VIEW POINT

ACCELERATING AI IN RISK AND COMPLIANCE

A BOARD-LEVEL IMPERATIVE FOR CANADIAN BANKS

Infosys
Navigate your next

AI in Risk and Compliance: The Canadian Context

Artificial Intelligence (AI) in risk and compliance has moved beyond a technology experiment and is fast becoming the defining factor between institutions that absorb risk and those that anticipate and contain it.

Yet most banks are getting it wrong. They're building models, running pilots and showcasing innovation, but they haven't put the basics in place to scale AI safely or deliver consistent outcomes.

Canadian banks are operating in an environment defined by three converging forces:

1. Rising regulatory scrutiny under OSFI Guidelines B-13 and E-23, and Québec Law 25
2. Heightened sophistication in cyber threats, fraud and financial crime
3. Growing data interconnectedness driven by the rollout of open banking from 2026

Despite the urgency, many Canadian financial institutions are finding it difficult to turn AI adoption into meaningful enterprise-wide impact. Statistics Canada reports that while 30.6% of finance and insurance organizations report some use of AI, it is most often limited to narrow applications like text analytics or virtual agents, rather than being embedded in core risk decision-making or control functions.

Canadian regulators are sending a consistent signal. AI adoption must be disciplined, governed and accountable.

- **OSFI Guideline B-13** mandates layered governance, operational resilience, and robust cyber controls for federally regulated financial institutions.
- **OSFI Guideline E-23** reinforces expectations for AI and Machine Learning (ML) model validation, signaling heightened scrutiny of advanced analytics in decision-making.
- **Québec Law 25** places transparency and explainability at the centre of automated decision-making
- **Open banking framework** (2026 rollout) will launch under the supervision of Bank of Canada, signalling a clear regulatory push toward financial modernization and controlled data sharing

The implication is straightforward. AI systems in risk and compliance must be explainable, auditable, controlled and enterprise-ready by design.

Investment Is Rising but Enterprise Impact Is Not

Despite significant investment, AI adoption in risk and compliance remains trapped in isolated pilots, fragmented data environments and disconnected control functions. The result is a familiar pattern. More models enter the pipeline, yet impact remains limited. False positives stay high, investigations move slowly, and reporting continues to rely heavily on manual effort.

The biggest issue is fragmentation. Data is scattered across fraud, AML, cyber and compliance teams and therefore AI ends up solving problems in silos rather than spotting patterns across risk domains. Ownership is often unclear too. AI efforts typically sit with innovation or data science groups, while the teams accountable for risk and controls remain on the sidelines. That disconnect makes it hard to move from a promising pilot to something that can actually run in production.

There is also a governance gap. Too often, explainability, traceability and validation are treated as afterthoughts, added late in the process once regulators, audit or security teams raise concerns. That slows deployment, reduces trust among investigators and compliance teams, and keeps manual workloads high.

Beyond governance failures, the operating model itself is the problem. Many Canadian financial institutions are still running risk and compliance in a reactive mode, responding after losses occur or after issues surface through alerts and complaints. In an environment where fraud sophistication is rising and cyber threats are evolving, that model will be inherently ill-equipped to keep pace

The fraud data reinforces the urgency. Canadian securities regulators reported over $310 million in investment fraud losses in 2024. The Canadian Anti-Fraud Centre estimates that 90-95% of fraud goes unreported which means an investigation-led, after-the-fact model will always capture only a fraction of the real harm. Canada's National Cyber Threat Assessment 2025–2026 further warns that the cyber environment is becoming more complex, and that ransomware threats continue to evolve.

Enforcement outcomes show the cost of weak monitoring. Financial Transactions and Reports Analysis Centre of Canada (FINTRAC) imposed an administrative monetary penalty on First Nations Bank of Canada in 2025, citing failures including suspicious transaction reporting, risk assessment and ongoing monitoring of business relationships.

The issue is not the lack of technology. The problem is that AI is still treated as an experiment, instead of being built into the bank's core risk and compliance controls. When AI is not linked to clear risk outcomes, embedded into everyday workflows, and governed like any other critical control, it fails to deliver real results.

From Fragmented Pilots to a Scalable, Regulator-Ready AI Operating Model

Closing the gap between experimentation and enterprise impact requires a deliberate shift from AI as an isolated capability to AI as a core control layer that is governed, integrated and measured against operational outcomes.

This shift rests on six foundations.

