



Command and control rewired:

Preparing AFC functions for agentic AI operations and governance



Contents

Executive summary	3
From human-led to agentic AML operations: Orchestrating AML with SRI	4
Building future-ready talent	5
Governance and oversight in an agentic AI world	6
Summary and reinforcing business benefits	7



Executive summary

Agentic AI is transforming anti-money laundering (AML) processes from fragmented, manual-heavy operations into integrated, enterprise-wide orchestration hubs to bring adaptive decision-making, continuous learning, and autonomous coordination across business functions. This evolution represents a structural shift in the way anti-financial crime works. Compliance is no longer a back-office activity but becomes a dynamic, risk-driven capability at the core of enterprise resilience.

Significant implications for leaders in AML

By coordinating AI agents that handle data ingestion, monitoring, escalation, and reporting, organizations can establish a new benchmark for AML operations that delivers measurable risk reduction, operational resilience, and regulatory confidence at unprecedented speed and scale. Crucially, the adaptive intelligence of agents ensures that financial crime controls remain responsive to emerging threats and evolving regulatory expectations without requiring constant manual intervention and updates, freeing compliance professionals to put their significant expertise to use in more high-value areas.

This shift requires investment in scalable data infrastructure, governance frameworks, and explainable outputs. The significant upside is that it positions financial institutions to meet intensifying financial crime risks with agility, transparency, and authority.

“ Agentic AI elevates AML from manual, fragmented processes to enterprise-wide orchestration, delivering faster detection, reduced friction, and stronger regulatory confidence.”



From human-led to agentic AML operations: Orchestrating AML with SRI

The shift from human-driven compliance to agentic orchestration marks a fundamental rewiring of AML operations. Where legacy models rely on analysts manually triaging alerts, piecing together investigations, and overseeing fragmented processes, the likes of Symphony Risk Intelligence (SRI) introduce a coordinated network of AI agents capable of managing these activities with speed, precision, and transparency.

What is SRI?

SRI is a cloud-native, AI-led anti-financial crime platform designed to reduce the total cost of compliance by embedding predictive, generative, and agentic AI into financial crime workflows. By intelligently automating and optimizing end-to-end processes, SRI empowers institutions to manage financial crime risk more efficiently and effectively, reinventing the compliance operating model. Respond rapidly to regulatory change, enhance detection, and drive faster business growth with a next generation risk and compliance platform.

A modular product, SRI allows for automating 50% of the workload with intelligent, scalable AI with build-your-own agent capabilities, freeing up human investigators to put their expertise into higher value areas. Able to integrate and augment from any data source, standardize and centralize data asset management, provide self-serve risk models, and unify operations, it is an evergreen platform that futureproofs your tech stack.

How do the AI agents work?

The AI agents become the operational drivers - ingesting and cleansing data, managing models, triaging alerts, and automating workflows, communicating with one another to enhance efficiency. Human oversight remains paramount with compliance leadership and governance committees setting risk tolerances, calibrating models, validating outputs, and managing escalation protocols. It's all about putting the best tool for the job in the right areas.

As such, final responsibility - whether for suspicious activity reports (SARs), regulatory engagement, or board reporting - remains firmly with compliance officers and executives. This 'human-in-command' approach ensures that regulators retain confidence in the fact that AI augments, but does not displace, institutional accountability.

SRI's agentic capabilities include:

- Adaptive data ingestion: Integrating transactions, sanctions, PEP lists, and external media with continuous updates
- Model management: Deploying, monitoring, and recalibrating models against defined risk appetites and regulatory standards
- Workflow automation: Dynamically prioritizing and triaging alerts, and initiating escalation protocols
- Investigation support: Compiling case files, generating narratives, and correlating behavioral data to support decisions
- Regulatory reporting: Producing auditable, explainable outputs aligned with compliance requirements

Example scenario:

When a suspicious transaction is flagged, SRI agents work together to automatically enrich the alert with relevant internal and external data, correlate behaviors against historical risk profiles, assign risk scores, and generate a full case file. Human investigators then validate the AI's recommendations, exercise judgment on edge cases, and determine escalation or SAR filing. Similarly, regulatory requests for information can be addressed rapidly via orchestrated AI-driven workflows, ensuring completeness, consistency, and traceability.

The model shifts oversight from manually reviewing individual transactions to continuously ensuring that AI-driven operations are reliable. It is built on strong governance, transparent audit trails, and outputs that can be easily explained and understood.

The result? A compliance function that is not only faster and more efficient, but also more risk-aligned, auditable, and resilient.

