

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 02 ISSUE: 02

RECEIVED: APRIL 05

REVISED: APRIL 24

ACCEPTED: MAY 11

PUBLISHED: MAY 21

Cost-Aware Event Screening for Financial Compliance

Daniel Wilson

Junior Research Assistant

Abstract— The Event-Stream Compliance Intelligence Framework presents a conceptual model for data-driven risk governance, with crime prevention and cost control as its central aims. It brings intelligence, monitoring, and screening together in a single system. By applying event-stream processing to internal and external data feeds in near real time, the framework supports timely event detection and risk governance. It assigns clear roles: who is accountable for compliance, who assesses newly identified violations, and at what point those assessments enter cost-aware risk management. Data normalization performed ahead of analysis ensures that the objects reported on are those most valuable for detecting crime and reducing losses. Together, these elements allow screening functions to be expressed in terms of minimizing risk-adjusted expected losses. Success is judged by detection accuracy, which is tied directly to total cost of ownership.

Financial screening for crime, fraud, and litigation is essential to protecting the organization's reputation and shareholder assets. Confirming that those who transact business on the organization's behalf comply with key regulations, laws, and best practices is the foundation for preventive crime controls. Monitoring intelligence gathers and refines information from diverse sources to detect a broad range of potential violations of restrictions, regulatory requirements, and business constraints. A further pillar of risk governance is engaging the appropriate governance team to assess suspected new crimes, frauds, or policy infractions. This step can be central to stopping some crimes before they occur and, at a minimum, to reducing their likelihood, severity, or financial impact.

Keywords Event Stream Compliance Intelligence Framework; Financial Screening; Risk Governance; Cost Awareness; Real Time Decision Making; Event Stream Processing; Regulatory Intelligence; Compliance Intelligence; Baseline Compliance Monitoring; Event Stream Processing.

1. INTRODUCTION

The event-stream compliance intelligence framework extends financial screening data collection across internal and external sources, prepares the data for screening, and governs screening decisions in a real-time context. Screening coverage, effectiveness, and decision-making performance are operationalized as success criteria.

Risk governance defines processes, personnel, and structures that provide oversight over risk decisions and related trade-offs. When supported by the closure of event streams associated with the screening of financial partners at the transaction-level, risk exposure and financial crime risk can be accounted for in risk governance discussions; however, these discussions are typically inefficiently conducted on a periodic basis and often fail to consider screening coverage or accuracy. In addition, when combined with automated cost-aware risk governance policies, screening performance becomes susceptible to optimization in real-time, enabling real-time decision-making.

2. Background and Motivation

Many regulations across the globe, such as the US Security and Exchange Commission Rule No. 17a-8, require firms in the financial services industry to ensure that their counterparties are also complying with relevant regulations. Nowadays, the counterparties of such firms can be located anywhere in the world and hence can be subject to any

regulation. In addition to being costly and time-consuming, having an in-house process to ensure compliance of every counterparty can also lead to complaints from customers. As a result, various intelligence companies have sprung up that, for a price, screen and advertise that a certain counterparty has been vetted. Such services are often considered the "gold standard" in compliance. Yet in reality, only a fraction of all transactions with respect to a given counterparty is monitored. As a result, many firms have begun to assume that a compliance service covering a simple "yes" or "no" assumption on the detected topic is reliable, leading to unexpected losses when the assumption proves inaccurate.

The following foundational aspects drive this work. Event-stream processing enables real-time access to information relevant to screening. A rich event-supply ecosystem is in place, and the associated costs, timeliness, reliability, and coverage indicators can be defined as a minimum cost to a minimum quality of service ratio. However, within the context of screening, having to perform the same process in similar respects for similar topics over short periods of time is costly and contributes little to risk governance. In dynamic planning or risk determination, keeping budgets in mind while using services at different costs for similar messages is possible using techniques from the operations research domain.

I. Notation and Formalization

Table I summarizes the notation used throughout this supplement.

Symbol	Definition
s	Real-time risk / screening score assigned to an event by the intelligence pipeline
τ	Screening decision threshold applied to s
c_{FN}, c_{FP}	Unit cost of a false negative (missed risk) and a false positive (over-screening)
$L(\tau)$	Risk-adjusted expected loss at threshold τ
C_{ops}, C_{esc}	Operational ingestion/processing cost; escalation-and-validation cost
$TCO(\tau)$	Total cost of ownership at threshold τ
B	Risk-governance budget allocated to escalation/validation over a period
$P(e), I(e)$	Estimated likelihood and impact of triggering event e
Q_k, V_k	Data-quality index and ingestion volume of source k

Table II summarizes the data-source categories described in Section III-A.

