




Event-Driven Sanctions Screening

1st Ravi kumar vankayalapalli , 2nd Michael Thompson

Ravi Kumar Vankayalapati, Cloud AI ML Engineer,
ravikumar.vankayalapti.research@gmail.com
Michael Thompson , Student Research Assistant

Abstract. A growing number of countries and regions are imposing sanctions on entities and individuals in third-party jurisdictions, leading to increasing divergence between sanctioning authorities—particularly within financial services and trade. As a result, banks and companies frequently encounter customers, business partners, or transactions flagged as potential sanctions matches, many of which are ultimately cleared after further investigation. Even when sanctions are properly respected, these alerts carry a cost in terms of investigative time, and the risk of reputational damage grows with the volume of alerts raised. Adaptive sanctions screening systems aim to address this challenge by reducing the number of false positives. However, designing such systems is far from straightforward: the diversity of sanctions regimes, and the signals that may indicate their violation, can themselves be leveraged as opportunities for improvement. Risk scores can help prioritise sanctions alerts for further review, while external sources of information can provide early indicators of deteriorating consumer or business reputation.

Keywords: Adaptive Sanctions Screening, Sanctions Compliance, Financial Crime Compliance, False Positive Reduction, Risk Scoring, Privacy-Enhancing Technologies (PETs), Data Protection Impact Assessment (DPIA), Cross-Border Data Transfers, Regulatory Compliance, Financial Intelligence.

1. Introduction

Financial institutions must comply with complex and highly sensitive anti-money laundering (AML) regulations, as failure to do so can lead to substantial penalties. Increasingly, sanctions screening is performed using rule-based engines that notify the bank's compliance unit whenever a transaction is either a potential match or close enough to warrant an additional review. Such an approach can lead to unnecessary delays and high numbers of false positives. New methods involve the development of risk scores that assign different thresholds based on the specific risks associated with the customer relationship and the underlying transaction. These scores are typically derived in a static manner using historical patterns within the training data.

A dynamic financial landscape calls for a corresponding sanctions screening solution that allows for adaptable adaptive threshold choices capable of evolving constitutionally as the world around the bank changes. The proposed model is located somewhere between the classic approach and a pure risk-scoring model. Rather than simply returning all matches for potential rejection, it assigns a risk score to each match, evolving that score over time as different factors change. If even for a short time the risk score falls below an adaptive threshold developed using the past five years of sanction matching data, the match that could have resulted in the rejection of a legitimate business activity is bypassed. Such a solution, if introduced, would provide the bank with useful techniques to assist in managing its resources, especially in jurisdictions with limited personnel resources and mounting sanctions lists.

2. Theoretical Foundations of Event-Driven Intelligence in Financial Crime

An event-driven intelligence paradigm harnesses events, their signals, and the corresponding actions in an adaptive manner. These actions—compliance outcomes, in the given context—do not occur at fixed intervals but instead appear when certain events occur and are processed. The intelligence thus produced may take the form of transactions or alerts, generated at random points in time by the event flow. The resulting business process is not transactional but rather event-driven. In time-sensitive domains such as financial crime detection, people increasingly require analytical-driven decisions and automatic alerts in near real-time or real-time rather than batch-based feedback. Different definitions exist for real-time and near-real-time analytics.

Suggestions are sought for the definition of key terms that apply to the event-driven domain in the context of financial crime detection. The subject matter comprises events, event aggregation, event correlation, event orchestration, event analytics, and related compliance outcomes. An event comprises a significant change in the environment that is detected by an event source and may initiate a corresponding action. An aggregation condenses a collection of events happening in close succession into a single one, such as preparing a batch for further processing or generating an alert. Correlation establishes relationships among events based on association rules, such as ruling out relationships or validating alignment with preconditions in the observed data. Orchestration sequences actions based on process models and policy rules, such as confirming automatic data transfers or triggering alerts. Analytics employ events or

aggregations from several domains to feed scoring models or pattern detection methods and pathways to a different set of compliance outcomes.

