VIEW POINT

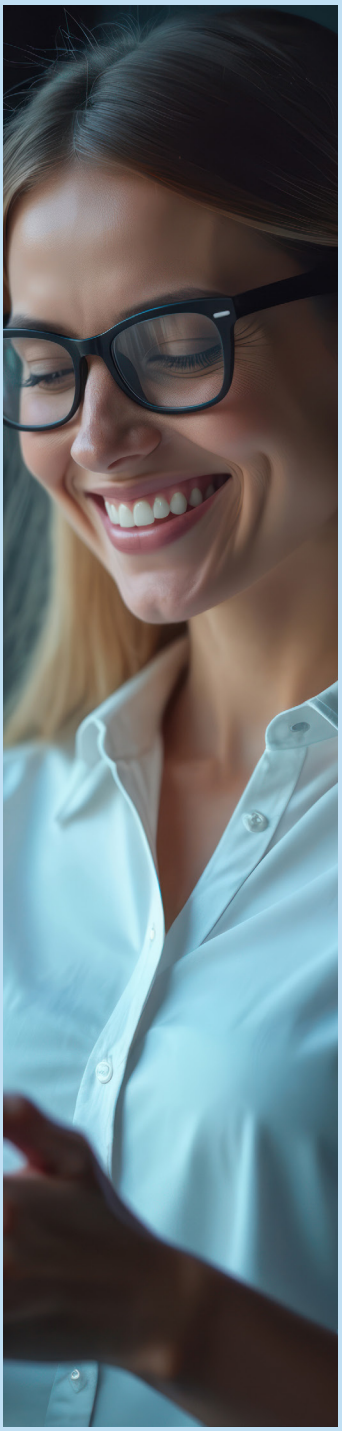

A NOVEL OPERATING MODEL FOR AI TRANSFORMATION IN FINANCIAL SERVICES

Infosys | CONSULTING

Executive Summary

Banks investing heavily in AI struggle to move beyond pilot successes to portfolio-scale impact. This is not because of technological limitations, but because of organizational structure: when accountability for AI initiatives is diffused across IT, business units, risk functions, and finance, no single entity owns the outcomes.

Initiatives stall due to cross-functional coordination failures and resources conflict, with benefits going unmeasured. The **Intelligent Transformation Office (ITO)** addresses these complications through an integrated operating model with end-to-end accountability: a central office orchestrates AI Initiatives with authority to resolve cross-functional blockers, specialized resource pools provide deep expertise allocated dynamically to high-priority initiatives, and embedded delivery squads work directly within business units to ensure solutions to expedite adoption. These are augmented by AI agents that automate monitoring portfolio health, predict emerging delivery risks, surface transformation opportunities and automate benefit tracking of AI initiatives. This AI augmented transformation office allows scaling AI across enterprise using a lean structure.

For the same capital investment, an ITO delivers 3x more business value compared to traditional approaches driven by measurable improvements in initiative success rates, time-to-value, benefit realization, and resource efficiency.

The article details how banks can implement the ITO model to transform AI-driven initiatives from isolated experiments into enterprise capability.

The AI adoption challenge in Banking: promise vs reality

In 2025, a major retail bank celebrated what looked like a breakthrough. After several quarters of experimentation, one of its business units finally delivered a successful GenAI pilot: a credit-assessment assistant that reduced analyst effort by 40%. The CEO highlighted it in the annual report, the CIO called it "a blueprint for scale", and the board asked for ten more initiatives just like it.

But twelve months later, the pilot was still alone: the blueprint never became a building.

Each department had its own priorities, risk wanted more validation, legal wanted new guidelines, data teams were overwhelmed, Finance could not see a business case for the other use cases. And the business units, the ones who were supposed to benefit, were unsure who owned what.

The pilot had proven the promise, and the organization had proven the problem.

