOW WELL GOVERNANCE, APPROVALS, AND AUDITABILITY ARE DESIGNED IN)

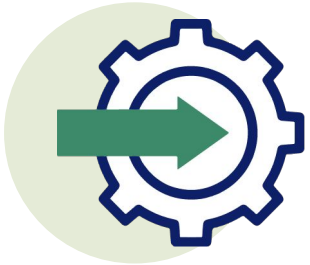

INTEGRATION EFFORT
(HOW MUCH WORK IS REQUIRED TO CONNECT RELIABLY TO CORE FINANCE SYSTEMS AND DATA)

Augment the existing tech stack

What it means

You add AI capability on top of current systems—typically as an AI layer, add-on module, or tightly integrated tool that connects to your ERP/FP&A/close/AP/AR processes.

Where it works best

Augment is usually the right starting point for pilots where the goal is to reduce manual effort without redesigning the underlying finance architecture, for example:

- close task orchestration and evidence chasing
- reconciliations and matching assistance
- variance analysis and first-draft commentary
- invoice coding suggestions and policy checks
- finance inbox triage and response drafting
- cash application and collections prioritisation support

Why CFO teams choose it

- lower disruption and faster deployment than replatforming
- easier to maintain a human-in-the-loop approach (recommend/draft, with approvals retained)
- can be deployed as “thin slices” that prove value quickly

Key risks to manage

- weak integration turns “augmentation” into another manual tool
- inconsistent master data (COA, dimensions, supplier/customer data) can undermine results
- auditability and evidence trails can be thin if AI outputs aren’t traceable to source data
- unclear permissioning or lack of segregation-of-duties support becomes a blocker as tools become more agentic

Typical outcome

Rapid productivity and cycle-time improvements in specific workflows—provided controls and evidence are designed into the operating model.

Buy a Specialist AI-native solution

What it means

You adopt a dedicated product that owns a workflow end-to-end (or a major portion of it), with AI embedded in the workflow and designed around exception handling and operational execution.

Where it works best

Buying a specialist solution makes sense where you have high volume, high friction, and clear ROI potential, such as:

- AP automation and exception management at scale
- AR cash application, collections workflow, and dispute management
- close acceleration (orchestration + reconciliations + audit pack readiness)
- forecasting refresh cycles that need better driver/assumption governance

Why CFO teams choose it

- deeper workflow ownership than many “overlay” tools
- often stronger exception handling and process discipline
- clearer operational ROI in high-volume processes

Key risks to manage

- integration complexity (ERP variants, entities, custom fields, approvals)
- risk of duplicated workflows unless roles and handoffs are redesigned
- “black box” automation if the system can’t show traceability and evidence
- commercial misalignment if pricing doesn’t map to your value drivers (users vs transactions vs entities)

Typical outcome

Meaningful step-change improvements in a specific process area, with stronger standardisation—if integration and controls are handled well.

Build (Selective, for differentiated needs)

What it means

You build internal AI capabilities using your data platform, integration tooling, and AI services—typically creating bespoke assistants, automations, or decision layers tailored to your business.

Where it works best

Building is most appropriate when the need is genuinely differentiating or too specific for off-the-shelf tools, for example:

- proprietary forecasting approaches or driver frameworks
- bespoke management reporting narratives grounded in internal KPIs and definitions
- industry-specific billing/revenue workflows
- finance policy assistants grounded in internal documentation and controls

Why CFO teams choose it

- control over data, logic, and governance design
- avoids vendor lock-in for unique processes
- can align tightly to internal definitions and operating model

Key risks to manage

- slower time-to-value and higher delivery risk
- ongoing ownership of governance (permissions, approvals, evidence retention, monitoring)
- maintaining accuracy as data definitions, policies, and processes change
- integration debt: reliability becomes your responsibility

Typical outcome

High fit for unique needs, but only sustainable when there is clear ownership, governance, and capacity to maintain it.

A practical way to decide

For most finance teams, the best route is sequential rather than ideological:
Start by augmenting to prove value quickly in a small number of high-impact workflows.
Buy specialist solutions where volume and friction justify deeper workflow ownership.
Build selectively only where differentiation or business-specific complexity makes it unavoidable.

