WHITE PAPER

SymphonyAI

LEADERS GUIDE

Agentic AI & Embedded Risk Intelligence

Transforming FinCrime Compliance Platforms

A Strategic Guide for Compliance & Financial Crime Leaders
Why Act Now | What to Prioritize | How to Lead the Transition

March 2026

Executive Summary

Agentic AI and embedded risk intelligence are how the compliance function moves from reactive obligation to proactive, intelligence-led advantage.

Financial crime has grown more sophisticated, faster-moving and harder to detect. Criminal networks exploit layered digital payment rails, cryptocurrency, trade finance and shell company structures simultaneously. Sanctions evasion demands real-time screening and contextual reasoning. Fraud typologies evolve in days. At the same time, the compliance function faces intensifying cost pressure, rising alert volumes and an accelerating regulatory agenda across FATF, AMLA, 6AMLD and FinCEN modernization initiatives.

Agentic AI is autonomous, goal-directed intelligence capable of reasoning, planning and executing multi-step tasks, and when combined with embedded risk intelligence is fundamentally changing what is possible. It is a structural reimagining of how compliance operates: from reactive detection to proactive prevention, from manual triage to intelligent orchestration, from compliance-as-cost to compliance-as-competitive-advantage.

The industry signal is clear. According to the SymphonyAI FinCrime Frontier 2025-26 survey of over 250 senior compliance professionals globally, 80% of institutions plan to innovate with AI in the next 12-18 months, and more than 70% still rely on manual operations for at least half of their AML activities. The opportunity to differentiate through purposeful AI adoption has never been greater.

Source: SymphonyAI FinCrime Frontier Survey 2025–26 (symphonyai.com)

This guide answers four questions for leaders who need to move from understanding to action:

- **WHY** - The forces driving urgency across regulatory, financial and operational dimensions
- **WHEN** - The signals that indicate readiness, and the compounding advantage of moving now
- **NOW** - Immediate entry points that build momentum, proof of value, and future-ready capability
- **HOW** - A practical roadmap covering people, process, technology and governance

Bottom Line Up Front

Organizations that successfully embed agentic AI into their FinCrime compliance platforms can expect to reduce false positive rates by 40-70%, cut SAR filing cycle times by over 50%, and improve analyst productivity by 30-60%, achieving a compelling return on investment within 8–18 months while materially strengthening their regulatory standing.

Note: This guide is designed for compliance and financial crime leaders who need to build the strategic case for agentic AI and lead the transition within their organization. It focuses on the why, when, and how of adoption — covering the business case, governance foundations, and implementation roadmap.

For a deeper operational walkthrough — including how AI agents work in practice, the financial crime agent ecosystem, real-world workflow scenarios across seven industry use cases, and the architecture underpinning Always-on Compliance — see the companion guide: **Agentic AI for Financial Crime Risk Management: A Guide for Risk and Compliance Professionals.**

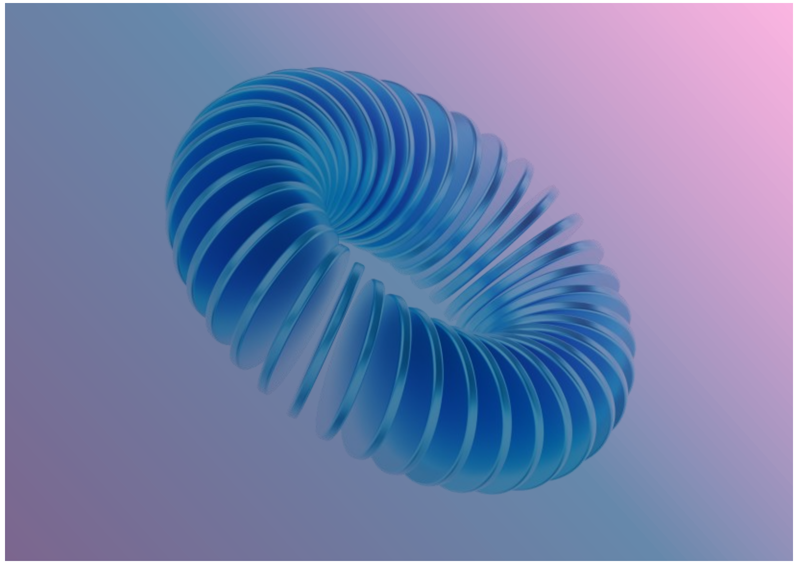

Introduction

The Case for Proactive Compliance

Financial crime does not pause. Compliance that does is already behind

The global financial system faces a level of criminal complexity that demands a fundamentally proactive response. The UN Office on Drugs and Crime estimates that between $800 billion and $2 trillion is laundered annually — equivalent to 2–5% of global GDP. Regulators across EMEA and North America are simultaneously expanding expectations: demanding greater transparency, stronger governance, and demonstrable effectiveness rather than mere process compliance.

