SalesChoice

Advancing Responsible AI in Global Financial Services

A leadership guide to turning the world's AI frameworks, from OSFI and the EU AI Act to Singapore, India, China, Taiwan, and the Philippines, into one working AI operating model, from intake to decommissioning, for banks, insurers, pension funds, and any government-regulated organization.

Contents

Executive Summary 3

The Strategic Imperative for Financial Services 6

The Global Regulatory Patchwork, in Brief 10

The SalesChoice AI Lifecycle: Intake to Decommissioning 19

Closing the Guardrails Gap: The Systems Behind Governance 27

AI Value Realization: Metrics That Matter 29

What the Frameworks Don't Tell You: Five Blind Spots 31

Case Studies: Responsible AI in Practice 33

SalesChoice: Our Track Record in Responsible AI 38

SalesChoice AI Advisory Services: How We Help 40

A 90-Day Roadmap to Get Started 42

Conclusion: Why This Matters, and Why It Must Be Led from the Top 43

About the Author 46

References 47

Executive Summary

Financial institutions around the world are making the fastest AI investments in their history at the moment however the rules for making them responsibly are diverging. Canada's OSFI has finalized a binding model risk guideline, E-23, that brings AI and machine learning into scope from May 2027. The United States rescinded SR 11-7 in April 2026 and replaced it with non-binding principles that exclude generative and agentic AI. Ontario's FSRA expects tested vendor resilience without yet naming AI accountabilities, and the EU AI Act comes closest to prescribing an operating model, though its high-risk obligations begin only in December 2027. Across Asia, five jurisdictions have each chosen a different mix of binding rules and supervisory principles. But none of these, nor NIST and OWASP's rigorous technical standards, guides financial service organizations on how to run an AI program that drives value outcomes.

The binding constraint is therefore neither regulation nor technology. It is the operating model. Institutions are accumulating policy faster than they are building the capability to act on it; and a policy in a compliance repository is not an intake process, an inventory an examiner can query, or a monitoring dashboard that catches drift early. RAND's 2024 research found that, by some estimates, more than 80 percent of AI projects fail, twice the rate of IT projects generally. In our extensive SalesChoice AI operating experience, most of that failure traces not to weak models but to the absence of a governed, repeatable path from first idea to retirement and to the absence of anyone senior enough insisting on a structured AI operating model.

> ***A poorly managed AI model, especially one used in critical decision-making, could even contribute to an institution's failure.***

— Peter Routledge, Superintendent, OSFI, October 30, 2025

October 30, 2025

This research paper offers four things an executive team can use directly. The operating model is built for banks, insurers, pension funds, and any government-regulated organization; only the regulatory anchors change.

- **A complete lifecycle, intake to decommissioning (Section 3).** Eight stages with inputs, outputs, an accountable owner, and a framework anchor, so every C-level executive works from one operating model.
- **The systems that turn policy into evidence (Section 4).** Seven categories of technical infrastructure that make governance auditable.
- **A point of view on AI value realization (Section 5).** A four-category value architecture, baselined before funding and re-tested after launch (see the measurement chapter of The AI Precipice).
- **Five blind spots no framework closes (Section 6).** From "zombie models" that no one retires to agentic workflows no one owns: gaps that someone must close, which is an executive decision.

Three anonymized SalesChoice experiences (Section 7) ground the argument, including a major UK bank whose Tier 1 C-suite took a week-long immersion before committing to strategy.

Our central message is for the executive team, not the compliance function: Responsible AI is a leadership discipline. Regulators in many countries have left the "how" to institutional leaders, so the institutions that lead will be those whose CEOs and boards own it, starting with the ninety-day plan in Section 10.

The Strategic Imperative for Financial Services

KEY TAKEAWAY

Five dynamics are accelerating AI adoption in financial services, shown here through North American data. Only one, the supervisory posture, tells you who is accountable when it goes wrong, and in the jurisdictions we reviewed the answer is consistent: you are.

Five dynamics are converging on banks, credit unions, and insurers at once. We document them with North American data, where the evidence is most complete; the same forces are visible in Europe and Asia.

> We are completely convinced the consequences will be extraordinary and possibly as transformational as some of the major technological inventions of the past several hundred years. Think the printing press, the steam engine, electricity, computing and the Internet, among others.
>
> — Jamie Dimon, Chairman and CEO, JPMorgan Chase, 2023 Annual Letter to Shareholders, April 2024

EXHIBIT A North American FSI Macro Pressures at a Glance

$5.00	44%	0.80→0.85%	~6.5%
Total cost per $1 lost to fraud, North America (up 25% in four years)	FSIs still running mostly or entirely manual fraud defenses	US credit union ROA forecast, 2026 → 2027	Forecast US credit union savings growth, 2026

Sources: LexisNexis Risk Solutions (2025); America's Credit Unions 2026 Economic Outlook.

1. Margin Recovery

Margins are recovering, but the memory of compression hasn't faded.

After several years of pressure, net interest margins across the US credit union system are improving and return on assets is forecast to rise from roughly 0.8% in 2026 to 0.85% in 2027 as investment portfolios reprice upward. But that recovery is arriving alongside savings growth of around 6.5%, as members move cash out of market investments and rebuild deposits, which means institutions are competing harder for the same funding even as they try to hold the margin gains. Non-interest income diversification (wealth management, insurance, small-business services) and branch networks repositioned as advisory centers are now standard strategic responses. Every one of those responses depends on data and AI to work at scale.

Implication: AI business cases will be judged against a thin, fragile margin recovery. Every dollar of AI spend needs a baseline before funding and a re-test after launch (Section 5).

2. Fraud Acceleration

Fraud is outrunning manual defenses.

Every dollar lost to fraud now costs North American financial institutions an average of $5.00 in total remediation cost, up 25% from $4.00 four years ago and mobile channels account for more than a third of total fraud losses. Yet 44% of North American financial institutions still rely mostly or entirely on manual fraud processes, while only 20% operate mostly or fully automated defenses. Fraud is already AI-accelerated; synthetic identities, deepfake-enabled social engineering, bot-driven account testing and the institutions still relying on manual review are absorbing the full multiplier effect.

Implication: Fraud detection is the clearest near-term place where AI value can be counted in dollars, and where an ungoverned model does the most damage.

3. Credit Risk Re-Rating

Trade and rate uncertainty is re-rating credit risk.

Tariff and trade-policy uncertainty through 2026 has pushed commercial and supply-chain-linked credit risk back onto underwriting teams' agendas, with the Federal Reserve's own research noting the link between trade-policy uncertainty and bank lending behavior. Combined with a Federal Reserve leadership transition and only modest rate-cut expectations for 2026 and 2027, institutions are underwriting into more macro uncertainty than they have faced in several years, exactly the environment in which AI-assisted risk models earn their keep, and in which a poorly governed one does the most damage.

Implication: Models that inform credit decisions are the ones examiners will ask about first. They need the full lifecycle (Section 3), not the lighter oversight a pilot usually gets.

4. Competition & Talent

The competitive and talent picture is forcing the issue.

Fintechs and megabanks continue to compress the window in which community institutions can differentiate on price alone, pushing member and customer experience to the top of the strategic agenda. At the same time, AI is already visibly reshaping labor markets and hiring patterns across financial services, adding workforce-transition questions to what was, a few years ago, purely a technology conversation.

Implication: Speed to production, not speed to pilot, is the competitive variable, and the workforce impact has to be designed into each use case at the value-case stage, not discovered after launch.

5. Supervisory Posture

Regulators expect rapid adoption, and expect leadership to own the consequences.

In Canada, the supervisor has been unusually explicit. Superintendent Peter Routledge said in October 2025 that AI "will change how financial services operate, and the pace of adoption will be rapid," and that OSFI "will allow institutions the space to innovate, while remaining vigilant in identifying and addressing material risks." A month later he told a Thomson Reuters conference that OSFI does not want to interfere in how institutions integrate AI, accepting the risk that it could be applied in ways that create unintended risk concentrations. In the United States, the rescission of SR 11-7 produced the same practical result by a different route: a thinner binding floor and more discretion for institutional leadership. Both supervisors are saying the same thing; how an institution governs AI is a management decision, and management will be held to it.

Implication: The "how" of AI governance has been left to institutional leaders. The operating model is theirs to build and theirs to defend. However, working with experienced and responsible AI leaders is crucial to build an AI operating model platform that achieves desired outcomes and connects the dots to accountability roles. Major change management enablements are key to ensure AI not only scales, but also delivers the intended value projected.

None of these pressures is new in isolation. What is new is that they arrive together, as the guardrails institutions have relied on for a decade are being rewritten and supervisors tell them the answer is theirs to find. That is the problem this white paper addresses.

The Global Regulatory Patchwork, in Brief

KEY TAKEAWAY

No single framework is enough. Canada has a binding floor, the US a voluntary one, the EU a prescriptive but deferred one, and Asia a spectrum from binding rules to supervisory principles, while NIST and OWASP are rigorous but silent on operations. Each financial institution must supply the operating model. If a model goes astray, you will be accountable, not the algorithm. I have seen many leaders hesitate to put the operating model in place, and let AI use cases run askew; the failures will continue unless value and responsible guardrails are integrated end to end.

Four bodies of guidance now shape how a North American institution should design, run, and secure its AI programs, and a fifth, the EU AI Act, matters to any institution with European customers, operations, or vendors. A sixth view, from Asia, follows later in this section: five jurisdictions that have each chosen a different way to supervise AI. Based on our research and experience, none of these frameworks alone is sufficient, and as the next section lays out, none of them tells a financial institution how to operationalize what it requires. But each sets a different, necessary piece of the foundation to ensure operational excellence.