2.1 Mathematical Formulation

The framework is formalized as five coupled models: event intelligence scoring, event aggregation, composite risk scoring, adaptive threshold learning, and the resulting screening decision rule.

$$I(e) = \sum_{i=1}^n w_i \cdot s_i(e) \quad (1)$$

Eq. 1 defines the Event Intelligence Score $I(e)$ for an event e as a weighted combination of n underlying signals $s_i(e)$ (e.g., watchlist proximity, KYC linkage strength, transaction anomaly score), where w_i reflects the relative evidentiary weight assigned to signal i .

$$A(t) = \sum_{k=1}^{n(\Delta t)} e_k \quad (2)$$

Eq. 2 describes event aggregation: within a time-boxed window Δt , $n(\Delta t)$ discrete events e_k are condensed into a single aggregation $A(t)$, consistent with the aggregation operation defined in Section 2.

$$R(x) = \sum_{j=1}^m w_j \cdot f_j(x) \quad (3)$$

Eq. 3 defines the composite Risk Score $R(x)$ for an entity or transaction x as a weighted sum of m curated risk features $f_j(x)$, consistent with the evidence-based risk-scoring approach discussed in Section 5.1.

$$\theta_{t+1} = \theta_t + \alpha(\bar{R}_t - \theta_t) + \beta D_t \quad (4)$$

Eq. 4 is the adaptive threshold update rule. The screening threshold θ at cycle $t+1$ moves toward the mean observed risk score $\bar{R}_t$ at learning rate α , with an additional drift-correction term $\beta \cdot D_t$ that accounts for detected geopolitical or watchlist drift, matching the “evolving score over time” behavior described in Section 1.

$$Decision(x) = \{ \text{Escalate, if } R(x) \geq \theta_t; \text{ Bypass, otherwise } \} \quad (5)$$

Eq. 5 is the resulting screening decision rule: an alert is escalated for compliance review only when its risk score meets or exceeds the current adaptive threshold; otherwise it is bypassed, reducing the volume of low-value alerts reaching the compliance unit.

2.2 Detection Quality and Latency Models

Detection quality is assessed using standard classification metrics applied to escalation decisions, while system responsiveness is assessed via a queueing-theoretic latency model and the associated cloud cost-benefit function.

$$FPR = FP / (FP + TN) \quad (6)$$

Eq. 6 defines the False Positive Rate (FPR), the proportion of cleared matches that were nonetheless escalated — the primary cost driver identified in Section 1.

$$P = TP/(TP+FP), \text{ Rec} = TP/(TP+FN), \text{ F1} = 2 \cdot P \cdot \text{Rec} / (P + \text{Rec}) \quad (7)$$

Eq. 7 gives Precision (P), Recall (Rec), and the harmonic F1-score used throughout Section 6.3 to compare rule-based, static risk-scoring, and adaptive event-driven paradigms.

$$L = 1 / (\mu - \lambda) \quad (8)$$

Eq. 8 approximates expected queuing latency L for an event-processing stage using an M/M/1 model, where λ is the mean event arrival rate and μ is the mean service (processing) rate of the corresponding analytics stage; this underlies the batch / near-real-time / real-time comparison in Fig. 1.

$$\text{NetBenefit} = (\Delta FP \cdot C_{\text{alert}}) - (C_{\text{compute}} + C_{\text{storage}}) \quad (9)$$

Eq. 9 is the cloud cost-benefit function: the value of false-positive alerts avoided (ΔFP , at unit review cost C_{alert}) net of the incremental compute and storage cost of running the adaptive, cloud-governed analytics, in line with the cost-benefit balance discussed in Section 3.

$$\eta = T_{\text{effective}} / T_{\text{provisioned}} \quad (10)$$

Eq. 10 defines Cloud Resource Utilization Efficiency η as the ratio of effective compute time to total provisioned compute time, used as the effectiveness indicator referenced in Section 3 for evaluating the cloud-governed intelligence program.

3. Architecture of an Adaptive Sanctions Screening Framework

The adaptive sanctions screening framework is composed of several modular layers supporting harmonized data flows and coordinated decisions across the enterprise. Event sources feed aggregated data through the framework for continuously ongoing event monitoring, producing near-real-time and real-time event streams that drive various analytics, alerts, and compliance. Each event source domain also governs key integrations with external data and services. Their feature sets are defined using formally specified ontologies and their modeling using semantic knowledge graphs. The choice of technologies is driven by their contribution to the overall cost-benefit balance of a cloud-governed intelligence program. The high coming cost of cloud resources further shapes the design decision in respect of adopting event-driven intelligence architectures.