Source: UNODC, Money Laundering Overview - unodc.org/unodc/en/money-laundering/overview.html

What this moment demands is a shift from reactive compliance, responding to what has already occurred, to proactive, intelligence-led compliance that anticipates, detects and disrupts financial crime before harm compounds. This requires technology that can reason, learn and act at the pace of the threat. Agentic AI and embedded risk intelligence are precisely that capability.

At the centre of this evolution is Always-on Compliance: rather than periodic reviews, batch-processed alerts and scheduled model updates, it means continuous monitoring, real-time risk evaluation and adaptive detection that operates around the clock, across every product, channel and customer relationship. When a new typology emerges, defences respond in days, not at the next model refresh.

Always-on Compliance Defined

Always on Compliance™ is a modern approach where compliance programs stay current by design as regulations, threats, and business activity change. SymphonyAI enables Always-on Compliance with Symphony Risk Intelligence (SRI), an AI native risk and compliance platform, purpose built for financial crime compliance. SRI combines unified risk intelligence with agent native orchestration to turn signals into decisions and decisions into action, automating repetitive work, reducing alert noise, and coordinating data, detection, and investigation as conditions change. A unified governance framework ensures transparency and control, while SRI's modular architecture supports flexible deployment- end to end, by domain, or to augment existing systems - without forcing a big bang rip and replace.

The success of this investment must be measurable. Leaders who embed agentic AI and embedded intelligence should track a clear set of performance outcomes. The table below defines what 'good looks like', and is grounded in data from the SymphonyAI FinCrime Frontier 2025-26 survey and industry benchmarks.

KPI	Baseline (Typical)	Target with Agentic AI	Timeframe
Alert false positive rate	>70% of institutions above 25% FP rate*	Targeted 40–70% reduction	3–9 months
SAR filing cycle time	10–20 days average	50–65% reduction	6–12 months
Analyst case throughput	Baseline volume	30–60% increase per analyst	6–12 months
Typology detection lag	Model refresh cycles (weeks)	Days to operationalize	Ongoing
Alert-to-SAR conversion	54% report <5% conversion*	Meaningful improvement in yield	6–12 months
Regulatory audit findings	Variable	Trend reduction over 2 cycles	12–18 months

* SymphonyAI FinCrime Frontier Survey 2025–26: 70%+ of institutions report false positive rates above 25%; 54% report alert-to-SAR conversion below 5%. Benchmark targets informed by ACAMS Technology in AML Survey 2023 and McKinsey & Company, 'The AI-powered compliance advantage', 2023.

SECTION 1

WHY: The Imperative for Change

1.1 A Compliance Landscape That Rewards the Proactive

Financial crime has grown more complex, more cross-border and more technology-enabled. Compliance programmes built on static rule sets and periodic review cycles remain valuable for the controls they established, and the evidence now shows that the pace and sophistication of the threat has overtaken what those controls can reliably detect alone.

The SymphonyAI FinCrime Frontier 2025-26 survey found that alert fatigue affects 30.8% of institutions and that poor data quality and integration remains the single biggest drag on AML efficiency, cited by nearly two-thirds (65.6%) of respondents. More than half (54%) report that fewer than 5% of alerts lead to case escalation or SAR filing. These are not signs of a broken function, they are signs of a function ready for intelligent enhancement.

Source: SymphonyAI FinCrime Frontier Survey 2025–26 — symphonyai.com

Conventional Compliance Model	Always-on Compliance Model
Rule-based detection with static thresholds	Dynamic, self-learning risk models that adapt to emerging typologies
Alert queues managed by human analysts	Intelligent triage — agents prioritize, contextualize and pre-investigate
Siloed customer, transaction and network data	Unified, entity-centric risk intelligence across all data sources and domains
Manual SAR drafting and narrative writing	AI-assisted narrative generation with analyst review and sign-off
Periodic model refresh cycles	Continuous, autonomous monitoring with real-time anomaly detection
Compliance as a cost centre	Compliance as a driver of smarter business growth

1.2 The Regulatory Direction Is Clear

Regulators globally - the FCA, FinCEN, FATF, the newly created EU Anti-Money Laundering Authority (AMLA) and others - are expanding expectations and demanding demonstrable effectiveness. The SymphonyAI survey found a notably positive industry response: 58% of respondents view upcoming regulations as an overall opportunity rather than a burden, and the regulatory landscape and volume of change ranks as the second-largest concern after the growing sophistication of criminal tactics (41.8%).

Supervisors increasingly focus on compliance effectiveness, not just compliance activity. A firm generating high volumes of marginal SARs is viewed less favourably than one producing fewer, higher-quality, intelligence-rich reports. Agentic AI enables the latter. Critically, regulators also expect institutions to engage proactively: the survey found that only 29.1% of firms maintain structured, regular engagement with supervisors, those that do are better positioned to shape the frameworks governing next-generation compliance.

