



Intelligent Finance Operations

From Manual Workflows to Supervised Execution

A GrowCFO white paper

Table of Contents

Executive summary	1-2
Ten findings for CFOs	3-5
1. The execution threshold: why 2026 is different	6-10
2. Where intelligence lives: embedded, suite and external agents	11-18
3. The intelligent transaction layer	19-22
4. Continuous close, control and assurance	23-25
5. The intelligent cash engine	26-28
6. The finance control plane	29-34
7. The new finance workforce	35-37
8. Proving value beyond headcount	38-40
9. Access and portability	41-44
10. The orchestration layer	45-48
Where this report sits in GrowCFO's 2026 coverage	49-50
Selecting the Right Solutions	51
Showcase demo day: turning this report into a vendor decision	52-63
Conclusions	64-65
Appendix: the 30-question CFO self-assessment	66-68
Source notes and bibliography	69-72

Executive Summary

Finance is crossing a threshold. Through 2025 and into 2026, generative assistance inside spreadsheets, chat windows and reporting packs has been joined by something more consequential: systems that can monitor a process, reason about what is happening, draft or prepare a response, and in bounded cases act on it. Autonomous finance has not arrived — but it has become operationally credible in specific, high-volume workflows: invoice capture and matching, cash application, reconciliation, variance narrative. Broader production maturity still lags vendor capability. Deloitte's finding that 63% of finance teams actively use AI but only 21% report clear, measurable ROI is the clearest single statistic of that gap [1].

The defining question for CFOs is no longer whether to use AI, but where intelligence should live — inside the transaction system, inside the corporate productivity suite, in an external cross-application agent, or in the orchestration layer that connects all three — and how each choice is governed. This report tests that framework against real 2026 evidence: production volumes, exception rates, cycle times and named accountable owners, not vendor roadmaps.



Executive Summary (cont'd)

If you are approving pilots but have not yet made a deliberate, function-wide decision about where AI should sit and who should govern it, you are not behind — most finance leaders are in the same position. To make this report useful in that position, not just informative, every one of the ten sections below closes with a short, practical checkpoint: three questions you should be able to answer about your own organization today. A full 30-question self-assessment checklist appears near the end for use as a standalone diagnostic.

The defining question is no longer whether finance will use AI. It is where intelligence should live, what it should be allowed to do, and how CFOs preserve control when work crosses systems, models and organizational boundaries.

10 Findings for CFOs



Ten findings for CFOs

The shift is from task automation to workflow ownership — the unit of transformation is an end-to-end outcome, not a single AI feature.

Embedded AI and external work agents are converging, not simply competing; the likely architecture is hybrid.

A Microsoft- or Google-only organization is not an AI laggard by definition — suite standardization trades frontier depth for governance and adoption speed.

Supervised autonomy is the credible 2026 position: further ahead in AP, matching and reconciliation; tighter review still warranted in tax, treasury and judgment-heavy accounting.

Control must evolve from user access to agent identity — owner, purpose, least privilege, approval thresholds, logs, evaluation and retirement.

ROI should be measured in cycle time, exception quality, working capital, accuracy and decision latency — not licenses or pilots [1].

Finance talent shifts from processor to designer, reviewer and supervisor.

Workflow debt is a recognized obstacle: agents layered over fragmented, human-paced processes produce limited gains without redesign first [23].

AI access is stratified — approved frontier agents, contracted suite assistants, and unsanctioned tools coexist, making access simultaneously an innovation, governance and workforce-equity issue [27][28].

The orchestration layer is becoming strategically important, connecting systems of record, models and human approvals — and risking shadow integration if poorly governed [29].

Ten findings at a glance

01

Workflow ownership, not task automation

02

Embedded AI and external agents converge, not compete

03

Suite-only is not an AI laggard by definition

04

Supervised autonomy is the credible 2026 position

05

Control shifts from user access to agent identity

06

ROI = cycle time and quality, not headcount

07

Talent shifts: processor to designer and supervisor

08

Workflow debt limits gains without redesign first

09

AI access is stratified - a governance issue too

10

The orchestration layer becomes strategically important

Visual index only — the full findings remain in the text.

01


The execution threshold: why 2026 is different



The execution threshold: why 2026 is different (cont'd)

Every prior wave of finance AI promised productivity; 2026 is the first in which credible production evidence exists for execution, not just assistance. Deloitte's 2026 Finance Trends survey of more than 1,300 global finance leaders found 63% actively using AI and 14% with fully integrated agents — but only 21% reporting clear, measurable ROI, with the strongest opportunities concentrated in sales and profitability management, working-capital optimization and expense management [1]. KPMG points the other way, reporting that 93% of US companies expect to deploy or scale AI in finance within eighteen months, and that half are planning multi-agent systems [6]. Forrester supplies the necessary counterweight: adoption intent is high, but meaningful production and scaled multi-agent systems remain rare, and the defining transition is still from chat and prototype toward enterprise execution.