Reposition AI as a Board-Level Agenda

The gap between experimentation and enterprise impact is rarely technical. It is an ownership problem. Too often, AI remains confined to data science teams instead of being anchored to business outcomes, risk appetite and accountable decision-making. Closing this gap starts with leadership clarity on what AI is expected to deliver:

- Detect fraud earlier and more accurately
- Accelerate investigations and case resolution
- Identify anomalies and emerging risks in near real time
- Strengthen regulatory reporting quality and responsiveness
- Reduce false positives and manual effort
- Build customer trust through consistent, explainable decisions

Integrated Risk Intelligence across Domains

AI changes the fundamental equation of risk management, shifting the operating model from reactive, after-the-fact detection toward anticipatory control. This shift depends on strong data foundations including high-quality transaction data, KYC records, alerts and behavioural signals, with clean data lineage, governed access and reusable data assets.

Fraud signals, AML red flags, conduct indicators and credit anomalies rarely occur in isolation A unified intelligence layer, anchored in these foundations and connecting the relevant domains, enables predictive models that identify emerging risks before they escalate and uncover cross-domain patterns that siloed functions inherently miss

The potential scale of impact enabled by these foundations is illustrated through two client engagements.

A large bank faced material weaknesses in fraud controls. More than 80% of fraud losses were going undetected, primarily from account takeover and false positives exceeded 80 - 95% across digital channels. Infosys implemented a centralised fraud data platform, standardising data across source systems and making fraud signals reusable across channels. Advanced analytics improved alert quality, automated the handling of low-risk alerts and introduced knowledge graphs to uncover fraud networks. The results were significant:

- 83% of previously undetected losses addressed through improved alerting and intelligence
- False positive rates in high-volume digital streams significantly reduced through alert prioritisation and rule refinement
- 99% detection coverage for ACH and 98% for wire transactions
- Improved linkage between alerts and claims, strengthening investigative effectiveness

In another case, a mid-sized bank was experiencing net fraud losses materially above peer benchmarks, with fragmented point solutions driving high operating costs and inconsistent controls. Infosys delivered a holistic anti-fraud transformation, rationalising existing solutions and replacing them with an Enterprise Fraud Management platform, extending fraud risk models consistently across products and channels, and transforming loss-prevention operations. An omnichannel customer communications service enabled real-time customer engagement for faster resolution. Outcomes included:

- $6.3M in combined technology and operations savings through platform rationalisation and efficiency gains
- $0.8M reduction in net fraud losses
- Enhanced customer experience through real-time, omnichannel communications

Responsible AI Embedded from the Outset

Responsible AI is not just a compliance checkbox but a competitive differentiator. Institutions that embed fairness, transparency and accountability into AI design can move into higher-impact use cases with greater confidence and less rework. This requires categorising AI use cases by impact and sensitivity, standardised bias testing and drift monitoring, and defined boundaries for human oversight. Validation templates, model cards and documentation should become operational artifacts, not afterthoughts.

Regulatory Readiness by Design

Under evolving OSFI guidance, AI systems must generate traceability and audit evidence as a native feature. Regulatory readiness cannot be retrofitted after deployment.

This requires:

- Full lifecycle documentation across design, training, validation and deployment
- Automated monitoring of bias and performance drift
- Seamless connection to existing risk and compliance workflows

Secure MLOps environments enable rapid experimentation while preserving control and auditability. Institutions that embed regulatory expectations by design reduce deployment risk and increase confidence in scaling AI across the enterprise.

A Scalable Operating Model

Scaling AI in risk and compliance requires safe, consistent and measurable expansion across the enterprise, rather than isolated tool adoption.

A scalable operating model requires:

- Clearly defined cross-functional accountability across business, compliance, data and technology
- Governance forums with decision-making and deployment authority
- Standardized release gates, validation workflows and approval criteria
- Strengthened third party AI vendor governance.

Reusable accelerators such as validation frameworks, MLOps pipelines and governance workflows reduce friction from prototype to production while preserving control and auditability.

Security as a Foundation, Not an Afterthought

In Canada's elevated cyber risk landscape, AI must strengthen the security posture of financial institutions while not complicating it. As AI systems become embedded across risk and compliance workflows, they also expand the attack surface and heighten the importance of robust technology risk governance under OSFI Guideline B-13.

