The AI Native Financial Institution:
Redesigning Operating Models for an Agent-Led World

Introduction

AI adoption in financial services is accelerating, but operating models have not always kept pace. While early deployments focused on copilots that assist employees, emerging agentic workflows can plan, coordinate and execute work end-to-end.

When AI executes decisions that affect customers, financial outcomes and regulatory obligations, traditional assumptions about ownership, controls and oversight no longer hold. This paper sets out how BFSI institutions are redesigning their operating models to support that shift, moving deliberately toward autonomy while preserving transparency, governance and accountability.

How to use this paper

Use this paper as a working guide to align leadership, select priority journeys and design governance that enables safe scale. Use the *Core Blueprint* to define autonomy tiers, decision rights, controls and evidence requirements for your first 2–3 journeys.

Apply *Appendix A* to pick candidate journeys and translate them into ownership, evidence and escalation requirements your teams can implement. Use *Appendix B* to align stakeholders on the operating-model shifts and capability gaps that typically emerge as you move from copilots to agentic workflows. Use *Appendix C* to align terminology and assumptions with your internal governance standards before making design decisions.

Key Takeaways

Treat AI as an operating model redesign:
Clarify end-to-end ownership, decision rights and evidence capture, not just tooling.

Scale autonomy through explicit tiers:
(assist → recommend → act), with controls-in-the-loop and stop-the-line rights.

Anchor governance in established supervisory expectations:
For resilience and model risk management and apply AI risk management practices across the lifecycle.

Start with 2–3 high-impact journeys and a 90-day plan:
Prove control quality and outcomes before expanding scope.

Why Operating Models Break When AI Moves from Assistance to Execution

When AI starts executing work, operating models—not just tools—must change. Copilots primarily improve productivity inside existing processes. Agentic workflows, by contrast, introduce a new reality: software can take actions that affect customers, risk and financial outcomes. That shift forces institutions to clarify end-to-end ownership, embed controls and evidence directly into workflows, and adopt an operating rhythm that treats monitoring, drift, incidents and remediation as business as usual.

- **Accountability changes:** journey ownership must be explicit, who owns outcomes, exceptions, and customer recovery end-to-end.
- **Controls change:** Audit-ready evidence must be produced by the workflow (inputs, policy checks, actions, overrides, versions), not reconstructed after the fact.
- **Operations change:** Telemetry, thresholds, and “stop-the-line” rights become core management practices as autonomy scales.

The AI Native Operating Model (Core Blueprint)

The maturity path: Copilots → agentic workflows → selective autonomy

Stage 1 - Copilots (assist humans):
Accelerate drafting, search, analysis, and coding while establishing baseline controls (data access, logging, acceptable-use).

Stage 2 - Agentic workflows (execute with guardrails):
Redesign priority journeys so agents can orchestrate end-to-end work, with explicit escalation thresholds and measurable KPIs.

Stage 3 - Selective autonomy (run continuously in defined domains):
Scale governed autonomy where it is appropriate, using continuous monitoring, incident response, and audit-ready evidence as standard operations.

Outlook (2026–2031)

Over the next several years, many institutions will move from isolated proofs of concept to repeatable, governed agentic workflows in priority journeys. The differentiator will be execution: redesigning end-to-end ownership, embedding controls and evidence capture into the workflow and operating with a discipline that supports scale.

Figure 1. Practical maturity path from copilots to agentic workflows to an autonomous operating model.

The blueprint: four operating-model building blocks

- **Journey ownership and measurable outcomes:** Clear accountability for customer, cost and risk results.
- **Autonomy tiers and decision rights:** Assist → recommend → act, with explicit approvals and escalation.
- **Controls-in-the-loop and evidence by design:** Auditability built into the workflow, not bolted on.
- **Agent-ready foundation:** Trusted data, secure tool access, and end-to-end observability.

Autonomous Operations Become the Norm

BFSI institutions are moving beyond basic automation to *fully autonomous processes.* AI agents now handle many core workflows end-to-end - from loan underwriting to trade settlement - executing multi-step tasks without waiting for human prompts. This goes beyond efficiency. It’s a structural change in how work gets done.

AI-First Organization & Talent

Leading banks are redesigning their **organizational structures** around cross-functional AI-centric teams. New roles (e.g. AI model governance lead, AI ethics officer) are emerging, and *human talent is being upskilled to work alongside AI.* Routine clerical roles are giving way to higher-skilled analytical and oversight positions.