EXHIBIT B The Diverse Framework Landscape at a Glance: North America and Europe

Framework	Jurisdiction / Reach	Status	Core Focus	Key Date
OSFI Guideline E-23	Canada: all federally regulated FIs	Binding	Enterprise-wide model risk management, including AI/ML and third-party models	Effective May 2027
FSRA Operational Risk & Resilience Guidance	Ontario: insurers, credit unions	Regulatory expectation (not codified)	Governance, outsourcing accountability, resilience; AI implicit, not yet named	Issued June 2026
NIST AI RMF 1.0 + Gen AI Profile	United States: voluntary, widely referenced	Voluntary standard	Govern / Map / Measure / Manage across the AI lifecycle	RMF 2023; Gen AI Profile 2024
OWASP LLM & Agentic Top 10	Global: technical/security community standard	Voluntary standard	Concrete attack patterns: prompt injection, excessive agency, identity abuse, supply chain	2026 revisions
SR 26-2 / OCC Bulletin 2026-13 (replacing SR 11-7)	United States: federally supervised banks	Non-enforceable guidance	Traditional statistical/ML models; explicitly excludes generative & agentic AI	Issued April 2026
NYDFS AI Circular Letter	New York: NY-regulated insurers	Regulatory guidance	AI and external consumer data in underwriting and pricing	Finalized July 2024
CRI Financial Services AI RMF	North America: industry-led, voluntary	Voluntary sector framework	Sector-specific AI risk management, filling the US federal gap	Published Feb 2026
EU AI Act (as amended by the Digital Omnibus, Reg. 2026/1744)	European Union: providers and deployers, incl. FIs	Binding regulation	Risk-based obligations; high-risk (e.g., credit scoring) requires a quality management system, risk management, monitoring, incident reporting	High-risk: Dec 2, 2027
ESAs Statement on Frontier AI & ICT Risk (EBA / EIOPA / ESMA)	European Union: financial entities	Supervisory statement	Robust governance and risk management for frontier-AI cyber risk; DORA third-party oversight	Issued July 2026

Note: Status reflects public guidance as of October 2026; confirm against each regulator's current publications.

OSFI Guideline E-23: Canada's De Facto AI Governance Regime

OSFI's Guideline E-23 is the most consequential development for Canadian institutions in this area in a generation. Finalized in September 2025 and binding for all federally regulated financial institutions from May 2027, it is deliberately technology-agnostic: it does not create a separate rulebook for AI and machine

learning, but folds them into a single, principles-based model risk framework covering every model an institution relies on, including third-party and vendor models. It requires multi-disciplinary governance, explicitly including legal and ethics expertise, not just quantitative risk; a formal model inventory, risk-tiered review intensity, and continuous monitoring across the full model lifecycle. Applied proportionately to size and complexity, E-23 is, in practice, Canada's binding national AI governance regime for the financial sector filling the vacuum left when the federal Artificial Intelligence and Data Act died with the prorogation of Parliament and the expiry of Bill C-27. OSFI's Financial Industry Forum on AI also produced four principles: Explainability, Data, Governance, and Ethics (EDGE) that examiners use as shorthand for what good looks like.

FSRA: Operational Risk, with AI Implicit

FSRA's new Operational Risk and Resilience Guidance for Ontario insurers, issued in June 2026, is a reminder that AI risk and operational risk are converging whether or not a regulator uses the word "AI." The guidance is explicit that accountability for outsourced functions, including cloud and AI vendors, remains with the institution regardless of the outsourcing arrangement, and it requires board-level governance, vendor concentration-risk management, and tested resilience plans. It does not yet name AI-specific risks, which is itself a signal: Ontario institutions must map AI vendor exposure onto a general resilience framework today, and should expect that gap to close.

NIST AI RMF: The Voluntary Reference Standard

The NIST AI Risk Management Framework (AI RMF 1.0), released in January 2023, remains the most widely referenced voluntary standard in North America, organized around four functions: Govern, Map, Measure, Manage, that map cleanly onto a model's lifecycle. NIST's July 2024 Generative AI Profile (NIST-AI-600-1) extended the framework to the risks specific to large language models and generative systems, and a Critical Infrastructure Profile is in development that will likely draw financial services more directly into NIST's orbit. NIST AI RMF alignment has become a de facto credibility marker: clients increasingly ask for it by name even when no regulation requires it.

OWASP: Where Governance Meets the Attack Surface

Where NIST and OSFI operate at the governance layer, OWASP operates at the control layer; the technical detail that determines whether a guardrail actually holds under attack. The OWASP Top 10 for LLM Applications (2026 revision) and the newer OWASP Top 10 for Agentic Applications (published December 2025) are the sharpest statements available of how generative and agentic AI actually fail in production. The 2026 LLM list elevated Excessive Agency from sixth to third place and renamed System Prompt Leakage to Hidden Context Exposure, both reflecting a shift toward authorization and access-control failures rather than pure content-manipulation risk. The companion Agentic list, covering risks like identity and privilege abuse, insecure inter-agent communication, and "rogue agents" that operate outside policy while appearing legitimate; is becoming required reading for any institution deploying multi-agent workflows in payments, trading, or customer service, and increasingly shows up inside auditors' questions about who approved an agent's permissions and how quickly it can be shut off. In April 2026, JPMorgan Chase's Jamie Dimon warned that AI agents make cybersecurity risk "far worse."

EU AI Act: The Closest Thing to a Prescribed Operating Model

For institutions with European customers, subsidiaries, or vendors and for any North American institution benchmarking against the most prescriptive regime in the world. The EU AI Act is the one framework that comes close to prescribing an AI operating model. For high-risk systems, which include creditworthiness assessment and credit scoring, Article 17 requires a documented quality management system spanning design control, testing and validation, data management, risk management, post-market monitoring, serious-incident reporting, record-keeping, supply-chain security, and a defined accountability framework; in effect, a regulatory description of the lifecycle in Section 3. Financial institutions receive partial relief; most of that system can be met through the internal governance arrangements they already maintain under EU financial services law, while risk management, post-market monitoring, and incident reporting still apply in their own right. The timing has moved, but the direction has not. The Digital Omnibus on AI (Regulation (EU) 2026/1744), in force since July 27, 2026, defers stand-alone high-risk obligations from August 2026 to December 2, 2027 and embedded-product systems to August 2028, while leaving transparency obligations on their original August 2026 date. On the supervisory side, the European Supervisory Authorities (EBA, EIOPA, and ESMA)

jointly called in July 2026 for robust governance and risk-management frameworks against frontier-AI cyber risk, applied consistently across sectors and linked to DORA oversight of critical ICT third-party providers. The practical lesson for North American institutions: a lifecycle built to the standard in Section 3 already has the shape the EU requires, and the deadline has been deferred, not cancelled.

The Widening US Gap

The honest summary is that the US federal guardrail is currently thinner than it was eighteen months ago. SR 11-7, the Federal Reserve/OCC/FDIC guidance examiners used for thirteen years to judge model risk, was jointly rescinded in April 2026 and replaced with SR 26-2 and OCC Bulletin 2026-13; both explicitly "non-enforceable" and both explicitly excluding generative and agentic AI from scope. New York's Department of Financial Services has partly filled that gap for its own regulated insurers with a 2024 circular letter on AI and external data in underwriting and pricing, and the industry-led Cyber Risk Institute published a voluntary Financial Services AI Risk Management Framework in February 2026 that is quickly becoming the default sector reference in the interim. For most US institutions, the binding floor has been replaced by a voluntary one, which raises, rather than lowers, the importance of internal discipline.

Asia: Five Jurisdictions, Five Supervisory Postures

Asia is not one regulatory story. The five jurisdictions below sit at different points on the spectrum from binding rules to voluntary principles, and most of them moved within the last twelve months. We focus on the instruments most relevant to a bank, insurer, or pension fund. This is a leadership briefing, not a legal survey, and institutions should confirm current requirements with their regulators and counsel.

Singapore: From Principles to Supervisory Expectations

Singapore has chosen supervision over legislation. There is no standalone AI statute. Instead, the Monetary Authority of Singapore (MAS) has built its position in layers: the FEAT principles of 2018 (Fairness, Ethics, Accountability, Transparency), the Veritas toolkits that put them into practice, a December 2024 information paper on good practices in AI model risk management, and its Guidelines on AI Risk Management, proposed in November 2025 and reported as finalized in October 2026 with

application phasing in over the following two years. As described in reporting on the draft, the expectations are proportionate, cover generative AI and AI agents, and ask for board and senior-management oversight, an accurate AI inventory, a risk materiality assessment based on impact, complexity, and reliance, controls across the lifecycle, and continued accountability for third-party AI. MAS has since added an operationalization handbook through Project MindForge (March 2026), and Singapore's IMDA issued a voluntary governance framework for agentic AI in January 2026.

India: Sector-Led, and Moving Quickly

India's approach is led by its sector regulators. The Reserve Bank of India's FREE-AI Committee published its framework in August 2025: seven guiding principles (which it calls sutras), six pillars, and 26 recommendations spanning governance, assurance, customer protection, and AI incident reporting. It is advisory, but a June 2026 draft of RBI model risk management guidance for all regulated entities would turn much of it into practice, with a board-approved framework, a model inventory with risk tiering, independent validation of vendor models, kill switches, and disclosure to customers that they are dealing with AI. In capital markets, SEBI's 2025 amendments already make regulated entities solely responsible for the outputs of their AI and machine-learning tools, for data security, and for legal compliance, and SEBI has said fuller guidelines will follow. The insurance regulator, IRDAI, convened an AI working group in June 2026. India's national AI Governance Guidelines (November 2025) are voluntary and state that a new AI law is not required at this stage, while the Digital Personal Data Protection Rules of 2025 phase in through 2027.

China: Binding, Purpose-Built Rules

China regulates AI through binding rules written for specific uses rather than through a single AI law. The Interim Measures for Generative AI Services (August 2023) require lawful training data and, for services with public influence, a security assessment and algorithm filing; the algorithm recommendation and deep synthesis provisions add filing duties of their own, and rules for labeling AI-generated content took effect in September 2025. The amended Cybersecurity Law, effective January 2026, wrote the first AI clause into national law. For banks and insurers, the National Financial Regulatory Administration (NFRA) issued guiding opinions in mid-2026 on the safe development and application of AI. As reported, they set out some 32

measures across governance, development, data, computing, risk, capability, and oversight, including lifecycle management, risk tiering, and the principle that whoever uses an AI system is responsible for it. A comprehensive AI law appears on 2026 legislative work plans, but no draft had been submitted to the legislature as of July 2026.

Taiwan: A Framework Law on Top of a Detailed Sector Guideline

Taiwan pairs a new framework law with an earlier, detailed sector guideline. The Financial Supervisory Commission (FSC) issued its Guidelines for the Use of AI in the Financial Industry in June 2024, built on six principles: governance and accountability, fairness and human-centric values, privacy and customer rights, system robustness and security, transparency and explainability, and sustainable development. They ask institutions to assess each AI use against six risk factors, manage it across four lifecycle stages, oversee third-party suppliers, and give customers recourse when AI produces an unfavorable outcome. The AI Basic Act, passed in December 2025 and promulgated in January 2026, sets seven national principles and leaves sector regulators such as the FSC to apply risk-based rules aligned to a national risk classification. An FSC survey released in May 2025 found AI in use at roughly a third of the 383 entities surveyed and at 87 percent of banks, and the FSC chair has signaled that agentic AI and risk classification will be added to the guidance.