AI-driven behavioral analytics and anomaly detection can significantly strengthen threat detection and expedite response when implemented within a well-governed architectural framework

Security, governance and AI cannot operate as siloed agendas. When designed together, they form a resilient control environment that enables institutions to scale AI with confidence.

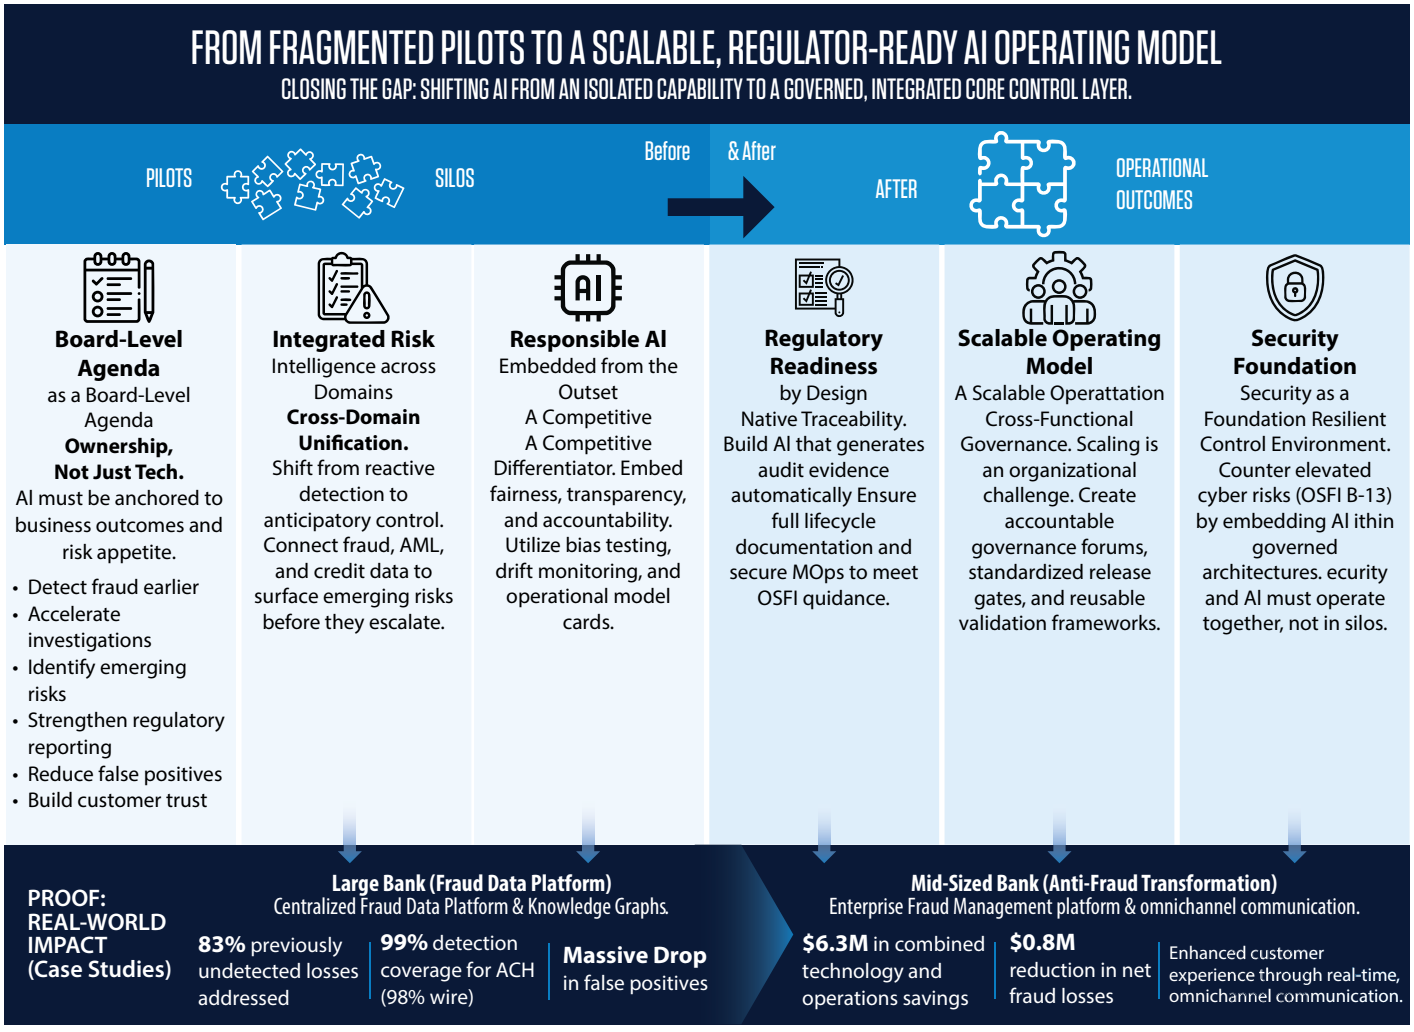

From Fragmented Pilots to a Scalable, Regulator-Ready AI Operating Model

The CXO Action Plan and Success Metrics - The Infosys Framework

Based on our global experience, Canadian banks face a clear imperative to move from isolated AI pilots to enterprise-scale adoption across risk and compliance, without compromising governance or control. This requires a structured sequence of executive decisions and measurable deliverables.

CXO Action Plan

1. Set governance foundations first

Define AI risk tiers, clarify enterprise-wide accountability, and establish non-negotiable controls, including traceability, access controls, audit logs and drift monitoring

2. Focus on a few high-impact use cases

Select a few measurable journeys, such as AML triage, fraud resolution, or regulatory reporting. Anchor them to clear targets for cycle-time reduction and risk containment. Demonstrate ROI before expanding.

3. Build a strong control backbone

Standardize validation, automate bias/drift monitoring, embed documentation within MLOps and integrate AI into existing risk and compliance workflows.

4. Pilot in controlled production environments

Deploy with defined guardrails, KPIs and oversight mechanisms. Scale only after governance thresholds and performance outcomes are demonstrably met