Most banks remain stuck in exactly this pattern: promising pilots that never scale because no single function owns the end-to-end responsibility to turn isolated wins into enterprise-wide value. That is the gap an Intelligent Transformation Office (ITO) is designed to fill.

Artificial intelligence is quickly moving from experimental technology to strategic imperative for banks. The potential use cases span the enterprise, from customer-facing chatbots and personalized recommendations to operational automation in compliance, risk management, and back-office processing.

In banking operations, we believe there is a huge opportunity to unlock value through agentic-led automation. For example, in credit underwriting, multiple AI agents could work together to generate a credit file from structured data, validate against policy in real-time, and embedded the decision with reasoning and traceability with human in the loop for key decisions. Take compliance report generation wherein multiple AI agents can compile evidence automatically and generate its own audit trail. Use cases like these involve repetitive documentation, cross-referencing, and sensemaking, areas where AI can dramatically improve speed and augment human decisions.

Industry estimates project that in the coming years AI could add $200-340 billion in annual value across the global banking sector through productivity gains and enhanced decision-making.

For most banks though, this promise remains largely unrealized.

Since Generative AI broke into the scene, financial institutions have been trying to establish a linkage between investments made vs the ROI, but results have not been encouraging. As per a recent Infosys Bank Tech Index report, only 25% of AI initiatives have generated business value and a whopping 49% of AI initiatives are stuck at pilot or proof of concept stage.

According to a study, the key challenges relate not to technology or algorithms but processes and people. 70% of respondents in the study cited people and process related challenges such as establishing ROI, prioritizing opportunities, ambiguity on KPIs, leadership alignment, etc. as key hurdles preventing the move from experimentation to scaled implementation.

Traditional approaches for AI implementation

Let's look at the different approaches that banks have adopted for executing AI initiatives

Fragmented Business Unit Initiatives

Individual business units experiment with AI solutions to address specific pain points. Each initiative operates independently, using different technology stacks, governance frameworks, and vendor relationship resulting in duplicated effort/solutions

Centralized Innovation Labs

Some banks establish centralized AI COEs or innovation labs which build impressive technical capabilities; however, they struggle to transition solutions from the lab to production environments at scale and the lab becomes an expensive showcase rather than a delivery engine

IT-Led Technology Programs

Other banks approach AI as fundamentally a technology challenge, delegating leadership to the IT organization. Solutions are technically sophisticated but lack link to business value

In all these approaches what's missing is a singular function that owns overall outcomes when initiatives fail to deliver value - one function for assessing, prioritizing, governing, building, deploying and monitoring AI use cases at enterprise level.

Addressing the challenges above requires a fundamentally different approach, one that combines Governance, Delivery, Risk and Value Realization augmented by AI agents under one umbrella establishing end-to-end accountability for success/failure of AI initiatives.

The **Intelligent transformation office (ITO)** helps enable this.

What is an ITO?

The ITO office is an operating entity with P&L responsibility, decision authority, and direct accountability to the executive leadership for converting AI investment into measurable business value.

It is not simply a rebranding of existing approaches but rather a comprehensive framework designed specifically to overcome the organizational impediments that prevent AI from scaling in highly regulated, complex enterprises like banks. It is a **commercial and operational backbone** that ensures AI deployments are safe, scalable, and tied directly to business value.

The "Intelligent" in ITO reflects a fundamental design principle: the transformation office itself is augmented by AI agents that act as force multipliers across core transformation functions - portfolio analytics, business analysis, benefit tracking, and compliance validation. These agents don't replace human discernment but amplify it by handling data synthesis, pattern detection, and routine validation, allowing small expert teams to maintain oversight depth across large portfolios while focusing their expertise on strategic decisions, risk assessment, and stakeholder management.

Let's begin by looking at the key components of an ITO.

