WHITE PAPER

SymphonyAI

The sanctions intelligence shift

Designing agentic, explainable screening under human-in-command governance

Contents

Executive summary

Sanctions compliance is entering a structural transformation.

Financial institutions face escalating geopolitical volatility, diverging jurisdictional programs, increasing designation volumes, convergence of AML and sanctions programs, and a decisive regulatory shift toward demonstrable effectiveness, not technical compliance alone.

Screening programs must now process structured and unstructured data, monitor dynamic ownership networks, and detect sophisticated evasion tactics – all while maintaining zero tolerance for missed matches.

Sanctions compliance can no longer operate in episodic cycles of list updates, rule recalibration, and retrospective tuning.

This shift enables what can be described as “Always-on” sanctions intelligence where exposure is continuously reassessed, rather than periodically reviewed.

In a world of escalating designation volumes, secondary sanctions, and complex circumvention strategies, institutions require a sanctions intelligence capability that continuously ingests regulatory updates, dynamically reassesses risk exposure, and adapts screening logic in near real time.

In sanctions screening, 99% accuracy is failure.

Financial institutions cannot afford black-box automation. They require AI systems that:

- Enhance detection without sacrificing defensibility
- Integrate into existing screening platforms
- Preserve human accountability
- Provide full auditability
- Withstand regulatory scrutiny
- Enable detection of sanctions evasion

Agentic AI represents the next evolution in sanctions compliance – not as autonomous decision-making, but as **orchestrated, explainable investigation intelligence** embedded within enterprise governance structures.

This paper examines how to:

- Close the gap between regulatory change and operational response
- Detect complex, network-based sanctions evasion beyond simple name matching
- Embed explainable AI within mission-critical screening environments
- Satisfy emerging AI governance expectations across jurisdictions
- Maintain human accountability while increasing investigative efficiency
- Transition to agentic screening capabilities safely and incrementally

The future of sanctions compliance is not autonomous AI.

It is AI-capable compliance that is explainable, governed, and human-in-command.

The sanctions inflection point

The current AML and sanctions outlook makes one reality clear: the landscape is growing more complex, and it is no longer about volume alone. It is about velocity, complexity, and effectiveness.
Key shifts include:

- Persistent regulatory scrutiny across jurisdictions
- Convergence of AML and sanctions functions
- Emphasis on risk-based effectiveness
- Growing use of AI in triage and investigations

Sanctions programs are no longer measured solely by process adherence. They are evaluated on their ability to:

- Screen structured and unstructured data
- Detect complex ownership and evasion networks
- Integrate open-source intelligence
- Automate alert triage with explainable AI

Traditional sanctions screening has historically operated on fragmented data sets and static name-matching logic, limiting contextual risk assessment.

While effective in an era of simpler designation regimes, this approach struggles to capture indirect exposure, evolving ownership structures, and increasingly sophisticated evasion tactics.

Sanctions screening is no longer a static matching exercise. It is a dynamic, intelligence-led process. Legacy rule-based systems struggle under these demands. They cannot keep pace with transliteration complexity, multi-jurisdictional sanctions divergence, real-time payments, and network-based circumvention strategies.

High false-positive volumes are not simply an operational burden, but often a symptom of screening models that lack contextual intelligence. When screening relies primarily on string comparison rather than multi-dimensional risk assessment, alert volumes rise while material risk remains obscured.

But replacing legacy controls recklessly might introduce unacceptable risk.

This is the new reality: higher expectations, more complexity, zero tolerance for error.