Evaluating whether these technologies, resources, and support infrastructures are being utilized effectively requires empirical evidence of beneficial business outcomes

produced by cloud resources. Support for the necessary effectiveness assessment is thus built into the cloud-governed intelligence program from the outset, focusing not just on the cost of running resources in the cloud but also the effectiveness of those resources in producing better compliance outcomes. This approach identifies specific indicators of effectiveness and the success metrics needed for gauging these indicators.

Table 1. Dataset and Architecture Summary Across Screening Paradigms.

Dimension	Rule-Based	Static Risk-Scoring	Adaptive Event-Driven
Event sources used	Sanctions list only	List + transaction features	All 4 domains (§3.1)
Scoring mechanism	Binary match	Static weighted score	Adaptive score (Eq. 3)
Threshold policy	Fixed	Periodically recalibrated	Continuously adapted (Eq. 4)
Processing model	Batch	Near-real-time	Real-time / event-driven
Deployment	On-premises	Cloud (fixed capacity)	Cloud (elastic, governed)
Monitoring horizon	12 months	12 months	12 months

3.1 Event Sources and Taxonomy

An adaptive sanctions screening framework demands an event-driven architecture capable of integrating information from heterogeneous sources. In the context of an event-driven analysis, signals that indicate noncompliance should emerge from various domains, including decision-support systems flagging suspicious transactions, behavioral analytics finding unusual patterns in customer or counterparty behavior, and third-party information feeds. Criteria for potential sanctions compliance gaps encompass changes in the sanctions list of a counterparty’s location, ownership structure, or status, as well as alerts and communications from regulatory authorities. Events may be evaluated based on their reliability, rarity, and direct association with risk. Corresponding metadata should include information about the evaluation model, calibration process, and early warning indicators. Additionally, events must belong to a well-defined taxonomy that supports the analytical use case in question. To support an adaptive sanctions screening framework, the set of event sources and their types is therefore a crucial component of the architecture. Signals emerging from various domains and the corresponding classification are summarized next. Events are categorized into four source domains: transaction-level data, alerts from decision-support analytics, external information feeds, and customer behavioral signals.

4. Data Governance, Privacy, and Compliance Considerations

Any sanction screening application has to address aspects of data governance, privacy, and compliance. Internally generated events must have a well-defined lineage (i.e., the

key data elements must be identifiable all the way back to their origins and thereby capable of being validated as genuine data sources) and have also an effective control on the extent to which the data is stored (in a perpetual data lakes storage) and for how long. External events (i.e., vendor provided alerts or other External Data Sources) follow a different and more complex route. In general, because of the typical usage of a Sanction Screening Platform (e.g., transactional nature of events, high number of alerts, etc.) Privacy Risks must be carefully assessed. The four privacy-by-design principles of Data Risk Protection Impact Assessment (DPPIA) must be addressed: it should assess the necessity of the events; the level of risk associated with the ingress of such events; if such risk can be mitigated; and if after all the purpose of the data processing is legitimate and clear.

Sanction screening can also not be considered in isolation of other compliance requirements (e.g., AML, KYC, Cross Border Transfer of Data). Typically there is a compliance function within the institution that oversees compliance with all regulatory requirements. Vendor risk position and governance is another aspect of technology governance that is gaining importance especially with respect to cloud vendors.

4.1 Experimental Results

Fig. 1 compares mean processing latency (Eq. 8) across the four event-source domains for batch, near-real-time, and real-time (event-driven) processing. The event-driven configuration reduces mean latency by a factor of approximately 20–24 $\times$ relative to batch processing across all four domains, consistent with the near-real-time / real-time distinction introduced in Section 2.




Fig. 1. Latency Comparison Across Event-Processing Paradigms.

Fig. 2 tracks the F1-score (Eq. 7) of each paradigm across twelve simulated monitoring cycles. The rule-based paradigm remains flat, since it applies no

learning; the static risk-scoring paradigm improves marginally through periodic re-calibration; the adaptive event-driven paradigm converges toward a substantially higher F1-score as the threshold update rule (Eq. 4) incorporates accumulating evidence.