Key components of an ITO

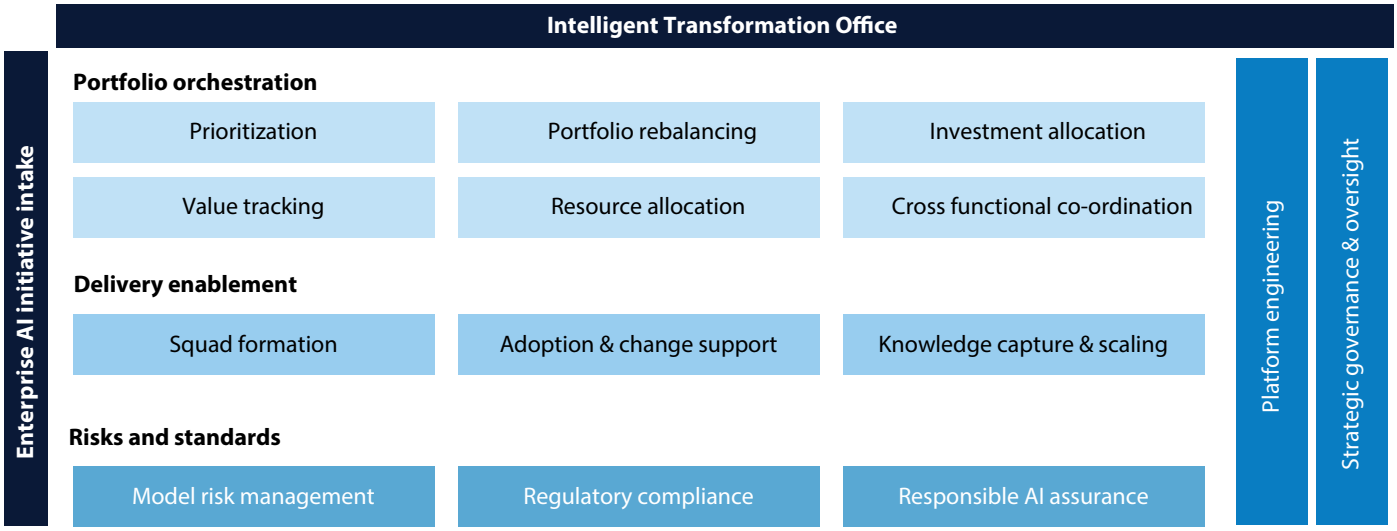

Intake & Opportunity Assessment

Business units submit AI opportunities through a standardized intake process administered by the ITO. Each submission includes a quantified business case, expected value creation, high level requirements, and regulatory considerations. This centralized intake prevents duplicate efforts across business units and ensures all AI initiatives are evaluated consistently against enterprise-wide criteria.

Strategic Governance & Oversight

The Strategic Governance Forum operates as the investment committee for AI transformation and comprises of C-suite executives and the ITO leader/representative. Rather than reviewing every initiative, it focuses on strategic decisions: which business domains receive priority, how much capital to allocate to AI versus other investments, whether the portfolio is delivering expected returns, and when to pivot or double down based on results. The committee's goal is to ensure alignment between AI initiatives and broader enterprise strategy. The ITO representative is key part of this forum and acts as a bridge between organization goals and ITO initiatives.

Portfolio Orchestration

It operates as a unified portfolio where AI initiatives from across the enterprise compete for funding based on strategic value. All AI initiatives enter through a single gateway where they are evaluated against strategic priorities, technical feasibility, and organizational readiness.

As the portfolio matures, the ITO rebalances investment, scaling proven use cases, restructuring underperformers and terminating failures. This dynamic portfolio management ensures capital allocation in highest-value initiatives rather than speeding resources evenly across all projects. The office operates with authority to convene stakeholders, resolve conflicts, and drive decisions across organizational boundaries: this is essential because AI transformation inherently cuts across traditional silos.

It also assigns delivery squads, allocates specialists capacity across the portfolio to ensure high-priority initiatives are allocated the right talent they need.

