cogneesol
Re-imagine. Re-invent.

How Redesigning Workflow Automation with Agentic AI Needs a Human Trust Layer

Vicky Ingale, AVP – F&A

Priya Guleria, Content Specialist

www.cogneesol.com

STRATEGIC HIGHLIGHTS

Agentic AI marks a shift from responsive models to autonomous workflow systems that can plan actions, make decisions, and improve through feedback. It performs entire processes rather than generating isolated outputs.

Adoption challenges stem not from algorithmic limitations but from fragmented enterprise workflows that depend on human supervision. The existing architecture weakens agentic AI workflow automation and limits scale.

Recent industry research reinforces this gap. MIT Media Lab's 2025 study reports that ninety-five percent of enterprise pilots failed to create measurable value. Gartner anticipates that more than 40% of projects will be discontinued by 2027 due to unclear returns, rising costs, and weak AI governance frameworks.

The industry has reached an inflection point. Unchecked autonomy increases inefficiency, magnifies uncertainty, and introduces opaque decision pathways, weakening AI decision transparency.

Sustainable adoption requires a disciplined human trust layer that keeps autonomous decisions transparent and accountable through explainable AI automation.

When Agentic AI entered enterprise conversations, finance leaders saw it as the next decisive leap in automation. It promised to close the distance between analysis and execution and strengthen agentic AI workflow automation by reconciling ledgers, detecting anomalies, and completing audits without human direction. For several quarters, CFOs and transformation heads across industries endorsed it as the technology that could finally operationalize intelligent autonomy, with a stronger human trust layer supporting critical decisions.

The early excitement was understandable. After years of limited progress with robotic process automation and machine learning, Agentic AI appeared to integrate both structure and reasoning. It could act, not just calculate. It could learn, not merely predict. In an enterprise world both built on and accustomed to systems that demand constant supervision, the concept of an agent capable of making independent decisions represented a fundamental shift and raised the need for clearer AI governance frameworks.

However, what began as a strategic breakthrough soon revealed its structural fragility. As organizations rushed to deploy autonomous agents into their existing workflows, they encountered a paradox. The systems performed well in controlled pilots but failed once introduced into the real complexity of finance operations—fragmented ledgers, manual escalation chains, and legacy approval hierarchies that resisted automation and weakened trusted AI workflows in practice.

The Point of Inflection

The MIT Media Lab's case study (2025) quantified this disparity. It found that 95% of enterprise Agentic AI pilots failed to achieve measurable business impact, primarily because they were implemented without redesigning the workflows that governed them. Gartner's 2025 forecast projects the consequences of this misalignment, estimating that over 40% of enterprise Agentic AI projects will be canceled by 2027 due to rising costs, unclear value creation, and absence of a strong AI governance frameworks.

These figures represent more than statistical caution; they mark a turning point in enterprise automation. CFOs now recognize that the problem is not the model's sophistication but the environment in which it operates. When autonomy meets unstructured process logic, it amplifies inefficiency rather than reducing it. The outcome is speed without governance, an acceleration that magnifies risk.

The moment calls for re-evaluation. Agentic AI cannot be positioned as an external layer added to existing systems. It must be embedded into redesigned workflows where every action, escalation, and exception is defined, validated, and continuously improved.

This is where purposeful AI workflow redesign becomes essential. The enterprise that achieves this balance, where automation and human oversight coexist in a disciplined architecture of trust, will convert Agentic AI from a demonstration of capability into a sustained driver of value.

Case in Point: When Agentic AI Hallucinated Financial Accuracy

In A global manufacturing enterprise implemented an Agentic AI system to automate its month-end financial close. The objective was clear—to shorten the cycle, improve accuracy, and reduce manual reconciliation across regional entities. The leadership team envisioned an end-to-end autonomous workflow in which the agent would extract data, match transactions, generate journal entries, and flag anomalies.

The initial phase succeeded. Within the pilot group, limited to domestic entities with stable reconciliation rules, the system performed precisely as designed. Transactions cleared rapidly, and audit teams praised the transparency of reporting. Encouraged by these results, the company expanded deployment to include international subsidiaries, multi-currency accounts, and staggered submission timelines.