This is a genuinely productive tension to hold rather than resolve. Capability has crossed the threshold into operational usefulness in bounded domains; enterprise readiness — data, controls, ownership — has not caught up everywhere. Treat every vendor claim and every internal pilot report with the same evidential discipline: ask for production volumes, touchless rates, approval rates, false positives, rollbacks and a named accountable owner, rather than accepting autonomous-finance predictions as fact.



The execution threshold: why 2026 is different (cont'd)


Bain's April 2026 survey adds the operating-model half of the story. While 83% of more than 100 surveyed CFOs plan to raise enterprise AI spending by more than 15% over two years, Bain estimates that only 15–25% have actually scaled AI across finance, and respondents cite speed — not the cost case that often justified the original investment — as the biggest realized win. Bain's prescription is blunt: pay down 'workflow debt' — fragmented handoffs, unclear ownership, duplicate controls, human-paced queues — before deploying agents on top of it [23]. An agent that accelerates a badly designed process improves a task metric while leaving the end-to-end outcome largely unchanged; use that distinction as a filter for every case study or vendor pitch you evaluate.

KPMG's parallel 2026 study of 1,013 senior finance leaders across 20 countries adds a further, less obvious finding: AI's strongest gains are appearing in judgment-heavy work — decision speed, decision quality and forecasting accuracy — rather than pure transaction processing. More than three-quarters of respondents now use AI in planning, reporting and commercial analysis, yet only 23% say it exceeds expectations. The standout statistic is governance-related: organizations able to produce AI audit evidence efficiently report three to six times the rate of significant improvement on selected measures compared with those that cannot [24]. Assurance readiness, in other words, is not simply a brake on adoption — it is a genuine performance lever, and it is one of this report's central, slightly counter-intuitive arguments.



The 2026 autonomy ladder: Which level are you on?

Report Section 1 — "agentic" is not a binary label; 2026 deployments sit at different rungs



1. The execution threshold — practical checkpoint


PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. Where in finance have we moved from AI-assisted drafting to real execution — and can we evidence production volumes, not pilots?
2. What share of our finance AI spend has a measured operating outcome attached to it, rather than none?
3. Which three finance processes are our strongest candidates for supervised autonomy in the next 12 months, and do we agree on them?


02

Where intelligence lives: embedded, suite and external agents



2. Where intelligence lives: embedded, suite and external agents

The comparison CFOs actually face is not 'core product AI versus ChatGPT or Claude' as a binary choice. There are three overlapping locations for intelligence — embedded product AI, the enterprise productivity suite, and external cross-application work agents — plus a fourth, enabling orchestration layer that connects all three. Use this four-part map as your organizing framework for every AI investment decision, rather than a single vendor-versus-vendor contrast.



Embedded AI: strongest where transaction truth and controls matter

AI embedded in ERP, close, AP, AR, billing or treasury products operates near authoritative data, configuration, workflow state and access controls, which makes it well suited to matching, coding, anomaly detection, reconciliation, cash application, approval routing and auditable suggestions. Its weaknesses matter equally: an embedded feature can:

- be confined to one application boundary,
- inherit poor master data,
- offer limited model choice,
- or simply become another silo.

Gartner forecasts that finance organizations using cloud ERP with embedded AI assistants will achieve a 30% faster financial close by 2028 — a genuine strategic proposition, but it is a forecast rather than a demonstrated 2026 benchmark, and should be treated as such [25].

External work agents: strongest where work crosses boundaries

Claude Cowork and ChatGPT Work represent the opposite proposition: a general agent that assembles context across files and applications, performs long-running analysis, creates spreadsheets and documents, and bridges the gaps between systems that embedded AI cannot see across. Anthropic documents finance use in forecasting, variance analysis, close and board-pack assembly [16]; OpenAI positions ChatGPT Work as a multi-hour agent producing finished work across apps and files [17]. Anthropic's own internal case is a useful, concrete illustration: one corporate-finance professional reports reclaiming 10–20 hours a week using Claude Cowork and Claude for Excel to reconcile board-deck numbers and narrative, draft variance commentary, diagnose financial models and retrieve cross-team context, with first-pass work also described in GL-to-subledger and bank reconciliations, acquisition screening, tax and treasury [26]. This is useful directional evidence, but it is a vendor describing its own product use, not independent customer validation — treat it accordingly. The flexibility that makes external agents powerful in the 'white space' between systems is the same property that raises the control burden: a parallel process outside the system of record, permissions used too broadly, or an artifact whose lineage is unclear. The right question is not whether an output looks right, but whether it is reproducible, reviewable and anchored to approved data.



