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.

This paper examines how AI-native operating models can be designed to support scale, resilience and accountability in organizations about ownership control and oversight. It draws on recent industry research and practitioner perspectives on redesigning team operating models to support that shift, moving deployment 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 decide what to automate, what to augment and what must remain human-led, then apply the operating model that fits your risk profile.

Apply Appendix A to pick candidate journeys and translate their recommendations into governance and operation requirements you have to 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 pressure-test vendor claims against the governance and control questions that actually matter in regulated environments.

Key Takeaways

- **Treat AI as an operating model redesign:** Clarify ownership and decision rights before scaling agents across workflows, not just tooling.
- **Scale autonomy through explicit tiers:** Assist → recommend → execute-with-approval → execute-and-log are the right lens.
- **Anchor governance in established supervisory expectations:** Model risk, third-party risk, cyber and consumer protection frameworks already apply to agentic deployments.
- **Start with 2-3 high-impact journeys and a 90-day plan:** Prove control quality and production readiness before expanding scope.

Why Operating Models Break When AI Moves from Assistance to Execution

When AI starts executing tools, operating models rest just tasks: change stage. Explicit primarily improve productivity inside existing workflows. Agentic systems can plan, coordinate and execute multi-step work across systems, which changes how decisions are made, ownership is assigned, and how exceptions are handled. Control shifts from static approvals to runtime policy enforcement, escalation paths and continuous monitoring. Organizations that treat agentic AI as a tool rollout rather than an operating model redesign often stall at pilot stage or accumulate risk faster than they can govern it.

- **Accountability changes:** Journey ownership must be explicit, who owns outcomes, exceptions and customer recovery end to end.
- **Controls change:** Audit, explainability and monitoring must operate across tools, agents, vendors, retrained, not retrained after the fact.
- **Operations change:** Telemetry, thresholds and "stop the line" rights become core management routines.

The AI-Native Operating Model (Core Blueprint)

The maturity path: Copilots + agentic workflows + selective autonomy

Stage 1 - Copilots (assist humans)
Accelerates drafting, search, analysis and decision preparation inside existing tools and workflows; scaling impact depends on adoption and workflow redesign.

Stage 2 - Agentic workflows (execute with approvals)
Redesigns priority journeys so agents can execute bounded steps across systems with escalation thresholds and measurable KPIs.

Stage 3 - Selective autonomy (execute-and-log) for a short-duration
Scale governed autonomy where it is appropriate, using continuous monitoring, incident response and audit-ready evidence as standard practice.

Autonomous Operations Become the Norm
Routine exceptions are resolved by long-term agents in fully monitored, governed environments with clear escalation paths and human oversight where risk warrants it.

AI-First Organization & Talent
Leadership models AI collaboration, hiring and upskilling reflect agentic workflows, and teams operate with clear ownership of agents. AI-native organizations lead in speed, cost structure and customer experience, driving sector-level disruption as incumbents struggle to retrofit legacy processes onto agentic foundations.

Hyper-Personalized Invisible Customer Experiences
Agents anticipate customer needs across products and channels, delivering personalized, proactive service that feels seamless rather than transactional, with governance ensuring fairness and transparency at scale.

Continuous Compliance & AI Governance
Regulators and organizations deploy AI-native supervisory tools for real-time monitoring and audit-ready evidence. Governance becomes a competitive advantage, enabling faster, safer deployment than manual oversight ever could.

Tech Services Reinvented to Era
Technology vendors and platforms pivot to focus **exclusively on AI-native** solutions, rewriting product roadmaps and business models to compete in an agentic economy. Services firms are accelerated by this than ever before, e.g., if software vendors ship agentic capabilities, consulting firms must rebuild delivery models and governance frameworks to stay relevant.

Outlook (2026-2031)

Over the next several years, many institutions will move from isolated pilots to governed agentic workflows in priority journeys. The shift depends less on model capability than on ownership, escalation paths and measurable outcomes rather than on adoption metrics.

The Blueprint four operating-model building blocks

- **Journey ownership and measurable outcomes:** Each journey has a customer outcome and a KPI owner.
- **Autonomy tiers and decision rights:** Assist → recommend → execute-with-approval → execute-and-log.
- **Controls in the loop and evidence by design:** Policy checks, logs and audit-ready evidence are standard.
- **Agent-ready foundation:** Trusted data, identity, permissions and observability.