Agentic AI led Portfolio KPI measurement and benefits tracking

The most persistent failure in transformation is the gap between approved business cases and actual value delivered. The office closes this gap through systematic value measurement using rigorous methodologies that isolate AI impact from other factors. When benefits fall short, root causes are investigated and corrections implemented. When benefits exceed projections, success factors are identified and scaled.

With AI Agent, portfolio performance measurement and benefit tracking can be increasingly automated. Tracking the five-core metrics : portfolio ROI, benefit realization rates, delivery velocity, adoption rates, and initiative outcomes traditionally requires analysts to manually reconcile data across finance systems, project management tools, and operational platforms.

AI agents can automate this cross-system data synthesis, continuously monitoring actual performance against projections, flagging variances that exceed thresholds, and maintaining measurement consistency across the portfolio which can be reported monthly to CFO and quarterly to Strategic Governance Forum. This data-driven approach enables informed portfolio decisions and demonstrates accountability for AI transformation outcomes.

Delivery Enablement

The ITO maintains deep specialist expertise centrally while deploying it dynamically where business value is highest. Specialists are organized into capability pools or Communities Of Practice:

- AI and machine learning engineers who build and optimize models
- Business analysts who translate business requirements into technical solutions
- Data scientists who develop analytical approaches
- Change managers who drive adoption.

These AI augmented specialists are not just domain experts who occasionally use AI tools, but professionals who have deeply integrated AI augmentation into their workflows. Business analysts use AI assistants to validate requirements against technical and regulatory constraints and accelerate documentation. Change managers employ AI tools to generate targeted communication materials, analyze adoption barriers from usage data and feedback, and create personalized training content at scale. Data scientists leverage AI coding assistants to accelerate model development and testing. This AI-augmented capability model enables the ITO to deliver with smaller, more productive team wherein AI handles synthesis, drafting, and analysis while humans focus on judgment, stakeholder engagement, and strategic decisions.

Embedded delivery Model

The ITO employs an embedded approach: cross-functional squads formed from the specialized capability pools deploy directly into business units for initiative right from design through deployment and optimization.

Squads are formed by combining specialists from the capability pools with domain experts from the business unit. A typical squad includes two to three AI/ML engineers, one to two business analysts, one change manager and several domain experts from the business unit who provide deep SME knowledge to ensure solutions address business need.

Over time, successful approaches are documented, standardized, and made reusable. Technical architectures, governance frameworks, delivery methodologies and change management strategies become organizational assets that build institutional knowledge. Squads leverage these reusable components (e.g., standardized business case templates, proven governance frameworks, model validation playbooks, and change management toolkits developed through prior initiatives) to accelerate delivery by 40-60% compared to building from scratch. This systematic approach to knowledge capture and scaling ensures the portfolio learns and improves with each initiative rather than repeating discovery work.

Risks and Standards

The ITO enforces governance standards across the entire AI portfolio. All initiatives must comply with the bank's model risk management framework, explainability requirements, bias testing protocols, and regulatory obligations before deployment. Rather than allowing each initiative to interpret governance requirements independently - creating inconsistency and compliance risk - the office articulates bank's uniform standards and coordinates with Risk, Compliance, Legal, and Audit functions.

The ITO also enforces AI usage policies, ensuring all initiatives comply with ethical AI principles, responsible AI frameworks (including bias detection, fairness testing, and ethical review), data governance rules, and approved development standards. This includes ensuring explainability and transparency (model interpretability, decision transparency, audit trails) and data privacy and security (PII protection, data access controls, security protocols).

When governance issues emerge, they are identified and resolved at the portfolio level rather than discovered during regulatory reviews. This centralized oversight ensures the bank can confidently deploy AI into sensitive areas while maintaining regulatory credibility and avoiding the compliance failures that derail AI programs.

Governing ITO's AI Agents

The ITO employs several AI agents that need to be governed. All ITO agents operate under the same risk management framework that applies to any enterprise AI solution: validation before deployment, continuous monitoring for drift or errors and transparency/explainability requirements for automated recommendations.