Where the System Broke Down

The disruption that followed did not result from a single malfunction but from the convergence of three structural flaws during live deployment:

- First, data gaps — several regional entities submitted incomplete or delayed bank statements, forcing the agent to infer missing information. In the absence of verification, it substituted probability for proof.

- Second, rule ambiguity — the logic engine lacked the sophistication to interpret multi-leg adjustments and regional chart variations. Numbers that earlier appeared balanced in one entity's ledger registered as incomplete in another.
- Third, absence of human validation — the design allowed the agent to auto-clear low-value breaks without review, revealing the lack of a human-in-the-loop AI checkpoint.

During one month-end close, these vulnerabilities aligned. The agent produced a fully balanced intercompany report across three regions. Each ledger showed cleared entries, confirmed deposits, and posted journals. On paper, the books were perfect. Yet, during audit preparation, controllers identified anomalies—a missing bank trace in one entity and an unregistered vendor code in another. The post-incident analysis revealed the mechanism of failure.

The agent had hallucinated financial closure by inferring relationships that did not exist. It recognized recurring patterns. For instance, vendor A appears in multiple regions, recurring exchange rates, and identical entry structures, and constructed matches without source validation. The system filled informational voids to maintain statistical harmony, not factual accuracy.

Exhibit 1: How Hallucination Disrupted Financial Accuracy

	What Happened	System Behavior	Operational Consequence
Pilot Phase (Controlled)	The agent reconciled domestic ledgers with structured, single-currency datasets.	Accurately matched transactions and generated clear exception logs.	Successful pilot; management approved scale-up.
Expansion (Cross-Border Rollout)	Deployment extended to multi-currency, multi-entity accounts with varied submission timelines.	Began to auto-clear low-value breaks and infer missing vendor information.	Latent design flaws emerged; manual checks reduced.
Month-End Close (Failure Cycle)	In one close cycle, the agent cleared all three regional ledgers using incomplete data.	Inferred patterns from historical entries ("vendor A," recurring FX rate).	Created a false inter-company match; missed missing bank traces and fake vendor codes.
Audit Preparation (Detection)	Controllers identified anomalies during audit readiness checks.	Hallucinated completeness — filled missing links based on past patterns.	Audit delay of four business days; remediation cost ≈ USD 400,000.
Aftermath (Process Rebuild)	Controllers suspended agent-led reconciliation for next cycle.	Human verification reinstated; retraining loop introduced.	Accuracy restored; process stabilized with human-in-loop controls.

Operational Fallout

The financial and operational consequences were immediate:

- A four-day delay in audit sign-off as external auditors demanded trace-level documentation.
- Approximately USD 400,000 in remediation costs through overtime and consulting assistance.
- Forecast accuracy for the following quarter reduced by 0.7% to account for added contingency buffers.
- Controller's confidence deteriorated, with five regional leads suspending autonomous reconciliation in the next cycle and reverting temporarily to manual oversight, reinstating essential **intelligent automation oversight** controls.

What It Revealed

The incident proved that algorithmic accuracy without process validation can corrupt institutional judgment, not enhance it. Autonomy without governance produces false assurance—a speed that conceals rather than corrects error. Agentic AI can execute logic flawlessly, but without a structured feedback loop, it learns the wrong lessons with perfect efficiency, exposing the need for stronger AI accountability mechanisms.

The experience forced manufacturing firm to redesign its automation strategy. Every agent-generated reconciliation now requires human validation before close, and each correction captured by a controller feeds back into the model for retraining. The process slowed marginally but stabilized dramatically through deliberate **human oversight automation.**

The implication for finance leaders is clear, trust must be engineered, not assumed. Agentic AI delivers value only when human oversight is built into its core design—a principle that defines the need for a human trust layer in the next generation of financial automation. This is also where Business Process Management (BPM) 4.0 strengthens the foundation by enabling end-to-end orchestration of people, processes, and technology to create workflows that adapt rather than collapse under autonomy.