Figure 1: Practical maturity path from copilots to agentic workflows with selective autonomy.

Figure 2: The four building blocks of an AI-native operating model.

Organization Design for Governed Autonomy

The minimum viable scale model: platform teams per priority journey

As AI takes on more routine decision-making, organizations need lean governance that scales without creating bottlenecks. The minimum viable scale model uses platform teams per priority journey to provide shared infrastructure, standards and control performance.

In practice, this shifts human capacity toward higher-value work: complex exceptions, edge cases, relationship building, customer recovery, policy interpretation and continuous improvement. Organizations that treat AI as headcount reduction alone miss the deeper opportunity: using freed capacity to strengthen the control and judgment layer that makes autonomy safe to scale.

Sit the Governance Committee Bridges risk, compliance, technology and business to review exceptions, approve escalation thresholds and govern changes to agent behaviors.

Sit the Operational Audit/Assurance Validates that evidence is complete and reproducible; tests controls and generates audit-ready reports for regulators.

Where to start: Prioritize 2-3 journeys where value is high and decision boundaries can be made explicit. Look for workflows with clear escalation paths, measurable outcomes and existing governance frameworks. Avoid RFP-starting pilots and the operating-model implications for decision rights, escalation and accountability are avoided or deferred.

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, auditable model looks like:

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

2nd line (Governance/risk/compliance): Define policy and control requirements, monitor exceptions, approve escalation thresholds, govern changes and maintain audit-ready evidence.

3rd line (Internal audit/assurance): Validate that evidence is complete and reproducible; test controls and generate reports for regulators and boards.

The goal is not "no humans." The goal is that human involvement at the right risk level, with clear accountability and unambiguous evidence, is what makes autonomy safe to scale. Organizations that can't show who approved what, why an exception was allowed and how a control performs will struggle to scale regardless of model capability.

- NIST's SP800-53 sets out model-agnostic, auditable control expectations and evidence requirements for models that increasingly act and interface decision-making.
- **Define autonomy tiers by decision criticality:** Assist → recommend → execute-with-approval → execute-and-log, with escalation thresholds and rollback triggers.
- **Set decision rights and escalation rules:** Confidence thresholds, stop-the-line triggers and named owners should be customer-impact.
- **Build auditability (decision evidence) by default:** Logs inputs, prompts, interventions and outcomes to support reconstruction, sampling and audit readiness.
- **Operationalize model risk management across the lifecycle:** Inventory, validation, monitoring and change control for models and agents; GenAI/agentic risk adds new failure modes (prompt injection, tool misuse, hallucinations) that must be tested and monitored.
- **Measure control quality alongside value:** False positive/negative, override rate, exception volume, time to remediation and control breaches should be first-class metrics.

Governance and Controls: How to Scale Without Losing Trust

As organizations move from pilots to production, governance becomes the operating system that allows teams to ship safely. Leaders need to see the agentic workflow, so they can explain decisions, attribute actions and manage exceptions consistently. In its 2026 Principles for operational resilience, the Bank of England highlights governance, accountability and operational resilience as foundations for managing AI adoption and broader digital change.

The goal is "not no humans." The goal is that human involvement at the right risk level, with clear accountability and unambiguous evidence, is what makes autonomy safe to scale. Organizations that can't show who approved what, why an exception was allowed and how a control performs will struggle to scale regardless of model capability.

- NIST's SP800-53 sets out model-agnostic, auditable control expectations and evidence requirements for models that increasingly act and interface decision-making.
- **Identity, consent and permissions:** Agents must inherit access rules and permissions so they can't exceed what a human is allowed to do.
- **Test adversarially:** Red-team for prompt injection, tool misuse, PII leakage, and edge cases before production.
- **Observability and telemetry:** Log every tool call, retrieval and decision so teams can reconstruct what happened and why.
- **Control performance by design:** Embed policy checks, approvals and evidence capture into the workflow, not after it.

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, organizations need an **agent-ready foundation** (governed identity, permissioned data access, and observability) so agents can act safely in real business systems. In its 2026 Artificial Intelligence Risk Management Framework (AI RMF 1.0), the National Institute of Standards and Technology (NIST) frames governance, mapping, measurement and management as the foundation for managing AI risk across the lifecycle.