The objective is not “more AI.” It is **measurable improvement** in finance outcomes—delivered with **finance-grade controls and auditability**.

Controls, Risk and Governance

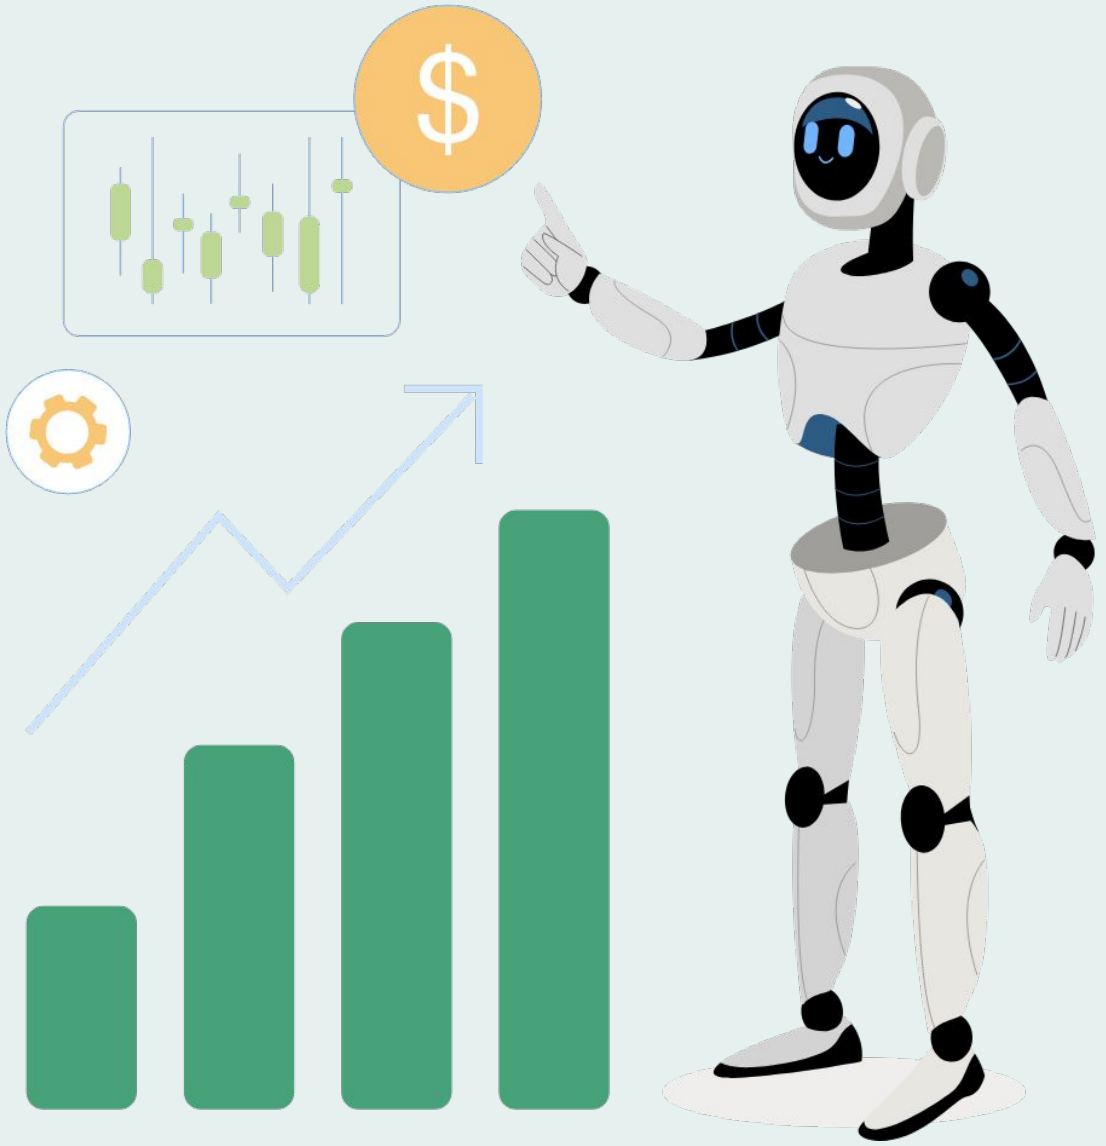

Automate work, not accountability

AI-native and agentic tools can compress cycle times in close, AP, collections and forecasting. But finance cannot adopt them on "productivity claims" alone. The CFO lens is always the same: **can we improve outcomes while strengthening control, auditability and accountability?**

This section sets out the minimum governance required to deploy AI safely in finance—without slowing innovation to a halt.

