

# Event-Driven Compliance Systems: Modernizing Financial Crime Detection Without Machine Intelligence

**P S L Narasimharao Davuluri**

*Senior Data Specialist Data Engineering, pslnarasimharao.davuluri@ieee.org,  
ORCID ID: 0009-0009-0820-8184*

## **Abstract**

Internal compliance processes for financial crime detection within banking and financial services organizations have, until now, relied heavily upon machine intelligence-oriented methods. While a significant number of suspicious matter reports are generated, the proportion of actionable matters from those reports is small and the reduction of false positives remains a challenge. Late-stage intervention on large and complex data classes is expensive and banks and governments recognize that substantial modernization is needed.

An alternative approach, rooted in event-driven architecture, enables a different form of compliance system that, while not using machine intelligence or analytics to drive detection, holds the potential for greater effectiveness, efficiency and reduced risk. Such a system is positioned as a minor enhancement to existing systems rather than a wholesale replacement. By incorporating high-frequency and low-cost detection strategies, real-time capabilities, triage workflows, governance and risk management, and compliance considerations that address the limitations of the event-driven system, tangible benefits can be sought.

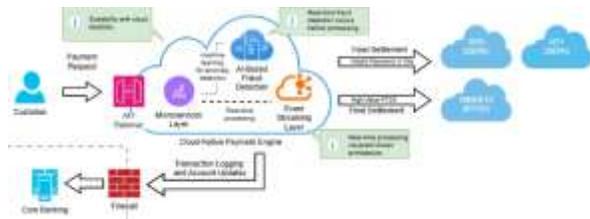
**Keywords :** Event-driven architecture; compliance; regulatory technology; supervisory technology; financial crime detection; business rules management systems; event processing; information retrieval; signal correlation; signalling pathways; financial services; governance; risk; compliance; governance risk and compliance; detection strategy; auditability; explainability; transparency.

## **1. Introduction**

The world faces an alarming threat from financial crime-related activities. While the related harm cascades through the economy and society, detection remains inefficient and ineffective. An exponential increase in transaction volume and unprecedented levels of investment has not reduced detection shortfalls. Financial institutions alone spend over US\$300 billion annually on compliance and contribute approximately US\$16 billion to US national and local treasuries, yet detection yields negligible national or global impact. The pursuit of machine intelligence continues to dominate despite limited recent success. Machine intelligence approaches, rather than revolutionising compliance monitoring, have dis-proportionately increased event streaming costs, complicated system management, and diminished detection efficacy.

The proposed architecture modernises compliance monitoring through event-driven design principles. Within a financial platform context, event streams are structured, captured, and analysed using well-established design patterns, without machine intelligence. Supervised learning and unsupervised clustering replace conventional machine intelligence-driven analytical methods. Reliable subject-matter expertise

becomes just as essential for successful detection of financial crime as it is for preventing other high-cost harms associated with fire, bankruptcy, or fatality. The discussion illustrates how the requirements of governance, risk, and compliance (GRC), critical to any successful enterprise, shape the event-driven financial-crime detection architecture.



**Fig 1: Event-Driven Architecture in Compliance**

**1.1. Background and Significance** Systems designed to detect financial crime have proliferated in the past twenty years, but with little effect on the incidence of money laundering or the rate of prosecution. The resulting compliance fatigue, compounded by the adverse impact of false positives on genuine customers, has exposed the yawning gulf between the expectations of regulators, the demands of business-as-usual, and the realities of detection system performance. In parallel, machine intelligence techniques have reached the stage where simple rule-based detection solutions may no longer seem sufficient to address complex detection requirements. The combination of these two trends suggests that detection systems could benefit from the same evolutionary leap as the rapidly-evolving user interfaces of retail banking, brokerage and insurance products.

Such modernization can reduce the reliance on machine intelligence for financial crime detection. Human analysis is increasingly being displaced by event detection systems capable of making independent decisions about detection and response. In compliance, systems of this type must sift complex signals against business-as-usual event streams, following the detection, or otherwise, of events of interest. In some systems, high user fatigue levels have been compounded by poor detection performance. A focus on poor investigation outcomes for true positives has also revealed the capacity of machine intelligence systems to exacerbate outcome quality, rather than improve it.

## 2. The Evolution of Financial Crime Compliance

Financial crime, especially money laundering and terrorist financing, is a major focus of regulatory authorities around the world. The governance, risk, and compliance (GRC) systems established in response to these demands have resulted in significant costs for businesses and for society as a whole, but have had only a limited impact on the efficacy of financial crime investigations. Although sceptics have long voiced their concerns, the inadequacies of the current approach are now provoking serious debate about how to achieve better outcomes at reasonable cost. An emerging consensus points to the need for a radical transformation of anti-financial crime GRC systems away from traditional batch-processing architectures towards more effective real-time or near-real-time compliance systems. These proposals are timely. Providers of artificial intelligence (AI) and machine-learning-based applications have recently begun to offer a vision of a transformed detection landscape in which traditional, rules-based compliance systems are complemented or replaced by systems that can learn and adapt to changing behaviour patterns in financial crime. The resulting hype and excitement, however, may obscure a vital point: developing effective real-time compliance systems need not involve machine intelligence. Financial crime detection can be transformed without incorporating any machine-learning techniques at all. Event-Driven Compliance Systems provide a complete set of compliance components capable of delivering modern, event-driven detection without any element of machine intelligence. They are able to detect money-

laundering and terrorist-financing signals by leveraging the analytical creativity of human minds—the very factor that makes a compliance team sound its alarm.

Event-Driven Compliance Systems represent a comprehensive set of tools for tackling the detection gap that remains in the shift to event-driven compliance. Such systems cover all aspects of the detection process—including real-time ingestion and event modelling, real-time processing and event streams, the definition of detection strategies based on standard rule, heuristic, and signal-detection approaches, implementation of triage workflows, and the definition of associated operational metrics—all while avoiding any machine-intelligence component. The growing body of evidence supporting a gradual drift away from reliance on traditional batch-processing architectures makes the exploration of these detection solutions timely. Rather than seeking to adapt compliance systems to take advantage of emerging machine-intelligence technology, events-driven compliance systems represent a coherent, end-to-end approach to developing realistic, effective, and explainable event-driven detection systems.

### Equation 1: Event system model (EDA) → formal definition

Let each event be a record:

$$e_i = (t_i, id_i, type_i, \mathbf{x}_i)$$

- $t_i$ : event time
- $id_i$ : entity key (customer/account/counterparty/case key)
- $type_i$ : event category (transaction, login, KYC update, sanction hit, etc.)
- $\mathbf{x}_i \in \mathbb{R}^d$ : attributes/features carried in the event payload

A stream is time-ordered:

$$\mathcal{S} = \{e_i\}_{i=1}^{\infty}, t_1 \leq t_2 \leq \dots$$

The paper states: an event model encapsulates structure and attributes; attributes are selected for monitoring; and each attribute is represented both as expressions and rules.