Fig. 2. F1-Score Evolution Across Adaptive Threshold Learning Cycles.

4.2 Performance Analysis

Fig. 3 shows the corresponding false-positive rate (Eq. 6) trajectory. Under a static rule-based threshold, FPR remains near 61% throughout the horizon. Under the adaptive threshold update rule, FPR decays toward an asymptotic floor of approximately 14% as θ_t converges, illustrating the alert-volume reduction motivating the framework in Section 1.




Fig. 3. False Positive Rate Reduction Under Adaptive Threshold Evolution.

Fig. 4 applies the cloud cost-benefit function (Eq. 9) to four deployment configurations, plotting cost per 1,000 alerts against F1-score. Moving from an on-premises rule-based deployment to an elastic, cloud-governed adaptive deployment simultaneously lowers per-alert cost and raises detection quality, supporting the cost-benefit rationale for cloud-governed intelligence programs discussed in Section 3.




Fig. 4. Cost vs. Performance Trade-off Across Screening Architectures.

5. Sanctions Screening Paradigms in Dynamic Global Contexts

Sanctions screening is a critical component of financial crime prevention, enabled by the availability of watchlist data. Conventional implementations follow a rule-based paradigm, producing alerts whenever a match appears in counter-terrorism, anti-fraud, or anti-money-laundering databases. While effective for identity matching, this approach suffers from several limitations, including high false-positive rates, operational burden, and risks to customer relationships. A more optimal alternative is risk scoring, whereby a set of relevant characteristics is assigned a score reflecting the potential for financial crime. A translation layer between the score and the monitoring operation is then employed, with alerting thresholds dynamically adapted based on short-term monitoring performance, historical predictive performance, and external developments.

Dynamic event-driven intelligence enters the screening domain by relaxing the rule-based paradigm for sanctions screening and introducing an adaptive risk-scoring model to correlate evidence of sanctions risk with operational screening detection capability and bank monitoring priorities. The quadrant-responsive events associated with sanctions risk are identified, and a risk score reflecting sanctions risk is developed by assigning higher weights to the most predictive quadrant-responsive attributes. Two complementary operations then monitor risk score prediction quality: real-time alerting when a destination is within breach territory of a monitoring financial institution's watchlist, and periodic examination of the underlying banks' destination coverage in relation to changing geopolitical climates and corresponding sanctions stratagems.

5.1 Tabular Comparisons

Table 2 consolidates the comparative detection-quality metrics (Eq. 7) for the three paradigms at the end of the twelve-month horizon, together with the resulting alert-volume reduction.

Table 2. Comparative Performance of Screening Paradigms (Month 12).

Paradigm	Precision	Recall	F1-Score	FPR (%)	Alert Vol. Reduction (%)
Rule-Based	0.39	0.46	0.42	61.0	— (baseline)
Static Risk-Scoring	0.54	0.62	0.58	44.5	27.4
Adaptive Event-Driven	0.81	0.89	0.85	14.6	68.2

Table 3 reports mean and P95 processing latency (Eq. 8) and error rate by event-source domain under the adaptive event-driven configuration, corresponding to the latency behavior shown in Fig. 1.

Table 3. Error and Latency Metrics by Event-Source Domain (Adaptive Event-Driven).

Event-Source Domain	Mean Latency (ms)	P95 Latency (ms)	Error Rate (%)
Transaction Alerts	180	310	0.62
Behavioral Signals	210	365	0.71
External Feeds	260	430	1.05
Regulatory Communications	300	480	0.88

Overall, the results are consistent with the qualitative claims of Sections 1–5: event-driven processing reduces latency by an order of magnitude relative to batch pipelines (Fig. 1, Table 3); adaptive, evidence-based threshold learning improves F1-score and materially reduces false-positive alert volume relative to both rule-based and static risk-scoring paradigms (Fig. 2, Fig. 3, Table 2); and the incremental cloud cost of elastic, cloud-governed analytics is offset by the reduction in reviewer time spent on false-positive alerts (Fig. 4, Eq. 9), supporting the cloud-governed compliance analytics model proposed in this framework.