The key principle:
AI can draft, classify, match, prioritise and even progress tasks. It cannot own accountability for financial statements, statutory reporting, payments, or policy compliance. That accountability remains with the CFO and delegated control owners. For that reason, "finance-ready AI" must be designed with:

- **clear ownership** (who is responsible for each workflow and control)
- **explicit approvals** (who authorises decisions and postings)
- **traceable evidence** (what happened, when, why, and by whom)

Key takeaway:
In finance, AI adoption succeeds when governance is designed in from day one. The objective is not to slow innovation, but to ensure the gains in speed are matched by gains in control and confidence.

Controls, risk & governance — Finance-safe AI blueprint

A) Finance-safe AI baseline (minimum controls)

Identity & access
- SSO/MFA + RBAC
- Least privilege scope
- Joiner–mover–leaver discipline

Approvals & SoD
- Explicit approval points
- Prepare ≠ approve ≠ execute
- Materiality thresholds

Monitoring & performance
- Usage/override/error dashboards
- Alerts for spikes/breaches
- Periodic review for drift

Evidence & auditability
- Immutable activity log
- Versioning for rules/workflows
- Traceability to transactions

Exceptions & reversibility
- Low-confidence handling
- Routing to owners + SLAs
- Rollback/correction logged

B) Human-in-the-loop permission model (recommended for pilots)

Level	What AI can do	Risk	Typical use
L0	Observe (read-only insight, summaries, explanations)	Low	Close status, variance explanations
L1	Suggest (recommendations; humans execute)	Low	Coding suggestions, matching proposals
L2	Draft & prepare (artefacts pending approval)	Med	Draft emails, journals pending, audit packs
L3	Execute with controls (limited, reversible actions)	High	Task routing, status updates, low-risk automations

Default recommendation: run pilots at L1–L2; use L3 only where actions are low-risk, reversible, monitored and approvals are

C) Three red lines (non-negotiable)

- No autonomous posting or payments without approvals + evidence trail
- No opaque outputs that cannot be traced to transactions/data sources
- No broad access by default (permissions must be scoped and audited)

D) Pilot-to-scale gate (go/no-go)

Outcomes
- KPI improvement proven vs baseline
- Stable across at least one cycle

Controls
- Approvals + SoD operate in practice
- Exceptions route correctly with SLAs
- Evidence logs exportable + reviewable

Risk
- RBAC + least privilege validated
- Retention/residency rules agreed
- Monitoring live (overrides, errors, anomalies)

Change control
- Named owners for changes
- Versioning + sign-off for material changes
- Rollback plan tested

E) What auditors will ask (be ready)

Authorisation:
Who approved, and where is it evidenced?

Traceability:
Can you reconstruct what happened and why?

Completeness and accuracy:
Resources were used and are they controlled?

Change control:
What changed (rules/workflows/config) and who signed it off?

