IMPACT OF AI — FINANCIAL INSTITUTION SERIES

Impact on AI on Risk in Financial Institutions

CapTech

As financial institutions accelerate the adoption of AI, risk and compliance functions are undergoing a fundamental shift in how work is performed. While some risk teams in financial institutions may be slower to fully adopt AI, its active use across development, customer service, and marketing functions makes AI ethics and governance an immediate reality for risk teams.

Applying today's AI tools, repetitive tasks like data reconciliation, policy review, and anomaly detection are being automated, shifting human effort away from information gathering and toward strategic judgment, oversight, and decision-making. However, risk professionals remain essential to reviewing AI outputs, investigating flagged issues, validating models, interpreting insights, and taking accountability for high stakes decisions. This "human in the loop" model ensures AI enhances — not replaces — risk expertise.

In the future, risk and compliance professionals will operate closer to the teams delivering business

"Teams are still thinking of risk in silos when risk is happening in horizontals."
— CapTech Director, Expert in Data at Financial Institutions

Key Takeaways for Financial Institutions:

CapTech's predictions and recommendations

"While institutions will experience cost reduction as a result of AI efficiencies, risk staff reduction should not be considered until AI risk optimization is achieved."
— CapTech Technical Director

outcomes. Risk professionals will continue to define standards, manage AI governance, and provide advanced analytics, while business unit level roles oversee compliance integrated directly into business processes. This operating model replaces after the fact controls with real-time collaboration to address concerns as processes are designed.

This transformation will not materially change the size of risk teams, initially, but it will significantly increase their effectiveness. Institutions that pair AI adoption with strong governance, clear accountability, and sustained investment in AI literacy will not only meet regulatory expectations but position risk management to achieve competitive advantages.

The Future of AI and Risk in Financial Institutions

In recent years, risk teams have become leaner by using machine learning, setting parameters, and configuring rules engines to automate select activities. While this shift reduced reliance on manual reviews and lengthy reports, it still left teams increasingly removed from work while it was taking shape — creating gaps between risk, controls, and execution. The next phase of AI will address this disconnect, fundamentally reshaping how risk is managed.

> **"An AI ecosystem should not be considered a tollgate. But rather an accelerator that builds trust."**
> **— CapTech Managing Director**

Risk Teams of the Future

Across the organization, three distinct groups shape outcomes:

- The enterprise that supports governance across the organization
- The risk and compliance teams responsible for oversight
- The business units who create value for customers

Each group has unique needs, but the risk function sits at the intersection of all three. To be effective, risk teams require strong AI model validation capabilities and data expertise to interpret regulatory guidance, signals from business units, and enterprise level data. Bringing these inputs together enables an automated, continuous learning flow — one where human judgment remains central but is informed by real-time insight from across the organization.

Roles will change as the working model moves out of the manual space and into more strategic work.

- AI model validators will be required and will set the standard for what is measured, what is decided, and what is alerted to the risk team.
- AI ethicists will join risk teams to ensure all AI across the organization is used ethically and within the bounds of the law.
- Data governance experts will need to understand what each model is solving for and the source data and perform continuous monitoring for up-to-date and accurate data.

- Senior risk officers and compliance officers will need to be conversant in AI to oversee these new tools.
- All risk professionals will need to be adept at interpreting AI output and spotting where the models might be wrong, providing a "sense-check" and ethical oversight.

Human in the Loop

Technology skills aside, risk expertise and the ability to recognize what looks correct (versus skewed) remain essential. AI will not succeed without people, and people will not succeed without AI. The most common human-in-the-loop activities will likely be:

- Review AI-prepared reports and documents: When AI drafts risk assessments, compliance reports, or policy updates, human experts must review, validate, and approve the final output.
- Investigate and adjudicate AI-flagged alerts or exceptions: AI systems detect anomalies while human analysts step in to investigate and make judgment calls on these high-risk cases.
- Validate and tune risk models: Machine learning can auto-adjust models, but human experts continuously review model outputs, back-test predictions, recalibrate, override, and audit parameters as needed.
- Confirm autonomous AI actions ("agentic AI"): When AI agents act on their own, a human risk officer reviews and validates the action.
- Interpret and contextualize AI-generated insights: Risk and compliance officers provide context, expertise, and judgment to the patterns and red flags that algorithms identify.
- Verify compliance decisions: High-stakes compliance and risk decisions remain in human hands, providing the necessary check on AI outputs and taking responsibility for decisions to regulators and stakeholders.

Roles and Team Composition

The next four sections comprise CapTech's hypotheses and recommendations for how risk teams, roles, and responsibilities will evolve in the next 2–5 years as risk team and AI usage matures.

In the future, risk management will be positioned closer to the teams delivering business outcomes.