Enterprise suites: the governed distribution layer

Many organizations will deliberately restrict staff to whatever AI is already covered by their Microsoft or Google enterprise contract — not for lack of ambition, but for reasons of security review, data residency, procurement leverage, information-protection integration, records management and the wish to avoid shadow AI. Microsoft's 2026 Finance Agent places ERP-grounded intelligence inside Excel, Outlook, Teams and Copilot chat, preserving familiar work patterns and central administration while ERP permissions continue to govern ERP data [19]. Google's Gemini Enterprise combines agent building, long-running work, shared context, a registry, identity and observability across Google Workspace, Microsoft 365 and partner systems [20][21]. Standardization can accelerate safe adoption and reduce integration sprawl, but finance may receive slower access to specialist capability or frontier models — a suite contract should be treated as a control boundary and distribution channel, not automatically as the complete finance AI strategy.




Where intelligence lives: side-by-side comparison

Dimension	Embedded product AI	Enterprise-suite AI	External work agent
Examples	SAP Joule, Oracle ERP agents, Workday, BlackLine, AP/AR/close vendors	Microsoft Finance Agent/Copilot; Gemini Enterprise/Workspace	Claude Cowork; ChatGPT Work and workspace agents
Native context	Deep transaction, configuration and process state	Email, meetings, documents, spreadsheets, identity graph	Broad files/apps/web; flexible context assembly
Best at	In-process controls and repeatable domain workflows	Everyday knowledge work inside the corporate suite	Cross-system research, reasoning, artifact creation and novel workflows
Primary advantage	Proximity to system of record and established permissions	Enterprise standardization, adoption and central administration	Frontier capability, speed of innovation and user-configurable agency
Primary risk	Vendor lock-in, uneven AI quality, process silos	Lowest-common-denominator capability; suite boundaries	Data leakage, shadow workflows, fragile integrations, unclear record of truth
Control question	Can the vendor prove lineage, approvals and rollback?	Are permissions and information-protection labels preserved end-to-end?	Exactly what can the agent read, infer, write and send — and under whose identity?

Where intelligence lives: the four-part map

Report Section 2 — three peer layers, one shared foundation, one connecting layer



2. Where intelligence lives — practical checkpoint

PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. For our highest-volume finance processes, have we deliberately chosen whether AI should sit in the ERP, the corporate suite or an external agent — or did that happen by default?
2. What data are we willing to let leave our system of record for an external agent to use, and who approved that boundary?
3. If our ERP or specialist vendor changed tomorrow, would we keep the AI capability we have built, or lose it?

03

The intelligent transaction layer




The intelligent transaction layer

Accounts payable is the clearest 2026 proof point for supervised autonomy, and it is the anchor case to test before extending outward to expenses, procurement and supplier communications. PwC estimates agents can already automate 25–40% of typical global-business-services tasks [4], while Forrester identifies capture, exception handling, matching, supplier communications, reporting and fraud/risk as the areas advancing furthest toward supervised autonomy [11]. PwC's own OpenAI-native finance build reports up to an 80% reduction in procure-to-pay cycle time in its example [4].

The more important implication is organizational rather than technical: teams increasingly manage exceptions, confidence thresholds and escalations rather than processing every transaction by hand. Test every vendor and internal claim in this area against straight-through processing rate, matching accuracy, exception cycle time, duplicate and fraud-detection precision, and supplier-query deflection — metrics that describe a redesigned operating model, not a faster version of the old one.



The supervised execution workflow



Referenced in Sections 3, 4 and 10 — the illustrative pattern behind AP, close and orchestrated workflows.

3. The intelligent transaction layer — practical checkpoint

PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. What is our actual straight-through processing rate in AP today, and what is specifically blocking it from rising?
2. Do we know our exception cycle time and duplicate/fraud-detection precision, or only anecdote?
3. Have we redesigned the underlying process, or simply automated the process we already had?

04

Continuous close, control and assurance



4. Continuous close, control and assurance

BCG distinguishes reactive copilots from systems that continuously sense signals, reason through causes and propose or execute a response; PwC describes a comparable shift toward a finance function that 'reasons continuously,' with agents executing and humans governing [8][22]. The practical effect is a move away from scheduled batches and end-of-period queues toward event-driven monitoring and intervention throughout the period: close becomes continuous readiness, with reconciliations, anomalies and supporting evidence monitored as they occur rather than assembled retrospectively; working capital becomes active, with agents prioritizing collection actions, supplier exceptions and cash risks as signals change; and reporting becomes investigation, with commentary generated from traceable drivers and humans validating the judgments that matter.

How auditors gain assurance over agent-produced work deserves particular attention — KPMG's finding that efficient AI audit evidence correlates with three to six times the improvement rate is the strongest available bridge between governance and performance [24]. Evidence to seek in every case study: close-days reduced, aged reconciling items, evidence completeness, control failures and audit effort.



4. Continuous close, control and assurance — practical checkpoint

PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. Could we produce audit-ready evidence for an AI-assisted reconciliation today, on demand?
2. Is close still a period-end event for us, or have we begun moving to continuous, evidence-as-you-go monitoring?
3. Who is accountable when an AI-generated variance explanation turns out to be wrong?