The 90 Day Deployment Blueprint

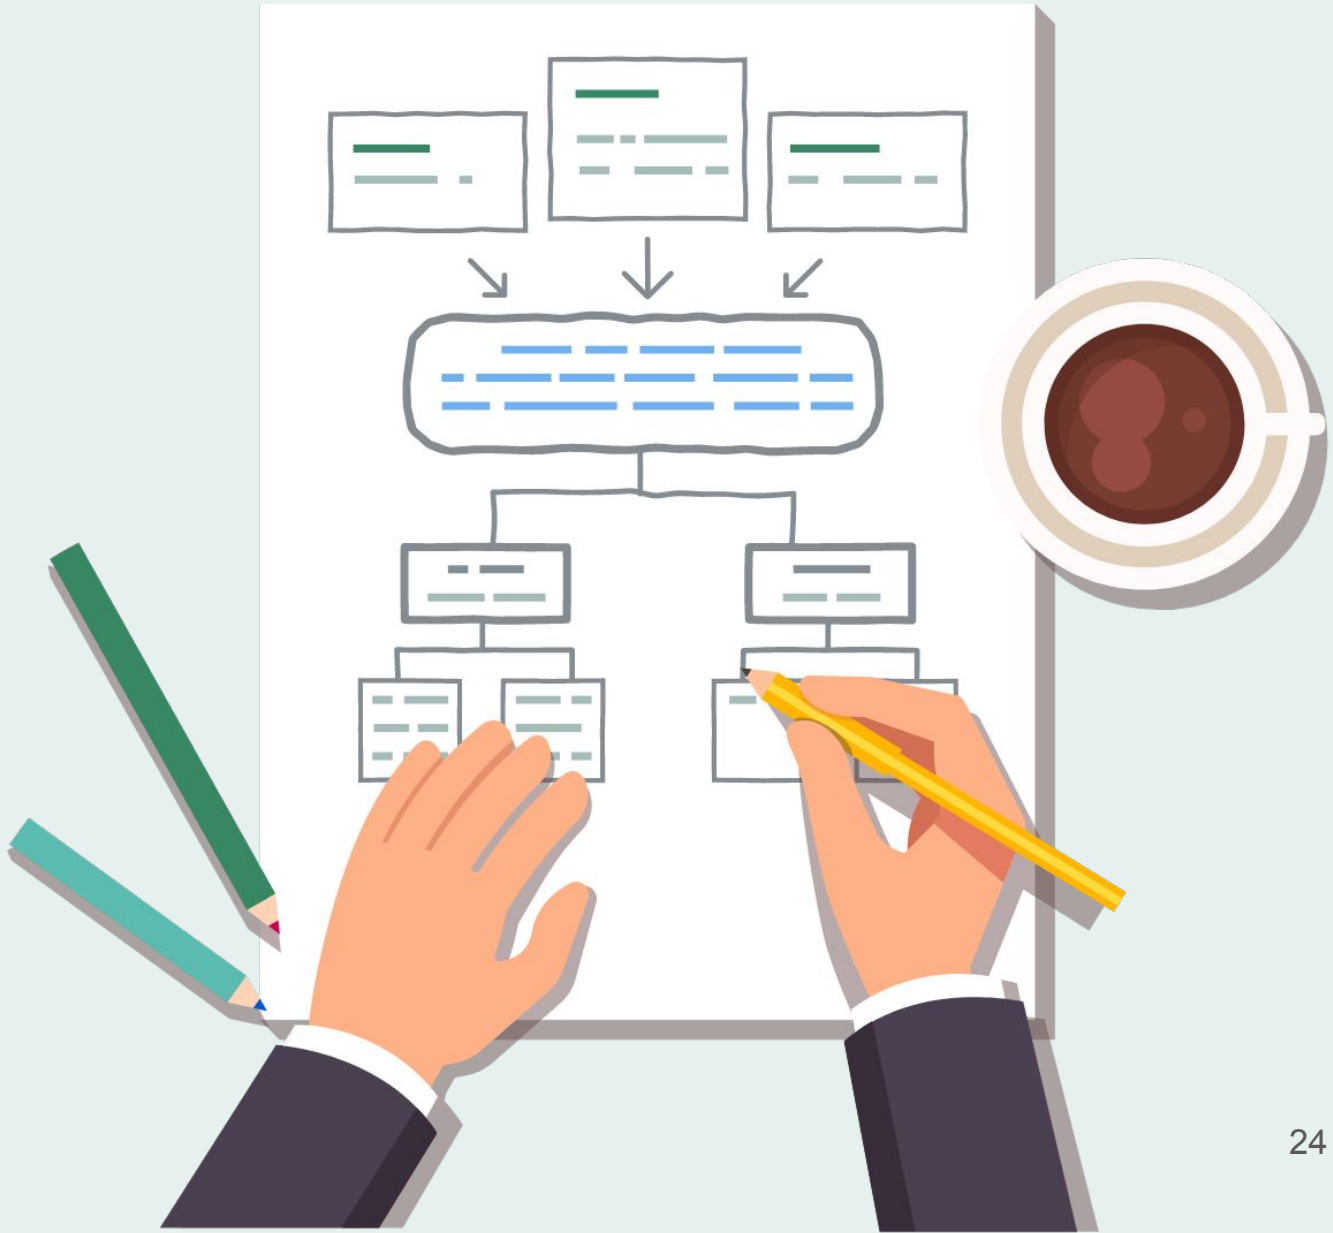

Implementation playbook (90-day, finance-safe pilot)

From demo to measurable impact in 90 days — without weakening controls.

Operating principle: Start thin-slice, keep humans in the loop, prove value weekly, then scale.

90-day deployment blueprint

Weeks 1–2 **Select, baseline, access**

- Pick one workflow; define scope boundaries and "no-go" actions.
- Baseline: cycle time, touch time, exceptions, rework (last month or 4–6 weeks).
- Confirm integrations/permissions (RBAC/SSO, least privilege). Outputs: pilot charter + baseline dashboard.

Weeks 3–6 **Configure, integrate, controls, UAT**

- Configure roles, approvals, thresholds, exception routes; align definitions/mappings.
- Integrate core systems; validate critical fields/statuses.
- Controls + UAT: approvals, evidence logging, low-confidence escalation; test real edge cases.

Weeks 7–10 **Limited rollout, measure, iterate**

- Roll out to one entity/team/segment; keep manual fallback defined.
- Weekly review: KPIs, overrides, exceptions, fixes; tune thresholds only when controls are strong.
- Output: week-10 scale/refine/stop decision.

Weeks 11–13 **Scale, document, governance**

- Scale one dimension at a time (volume → entities → users).
- Operationalise SOPs + exception playbooks; establish monitoring and changecontrol.

Suggested starting bundles

- Fast Close: close orchestration + reconciliations + variance commentary.
- AP Throughput: AP coding + exception triage + supplier/internal inbox agent.
- Cash Acceleration: cash application assistance + collections prioritization.

Measurement framework (pilot dashboard)

Cycle time:	Close days; invoice cycle time; dispute turnaround; forecast refresh time.
Touch time:	Touch time per item; handoffs per case; chasing time (evidence/approvals/queries).
Cash outcomes:	DSO and overdues; unapplied cash; collections productivity and resolution rate.
Quality & control:	Exception + override rate; rework rate; audit queries/adjustments; policy breaches resolved.