Category	Examples	Typical Mode	Governance Attribute
Internal	Transactions, accounts, smart contracts, wallets	Real-time / batch	Compliance-benchmark filtering

Category	Examples	Typical Mode	Governance Attribute
External – lists	Sanctions/bl acklists, PEP registries	Batch	Licensor provenance tracking
External – intel	Threat intelligence, negative news	Near real-time	Source reliability scoring
External – market	OTC trade / market data	Batch	Trade-validation checks

3. The Event-Stream Compliance Intelligence Framework

The integrated framework conceptual model, components, interfaces, and interactions are presented. Governance roles, data lineage, and decision workflows are defined. Framework objectives, success criteria, and measurement metrics are articulated.

The Event-Stream Compliance Intelligence Framework enables efficient monitoring of cost-sensitive real-time events that warrant financial screening. During incident-free periods or positive event charges, triggering events may materialize with a non-negligible likelihood and impact. Budgeting mechanisms based on cost-aware policies assign monetary limits to screening capabilities, while cost optimization techniques define computing-resource thresholds that minimize the risk-adjusted expected loss. The composite non-exceedance condition of the total cost of ownership governs the decision path associated with real-time risk signals.

The integrated framework can also be viewed from the lens of risk appetite alignment. Role players—data producers, consumers, handoff authorities, and expert risk officers—are configured to carry out specific functions within the overall environment. The deployment landscape encompasses event originators, an Event-Stream Compliance Intelligence data-processing back office, exposing entities, event monitoring signals, financial screening engines, and screening-request approval authorities. Request-validation controls manage the legitimacy of screening demands, decision-rights exclusivity is gauged in real time, and cost-relevant event domains calibrate supervisory monitoring.

Mathematical Formulas:

The risk-adjusted expected loss combines the cost of missed events (false negatives) with the cost of unnecessary escalations (false positives) at a given threshold:

$$L(\tau) = c_{FN} \cdot P(y=1, s < \tau) + c_{FP} \cdot P(y=0, s \geq \tau) \quad (1)$$

The total cost of ownership adds operational and escalation costs to the risk-adjusted loss, in line with the paper's framing of TCO as the governing decision criterion:

$$TCO(\tau) = L(\tau) + C_{ops}(\tau) + C_{esc}(\tau) \quad (2)$$

The framework's screening threshold is selected to minimize total cost of ownership, formalizing the cost-optimization objective described in Section V:

$$\tau^* = \arg \min_{\tau \in [0,1]} \text{TCO}(\tau) \quad (3)$$

A budget constraint limits the number of events escalated to specialized-analyst validation within a governance period, reflecting the paper's budgeting-mechanism discussion:

$$\sum_{i=1}^N c_{\text{esc},i} \cdot \mathbf{1}[s_i \geq \tau] \leq B \quad (4)$$

Event prioritization within the monitoring layer combines likelihood and impact, consistent with the description of triggering events materializing with 'non-negligible likelihood and impact':

$$U(e) = w_p \cdot P(e) + w_i \cdot I(e), \quad w_p + w_i = 1 \quad (5)$$

The composite non-exceedance condition that governs the real-time decision path (Section III) can be stated as a probabilistic budget-adherence constraint:

$$\Pr[\text{TCO}(\tau) \leq \text{TCO}_{\max}] \geq 1 - \varepsilon \quad (6)$$

A. Data Sources and Ingestion

Internal data sources include transactions, accounts, users, and smart-contract deployments registered on a blockchain node, as well as associated wallet addresses for custodians or off-chain interactions. Through filtering with compliance benchmarks, data transfers may also prompt governance processes. Future capital-market activities will generate additional data for tracking. External sources include blacklists from compliance licensors, threat intelligence from partners, negative news from commercial providers, reputation scoring from service providers, and market data for over-the-counter trade validation. Data quality is paramount to the accuracy of the compliance activity, and any erroneous data must be rectified during the ingestion phase; neglecting input data quality may result in abundant false positives and ineffective utilization of compliance resources.