Event-Driven Compliance Systems...

Formally, define an event type schema:

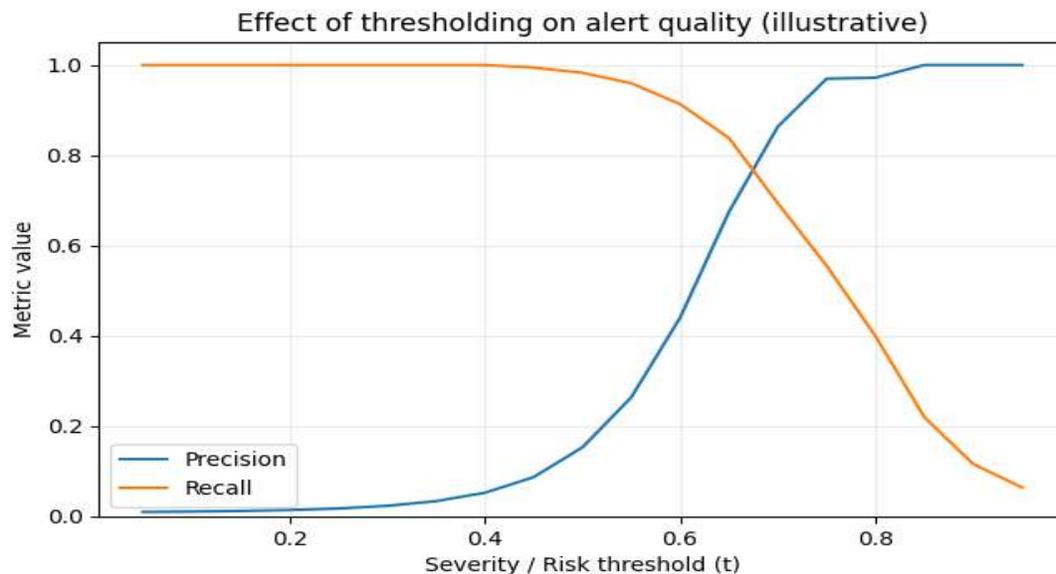
$$\mathcal{M}_{\text{type}} = \{a_1, \dots, a_d\}$$

and an attribute extractor:

$$x_{i,j} = a_j(e_i)$$

**2.1. Research design** Research was conducted using the design science framework, which emphasizes the construction and evaluation of innovative artifacts in response to recognized problems. The artifact in this case is an event-driven architecture for financial crime compliance systems, with special focus on a detection component. In this context, design science offers a structured process for creating an artifact that can alleviate the identified problem, while simultaneously ascribing the resulting artifact with sufficient rigor and persistence to be considered a candidate contribution to the body of knowledge.

The event-driven compliance system is conceived as a remedy to the three limitations of established systems—namely that they do not exploit predictive power, generate excessive false positives, and lack transparency and explainability—that result from their reliance on machine intelligence for detection. Key features include dedicated ingest capability for predictive signals, a specialized event modeling layer, a real-time processing engine that triggers alerts and creates analysis-ready cases, a signal correlation capability in the detection layer, thresholding for sensitivity tuning, and triage workflows for both temporal and non-temporal events.



**Fig 2: Functional Layers of the Event-Driven Detection Artifact**

### 3. Event-Driven Architecture in Compliance

The application of event-driven architecture to compliance is described in this section. Compliance is modeled as a distribution of responsibilities across a set of stakeholders, a clear definition enables identification of event-based equivalents of compliance services and functions. The approach is further refined by conceptualizing the kind of events that should be monitored, the source components of detection engines, and how the outputs should be combined to execute detection-related workflows. These sub-models guide the subsequent design of an event-driven architecture for compliance that leverages continuous, scalable, and real-time processing for the distribution, sharing, and execution of detection-related responsibilities.

The design of event-driven systems begins with modeling the application domain as an event system—in other words, a system that generates and consumes events. An event system consists of event sources, generating and publishing events, and event sinks, subscribing to and processing events. An event-driven architecture for compliance can be identified by explicitly modeling the domain as a distribution of responsibilities across stakeholders, that is, industry participants expecting delineated compliance capabilities from others. The respective responsibilities for compliance can be separated into distinct services whose provision relies on the notifications generated by other stakeholders. Such a clear functional separation greatly aids the detection of signal patterns in multi-party settings.

**3.1. Core Concepts and Components** In event-driven architecture (EDA), the system handles events to react to business activities. An event is a significant change in the state of something in the business domain and can relate to a transaction or an interaction. An event-driven compliance system is composed of a set

of integrated components performing specific functions at various points of control. These components include an event model, event sculptors focused on monitoring and controlling data ingestion, event processors concentrating on analyzing event streams in real time and detecting operational escalations; a signal correlation service for connecting detected events based on correlation logic and a case management application for the triage of signals and escalated alerts.

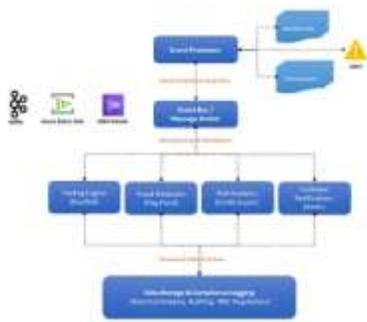
An event model encapsulates the structure and associated attributes of each data category crossing an EDA system. Attributes are selected for monitoring and analysis according to the purposes of the event-driven detection strategy. Each attribute of interest is modeled in two forms: as a model expression and as model rules that enable detection at runtime. Model expressions enable thresholds to be defined for drowning alerts that surface when the attribute value crosses (or reaches) the severity threshold defined at runtime, thus surfacing events that need operational attention. Model rules are judgmental conditions and support the enabling/disabling of such alerts for selected data categories and periods.

**3.2. Data Ingestion and Event Modeling** Rules, heuristics, correlated signals with thresholds, and management of Cases and Triage workflows function at or near real-time. An event-driven compliance system builds detection capabilities based on these elements and associated workflows to identify financial crime-related events. Within the homicide context analogy, it leaves the identification of who committed the crime, when, where, how and for what motivation to humans without the aid of machine intelligence while triggering alerts by identifying victims, witnesses, persons of interest, and other entities that merit closer scrutiny.

Events are modeled as data structures and diverse external sources feed primary and secondary events into event streams. Sources are designed to encourage separate data preparation paths for high-volume and often-low-quality wiring-gapping compliance, such as Alerts from transaction monitoring systems, and data profiling for support functions like Manual transaction monitoring or Media watch for presence in newsfeeds and unplanned events like asset freezes. The arrival of data from these sources is managed through streaming platforms like Kafka or Services like AWS Kinesis or Azure Event Hubs. These systems offer linear scaling and support a high number of production-producing and consumption-consumer Applications, reducing the structural development phase for research and experimentation based on rapidly growing datasets.