Hyper-Personalized, Invisible Customer Experiences

Customer engagement is shifting to **agent-driven, hyper-personalized** experiences. AI-driven services anticipate needs and act proactively - often without a traditional interface - delivering “invisible banking” where smart agents initiate actions for customers (e.g. automatic savings moves or loan offers).

Continuous Compliance & AI Governance

Regulatory compliance and risk management are evolving into **real-time, AI-monitored functions**. AI systems monitor transactions 24/7, flag anomalies instantly, and adapt to regulatory changes on the fly. Financial regulators themselves are deploying “SupTech” for active oversight, forcing institutions to build explainable, auditable AI into their operating DNA.

Tech Services Reinvented AI Era

Technology service providers (e.g. consulting and IT firms) must **transform their go-to-market**. In the AI-native era, BFSI clients demand strategic AI partnership, ready-to-deploy AI solutions, and co-innovation. Services firms are incorporating AI into their own delivery (e.g. AI software engineers alongside human teams) and shifting offerings toward AI-powered platforms, data governance, and talent enablement to stay relevant.

Organization Design for Governed Autonomy

The minimum viable org model: platform teams per priority journey

As AI takes on more routine decisions, leading institutions must reorganize journey-level ownership. Cross-functional platform teams (product/process, engineering, data/AI, operations and controls) own outcomes end-to-end—cycle time, quality, customer experience and control performance.

In practice, this shifts human capacity toward higher-value work: complex judgment calls, exception handling, customer recovery, policy tuning and oversight of automated decisions. The institutions that scale fastest treat these roles—and the training and handoffs that support them—as part of the product design. In its 2025 Future of Jobs Report, the World Economic Forum highlights the scale of expected role change and the reskilling required as AI and automation reshape work.

Who Does What: Delivery, Guardrails, and Assurance

Governed autonomy works when lines of responsibility are explicit and embedded in delivery, not managed through ad-hoc steering. A simple, scalable model looks like this:

1st line (journey/platform teams): Build and run agentic workflows; own KPIs, exceptions, and service recovery; produce evidence-by-design.

2nd line (risk/compliance): Define guardrails, approval gates, and control requirements by autonomy tier; pressure-test exceptions; monitor adherence.

3rd line (internal audit/assurance): Validate that evidence is complete and reproducible; test controls and governance; provide independent challenge.

Where to start: Prioritize 2–3 journeys where value is high and decision boundaries can be made explicit. Look for workflows with clear decision points, measurable outcomes and controllable exception paths. Common BFSI starting points, and the operating-model implications for decision rights, evidence and ownership, are provided in *Appendix A.*

Governance and Controls: How to Scale Without Losing Trust

In regulated environments, autonomy creates a governance design problem. Institutions scale more confidently when they build controls into the agentic workflow, so they can explain decisions, attribute actions and manage exceptions consistently. In its 2021 Principles for operational resilience, the Basel Committee on Banking Supervision emphasizes disciplined operational risk management, dependency mapping and incident response capabilities that support sustained service delivery under stress.

The goal is not “no humans,” but the right human involvement at the right risk level, with clear accountability and unambiguous evidence that stands up to internal audit and supervisory scrutiny. In the United States, banks commonly anchor model governance and validation expectations in the Federal Reserve’s SR 11-7 guidance on model risk management. In the United Kingdom, the Prudential Regulation Authority’s SS1/23 sets out model risk management principles, including governance, validation and expectations relevant to models that incorporate AI and machine learning techniques.

Define autonomy tiers by decision criticality:
(e.g., Assist → recommend → act with approval → act with post-review) and link each tier to required controls and signoffs.

Set decision rights and escalation rules:
Confidence thresholds, “stop-the-line” triggers, override authority, and service recovery steps for customer impact.

Build auditability (decision evidence) by default:
Log inputs, prompts, intermediate steps where appropriate, outputs, actions taken, human overrides, and versioning of models/policies.

Operationalize model risk management across the lifecycle:
(Testing, validation, drift monitoring, incident management, rollback) with clear ownership across 1st/2nd/3rd line. Align the operating cadence to established supervisory expectations, such as the Federal Reserve’s SR 11-7 guidance and the Prudential Regulation Authority’s SS1/23.