The Philippines: A Voluntary Memorandum Treated as a Floor

The Bangko Sentral ng Pilipinas (BSP) moved in mid-2026. Its Memorandum M-2026-031, Governance Principles for Artificial Intelligence in Financial Services, sets five principles under the acronym STARS: Sustainability, Transparency, Accountability, Responsibility, and Security. It is voluntary in form, but BSP describes it as its minimum supervisory expectations. It applies proportionately to every supervised institution, including outsourced providers and vendors, and follows the lifecycle from planning through development, validation, deployment, and monitoring. The controls it describes include an AI inventory, disclosure when AI influences decisions, audit documentation, human accountability, adversarial testing, and continuous bias monitoring. The National Privacy Commission's Advisory 2024-04 adds privacy impact assessments, human intervention, and a right to contest automated decisions wherever AI processes personal data. A consolidated, risk-based AI bill was

under discussion in the House of Representatives as of July 2026, but nothing has been enacted.

EXHIBIT C The Asian Framework Landscape at a Glance

Jurisdiction	Principal Instruments	Legal Force	Core Focus	Key Dates
Singapore	MAS Guidelines on AI Risk Management; FEAT principles; IMDA agentic AI framework	Supervisory expectations; no AI statute	Board oversight, AI inventory, risk materiality, lifecycle controls, third-party AI	FEAT 2018; Guidelines finalized Oct 2026, phased from Oct 2027 (as reported)
India	RBI FREE-AI framework; RBI draft model risk guidance; SEBI AI/ML rules	Advisory (RBI); draft (RBI); binding (SEBI accountability rule)	Seven principles and 26 recommendations; model inventory, tiering, kill switch, customer disclosure	SEBI Feb 2025; FREE-AI Aug 2025; RBI draft Jun 2026
China	Generative AI Interim Measures; amended Cybersecurity Law; NFRA guiding opinions for banking and insurance	Binding rules (CAC); sector guidance (NFRA)	Algorithm filing, security assessment, content labeling, lifecycle management, risk tiering	Aug 2023; labeling Sep 2025; Cybersecurity Law Jan 2026; NFRA Jun 2026
Taiwan	AI Basic Act; FSC Guidelines for the Use of AI in the Financial Industry	Framework law; supervisory guidelines	Seven national principles; six-factor risk assessment, four lifecycle stages, supplier oversight, customer recourse	FSC Jun 2024; AI Basic Act promulgated Jan 2026
Philippines	BSP Memorandum M-2026-031 (STARS principles); NPC Advisory 2024-04	Voluntary in form, treated as minimum supervisory expectations	Sustainability, transparency, accountability, responsibility, security; AI inventory and bias monitoring	NPC Dec 2024; BSP Jun 2026

Status reflects public sources as of October 2026; several 2026 items are drawn from secondary reporting and should be confirmed against each regulator's own publications.

What the Global View Teaches

Read together, these instruments differ in legal force but converge in substance. Whether the supervisor is OSFI, the European Commission, the Monetary Authority of Singapore, the Reserve Bank of India, Taiwan's FSC, or the BSP, the same seven expectations recur: know every AI system you run, tier each by risk, validate before release, keep a human able to intervene, hold third parties to account, monitor and report incidents, and name an executive who is answerable. None of them says how to run the machinery that delivers those expectations week after week. That is the work of an operating model, and it is why the lifecycle in Section 3 is organized

around the expectations rather than around any single rulebook. Exhibit D shows where each expectation lands.

EXHIBIT D One Operating Model, Many Rulebooks: Common Expectations Across Jurisdictions

Common Expectation	Where It Appears	Lifecycle Stage
Know every AI system	OSFI E-23 model inventory; MAS AI inventory; RBI draft model inventory; BSP AI inventory; EU AI Act quality management	Stage 1
Tier each use by risk	OSFI risk-tiered review; MAS materiality assessment; RBI draft tiering; Taiwan FSC six-factor assessment; China NFRA risk tiering; EU high-risk classification	Stages 1 and 3
Validate before release	OSFI E-23 lifecycle controls; RBI draft (including vendor models); BSP validation and adversarial testing; NIST Measure function	Stage 4
Keep humans in control and tell customers	EU AI Act human oversight; RBI FREE-AI and draft disclosure and kill switch; Taiwan FSC recourse; BSP disclosure; China NFRA human oversight (as reported)	Stage 5
Hold third parties to account	FSRA outsourcing accountability; MAS third-party AI; SEBI sole responsibility; Taiwan FSC supplier oversight; BSP vendors; China NFRA outsourcing risk	Stages 1, 4, and 6
Monitor, log, and report incidents	OSFI continuous monitoring; EU AI Act post-market monitoring and incident reporting; RBI FREE-AI incident reporting; BSP continuous bias monitoring	Stage 6
Name an accountable executive	SEBI; China NFRA (whoever uses is responsible); MAS board and senior management; OSFI and FSRA board-level governance	All stages (AIPOM)

Illustrative mapping from published instruments to the SalesChoice lifecycle; it is not a compliance determination for any jurisdiction.

The SalesChoice AI Lifecycle: Intake to Decommissioning

Every AI capability, built, bought, or embedded, should move through the same eight gates, each with a named owner, defined inputs, and required outputs. That is what turns policy into a repeatable operating model.

EXHIBIT E The SalesChoice AI Lifecycle at a Glance: Eight Stages, Three Feedback Loops, One Operating Model