The integration process normalizes the data from different sources into a canonical representation, particularly for source-agnostic features such as addresses and companies. Data schemas are dynamically controlled, with automated updates to consumers and persistence artifacts. Given the wide variety of sources, data deduplication is applied where appropriate. Data ingestion into the event-stream processing pattern can occur in real time or batch mode, depending on latency requirements, resource consumption, and associated governance constraints. A monitoring solution provides transparency on data provisioners and their processing timeliness. Cost constraints and the burden of technical debt—the effort of adding further rule-based filtering to the various compliance engines—prohibit constant operation at maximum capacity. Consequently, certain sources operate in batch mode, with explicit workflows on system utilization. Data provenance, timeliness, and compliance with regulatory constraints further guide ingestion operation.

Table III summarizes the governance role matrix described in Section III.

Role	Primary Function	Decision Right
Data producer	Originates internal/external event feeds	None (supply only)
Data consumer	Applies screening logic to normalized events	Flag / no-flag
Handoff authority	Routes flagged events per escalation path	Route / defer
Risk officer (expert)	Validates classification of escalated events	Approve / reject
Screening-approval authority	Authorizes screening-request legitimacy	Approve / deny request

Table IV contrasts the two error-cost regimes central to the threshold-optimization objective in Eq. (2).

Error Type	Governance Consequence	Cost Driver
False negative	Missed risk event proceeds unscreened	$c_{\text{FN}} \times \text{severity/likelihood}$
False positive	Legitimate event escalated for validation	$c_{\text{FP}} \times \text{analyst/validation cost}$

4. Architecture and System Design

The overall architecture consists of multiple modular layers, enabling easy scalability. The resulting design can be easily enhanced with fault tolerance measures. Various security controls are incorporated to safeguard sensitive data. Furthermore, the design ensures interoperability with existing platforms, allowing seamless integration into the organization's technology landscape. In addition to the functional requirements, the architecture addresses other non-functional objectives, such as provisioning and connecting the data processing pipelines to Governance Overlays.

The proposed system design provides a high-level overview of the data processing pipelines that constitute the heart of the system. These pipelines collectively implement the data extraction, transformation, enrichment, validation, and feature engineering steps of the data-processing lifecycle. The architecture supports modern streaming technologies and implements event-time-based features, such as windowing operations. Comprehensive backpressure handling mechanisms safeguard pipeline performance without compromising the end-to-end ingestion latency. Finally, monitoring and traceability mechanisms are integrated to assess data quality and completeness.

A. Data Processing Pipelines

Data processing pipelines are built for a specific use case, tuning and utilization may vary across specific projects. Data processing pipelines for a compliance-related screening project require special care in the steps involved. Main items that need to be considered during the design and build of these pipelines are:

- At the start of the pipeline, data is extracted from one or multiple sources. A process for validating the collected documents is built and put into practice. With well-known data sources, the required information can be validated against known criteria. Information with poor data quality can go through an enrichment process. If an enrichment process cannot provide the required information, that particular instance can be flagged and escalated to a manual validation step.

- Validation processes that ensure data quality are put in place before the data can go into a machine learning model.

- Features of the data related to all the analysed products and the prediction model are engineered to allow data from both the model and the applied algorithms to be combined during the backtesting.

Data processing pipelines gain importance when a real-time-response-oriented architecture is needed. The need for a stream-processing architecture usually stems from the desire to allow immediate responses or warnings. There are some important recommendations for working with such an architecture.




Fig. 1. Illustrative cost-threshold trade-off implied by Eqs. (1)–(3): the optimal τ^* balances over-screening cost against risk-adjusted expected loss.




Fig. 2. Illustrative example allocation of a fixed screening budget B across event categories (schematic, not derived from operational data).

5. Risk Governance with Cost Awareness

Risk governance objectives, decision rights, and escalation paths are aligned with the organizational governance model and supported by a cost-aware screening policy. Budgeting controls permit the allocation of dedicated resources for dealing with classified risk events and optimize the screening threshold to minimize both the risk-adjusted expected loss and the total cost of ownership.

The validation of risk classification by a specialized analyst incurs an associated cost. Depending on the organization's internal structure, such analysts may be incorporated within the risk management team or may belong to other departments (e.g. finance or human resources). A cost-aware strategy thus enables the provisioning of cost limits for those risk events flagged by the compliance screening that require further validation. Such limits designate the risk events that, if indeed materialized, should bypass validation. Consequently, a budget can be defined and allocated to the approval of external and structured decisions with risk implications. The screening threshold for a specific risk signature is therefore adjusted by a cost-aware strategy that aims to minimize the total cost of ownership corresponding to the approval of false positives and the risk-adjusted expected loss related to the approval of false negatives.