Sources: SymphonyAI FinCrime Frontier Survey 2025–26; FCA Innovation Pathways — fca.org.uk/innovation; FATF Recommendation 15 - fatf-gafi.org; EU AMLA — eur-lex.europa.eu

1.3 The Business Case Is Well Established

Early adopters in larger banking, global payments and specialist FinTech have moved well beyond pilots. Compliance technology investment, particularly in AI-enhanced monitoring and investigation platforms, delivers measurable returns. The SymphonyAI survey confirms that the top five KPIs institutions now track for AI ROI are: false positive rate reduction (27%), case resolution time (23%), detection accuracy (18%), automation gains (15%) and customer onboarding experience (9%) - together accounting for over 90% of all KPIs cited.

Source: SymphonyAI FinCrime Frontier Survey 2025–26 - symphonyai.com

Value Dimension	Indicative Impact
Alert false positive reduction	40-70% reduction targeted; frees significant analyst capacity
SAR filing cycle time	50-65% reduction from alert to filed report
Analyst throughput	30-60% increase in cases reviewed per analyst per day
Typology detection	Days rather than model-refresh cycles to operationalize new patterns
Staff retention	Analysts focus on complex judgement work, reducing attrition
Customer experience	Fewer false declines and reduced CDD friction for low-risk customers

Benchmarks: ACAMS Technology in AML Survey 2023 (acams.org); McKinsey & Company, 'The AI-powered compliance advantage', 2023 (mckinsey.com); Deloitte, 'AI in Financial Crime Compliance', 2024 (deloitte.com)