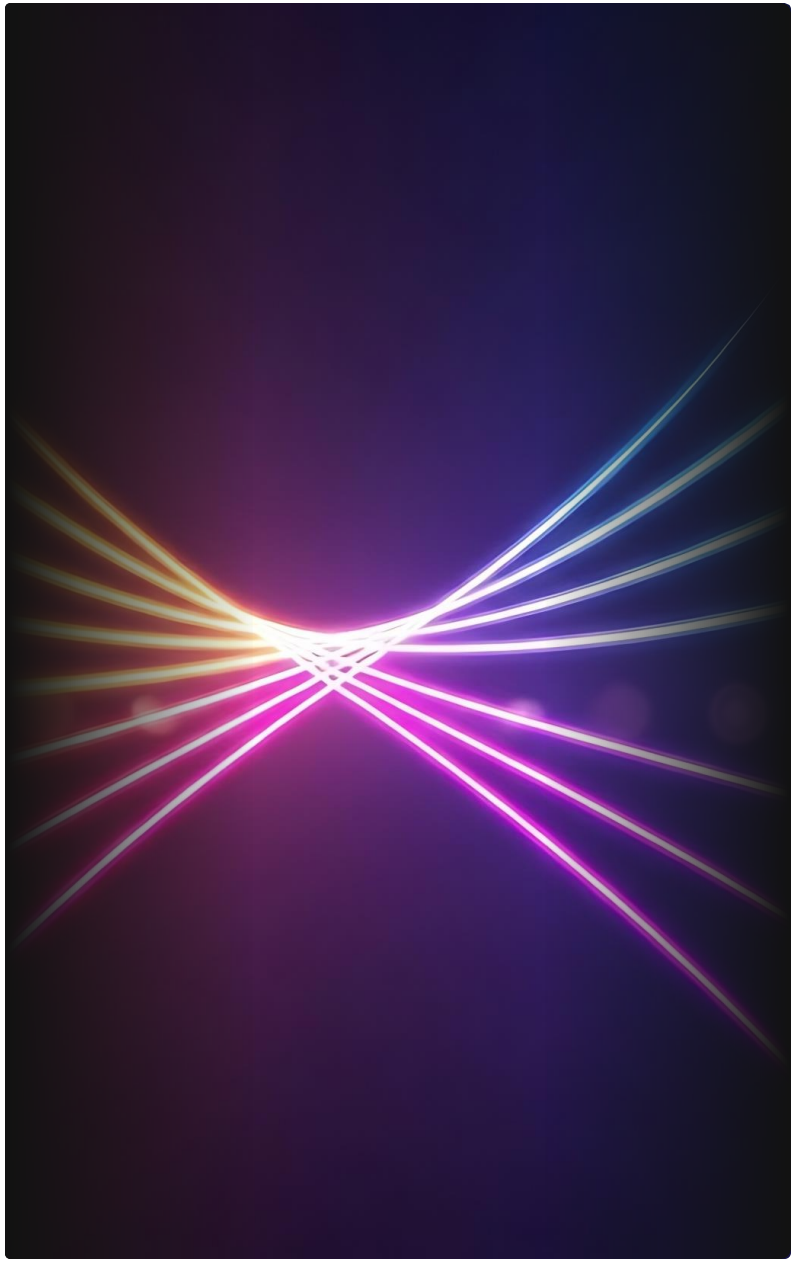

Why 99% accuracy fails in sanctions

In fraud detection, 80–90% accuracy can still generate ROI.
In sanctions screening, **99% is not good enough**.
If 1% of sanctioned parties are missed:

- Institutions must explain exactly why
- Regulators will scrutinize model design and governance
- Enforcement risk becomes material

Explainability therefore is not optional. AI systems in financial crime must provide transparency, accountability, and regulatory defensibility.

This includes:

- Clear rationale for automated decisions
- Auditability of model logic
- Bias detection and oversight
- Alignment with GDPR and emerging AI regulations

Sanctions compliance demands: 100% defensibility,
not just high performance.

Sanctions compliance demands: 100% defensibility, not just high performance

What agentic AI really means in sanctions screening

Agentic AI transforms sanctions screening from a static control into a continuously adaptive intelligence layer. Traditional sanctions programs rely on periodic list updates, batch rescreening, and manual rule recalibration creating operational lag between regulatory change and detection logic.

By contrast, agentic systems dynamically:

- Reassess exposure when new designations are published
- Enrich historical transactions with updated risk intelligence
- Identify emerging ownership linkages
- Adjust triage prioritization in response to geopolitical developments

This shifts sanctions screening from a static matching exercise to an “Always-on” intelligence model, where detection logic evolves alongside regulatory change.