**3.3. Real-Time Processing and Event Streams** Events can include the detection of anomalous or suspicious activity within a detection engine. Groups of alerts or cases are thus produced in real time and delivered to analysts for investigation. Event stream processing can improve the exploratory analysis of potential signals by examining the interaction of systems and entities. Event collection tools or notification hubs can centralize alert, anomaly, and case notification from disparate detection engines. Related alerts can be organized into case-like structures and provided as an investigative APT, supporting triage and more efficient analysis.

Event-driven analysis tools can be used to investigate pairs or small groups of entities—typically, clients and counterparties associated with a user of the tool—across a set of common events, generating content suitable for human exploration. Tool output can cover checklists for different issues, including volume increase and sudden spikes, along with previously defined and tested ad hoc signals. Checks enabled by APTs can range from high-volume flows, sudden changes in transaction direction, change or introduction of high-risk counterparties (particularly those using different identified types, such as cash), to transaction volume correlations over selectable time periods. The output can be presented on an interactive dashboard, integrating situation-scoped visualizations to enable simple exploration.



**Fig 3: Real-Time Processing and Event Streams**

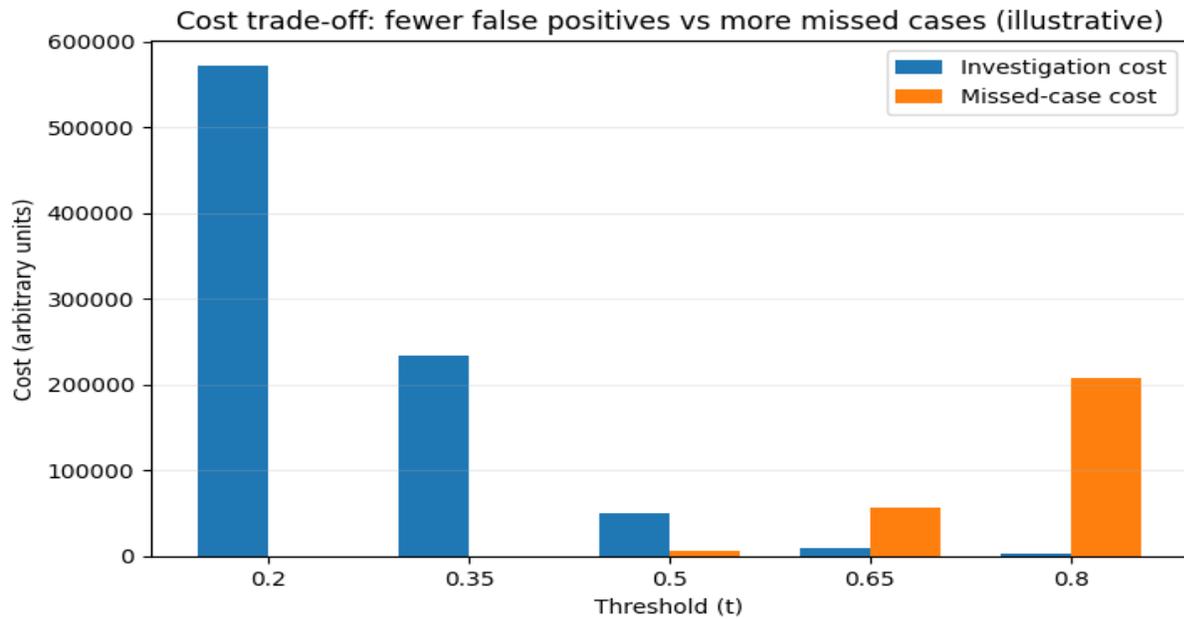
#### 4. Detection Strategies Without Machine Intelligence

Despite the awareness that detection strategies need not involve machine intelligence, there are few detailed instances describing such approaches in the literature. Two categories of detection methods can be readily identified: those based on explicit and semi-explicit specifications, for which the implementation is largely self-evident, and those that rely on less precise descriptions and for which publication of patterns would be counterproductive.

Detection strategies based on explicit or semi-explicit specifications take the form of rules, heuristics, frequently occurring activity patterns, and thresholds on control signals. All security technologies generate copious log data, a small but critical fraction of which pertain to security events. The ability to correlate segments of logs that appear innocuous—perhaps even benign on a narrow bandwidth—in isolation enables Discovery Security Operations Systems to reveal dangerous activity of significant duration, or climaxing in unexpected situations. For instance, a surge of employee access to payroll data will likely indicate a malicious insider, especially when the frequency far exceeds an established baseline and there is no internal corporate trigger, such as a merger or acquisition. Increasingly, security teams are fusing Signals-of-Intent from the Dark Web, Crowd Intelligence data, and tradecraft-compliant cyber operations with information from their SIEMs and SOAR solutions to enhance decision-making capabilities.

**4.1. Rule-Based and Heuristic Methods** Most major categories of model risk in compliance systems rely on rules, heuristics, or combinations of the two as detection methods. A rule represents a pattern of behavior that is indicative of potential financial crime across a limited set of parameters (Balko et al. 2020). Rules trigger events for specific types of financial crime within narrowly defined context. Despite the large number of rules maintained by most institutions, these still tend to capture only a small proportion of a firm’s financial crime risk. The low incremental case generation rate means that many of these rules have not been working effectively, as support for human resources on case investigations has been relatively stable. Consequently, adding new ones does not provide greater support for investigations. However, with a continual process of deletion, invalid rules could be pruned and attention could shift toward those that are known or believed to be working effectively.

Heuristics are slightly broader than rules, used to capture a wider spectrum of risk. A heuristic might specify problematic combinations of flows or contrasting flows between two entities. Relying on detection strategies with these patterns has generated good case generation rates. The application of heuristics to cases flagged for review has also provided incremental support for investigations and improved results.



**Fig 4: Operational Flow of Rule and Heuristic Signal Generation**

**4.2. Signal Correlation and Thresholding** Comparison, correlation, and thresholding of relevant events signals facilitate detection of financial crime by identifying high-density zones, measuring intensity and flow, and comparing disparate events. In rules-based systems, similar methods fulfil an auxiliary role, serving to identify specific operations to trigger further detection, investigation, and potential case creation.

Commerce, finance, and the global web of finance have been affected by intermittent crises of fraud since at least the tulip mania of the 1630s; later episodes include the 1720 South Sea Bubble, the Netherlands’ 1763–1774 financial scandal, and the 2008 financial crisis; today’s bubbles and breakouts. These and other financial bubbles test the plausibility of Game’s Thief-in-the-Temple Model and Theorem, a generic formal model that holds for the outright theft of funds and securities as well as for other forms of financial fraud. The events described by the model include manic euphoria, the appearance of novel financial schemes, gathering signs of undenture, suspicious transactions (real or apparent), public warnings, and opportunist burglaries.