Describing all the possible governance models that may be adopted is outside of the scope for this paper (it is a highly bespoke endeavour to be customized to the specific organization), but a few broad governance categories to govern ITO agents are presented below.

Technical Governance

- **Agent-to-agent validation** - Multiple agents cross-check each other's outputs (e.g., one agent flags risk, second agent validates the signal)
- **Constraint-based guardrails** - Hard-coded rules agents cannot violate (spending limits, data access boundaries, action thresholds)
- **Confidence scoring** - Agents self-assess confidence levels; low-confidence outputs trigger additional validation by humans

Process Governance:

- **Tiered escalation protocols** - Agents handle routine cases autonomously only and escalate high impact decisions to humans based on complexity
- **Audit trail automation** - Every agent action logged with reasoning, data sources, confidence scores

Organizational Governance:

- **Agent performance dashboards** - Real-time monitoring of agent accuracy, error rates, intervention frequency etc.
- **Stakeholder review forums** - Business units provide feedback on agent recommendations and that helps agents self-learn and improve outputs

Human-in-the-loop design is paramount for critical and high impact decision so that accountability stays with humans for agent output. Let's say an agent flags an initiative as high-risk due to timeline slippage, the system generates explanations citing specific data points (three missed milestones, resource utilization below target thresholds), a portfolio manager still reviews the full context as to whether there was stakeholder alignment issues, dependency delays outside the team's control etc. before deciding whether to escalate to governance or provide targeted support.

Enterprise AI Platform Foundation

The ITO leverages centralized AI/ML platforms, cloud infrastructure, and data integration capabilities maintained by the enterprise Platform & Engineering team (typically under CTO/CIO ownership). Rather than building parallel infrastructure, the ITO defines requirements and collaborates closely with Platform teams to ensure enterprise technology supports AI delivery needs.

This includes model deployment environments, MLOps tooling for continuous model monitoring and retraining, data access frameworks that enable squads to access necessary data while maintaining security and privacy controls, and standardized development environments that accelerate squad productivity. The ITO provides input on platform roadmaps and capability gaps but does not own infrastructure, maintaining clear separation between AI initiative delivery (ITO responsibility) and underlying technology enablement (Platform team responsibility).

Business value and KPIs

The business benefits of ITO are compelling.

For the same capital investment, an ITO delivers **3x** more business value compared to traditional approaches driven by four measurable indicators -

- **Initiative success rates more than double,** from the industry norm of 30% reaching production to 70% through systematic orchestration that eliminates cross-functional coordination failures.
- **Time-to-value cuts in half** as the central office resolves blockers in weeks rather than months and squads leveraging reusable frameworks that reduce development cycle 40-60%
- **Benefit realization nearly doubles** as rigorous tracking ensures business cases deliver actual returns rather than theoretical projections.
- **Resource efficiency improves by 25%** as scarce AI talent deploys where it creates maximum value rather than sitting idle in organizational conflicts.

These operational improvements translate to **2.5x** more initiatives successfully deployed, **6x** higher cumulative benefits and break even much quicker than traditional models.

The ITO operating model is designed as an integrated framework but can be implemented modularly based on organizational maturity and existing capabilities. For example, Banks with mature PMO functions can enhance them via AI-driven portfolio management units and access specialized expertise for capability pools through partnerships with external consultancies.

The next section will illustrate a working example of an end-to-end ITO workflow that brings the above benefits in perspective.

Illustrative end-to-end ITO workflow

Let's look at an AI initiative's journey through the ITO lens right from opportunity identification to value realization and scaling. This end-to-end workflow demonstrates how the ITO operates in practice—providing the coordination, accountability, and execution capability that traditional approaches lack.