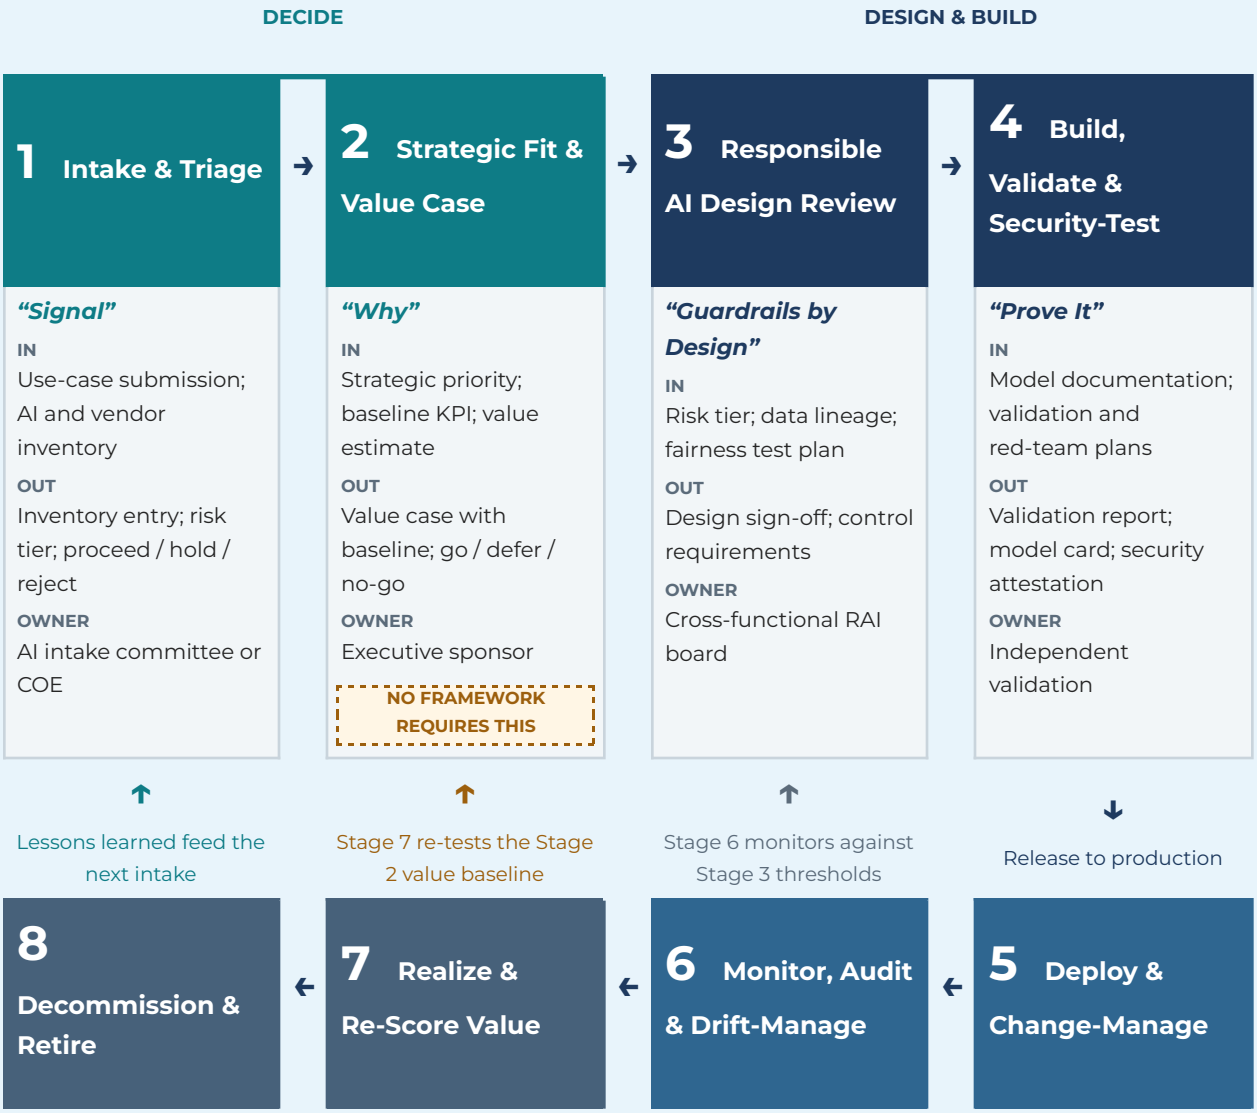

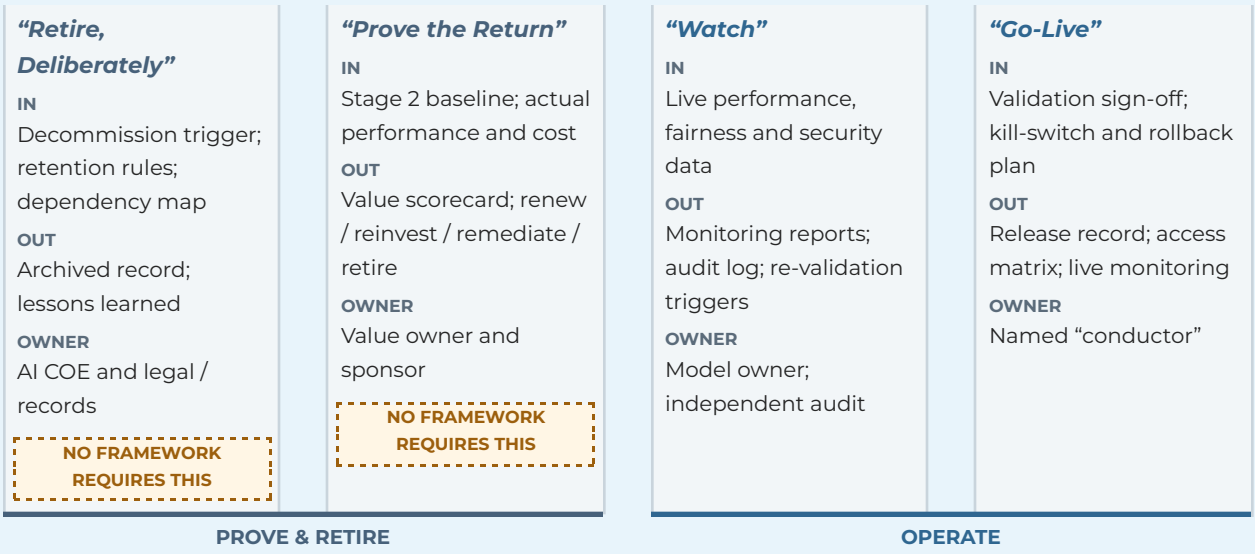

AIPOM AI Production Operating Model

The signalling, governance and audit layer that runs underneath every stage, for every AI capability, whether built in-house, bought, or embedded, regardless of which team owns it or which framework it answers to.

→ Lifecycle flow ↑ Feedback loop NO FRAMEWORK REQUIRES THIS SalesChoice-added control

Read left to right across the top row, then right to left across the bottom. Upward arrows mark where a later stage returns to what an earlier stage set.

Every framework above tells you what a responsible AI program must achieve. None of them tells you how to run one on a Tuesday morning when a product owner submits a new generative-AI use case for a contact center pilot. That operational gap between policy and practice is where we see the most AI programs stall, and it is where SalesChoice spends a great deal of our advisory time.

We call the connective tissue that closes that gap the AIPOM: the AI Production Operating Model: the signalling, governance, and audit layer that runs underneath every AI capability an institution operates, end to end, regardless of which team built it or which framework it answers to. Practically, the AIPOM is expressed as a single lifecycle with eight stages, each with a defined owner, a defined entry requirement (inputs), and a defined exit deliverable (outputs). An institution does not need all eight stages to be sophisticated on day one; however, it needs all eight stages to exist, even in a lightweight form, so that nothing enters production through a side door. The risks without relevant controls for AI use cases cannot be underestimated.

One Model, Any Regulated Organization

Although this paper is written for financial services, nothing in the lifecycle depends on the type of institution. A bank applies it to credit decisioning, fraud detection, and customer service. An insurer applies it to underwriting, pricing, and claims. A pension fund applies it to investment analytics, member communications, and benefit administration, where fiduciary duty raises the bar on explainability. Any other organization overseen by a government regulator, whether a securities dealer, a payments company, a utility, a healthcare provider, or a public agency, applies it to whatever its own decisions do to customers, members, or citizens. What changes from one organization to the next is the framework anchor at each stage and the thresholds that set the risk tier. The eight gates, the named owners, and the evidence do not change.

1. Intake & Triage “Signal”

Every proposed AI use case, whether built in-house, bought from a vendor, or embedded inside a SaaS tool the business already owns is logged in one place before any work begins. Triage assigns a provisional risk tier based on the population affected, the type of decision involved, and whether the model is generative, agentic, or traditional. This single step is what makes every later control possible: an institution cannot govern, monitor, or decommission a model it does not know exists.

Inputs: Use case submission (sponsor, problem statement, proposed data sources, build-vs-buy decision, affected customer/member population); current AI/vendor inventory; enterprise risk taxonomy.

Outputs: A logged, uniquely identified inventory entry; a provisional risk tier; a proceed / hold / reject intake decision.

Accountable: AI intake committee or AI COE triage lead Framework anchor: OSFI E-23 model inventory requirement; NIST “Map” function

2. Strategic Fit & Value Case “Why”

Before any technical work is funded, the use case has to earn its place against a defined value hypothesis and a strategic priority, not a vague sense that “AI could help here.” This is also where the use case is scored against the institution's broader AI portfolio, so that leadership can sequence investment deliberately rather than fund whichever proposal arrived first. No regulatory framework requires this step; it is the step that turns compliance into competitive advantage, which is why we treat it as non-negotiable.

Inputs: Strategic priority mapping; baseline KPI and current-state cost or risk; projected value estimate; feasibility and data-readiness score.

Outputs: An approved value case with a baseline metric; a prioritization score; a go / defer / no-go decision documented against portfolio sequencing.

Accountable: Executive sponsor and Value/ROI architecture owner

Framework anchor: No framework requirement; a SalesChoice-recommended control (see Section 5)

3. Responsible AI Design Review “Guardrails by Design”

A cross-functional review requires legal, compliance, privacy, security, and data science together, not in sequence, and happens before build or procurement, not after. This is where explainability requirements, bias and fairness test plans, data provenance expectations, and human-in-the-loop checkpoints are set as design constraints rather than retrofitted. For vendor and embedded AI, this is also where a third-party AI risk questionnaire is completed, closing the accountability gap FSRA's resilience guidance makes explicit: outsourcing the model does not outsource the responsibility.

Inputs: Confirmed risk tier; data lineage and provenance documentation; bias/fairness test plan; explainability requirement for the use case; vendor AI risk questionnaire (if applicable).

Outputs: A documented design review sign-off; specified control requirements (explainability method, monitoring thresholds, human-review triggers); a completed privacy impact assessment.

Accountable: Cross-functional Responsible AI review board Framework anchor: OSFI E-23 multi-disciplinary governance; OWASP supply-chain and data-poisoning categories; EU AI Act Art. 9 (risk management)

4. Build, Validate & Security-Test "Prove It"

Independent validation is a second-line function with no stake in the build, and tests the model against its documented purpose before it touches production, alongside security testing mapped explicitly to the OWASP LLM and Agentic categories most relevant to the use case (prompt injection and excessive agency for a conversational agent; identity and privilege abuse for anything with tool access). The output is not a verbal sign-off; it is a model card and a validation report an examiner can read without a translator.

Inputs: Model documentation (architecture, training data, parameters); validation test plan; security/red-team test plan scoped to relevant OWASP categories; **challenger**-model or benchmark comparison.

Outputs: An independent validation report; a security attestation; a completed model card; an approved deployment package.

Accountable: Independent model validation function (2nd line) and security/red-team Framework anchor: OSFI E-23 validation requirement; OWASP LLM & Agentic Top 10; EU AI Act Art. 17 (quality management)

5. Deploy & Change-Manage "Go-Live"

Production release is controlled, reversible, and accompanied by trained users, this is not a code push on a Friday afternoon. A kill-switch and rollback plan is a release requirement, not an afterthought, and every capability has a single named accountable owner (what we call its "conductor" in client engagements) so that when something goes wrong, there is exactly one person who is not surprised by the question.

Inputs: Validation sign-off; deployment runbook; rollback/kill-switch plan; staff training plan; internal communication plan.

Outputs: A production release record; an access-control matrix; an activated monitoring dashboard; a trained user population.

Accountable: Named business/technology co-owner (“conductor”) Framework anchor: NIST “Manage” function; FSRA operational resilience expectations

6. Monitor, Audit & Drift-Manage “Watch”

This is the AIPOM's signalling and audit layer made continuous: performance, fairness, and security metrics are tracked against the thresholds set in Stage 3, drift triggers automatic re-review rather than waiting for the next scheduled audit, and every material event is logged in a form a third-line audit function, or an examiner, can retrieve without a special request. Most institutions already do some version of this for traditional credit models; the gap is almost always generative and agentic AI, which are monitored far less consistently.

Inputs: Live performance and fairness data; drift-detection alerts; incident log; defined audit schedule and thresholds.

Outputs: Periodic monitoring reports; an audit findings log; re-validation triggers; an exam-ready evidence package.

Accountable: Model owner (1st line) and independent audit (3rd line) Framework anchor: OSFI E-23 ongoing monitoring; NIST “Measure” and “Manage” functions; EU AI Act Art. 72 (post-market monitoring)

7. Realize & Re-Score Value “Prove the Return”

On a defined cadence, quarterly for high-impact use cases, annually otherwise. Actual performance is measured against the value case approved in Stage 2. This is the step almost every institution we work with is missing, and it is the single highest-leverage addition a CEO or CFO can make to an existing governance program, because it is the only point in the lifecycle that asks whether the AI investment is actually working, not just whether it is still compliant.

Inputs: Baseline KPIs from Stage 2; actual performance, cost-to-run, and adoption data for the review period.

Outputs: A value realization scorecard; a renew / reinvest / remediate / retire recommendation to the executive sponsor.

Accountable: Value/ROI architecture owner and executive sponsor Framework anchor: No framework requirement; a SalesChoice-recommended control (see Section 5)

8. Decommission & Retire “Retire, Deliberately”

A model or agent is retired through the same discipline it was launched with: data retention and legal-hold requirements are honored, downstream dependencies are remediated so nothing silently breaks, and the model's full history including: design review, validation, monitoring record is archived rather than deleted. Critically, what was learned feeds back into how the next Stage 1 proposal is scored, closing the loop rather than starting every new use case from a blank page. This is the stage every framework in Section 2 is silent on the first of the five blind spots examined in Section 6.

Inputs: A decommission trigger (value failure, replacement system, vendor end-of-life, regulatory obsolescence); data-retention requirements; a dependency map.

Outputs: An archived decommission record; a dependency remediation plan; an updated inventory; a lessons-learned entry feeding back into Stage 1.