05

The intelligent cash engine



The intelligent cash engine

Revenue management belongs here, as part of a broader cash and revenue-assurance story rather than a stand-alone legacy RevOps theme. PwC's 2026 model of an 'agentic revenue engine' for lead-to-cash shows agents coordinated across system seams, with identity, memory, tool access, observability and auditability mattering more than the underlying model alone [22]. Connect billing accuracy and revenue leakage, collections prioritization and payment prediction, and short-term cash visibility into a single coherent narrative, using DSO, promise-to-pay performance, leakage prevented, cash-application rate and dispute-resolution time as the evidence rubric.



5. The intelligent cash engine — practical checkpoint


PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. Do we have one connected view of DSO, promise-to-pay performance and cash-application rate, or three separate reports?
2. Where is revenue leakage most likely occurring in our billing and contract-to-cash process — and have we actually tested for it?
3. Are we ready to let AI prioritize collections actions on our lowest-risk accounts without a human touch?

06

The finance control plane



The finance control plane

Accenture calls for an AgentOps control tower; Forrester argues every agent should be treated as a governed identity; Gartner centers oversight, system constraints and continuous evaluation; and Google, Microsoft and OpenAI are all now exposing administrative controls, agent registries, action restrictions or observability of their own [10][13][14][15][19][20]. The consistent message across every source is that the control plane — not the chat interface — is becoming the strategic layer. Organize governance around the lifecycle of an individual agent rather than a one-time model approval.



The finance control plane: a lifecycle framework

Control layer	Minimum design question
Purpose	A bounded outcome, named business owner and prohibited uses.
Identity	Unique agent/service identity; no shared human credentials; least privilege.
Context	Approved sources, freshness rules, classification handling and lineage.
Reasoning and tools	Allowed models, tools and deterministic calculations; prompt/skill versioning.
Action	Read/write boundary, financial thresholds, segregation of duties and approval gates.
Evaluation	Accuracy, completeness, exception quality, bias, security and adversarial tests.
Operations	Logging, observability, cost limits, incident response, rollback and kill switch.
Assurance	Evidence retained for management, internal audit, external audit and regulators.
Lifecycle	Periodic reapproval, change control, owner transfer and retirement.



The finance control plane — the diagnostic test

This framework gives you a concrete diagnostic: can you inventory every agent operating in your finance function and its owner? If you cannot answer that question in specific terms — purpose, identity, evaluation, retirement — you have not yet reached governed production, regardless of how sophisticated the underlying model is. The same test applies to any vendor claiming a mature offering.

The finance control plane lifecycle

Report Section 6 – govern life cycle of every agent, not a one-time model approval



6. The finance control plane — practical checkpoint

PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. Could we produce, right now, a complete inventory of every AI agent operating in finance and who owns each one?
2. Does every agent operate under its own identity and least-privilege access, or are agents still using shared human credentials?
3. What is our kill switch and rollback process if an agent acts incorrectly — and has it ever actually been tested?

07

The new finance workforce



The new finance workforce

As agents take on more of the transaction and exception load, finance roles shift from processor to designer, reviewer and supervisor. The scarce skills are process architecture, data fluency, model evaluation, control design and business judgment — not prompt writing. Treat deskilling and over-reliance as live risks, not footnotes, and look for redesigned roles, capacity redeployed, adoption by role, override patterns, and training and competency assessment as concrete evidence that your workforce is genuinely changing rather than simply being asked to do more with the same tools.



7. The new finance workforce — practical checkpoint

PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. Which roles have genuinely been redesigned around supervision and exception-handling, versus simply handed a new tool?
2. Do we have a training and competency plan for staff who now review AI output more than they produce work themselves?
3. Is there a clear, known process for a finance professional to challenge or override an AI recommendation?

08

Proving value beyond headcount



Proving value beyond headcount

Deloitte's adoption/ROI gap — 63% active use, 21% clear ROI — is a standing warning against counting licenses or pilots as value [1]. Insist on value measured in cycle time, exception quality, working capital, accuracy and decision latency, with a credible baseline captured before deployment and full run cost — including human review effort — counted alongside any gain. Bain's finding that only 15–25% of CFOs have actually scaled AI across finance, despite 83% planning double-digit spending increases, is good reason to be skeptical of vendor-reported adoption figures that are not tied to an operating metric [23].



8. Proving value beyond headcount — practical checkpoint

PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. Did we capture a proper baseline before deploying AI, or are we estimating the benefit after the fact?
2. Are we counting the full cost of AI adoption — including human review time — or only the license fee?
3. Which metric, other than headcount, has genuinely moved this year as a result of AI in finance?

09

Access and portability



Access and portability


Corporate procurement increasingly determines which AI capability a finance professional can actually use, independent of what the market's frontier tools can do. Google has included premium AI in eligible Workspace Business and Enterprise plans since January 2025, making the contracted suite a default distribution channel for workplace AI whether or not a given user asked for it [27]; Microsoft can use Purview and Edge for Business as an enforcement point to prevent sensitive data leaving for unmanaged AI applications [28]. The result is a three-part access divide: teams with an approved specialist or frontier work agent; teams limited to the contracted Microsoft or Google assistant; and staff using unsanctioned tools despite policy.