Agentic AI is not a single model. It is an orchestrated network of purpose-driven agents coordinated through governance layers. These agents operate across multiple data environments, structured and unstructured data sources to support investigators in real time.

They can:

- Ingest and cleanse internal and external data
- Perform contextual risk scoring
- Triage alerts based on multi-factor assessment
- Map beneficial ownership and indirect exposure
- Correlate network relationships across counterparties
- Extract intelligence from free-text messages and trade documentation
- Generate structured investigation narratives
- Log documented decision chains for audit review

Importantly: Agentic AI enhances human judgment.

This represents a structural shift:

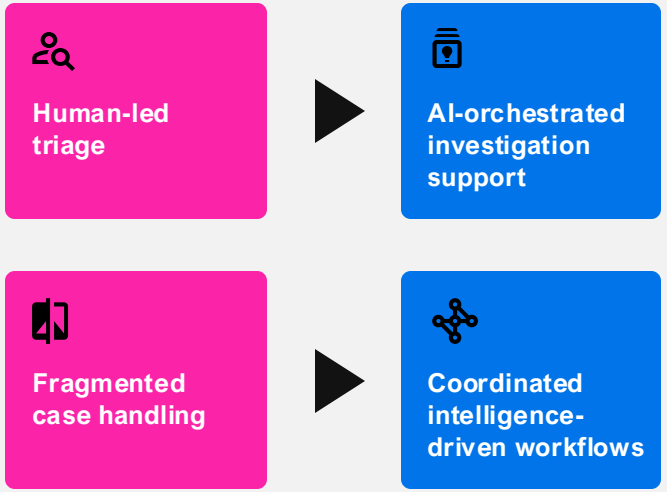

Multi-agent reasoning: the real value

Traditional systems produce alerts. Agentic systems produce reasoned investigation packages. That difference matters most when dealing with sanctions evasion. In sanctions screening, evasion rarely presents as a simple name match. It is increasingly network-based, ownership-layered, trade-masked, and digitally facilitated.

Periodic screening approaches can identify direct matches but struggle to surface indirect exposure or emerging evasion tactics.

Specialized agents evaluate each transaction and entity from multiple perspectives:

- Name similarity, aliases, and transliteration complexity
- Behavioral anomalies and transaction patterns linked to evasion typologies
- Beneficial ownership structures and layered control
- Network exposure across related entities
- Trade and payment context
- Regulatory and jurisdictional context relevant to the decision

The output is not just a score.
It is:

- A structured case file
- A clear and explainable rationale
- A documented decision chain suitable for audit review
- An investigator-ready narrative that highlights material risk indicators

Critically, this approach supports evasion detection by enabling:

- Continuous network reassessment as intelligence changes
- Dynamic beneficial ownership mapping to surface indirect exposure
- Real-time contextual enrichment from internal and external sources
- Adaptive alert prioritization in response to geopolitical and regulatory developments

Agents can also translate technical model outputs (e.g., SHAP values) into natural language investigators understand. This enhances transparency, investigator efficiency and regulatory defensibility.

Explainability by design

Explainability in sanctions screening must operate at three levels:

1. **Prediction transparency:** Why was this match scored as high or low?
2. **Feature interpretability:** Which data points influenced the decision?
3. **Audit traceability:** Can the decision be reconstructed and defended?

Agentic AI strengthens all three.

By combining contextual enrichment with structured reasoning, institutions gain:

- Faster investigations without sacrificing oversight
- Reduced investigator fatigue
- Improved decision consistency
- Enhanced regulatory defensibility

The value is not automation alone.
It is intelligent orchestration within defined governance frameworks that enhances control and detection.

What it really takes to deploy agentic AI in sanctions screening

Deploying AI into mission-critical sanctions screening is not about replacing systems. It is a governance and operating model challenge. Successful deployment requires integration without disruption.

Agentic AI will:

- Integrate into existing screening platforms
- Operate as an AI overlay
- Preserve core rules engines
- Avoid disrupting mission-critical workflows

Institutions can start small and scale gradually, without ripping out mission-critical systems.