Accountable: AI COE and legal/records management Framework anchor: No framework requirement; the lifecycle's most significant governance gap

Read end to end, the eight stages are deliberately unglamorous. That is the point: the institutions that handle AI governance best are not the ones with the most sophisticated model, but the ones where every AI model ; simple or sophisticated moves through the same eight gates every time.

Exhibit F, on the following (landscape) page, summarizes all eight stages in a single reference view.

EXHIBIT F The AI Lifecycle, Intake to Decommissioning: Inputs, Outputs & Ownership at a Glance

Stage	Key Inputs	Key Outputs	Accountable Role	Primary Framework Anchor
Intake & Triage	Use case submission (sponsor, problem statement, proposed data sources, build-vs-buy decision, affected customer/member population); current AI/vendor inventory; enterprise risk taxonomy.	A logged, uniquely identified inventory entry; a provisional risk tier; a proceed / hold / reject intake decision.	AI intake committee or AI COE triage lead	OSFI E-23 model inventory requirement; NIST "Map" function
Strategic Fit & Value Case	Strategic priority mapping; baseline KPI and current-state cost or risk; projected value estimate; feasibility and data-readiness score.	An approved value case with a baseline metric; a prioritization score; a go / defer / no-go decision documented against portfolio sequencing.	Executive sponsor and Value/ROI architecture owner	No framework requirement; a SalesChoice-recommended control (see Section 5)
Responsible AI Design Review	Confirmed risk tier; data lineage and provenance documentation; bias/fairness test plan; explainability requirement for the use case; vendor AI risk questionnaire (if applicable).	A documented design review sign-off; specified control requirements (explainability method, monitoring thresholds, human-review triggers); a completed privacy impact assessment.	Cross-functional Responsible AI review board	OSFI E-23 multi-disciplinary governance; OWASP supply-chain and data-poisoning categories; EU AI Act Art. 9 (risk management)
Build, Validate & Security-Test	Model documentation (architecture, training data, parameters); validation test plan; security/red-team test plan scoped to relevant OWASP categories; challenger-model or benchmark comparison.	An independent validation report; a security attestation; a completed model card; an approved deployment package.	Independent model validation function (2nd line) and security/red-team	OSFI E-23 validation requirement; OWASP LLM & Agentic Top 10; EU AI Act Art. 17 (quality management)
Deploy & Change-Manage	Validation sign-off; deployment runbook; rollback/kill-switch plan; staff training plan; internal communication plan.	A production release record; an access-control matrix; an activated monitoring dashboard; a trained user population.	Named business/technology co-owner ("conductor")	NIST "Manage" function; FSRA operational resilience expectations
Monitor, Audit & Drift-Manage	Live performance and fairness data; drift-detection alerts; incident log; defined audit schedule and thresholds.	Periodic monitoring reports; an audit findings log; re-validation triggers; an exam-ready evidence package.	Model owner (1st line) and independent audit (3rd line)	OSFI E-23 ongoing monitoring; NIST "Measure" and "Manage" functions; EU AI Act Art. 72 (post-market monitoring)
Realize & Re-Score Value	Baseline KPIs from Stage 2; actual performance, cost-to-run, and adoption data for the review period.	A value realization scorecard; a renew / reinvest / remediate / retire recommendation to the executive sponsor.	Value/ROI architecture owner and executive sponsor	No framework requirement; a SalesChoice-recommended control (see Section 5)
Decommission & Retire	A decommission trigger (value failure, replacement system, vendor end-of-life, regulatory obsolescence); data-retention requirements; a dependency map.	An archived decommission record; a dependency remediation plan; an updated inventory; a lessons-learned entry feeding back into Stage 1.	AI COE and legal/records management	No framework requirement; the lifecycle's most significant governance gap

Each stage gate above is a control point, not a formality: an institution should be able to name the specific person accountable for the decision at every row, and produce the listed outputs on request.

Closing the Guardrails Gap: The Systems Behind Governance

KEY TAKEAWAY

A guardrail is real only if you can name the system that enforces it. Seven categories of technical infrastructure turn governance policy into audit-ready evidence.

Ask most FSI governance programs to list, within an hour, every generative-AI tool touching customer data, its validation record, and its last fairness test and most cannot. The policy exists; the plumbing does not. Operationalizing the AIPOM lifecycle in Section 3 requires seven categories of technical infrastructure, and an institution should be able to name the specific system that satisfies each.

1. AI & Model Inventory / Registry

The single system of record that every Stage 1 intake populates and every later stage updates, and required explicitly by OSFI E-23, and the one most institutions try to maintain in a spreadsheet, which fails the moment a second team adds to it. An institution cannot monitor, audit, or decommission what isn't inventoried.

2. Decision & Audit Logging

Immutable, retrievable logs of model decisions for every Tier 3 and Tier 4 use case; the inputs, model version, and confidence score behind each one. This turns “we believe our AI is fair” into evidence an examiner, a plaintiff's attorney, or a board member can inspect.

3. Data & Model Lineage / Versioning

Provenance tracking from raw data source through every model version that touched it, linked to the validation report each version passed. Without it, a Stage 6 drift investigation cannot tell whether the problem is the data, the model, or a silent upstream change; the question on which most post-incident reviews stall.

4. Real-Time Monitoring & Drift Detection

Dashboards that track performance, fairness, and security signals continuously against Stage 3's thresholds, with automatic escalation rather than a quarterly manual pull, the system that makes Stage 6 a continuous discipline, not a pre-exam fire drill.

5. Identity & Access Control for Models and Agents

Every model and autonomous agent needs its own identity and scoped permissions, distinct from its builders, directly addressing OWASP's Identity & Privilege Abuse and Excessive Agency categories, both elevated in 2026 because institutions are deploying agents with shared or overly broad credentials.

6. Third-Party & Vendor AI Risk Register

A dedicated register, separate from the general vendor-risk program, tracking every embedded or purchased AI capability, its risk tier, last review date, and concentration exposure. FSRA's resilience guidance makes it necessary even without saying “AI”: accountability for an outsourced model stays with the institution, so the institution needs its own record.

7. Human-in-the-Loop Workflow & Escalation Tooling

An auditable mechanism for routing exceptions, low-confidence predictions, and policy-flagged cases to a human reviewer, with the decision logged back to the audit trail. It is the system-level answer to Excessive Agency: not a policy saying “a human must approve this,” but a workflow that makes proceeding without one impossible.

A Note for Smaller Institutions

A 60-person community credit union might object that this looks like a big-bank stack. It need not be: for a handful of AI use cases, a well-configured spreadsheet-plus-ticketing combination can cover all seven categories. What matters is that every function exists and is consistently used, and the mistake we see most often is not under-investment, but skipping categories because none feels urgent on its own.

AI Value Realization: Metrics That Matter

KEY TAKEAWAY

Measure value, not activity. Four categories, a baseline set before funding and re-tested after launch, make AI clear the same bar as every other investment the board approves.

By some estimates, more than 80 percent of AI projects fail (RAND, 2024), and in our experience the cause is rarely that the models don’t work. It is that almost nobody measures whether they are working against the right definition of success. The most common failure we see is measuring activity instead of value; models deployed, queries processed, pilots launched, none of which tells a CFO whether the institution is better off or not.

SalesChoice's Value ROI Architecture organizes AI metrics into four categories that mirror how a board already evaluates every other investment so an AI business case clears the same bar as a branch expansion or core-system replacement, not a lower, AI-specific one.

EXHIBIT G The SalesChoice Value ROI Architecture: Four Categories of AI Value

Category	What It Captures
A. Financial Value	Revenue lift attributable to the capability (new-member conversion, cross-sell, retention), cost avoidance and productivity hours reclaimed, and direct margin impact. This is the category every AI business case claims; it is also the category almost none revisit after launch.
B. Risk-Adjusted Value	Fraud and loss dollars prevented, valued with the $5-per-dollar multiplier from Section 1, so every prevented loss has a defensible, board-ready figure, plus compliance cost avoided and capital relief from improved risk models. It makes a risk-management investment legible in financial terms, not just regulatory ones.
C. Adoption & Trust	Active usage rate, override or ignore rate (how often frontline staff disregard the model's recommendation, and why), and complaint or escalation rate. Override and adoption data are the earliest warning signs of a value problem, well before the financial numbers move.
D. Velocity	Time-to-decision, cycle-time reduction, and the time from Stage 2's approved value case to Stage 5's go-live. Institutions that track this almost always find their AI pipeline slower than assumed, with the delay concentrated in Stage 3's design review rather than the technical build.

What makes this more than a metrics taxonomy is that the loop closes structurally: the Stage 2 baseline is the baseline re-tested in Stage 7, and the Stage 7 verdict (renew, reinvest, remediate, or retire) feeds Stage 8. Value realization is not a reporting exercise bolted on; it is two of the lifecycle's eight gates.

What the Frameworks Don’t Tell You: Five Blind Spots

KEY TAKEAWAY

The frameworks are safety gates. None is a value gate, a retirement policy, or a learning loop. Here are the five blind spots.

Sections 2 through 5 synthesize what the frameworks and disciplined practice already tell us. This section is SalesChoice's own point of view, drawn from advisory engagements, on what the frameworks still don't address. We offer it to sharpen board-level conversation.

1. Nobody defines decommissioning, so institutions accumulate “zombie models.”

OSFI E-23, NIST, OWASP, and FSRA's guidance all describe, in detail, how a model should enter production and be monitored once there. None describes when or how it should leave. The predictable result is models that outlive their business sponsor, their original purpose, and often anyone's active awareness that they are still making decisions. Stage 8 of our lifecycle exists specifically to close it.

2. Compliance frameworks are safety gates, not value gates.

Every framework in this paper answers one question: is this AI system safe enough to run? None of them asks the second, equally important question: is it worth running? An institution can be fully E-23 compliant on a model that has quietly stopped delivering value for eighteen months. Compliance and value are different axes, and an operating model that only tracks the first one will keep well-governed, worthless AI in production indefinitely.

3. “Proportionality” is asserted, not operationalized.

OSFI explicitly scales E-23 to institutional size and complexity; NIST's RMF is voluntary by design partly for the same reason. But none of the frameworks tells a 40-person community credit union what a right-sized version of multi-disciplinary governance actually looks like in practice. In our experience, smaller institutions either import a big-bank process wholesale and grind to a halt, or quietly skip governance steps because no scaled-down version was ever specified. A right-sized AIPOM, lightweight at every stage, but with no stage skipped, is the practical answer.