“With SRI's AI agents, AML evolves with AI taking on the lower-level tasks but keeping humans in control. This approach improves efficiency, auditability, and resilience across the organization.”

Building future-ready talent

It's evident that agentic AI completely reshapes the human role in financial crime compliance. As AI agents assume more responsibility for the likes of data ingestion and monitoring, compliance professionals' roles must also evolve into positions that blend regulatory expertise with technical fluency and governance capabilities. Here are a few examples:

- **From manual review to validation:** Instead of clearing alerts line by line, staff provide oversight and governance of AI outputs, monitor model performance, and ensure every decision remains defensible to regulators
- **From static rules to model stewardship:** Compliance teams must develop proficiency in AI model risk management, bias detection, explainability, and continuous calibration of agentic workflows
- **From isolated expertise to coordinated teamwork:** With AI agents operating across areas like AML, sanctions, and fraud, human oversight needs to cover all these domains to ensure risks are managed consistently and strategies are aligned
- **From execution to ethical judgment:** Humans will increasingly focus on edge cases, escalations that come about due to specific situations, and the ethical implications of outcomes provided by AI.

Training in an agentic AI environment

AI literacy will be vital at all levels of the compliance function. Training programs should explain how agentic systems work, what their limitations are, and how to interpret their outputs. Additionally, staff need to be upskilled in responsible AI, transparency standards, and changing regulatory requirements, using scenario-based exercises to prepare them for real-life situations. These might include AI misclassifying risks, handling new types of cases, or relying on human judgment for sensitive decisions.

With this in mind, institutions will increasingly seek professionals who combine financial crime expertise with data science and governance knowledge. Alongside investigators and policy experts, key roles will include model risk managers, AI governance specialists, and digital ethics officers. It's not unrealistic to imagine that some of these roles might be combined with investigators expected to have multiple areas of expertise. Hiring managers will therefore favor those that can translate AI insights for regulators and senior leadership, design and refine AI-driven workflows, and lead in the cultural adoption of new technologies while maintaining trust. Ultimately, the compliance workforce of the future will be equipped to manage complex ecosystems of autonomous agents, ensure transparency, and reinforce institutional accountability.

“ Compliance roles require a new blend of regulatory expertise, AI fluency, and ethical judgment to help provide transparent, accountable operations.”

Governance and oversight in an agentic AI world

As financial institutions embrace using agentic AI in AML, the primary challenge is balancing automation with regulatory expectations for human oversight and accountability. Decisions cannot be a black box, and regulators have made clear that while AI can drive efficiency and scale, ultimate responsibility cannot and should not be delegated to machines. They must remain explainable with risks actively monitored and accountability staying within human governance structures.

There are several governance models that can be effective in maintaining human accountability:

- **Operational level** – frontline monitoring: AI agents detect unusual transactions and prioritise high-risk alerts. Human compliance analysts then validate these findings and determine escalation
- **Tactical level** – model governance and risk management: AI agents propose model adjustments or simulate scenarios, but all changes require human sign-off with workers providing continuous validation, drift monitoring, and override mechanisms to ensure ongoing defensibility
- **Strategic level** – board and executive oversight: Senior leadership reviews AI-driven dashboards, trend analysis, and audit outcomes. Human commentary provides context, limitations, and remediation plans, reinforcing that accountability resides with the institution, not the technology

Strong governance relies on layered mechanisms that reinforce trust in AI systems. These include centralized audit logs that capture every agentic decision chain, continuous model validation with independent reviews, and clearly defined escalation paths that mandate human intervention in high-risk cases. Stress-testing governance models through simulated attacks and maintaining transparent reporting to senior management and regulators are also essential. By integrating these controls, organizations ensure that AI augments, but never replaces, human responsibility. Compliance leaders, executives, and Boards retain ultimate accountability while making the best use of AI to enhance scale, precision, and predictive insights.

This approach makes sure that governance is both regulator friendly and operationally resilient, embedding trust, transparency, and control into the heart of agentic AML operations.

“ Strong governance in agentic AI ensures decisions remain explainable, auditable, and defensible, embedding regulator trust while keeping accountability squarely with the institution.”

Summary and reinforcing business benefits

The transition to using agentic AI in AML represents more than a technological upgrade. It is a strategic shift that redefines how institutions manage financial crime risk. By moving from fragmented, manual-heavy processes to coordinated AI-driven operations, firms can strengthen their control environments while achieving significant efficiency gains.