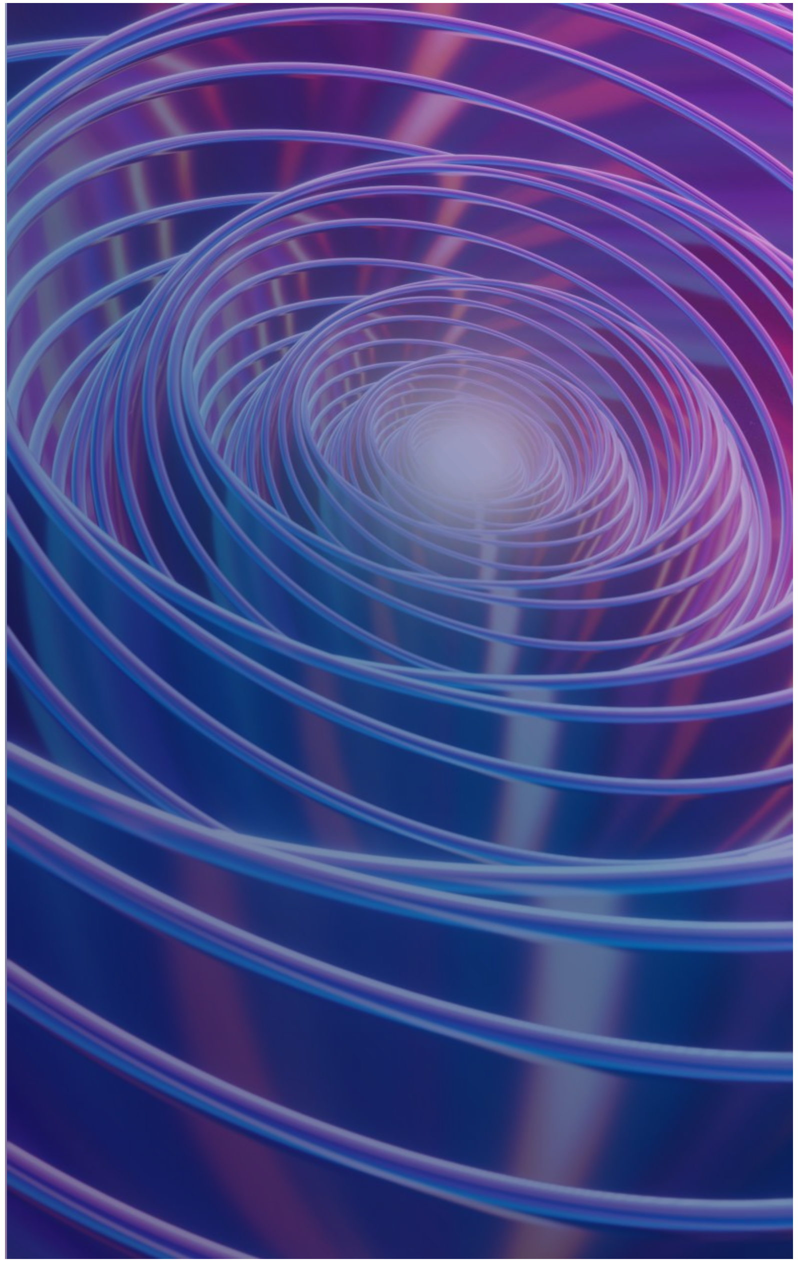

SECTION 2

WHEN: Recognizing the Right Moment

2.1 Signals That Your Organization Is Ready

The organizations best positioned for successful adoption share common characteristics. The window for thoughtful, planned adoption is open now, those who move purposefully in the next 12 months will establish a performance and governance standard that compounds over time. The SymphonyAI survey shows 80% of institutions plan to innovate with AI in the next 12-18 months, with nearly nine in ten already implementing or planning AI by the end of 2026. Early movers gain compounding advantage: better models, more capable analysts, deeper regulatory trust.

The following signals indicate your organization has reached the threshold for meaningful adoption:

- Your alert-to-SAR conversion rate sits below 5%:- the majority position in the survey, indicating capacity absorbed by low-yield investigations
- Analyst attrition or burnout is measurable and linked to high-volume, repetitive review work
- Your risk models have not been materially updated in 12+ months and are not capturing recently observed typologies
- Customer onboarding friction is driving measurable drop-off for legitimate customers due to over-escalation in CDD
- Regulatory feedback has referenced the quality, timeliness or consistency of your suspicious activity reporting
- Your data infrastructure is sufficiently consolidated:- or on a roadmap to consolidation, to support cross-system intelligence queries

Leadership Diagnostic

Ask your Head of Financial Crime: 'If a new layering typology emerged today targeting our product set, how long before our current models would reliably detect it?' If the answer is measured in weeks or model refresh cycles rather than days, the operational case for agentic AI is already compelling.

2.2 The Compounding Advantage of Acting Now

Timing matters in AI adoption because the benefits compound. Agentic systems learn from feedback, improve with use and generate organizational capability - in people, process and data quality - that deepens over time.

Every month of purposeful deployment builds an advantage a later adopter cannot quickly replicate.

Consider what builds during a structured adoption programme: model performance improves as agents accumulate labelled analyst decisions; analyst capability grows as practitioners develop fluency in interrogating and challenging AI outputs; regulatory trust develops through demonstrated governance; data quality improves as embedding AI surfaces and forces resolution of underlying gaps; and the innovation pipeline opens as foundational capability enables more ambitious use cases, custom agent development, cross-institution intelligence sharing, predictive risk modelling.

Conversely, deferral has its own compounding cost. Technology debt deepens as legacy platforms diverge from next-generation infrastructure. Talented analysts gravitate toward organizations with more sophisticated tooling. The survey found that 72% of compliance teams have never conducted a formal ROI assessment for compliance technology - those who begin now will build the measurement discipline that justifies and accelerates every subsequent investment.

SECTION 3

NOW: Build Today, Lead Tomorrow

Starting well matters more than starting perfectly. Each entry point below builds a capability that enables the next → from doing compliance better today to building the Always-on compliance function of tomorrow.

3.1 Immediate Entry Points: From Reduction to Intelligence

Entry Point 1 - Now: Intelligent Alert Triage

Deploy agentic triage models that evaluate alerts in full context - customer history, peer behavior, product risk, channel signals and network relationships - before an analyst touches them. False positive rates can reduce by 30–50% within months, releasing analyst time for genuinely complex investigation. The survey confirms alert fatigue (30.8%) and poor data quality (65.6%) as the dominant AML efficiency blockers, intelligent triage directly addresses both.

Future state enabled: triage data becomes the training signal for next-generation detection models. You move from reducing noise to continuously sharpening signal.

Entry Point 2 - Now: AI-Assisted SAR Narrative Generation

Use embedded AI to draft SAR narratives pre-populated with relevant facts, typology references and risk indicators, ready for analyst review and sign-off. The analyst retains full ownership of quality and judgement. Cycle time reductions of 40-60% are consistently achievable, and filing quality improves measurably. The survey identifies SAR drafting as one of the most manually intensive areas in the compliance lifecycle.

Future state enabled: consistently structured, high-quality SAR outputs feed more valuable intelligence back to FIUs and law enforcement, strengthening your contribution to the financial intelligence ecosystem.

Entry Point 3 - Soon: Dynamic Customer Risk Scoring

Move from static, periodic customer risk scores to dynamic, event-driven risk profiles incorporating real-time behavioral signals, network relationships, adverse media and cross-product activity. Agentic AI continuously re-evaluates risk as new information becomes available, reducing both under-flagging of genuine risk and over-restriction of low-risk customers.

Future state enabled: dynamic scoring powers predictive risk management. Rather than responding to indicators after they materialize, your compliance function begins identifying elevated-risk trajectories before harm occurs.