Agentic AI does not override screening decisions. It enhances them.

In sanctions screening agentic AI:

- Operates as a deep-dive transaction and entity analysis layer
- Extracts context from unstructured data (SWIFT messages, trade documents)
- Analyzes transliterations, aliases, ownership structures
- Generates probability scores with human-readable rationale

This improves match quality while maintaining transparency. Crucially, agentic AI does not make final sanctions decisions but provides structured feedback.

Human investigators remain accountable.

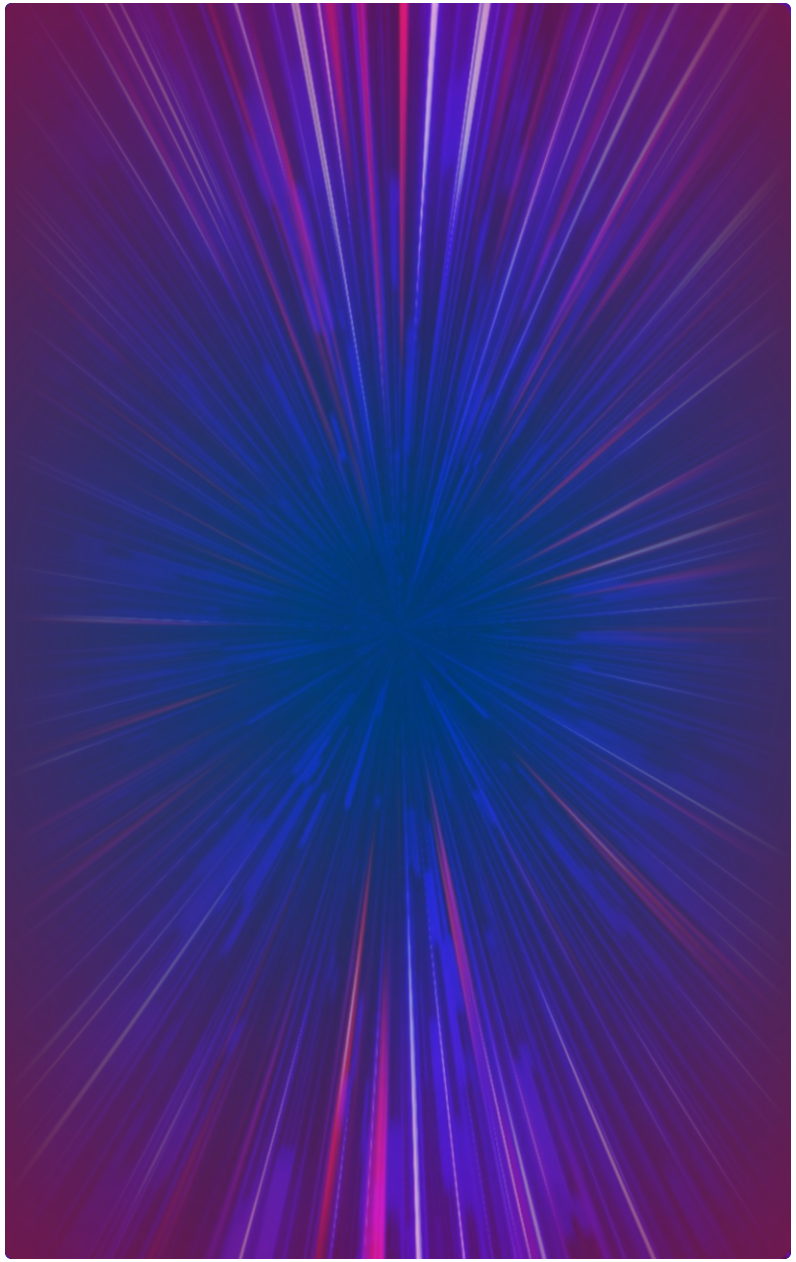

Designing regulator-trusted agentic sanctions screening

Regulators are increasingly open to AI, but only if it is explainable, controlled, and defensible.

