

Beyond Bots: Moving Financial Services from AI Experimentation to Agentic Operations

Leigh Varnham

Chief Growth Officer – UK and EU, WNS-Vuram



[Banking and Financial Services](#) (BFS) has moved from asking what Artificial Intelligence (AI) can do, to asking where it can create measurable enterprise value. Early [Generative AI](#) (Gen AI) adoption was largely exploratory, with financial institutions testing use cases while navigating model risk, data, governance and explainability concerns. The next phase is different: The priority is to move from isolated AI tools and pilots to [Agentic AI](#) workflows that can operate across systems, functions and decision points — with the Human-in-the-Loop (HITL) controls required for a regulated environment.

A recent WNS live poll of BFS leaders points to [an uneven transition](#): Exactly one-third of respondents were experimenting through ad hoc trials, another third were running limited departmental pilots and a final third had achieved full automation in select processes. Two-thirds therefore remained in what the poll characterized as a transition zone.

The implication is not that financial institutions lack AI capability. It is that scaling AI requires a different operating model — one that treats workflow re-design, orchestration, governance and human accountability as part of the transformation rather than as layers added after the technology is selected.

The Barriers to Scale: Why Agentic AI Is Hard to Operationalize

Financial services workflows often reflect the constraints of legacy platforms, fragmented data and processes built around sequential human handoffs. Following decades of mergers and acquisitions, many institutions now operate across complex technology estates, making it difficult to move information cleanly from one stage of a journey to the next. Where a workflow spans multiple systems, simply inserting an AI agent into one step can improve local efficiency without removing the underlying friction.

Regulation raises the bar further. Financial institutions need governance that can establish accountability, monitor AI use and provide appropriate transparency and explainability for material decisions. The Financial Stability Board's 2026 guidance¹ emphasizes organization-wide AI

governance, lifecycle risk management and explainability, while the Bank for International Standards (BIS)² has stressed that human responsibility and accountability must remain central to AI-enabled financial decision-making.

This makes agent sprawl more than a technology-management issue. When business units adopt AI agents independently, organizations can lose visibility into what those agents are doing, what data they can access and how decisions move across an end-to-end process.

The strategic response is not to eliminate experimentation, but to create a governed path from experimentation to production.



¹<https://www.fsb.org/2026/06/sound-practices-for-responsible-adoption-of-artificial-intelligence-ai-consultation-report/>


²<https://www.bis.org/speeches/20260416-ai-finance-what-can-change-what-must-never-change>

Where to Start:

Re-design the Workflow, Not Just the Task

A useful starting point is the workflow rather than the individual task. The strongest candidates for Agentic transformation are processes that are long-running, document-intensive and dependent on multiple systems or teams. In financial services, these often occur at the handoff between client-facing teams and middle- or back-office operations.

Re-imagining Legacy Processes with Agentic Workflows



Potential use cases for this include:

Customer Onboarding

Coordinating information gathering, Know Your Customer (KYC) / Anti-Money Laundering (AML) checks, document review and account setup across multiple platforms.

Life-event Servicing

Orchestrating changes such as beneficiary updates or account changes across lending, investment and credit relationships, with appropriate advisor notification.

Lending Documentation

Gathering, extracting and validating information from complex document packages before human credit decisions are made.

The strategic shift is to ask what the ideal agent-supported workflow should look like if the institution were designing it today – not how to automate each step of a process inherited from the past. That distinction separates task automation from operating-model transformation.

Re-defining ROI in the Agentic Era

Traditional automation programs often emphasized measures such as Full-Time Equivalent (FTE) reduction, cost per transaction and processing speed. Those metrics remain relevant but do not capture the full value of an Agentic operating model. A broader ROI framework should connect operational performance to enterprise outcomes.

A practical framework can assess value across five dimensions

 Revenue Expansion	 Customer Experience	 Risk and Compliance	 Employee Experience	 Organizational Agility
Faster cycle times, stronger relationship-manager capacity and more opportunities to act on customer needs	Lower customer effort and shorter waits for decisions or servicing outcomes	More consistent controls, stronger decision trails and fewer manual data-entry errors	Less administrative chasing and greater capacity for judgment-intensive work	Faster adaptation when products, policies, market conditions or regulatory requirements change

Competitive advantage does not rely simply on how much labor and costs an institution can save. It lies in what additional value can be created when scarce expert capacity is re-directed from administrative coordination toward judgment, relationships and higher-value work.



A Case Study:

Automating the Investment Lifecycle

A global investment management firm faced operational friction owing to fragmented, manual workflows across spreadsheets, legacy systems and e-mail. Core activities such as investor onboarding, record maintenance, approvals and reporting were time-consuming and lacked end-to-end visibility.



[WNS-Vuram](#) implemented an Appian-powered automation solution across more than 15 applications, streamlining workflows, approvals and investor record management. The solution enabled:



Rapid onboarding of more than **400 investors**



Automated generation and distribution of over **150 investor reports**, reducing manual effort and delivery errors



Automated reading, interpretation and routing of **vendor e-mails**, reducing administrative delays



Substantial **process and productivity improvements** through the integration of Appian applications into daily operations

This real-world example illustrates an important distinction:

The value came from connecting multiple activities into a coherent workflow, rather than automating a single task in isolation.


Keeping Humans in the Loop: Match Autonomy to Risk

Agentic automation should not be treated as an all-or-nothing decision. A more practical model is adjustable autonomy: Agents can begin by gathering information, preparing recommendations or executing low-risk actions while human professionals validate material decisions. As evidence accumulates, organizations can expand autonomy where controls and performance support it.


For high-impact decisions such as commercial lending, agents may be well suited to gather, synthesize and extract information from large documentation sets. Final credit decisions, however, should remain subject to the institution's risk framework, applicable controls and accountable human judgment. This is consistent with emerging guidance that emphasizes human responsibility, explainability and appropriate oversight in financial AI.

The Human-in-the-Loop Autonomy Dial

Adjusting the right level of autonomy based on risk, complexity, and business impact



Technology fit matters too. Large Language Models (LLM) and Agentic systems are strong at language-heavy work such as document interpretation, information retrieval and conversational interaction. Deterministic calculations, credit-scoring logic and rule-based compliance controls may require other technologies. The target architecture is therefore not an LLM-only stack, but a hybrid model in which Agentic intelligence orchestrates specialized tools.



About WNS-VURAM

WNS-Vuram, part of Capgemini, is a global leader in hyperautomation, empowering organizations to accelerate digital transformation and realize measurable business value through an 'AI-first' approach. We build custom applications, agentic automation, and industry-aligned solutions that address real-world business challenges. With deep expertise across platforms such as Appian, ServiceNow, UiPath, Microsoft, AWS, and Automation Anywhere, we design and deliver cross-technology solutions tailored to client needs. Our ability to seamlessly integrate multiple platforms enables enterprises to drive agility, scale innovation, and reimagine customer and employee experiences.

For over 13 years, more than 200 global enterprises, spanning BFSI, Healthcare, Shipping, Logistics, and Travel and Hospitality, have trusted WNS-Vuram to deliver outcomes that matter. We combine domain expertise, technical excellence, and a robust partner ecosystem to help enterprises achieve sustainable innovation. Our Gen AI and agentic automation capabilities fuel smarter processes, faster decision-making, and long-term value creation. At WNS-Vuram, we go beyond transforming processes—we co-create intelligent enterprises with our clients, delivering better outcomes for their people, customers, and the communities they serve.

For more information, visit <https://www.vuram.com>.