Measure control quality alongside value:
False positives/negatives, override rate, exception volume, time-to-remediate, and control breaches should be first-class metrics

The Agent-ready Foundation

Agentic workflows are only as dependable as the systems they can call upon and the data they can trust. To move from prototypes to production, institutions need an **agent-ready foundation;** governed data products, secure tool access and end-to-end observability so agents can act safely in real business systems. In its 2023 Artificial Intelligence Risk Management Framework (AI RMF 1.0), the National Institute of Standards and Technology emphasizes governance, mapping context, measuring risk and managing risk throughout the AI lifecycle—an approach that aligns well to regulated delivery.

- **Data products and semantics:** Standardized definitions, lineage, quality SLAs, and controlled access to features used in underwriting, servicing, and compliance decisions.
- **Identity, consent, and entitlements:** Fine-grained permissions for what agents can see and do, separation of duties for high-risk actions, and strong customer consent controls.
- **Tool orchestration:** Safe patterns for agents to call internal services (APIs, workflows, case systems) with guardrails (policy checks, rate limits, approvals).
- **Observability and telemetry:** End-to-end tracing across prompts, tool calls, decisions, and downstream system updates for rapid detection and remediation.
- **LLMOps/MLOps for agents:** Evaluation suites, red-teaming, version control, drift monitoring, and rollback aligned to regulated change management.

Next 90 Days: A Practical Path to Production

Days 0–30:
Pick journeys, agree on autonomy-tier targets, map controls/evidence requirements, and establish KPI baselines.

Days 31–60:
Build the agentic workflow MVP, implement logging/monitoring, run validation checks, and pressure-test exception handling with risk/compliance.

Days 61–90:
Deploy into a controlled production slice, measure lift vs. baseline, remediate control gaps, then decide to scale, pause or redesign.

Conclusion: The Executive Call to Action

AI native BFSI institutions will differentiate by how quickly they can redesign priority journeys into governed agentic workflows—while maintaining auditability, resilience, and clear accountability.

The practical path forward is to start small: select a few high-impact journeys, define autonomy tiers and decision rights, embed controls and evidence capture into the workflow, and instrument outcomes and guardrails so leaders can make scale/stop decisions quickly. Over time, institutions that industrialize repeatable patterns—platform teams, operating-speed governance and an agent-ready foundation—will be best positioned to expand autonomy safely and sustainably.

If additional support is helpful, Synechron can assist with journey selection, control-by-design requirements and implementation support so teams can validate outcomes and governance before scaling.

Appendices

Appendix A — Priority Journeys and Operating-Model Implications

Priority Process	Operating-Model Implication
Underwriting & credit decisioning	Clarify decision rights (auto / approve / escalate), standardize evidence packs (data, rationale, policy checks), and assign end-to-end ownership for cycle time and credit quality.
Customer service & engagement	Redesign the service journey with clear thresholds (complaints, vulnerable customers, high-value cases), and make conversation/action logging and case notes audit-ready by default.
Claims processing & fraud detection	Define routing and approval gates by severity, ensure investigation packets are explainable (signals, sources, rationale), and assign owners for leakage, customer recovery, and control quality.
KYC/AML & financial crime operations	Standardize case packs (alerts, typology, rationale, model/version traceability), codify escalation and filing decision rights, and align 1st/2nd line workflows to the same operating rhythm.
Payments, disputes & operational exceptions	Define which actions an agent may initiate (holds, recalls, returns), set approval thresholds, and run end-to-end traceability so operations can manage risk, SLAs, and customer recovery.
Collections & delinquency management	Embed conduct and fairness controls (e.g., UDAAP/fair lending), define escalation for hardship/vulnerability, and shift human capacity to negotiations, complaints, and regulatory-sensitive accounts. In its UDAAP examination procedures, the Consumer Financial Protection Bureau outlines supervisory expectations for identifying and managing unfair, deceptive, or abusive acts or practices.
Risk, finance & regulatory reporting	Assign accountable owners for data lineage and reconciliations, standardize controls and sign-off, and ensure disclosures remain reproducible (inputs, transformations, versions) for audit and supervisors.

Appendix B — Operating Model Shifts

The table below summarizes the most important operating-model shifts driven by AI and agentic workflows, along with practical implications for BFSI institutions and their technology partners.