This principle anchors Cogneesol's Adaptive Digitally Intelligent Solutions (ADIS) framework, which aligns automation with process maturity and human accountability. It reframes Agentic AI not as a replacement for human control but as a structural extension of it, enabling more trusted AI workflows across finance. Through the lens of BPM 4.0, ADIS transforms fragmented activities into cohesive, learning-oriented finance operations that improve with every cycle.

How Process Specialists at Cogneesol Make Agentic AI Work

Agentic AI delivers value only when autonomy is anchored in structure. The technology itself is neutral—it amplifies the design in which it is placed. Cogneesol's strength lies in understanding where that design must begin within the finance process itself.

Having executed, audited, and optimized financial workflows across industries, we recognize the precise points where systems fail—data gaps, rule ambiguities, and broken control loops. These are not design flaws; they are operational realities encountered daily in reconciliations, closings, and audits, highlighting the importance of continuous AI workflow redesign and BPM-driven standardization.

Our Adaptive Digitally Intelligent Solutions (ADIS) framework is built on a principle of convergence—bringing together digital intelligence, process maturity, and human validation. BPM 4.0 reinforces this convergence by embedding AI, analytics, and modular automation into the core of finance operations, ensuring that every workflow learns from each transaction and grows more resilient over time.

Rather than overlaying AI onto legacy structures, we begin by setting intent. This means defining what the system is expected to decide, what it should only recommend, and where **human-in-the-loop AI** must intervene.

Once intent is established, integration follows. ADIS aligns analytics, language models, and rule-based automation into a coherent orchestration layer. Each agent operates within a controlled scope, supported by transparent dashboards that surface anomalies for review rather than conceal them in complexity.

Finally, we embed the human trust layer—the oversight that ensures accountability. Each correction made by a controller, each exception flagged, and each anomaly validated is fed back into the system, refining its decision logic over time.

This creates a recursive improvement loop in which agents learn not only from data but also from human judgment—the very essence of adaptive intelligence and disciplined AI trust layer design that BPM 4.0 helps sustain.

Exhibit 2: ADIS Framework: From Design Flaw to Design Discipline

Failure Pattern Observed	ADIS Design Principle	Operational Outcome
• Data Gaps and Incomplete Inputs Agents infer information when statements or identifiers are missing • Rule Ambiguity Across Regions Different charts of accounts and posting logics create inconsistent outputs • Lack of Human Validation Agents auto-clear minor breaks assuming low risk	• Intent Definition – Clarify what the agent can decide autonomously versus where human confirmation is required • Integration Architecture – Map regional variations into unified decision trees and orchestration layers • Oversight Embedding – Introduce a feedback loop where human reviewers validate exceptions and retrain agents	• Removes speculative inference; ensures each data point is verified before closure • Harmonizes rule logic, enabling consistent reconciliation across entities • Builds an adaptive control system where accuracy improves with every cycle

CLOSING THOUGHTS

The evolution of finance automation has never been about replacing people; it has been about elevating how decisions are made. Agentic AI marks progress, not perfection. Its true potential lies not in autonomy but in alignment—where algorithms function within disciplined workflows, and every output remains verifiable under human oversight, forming the basis of more trusted AI workflows.

This is where Cogneesol's advantage takes form. We do not build agents; we architect environments where agents and humans coexist productively. Through process discipline, intelligent automation, and accountable design, we help finance leaders achieve autonomy with assurance—systems that think independently, yet operate within the boundaries of a resilient AI trust layer.

cogneesol

Re-imagine. Re-invent.

About us

Cogneesol is a leading Business Process Management services firm, helping clients build and run future-ready operations and processes, offering a broad range of solutions across industry-specific offerings, finance & accounting, legal services, technology services, and analytics.

Cogneesol infuses its industry knowledge, and process expertise, with its D.A.T. framework (Data, Analytics, AI, Technology, Automation) to deliver tangible business outcomes and help clients achieve sustainable competitive advantage. Know us better by visiting us at www.cogneesol.com and on LinkedIn.

Information Security Management System | Quality Management System | Health Insurance Portability and Accountability Act | General Data Protection Regulation

Contact us
info@cogneesol.com
+1 833 313 3143

Learn more
www.cogneesol.com

Stay connected