Regulation will be integrated into daily operations, shaping how work is performed rather than operating as a standalone process. While a core team remains responsible for enterprise AI model risk and advanced data analysis, most compliance and risk analysts will operate directly alongside the teams they support.

CapTech's work with financial institutions shows that this model preserves a strong enterprise risk function — responsible for setting standards, managing AI governance, and delivering advanced analytics — while enabling business-unit risk roles to oversee compliance as part of day-to-day execution. This shift moves risk out of silos and into the everyday vernacular of the enterprise.

The table on the next page illustrates how risk roles evolve from today's centralized team to a future-state model distributed across enterprise and business units. It highlights where AI is currently applied and the skills risk professionals will need — either to succeed in their existing roles or to transition into new roles within the risk domain. Regardless of role or location, risk expertise remains essential. Human judgment is what allows AI to function as an effective co-laborer; without it, AI lacks the correction and tuning required to deliver reliable outcomes.

Future-state skills represent the capabilities risk professionals should develop to operate effectively in an AI-centered environment.

ROLE	CURRENT SKILLS & ACTIVITIES	FUTURE SKILLS & ACTIVITIES
Compliance Officers	*6 roles using 10% AI* • Understand all aspects of regulations: federal, state, city, bank policy • Interface with colleagues, executives, and regulatory agencies to uphold and provide documentation of the high standards the organization holds • Present recommended policy and/or procedure updates for approval	*2 roles at the enterprise level using 65% AI* • Create an agent to pull all latest regulations and cross-align regulations with current policies and practices • Automate internal information gathering surrounding a specific policy • Use AI to draft policy in relation to updates in regulations • Function as the AI governance lead in tandem with business, IT, and executive leadership
Risk Officers	*6 roles using 10% AI* • Share any regulatory changes with the enterprise • Understand the landscape of risks, incidents, failed controls, and investigations • Interpret incident and risk themes, identify root causes	*2 roles at the enterprise level using 50% AI* • Identify systemic risk mitigation, reducing manual burden • Create a model to identify risk of compliance leaks within a process (proactive, systemic) • Identify potential areas where risk or compliance could be inherent in an enterprise process and create a strategy to overcome (ongoing real-time auditing)
Compliance Analyst	*6 roles using 10% AI* • Conduct audits in enterprise units • Prepare compliance documentation • Review policies for discrepancies and/or violations • Review transactions for violations or issues	*1 role for every 3–5 business units using 90% AI* • Interpret regulatory changes and AI-generated alerts • Create AI agent for auditing, documentation, and policy reviews • Ensure automated controls align with regulations • Make judgments on complex systems and where AI could/should be used and how it impacts controls • Function as compliance lead for several business units
Risk Analyst	*6 roles using 10% AI* • Perform incident tracking, control testing, investigations • Understand business processes and controls • Compile incident reports and records • Define risk mitigation measures	*1 role for every business unit using 60% AI* • Train, monitor, and tune models to compile incident reports and themes • Identify risk in new business processes and how to automate risk/control within the business process • Create/contribute to a library of AI tools to encourage fewer AI tools meeting the same needs • Use tech and data skills to work alongside IT to configure rule-based monitoring systems • Build strong business knowledge without compromising the risk lens
Data Scientist	*2 roles using 30% AI* • Compile reports and data to create a picture of risks, issues, and controls • Possess advanced skills in financial risk domains and regulations • Possess advanced data science skills	*3 roles in risk & compliance team using 90% AI, may drop to 2 once efficiency at scale reached* • Design and deploy machine learning models and analytics tools • Develop predictive risk models • Integrate data from siloed systems • Continuously train and improve detection algorithms to uncover complex risk patterns • Possess advanced skills in programming, AI/ML, financial risk domains and regulations • Possess advanced data science skills • Note: 3 roles in data science considered initially for a medium-large size organization due to the scale/amount of enterprise level models that could be developed and require governance, tuning, and auditing. This could revert to 2 when AI becomes more mature at the organization.

ROLE	CURRENT SKILLS & ACTIVITIES	FUTURE SKILLS & ACTIVITIES
AI Trainers/ AI Operations Specialists	*Does not exist today*	*3 roles in risk & compliance using 90% AI* • Curate and label datasets • Configure AI model parameters • Iterate on prompts and rules • Possess risk and ML expertise and serve as bridge between SMEs and data
Data Governance & AI Ethics Specialists	*Does not exist today*	*3 roles at enterprise level using 70% AI* • Oversee data quality, data governance, and ethical AI use • Create policies for data governance • Develop frameworks for responsible AI • Specialize in data, compliance, risk, and AI • Collaborate with IT, data, and risk SMEs to meet needs without breaching trust or creating new risks

In the graphic below, the work of the teams becomes spread across the organization. Continuous learning occurs through dynamic communication where the risk team broadcasts regulatory updates, while frontline teams surface emerging risks in real time. The enterprise team updates its models to adapt to new learning. This feedback loop allows the organization to refine controls and models based on real-world conditions.