- **Data products and semantics:** Standardized definitions, lineage, quality SLAs and controlled access to features used in underlying, governed business systems.
- **Identity, consent and permissions:** Fine-grained access so agents inherit only what they are allowed to do, with logging and revocation controls.
- **Test and evaluation:** Pre-deployment testing for accuracy, robustness, bias, security (prompt injection, tool misuse) and failure modes.
- **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

Pick journeys, agree on autonomy tier, ship with guardrails, measure impact, iterate. Organizations that move from pilots to production typically follow a 90-day sprint that ties ownership, controls and measurement to a few high-value journeys.

Days 0-30:
Pick 2-3 journeys, define success metrics, map data and permissions, agree on autonomy tiers, draft control requirements, and assign owners.

Days 31-60:
Build agentic workflows with retrieval, tool calling, and human-in-the-loop approvals, integrate with identity and logging, and run pre-production evaluations and red-teaming.

Days 61-90:
Deploy into a controlled production slice, measure KPIs and incidents, refine controls and escalation thresholds, and prepare a scaled rollout plan with governance checkpoints.

Conclusion: The Executive Call to Action

A pilot is AI that proves it can work. Production is AI that earns the right to keep running. The gap between those two states is almost always operational, not technical. Organizations that treat governance as a tollgate slow everything down. Those that build it into the workflow, with clear ownership, escalation paths and evidence collected by design, can move faster and still sleep at night.

The practical path forward is to start with a few high-impact journeys, define autonomy tiers and decision rights, embed controls and evidence capture into the workflow, and measure control performance alongside business outcomes. The organizations that will lead in this decade are not the ones with the most AI pilots. They are the ones that can show a regulator exactly what an agent did last Tuesday, why it was allowed to do it, and who was accountable for the outcome.

If additional support is needed, Synechron can assist with process deep-dives, control by design requirements and implementation support as 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 and lines; separate mandatory standards by product/jurisdiction; pilots, retrained, policy checks and change control and auditability for agents (risk and audit).
Customer service & engagement	Redesign the service journey with clear responsibility (complaints), escalation, customer recovery, high-value cases and measurable outcomes; agents assist and execute bounded tasks.
Claims processing & fraud detection	Define routing and approval gates by severity; ensure investigator capacity and explainable signals; access, retrained, policy checks and change control and auditability for agents (risk and audit).
KYC/AML & financial crime operations	Automate evidence collection and monitoring; maintain explainability and audit readiness; centralize ownership and decision rights and align risk/compliance to the same operating rhythm.
Payments, disputes & operational exceptions	Define exception ownership and root-cause analysis; escalate, access, retrained, policy checks and change control and auditability for agents (risk and audit).
Collections & delinquency management	Define escalation and consent rules; ensure fair treatment and auditability; access, retrained, policy checks and change control and auditability for agents (risk and audit).
Risk, finance & regulatory reporting	Assign accountable owners for data lineage and reconciliations; standardize evidence and change control; ensure traceability from inputs to reported numbers.

Appendix B — Operating Model Shifts

The table below summarizes the most important operating-model shifts across AI-native and agentic workflows, along with practical implications for leadership teams and their technology partners.

AI-driven shift	Capabilities required
Autonomous processes and agents in core workflows	Integration of autonomous agents into existing BPM/case tools; underwriting and claims with production grade controls
Cross-functional platform teams and agile org design	Product-aligned teams with end-to-end ownership; fewer handoffs; faster iteration on customer journeys
Continuous risk monitoring and AI-native audit	Automated evidence capture and monitoring; bias, drift, hallucination and exception dashboards; audit-ready logs by design
Rapid legacy modernization with AI and cloud	Use AI to refactor legacy systems; cloud-native platforms for scalable deployment; reduce technical debt blocking AI adoption
Outcome-focused, pay-as-service operating partnerships	Vendors co-owning outcomes; outcome-based contracts; shared accountability for agent performance and compliance

These shifts are most manageable when treated as operating model design decisions (ownership, decision rights, controls and evidence), not as tool upgrades alone.

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

Method, Definitions and Assumptions

Method
This white paper synthesizes publicly available regulatory guidance, industry research with operating-model patterns observed across BFSI transformation programs. Regulatory references include the Bank of England's Principles for operational resilience (2026), NIST's AI Risk Management Framework (AI RMF 1.0), the European Banking Authority's ICT and security risk management guidelines, and the European Banking Authority's AI and machine learning principles. Industry sources include McKinsey's State of AI in Banking and Capital Markets (2026), the World Economic Forum's Financial Services AI Governance Framework (2026), and Deloitte's Tech Trends 2026. The blueprint and maturity model are practitioner frameworks designed to translate these into implementable operating-model decisions.