Entry Point 4 - Soon: Dynamic Typology Ingestion

Embed intelligence feeds from FATF, national FIUs, the Egmont Group and law enforcement advisories, directly into your detection model layer so that new typologies are operationalized within days rather than model refresh cycles. Agentic AI reasons about how an emerging typology maps to your customer base and product set in near real time.

Future state enabled: your compliance function shifts from a consumer of threat intelligence to an active contributor, sharing enriched, AI-generated signals back into public-private partnerships and FIU networks.

Entry Point 5 - Build for the Future: Custom Agent Development

As SaaS-deployed capabilities demonstrate value and your team builds AI fluency, the next horizon is developing proprietary agents tailored to your specific risk profile, product set and customer base. The SymphonyAI FinCrime Frontier survey confirms the industry's direction: the next phase moves beyond rule-based automation toward agentic AI - intelligent agents that make decisions, plan and learn from previous experiences, collaborating with one another to complete complex compliance workflows end to end.

This is the destination that day-one investments build toward. Organizations that start with intelligent triage, earn analyst trust, develop governance frameworks and invest in data quality are positioned to build and sustain proprietary agentic capability.

The Capability Sequence

Reduce (triage and SAR efficiency) → **Enrich** (dynamic scoring and typology intelligence) → **Anticipate** (predictive risk) → **Innovate** (custom agents and ecosystem contribution)

3.2 Governance From Day One

Regulatory trust in AI-driven compliance depends above all on explainability and accountability. The SymphonyAI survey found that only 17.2% of institutions have a fully defined and operational AI governance framework, while 22.8% have no governance structure at all. This is the single greatest differentiator between leaders and laggards as AI scales. Governance infrastructure built in from the start (not retrofitted) is what makes AI-powered compliance durable.

- Appoint a named AI Model Risk Owner with accountability for performance, bias monitoring and model drift
- Define Human-in-the-Loop requirements for each agentic workflow:- the survey found 69% of institutions rely on human reviewers to validate alert relevance, reflecting the continued importance of analyst judgement
- Establish a Model Validation framework aligned to SR 11-7 or your relevant regulatory model risk standard before any system enters production
- Create an AI Compliance Audit Trail:- every agent decision must be logged with sufficient context for supervisory review
- Define an Escalation Protocol:- when an agentic system encounters scenarios outside its training distribution, it escalates rather than decides

Reference: Federal Reserve SR 11-7, Guidance on Model Risk Management - federalreserve.gov/supervisionreg/srletters/sr1107.htm