4. Every framework still assumes “one model, one owner.”

E-23's model inventory, NIST's lifecycle functions, and even the original OWASP LLM Top 10 were all conceived around a single model making a single class of decision. Agentic AI breaks that assumption: when five specialized agents collaborate on one underwriting decision, “who owns this model” is the wrong question. OWASP's own 2026 Agentic Top 10 gestures at this with categories like Insecure Inter-Agent Communication and Rogue Agents, but no framework yet specifies how to assign accountability at the workflow level rather than the model level.

5. Nothing closes the loop from decommissioning back to intake.

When an AI use case fails, underdelivers, gets retired, or never leaves the pilot, that failure almost never gets formally scored against the next proposal entering Stage 1. Each new use case is evaluated on its own merits, with no structural memory of what the institution has already learned the hard way. Stage 8's “lessons-learned” output is designed to prevent this, but building the feedback loop itself, not just documenting that one should exist, is work almost no institution we've encountered has done.

Case Studies: Responsible AI in Practice

Client details are anonymized at each client's request; the methodology and results described are accurate. Each case study follows SalesChoice's standard case-study structure (the Challenge, the Solution, and the Results), which is used consistently across saleschoice.com.

Case Study One: Sequencing AI Investment

A North American FSI needed a board-defensible way to decide which AI investment to fund first.

The Challenge

A North American FSI engaged SalesChoice to build the business case for its next three-year AI investment strategy. Like most regional institutions, it faced the full weight of the pressures described in Section 1 of this paper: margin and deposit competition, rising fraud exposure, and a membership base expecting the digital experience of a national bank from an institution a fraction of the size. Leadership's problem was not a shortage of AI ideas, rather it was too many of them, with no defensible way to sequence investment or reassure a risk-averse board that the sequencing was sound.

The Solution

SalesChoice built a capability map of thirteen core banking functions, from digital channels and member acquisition to AML and core technology operations, and scored every one of them through two deliberately independent lenses. The first lens plotted each capability by projected growth impact against projected productivity impact, the classic "where is the value" question. The second, run as a genuinely separate exercise rather than a restatement of the first, plotted the same thirteen capabilities by feasibility today against the scale of transformation each would require focused on the "can we actually do this, and how disruptive is it" question. Where a capability scored strongly on both lenses, independently, the recommendation to prioritize it stopped being one consultant's opinion and became a corroborated, board-defensible finding.

The Results

Three capabilities identified from both lenses: a contact-center AI copilot, AI-powered member acquisition and CRM, and AML/financial-crime detection AI.

- Each priority capability was assigned a single accountable executive sponsor. what SalesChoice calls its “conductor” and was scoped from day one against the AIPOM lifecycle described in Section 3 of this paper, rather than being built first and governed later.
- Leadership received a decomposition of each priority capability into its concrete sub-processes, benchmarked against named peer institutions already running comparable AI in production, bringing best practices to the organization intelligence.
- Two independent lenses arriving at the same three answers gave the leadership team something qualitatively different from a single prioritization framework: not a recommendation to trust, but a finding to act on; the element of this methodology SalesChoice has carried into every FSI engagement since.

Case Study Two: Designing an AI Production Operating Model

A North American credit union had live AI pilots running but did not have an operating model or risk framework to govern them.

The Challenge

A second North American FSI had already launched three live AI production pilots, but had no production operating model or risk framework to govern them. In the institution's own words at the start of the engagement: “[We do] not have an AI production operating model and production risk framework to ensure AI value realization outcomes are achieved.” Pilots were running without agreed accountabilities for model and drift monitoring, human-override controls, regulatory evidence trails, or performance measurement, precisely the guardrail gap described in Section 4 of this paper.

The Solution

SalesChoice designed a best-practice AI Production Operating Model (AIPOM) anchored on the institution's three live pilots, through stakeholder interviews and a discovery-and-alignment workshop. The engagement produced a production risk framework; a roles-and-responsibilities framework defining new accountabilities, including model and drift monitoring, human-in-the-loop verification, and performance and value measurement, spanning both internally built and vendor-developed AI; a documented Case for Action; and an as-is-versus-to-be communication framework the institution could take directly to its own board. Advisory support then translated the design into a gap-closure roadmap, new job descriptions, and a journey roadmap for closing the gap between current and target state.

The Results

For the first time, the institution had a named, board-ready AI Production Operating Model connecting its governance design to the three pilots actually running in production; not a policy sitting alongside them.

- New accountabilities were defined and assigned across both people (model and drift monitoring, human-in-the-loop verification, performance and value analysts) and technical controls (continuous uptime, regulatory evidence trails, human override controls, cybersecurity); the same categories this paper argues belong in every institution's guardrail stack (Section 4).
- The design phase's Case for Action gave leadership a documented as-is-versus-to-be view it could take directly to its board, closing the gap between having an AI policy and being able to demonstrate one in practice.
- The strength of the design phase's outcomes led the institution to commission a second, follow-on engagement extending the operating model into new workforce capabilities and an Agentic AI process-retrofitting framework, the clearest proof available that an operating model, once built, gets used rather than shelved.

Case Study Three: A C-Suite AI Immersion Experience

A major UK bank needed its most senior leaders to share one firsthand understanding of world-class AI practice, not just individual exposure to it.

The Challenge

A major UK-headquartered bank, with more than £18 billion in annual income and roughly 60,000 employees, among the largest retail and commercial banking groups in the United Kingdom faced a problem common to large, regulated institutions at every stage of AI maturity: its Tier 1 C-suite did not share a common, firsthand understanding of what world-class AI practice actually looks like. Individual executives had very different levels of AI fluency, shaped by whichever conferences, vendors, or internal briefings each happened to have been exposed to. That unevenness was slowing the one conversation that mattered most, aligning the entire leadership team behind a shared AI strategy because the group had no common reference point to align around.

The Solution

Rather than another internal workshop or slide deck, SalesChoice partnered with a leading global market research firm to design and deliver a curated five-day AI leadership immersion experience in Canada for the bank's entire Tier 1 C-suite. The itinerary took executives directly to some of the institutions most advanced in applied AI, including the Vector Institute. one of the world's leading AI research institutes alongside leading healthcare organizations and other banking institutions further along their own AI journeys, giving the group a cross-industry view of mature AI practice rather than a single vendor's pitch. Each site visit was paired with structured working sessions, designed and facilitated by SalesChoice, to translate what the group was seeing into a prioritized shortlist of AI use cases specific to the bank's own business. University design and AI foresight students joined as live visual artistic scribes, turning each day's discussion into graphic-recorded artifacts the C-suite could review, annotate, and carry directly into the next day's conversation rather than waiting only a written summary. A picture is like a thousand words.

The Results

The week produced something a conventional briefing never could: a Tier 1 C-suite that had watched AI work, in person, across three different industries, together. and could point to the same examples once they got home.

- Every member of the Tier 1 C-suite left with a shared mental model of production-grade AI, built from the same five days of firsthand exposure, replacing fragmented, individually-acquired AI literacy with one common reference point for every subsequent investment conversation.
- The facilitated working sessions produced a prioritized shortlist of AI use cases specific to the bank's own operations, each already pressure-tested against what the group had just seen working in practice at the Vector Institute and peer institutions; the C-suite arrived home with clearly agreed upon AI pilots and advancing 3 big AI bets to build more AI muscle.
- The graphic-recorded artifacts from each day gave the bank a tangible, visual record of the week it could take directly into board and investment-committee discussions, turning a week of travel into a reusable internal asset rather than a trip report nobody reopens.
- The immersion became the reference point the bank's leadership returned to in subsequent AI governance and investment conversations, evidence that closing the literacy gap at the very top of an organization is sometimes the highest-leverage Responsible AI intervention available, ahead of any framework or policy.

SalesChoice: Our Track Record in Responsible AI

SalesChoice's position on Responsible AI is not a marketing layer added after the fact, it is built into the product architecture and the advisory methodology alike.

- SalesChoice's platform is built around what the company calls Human-Centered Autonomous Intelligence: AI positioned explicitly as decision support, with humans retaining accountability for consequential business decisions, and every recommendation accompanied by the business signals and model parameters that produced it, not just the output itself.
- SalesChoice's explainability and governance practices are mapped directly to the NIST AI Risk Management Framework, ISO 42001 (AI management systems), OSFI's regulatory expectations, the OECD AI Principles, and the EU AI Act and the same frameworks discussed in Section 2 of this paper, applied to our own platform before we recommend them to clients.
- SalesChoice was recognized as Predictive Sales Analytics Advisory Firm of the Year in Canada at the Corporate INTL Global Awards in both 2026 and 2025, and has been named among Canada's leading technology companies.
- SalesChoice's AI Center of Excellence governance methodology, covering centralized, decentralized, hub-and-spoke, and federated operating models is published and in active use with enterprise clients, built on the same premise this paper argues throughout: that an AI Center of Excellence is now a required organizational structure for moving beyond pilots, not an optional maturity milestone.
- SalesChoice is led by Dr. Cindy Gordon, this paper's author, a globally recognized AI and digital-transformation advisor, certified board director (ICD.D) with a focus on responsible AI, member of multiple global AI advisory boards including the United States Artificial Intelligence Institute (USAII), The AI Forum Board of Advisors, and author of sixteen books, most recently The AI Precipice (Routledge, 2026), co-authored with Cathy Cobey, Global EY AI Assurance Head. Dr. Gordon's background spans senior roles at Accenture, Citicorp, and Xerox, and she also teaches AI strategy and digital transformation at the United States Artificial Intelligence Institute.

From the Author's Desk: The AI Precipice

SalesChoice's approach to Responsible AI is informed by Dr. Gordon's new book, The AI Precipice: An Executive Guide to Balancing AI Innovation and Safety, co-authored with Cathy Cobey and published by Routledge in 2026. Written to equip business leaders with frameworks for AI strategy, governance, and adoption, the book balances the pressure to innovate quickly against the requirement to do so safely, embedding safeguards that build organizational trust and resilience. It is organized in three parts: The Courage Needed for Balanced AI, which establishes responsible practice from inception; a Roadmap to Achieving Responsible AI, covering practical data monitoring and cost-efficient controls; and Utopia or Dystopia?, which looks ahead to 2035 scenarios across finance, government, healthcare, and telecom.

Don Tapscott calls it "a must-read for leaders."

This research white paper you are reading applies the same premise to the specific operating realities of financial services worldwide: that safety and innovation reinforce one another rather than compete. More information is available at aiprecipice.com.