Continuous learning at all levels keeps the flow of knowledge and improvement moving throughout the organization.

Enterprise Risk Make Up

As AI adoption expands, Compliance and Risk Officers will move to the enterprise level under the CCO or CRO. Centralized AI governance will be required to avoid fragmented tooling and multiple solutions addressing the same risks in parallel. Together with data governance and AI ethics specialists, this group will define approved use cases, tool ownership, and oversight — partnering closely with the Chief Data Officer to ensure disciplined data and intelligence governance. The pace of AI advancement makes it clear that data and AI governance will be among the largest risk areas organizations face in the next few years.

Risk and Compliance Team

With the adoption of AI, the risk and compliance function becomes the intelligence hub of the organization. Comprised of data scientists and AI trainers/model operators, this team serves as the analytical backbone of risk. It aggregates signals from both business units and enterprise systems to measure and manage risk, as intended. The team monitors outputs, continuously tunes centralized models, and recommends enterprise wide controls to prevent risk silos from forming or persisting. While risk data may be widely available across the organization, this group ensures the data is contextualized, consistent, and suitable for strategic decision making.

Business Unit Make Up

Historically, business units and risk teams operated as separate functions. Business processes were designed and executed first, then reviewed by risk and compliance, which identified gaps and required additional controls, such as checklists or end-of-process approvals. These requirements were then layered back onto the original process. As a result, risk management functioned as an add-on rather than an integrated component, often creating unorganized workflows and inconsistent customer experience. When risk and compliance professionals work directly with business units, concerns are addressed in real time. Developers can mitigate risk as processes are built, and product owners can clearly communicate impacts to business stakeholders — avoiding manual, end-of-process controls that slow delivery and detract from the customer experience.

Evolving Workforce and Culture

The overall size of risk and compliance teams is unlikely to change materially, at least in the near term. Instead, teams will operate with significantly greater effectiveness, taking on new skills and roles within the same team. Regulatory changes will be implemented more quickly, risk will be less likely to sit on the critical path for new products and services, and risk events will be identified and resolved earlier with reduced impact. The result is a more resilient organization.

Fewer administrative roles and more technology-focused roles are required for effective risk management of the future:

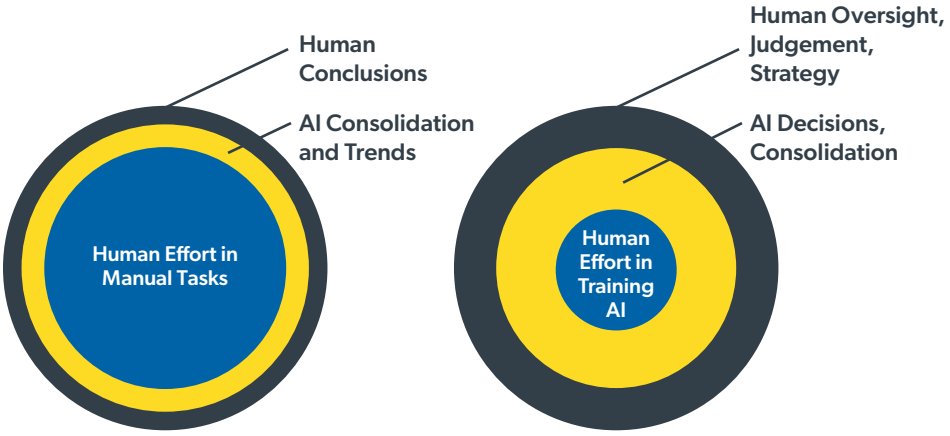

Risk teams of the future will move from manual tasks to AI training and strategy.

"…a major bank replace[d] a manual compliance document workflow with a generative AI solution that doubled extraction accuracy at just 5% of the original cost (saving about $2 million per year), which freed the compliance team to repurpose hundreds of hours toward higher-value analysis and oversight."
— CapTech Delivery Expert