**Equation 2: Model expressions → thresholds (“severity threshold”) step-by-step**

Pick an attribute  $a_j$  and a threshold  $\theta_j$ .

**Step 1: compute attribute value**

$$v_{i,j} = a_j(e_i)$$

**Step 2: compare to threshold**

$$\text{alert}_{i,j} = \mathbb{1}[v_{i,j} \geq \theta_j]$$

where  $\mathbb{1}[\cdot]$  is the indicator function (1 if true, else 0).

Many compliance signals are “over the last  $W$  minutes/days”. Define a time window for entity  $id$ :

$$\mathcal{W}(id, t; \Delta) = \{e_k : id_k = id, t - \Delta < t_k \leq t\}$$

Example: total outgoing amount over  $\Delta$ :

**Step 1: aggregate**

$$A(id, t) = \sum_{e_k \in \mathcal{W}(id, t; \Delta)} \text{amount}(e_k)$$

**Step 2: threshold**

$$\text{alert}(id, t) = \mathbb{1}[A(id, t) \geq \theta]$$

**4.3. Case Management and Triage Workflows** Signal detection systems often generate false positives. Focusing on the correct assignment and progressions of those signals into triage workflows is important. Composite signals recognize similar or recurring scenarios across entities, activities, and time that need analysis. Examination can clarify whether what is detected is straight-through processing for communication, systems operation, or other purposes. Business application-system custom case signals can smoothly integrate system-specific operational concerns, either with rules, heuristics, or cross-entity correlation signals. Cost-based signals show overhead scaling with volume or activity levels and are tied to analytical skill areas needing examination.

Implicit detection systems may generate clutter—a warning to the processes that operational staff traverse daily. Redundancy reduction or navigation assistance supports the objective. Critical empty areas of systems can be positively regarded as too low a data load to trigger. Signals flagging missing cases or those that are abnormal or rare can assist or warn users. SISs track area-domain or skill-area attributes to avoid over-editing or relay for others’ inspection.

Non-MIS-detection systems cover risk areas best. Application-specific signals and triage management assess the need for investigation and relay the alerts. Incremental cost alerts can detail curved-relation breakages. Retro-DA signals indicate system abuse, and GRC-SISs capture often-overlooked business process considerations and typical system operation concerns. Operational-knowledge signals on behavior warranting review are alerts reserved for limited parties.

**5. Governance, Risk, and Compliance Considerations**

Organizational, technical, and procedural best practices in Governance, Risk, and Compliance (GRC) enhance the feasibility, effectiveness, efficiency, and trustworthiness of Event-Driven Compliance Systems using non-Machine Intelligence detection methods. Alignment with regulations and laws, an auditable process, and proper documentation are foundational. Additional principles of explainability and transparency support stakeholders in understanding GRC implications for Event-Driven Systems, particularly the management of signal and false-positive data quality and handling the consequences of privacy violations.

Equipped with adequate GRC measures, Event-Driven Systems can reasonably expect to mitigate false positives at a lower operational cost than current systems. The impact on detection sensitivity remains uncertain, as the role of non-detection during investigations has not been systematically evaluated. Event-Driven Systems using machine learning models flow through the organization, particularly in areas

characterized by a low signal-to-noise ratio, where the case-management system is a potential choke point. Here, human involvement—especially the synthesis of evidence from seemingly unrelated signals across disparate business processes—may strengthen the pipeline.



**Fig 5: AI Governance and Risk Management Framework**

**5.1. Regulatory Alignment and Documentation** Emerging frameworks in Machine Learning and Natural Language Processing open up unprecedented avenues for organisations engaged in financial crime detection. In contrast with these techniques, Event-Driven Compliance Systems seek to modernise these detection environments in the absence of Machine Intelligence along three parallel fronts: the use of real-time processing paradigms, a more structured and comprehensive method of data ingest within a streaming environment, and an overarching emphasis on development, governance, and support strategies for dwell-time minimisation.

To achieve harmonisation with requirements imposed by governing bodies, the traditional principles of Regulatory Alignment and Documentation are augmented with specialised Theories of Governance, Risk, and Compliance.

When introducing an Event-Driven Compliance System for financial crime detection, alignment with the guidelines provided by governing entities is essential. Supporting documentation must cover prevention or mitigation measures related to any conceivable risk, fraud, or misbehaviour statement in an organisation’s associated regulatory framework such as the UK Money Laundering Regulations, the Dodd-Frank Act, or the Payment Services Directive. Security refers to people, policies, and resources dedicated to shielding an organisation’s technical infrastructure. From a Data Perspective, substantial levels of data quality, privacy, and security must also be guaranteed for these compliance systems.

**Table1:** Event-Driven Compliance System Overview

Element	What It Does	Why It Matters
<b>Event-Driven Processing</b>	Monitors financial events in real time	Faster detection of suspicious activity
<b>Detection Approach</b>	Uses rules, heuristics, thresholds, and signal correlation	Clear, explainable, and auditable decisions
<b>Case &amp; Triage System</b>	Groups alerts and routes them for investigation	Reduces analyst overload
<b>Governance &amp; Metrics</b>	Applies compliance controls and performance measures	Improves regulatory alignment and efficiency

**5.2. Auditability, explainability, and transparency** Event-driven detection methods enable evidence-based supervision of the detection systems underpinning financial crime compliance. Every detected signal should invoke a clearly articulated detection specification together with a relevant pre-event history. Moreover, the specification and supporting data should be rapidly accessible so that responsible stakeholders – typically compliance officers, external auditors, or regulators – can independently assess the validity, relevance, and importance of any detected case. Detection specifications built from heuristics or business rules can often be directly audited by the stakeholders, while such clarity typically diminishes for detections based on correlation analysis or continuous signal thresholding.

A related primer question concerns the definitional nature of any detected signal. A case detection is fully FAEs if it establishes an unambiguous fact or event – e.g., the execution of a fraudulent transaction by an account, instantiation of a relevant cash transaction surveillance case, or country sanction violation. Signals based on discreet and fully-FAE case definitions warrant the highest degree of auditability and any associated audit trails are expected to address key queries with high precision. Signals marking detected anomalies tend to be less FAE in nature, as they do not denote the existence of a specific fact or event. Auditability is correspondingly weaker, with the responsible personnel formally adjudicating on these detected violations to maintain governance. Detections undoing alerts raised by other compliance detection systems might be similarly classified, with human intervention establishing whether a genuine compliance violation is present or has been nullified.

**5.3. Data Quality, Privacy, and Security** An event-driven compliance system that fulfills traditional detection and alerting functions can address regulatory requirements such as an auditable record of detection, justifications for escalation and closure, and supporting documentation for the decision-making process. Compliance is also supported by an explanation of why the methodology does not rely heavily on machine intelligence, thus enhancing process transparency. However, these benefits are dependent on the quality of the data used by the system, particularly given that cleansing and processing activities are not automated. These data quality considerations, together with privacy and security issues connected to the handling of sensitive data, must therefore be carefully addressed.