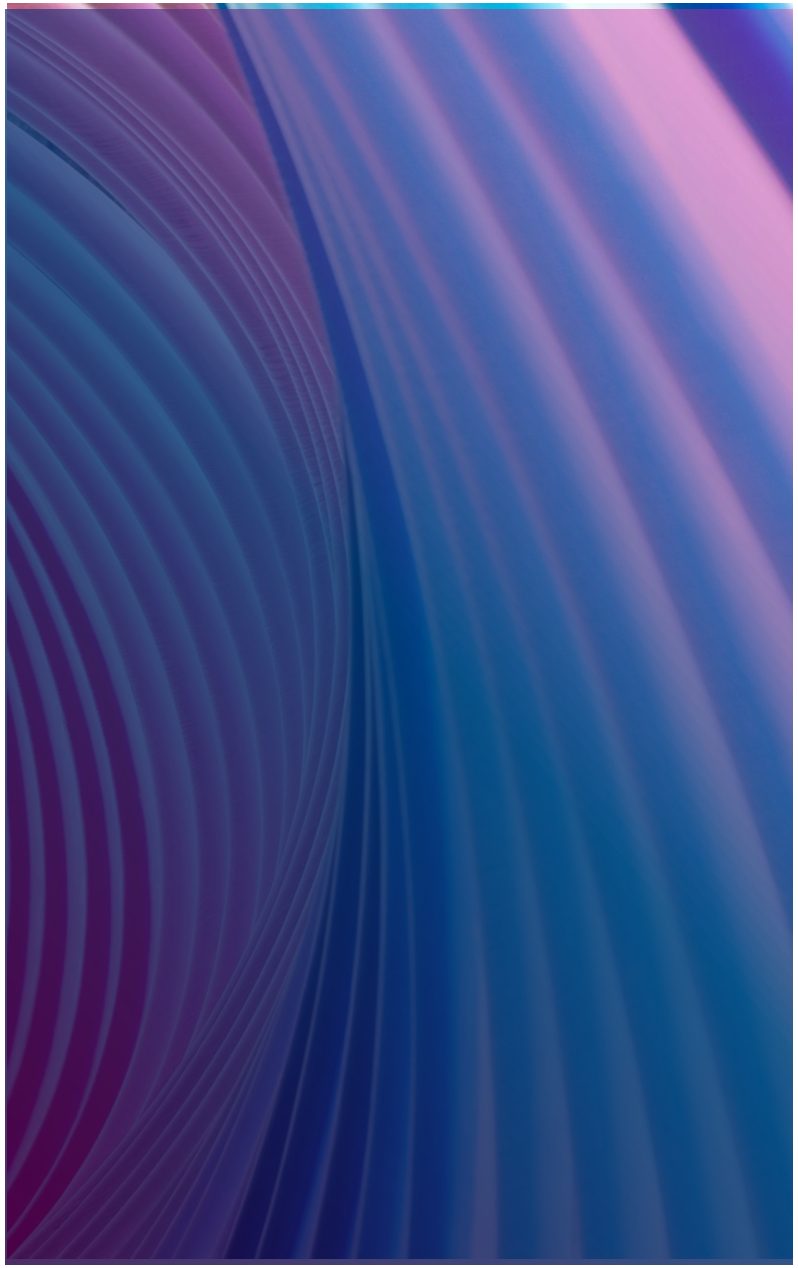

SECTION 4

HOW: Operationalizing Embedded Risk Intelligence

4.1 The Four Pillars of Operationalization

These four pillars are how organizations realise the five value outcomes of Always-on Compliance™ - closing coverage gaps, reducing effort and cost, unifying risk intelligence, ensuring regulator-ready AI, and building strategic agility. Each pillar is required; neglect any one and the value themes remain aspirational rather than operational.

Successfully embedding agentic AI is an organizational transformation, not a technology procurement. It requires investment across four pillars simultaneously. Neglecting any one pillar weakens the others.

The four pillars of operationalization

People

Build AI literacy and reskill compliance teams

New roles: AI Specialist, Typology Analyst, Model Reviewer

Process

Redesign workflows around intelligence, not automation

Build feedback loops so analyst decisions continuously improve agent behavior

Technology

Unified data foundation as prerequisite

Modular agents, real-time feeds, traceable inference, cloud-native scale

Governance

Model Risk Owner, validation framework, audit trail

Human-in-the-loop at every agentic workflow; escalation protocols defined

Together, these pillars realize the five value outcomes of Always-on Compliance™

Neglecting any one pillar weakens the others – all four must be invested in simultaneously

Pillar 1 - People: Building a Compliance Intelligence Workforce

Educating compliance professionals on what agentic AI changes are, will do and how it elevates the value of their expertise is key. The analyst of the future is a risk intelligence practitioner: interrogating AI outputs, exercising judgement on complex scenarios, communicating AI-driven findings to regulators and continuously improving the models they work alongside. The SymphonyAI survey found 56.8% of institutions plan to upskill existing staff as their primary investment, reinforcing that human capability remains central.

- Reskilling programmes developing AI literacy, critical evaluation of model outputs and data interpretation skills
- New role definitions:- Compliance AI Specialist, Typology Intelligence Analyst, Model Performance Reviewer
- Senior capability to create anamed AI Model Risk Owner with board-level reporting accountability
- Culture that positions AI as a tool enhancing professional judgement - critical for analyst buy-in and regulatory confidence

Pillar 2 - Process: Redesigning Workflows Around Intelligence

Most compliance workflows were designed around alert volumes and manual handoffs. Agentic AI makes a different architecture possible and necessary. Redesign workflows around the intelligence AI generates, not simply to automate existing steps. Move from sequential, linear case management to parallel, multi-signal intelligence assembly. Redesign the SAR lifecycle around AI-generated structured analysis with human quality assurance at key decision gates. Build feedback loops into every workflow so analyst decisions continuously improve agent behaviour.

Pillar 3 - Technology: The Embedded Intelligence Architecture

Architecture principles for embedded risk intelligence include a consolidated data foundation as prerequisite. The Symphony AI survey found poor data quality and integration cited by 65.6% of respondents as the top AML efficiency blocker.

Build modular agents for specific, bounded tasks that compose into workflows. Consume and operationalize external typology, sanctions and adverse media feeds in near real time. Ensure every inference is traceable to input features and model logic. Scale cloud-natively for event-driven volume surges.

Pillar 4 - Governance: Making AI Trustworthy by Design

Governance is the foundation on which all other investment rests. Formal validation, stress testing and ongoing monitoring against defined thresholds for every deployed model. Regular bias and fairness audits. Data lineage, quality standards and access controls. Proactive communication with supervisors - the survey confirms that institutions engaging regularly with regulators on AI are better positioned to help shape evolving frameworks. Quarterly Board reporting on model performance, incidents and the innovation pipeline.

4.2 The Investment Case: ROI and Strategic Value

Investment in agentic AI and embedded risk intelligence should be framed as a strategic business decision. The financial return is real and typically achievable within 8 - 18 months of production deployment, depending on data estate maturity, deployment model (SaaS or build), and initial use case scope.