Cadence:	Weekly: KPI trend + top exceptions + actions. Monthly: value statement (hours saved + control gains).

Change and adoption plan

Training

- Role-based sessions (30 mins): preparers, approvers, controllers/FP&A.
- Guidance: what to trust vs verify; how to document overrides.

Human-in-the-loop rules

- Define allowed actions: Recommend → Draft → Prepare → Execute (with approvals).
- Set materiality thresholds; log every override with reason + owner.

Exception handling

- Pre-define top exception types, owners, SLAs, escalation paths.
- Low confidence/conflicting data → pause and escalate (don't guess).

Edge-case playbooks

- For each case: symptoms → evidence → options → approval → logging steps.

Upskilling your team for AI

Upskilling the finance team to get the most out of AI-native systems

AI-native finance systems don't just "automate the old process." They change how work is done: less time spent navigating screens, copying data, chasing approvals, and reconciling spreadsheets; more time spent defining the question, validating outputs, resolving exceptions, and improving controls. To realise the value, finance teams need new habits and a small set of practical skills that sit between finance expertise and product/AI literacy.

The shift in roles: from processing to orchestration

In an AI-native environment, the system can draft journals, code invoices, propose accruals, reconcile transactions, write variance commentary, and prepare first-pass forecasts. Finance's role shifts toward:

- **Setting intent** (what good looks like: policies, thresholds, business rules, definitions)
- **Supervising outcomes** (reviewing, sampling, approving, and escalating exceptions)
- **Improving performance** (tuning workflows, reducing error rates, tightening controls)
- **Building trust** (auditability, evidence capture, documentation, governance)

This is not a "data science" requirement. It is a capability uplift in **judgement + control + process ownership**, enabled by better tools.