Compliance is assisted by the absence of machine intelligence: detection methods are simple and explicable, and the majority of the processing is rule-based. Removal of data quality concerns is crucial because event-driven systems lack machine intelligence. High-quality event data streams are essential, and every stream used by the method must have its source fully understood. Any requirements around data ethics and algorithmic auditing apply. A significant caveat concerns the signal correlation approach, where a disparate accumulation of signals makes the combination more difficult to explain. Data accountability challenges are also pronounced because of the storage of both AML and fraud data. Data privacy requirements are particularly sensitive given the personal nature of much of the data involved in AML and fraud detection.

## **6. Evaluation of Effectiveness**

A broad range of measures can assess event-driven systems that detect and react to potential monetary crimes and help reduce the number of false signals sent for investigation. These measures balance the detection of actual criminal activity (number of true positives) against the operational costs associated with investigating flag signals. In simple terms, how many criminal cases were caught, and how many were false positives? The dollar amount of criminal involvement is a secondary concern for these measures; in fact, the dollar value of flagged cases should generally remain a secret.

The architecture has been tested in industry use cases to yield several measures of operational effectiveness: supervision, alert, and investigation rates; true positive, false positive, and negative ratios; alert and investigation dollar amounts; false alert rate by investigation outcome; investigation duration and resource consumption; case closure rate; supervision commitments; investigation throughput; and external reporting

fees. Together, these metrics convey the trade-offs between the need for investigations of alert signals and the valid criminal activity being reported. Additionally, the external investigations highlight how the audit law and Anti-Money Laundering laws also serve to deter, detect, and formally punish financial crime when other measures are insufficient.

### Equation 3: Signal correlation → from isolated alerts to cases

Let each “atomic” detection produce a signal:

$$\sigma_k = (t_k, \text{id}_k, \text{signal\_type}_k, \text{severity}_k)$$

Build a graph  $G = (V, E)$  where vertices are signals and edges connect “related” signals:

$$(\sigma_p, \sigma_q) \in E \Leftrightarrow \phi(\sigma_p, \sigma_q) = 1$$

where  $\phi$  is correlation logic such as:

- same customer id within  $\Delta$  time
- shared counterparty
- same device/IP
- same beneficiary across many senders, etc.

For an entity id and time  $t$ , define the set of recent signals:

$$\Sigma(\text{id}, t) = \{\sigma_k : \text{id}_k = \text{id}, t - \Delta < t_k \leq t\}$$

Aggregate severity:

$$S(\text{id}, t) = \sum_{\sigma_k \in \Sigma(\text{id}, t)} \text{severity}_k$$

Create a **case** when the correlated score crosses a case threshold:

$$\text{CASE}(\text{id}, t) = \mathbb{1}[S(\text{id}, t) \geq \theta]$$

**6.1. Metrics for Detection and False Positive Reduction** Modernizing financial crime compliance systems can occur without reliance on specialized machine intelligence to detect potential financial crime. The detection of signals or occurrences of potential illicit behavior depends upon the strategy design and implementation, which can comprise a hybrid of sound heuristics coupled with expert domain knowledge. Adopted detection methods can include rule-based, signature-based, heuristic-based, pattern-matching, or coordination proposal strategies. For timeliness, these detection methods should strive to operate outside of conventional batch windows, reducing the time lag for detection. Nevertheless, operating outside the limits of machine intelligence also entails the risk of increased false positives.

There is a tension between timely detection and false positive reduction: lower false positive rates often come with added delays. Within an event-driven compliance approach, deployed detection mechanisms must possess sufficient effectiveness metrics to limit, or at least not increase, operational burden. Although any operational cost is an important metric, the amount of possible false positives can also be expressed as

a detection costs rather than a pure operational cost. Where hours of expert analyst time are invested in analyzing potential false signals, the cost of such investigations can be deemed a financial detection cost against which the operational effectiveness metric is defined. The rationale is that an event-driven architecture may not require machine intelligence but will nonetheless generate a non-zero analyst time detection cost that will grow with additional false positives.

**6.2. Operational Metrics and Cost Considerations** Operational metrics inform the efficiency and adequacy for statutory and organizational purposes of event-driven compliance systems. They also govern costs. Two kinds of cost are relevant: recurrent costs for maintaining the system, including the associated work and automated resources, and unit costs for inspections and audits generated by the system. Since the systems assume a self-informing organizational posture, detection resources must ensure that the investment in inspection resources typically externalized is reasonable. One means of ensuring reasonable investment is to cap false positive rates and the other is to interrogate the bulk of detections. A third, usually implicit, test relates to how quickly a compliance system has a fair chance at announcing a detection of required action (an alert) without generating excessive false alarms.

Evaluations conducted in real time of two event-driven compliance systems illustrate some implications for recurrent costs and operational metrics. False peak rates provide an idea of whether the operational staff deployments are reasonable for keeping the cost of external inspection within reason. Detection rates per month are thus divided by the size of the current engagement and also weighed against a cap of 10–15 % of the total outgoing to third parties. Lastly, the ratio of alerts to other detections indicates the effectiveness of ongoing tuning efforts in scaling-up cover for cases requiring immediate actions.

### **6.3. Post-Detection Investigation Outcomes**

Evaluation of Effectiveness: Post-Detection Investigation Outcomes

For any detection strategy, the ultimate goal is to prevent financial crime in the broader interest of society. A useful way to evaluate detection systems that do not rely on machine intelligence is to analyze outcomes of post-detection investigations. This is an obvious yet often neglected line of inquiry. For example, events flagged as potential financial crime by established detection systems — those using machine intelligence or rule-based methods — represent a large subset of all detected instances. The next stage of these detections most commonly involves case management and assignment for investigation and review. In such cases, the higher the number of investigated detections that are rated as “false positives” (with no evidence of financial crime), the more ineffective the system may appear.

Even with the high number of investigated negative detections, an uptick in detection assignment volume may support the need for deeper analysis on subsequently detected signals. As detectives and investigators manually explore more candidate signals, perhaps flagged potential financial activity representing a sudden accumulating capital position towards a specific known high-risk address or address generating unusual wealth, trends may begin to emerge. Neither the triggering rule nor the triggering heuristic detection heuristic may have raised the combination of signals to a detection borderline, yet the underlying financial traffic may nonetheless be associated with flagged behaviour. Over time, conversation with detectives and management may even result in creating a high-risk category, growing rules pertaining to higher-risk-change, or developing new heuristic checks.

## **7. Comparative Analysis: Event-Driven Versus Machine Intelligence Approaches**

Live-world developments, particularly advancements in artificial intelligence and machine learning, have prompted much interest in machine intelligence-based detection solutions. This fascination is not surprising, given the apparent potential. Yet the application of machine intelligence requires time and resources: data must be prepared, clear rules defined, and models trained. These demands may not always

be feasible or necessary. During the past five years, several events highlighted or prompted the addition of new detection rules — technically straightforward additions to an event-driven compliance architecture — and the immediate need for enforcement teams to work faster and smarter. During these increasingly urgent periods, organizations turned to the event-driven detection framework.