Once the overall risk budget is spent, the monitoring mechanism can issue alerts regarding further risk events but should not trigger the compliance screening subcomponent until it is recharged. The compatibility of the framework with the risk appetite and external constraints of the organization enables an enhanced-level alignment of the decision with requirements that extend beyond pure risk exposure. The association of risk appetite costs with thresholds for monitored risk signatures further guarantees a non-marginal effect on the deliverables generated by the Event-Stream Compliance Intelligence Framework during periods in which compliance monitoring covers a broad scope.

A. Risk Appetite Alignment

To operate in a profitable manner, organizations need to set and respect risk appetite levels. However, stakeholder expectations as well as the associated risk costs constantly shift, making it necessary to adjust thresholds for risk screening, the levels of required approval, and the frequency of impact monitoring. Closing business processes may, for instance, require risk screening for customers. Yet, in periods of low business demand, such low costs of trade price may no longer justify the need to perform a risk screening. Whenever a cost that has a risk mitigation impact incurs, risk appetite and thus also risk cost needs to be reflected.

Regulatory governing bodies, in particular, expect organizations to incorporate cost considerations in their decision-making. Non-compliance represents a possible cost that may incur. Using reduced thresholds close to zero for all

screening variables and closely monitoring the decisions for false positives would ensure mitigation of non-compliance cost. The cost of possible False Negative decisions, therefore, needs to be meticulously compared with the total cost of governing the decision making, thus including all possible error costs involved in risk governance. Reasoning that cost causes cost, for an organization it is slightly more economic to have make use of generally more expensive offerors for services that need risk approval, thus causing it to be more expensive for any possible malicious customer, making it more likely for the customer to decide not to trade with this switch and to rather abuse a different switch without risk approvals. The same oversights can also be applied to any False Positive cost connection, like insurance approval costing monitoring against the total costs of insurance decisions and operation.

6. Conclusion

The Event-Stream Compliance Intelligence Framework comprises a conceptual model and components that address event-stream-driven risk governance in relation to cost awareness and real-time decision-making. Risk governance objectives related to screening and escalation are specified, and cost concerns are addressed through budget rules and sensitivity analysis of screening thresholds. The framework is linked to the broader risk appetite and potential approval mechanisms for deviations and alerts.

The contributions provide evidence that an appropriately designed event-stream compliance intelligence capability can enhance compliance screening accuracy while reducing overall cost. Existing screening controls for money laundering, fraud, and related-party transactions can be adopted and adapted to forensically investigate tipping points for other forms of inappropriate behaviour, particularly in relation to ethical investing, but ultimately in relation to the integrity of the balance sheet, income statement, and cash flow position of an organisation and its determined-risk appetite. Event-stream processing has not traditionally been evident in risk governance discussions, but the absence of a clear risk tolerance has also limited cost-aware screening decisions.