- AI must be subject to model risk management frameworks
- Transparency and explainability are mandatory
- Vendor risk and governance frameworks must mature alongside technological innovation

In sanctions screening, where enforcement exposure is material, experimental or opaque AI is not acceptable. What supervisors expect is disciplined innovation: intelligent systems operating within clearly defined boundaries.

Full autonomy is not viable in enterprise compliance environments. Instead, agentic systems should operate through orchestrated layers that:

- Limit scope of agent actions
- Monitor tool usage and decision steps
- Log reasoning chains
- Enable intervention at defined checkpoints

Guardrails ensure that intelligence enhances control rather than weakening it.

These safeguards include:

- Limited autonomy calibrated to risk appetite
- Orchestration layers governing how agents interact with data and systems
- Continuous monitoring of outputs and tool calls
- Explicit adherence to accountability, transparency, reliability, safety, and privacy principles

In sanctions, the threshold is simple: if a decision cannot be explained and reconstructed, it cannot be defended.

Instead, institutions must adopt a human-in-the-loop model:

AI proposes → Humans validate → Institutions remain accountable

Agentic systems may triage alerts, enrich investigations, or generate structured case summaries, but final determinations remain with compliance professionals.

Escalation decisions, suspicious activity reports (SARs), regulatory disclosures, and board reporting continue to sit squarely with compliance leadership.

Human oversight is not a fallback mechanism. It is a structural design principle. It aligns with supervisory expectations for continuous monitoring, adaptive controls, and ongoing model validation effectively moving sanctions programs toward an “Always-on” operating posture.

This design also aligns with emerging AI regulatory frameworks, such as the EU AI Act, the U.S. NIST AI Risk Management Framework, and MAS FEAT principles in Singapore, which emphasize transparency, governance, documentation, and oversight for high-risk systems. The objective is not autonomous compliance. It is AI-capable compliance that is governed, explainable, and defensible by design.

Business Outcomes:

100% audit pass

90% reduction in audit prep

Reduced model and operational risk

Stronger governance frameworks

Increased internal and regulatory confidence

Layered governance in agentic sanctions screening

Regulator trust is reinforced when governance operates across multiple levels:

Operational oversight

- Analysts review and validate AI-generated outputs
- Overrides are logged and analyzed
- High-risk cases trigger mandatory human escalation

Tactical model governance

- Ongoing model validation and drift monitoring
- Performance metrics reviewed against risk appetite
- Independent review and recalibration processes

Strategic accountability

- Executive dashboards provide AI performance transparency
- Boards review risk trends and audit findings
- Accountability remains institutional, not technological

This layered structure ensures that AI is continuously supervised.

Controlled transparency in agentic AI is about:

- Clear rationale for recommendations
- Documented decision chains
- Documented and traceable investigations
- Measurable performance metrics
- Defined override mechanisms

When properly designed, agentic AI can enhance regulatory confidence. Unlike static heuristic rules that may lack contextual adaptability, agentic systems can provide structured reasoning that is easier to document, audit, and defend.

Measuring agentic AI success in sanctions screening

Governance requires metrics. Regulators care about:

- Drift monitoring
- Accuracy
- Escalation logic
- Override rates

Agentic AI performance should be tracked through:

- Task completion rate
- Processing time
- Output accuracy
- Tool call success rate
- Acceptance/rejection rate

Business Impact

When responsibly deployed, agentic AI delivers measurable impact in sanctions screening – not by replacing compliance teams, but by **strengthening their ability to respond to evolving risk.**

Continuous sanctions intelligence reduces the lag between regulatory change and operational response. Institutions no longer rely solely on manual rescreening cycles, periodic list updates, or reactive rule tuning.

Instead, agentic systems continuously evaluate exposure, reassess emerging relationships, prioritize material risk, and support investigators with dynamically enriched contextual intelligence, reducing both regulatory exposure and operational strain.