SalesChoice AI Advisory Services: How We Help

SalesChoice's AI Advisory practice is organized around the same lifecycle this paper describes, each engagement type maps to specific stages of the AIPOM, so that clients can enter wherever they are furthest behind rather than restart the whole program.

AI Strategy & Readiness Assessment maps to Lifecycle Stages 1–2

A structured evaluation of governance, strategy, data, process, technology, and talent readiness, producing a prioritized capability roadmap rather than a generic maturity score. This is typically the entry point for institutions that have AI activity underway but no consolidated view of it.

AI Intake Operating Model & Center of Excellence Design maps to Lifecycle Stages 1, 3, 6

Design of the institution's AIPOM itself: selecting and standing up the right governance operating model (centralized, decentralized, hub-and-spoke, or federated), defining stage gates and accountable roles, and specifying the minimum viable technical stack from Section 4 sized to the institution.

AI Use Case Prioritization & Value ROI Architecture maps to Lifecycle Stages 2, 7

The two-lens prioritization methodology demonstrated in Section 7's first case study, paired with the four-category value measurement framework from Section 5, so that sequencing decisions and value tracking use one consistent architecture rather than two disconnected exercises.

Responsible AI Design & Compliance Enablement maps to Lifecycle Stage 3; Section 4

Translation of FSRA, OSFI, NIST, and OWASP requirements into concrete design-review checklists, explainability and bias-testing plans, and vendor AI risk questionnaires, built once, reused for every use case entering Stage 3.

AI Application Development maps to Lifecycle Stage 4

Design and build of custom AI and generative AI applications, through SalesChoice's InsightEngine™ platform or the client's own technology stack, scoped against the design-review requirements set in Stage 3.

ROI Sustainability & Optimization maps to Lifecycle Stages 6–7

Ongoing monitoring of deployed models and generative AI applications: performance, drift, and MLOps infrastructure paired with the value-realization scorecards that turn monitoring data into renew / reinvest / retire recommendations.

Process Modernization & Data-Led Transformation maps to Cross-cutting enabler

Data lineage, labeling, and process-modernization work that underpins every other service line; most AI programs stall on data readiness well before they stall on model quality.

Executive AI Immersion & Capability-Building Experiences maps to Lifecycle Stage 1

Curated, multi-day learning experiences that take senior leaders directly to world-class AI practice: research institutes, peer institutions, cross-industry site visits, paired with facilitated sessions that turn what they see into a prioritized use-case shortlist. Often the highest-leverage starting point when the blocker is executive fluency rather than technology (see Case Study Three).

A 90-Day Roadmap to Get Started

An institution does not need to build all eight lifecycle stages and all seven guardrail systems before it can start. It needs ninety days and executive sponsorship.

Days 1–30: Stand up Intake

Catalog every AI, machine learning, and generative AI initiative already in flight anywhere in the institution, including embedded and vendor tools nobody has formally reviewed. Assign a provisional risk tier to each. This alone typically surfaces more active AI use than leadership believed existed and is, by itself, the single highest-value ninety-day outcome available to most institutions.

Days 31–60: Design the AIPOM

Select a governance operating model sized to the institution, name accountable "conductors" for every capability already in production, define the eight stage gates in a form specific to your organization, and specify the minimum viable version of the seven guardrail systems from Section 4.

Days 61–90: Pilot the Full Lifecycle

Run one or two live use cases through the complete lifecycle, at minimum through Stage 4 with a documented value case from Stage 2 already in place. Brief the board using the two-lens prioritization approach from Section 7, so the next investment conversation starts from evidence rather than enthusiasm.

Conclusion: Why This Matters, and Why It Must Be Led from the Top

The research in this white paper points to one finding that cuts across every framework, every jurisdiction, and every case study: the gap between AI ambition and AI governance in financial services, in North America and in Asia alike, is an operating-model gap, not a technology gap, and not a regulatory one.

Why This Research Matters

First, the regulatory floor is diverging, not converging. Canada now has a binding model risk guideline that covers AI; the United States has replaced a binding one with a voluntary one; and the European Union has deferred, but not cancelled, binding obligations for high-risk AI until December 2027; and across Asia, supervisors range from binding rules in China and in India's securities market to supervisory principles in Singapore, Taiwan, and the Philippines. An institution that waits for a rulebook to tell it what good looks like will wait indefinitely, and will be judged against the standard it should have set for itself.

Second, the cost of the gap is concrete. Every dollar lost to fraud now costs North American institutions $5.00 to remediate, while 44 percent still rely mostly or entirely on manual defenses; and RAND's research suggests that more than four in five AI projects fail to deliver. Ungoverned AI is not a neutral risk. It is a drag on every other AI investment the institution makes.

Third, the frameworks are silent exactly where value is created or destroyed: who retires a model, who measures whether it paid back, who owns a workflow that five agents share. These are the five blind spots of Section 6, and they are closed by operating-model decisions, not by waiting for guidance.

Why Executive Leadership Cannot Be Delegated

Responsible AI is too often handed to compliance, risk, or technology as a project to be completed. The evidence in this paper argues for the opposite. Accountability for AI cannot be delegated: OSFI's E-23 requires multi-disciplinary governance, FSRA expects board-level governance and accountability for outsourced functions including AI vendors, and US supervisors have stepped back in a way that places the

judgment with management. The decisions that determine whether AI creates value, which capability to fund first, which to sequence behind it, which to retire, are capital-allocation decisions that only an executive team can make. The conflicts that stall AI programs, legal against product, security against speed can be resolved only by someone with authority over both. And leaders cannot govern what they have not seen work, which is why a UK bank of 60,000 employees brought its entire Tier 1 C-suite to Canada for a week before committing to a strategy (Case Study Three). The regulator's own words describe the arrangement plainly.

> We could try and dream up how to regulate this extraordinary new technology, but I think if we did, we would probably be acting too early and subject ourselves to the risk of doing more harm than good.
>
> — Peter Routledge, Superintendent, OSFI, Thomson Reuters conference, November 27, 2025

When the supervisor declines to write the rulebook for AI, it is handing the question to the people who run the institution. The institutions that treat that as an invitation to lead, rather than a gap to be navigated, will define what good looks like for everyone else.

Five Questions Every CEO Should Be Able to Answer This Quarter

- How many AI use cases, built, bought, and embedded in tools we already own are running in production today, and who is the single named executive accountable for each?
- Within an hour, could we produce the validation record and the most recent fairness test for every model that touches customer data?
- What value case justified each AI investment, and when was it last re-tested against actual results?

- Which of our AI systems would we retire first, and what process would make that decision?
- If our regulator asked tomorrow how the board oversees AI, what would we put on the table?

Each of these questions maps to a stage of the lifecycle in Section 3, and each can be answered credibly within ninety days using the roadmap in Section 10.

An institution that can answer all five has built an operating model robust enough to absorb whatever comes next from OSFI, FSRA, NIST, OWASP, or whichever body fills the gap SR 11-7's rescission has opened, because it was never built to satisfy a single framework in the first place. That is the difference between managing AI compliance and building AI advantage. SalesChoice works with executive teams at exactly this starting point, and we welcome the conversation.

> Responsible AI is not the brake on innovation; it is what allows an institution to move faster with confidence. The leaders who will define this decade are the ones with the courage to build the guardrails first: an operating model owned at the top, measured for value, and honest enough to retire what no longer serves its customers.
>
> — Dr. Cindy Gordon, CEO, SalesChoice, and co-author, The AI Precipice: An Executive Guide to Balancing AI Innovation and Safety.

About the Author

Dr. Cindy Gordon, CEO & Founder

Dr. Cindy Gordon is a globally recognized thought leader in AI and digital transformation. Her background spans financial services, telecommunications, professional services, and high technology, with senior executive roles at Accenture, Citicorp, Xerox, and XDLI, a venture capital firm, and she has led major transformation programs for global organizations in both the for-profit and not-for-profit sectors.

A serial entrepreneur and angel investor, she has supported companies including Eloqua and advises leading AI firms. She is a certified board director (ICD.D) with a strong focus on responsible AI and ethical governance, and serves on multiple global AI advisory boards.

Follow Cindy on **LinkedIn →**

References

Allen & Gledhill LLP. “MAS Consults on Proposed Guidelines for Artificial Intelligence Risk Management.” 2025. https://www.allenandgledhill.com/sg/publication/articles/31741/mas-consults-on-proposed-guidelines-for-artificial-intelligence-risk-management

Allen & Gledhill LLP. “MAS Publishes Information Paper on AI Model Risk Management.” December 2024. https://www.allenandgledhill.com/sg/publication/articles/29565/mas-publishes-information-paper-on-ai-model-risk-management

Alston & Bird LLP. “NYDFS Issues Final Circular Letter Guidance on Use of AI in Insurance Underwriting and Pricing.” August 2024. https://www.alston.com/en/insights/publications/2024/08/nydfs-issues-final-circular-letter-guidance-on-use

America's Credit Unions. “What to Expect from the Economy: 2026 Outlook.” 2026. https://www.americascreditunions.org/blogs/americas-credit-unions/what-expect-economy-2026-outlook

Axios. “Behind the Curtain: Jamie Dimon's warning.” April 2, 2026. https://www.axios.com/2026/04/02/jamie-dimon-ai-geopolitical-risk-axios-show

AZB & Partners. “SEBI Introduces Provisions in Relation to Usage of Artificial Intelligence and Machine Learning Tools by Market Infrastructure Institutions and Registered Intermediaries.” April 28, 2025. https://www.azbpartners.com/bank/sebi-introduces-provisions-in-relation-to-usage-of-artificial-intelligence-and-machine-learning-tools-by-market-infrastructure-institutions-and-registered-intermediaries/

Baker McKenzie. “Philippines: BSP Releases AI Governance Framework.” July 2026. https://www.bakermckenzie.com/en/insight/publications/2026/07/philippines-bsp-releases-ai-governance-framework

Baker McKenzie. “Taiwan AI Basic Act.” January 2026. https://www.bakermckenzie.com/en/insight/publications/2026/01/taiwan-ai-basic-act

Blake, Cassels & Graydon LLP. “OSFI Releases Final Guideline E-23 for Model Risk Management and AI Use by FRFIs.” 2025. https://www.blakes.com/insights/osfi-releases-final-guideline-e-23-for-model-risk-management-and-ai-use-by-frfis/