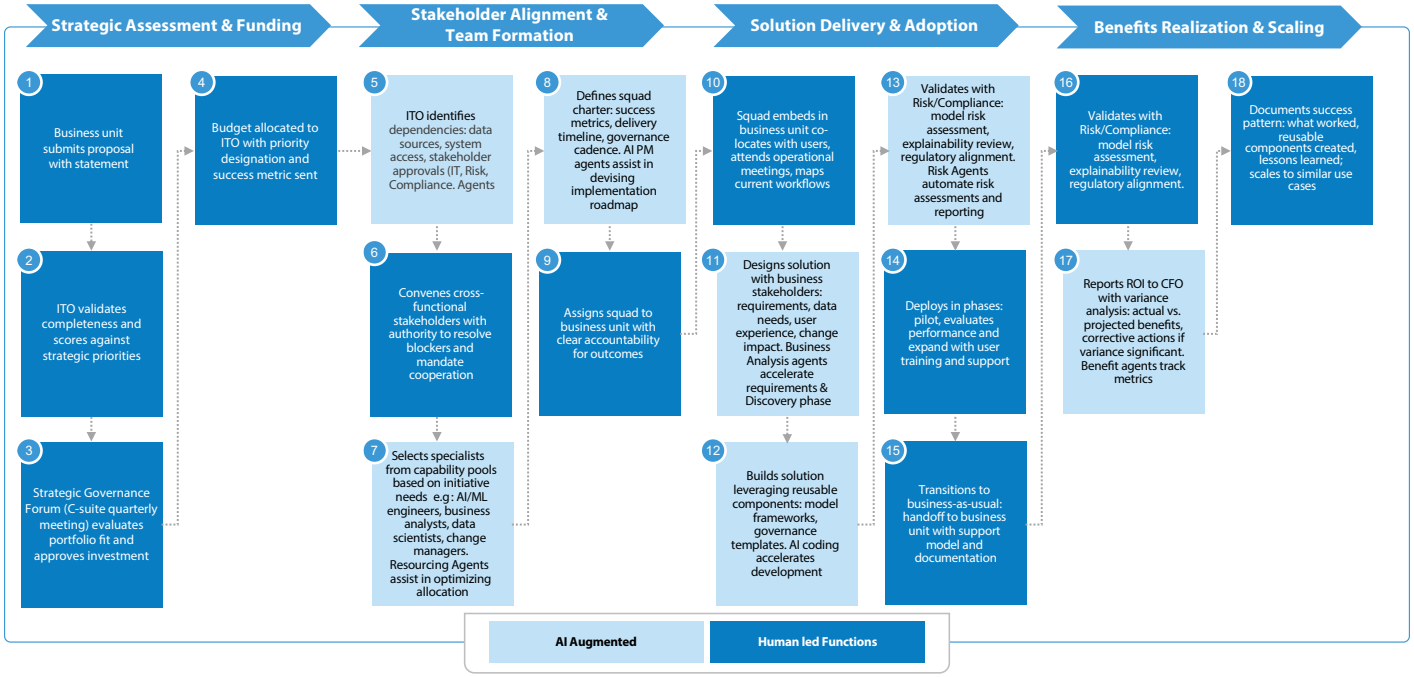

Use case: A bank's compliance function generates 15,000 monthly transaction monitoring alerts, but only 5% prove genuine after investigation. Compliance analysts spend 80% of their time on false positives, delaying detection of actual suspicious activity. Previous attempts to deploy AI stalled due to fragmented ownership: IT lacked compliance expertise, compliance lacked AI capability, and no function could coordinate the required data access across Operations, Retail Banking, and Risk.