Machine intelligence-based detection systems support rules, thresholds, and portals for the creation of schematic visualizations. Yet real-time detection architectures more commonly employ rules and thresholds as they process combinations of event streams. Detection in an event-driven architecture can incorporate organizational knowledge through both heuristic rules and combinations that correlate and threshold other signals — a powerful complement and precursor to exploratory query-based machine intelligence systems. Human beings may be placed within the loop when the potential consequences call for multi-dimensional consideration. Moreover, event-driven detection architectures are often coupled with case management and triage workflows.



**Fig 6: Event-Driven Security**

**7.1. Strengths and Limitations of Event-Driven Systems** Event-driven compliance systems offer unique strengths and limitations, distinguishing them from approaches employing machine intelligence. Key strengths include improved detection performance, increased cost-effectiveness, and enhanced governance, risk, and compliance alignment; meanwhile, limitations encompass a smaller volume and scope of detected signals, a reliance on design breadth and depth, and heightened operational strain.

Signals identified by an event-driven strategy generally benefit from operational-level correlation. Conversely, an emphasis on qualitative data analysis over quantitatively driven decision-assistance techniques often results in a higher total human cognition effort than would an alternative approach that discerns fewer but clearer signals. Event-driven compliance thus often functions as a highly efficient human-in-the-loop analytical process instead of a fully autonomous machine-intelligence-based system. Investigative performance efficiency is likely to improve when either the quantity of detected signals approaches operational limit or a non-compliance signal detection is genuinely informative at a depth beyond the operational-layer processing thresholds deployed.

Signal detection scenarios where event-driven compliance systems shine are typically those aligned with the known detection advantages of human intelligence in cognition-demanding contexts—especially those potentially benefiting from qualitative hindsight exploring easy-to-gather but usually viewed as too-little-available (or potentially negative-value-adding) situational details. These could equally include human-in-the-loop qualification, validation, and explanation beyond a core case management platform. Satisfactory performance levels at cost-effective levels distinguish event-driven financial crime compliance systems, especially within horizontal levels of detection where either the machine-intelligence-based alternatives are commercially unviable or the supporting financing and risk appetite are not apparent.

**7.2. Scenarios Where Human-In-The-Loop Enhances Outcomes** Consider a fictitious financial services organization that has implemented an event-driven compliance system to complement existing machine intelligence technology. The hypothetical institution serves domestic customers as well as various authorized overseas locations. The conventional approach executes batch model-based detection in a separate country to identify remittance suspicious activity.

The event-driven system deploys the detection in the source country in real-time. The event is manually marked as suspicious, then executed. Certification of all remittances above \$2,500 is automatic; transactions between friends and family are not automatically marked as suspicious; remittances to prohibitive country are monitored. A combination of risk rating of customer relationship, volume of remittance, and sum of remittance over three month period is used to detect unusual behaviour.

The event-driven system allows for a greater number of alerts on remittance, at a far lesser false positive rate; and this would potentially lead to smarter investigations — as now all these higher risk transactions are scrutinized, rather than filtering with a batch approach, followed by an automatic execution.

## 8. Conclusions

The financial crime detection landscape has evolved considerably over the past half-century, yet the desire for a ‘machine’ capable of performing the task with both high accuracy and low cost endures. However, historical data and lessons from other domains suggest that the task may not be amenable to totally autonomous solutions. A case is made for modernizing detection capabilities without the use of machine intelligence by framing the problem as one of event-driven compliance.

The key strength of event-driven compliance lies in its combination of rule-based detection strategies, human expertise, real-time processing, and operational resources capable of managing high volumes of alerts. It allows organizations to strategically reduce noise while remaining compliant with external governance requirements, thereby reallocating costs toward innovation and developing more resilient business models. Financial services organizations pursuing event-driven compliance systems can aim to retain detection capabilities for regulatory alignment and external assurance, but seek to flag, record, and prioritize alerts rather than independent signals. When combined with automated audit trails, organization-wide data-quality constraints, and carefully curated datasources, these measures serve to foster greater alignment with the principles of auditability, explainability, and transparency now expected for other third-party-deployed machine-intelligence-based services.

**8.1. Emerging Trends** Emerging trends offer a particular opportunity for financial crime detection systems. eID, or verified identity attributes sourced from a trusted origin, are expected to increase in adoption. For example, using Verifiable Credentials to retrieve definitive identity information from eIDAS allows detection rules, heuristics, and thresholds to change from best effort to immutable. In addition, more actuarially sound crime risk indicators are emerging. Calibrated data from prudential supervision and financial consumer protection agencies, combined with data from law enforcement, allow for creation of risk signals that predict when FIs are likely to perpetrate, and not just when they are likely to be victims of, financial crime. Finally, the strategies of first-line defence are becoming available in product form, allowing clients to procure crime-detection Signals, Alerts, and Triage Workflows as dedicated services.

In conclusion, Event-Driven Compliance Systems propose a modernised detection architecture that does not use machine intelligence. The adoption of EDA dovetails with risk-based approaches, reduces false positives without sacrificing detection performance, and enables predictive crime-detection Signals. Emerging trends in verified identity information, crime-risk indicators and first-line-defence strategies enrich the system’s offerings.