Resist assuming universal access to best-of-breed agents. Examine whether your sanctioned platform functions as an effective innovation substrate, a capability ceiling, or both — and whether restrictive policy is itself creating the pressure that drives shadow workflows. A related and largely unexamined question worth asking of your own organization: if an embedded agent has learned from years of exception decisions and user corrections, does that operating knowledge belong to you, and can it be exported if the ERP or specialist vendor changes?



The AI access divide

Report Section 9 — three groups exist today, and shadow use crosses both of the sanctioned ones



9. Access and portability — practical checkpoint

PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. Do we know precisely which AI tools our finance staff are contractually allowed to use, and which they are actually using?
2. Is our current AI access a deliberate strategic choice, or simply whatever came bundled with our Microsoft or Google contract?
3. Do we have visibility into shadow AI use in finance, and a plan to bring it into a sanctioned channel rather than just banning it?

10

The orchestration layer



The orchestration layer

A fourth layer sits between systems of record, embedded intelligence, corporate suites and external agents: workflow automation and agent orchestration. Representative platforms include n8n, Microsoft Power Automate and Copilot Studio, Workato, UiPath, Zapier and Make. Their role is not to replace the ledger or the model but to coordinate triggers, data movement, deterministic rules, model calls, approvals, actions and evidence across application boundaries. n8n illustrates the category well: more than 400 integrations, a visual workflow builder, code extensions, model choice, AI-agent nodes, execution replay, logging, testing and deployment environments, allowing traditional automation and AI reasoning to be combined while retaining explicit workflow structure [29].



The orchestration layer (cont'd)

This creates a genuine architectural choice — orchestration may belong to an ERP vendor, to Microsoft or Google, to a frontier-agent platform, or to an independent automation product. Independent orchestration can reduce platform dependence and fill genuine cross-system gaps, but BCG's warning applies directly here: without governed credentials, version control, testing, monitoring, segregation of duties and durable technical ownership, finance risks replacing shadow Excel with shadow code, or shadow integration [9]. A representative pattern looks like this: an invoice triggers data retrieval from procurement and ERP systems; deterministic rules perform validation and matching; AI reasons about ambiguous exceptions; a human approves items above a confidence or monetary threshold; the approved action is posted back to the system of record; and the workflow retains execution evidence throughout.

10. The orchestration layer — practical checkpoint

PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. Do we know how many finance workflows already depend on automation platforms outside formal IT ownership?
2. Does every automated workflow that touches money run under a governed service identity, with version control and a rollback plan?
3. Who actually owns production incident response if an orchestrated finance workflow fails?




Where this report sits in GrowCFO's 2026 coverage

This report sits deliberately between GrowCFO's Q1 2026 examination of AI-native product architecture and its planned Q4 2026 focus on AI-enabled FP&A. Its distinctive territory is the redesign and governance of finance work itself.

The Q3 2026 editorial calendar in context

2026 editorial calendar

Three quarters, three distinct questions about the same underlying shift



Selecting the Right Solutions



Showcase demo day: turning this report into a vendor decision

This report is designed to prepare CFOs for a room full of vendors, not just to describe the market. A strong demonstration day tests architecture and operating outcomes rather than rewarding the most fluent chat interface, using the same scenario and evidence rubric across every vendor. This section turns that principle into the checklist, scorecard and questions to carry into the room.

Intelligent Finance Operations Tech Showcase



23 September 2026



3-5pm BST



Virtual

Register for free

FEATURING



zenskar

Before the room: five things to know before you watch a single demo

- Name your three highest-volume, highest-pain finance processes before you arrive — the ones identified against Section 1 — so you score every demo against your own problem, not the vendor's chosen script.
- Inventory what you already have switched on: embedded features already live in your ERP, close or AP tools, and what your Microsoft or Google contract already includes. A demo that only repeats a capability you already pay for and do not use is not a finding.
- Set your weighted decision criteria, and who signs off on them, before the first demo starts — not after the most polished one finishes.
- Ask every vendor in the same category to run the same anonymized transaction, invoice or dataset live, rather than accepting a pre-built walkthrough.
- Decide in advance who at the table holds veto rights on data residency, security and identity — and give them explicit authority to use it in the room.

Smart teams deserve intelligent software.

Growing organizations shouldn't have to choose between efficiency and control. **iplicit brings financial data, processes and reporting together in one cloud platform**, automating routine work, giving real-time visibility and better information for confident decisions.



Real results from real teams

66%

Faster month-end close

du Boulay cut month-end close from three weeks to under one week with iplicit.

50%

Less time spent on reporting

Dawn Capital halved management reporting time, from 60-80 hours to around 30 hours a month, with iplicit.

3,000+

UK organizations use iplicit

Triple ISO certified. Built to be trusted.