CXO Decision Framework for Scaling AI in Risk & Compliance

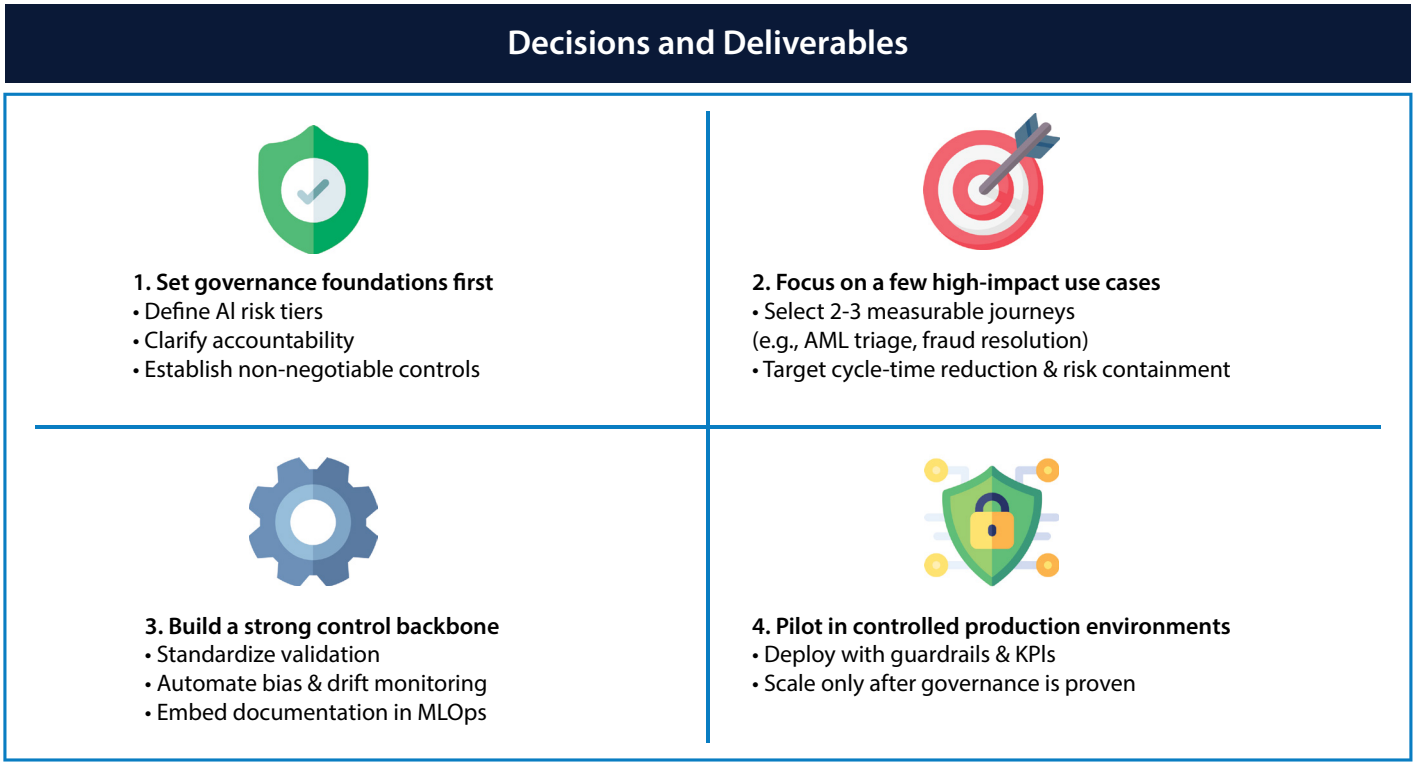

Metrics That Matter

AI in risk and compliance must be measured in operational and control outcomes, instead of experimentation volume. The metrics that matter are:

- Fraud detection accuracy — improvement in identifying fraud patterns and anomalous behaviour.
- False positive reduction — decrease in unnecessary alerts and manual reviews
- Investigation throughput — increase in cases processed per analyst or per cycle
- Cycle-time reduction — shorter case handling time (median and 90th percentile)
- Model performance — improved precision, stability, and effective drift management
- Operational efficiency — reduction in manual effort across testing, oversight, and documentation
- Early risk detection — faster identification of emerging threats and anomalies.

Strategic Outcome

When AI is embedded with governance, integrated into workflows, and tied to measurable impact, it shifts from experimental capability to institutional advantage.

This is not a theoretical leap. In markets like the UK, banks that moved early to embed AI into fraud detection, transaction monitoring and supervisory reporting were able to reduce false positives, shorten investigation cycles and improve regulator engagement at the same time. The real differentiator was the ability to embed intelligence directly into everyday decision-making while preserving transparency and explainability. Institutions that made AI part of how they operated were able to move faster and act with greater confidence than those that treated it as a peripheral capability.

The same opportunity exists for Canadian banks. As financial ecosystems become more interconnected through open banking, real-time payments and third-party data access, risk events will surface earlier and spread faster. Those that embed AI into their risk and compliance workflows will be positioned to detect emerging issues sooner, respond with greater precision and demonstrate control effectiveness continuously.