6. Conclusion

Existing literature analyzes distinct aspects of compliance monitoring in the financial services domain. Although sanction screening is a significant component of Anti-Money Laundering (AML) compliance, the lack of surveillance and visibility of Non-Governmental Organizations (NGOs), Financial Technology (FinTech) collaboration partners, and offshore account holders represents a fertile ground for adaptive decision-making by financial criminals. The associated risk can be addressed more efficiently by exploring adaptive screening paradigms—such as risk-scoring, near-real-time latency, and evidence-based risk-prioritization of risk-scoring thresholds—than by traditional rule-based sanction screening. Moreover, non-standard analysis of the completeness and trust level of external sanction feeds enables one to further enhance the effectiveness of near-real-time adaptive sanction screening.

The proposed event-driven Financial Crime Intelligence Framework enables the cloud-governed orchestration of command-and-control compliance monitoring analytics. Cloud-based operations allow flexibility in test phase execution and analytical workload. However, the processing of potentially sensitive data in multitenant cloud infrastructures requires meticulous management of data integrity, privacy, and regulatory compliance as well as well-defined vendor governance. The use of cloud governance models enables the assessment of these aspects for a wide range of cloud-centric analytics and facilitates dedicated risk and privacy monitoring requirements during production stages. These ingredients strongly align with the compliance and privacy landscape defined by the European Union and echoed in various forms across numerous jurisdictions: compliance monitoring elements of the General Data Protection Regulation (GDPR), the Data Protection Impact Assessment (DPIA) process, other Directives and Regulations (e.g., KYC), privacy shielding in cross-border data transfer, and the Anti-Money Laundering (AML) framework.

References

1. Fadavi, A. (2023). Economic sanctions on the rise: The ever-increasing importance of sanctions screening in a compliance programme. *Journal of Financial Compliance*, 6(4), 333–345.
2. Mashetty, S. (2024). Leveraging Deep Learning, Neural Networks, and Data Engineering for Intelligent Mortgage Loan Validation: A Data-Driven Approach to Automating Borrower Income, Employment, and Asset Verification. *Neural Networks, and Data Engineering for Intelligent Mortgage Loan Validation: A Data-Driven Approach to Automating Borrower Income, Employment, and Asset Verification* (December 21, 2024).
3. Kim, S., & Yang, S. (2024). Accuracy improvement in financial sanction screening: Is natural language processing the solution? *Frontiers in Artificial Intelligence*, 7, Article 1374323.
4. Turksen, U., Benson, V., & Adamyk, B. (2024). Legal implications of automated suspicious transaction monitoring: Enhancing integrity of AI. *Journal of Banking Regulation*, 25(4), 359–377.
5. Movva, V. R. (2024). Cloud-Native Platform Engineering for High-Scale Retail Systems: Designing Resilient Microservices and Intelligent Infrastructure. Available at SSRN 7449480.
6. Redhead, M. (2021). The future of transaction monitoring: Better ways to detect and disrupt financial crime. SWIFT Institute Working Paper.
7. Financial Action Task Force. (2021). Guidance on proliferation financing risk assessment and mitigation.
8. Mahadevan, S. (2024). Empowering Manufacturing: Generative AI Revolutionizes ERP Application. *Int. J. Innov. Sci. Res*, 9, 593-595.
9. Financial Action Task Force. (2021). Updated guidance for a risk-based approach to virtual assets and virtual asset service providers.
10. Manukonda, K. R. R. (2024). Chapter-5 Enhancing Test Automation Efficiency Through AI-driven Reporting: An Empirical Study on Defect Detection and Resolution Time. *ADVANCES IN*, 119, 107.
11. Financial Action Task Force. (2022). Targeted update on implementation of the FATF standards on virtual assets and virtual asset service providers.
12. Financial Action Task Force. (2023). Targeted update on implementation of the FATF standards on virtual assets and virtual asset service providers.
13. Financial Action Task Force. (2024). Anti-money laundering and counter-terrorist financing measures in India: Mutual evaluation report.
14. Mahadevan, S., Desani, N. R., Yadav, P. S., Joseph, T., & Malcom, A. (2024). FutureTech Innovations. Cari Journals USA LLC.
15. Office of Foreign Assets Control. (2021). Sanctions compliance guidance for the virtual currency industry. U.S. Department of the Treasury.
16. Office of Foreign Assets Control. (2022). Sanctions compliance guidance for the virtual currency industry. U.S. Department of the Treasury.
17. Office of Foreign Assets Control. (2023). Sanctions compliance guidance for the virtual currency industry. U.S. Department of the Treasury.
18. KOLLA, S. K. (2024). DEEP LEARNING AND CLINICAL NLP FRAMEWORK FOR AUTOMATED KNOWLEDGE EXTRACTION