AI-driven shift	Capabilities required
Autonomous processes and agents in core workflows	• Implement AI-powered automation for priority BFSI journeys (e.g., underwriting and claims) with production-grade controls. • Measure and manage outcomes end-to-end (cycle time, accuracy, cost-to-serve) to support outcome-based delivery expectations. • Run AI-enabled delivery workflows (engineering, testing, deployment) to match compressed operating timelines without sacrificing quality.
Cross-functional platform teams and agile org design	• Support product- and platform-team ways of working (agile delivery, cross-functional collaboration, and shared ownership). • Execute organization and change interventions that align operating model, talent, and governance to AI-native delivery. • Field integrated teams that combine domain expertise, data/AI engineering, and controls.
AI-empowered workforce and new talent models	• Maintain an AI-skilled workforce (data engineering, model development, MLOps/LLMOps, security) with strong BFSI domain fluency. • Sustain continuous learning mechanisms (curricula, certifications, communities of practice) to keep skills current. • Enable workforce adoption and operating change (training, role redesign, new ways of working) so AI scales beyond pilots.
Hyper-personalization and embedded AI in customer experience	• Engineer personalization and decisioning (recommendation, next-best action) with governed data and measurable uplift. • Integrate BFSI services into third-party channels securely (APIs, identity, consent, monitoring). • Design agent-driven customer experiences across channels (conversation design, escalation paths, and service recovery).
Continuous risk monitoring and AI-enabled compliance	• Operationalize AI governance and compliance (validation, bias testing, audit evidence, and regulatory alignment) for regulated use cases. • Automate KYC/AML workflows with AI while managing false positives, evidence, and escalation. • Implement ethical and explainable AI frameworks as standard operating controls.
Rapid legacy modernization with AI and cloud	• Modernize cloud and data foundations to support AI workloads (scalability, resilience, and cost discipline). • Use AI-assisted engineering to accelerate legacy renovation and testing while maintaining control and traceability. • Modernize core workflows end-to-end (APIs, orchestration, eventing) to make systems agent-ready.
Outcome-focused, as-a-service operating philosophy	• Align delivery to business outcomes (customer experience, risk reduction, efficiency) with shared measurement and accountability. • Operate AI solutions continuously (monitor, improve, and govern) rather than treating deployment as the finish line. • Support modern commercial models (managed outcomes, gainsharing, transaction-based pricing) that match continuous value delivery.

These shifts are most manageable when treated as operating-model design decisions (ownership, decision rights, controls, and evidence) rather than isolated technology deployments.

Appendix C — References, Method, Definitions, Assumptions, and Glossary

Method, Definitions, and Assumptions

Method
This white paper synthesizes publicly available regulatory guidance and industry research with operating-model patterns observed across BFSI transformation programs. Representative sources include the Basel Committee on Banking Supervision’s 2021 *Principles for operational resilience;* the Federal Reserve’s SR 11-7 guidance on model risk management; the Prudential Regulation Authority’s SS1/23 model risk management principles; and the National Institute of Standards and Technology’s 2023 *AI Risk Management Framework (AI RMF 1.0).* The intent is to provide a practical, executive-level framework and language that BFSI institutions can adapt to their own strategy, risk appetite and delivery constraints.

Definitions
Terms such as *AI-native, agentic workflows, autonomous operating model,* and *auditability* are used in the specific sense defined in the Glossary of Key Terms. Where terms have different meanings across jurisdictions or institutions, readers should use their internal definitions and control standards for formal governance decisions.

Assumptions and limitations
The recommendations in this paper generally assume:

A regulated BFSI context where model risk management, operational resilience, privacy, and conduct requirements apply, and where humans remain accountable for material decisions.

Availability (or planned buildout) of baseline capabilities such as identity and access controls, data governance (lineage/quality), logging/telemetry, and change management suitable for production AI.

That agentic workflows are introduced incrementally (starting with narrow scopes and clear guardrails) and scaled only when outcomes improve and control evidence is sufficient for internal assurance.