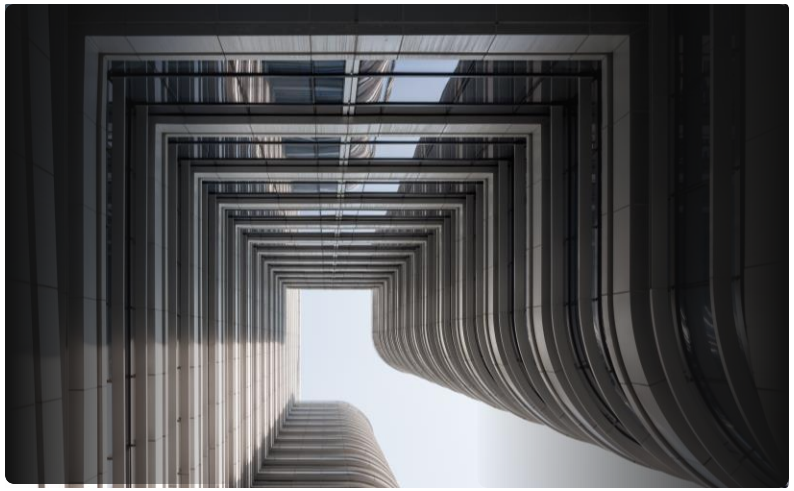

In practical terms, this means:

- **Enhanced detection of complex sanctions evasion**, including indirect ownership and network-based exposure
- **Reduced false positives**, through contextual and multi-factor analysis rather than static name matching
- **Lower operational costs**, by automating enrichment and triage within governed workflows
- **Faster investigations**, supported by structured case files and explainable rationale
- **Stronger audit defensibility**, with documented decision chains and layered oversight
- **Increased regulatory confidence**, through transparent, human-in-command AI design

Agentic AI reinforces institutional control while equipping sanctions teams to operate with greater precision, speed, and resilience in a continuously shifting geopolitical environment.

Case study

From hours to minutes – agentic AI in sanctions screening

Transforming sanctions investigations at a major U.S. financial institution

The challenge

A major U.S. financial institution was processing extremely high daily transaction volumes, resulting in a substantial number of sanctions alerts. Its screening system relied heavily on broad name-matching logic, generating high false-positive rates and placing significant strain on investigation teams.

Each alert required a manual, multi-step investigation process. Analysts were responsible for:

- Reviewing transaction details and counterparty information
- Conducting entity background checks and historical analysis
- Performing extensive web and media research
- Managing Requests for Information (RFIs)

Individual investigations could take over 100 minutes to complete.

A further challenge was the need to review large volumes of adverse media. Investigators often had to assess dozens of articles per alert to determine whether the information related to the correct individual or entity which is a time-intensive and error-prone process.

The institution required a solution that could reduce false positives, accelerate investigations, and improve consistency without compromising regulatory defensibility.

The approach

SymphonyAI deployed its agentic AI solution in a proof-of-concept (PoC) to demonstrate how agentic AI could augment existing sanctions screening workflows.

The agents were trained on the institution's internal policies and procedures and operated as an intelligence layer alongside existing systems – not as a replacement.

Agentic AI solution automated key investigative tasks, including:

Entity resolution and disambiguation
Agents analyzed transaction participants (senders, receivers, and matched entities) by gathering and correlating publicly available information to determine true matches versus false positives

Contextual research and enrichment
Agents conducted web-based analysis, including relationship mapping, background verification, and adverse media assessment to provide contextual intelligence

Prioritization and escalation
Where sufficient evidence could not be established, agents escalated cases to investigators with structured summaries and policy-aligned justifications

Each agent-generated output included a clear rationale, supporting evidence, and alignment to internal compliance frameworks ensuring transparency and auditability.

The outcome

The PoC demonstrated significant improvements in both efficiency and effectiveness:

99% **reduction in false positives,** driven by contextual entity resolution and intelligent disambiguation

90% **reduction in manual effort**, with average alert review time reduced by up to 10x

Over 98% **agreement between agent outputs and investigator decisions**, demonstrating high alignment and reliability

To validate scalability, 50 agents were deployed in parallel:

- Over 300 alerts adjudicated per hour
- Thousands of investigator hours saved
- Consistent, high-quality investigation outputs across all cases

Each alert was accompanied by:

- A structured investigation summary
- A clear adjudication recommendation
- Supporting evidence and linked sources
- Documented rationale aligned with internal policies

This significantly reduced investigator workload while strengthening decision consistency and audit defensibility.