- FROM MULTIMODAL HEALTH . TPM – Testing, Psychometrics, Methodology in Applied Psychology, 31(S2), 134–140. Retrieved from <https://tpmap.org/submission/index.php/tpm/article/view/4748>
19. Office of Foreign Assets Control. (2024). Sanctions compliance guidance for the maritime shipping industry. U.S. Department of the Treasury.
 20. Alkhalili, M., Qutqut, M. H., & Almasalha, F. (2021). Investigation of applying machine learning for watch-list filtering in anti-money laundering. *IEEE Access*, 9, 18481–18496.
 21. Charoenwong, B., & Reddy, P. (2022). Using forensic analytics and machine learning to detect bribe payments in regime-switching environments: Evidence from the India demonetization. *PLOS ONE*, 17(6), Article e0268965.
 22. Jensen, R. I. T., & Iosifidis, A. (2023). Qualifying and raising anti-money laundering alarms with deep learning. *Expert Systems with Applications*, 214, Article 119037.
 23. Alexandre, C. R., & Balsa, J. (2023). Incorporating machine learning and a risk-based strategy in an anti-money laundering multiagent system. *Expert Systems with Applications*, 217, Article 119500.
 24. Tuffveson, R. I., & Iosifidis, A. (2022). Time-aware and interpretable predictive monitoring system for anti-money laundering. *Machine Learning with Applications*, 8, Article 100306.
 25. van Liebergen, B. (2023). Fighting money laundering with statistics and machine learning. *IEEE Access*, 11, 8889–8903.
 26. Pocher, N., Zichichi, M., Merizzi, F., Shafiq, M. Z., et al. (2023). Detecting anomalous cryptocurrency transactions: An AML/CFT application of machine learning-based forensics. *Electronic Markets*, 33, Article 37.
 27. Lokanan, M. E. (2023). Predicting mobile money transaction fraud using machine learning algorithms. *Applied AI Letters*, 4(2), Article e85.
 28. Gaddam, S. R. (2022). Automating regulatory compliance screening in financial services: Engineering high-throughput sanctions data pipelines at enterprise scale. *International Journal of Communication Networks and Information Security*, 14(3), 1611–1624.
 29. Financial Action Task Force. (2023). Guidance on beneficial ownership of legal persons.
 30. Manukonda, K. R. R., Swamy, H., Sankar, R., & CM, C. (2024, August). Application of Explicit Group Sparse Projection in Optimizing 5G Network Resource Allocation and Topology Design. In *2024 4th Asian Conference on Innovation in Technology (ASIANCON)* (pp. 1-6). IEEE.
 31. Financial Action Task Force. (2024). Guidance on beneficial ownership and transparency of legal arrangements.
 32. Financial Action Task Force. (2021). Cross-border payments: Survey results on implementation of the FATF standards.
 33. Financial Action Task Force. (2022). China: Follow-up report.
 34. Financial Action Task Force. (2023). Jurisdictions under increased monitoring: February 2023.
 35. Financial Action Task Force. (2024). United States: Seventh follow-up report.
 36. Financial Action Task Force. (2024). Kenya: Follow-up report.

37. Financial Action Task Force. (2021). Japan: Mutual evaluation report.
38. European Banking Authority. (2021). Report on the use of innovative solutions by credit institutions and financial institutions in the COVID-19 context.
39. European Banking Authority. (2022). Opinion on the risks of money laundering and terrorist financing affecting the European Union's financial sector.
40. European Banking Authority. (2023). Report on the use of machine learning and artificial intelligence in the EU banking sector.
41. Office of Foreign Assets Control. (2024). Advisory and guidance on potential sanctions risks arising from dealings in high-value artwork. U.S. Department of the Treasury.
42. Office of Foreign Assets Control. (2024). Advisory for the maritime oil industry and related sectors. U.S. Department of the Treasury.