Transitioning to a more tech-focused approach using products like SRI doesn't need to happen all in one go and requires hitting some key steps to ease the process:

- Establish governance foundations: Define clear ownership of data quality, model performance, and escalation protocols before introducing agentic systems. Ensure oversight structures, auditability, and explainability are built in from the start
- Incremental adoption: Begin with high-value workflows such as alert triage, case file compilation, or sanctions screening before slowly expanding. Each stage should deliver measurable improvements in accuracy, efficiency, and risk coverage
- Invest in skills and culture: Equip compliance teams with AI knowledge, model oversight, and governance skills to supervise agentic operations. Recruiting hybrid profiles that combine compliance expertise with data and risk management capabilities will help to accelerate adoption
- Embed continuous assurance: Implement monitoring, validation, and simulated attacks (red-teaming) to ensure models remain aligned with regulatory standards and evolving risks

Business benefits

Taking such an approach and implementing agentic AI in compliance operations offers many benefits, including enhanced risk detection and mitigation as the technology continuously learns, identifies emerging threats, and adapts faster than human-only teams. Alongside this, automation of routine tasks boosts operational efficiency by reducing false positives, lowering costs, and enabling investigators to manage higher volumes without compromising oversight.

It doesn't stop there. Transparent audit trails, layered governance, and human-in-command accountability instill regulatory confidence and provide defensibility with supervisors and auditors. Furthermore, by relieving staff of repetitive tasks, compliance professionals can be strategically redeployed to more complex and value-added areas, helping to shape anti-financial crime strategies. Together, these benefits not only improve daily operations but also contribute to stronger long-term strategic outcomes.

Agentic AI is not simply about doing today's work faster; it is about positioning compliance as a dynamic, capability offering advantages across the enterprise. Institutions that embrace this evolution by using the likes of SRI will achieve greater resilience, agility, and regulatory trust. Those that delay risk being left with outdated, resource-intensive models that cannot keep pace with financial crime's complexity.



About Symphony Risk Intelligence

Symphony Risk Intelligence (SRI) is an AI-native FinCrime platform that's designed to help financial institutions move from reactive to proactive risk management. It centralises and optimizes every aspect of compliance operations, from data ingestion and model management to investigation and reporting.

Every layer of the cloud-native platform has been built to leverage the latest predictive, generative, and agentic AI technology to enable greater automation, improved detection effectiveness, and significant operational efficiency savings.

SRI represents a reimagining of the compliance operating model, built to eliminate legacy inefficiencies, unify fragmented systems, and turn compliance into a driver of business growth. With the ability to automate a wide range of processes and operational tasks, SRI halves the compliance workload while equipping teams with powerful AI-driven tools to accelerate workflows and maximise performance.

SRI allows financial institutions to begin their transformational journey to a next generation, evergreen infrastructure that futureproofs payment ecosystems and prepares businesses to scale with speed, efficiency, and regulatory resilience.



symphonyai.com

About SymphonyAI

SymphonyAI, 2024 Microsoft Partner of the Year for Business Transformation– AI Innovation, is building the leading enterprise AI SaaS company for digital transformation across the most critical and resilient growth verticals, including retail, consumer packaged goods, finance, manufacturing, media, and IT/enterprise service management. SymphonyAI verticals have many leading enterprises as clients. Since its founding in 2017, SymphonyAI has grown rapidly to 3,000 talented leaders, data scientists, and other professionals. SymphonyAI is an SAI Group company, backed by a \$1 billion commitment from Dr. Romesh Wadhvani, a successful entrepreneur and philanthropist.

SymphonyAI provides AI-focused SaaS solutions developed from the ground up for fighting financial crime. An end-to-end offering compiled from more than 25 years of knowledge in tackling money laundering and fraud, our award-winning product ecosystem comprises an innovative and modern approach to financial crime investigation using world-leading agentic, predictive, and gen AI. From KYC/CDD and transaction monitoring through to payments fraud, entity resolution, and sanctions, PEP, and adverse media screening, SymphonyAI is a trusted name in finance technology with more than a third of the world's top 100 banks enjoying the power, efficiency, effectiveness and productivity that our solutions provide.


Authors



Craig Robertson

Financial Crime &
Compliance SME, APAC


craig.robertson@symphonyai.com



Charmian Simmons

Financial Crime &
Compliance SME, EMEA

charmian.simmons@symphonyai.com



Elizabeth Callan

Financial Crime &
Compliance SME, NA

elizabeth.callan@symphonyai.com