Glossary of Key Terms

Key Term	Definition (in this paper)
AI-native	An operating model where AI is embedded into core decisions and workflows (not bolted on), with governance, data, and controls designed to support continuous automation at scale.
Autonomous operating model	The enterprise-wide end state where many governed agentic workflows and autonomous agents run core operations continuously, with defined decision rights, escalation paths, monitoring and audit evidence.
Autonomous agent	A software component that can plan and take actions toward a goal (e.g., gather data, execute workflow steps), within defined guardrails, controls, and escalation rules.
BFSI	Banking, Financial Services, and Insurance.
Copilot	An AI assistant that augments an employee’s work (e.g., drafting, summarizing, coding) but typically does not execute business actions autonomously.
Embedded finance	Delivery of financial products and services within non-bank channels (e.g., commerce platforms), increasingly enabled by real-time AI decisioning.
Explainability	The ability to communicate why an AI system produced a result in a way that supports human review, regulatory expectations, and appropriate customer transparency.
Human-in-the-loop (HITL)	A control model where humans review, approve, or override AI outputs/actions—especially for high-risk decisions, exceptions, or low-confidence cases.
MLOps / LLMOps	Operational practices to build, deploy, monitor, evaluate, and govern models in production (including drift, testing, rollback, and lifecycle management).
Model risk management (MRM)	Policies, controls, and oversight used to manage model risks (performance, bias, explainability, compliance) across the model lifecycle.
Outcome-based delivery	Operating and funding model where programs are measured by business results (e.g., cycle time, loss reduction, false-positive reduction) rather than activity or effort alone.
Platform team	A cross-functional team that owns a product or domain end-to-end (technology, data, controls, KPIs), enabling faster delivery and clearer accountability than siloed functions.
Straight-through processing (STP)	Automated processing from initiation to completion with minimal manual intervention; humans focus on exceptions and escalations.
SupTech	Supervisory technology used by regulators to analyze risk and compliance data more continuously and at higher scale.
UDAAP	Unfair, Deceptive, or Abusive Acts or Practices—U.S. consumer-protection standard (Dodd-Frank Act) that prohibits unfair, deceptive, or abusive conduct in connection with consumer financial products or services.

Sources

Bank of America. “Bank of America’s Erica surpasses 2 billion interactions, helping 42 million clients since launch.” Press release, April 8, 2024.

Basel Committee on Banking Supervision. Principles for operational resilience. Bank for International Settlements, March 31, 2021.

Basel Committee on Banking Supervision. Progress in adopting the Principles for effective risk data aggregation and risk reporting. Bank for International Settlements, Nov. 28, 2023.

Board of Governors of the Federal Reserve System. “SR 11-7: Guidance on model risk management.” Supervisory letter, April 4, 2011

Consumer Financial Protection Bureau. “Unfair, deceptive, or abusive acts or practices (UDAAPs) examination procedures.” Oct. 1, 2012; updated Sept. 25, 2023.

DBS Bank. “DBS’ AI-powered digital transformation.” Aug. 7, 2023.

Gandharv, K. “DBS Bank leverages data to raise operational efficiency and customer engagement.” iTnews Asia, Nov. 12, 2024.

European Banking Authority. Guidelines on ICT and security risk management (EBA/GL/2019/04). Nov. 29, 2019.

EY-Parthenon. Generative AI in retail and commercial banking. Report, 2025.

HCLTech. “Looking ahead to an autonomous, intelligent and hyper-personal future for financial services in 2026.” Company publication, Dec. 22, 2025.

International Organization of Securities Commissions. The use of artificial intelligence and machine learning by market intermediaries and asset managers. Final report FR06/2021, Sept. 2021.

McKinsey Global Institute. “Human skills will matter more than ever in the age of AI.” Jan. 6, 2026.

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). Jan. 26, 2023. DOI: 10.6028/NIST.AI.100-1.

Organisation for Economic Co-operation and Development. Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449). Adopted May 22, 2019.

Prudential Regulation Authority. SS1/23: Model risk management principles for banks. Supervisory statement, May 17, 2023.

World Economic Forum. The Future of Jobs Report 2025. Report, Jan. 7, 2025.

World Economic Forum. “Investing in people: The power of human capital in banking’s AI era.” Jan. 10, 2025.

Author

Frances Cruz is Head of Consulting, Americas at Synechron. She advises financial services leaders on operating-model transformation across digital, data and AI, connecting measurable business outcomes with practical governance, risk and control frameworks needed to scale safely.

Disclaimer

This white paper is provided for general informational purposes only and reflects the author’s views as of the publication date. It does not constitute legal, regulatory, tax, investment, or other professional advice, and it should not be relied upon as a substitute for consultation with qualified advisors. References to laws, regulations, supervisory expectations, and industry practices are generalized and may vary over time and by jurisdiction. Any examples, use cases, metrics, figures, or third-party references are illustrative; readers should validate assumptions and results using primary sources and their own data, controls, and risk appetite. The content is provided “as is,” without warranties of any kind, and Synechron disclaims liability for actions taken based on this material.

No client endorsement or affiliation is implied by any references to third-party organizations, institutions, products, or services.