The future state

Following the success of the PoC, the institution-initiated plans to deploy agentic capabilities into production environments.

With access to internal data sources, agent performance is expected to further improve enhancing detection accuracy, contextual understanding, and investigative depth.

The institution is now expanding agentic capabilities across additional compliance workflows, including:

- KYC onboarding
- SAR narrative generation
- Investigation summarization

This progression reflects a broader shift toward embedding agentic AI across the compliance operating model enabling scalable, explainable, and governed intelligence across financial crime functions.

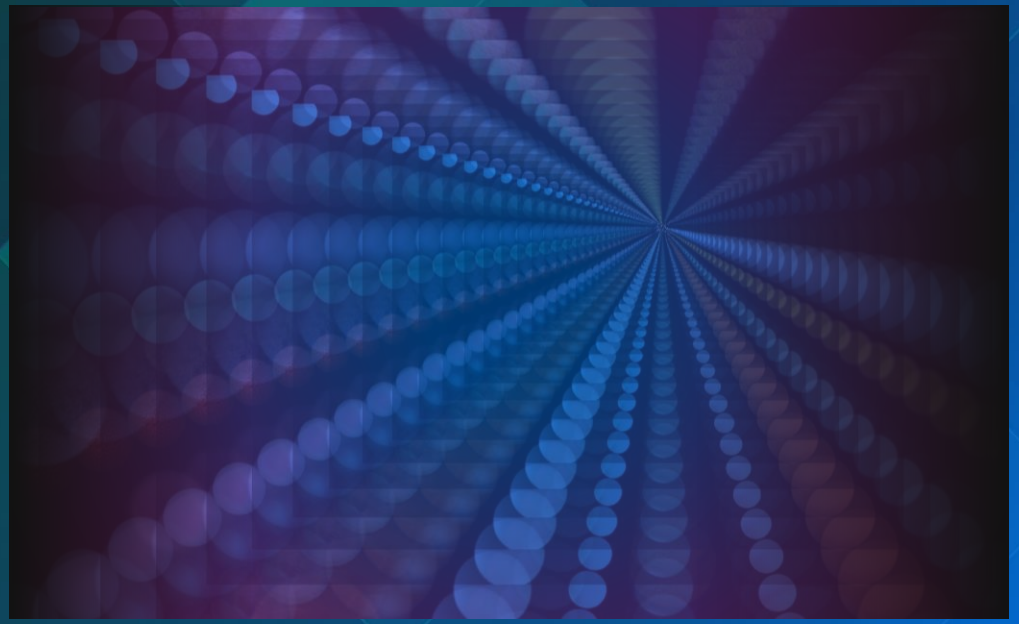

Conclusion

Sanctions compliance is evolving from:

Rules → Predictive AI → Explainable AI → Agentic Investigation Orchestration

As geopolitical volatility accelerates and evasion tactics grow more sophisticated, sanctions screening can no longer rely on periodic updates and static controls. The future lies in sanctions intelligence capability that continuously reassesses exposure, adapts to regulatory change, and surfaces material risk in real time.

Agentic AI, when designed responsibly, enables this shift without destabilizing core systems or compromising regulatory trust. By embedding intelligence within existing workflows and keeping humans firmly in command, institutions can modernize sanctions operations while preserving accountability.

This is not automation for its own sake. It is:

- Controlled, explainable orchestration
- Multi-agent reasoning operating within layered governance
- Continuous enrichment without uncontrolled autonomy

The result is measurable operational relief coupled with stronger regulatory confidence – reduced investigative burden, improved decision consistency, enhanced audit defensibility, and greater resilience in the face of evolving sanctions risk.

The future of sanctions compliance is not autonomous AI.
It is governed, explainable intelligence, with humans firmly in control.

Sanctions programs must now operate with greater speed, precision, and defensibility than ever before. The challenge is not whether to modernize, but how to do so without compromising regulatory confidence.

Agentic AI enables institutions to modernize sanctions screening within existing governance frameworks, strengthening control rather than introducing uncertainty.