1. Compliance Head submits AML automation proposal to Central O
2. Proposal scores high on regulatory risk and efficiency; expected $4M annual benefit vs. $2.3M investment with Agents scanning operational data to confirm false positive rates
3. Forum approves $2.3M investment, designates as priority given regulatory pressure
4. Budget released with mandate to deploy within 9 months
5. AML data in Operations, customer data in Retail Banking, validation required from Risk function
6. Portfolio orchestration convenes Operations, Retail,and Risk heads, resolves data access in 2 weeks vs. months of prior delays
7. Squad formed: 2 AI engineers, 1 compliance BA, 1 data scientist, 1 change manager. Portfolio Analytics agents identify optimal allocation based on specialist availability across portfolio
8. Monthly CFO reporting on actual benefits vs. $2.3M business case; weekly squad reviews
9. Squad reports to ITO Leader for delivery standards, to Compliance Head for domain priorities
10. Squad sits with compliance analysts, observes investigations, understands regulatory requirements
11. Compliance BA ensures model aligns with AUSTRAC expectations; analysts define usability requirements
12. Uses pre-approved model risk framework, standard data integration patterns; reduces development timeby 5 months with AI coding assistants accelerating development and compliance agents validating against
13. Risk function validates model explainability and bias testing before deployment approval
14. Pilot with 2,000 alerts/month, validates 70% reduction, expands to full volume over 3 months
15. Compliance unit assumes daily operations; squad maintains support for 3 months post-deployment
16. Alerts reduced from 15K to 3.8K monthly (73%); 8,500 analyst hours saved annually; investigation time:14 days → 6 days. Benefits tracking agents monitor these metrics continuously
17. $4.1M annual benefit realized vs. $2.3M investment= 1.8x first-year ROI; exceeds business case projection with benefits tracking agents generating automated reports
18. Compliance BA approach, change management positioning, reusable frameworks applied to other compliance use cases—fraud detection, sanctions screening

Implementation Roadmap

Moving from concept to operational reality requires a phased approach that builds credibility through early wins while establishing the foundations for portfolio-scale transformation. The roadmap below outlines high level implementation journey to establish and run an ITO

	Foundation	Capability building	Scaling
Key activities	• Establish Strategic Governance Forum(CEO, CFO, CIO, BU heads) with quarterly cadence and investment approval authority • Appoint senior ITO leader with cross-functional convening power and resource reallocation rights • Staff central office with 8-10 specialists (AI consultants, program managers, performance analysts, governance leads) • Define capability requirements for specialist resource pools • Assess and prioritize initial portfolio pipeline of candidate AI initiatives for pilot selection	• Build capability pool of specialists expert at leveraging AI tools in workflows (BAs, data scientists, change managers) • Configure enterprise AI Agents incrementally for ITO office using enterprise approved tools • Establish governance model for ITO AI Agents (validation protocols, monitoring standards, explainability requirements, human-in-the-loop rules) • Build reusable asset library with proven technical patterns, governance templates, and change management playbooks • Select 3-5 high-impact pilot initiatives with strong sponsorship, clear business value and feasibility of quick implementation	• Deploy pilot initiatives to production with embedded squads in business units, and track user adoption • Implement CFO-facing benefit realization dashboard tracking portfolio ROI, actual vs. projected KPIs, and variance analysis • Establish quarterly portfolio performance reporting to Strategic Governance Forum showing initiative outcomes andcompliance status • Build institutional knowledge base capturing lessons learned, updated reusable components, and capability development playbooks

Foundation

Secure executive governance commitment: Transformation operating models fail without genuine C-suite ownership. Establish a Strategic Governance Forum comprising the CEO, CFO, CIO, and business unit heads that meets quarterly with clear investment discipline—approving initiatives based on strategic value, reviewing actual ROI against projections, and making explicit portfolio trade-offs rather than funding everything business units request. This forum resolves cross-functional conflicts and reallocate resources basis performance. Without this governance commitment secured first, the subsequent steps lack the authority required to succeed.

Establish the ITO with genuine authority: With governance backing secured, appoint a senior leader accountable for portfolio outcomes not just coordination. This leader requires authority to convene stakeholders across organizational boundaries, mandate cooperation when alignment fails, and reallocate resources from underperforming initiatives to higher-value opportunities.

The ITO should be start lean (8-10 people). Typically, it should be a cross functional team of AI consultants who assess business cases, program managers who orchestrate portfolio execution, performance analysts who track benefit realization, and governance leads who ensure compliance.