Investment Area	Expected Return / Strategic Benefit
AI-enhanced alert triage	40–70% false positive reduction; 30%+ analyst capacity release
Embedded SAR assistance	50–65% SAR cycle time reduction; improved filing quality
Dynamic customer risk scoring	Reduced CDD over-escalation; improved customer experience
Typology intelligence ingestion	Days rather than weeks to operationalize new crime patterns
Analyst reskilling programme	Reduced attrition; stronger talent pipeline for complex work
Regulatory engagement investment	Reduced supervisory friction; enhanced standing as responsible innovator

Strategic value beyond the financial model: firms with demonstrably effective AI-driven compliance programmes attract more constructive supervisory relationships; in wholesale and correspondent banking, the quality of your AML programme is increasingly a factor in partner selection; a compliance function demonstrating intelligent, data-driven risk management gives the Board assurance to pursue growth strategies with reduced reputational anxiety.

4.3 The Adoption Roadmap: 8 - 18 Months to Full Advantage

The timeline to full operational advantage ranges from 8 to 18 months, driven by three variables: data estate maturity; SaaS deployment (faster) vs. proprietary build (more tailored); and initial agent scope. The roadmap below reflects this range with decision points that allow your programme to move at the pace your readiness supports.

Phase	Key Actions & Milestones	Timeline
Phase 1: Foundation	Data estate audit; baseline KPI measurement; governance framework design; SaaS vendor evaluation; Board mandate secured; AI Model Risk Owner appointed	Months 1-2
Phase 2: Pilot	Intelligent alert triage (SaaS) deployed; AI-assisted SAR narrative generation live; Model Risk Management framework established; 90-day success metrics defined; first regulatory engagement on AI program	Months 2-5
Phase 3: Evaluate & Scale	Pilot outcomes measured against KPIs; validated use cases expanded; dynamic customer risk scoring in production; reskilling program active; data quality program underway. Fast-track: SaaS + consolidated data estate achieves Phase 3 KPI improvement by Month 8	Months 5-10
Phase 4: Embed & Govern	Full production across core workflows; typology intelligence live; Board AI performance reporting established; regulator briefed; analyst feedback loops driving model improvement	Months 10-15
Phase 5: Innovate	Custom agent development initiated for proprietary risk scenarios; customer-built agents capability explored; Always-on Compliance operating state achieved; innovation pipeline defined	Months 15-18+