Definitions
Terms such as AI native, agentic workflows, autonomous execution, governance and auditability are used in the specific senses defined in the Glossary of Key Terms. Where terms have different interpretations in regulatory or analytical contexts, readers should use their internal definitions and map them to those used in this paper before applying the framework.

Assumptions and Limitations
The recommendations assume a regulated financial services environment.

- A regulated BFSI context where model risk management, operational resilience, privacy and conduct requirements apply and where human oversight is required for material decisions.
- Availability of data and identity infrastructure sufficient to support governed agent deployment, including permissioned data access, logging and observability.
- That agentic workflows are monitored and incrementally governed, with escalation paths and rollback capabilities and evidence collected by design, rather than deployed as fully autonomous systems without controls.

Glossary of Key Terms

Key Term	Definition (in this paper)
AI-native	An operating model where AI is embedded into decision and workflow design by default, with governance, evidence and accountability built in from the start, rather retrofitted onto existing processes.
Autonomous operating model	The enterprise operating model that combines human-led governance, certified and autonomous agents and supervisory controls, with defined decision rights, escalation paths, monitoring and audit evidence.
Autonomous agent	A software component that can plan and take actions toward a goal by using tools and context, within defined boundaries and escalation rules.
BPM	Business Process Service and Insurance.
Copilot	An AI assistant that augments an employee's work (e.g., drafting, summarizing, coding) but typically does not execute actions independently.
Embedded finance	Delivery of financial products and services within non-bank channels (e.g., commerce platforms, messaging platforms, mobility apps).
Explainability	The ability to communicate why an AI system produced a result in a way that a human can review, challenge, and document for governance and audit purposes.
Human-in-the-loop (HITL)	A control model where humans review, approve or override AI recommendations, especially for high-risk decisions, exceptions and boundary conditions.
MLOps / LLMOps	Operational practices for deploying, monitoring, evaluating and governing models in production (including drift, quality, cost, latency and incident response).
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 pricing model where progress is measured by business results (e.g., cycle time, loss reduction, fraud avoided reduction) rather than activity or effort alone.
Platform team	A cross-functional team that owns a product or domain service and develops/runs the tools, data, controls, APIs and observability that other teams use to ship safely at scale.
Straight-through processing (STP)	Automated processing from intake to completion with minimal manual intervention, humans focus on exceptions and oversight.
SupTech	Technology used by supervisors and regulators to analyze risk and compliance data across firms and markets at scale.
UiPath®	UiPath, Description or describe AI or Platform: n/a), consumer protection standard (Deals Hub & Cost Recoveries UiPath), detection or remediation (Deals Hub & Cost Recoveries UiPath), and governance and audit (Platform).

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, March 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. 30, 2023.

DBS Bank, "DBS AI powered digital transformation," Aug. 7, 2025.

Cgnalysis, n. "DBS Bank leverages data to raise operational efficiency and customer engagement," Press Rels, Nov. 12, 2016.

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, August, 2025.

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

International Organization of Securities Commissions. The use of artificial intelligence and machine learning by market intermediaries and asset managers. Final report FINR6/2021, Sept. 2021.

McKinsey Global Institute, "Human skills will matter more than ever in the age of AI," Nov. 5, 2026.

National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), (NIST AI 100-1), Jan. 26, 2023, updated March 30, 2023.

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

Prudential Regulation Authority, SS23: 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 talent in banking 5 AI era," Jan. 10, 2025.

Author

Prasanna Lohar is Head of Consulting, Americas at Synechron. He advises financial services leaders on operating model transformation across digital, data and AI, governance and control design, and delivery execution. He has worked across banking and capital markets, insurance and wealth.

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, financial or investment advice and should not be relied upon as such. Readers should consult with their own legal, regulatory, or other qualified advisers. References to laws, regulations, supervisory expectations and industry practices are generalized and may vary over time and across jurisdictions. Implementation of the concepts described in this paper does not guarantee regulatory compliance, supervisory approval, or avoidance of enforcement actions. Synechron does not provide legal or regulatory advice. Any examples, case studies or illustrations are for educational purposes only and are not intended to be binding or exhaustive.

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