"iplicit has given me peace of mind. I feel we've future-proofed the financial team's way of working."

Diane Foote · Head of finance · Beyondly

In their own words

"We can now do at least 3.5 days of the old way of things in less than a single afternoon."

| Profiles Creative replaced manual processes with days won back every week.

Lee Cook, Profiles Creative

"iplicit is so easy to use that we've got a better eye on things now than we've ever had."

| 1-2 days off monthly management accounts, 250 users with live visibility.

Carl Leah, Financial controller, Field Studies Council

"It enables reporting across all entities at a fraction of the cost."

| Real-time multi-entity consolidation without the enterprise price tag.

Colin H., CFO ★★★★★ G2

One source of truth

iplicit connects to the tools you already use, keeping everything in sync.

You see further

Your whole organization, live. No exports, no waiting.

You do less chasing

Approvals route, escalate and complete automatically.

Grow without the overhead

Scale your organization without adding headcount.



Modern finance Unlocked

Cloud accounting for the mid-market

See iplicit in action

iplicit.com/book-a-demo


zenskar | AI-Native Revenue Automation



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."

Sardine

Prove

swift
NAVIGATION

Luma

posh

CYFLARE

thrive

dataweb

asigra

Gremlin

[Book a Demo →](#)

Use cases across operational finance


When selecting software, ask vendors to demonstrate:

- embedded intelligence in AP/spend, ingesting an invoice, validating and matching it, resolving one exception and exposing the audit rationale;
- embedded intelligence in close/control, reconciling a material account, investigating an anomaly, assembling evidence and showing approval and rollback;
- revenue/AR/cash orchestration, identifying a collection or billing risk and coordinating actions across systems while preserving customer and finance controls;
- and a cross-application work agent, perhaps starting from an approved finance data pack and contract set, investigating a variance, updating a workbook and drafting a review-ready narrative — showing lineage and permissions throughout.

An orchestration-platform demonstration can be folded into the fourth session or run as a fifth, showing a governed cross-system workflow rather than a prebuilt community template, with deterministic rules, model calls, approval gates, credentials, logs and rollback all made visible.

The showcase demo day journey

Demo day toolkit – one coherent operational story, not four unrelated pitches



Questions to put directly to each vendor

AP / spend (embedded transaction layer)

- Show us a real exception, not the clean example — what happens next, and who approves it?
- What is your live customer's straight-through processing rate, not your target rate?
- What data leaves our environment, and where is it processed and stored?

Close / control (embedded assurance layer)

- Walk us through a material reconciliation from anomaly to approved evidence, live.
- What audit evidence does this produce automatically, and has an external auditor accepted it?
- What is your rollback process if the system posts something incorrectly?

Revenue / AR / cash (orchestration across systems)

- Show the same customer record moving across billing, collections and cash — not three separate screens.
- How do you handle a dispute or an ambiguous payment match?
- What identity does this run under when it writes back to our systems of record?

Cross-application work agent

- Give it an unfamiliar task live, using our data pack — not a rehearsed one.
- What can it read, infer, write and send, and under whose permissions?
- If we ask it to stop and explain itself, can it show its reasoning and sources?

Vendor evaluation scorecard

Criterion	What good looks like	Score (1–5)	Notes
Business outcome	Solves one full, named finance problem end-to-end, not a feature shown in isolation.		
Truth and lineage	Every number and assertion can be traced live to an approved source, on request.		
Autonomy boundary	Vendor states plainly what it does unaided, what needs approval, and what it will never do.		
Exception quality	Shown handling an ambiguous case, not just the clean example.		
Control evidence	Identity, permissions, approvals and model/tool versions are logged and can be shown, not just described.		
Interoperability	Demonstrated working across your actual stack, not a single-vendor sandbox.		
Measurable value	A credible baseline, an operating metric and a total cost, not a productivity claim alone.		
Human adoption	The reviewer's job is visibly clearer and better, not just faster.		

Red flags to watch for in vendor demonstrations

- The demo never touches an ambiguous or messy case — only the clean path.
- No named production customer or metric is offered, only a roadmap or a lab result.
- Vague or evasive answers on data residency, retention or model-training use.
- No clear owner-of-record when the agent writes back to a system.
- Pricing is described only as 'contact us', with no consumption or usage-based detail.
- The presenter, not the product, is doing the reasoning live.

From scorecards to a decision

Turn the scorecard into a weighted shortlist, not a gut feeling: agree the weighting for each criterion before the first vendor presents, and total the scores immediately after each session while the room's memory is fresh. Before any commercial conversation, insist on a reference call with a production customer of comparable size and complexity — a case study is marketing; a reference call is evidence. Design the pilot before signing anything: time-boxed, using real data and real exception volume, measured against the baseline metrics you set out in Section 8, not against the vendor's own claims. And insist on a data and configuration portability clause: if the relationship ends, what specifically can be taken away, in what format, and on what notice — the practical answer to the portability question raised in Section 2.