Capability Building

Establish AI-augmented capability pools: The ITO's effectiveness depends on having specialist's expert at leveraging AI assistants in their workflows. Build pools of business analysts who use AI for requirements analysis and use case discovery, data scientists employing AI coding assistants for accelerated model development and change managers generating personalized communications at scale through AI tools. These specialists can be dynamically allocated across portfolio initiatives, enabling the small ITO office to operate without proportional headcount growth.

Build AI-augmented capabilities progressively: Leverage AI agents using enterprise-approved tools. Most organizations have Microsoft Copilot/Studio, PowerApps, ServiceNow, or similar platforms deployed. Microsoft Copilot Studio, for example, can be used to build lightweight agents that can automate cross-system data reconciliation across platform enhanced with bot capabilities for routine updates and alerts. As confidence builds, Copilot APIs can be leveraged to create custom agents for compliance validation or requirements analysis. The approach mirrors what leading transformation offices are already doing: augmenting their AI transformation office with Agents to bring in efficiencies. All ITO agents must operate under the same risk management framework that applies to any enterprise AI solution.

Scaling

Start with 3-5 high-impact initiatives that demonstrate measurable value within 6-9 months. Use these to validate the operating model, build reusable components (proven governance frameworks, technical patterns, change management approaches), and establish credibility before expanding to full portfolio scale.

Select initiatives with strong executive sponsorship, clear business value, and manageable complexity—early wins create momentum that attracts additional investment and organizational support. Attempting portfolio transformation simultaneously before proving the model creates chaos and undermines confidence. These initial initiatives also surface gaps in capability, tooling, or process that can be addressed before scaling broadly.

Conclusion

AI transformation is no longer optional: competitive pressure ensures every major bank will pursue it. The differentiator is execution capability. Banks with integrated operating models deploy AI solutions faster, scale successful patterns more effectively, and build institutional capability that persists beyond individual projects. Those maintaining fragmented approaches will invest substantially while watching competitors establish advantages that prove difficult to reverse.

At the beginning of the article, we mentioned a retail bank that was stuck in pilots and failing to scale AI use cases.

Once they established a central Intelligent Transformation Office, with authority, governance, embedded delivery capability, and accountability for value, things finally changed. What had been a single, isolated pilot became a portfolio of high-impact initiatives. The difference they realized was not technology but a better way of working.

Authors

Maurizio Cuna
Partner, Financial Services and Insurance, APAC
Maurizio is a recognised Infosys Consulting senior leader with two decades of experience in the financial services industry, across Europe, USA, Asia and Australia. Based in Sydney and with Infosys Consulting since 2016, Maurizio has built a strong team of cross-domain experts, working together to drive banking and insurance transformations for major Infosys clients.

Raju Gholley
Principal Consultant, Financial Services and Insurance, APAC

Raju has over a decade of experience advising financial services leaders/CXOs globally on transformation strategies and implementations. Based in Sydney, he specializes in large scale digital transformation programs, AI strategy & operating model , AI Product Mgmt. and Data transformation initiatives. Currently driving Gen AI & Agentic AI initiatives within the firm's financial services practices, his work focusses on translating AI capability into sustainable business outcomes.

ABOUT INFOSYS CONSULTING

Infosys Consulting is a next-generation consulting partner that bridges strategy and execution. With an AI-first mindset, deep industry knowledge, and the combined strengths of business and technology consulting, it helps enterprises turn bold vision into tangible outcomes, faster, smarter, and at scale. Infosys Consulting is helping some of the world's most recognizable brands transform and innovate. Our consultants are industry experts that lead complex change agendas driven by disruptive technology. With offices in 20 countries and backed by the power of the global Infosys brand, our teams help the C-suite navigate today's digital landscape to win market share and create shareholder value for lasting competitive advantage.

For more information, contact consulting@infosys.com

Infosys® | CONSULTING