REFERENCES

1. Baesens, B., Höppner, S., & Verdonck, T. (2021). Data engineering for fraud detection. *Decision Support Systems*, 150, 113492.
2. Błaszczyński, J., de Almeida Filho, A. T., Matuszyk, A., Szeląg, M., & Słowiński, R. (2021). Auto loan fraud detection using dominance-based rough set approach versus machine learning methods. *Expert Systems with Applications*, 163, 113740.
3. Canhoto, A. I. (2021). Leveraging machine learning in the global fight against money laundering and terrorism financing: An affordances perspective. *Journal of Business Research*, 131, 441–452.
4. Srinivas Kalisetty. (2023). Big Data–Driven Cloud Collaboration Models for Enhancing Supplier–Retailer Synchronization in Modern Manufacturing Supply Chains. *Journal of Computational Analysis and Applications (JoCAAA)*, 31(4), 2188–2205. Retrieved from <https://eudoxuspress.com/index.php/pub/article/view/4232>
5. Cheng, C.-H., Kao, Y.-F., & Lin, H.-P. (2021). A financial statement fraud model based on synthesized attribute selection and a dataset with missing values and imbalanced classes. *Applied Soft Computing*, 108, 107487.
6. Domashova, J., & Mikhailina, N. (2021). Usage of machine learning methods for early detection of money laundering schemes. *Procedia Computer Science*, 190, 184–192.
7. Pamisetty, V. (2022). Making IoT Data Processing Scalable: Architectures That Actually Work. *International Journal of Scientific Research and Modern Technology*, 281. <https://doi.org/10.38124/ijrmt.v1i12.1187>
8. Gupta, A., Dwivedi, D. N., Shah, J., & Jain, A. (2021). Data quality issues leading to sub optimal machine learning for money laundering models. *Journal of Money Laundering Control*, 25(3), 551–555.
9. Izotova, A., & Valiullin, A. (2021). Comparison of Poisson process and machine learning algorithms approach for credit card fraud detection. *Procedia Computer Science*, 186, 721–726.
10. Shou, M., Bao, X., & Yu, J. (2021). An optimal weighted machine learning model for detecting financial fraud. *Applied Economics Letters*, 30(4), 410–415.
11. Stojanović, B., Božić, J., Hofer-Schmitz, K., Nahrgang, K., Weber, A., Badii, A., Sundaram, M., Jordan, E., & Runevic, J. (2021). Follow the trail: Machine learning for fraud detection in FinTech applications. *Sensors*, 21(5), 1594.
12. Danda, R. R., Maguluri, K. K., Yasmeen, Z., Mandala, G., & Dileep, V. (2023). Intelligent Healthcare Systems: Harnessing Ai and Ml To Revolutionize Patient Care And Clinical Decision-Making. *International Journal of Applied Engineering and Technology* <https://papers.ssrn.com/sol3/papers.cfm>.
13. Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications*, 193, 116429.
14. Hayble-Gomes, E. (2022). The use of predictive modeling to identify relevant features for suspicious activity reporting. *Journal of Money Laundering Control*, 26(4), 806–830.
15. Jensen, R. I. T., & Iosifidis, A. (2022). Fighting money laundering with statistics and machine learning. *IEEE Access*, 11, 43887–43902.
16. Jensen, R. I. T., & Iosifidis, A. (2023). Qualifying and raising anti-money laundering alarms with deep learning. *Expert Systems with Applications*, 214, 119037.
17. Meda, R., & Pamisetty, A. (2023). Intelligent Infrastructure for Real-Time Inventory and Logistics in Retail Supply Chains. *Educational Administration: Theory and Practice*, 29 (4), 5215–5233.
18. Lokanan, M. E. (2022). Predicting money laundering using machine learning and artificial neural networks algorithms in banks. *Journal of Applied Security Research*, 19(1), 20–44.

19. Ouazzane, K., Polykarpou, T., Patel, Y., & Li, J. (2022). An integrated machine learning framework for fraud detection: A comparative and comprehensive approach. *International Journal of Information Security and Privacy*, 16(1), 1–17.
20. Vorobyev, I., & Krivitskaya, A. (2022). Reducing false positives in bank anti-fraud systems based on rule induction in distributed tree-based models. *Computers & Security*, 120, 102786.
21. Achakzai, M. A. K., & Juan, P. (2022). Using machine learning Meta-Classifiers to detect financial frauds. *Finance Research Letters*, 48, 102915.
22. Khan, A. T., Cao, X., Li, S., Katsikis, V. N., Brajevic, I., & Stanimirovic, P. S. (2022). Fraud detection in publicly traded U.S. firms using Beetle Antennae Search: A machine learning approach. *Expert Systems with Applications*, 191, 116148.
23. Pandugula, C., & Nampalli, R. C. R. Optimizing Retail Performance: Cloud-Enabled Big Data Strategies for Enhanced Consumer Insights.
24. Lin, K., & Gao, Y. (2022). Model interpretability of financial fraud detection by group SHAP. *Expert Systems with Applications*, 210, 118354.
25. Tuffveson Jensen, R. I., & Iosifidis, A. (2022). Time-aware and interpretable predictive monitoring system for anti-money laundering. *Machine Learning with Applications*, 8, 100306.
26. Naveed, N., Munawar, S., & Usman, A. (2023). Intelligent anti-money laundering fraud control using graph-based machine learning model for the financial domain. *Journal of Cases on Information Technology*, 25(1).
27. 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, 119500.
28. Pavlidis, G. (2023). Deploying artificial intelligence for anti-money laundering and asset recovery: The dawn of a new era. *Journal of Money Laundering Control*, 26(7), 155–166.
29. Yandamuri, U. S. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. Zenodo.
30. Tong, G., & Shen, J. (2023). Financial transaction fraud detector based on imbalance learning and graph neural network. *Applied Soft Computing*, 149, 110984.
31. Yi, Z., Cao, X., Pu, X., Wu, Y., Chen, Z., Khan, A. T., Francis, A., & Li, S. (2023). Fraud detection in capital markets: A novel machine learning approach. *Expert Systems with Applications*, 231, 120760.