From scorecards to a decision — practical checkpoint

PRACTICAL CHECKPOINT

Three questions CFOs should be answering

1. Have we defined and weighted our decision criteria before the first demo, or will we decide based on which one looked best?
2. Are we asking every vendor to run the same live scenario on our own data, or accepting a rehearsed walkthrough?
3. Does our shortlist process require a production reference call and a data-portability clause, not just a scored demo?



Conclusions

Conclusions

The most defensible position is neither 'autonomous finance has arrived' nor 'agents are hype.' Production-capable systems are emerging in bounded, high-volume workflows, while the wider finance function is still learning to orchestrate them safely. The decisive advantage comes from redesigning work and controls around outcomes, not from buying the most AI-branded feature set.

Technology placement is a CFO decision, not a vendor category. Embedded product AI may be the safest executor; Copilot or Gemini may be your organization's governed distribution layer; an orchestration platform such as n8n or Power Automate may connect deterministic workflows across systems; and Claude Cowork or ChatGPT Work may be the most flexible cross-application colleague. Most organizations will use all four. Your task — and this report's central argument — is to decide which layer owns context, reasoning, action and evidence for each workflow.

The intelligent finance function is not one in which AI does everything. It is one in which every task has a deliberate home, every autonomous action has an accountable owner, and humans spend their judgment where it changes the outcome.

Appendix: the 30-question CFO self-assessment

Every section of this report closes with three practical questions. They are collected here as a single diagnostic that a CFO or finance leadership team can work through directly — section by section, with an honest yes/no/partial for each.

1. The execution threshold

- Where in finance have we moved from AI-assisted drafting to real execution — and can we evidence production volumes, not pilots?
- What share of our finance AI spend has a measured operating outcome attached to it, rather than none?
- Which three finance processes are our strongest candidates for supervised autonomy in the next 12 months, and do we agree on them?

2. Where intelligence lives

- For our highest-volume finance processes, have we deliberately chosen whether AI should sit in the ERP, the corporate suite or an external agent — or did that happen by default?
- What data are we willing to let leave our system of record for an external agent to use, and who approved that boundary?
- If our ERP or specialist vendor changed tomorrow, would we keep the AI capability we have built, or lose it?

3. The intelligent transaction layer

- What is our actual straight-through processing rate in AP today, and what is specifically blocking it from rising?
- Do we know our exception cycle time and duplicate/fraud-detection precision, or only anecdote?
- Have we redesigned the underlying process, or simply automated the process we already had?

Appendix: the 30-question CFO self-assessment (cont'd)

4. Continuous close, control and assurance

- Could we produce audit-ready evidence for an AI-assisted reconciliation today, on demand?
- Is close still a period-end event for us, or have we begun moving to continuous, evidence-as-you-go monitoring?
- Who is accountable when an AI-generated variance explanation turns out to be wrong?

5. The intelligent cash engine

- Do we have one connected view of DSO, promise-to-pay performance and cash-application rate, or three separate reports?
- Where is revenue leakage most likely occurring in our billing and contract-to-cash process — and have we actually tested for it?
- Are we ready to let AI prioritize collections actions on our lowest-risk accounts without a human touch?

6. The finance control plane

- Could we produce, right now, a complete inventory of every AI agent operating in finance and who owns each one?
- Does every agent operate under its own identity and least-privilege access, or are agents still using shared human credentials?
- What is our kill switch and rollback process if an agent acts incorrectly — and has it ever actually been tested?

Appendix: the 30-question CFO self-assessment (cont'd)

7. The new finance workforce

- Which roles have genuinely been redesigned around supervision and exception-handling, versus simply handed a new tool?
- Do we have a training and competency plan for staff who now review AI output more than they produce work themselves?
- Is there a clear, known process for a finance professional to challenge or override an AI recommendation?

8. Proving value beyond headcount

- Did we capture a proper baseline before deploying AI, or are we estimating the benefit after the fact?
- Are we counting the full cost of AI adoption — including human review time — or only the license fee?
- Which metric, other than headcount, has genuinely moved this year as a result of AI in finance?

9. Access and portability

- Do we know precisely which AI tools our finance staff are contractually allowed to use, and which they are actually using?
- Is our current AI access a deliberate strategic choice, or simply whatever came bundled with our Microsoft or Google contract?
- Do we have visibility into shadow AI use in finance, and a plan to bring it into a sanctioned channel rather than just banning it?

10. The orchestration layer

- Do we know how many finance workflows already depend on automation platforms outside formal IT ownership?
- Does every automated workflow that touches money run under a governed service identity, with version control and a rollback plan?
- Who actually owns production incident response if an orchestrated finance workflow fails?

Source notes and bibliography

Sources were selected for relevance, recency and institutional credibility. Survey figures retain the source's own population and wording and should not be compared as if methodologies were identical. Vendor materials are used to describe product direction and should be validated through disclosed customer evidence where possible. Accessed 6 August 2026.