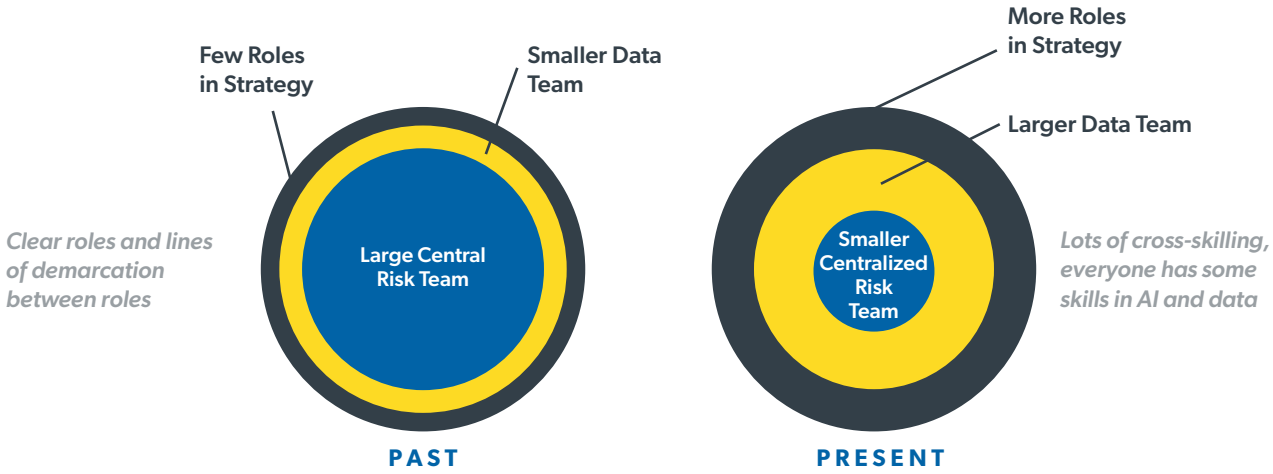

Training programs will be crucial — not only to upskill current staff, but to instill a culture of collaboration between humans and AI. CapTech's client experiences show that successful AI adoption always hinges on change management and training. For instance, in one large organization rolling out an AI assistant to employees, adoption was initially slow until an enablement program clarified how to use the tool and set clear guidelines for its use. Successful financial institutions focus on "AI literacy" for everyone and ensure clarity on the boundaries of AI decision-making versus human decision-making.

Banks must be ready to demonstrate control over their AI models — documenting how models are validated, how biases are checked, and how outcomes can be explained. Over the next two years, establishing a strong AI governance framework will be as important as the technology itself. This includes formal structures (e.g. AI risk committees, ethical AI guidelines, model documentation standards) and technical tools (such as bias detection software and robust monitoring of model performance in production). The firms that get this right will not only satisfy regulators but will gain confidence to push AI capabilities further. One CapTech expert described it as a need to:

> **"…rethink where humans are—which is judgment—and [where] machines are—processing data. The two need to work hand in hand."**

Identifying and Mitigating Risks in the AI Era

The marriage of AI and risk management demands strong governance frameworks. If not carefully managed, AI can introduce new kinds of risks, from model bias and errors to compliance issues with data privacy. CapTech's experience emphasizes that simply deploying more AI tools is not a solution; in fact, without alignment and oversight, it can create chaos.

One pitfall observed is the proliferation of algorithms or "agents" without a unifying strategy, leading to duplicative efforts and contradictory outputs (two models giving different answers for different sources of truth). Leading institutions are tackling this by establishing clear AI governance bodies and processes. Many have set up AI governance committees or AI centers of excellence that include risk, compliance, IT, and business leaders to evaluate and approve new AI use cases, monitor model performance, and enforce standards for model validation and documentation. This ensures that as AI capabilities scale up, they do so in a controlled, transparent manner.

> **"Governance and innovation must advance together: if they become misaligned, the result can be an erosion of trust in both the AI solutions and the institution's risk management itself."**
>
> **— CapTech Director**

Conclusion

The future of risk management is not defined by replacing people with AI, but by redefining how humans and machines work together. AI enables scale, speed, and pattern recognition that were previously unattainable, while risk professionals provide the judgment, context, and accountability that regulators and stakeholders expect.

As risk and compliance functions move closer to the business, controls become part of how work is designed and executed rather than layered on afterward. This shift enables faster response to regulatory changes, earlier identification of risk, and smoother delivery of products and services. With strong governance, continuous learning, and high AI literacy across teams, risk management evolves from a reactive function into a strategic enabler of the enterprise. Ultimately, organizations that get this right will not simply manage risk more efficiently, but more effectively.

References

[1] CapTech Ventures. (2025). Interviews with senior CapTech leaders specializing in risk management, compliance, AI governance, and data strategy. [Personal interviews].

[2] Hernández de Cos, P. (2024, April 17). Managing AI in banking: Are we ready to cooperate? Bank for International Settlements. https://www.bis.org/speeches/sp240417.pdf

Contributors

This paper reflects the collective experience of CapTech practitioners working with financial institutions across risk management, compliance, data, and AI governance. We are grateful for their contributions and insights: Emerson Beck, Ethan Trifari, Hampton Cobb, Jai Oberoi, Nathanael Jewell, Victoria McMahon, Ashley Smith, Priya Patel

Ashley Smith

Senior Manager

asmith@captechconsulting.com
614.596.2185

Priya Patel

Senior Consultant

ppatel@captechconsulting.com
571.420.7419