What skills matter most

Upskilling works best when it's framed as *job-to-be-done capability*, not "AI training". The aim is confident, consistent use of AI-native features while maintaining strong financial control.

1) AI fluency for finance (everyone)

Understanding what AI-native systems can and can't do (probabilistic outputs, failure modes)
Knowing when to trust, when to verify, and how to challenge results
Using copilots effectively (clear prompts, good context, structured requests)

2) Data & process literacy (most of the team)

Basic understanding of data quality drivers (master data, mappings, chart of accounts discipline)
Knowing which upstream processes affect AI outcomes (POs, approvals, coding, project tags)
Being able to diagnose "why the system is wrong" without defaulting to spreadsheets

3) Controls, risk, and audit readiness (core finance + controllers)

Designing "human-in-the-loop" review rules (thresholds, sampling, exception types)
Evidence capture by default (why a decision was made, what changed, who approved)
Segregation of duties in an automated world (who configures vs who approves)

4) Analytics & storytelling (FP&A and business partnering)
Moving from report production to insight production (drivers, sensitivity, scenarios)
Interpreting AI-generated commentary and improving it with business context
Communicating uncertainty clearly (confidence ranges, assumptions, limitations)

5) Product ownership & continuous improvement (a small cohort)
Acting as "AI-native process owners" for close, P2P, O2C, FP&A
Maintaining workflow configuration, exception queues, and user adoption
Coordinating with IT/security on access, change control, and governance

A simple capability model (so it sticks)

Rather than training everyone on everything, define three capability tiers:
All users (baseline): confident use of copilots + safe handling of outputs
Power users (role-based): exception management, review rules, workflow tuning
Process owners (few): configuration governance, metrics, continuous improvement, controls

This prevents "AI champions" becoming a bottleneck, while ensuring ownership doesn't drift to IT or vendors.

GrowCFO AI Courses

These two upcoming live AI workshops are designed to be highly applied, focusing on how people actually use AI day to day rather than theory or hype.

9 March (4pm GMT)

Process Optimization with AI

How to use AI to improve workflows, decision-making and productivity, with real examples and practical takeaways you can apply immediately.

23 March (4pm GMT)

Getting Your Team to Adopt AI

How to move from experimentation to real adoption. This session covers capability, behaviors, and what actually helps AI stick in an organization.

Both sessions are led by **Oli Deacon** (ICAEW FCA, PCC-certified coach); former Microsoft Finance Director and one of our lead AI trainers.

Access:
These workshops are included with **GrowCFO Premium.**

Click the button below for access details.

LEARN MORE

Key Vendors and Platforms

AI Native across the Tech Stack

AI-native Accounting & ERP

New ledgers and ERP layers built for automation, faster close, and AI-led

Digits · Puzzle · DualEntry · Rillet

Examples (not exhaustive)

Close, Recs & Consolidation

Agentic controllership tools that prepare recs, variance/flux,

Numeric · Stacks · Ledge · Truewind · Nominal

Examples (not exhaustive)

AP, Spend & Procurement

Automate intake-to-pay: invoice capture, matching, coding, approvals

Vic.ai · Ramp · Omnea · Levelpath · Coupa · Payflows

Examples (not exhaustive)

FP&A & Business Planning

AI-native planning that speeds forecasting cycles, scenario

Drivetrain · Abacum · Aleph · Clockwork · Compass

Examples (not exhaustive)

Treasury & Cash Operations

Cash visibility, bank connectivity, payments orchestration and

Atlar · Round Treasury · Payflows Treasury

Examples (not exhaustive)

AR, SaaS Billing & Revenue Mgmt

Order-to-cash automation for subscription + usage billing,

Vic.ai · Ramp

Examples (not exhaustive)

Tax, Compliance & Technical Accounting

Automation for tax/compliance research, determinations, documentation and audit readiness in specialist accounting areas.