[1] Deloitte. Finance Trends 2026. Oct. 8, 2025. Survey of 1,300+ global finance leaders: 63% actively use AI; 14% have fully integrated agents; only 21% report clear, measurable ROI.

[2] Deloitte. The CFO Guide to Tech Trends 2026. Mar. 24, 2026. Agentic readiness constrained by legacy integration, data architecture and governance/control frameworks.

[3] PwC. PwC and OpenAI build an OpenAI-native finance function. May 5, 2026. Agent build across procurement, planning, reporting, payments, treasury, tax and close, with human supervision.

[4] PwC. How AI agents help drive a new finance operating model. 2025–26. Frames agents as an operating-model lever; reports up to 80% procure-to-pay cycle-time reduction in its example.

[5] PwC. Banking CFO insights and trends 2026. Jul. 2026. Restrictive internal tool access can slow CFO-office experimentation and agent development.

[6] KPMG. AI in Finance 2026 survey. May 11, 2026. 93% of US companies expect to deploy or scale AI in finance within 18 months; half plan multi-agent systems.

[7] EY. 2026 Global DNA of the CFO Survey. Jun. 3, 2026. Only 21% describe finance as advanced or leading in readiness.

[8] BCG. The CFO's AI Agenda: From Automation to Advantage. May 12, 2026. Distinguishes passive copilots from sense–reason–act agents.

[9] BCG. Vibe Coding Is Coming to Finance. CFOs Need Guardrails. Jun. 4, 2026. Warns finance risks replacing shadow Excel with shadow code.

[10] Forrester. The State of Agentic AI in 2026. Jun. 3, 2026. Adoption intent high; scaled multi-agent production remains rare.

[11] Forrester. Top Agentic AI Use Cases for AP Automation in 2026. Apr. 21, 2026. Capture, exceptions, matching and fraud/risk advancing toward supervised autonomy.

Source notes and bibliography (cont'd)

[12] Forrester. Agentic AI Enters Its Enterprise Execution Era. May 21, 2026. Transition from chat/prototype to enterprise execution.

[13] Accenture. Top Banking Trends for 2026. 2026. Advocates an AgentOps control tower.

[14] Gartner. CFO Guide to Governing Agentic AI. Mar. 9, 2026. Oversight, controls and continuous evaluation.

[15] Gartner. A CFO's Guide to Building a Finance AI Agent. May 18, 2026. Use the first agent to test governance before scaling.

[16] Anthropic. How finance teams use Claude Cowork. Jun. 2, 2026. Forecast updates, variance commentary, flux analysis, board-packet assembly.

[17] OpenAI. ChatGPT Work. Jul. 9, 2026. Multi-hour agent producing finished work across apps and files.

[18] OpenAI. ChatGPT for Excel and financial data integrations. 2026. Spreadsheet-native modeling and approved data-source access.

[19] Microsoft. Finance agents in Microsoft 365: 2026 release wave 1. Mar. 18, 2026. Finance Agent in Excel, Outlook, Teams and Copilot.

[20] Google Cloud. Gemini Enterprise for the agentic task force. Apr. 22, 2026. Long-running agents, shared workspaces, identity, registry, observability.

[21] Google Cloud. Gemini Enterprise Agent Platform. Apr. 22, 2026. Agent development, orchestration, governance and security.

[22] PwC. The agentic revenue engine: reinventing lead-to-cash. Jul. 20, 2026. Identity, policy, memory, tool access and auditability across lead-to-cash.

[23] Bain & Company. 42% of CFOs plan to increase AI investment by over 30% within two years. Apr. 13, 2026. Only 15–25% have scaled AI across finance; workflow debt must be paid down first.

[24] KPMG. 2026 Global AI in Finance Report: The Decision Advantage. 2026. AI gains strongest in judgment-heavy work; assurance readiness a performance differentiator.

Source notes and bibliography (cont'd)

[25] Gartner. Embedded AI in cloud ERP will drive a 30% faster financial close by 2028. Feb. 24, 2026. Forecast, not a demonstrated 2026 outcome.

[26] Anthropic. How Anthropic's finance team uses Claude to shape the narrative behind the numbers. May 22, 2026. First-party case: 10–20 hours reclaimed weekly.

[27] Google Workspace. Get started with Google Workspace with Gemini. 2025–26. Premium AI included in eligible plans since January 2025.

[28] Microsoft. Help prevent sharing to unmanaged AI apps via Edge for Business. Jul. 2026. Purview/Edge as enforcement point against unmanaged AI.

[29] n8n. Workflow automation platform features. Accessed Aug. 6, 2026. 400+ integrations, deterministic logic, AI-agent nodes, execution monitoring.

Research limitations

This report draws on published analyst, consultancy and vendor research rather than a systematic literature review. Some analyst research is available only as a public abstract. Claims about product capability and benefits are not treated as independent proof unless supported by disclosed customer evidence.

Vendor-sourced material is flagged as such throughout. Survey figures retain each source's own population and wording and are not directly comparable across firms.