Borden Ladner Gervais LLP. “FSRA's New Operational Risk and Resilience Guidance for Ontario Insurers.” July 2026. https://www.blg.com/en/insights/2026/07/fsras-new-operational-risk-and-resilience-guidance-for-ontario-insurers

Cobey, Cathy, and Cindy Gordon. The AI Precipice: An Executive Guide to Balancing AI Innovation and Safety. Routledge, 2026. ISBN 9781041020196. https://www.aiprecipice.com

Covington & Burling LLP (Inside Privacy). “China Amends Cybersecurity Law and Incident Reporting Regime to Address AI and Infrastructure Risks.” October 28, 2025. https://www.insideprivacy.com/cybersecurity-2/china-amends-cybersecurity-law-and-incident-reporting-regime-to-address-ai-and-infrastructure-risks

Cyberspace Administration of China. “Interim Measures for the Management of Generative Artificial Intelligence Services.” July 2023 (effective August 15, 2023). https://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm

Cyberspace Administration of China. “Measures for Labeling AI-Generated Synthetic Content.” March 2025 (effective September 1, 2025). https://www.cac.gov.cn/2025-03/14/c_1743654684782215.htm

Cycode. “OWASP Top 10 for Agentic Applications 2026.” December 2025. https://cycode.com/blog/owasp-top-10-agentic-applications/

Dimon, Jamie. 2023 Annual Letter to Shareholders. JPMorgan Chase & Co., April 2024. Quoted in Fortune, April 8, 2024. https://www.fortune.com/2024/04/08/jpmorgan-ceo-jamie-dimon-compares-ai-electricity-steam-engine-tech-augment-every-job

Elevate Consulting. “SR 11-7 Rescinded: What Replaced It and the AI Gap.” 2026. https://elevateconsult.com/insights/sr-11-7-rescinded-model-risk-guidance-ai-gap/

European Insurance and Occupational Pensions Authority (with EBA and ESMA). “EBA, EIOPA and ESMA Call for Enhanced Governance and Consistent Supervision to Mitigate ICT Risks from Frontier AI.” July 31, 2026. https://www.eiopa.europa.eu/eba-eiopa-and-esma-call-enhanced-governance-and-consistent-supervision-mitigate-ict-risks-frontier-2026-07-31_en

European Union. Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 17: Quality Management System. Official Journal of the European Union, July 12, 2024. https://artificialintelligenceact.eu/article/17/

Executive Yuan, Taiwan. Promulgation of the Artificial Intelligence Basic Act. January 14, 2026. https://www.ey.gov.tw/Page/9277F759E41CCD91/999df6da-6539-4941-913f-5125d684c58d

Federal Reserve Bank research staff. “Supply Chain Risk and Bank Lending Amid Trade Policy Uncertainty.” FEDS Notes, January 2026. https://www.federalreserve.gov/econres/notes/feds-notes/supply-chain-risk-and-bank-lending-amid-trade-policy-uncertainty-20260130.html

Financial Supervisory Commission, Taiwan. Press release on its survey of AI use in the financial sector. May 20, 2025. https://www.fsc.gov.tw/ch/home.jsp?id=96&parentpath=0%2C2&mcustomize=news_view.jsp&dataserno=202505200001&dtable=News

Finopotamus. “Five Trends Shaping Credit Union Strategy in the Second Half of 2026.” 2026. https://www.finopotamus.com/post/five-trends-shaping-credit-union-strategy-in-the-second-half-of-2026

Gibson, Dunn & Crutcher LLP. “EU AI Act Omnibus Agreement: Postponed High-Risk Deadlines and Other Key Changes.” May 2026. https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/

Hunton Andrews Kurth LLP. “EU Digital Omnibus on AI Enters Into Force.” 2026. https://www.hunton.com/privacy-and-cybersecurity-law-blog/eu-digital-omnibus-on-ai-enters-into-force

Insurance Business Canada. “OSFI not thinking of touching AI regulation for now, says superintendent Peter Routledge.” November 27, 2025. https://www.insurancebusinessmag.com/ca/news/technology/osfi-not-thinking-of-touching-ai-regulation-for-now-says-superintendent-peter-routledge-557990.aspx

LeapRate. “MAS Finalises AI Risk Guidelines for Financial Firms.” October 7, 2026. https://www.leaprate.com/regulation/mas-finalises-ai-risk-guidelines-for-financial-firms

Lee and Li, Attorneys-at-Law. Newsletter on the FSC Guidelines for the Use of AI in the Financial Industry. July 4, 2024. https://www.leeandli.com/EN/Newsletters/7287.htm

LexisNexis Risk Solutions. “Every Dollar Lost to Fraud Costs North America's Financial Institutions $5.” September 2025. https://risk.lexisnexis.com/about-us/press-room/press-release/20250910-fraud-multiplier

Linklaters LLP. “Singapore: MAS Proposes Comprehensive AI Risk Management Guidelines for Financial Institutions.” November 2025. https://financialregulation.linklaters.com/post/102lw8a/singapore-mas-proposes-comprehensive-ai-risk-management-guidelines-for-financial

Ministry of Digital Development and Information, Singapore. “Singapore Launches New Model AI Governance Framework for Agentic AI.” January 22, 2026. https://www.mddi.gov.sg/newsroom/singapore-launches-new-model-ai-governance-framework-for-agentic-ai--/

Ministry of Electronics and Information Technology, India (via Press Information Bureau). “MeitY Unveils India AI Governance Guidelines.” November 5, 2025. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2186639

Montreal AI Ethics Institute. “The Death of Canada's Artificial Intelligence and Data Act: What Happened, and What's Next for AI Regulation in Canada?” 2025. https://montrealethics.ai/the-death-of-canadas-artificial-intelligence-and-data-act-what-happened-and-whats-next-for-ai-regulation-in-canada/

National Financial Regulatory Administration (China). Guiding Opinions on the Safe Development and Application of Artificial Intelligence in the Banking and Insurance Industries (关于银行业保险业人工智能安全开发应用的指导意见), June 2026 (English title is our translation). Summarized in AML Network, August 6, 2026, which we relied on because we could not read the full text. https://amlnetwork.org/aml-news/china-ai-finance-rules-nfras-2026-bank-and-insurer-framework/

National Institute of Standards and Technology. “AI Risk Management Framework (AI RMF 1.0).” January 2023. https://www.nist.gov/itl/ai-risk-management-framework

National Institute of Standards and Technology. “Generative AI Profile (NIST-AI-600-1).” July 2024. https://airc.nist.gov/

National Privacy Commission, Philippines. “Advisory No. 2024-04: Guidelines on the Application of the Data Privacy Act to Artificial Intelligence Systems Processing Personal Data.” December 19, 2024. https://privacy.gov.ph/wp-content/uploads/2024/12/Advisory-2024.12.19-Guidelines-on-Artificial-Intelligence-w-SGD.pdf

Office of the Superintendent of Financial Institutions. “Backgrounder: Guideline E-23: Model Risk Management.” September 2025. https://www.osfi-bsif.gc.ca/en/news/backgrounder-guideline-e-23-model-risk-management

Office of the Superintendent of Financial Institutions. “Superintendent Peter Routledge participates in a fireside chat at Morningstar DBRS.” October 30, 2025. https://www.osfi-bsif.gc.ca/en/news/superintendent-peter-routledge-participates-fireside-chat-morningstar-dbrs

OWASP Foundation / Aembit. “The OWASP Top 10 for LLM Applications (2026): What Changed and Why It Matters.” 2026. https://aembit.io/blog/the-owasp-top-10-for-llm-applications-2026-what-changed-and-why-it-matters/

Philstar. “Proposed AI Bill Gains Momentum.” July 14, 2026.
https://www.philstar.com/opinion/2026/07/14/2542095/proposed-ai-bill-gains-momentum-congress

RAND Corporation (Ryseff, J., De Bruhl, B. F., and Newberry, S. J.). “The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed.” Research Report RR-A2680-1, 2024.
https://www.rand.org/pubs/research_reports/RRA2680-1.html

Reserve Bank of India. “FREE-AI Committee Report: Framework for Responsible and Ethical Enablement of Artificial Intelligence.” August 13, 2025. https://rbi.org.in/scripts/PublicationReportDetails.aspx?ID=1306

Reserve Bank of India. “Guidance on Regulatory Principles for Model Risk Management, 2026” (draft for comment). June 2026. https://www.rbi.org.in/Scripts/bs_viewcontent.aspx?Id=5089

SalesChoice Inc. “AI & Data Science Advisory Services.” https://www.saleschoice.com/ai-data-science-services/

SalesChoice Inc. “Case Study: How SalesChoice Delivers Explainable AI in an Era of Increasing Regulation.”
https://www.saleschoice.com/case-study-how-saleschoice-delivers-explainable-ai-in-an-era-of-increasing-regulation/

SalesChoice Inc. “Scaling AI Responsibly: A Guide to AI Center of Excellence (COE) Governance.” 2026.
https://www.saleschoice.com/scaling-ai-responsibly-guide-to-ai-center-of-excellence-coe-governance-whitepaper-by-saleschoice/

Securities and Exchange Board of India. “Consultation Paper on Guidelines for Responsible Usage of AI/ML in Indian Securities Markets.” June 20, 2025.
https://www.sebi.gov.in/reports-and-statistics/reports/jun-2025/consultation-paper-on-guidelines-for-responsible-usage-of-ai-ml-in-indian-securities-markets_94687.html

Sunstar. “BSP Issues AI Guidelines for Financial Firms.” June 2026.
https://www.sunstar.com.ph/davao/business/bsp-issues-ai-guidelines-for-financial-firms

TaxGuru. Report on the IRDAI office order of June 17, 2026 constituting a Working Group on AI. June 2026.
https://taxguru.in/corporate-law/irdai-forms-ai-working-group-strengthen-governance-insurance-sector.html

Let's Talk AI

THAT WORKS FOR YOU

Get in touch to learn how SalesChoice can help you:

- ✔ Plan your AI journey to future-proof your investments
- ✔ Forecast sales with increased confidence
- ✔ Clean and modernize your Salesforce
- ✔ Capture the opportunity of AI with real-time coaching
- ✔ **Unleash the latent power of your workforce with MoodInsights™**
- ✔ Make data-driven decisions you trust
- ✔ Build agentic AI routines for smarter workflows
- ✔ Ensure explainable and responsible AI

Book a call at **SalesChoice.com**

Prefer to email? Say hi at **hello@saleschoice.com**