Accelerating AI adoption in risk and compliance is ultimately about resilience and trust. It gives institutions the ability to operate with speed without sacrificing oversight and to meet rising supervisory expectations while delivering better outcomes for customers. In an environment where confidence is hard-won and easily lost, that combination becomes a strategic differentiator rather than a future aspiration.

Authors

Archana Ashok
Principal – Business Consulting

Parantap Chakrabarti
Senior Consultant – Business Consulting

Sanket R Singhania
AVP - Infosys Financial Services, Canada

Sonal Pothapragada
Consultant – Business Consulting

References

1. BCG: AI adoption in 2024 — 74% of companies struggle to achieve and scale value (Oct 24, 2024). https://www.bcg.com/press/24october2024-ai-adoption-in-2024-74-of-companies-struggle-to-achieve-and-scale-value
2. OSFI: Technology and Cyber Risk Management. https://www.osfi-bsif.gc.ca/en/risks/technology-cyber-risk-management
3. Raymond Chabot Grant Thornton: Law 25 — The issue of automated decisions. https://www.rcgt.com/en/insights/expert-advice/law-25-issue-automated-decisions/
4. McMillan LLP: Balancing innovation and risk — OSFI's principles for responsible AI in finance. https://mcmillan.ca/insights/publications/balancing-innovation-and-risk-osfis-principles-for-responsible-ai-in-finance/
5. Canadian Financial Crime Academy: The role of technology in combating financial crime in Canada. https://www.canadianfinancialcrimeacademy.ca/financial-crime-articles/the-role-of-technology-in-combating-financial-crime-in-canada
6. https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024
7. TD Bank API & Data Solutions | Plaid
8. https://www.rbcroyalbank.com/business/api/index.html
9. https://developer.rbc.com/
10. [cyber.gc.ca] National cyber threat assessment 2025–2026
11. Analysis on artificial intelligence use by businesses in Canada, second quarter of 2025

For more information, contact askus@infosys.com

Infosys®
Navigate your next