## 9. References

- [1] Abbasi, M., Alhussaini, M., & Aljohani, N. R. Optimizing database performance in complex event processing systems: Indexing strategies and real-time stream analytics. *Data*, 9(8), 93.
- [2] Kummari, D. N. (2021). Smart Infrastructure Auditing: Integrating AI to Streamline Manufacturing Compliance Processes. *Journal of International Crisis and Risk Communication Research*, 168-193.
- [3] Allen, J. S., Clark, J., & O'Neill, T. Can LLMs improve sanctions screening in the financial system? (Finance and Economics Discussion Series). Board of Governors of the Federal Reserve System.
- [4] Arner, D. W., Barberis, J., & Buckley, R. P. (2017). FinTech, RegTech, and the reconceptualization of financial regulation. *Northwestern Journal of International Law & Business*, 37(3), 371–413.
- [5] Sheelam, G. K., & Nandan, B. P. (2021). Machine Learning Integration in Semiconductor Research and Manufacturing Pipelines. *International Journal of Advanced Research in Computer and Communication Engineering (IJARCCE)*, DOI, 10.
- [6] Association for Financial Markets in Europe. (2021). Anti-money laundering transaction monitoring in the capital markets: A survey of current practices and emerging trends. AFME.
- [7] Australian Transaction Reports and Analysis Centre. (2020). Anti-money laundering and counter-terrorism financing: Guidance on transaction monitoring and suspicious matter reporting. AUSTRAC.
- [8] Inala, R. (2021). A New Paradigm in Retirement Solution Platforms: Leveraging Data Governance to Build AI-Ready Data Products. *Journal of International Crisis and Risk Communication Research*, 286-310.
- [9] Basel Committee on Banking Supervision. (2021). Principles for operational resilience. Bank for International Settlements.
- [10] Basel Committee on Banking Supervision. Basel Core Principles for effective banking supervision (updated). Bank for International Settlements.
- [11] Meda, R. (2021). Digital Infrastructure for Predictive Inventory Management in Retail Using Machine Learning. *International Journal of Advanced Research in Computer and Communication Engineering (IJARCCE)*, DOI, 10.
- [11] Carbone, P., Katsifodimos, A., Ewen, S., Markl, V., Haridi, S., & Tzoumas, K. (2015). Apache Flink: Stream and batch processing in a single engine. *IEEE Data Engineering Bulletin*, 38(4), 28–38.
- [12] Aitha, A. R. (2021). Optimizing Data Warehousing for Large Scale Policy Management Using Advanced ETL Frameworks.
- [13] Council of the European Union. (2018). Directive (EU) 2018/843 of the European Parliament and of the Council (Fifth Anti-Money Laundering Directive). Official Journal of the European Union.
- [14] Council of the European Union. Regulation (EU) on the prevention of the use of the financial system for the purposes of money laundering or terrorist financing. Official Journal of the European Union.
- [15] Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
- [16] Egmont Group of Financial Intelligence Units. (2020). Operational guidance for FIU information exchange. Egmont Group.
- [17] Egmont Group of Financial Intelligence Units. (2021). FIU-FinTech cooperation: Typologies and case studies. Egmont Group.
- [18] European Banking Authority. (2021). Guidelines on risk-based supervision in the context of anti-money laundering and countering the financing of terrorism. EBA.
- [19] European Banking Authority. . Guidelines on the use of remote customer onboarding solutions under AML/CFT. EBA.

- [20] Gottimukkala, V. R. R. (2021). Digital Signal Processing Challenges in Financial Messaging Systems: Case Studies in High-Volume SWIFT Flows.
- [21] European Central Bank. Supervisory priorities ECB Banking Supervision.
- [22] European Commission. (2021). A package to fight money laundering and terrorist financing: Proposal and impact assessment. European Commission.
- [23] European Commission. EU anti-money laundering authority (AMLA): Establishment and mandate (policy paper). European Commission.
- [24] Europol. (2020). Internet organised crime threat assessment (IOCTA) 2020. Europol.
- [25] Europol. Financial intelligence public-private partnership: Strategic report. Europol.
- [26] Financial Action Task Force. (2019). Risk-based approach guidance for the banking sector. FATF.
- [27] Financial Action Task Force. (2020). Guidance on digital identity. FATF.
- [28] Financial Action Task Force. (2021). Updated guidance for a risk-based approach to virtual assets and virtual asset service providers. FATF.
- [29] Financial Action Task Force. (2021). Opportunities and challenges of new technologies for AML/CFT. FATF.
- [30] Financial Action Task Force. Best practices on the use of legal arrangements. FATF.
- [31] Gottimukkala, V. R. R. (2020). Energy-Efficient Design Patterns for Large-Scale Banking Applications Deployed on AWS Cloud. *power*, 9(12).
- [32] Financial Action Task Force. International standards on combating money laundering and the financing of terrorism & proliferation: The FATF recommendations (updated). FATF.
- [33] Financial Conduct Authority. (2021). Financial crime guide: A firm's guide to countering financial crime risks (updated). FCA.
- [34] Financial Conduct Authority. Thematic review: Transaction monitoring and sanctions screening in retail banking (selected findings). FCA.
- [35] Financial Crimes Enforcement Network. (2020). Advisory on ransomware and the use of the financial system to facilitate ransom payments. FinCEN.
- [36] Financial Crimes Enforcement Network. (2021). Advisory on kleptocracy and foreign public corruption. FinCEN.
- [37] Segireddy, A. R. (2020). Cloud Migration Strategies for High-Volume Financial Messaging Systems.
- [38] Financial Stability Board. (2020). The use of RegTech and SupTech: A primer. FSB.
- [39] Fragkoulis, M., Carbone, P., Kalavri, V., & Katsifodimos, A. A survey on the evolution of stream processing systems. *The VLDB Journal*, 33(1), 1–45.
- [40] Gartner. (2020). Market guide for anti-money laundering solutions. Gartner.
- [41] Giatrakos, N., Alevizos, E., Artikis, A., Deligiannakis, A., Garofalakis, M., & Paliouras, G. (2020). Complex event recognition in the big data era: A survey. *The VLDB Journal*, 29(1), 313–352.
- [42] Global Legal Entity Identifier Foundation. LEI in the fight against financial crime: Policy report. GLEIF.
- [43] Goodell, J. W. (2020). COVID-19 and finance: Agendas for future research. *Finance Research Letters*, 35, 101512.
- [44] Group of Thirty. (2020). Digital currencies and stablecoins: Risks, opportunities, and policy priorities. Group of Thirty.
- [45] Herley, C., & Florêncio, D. (2010). Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. *Economics of Information Security and Privacy*, 33–53.
- [46] Segireddy, A. R. (2021). Containerization and Microservices in Payment Systems: A Study of Kubernetes and Docker in Financial Applications. *Universal Journal of Business and Management*, 1(1), 1–17. Retrieved from <https://www.scipublications.com/journal/index.php/ujbm/article/view/1352>
- [47] International Organization for Standardization. ISO 37301: Compliance management systems—Requirements with guidance for use. ISO.