With SymphonyAI's agentic AI solution, institutions can transform their end-to-end operations enabling radical effort reduction while improving process consistency and effectiveness:

- Detect complex evasion patterns beyond static name matching
- Reduce false positives through contextual, multi-factor analysis
- Shorten investigation cycles with structured, explainable case narratives
- Strengthen audit defensibility with documented decision chains
- Integrate intelligent workflows into existing screening environments without disrupting core systems

Business outcomes

- 50-70% faster case resolution
- Lower cost per alert and investigation
- 95% of alert noise is cleared before an investigator touches it
- 6x increase in investigator productivity
- 80% fewer false positives
- Near-zero SAR rework
- Alert to SAR in hours, not days
- New detection models live in weeks, not months
- Significant operational efficiency gains
- Measurable ROI from AI-driven automation

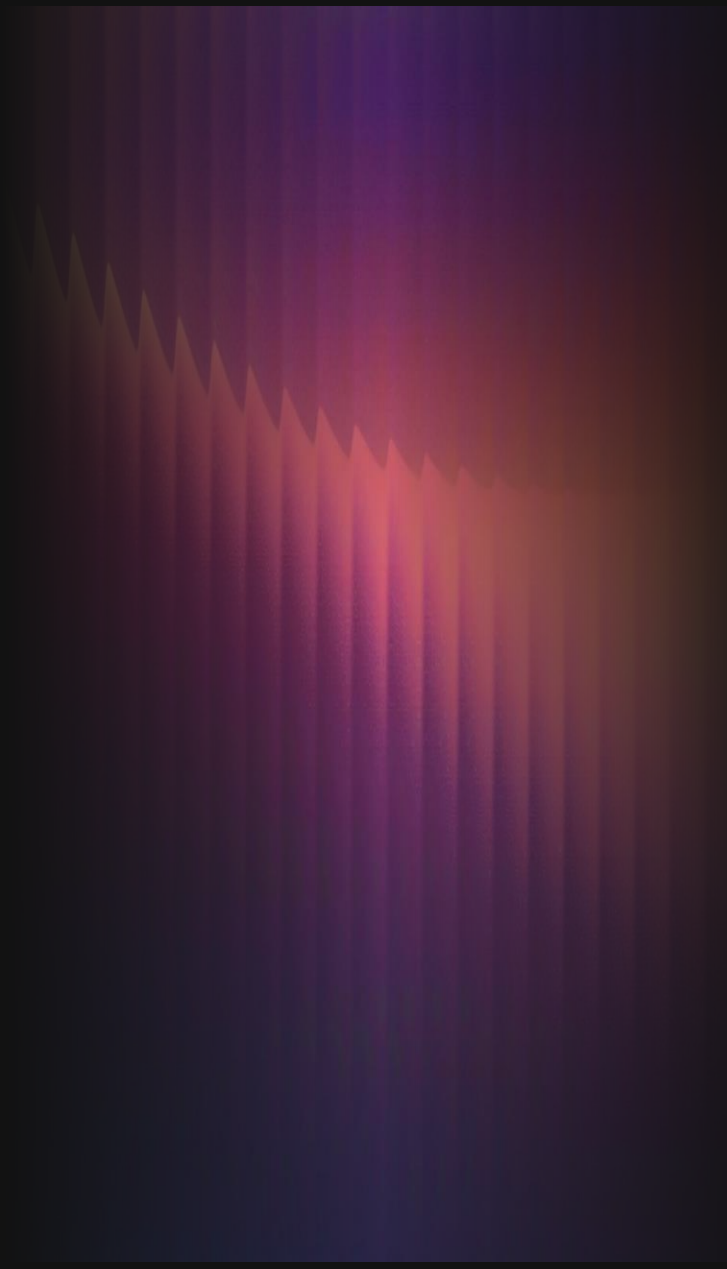

About SymphonyAI

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.

SymphonyAI

www.symphonyai.com

Get in touch with us to assess how agentic AI can enhance your sanctions screening framework responsibly, defensibly, and under full governance control.

Email us directly at connect.fs@symphonyai.com