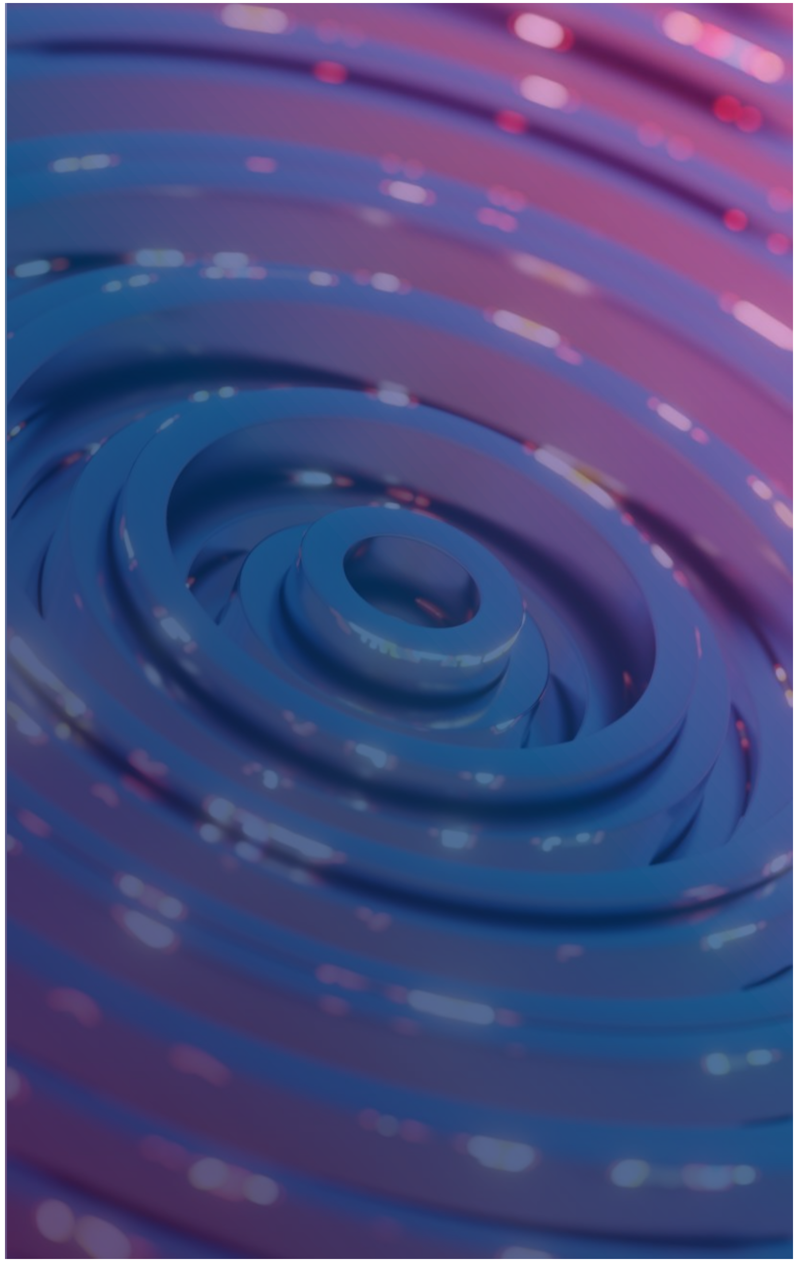

A Word to the Leader

This guide is not a technology manual. It is a call to leadership.

The [SymphonyAI FinCrime Frontier 2025-26 survey] confirms the industry is at an inflection point. Nearly 80% of institutions plan to innovate with AI in the next 12-18 months. The question has shifted from whether to adopt, to how purposefully and how well. The organizations that move with deliberate strategy - investing in people, governance and technology in equal measure - will set the standard that others follow.

Agentic AI and embedded risk intelligence will not solve every financial crime challenge. Criminal networks will continue to adapt. Typologies will continue to evolve. The regulatory bar will continue to rise. But institutions with intelligent, adaptive, well-governed compliance platforms will be substantially better placed to meet those challenges, and to demonstrate to supervisors, boards, customers and counterparties that they take the integrity of the financial system seriously.

The cost of getting this right is real. The return on getting it right is greater, more durable, and more strategically significant than any other investment available to the compliance function today.

The Leader's Commitment

Commit to three things simultaneously: invest in the technology, invest in the people, and invest in the governance. The compliance function of the future is built on all three — with your leadership visible and sustained throughout.

The capabilities described in this guide - intelligent triage, embedded risk intelligence, agentic orchestration, and always-on governance - are delivered today through Symphony Risk Intelligence (SRI). To explore how SRI can accelerate your organization's journey to Always-on Compliance™, visit symphonyai.com.

Quick Reference: Key Definitions

Term	Definition
Agentic AI	AI systems capable of autonomous goal-directed behaviour — planning, reasoning and executing multi-step tasks without continuous human instruction
Always-on Compliance	An operating state where risk intelligence is continuously generated, applied and refined across monitoring, investigation and reporting — without dependency on manual cycles or scheduled reviews
Embedded Risk Intelligence	The integration of AI-generated risk signals and insights directly into compliance workflows, rather than as separate analytical outputs
Human-in-the-Loop	Governance principle ensuring human judgement is a required input at defined points in any AI-driven decision process
Model Drift	Degradation of an AI model's predictive performance over time as real-world patterns evolve beyond those it was trained on
Typology	A documented pattern or method of financial crime — regulators and FIUs publish typology guidance that compliance systems should operationalize in near real time
Model Risk Management	The formal framework for validating, monitoring and governing AI and statistical models — SR 11-7 is the global benchmark standard

Key Sources & Further Reading

- SymphonyAI FinCrime Frontier Survey 2025-26: symphonyai.com
- UNODC - Money Laundering Overview: unodc.org/unodc/en/money-laundering/overview.html
- FATF - Recommendation 15 on New Technologies: fatf-gafi.org
- FCA - Innovation Pathways and SupTech Strategy: fca.org.uk/innovation
- EU AMLA - AML/CFT Framework: eur-lex.europa.eu
- FinCEN - SAR Filing Statistics and FinCEN Exchange: fincen.gov
- Egmont Group of Financial Intelligence Units: egmontgroup.org
- Federal Reserve SR 11-7 - Model Risk Management: federalreserve.gov/supervisionreg/srletters/sr1107.htm
- ACAMS - Technology Adoption in AML Survey 2023: acams.org
- McKinsey & Company - The AI-powered compliance advantage 2023: mckinsey.com
- Deloitte - AI in Financial Crime Compliance 2024: deloitte.com

About SymphonyAI

SymphonyAI, 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 2000 leading enterprises as clients globally. 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 Wadhwani, a successful entrepreneur and philanthropist.

SymphonyAI Financial Services provide AI-focused SaaS solutions developed from the ground up for compliance and financial crime risk management. 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 200 financial institutions enjoying the power, efficiency, effectiveness and productivity that our solutions provide.

Get in touch with us to explore how agentic AI can enhance your risk-based approach and power effectiveness throughout your financial crimes risk management program.

Email us directly at
connect.fs@symphonyai.com
Visit www.symphonyai.com

SymphonyAI