Avalara · Thomson Reuters CoCounsel Tax · Trullion · Sage Copilot

Key considerations for CFOs evaluating AI-native applications

Maturity of AI features: Some vendors talk "AI-native" but may have fewer live deployments in the mid-market; check case studies, third-party reviews.

Data readiness & process maturity: AI works best when data is clean, processes are well-defined, and you have fewer legacy silos. Upgrading to an AI-native system requires a baseline of process discipline.

Change management & skills: Moving to an AI-infused system changes the role of the finance team (more insight, less manual). Budget for training, role re-definition.

Implementation speed vs risk: Some of these platforms promise rapid deployment (DualEntry claims "migrate in 24 hours" for legacy data) – but evaluate integration scope, customization needs, data migration risk.

Ecosystem & support: A vendor's partner network, add-on ecosystem, localization (tax/currency) still matters. Emerging AI-native vendors may have smaller ecosystems than established mid-market players.

Vendor roadmap and stability: Since many of these are newer, check their financial backing, roadmap, customer support model, and how they compare with more established systems in terms of long-term viability.

zenskar | AI-Native Order-to-Cash

4.9 ★ 5.0 ★

FALL 2025 Leader | FALL 2025 Momentum Leader | FALL 2025 High Performer | FALL 2025 Easiest Admin | FALL 2025 Easiest To Use

ZERO-TOUCH AUTOMATION

Send your contracts and usage data, Zenskar's AI automates everything downstream - invoices, collections, journal entries and analytics.

Whether you bill by subscriptions, usage, or hybrid models, Zenskar handles the complexity so your team doesn't have to.

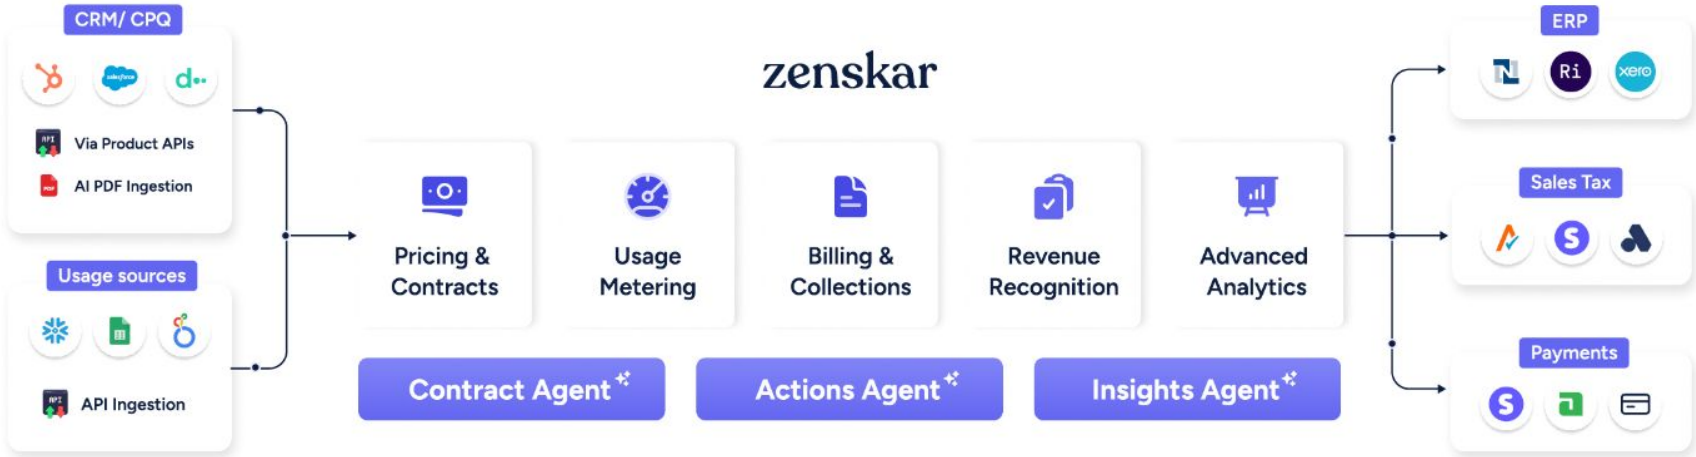

BUILT FOR IMPACT

- **100+ hours** saved for finance
- **2-3% uplift** in revenue
- **Zero day** close process
- **95%** dev time saved
- **50%** faster collections
- **6X** faster migration

TRUSTED BY INDUSTRY LEADERS

Matt Barnard
VP Finance, Vertice
4.5/5

70% faster month-end closing

"Zenskar automates revenue recognition accurately for our value-based billing, with AI reducing the manual hours spent by 70%."