- [48] International Organization for Standardization. ISO/IEC 27001: Information security management systems—Requirements (3rd ed.). ISO.
- [49] International Telecommunication Union. (2020). Security aspects of digital identity (DID) ecosystems: Technical report. ITU.
- [50] Amistapuram, K. (2021). Digital Transformation in Insurance: Migrating Enterprise Policy Systems to .NET Core. *Universal Journal of Computer Sciences and Communications*, 1(1), 1–17. Retrieved from <https://www.scipublications.com/journal/index.php/ujcsc/article/view/1348>
- [51] International Monetary Fund. Anti-money laundering and combating the financing of terrorism: Review of the Fund's strategy. IMF.
- [52] Jain, R., Mahanti, A., & Dutta, K. (2021). Governance and accountability in algorithmic decision-making systems: A review. *Information Systems Frontiers*, 23(6), 1331–1350.
- [53] Rongali, S. K. (2020). Predictive Modeling and Machine Learning Frameworks for Early Disease Detection in Healthcare Data Systems. *Current Research in Public Health*, 1(1), 1-15.
- [54] Kleppmann, M. (2017). Designing data-intensive applications: The big ideas behind reliable, scalable, and maintainable systems. O'Reilly Media.
- [55] Rongali, S. K. (2021). Cloud-Native API-Led Integration Using MuleSoft and .NET for Scalable Healthcare Interoperability. *Journal for ReAttach Therapy and Developmental Diversities*, 4(2), 181-192.
- [56] Kshetri, N. (2021). Cybersecurity management: An organizational and strategic perspective. *Telecommunications Policy*, 45(2), 102122.
- [57] Lazzari, L., Farias, K., & de Souza, C. Event-driven architecture and REST architectural style: A comparative study for distributed systems integration. *Journal of Applied Research and Technology*, 20(2), 1–15.
- [58] Vadisetty, R., Polamarasetti, A., Guntupalli, R., Raghunath, V., Jyothi, V. K., & Kudithipudi, K. (2021). Privacy-Preserving Gen AI in Multi-Tenant Cloud Environments. Sateesh kumar and Raghunath, Vedaprada and Jyothi, Vinaya Kumar and Kudithipudi, Karthik, Privacy-Preserving Gen AI in Multi-Tenant Cloud Environments (January 20, 2021).
- [59] Lewis, R., McKee, K., & Carta, S. (2020). Explainability in risk and compliance systems: A practical survey. *Journal of Financial Regulation and Compliance*, 28(4), 523–540.
- [60] Varri, D. B. S. (2021). Cloud-Native Security Architecture for Hybrid Healthcare Infrastructure. Available at SSRN 5785982.
- [61] National Institute of Standards and Technology. (2020). Security and privacy controls for information systems and organizations (SP 800-53 Rev. 5). NIST.
- [63] Vadisetty, R., Polamarasetti, A., Guntupalli, R., Rongali, S. K., Raghunath, V., Jyothi, V. K., & Kudithipudi, K. (2021). Legal and Ethical Considerations for Hosting GenAI on the Cloud. *International Journal of AI, BigData, Computational and Management Studies*, 2(2), 28-34.
- [64] Office of Financial Sanctions Implementation. Monetary penalties for breaches of financial sanctions: Guidance (updated). HM Treasury (UK).
- [65] Organisation for Economic Co-operation and Development. (2021). The role of financial intelligence units in combating tax crimes and other financial crimes. OECD.
- [66] Oztas, B. Transaction monitoring in anti-money laundering: Challenges, requirements, and future trends. *Future Generation Computer Systems*, 158, 1–16.
- [67] Varri, D. B. S. (2020). Automated Vulnerability Detection and Remediation Framework for Enterprise Databases. Available at SSRN 5774865.
- [68] Rossi, M., Mueller, J., & Huber, M. (2021). Process mining in financial services: Compliance monitoring and operational risk insights. *Decision Support Systems*, 150, 113561.
- [69] Roldán, J., Boubeta-Puig, J., Martínez, J. L., & Ortiz, G. (2020). Integrating complex event processing and machine learning: An intelligent architecture for detecting IoT security attacks. *Expert Systems with Applications*, 149, 113251.
- [70] Securities and Exchange Board of India. Master circular on anti-money laundering standards and combating financing of terrorism. SEBI.

- [71] Yandamuri, U. S. (2021). A Comparative Study of Traditional Reporting Systems versus Real-Time Analytics Dashboards in Enterprise Operations. *Universal Journal of Business and Management*.
- [72] Sivarajah, U., Irani, Z., Gupta, S., & Mahroof, K. (2020). Role of big data and analytics in compliance and risk management: A review and research agenda. *Information Systems Frontiers*, 22(3), 1–24.
- [73] Sweeney, S., & Cuganesan, S. (2019). RegTech and the transformation of compliance: A socio-technical perspective. *Accounting, Auditing & Accountability Journal*, 32(8), 1–28.
- [74] The Wolfsberg Group. (2019). Wolfsberg guidance on correspondent banking due diligence. Wolfsberg Group.
- [75] Keerthi Amistapuram , "Energy-Efficient System Design for High-Volume Insurance Applications in Cloud-Native Environments," *International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering (IJREEICE)*, DOI 10.17148/IJREEICE.2020.81209
- [76] The Wolfsberg Group. Wolfsberg guidance on sanctions screening. Wolfsberg Group.
- [77] Botlagunta, P. N., & Sheelam, G. K. (2020). Data-Driven Design and Validation Techniques in Advanced Chip Engineering. *Global Research Development (GRD)* ISSN: 2455-5703, 5(12), 243-260.
- [78] Toshniwal, A., Taneja, S., Shukla, A., Ramasamy, K., Patel, J. M., Kulkarni, S., Jackson, J., Gade, K., Fu, M., Donham, J., Bhagat, N., Mittal, S., & Ryaboy, D. (2014). Storm@Twitter. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data* (pp. 147–156). ACM.
- [79] United Nations Office on Drugs and Crime. (2020). Global study on illicit financial flows. UNODC.
- [80] Aitha, A. R. (2021). Dev Ops Driven Digital Transformation: Accelerating Innovation In The Insurance Industry. Available at SSRN 5622190.
- [81] United States Department of the Treasury. (2021). Sanctions compliance guidance for the virtual currency industry. Treasury/OFAC.
- [82] Inala, R. (2020). Building Foundational Data Products for Financial Services: A MDM-Based Approach to Customer, and Product Data Integration. *Universal Journal of Finance and Economics*, 1(1), 1-18.
- [83] Vasylykiv, N., & Hrytsenko, V. Fuzzy compliance risk monitoring system. In *Proceedings of the CEUR Workshop Proceedings* (Vol. 3974, pp. 1–12). CEUR-WS.
- [84] Meda, R. (2020). Data Engineering Architectures for Real-Time Quality Monitoring in Paint Production Lines. *International Journal Of Engineering And Computer Science*, 9(12).
- [85] Wamba, S. F., Queiroz, M. M., & Wu, L. (2020). Big data analytics-enabled sensing capability and operational performance: Evidence from financial services. *International Journal of Information Management*, 50, 1–13.
- [86] Wang, D., Franklin, M. J., & Garofalakis, M. (2011). Active complex event processing over event streams. *Proceedings of the VLDB Endowment*, 4(10), 634–645.
- [87] World Bank. (2020). Financial consumer protection and new forms of digital finance: Supervisory and compliance considerations. World Bank.
- [88] Kummari, D. N. (2021). A Framework for Risk-Based Auditing in Intelligent Manufacturing Infrastructures. *International Journal on Recent and Innovation Trends in Computing and Communication*, 9(12), 245-262.
- [89] Yıldız, A., Demirörs, O., & Sözer, H. Event-driven analysis and design method for microservices. *Procedia Computer Science*, 232, 1–12.
- [90] Zetzsche, D. A., Buckley, R. P., Arner, D. W., & Barberis, J. N. (2020). Regulating Libra: The transformational potential of digital stablecoins. *Oxford Journal of Legal Studies*, 40(1), 1–30.