Noy Kalansky
Controller, Pontera
5/5

Saved 200+ hours of grunt work

"We're saving 200+ hours/quarter on invoicing and receivables management by completely automating our recurring billing."

100ms · Pontera · Gremlin · YEMBO · Sardine · INDIGOV · Adelaide · Dfy · SquadStack · LOOP

Backed by Bessemer Venture Partners

Book a Demo >

Drivetrain: AI-native Business Planning Platform

Drivetrain brings the speed of AI and the clarity of a modern planning platform; helping finance teams forecast faster and decide smarter.

Your Data. Unified. Instantly: With 800+ integrations, Drivetrain provides real-time, analysis-ready data; no manual work needed.

AI Where It Matters Most: Build models, and generate reports in natural language. AI Agents handle everything from complex data transforms to multidimensional models.

Fast Time to Value: Rapid deployment, quick ramp-up, and an intuitive self-serve interface deliver the lowest total cost of ownership in the category.

- Cash Flow Forecasting
- Revenue Planning
- Sales Capacity Planning
- Financial Consolidations
- Headcount Planning
- 3-Statement Reporting
- Investor Reporting
- Expense Planning

98%
avg reduction in data consolidation time

5+
days saved every month on reporting

6-8
weeks average go-live time (G2)

5.7
months average payback period (G2)

Trusted by finance teams across 17+ countries

OBSERVE.AI · POSTMAN · Apollo · teikametrics · dental intelligence · TruHearing · matillion · Graham Allen

See Drivetrain in Action

Selecting the Right Solution

How to find the right solution for you

- Be clear on your requirements and what is important to you.
- Examine all the alternatives. Read reviews, ask for recommendations from fellow finance leaders.
- Software comparisons are available in a number of places.
  - The reviews on G2 are particularly useful:
    - https://www.g2.com/categories/accounting
- Make a short list:
  - Ask vendors on your shortlist targeted questions that will determine whether they meet your needs.
  - Consider whether a demo from the vendor would be useful, and don't be afraid to ask the vendor to tailor the demo to show how their product specifically meets your key needs.

CFO checklist — is this “finance-ready” AI?

Use this checklist in vendor demos and selection.

1) Controls & permissions

- [ ] Role Based Access Control (RBAC) / SSO and least-privilege access
- [ ] Segregation of duties support (prepare vs approve vs post/pay)
- [ ] Clear human approval steps for postings, payments, and master data changes

2) Evidence & auditability

- [ ] Immutable activity log (inputs, outputs, actions taken)
- [ ] Version history for rules, workflows, and AI configuration
- [ ] Traceability from AI output back to underlying transactions/data sources

3) Workflow reality

- [ ] Integrated with core systems (ERP/billing/banking/CRM/procurement) via APIs/connectors
- [ ] Exception handling (what happens when confidence is low or data conflicts?)
- [ ] Ownership routing (tasks land with the right person, with SLAs)

4) Risk management

- [ ] Materiality thresholds (what can run without approval?)
- [ ] Monitoring dashboards (activity, overrides, error rates, drift)
- [ ] Reversibility (rollback/correction process + documentation)

Decision rule:

If a vendor cannot clearly demonstrate **controls + evidence**, treat it as **AI-enabled experimentation**, not a finance-ready deployment.

AI Native Solutions Tech Showcase

Live demos from AI-native solutions redefining finance teams with intelligent automation, real-time insights, and faster decision-making.

Wednesday | 25 March 2026
3:00 PM GMT

FEATURING:

zenskar

Drivetrain

DualEntry

REGISTER NOW