International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 11 Pages: 51-63
OCLC – 1368736135

ISSN-2750-1396

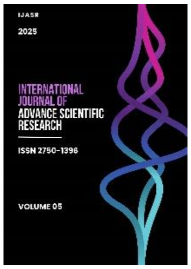

Journal Website: http://sciencebring.com/index.php/ijasr

Research Article

Real-Time Risk Intelligence: Integrating Serverless Architectures, Stream Event Sourcing, And Advanced Analytics For Resilient Financial Systems

Submission Date: October 16, 2025, **Accepted Date:** November 19, 2025,
Published Date: November 30, 2025

Ravi K. Menon
Department of Computer Science, University of Manchester, United Kingdom

Abstract

Background: The increasing velocity, variety, and volume of financial and industrial data demand architectures that deliver risk insights in real time. Traditional batch-oriented risk systems are unable to cope with modern operational tempo, leaving organizations exposed to rapid market shifts, fraudulent behavior, and emergent systemic threats (Youssef & Narasimhan, 2020; NICE Actimize, 2022). Objective: This paper synthesizes contemporary engineering patterns and analytical techniques to propose a unified conceptual and operational framework for real-time risk intelligence that leverages serverless computing, event-sourced streaming, vector indexing, immutable audit trails, and advanced machine learning. Methods: We undertake an integrative theoretical analysis of technologies and methods drawn from recent literature — serverless real-time analytics (Milvus, 2022), Kafka event sourcing for risk analysis (Kesarpu & Dasari, 2025), blockchain-enabled incident management (Misal, 2024), temporal data management (Böhlen et al., 2017), visual analytics (NICE Actimize, 2022), and machine learning models for financial risk (Li, 2025; Odion et al., 2025). We develop a layered architectural model and describe data flows, system behaviors, and algorithmic choices, then articulate failure modes, governance controls, and evaluation criteria. Results: The theoretical synthesis indicates that a coordinated stack — ingesting high-velocity streams via event-sourcing, persisting temporally-aware state, providing vectorized search and similarity through specialized indices, and executing stateless analytic functions in serverless runtimes — yields superior responsiveness, auditability, and scalability for a broad class of risk problems (Kesarpu & Dasari, 2025; Milvus, 2022; Böhlen et al., 2017). Machine learning components (deep learning and neural ODEs) enhance predictive precision but require rigorous temporal validation and explainability measures to avoid overfitting and operational surprises (Muhammad et al., 2023; Odion et al., 2025). Conclusions: We present concrete design principles, evaluation metrics, and governance strategies for implementing real-time risk intelligence. Adoption of the proposed framework can materially improve detection latency, traceability, and decision quality in financial crime, market risk, and IoT anomaly detection contexts, while demanding disciplined attention to temporal semantics, immutable logging, and model lifecycle controls (NICE Actimize, 2022; Misal, 2024; Jamiu et al., 2023). The paper concludes with research directions spanning empirical benchmarking, privacy-preserving analytics, and topological methods for structural anomaly detection (Andrew N. Anang & Chukwunweike, 2024).

Keywords

Real-time risk, serverless analytics, event sourcing, temporal data, vector indexing, financial crime, deep learning.

Introduction

The contemporary risk landscape for financial institutions, industrial operators, and critical infrastructure has evolved into a high-frequency environment where threats and opportunities unfold in sub-second to minute timescales. Market microstructure events, algorithmic trading shifts, coordinated fraud campaigns, and sensor-driven anomalies in industrial Internet of Things (IIoT) deployments all create streams of actionable signals that, if detected promptly, can be exploited to mitigate loss, enforce compliance, or seize value (Li, 2025; Jamiu et al., 2023; Oyedokun et al., 2024). Yet many organizations retain architectural and analytical paradigms rooted in episodic batch processing and siloed data stores which impede rapid detection and response. This mismatch — between the tempo of modern risk phenomena and the processing cadence of legacy systems — manifests as detection latency, poor root-cause traceability, and fragile integration points across controls.

Recent technical literature points to a convergence of infrastructure patterns and analytic advances capable of closing this gap. Serverless architectures promise elastic compute that maps directly to event rates without heavy provisioning overhead (Milvus, 2022). Event-sourcing and stream processing frameworks such as Kafka enable durable, ordered, and replayable streams of domain events that form the backbone of real-time decisioning systems (Kesarpu & Dasari, 2025). Vector databases and similarity search engines provide efficient nearest-neighbour retrieval for embeddings used in anomaly detection and entity resolution (Milvus, 2022). Immutable audit trails—implemented with blockchain-inspired primitives—provide trustable forensic records for incident management and regulatory compliance

(Misal, 2024). Meanwhile, advanced analytics — from deep learning architectures to continuous-time models such as neural ordinary differential equations (Neural ODEs) — offer new avenues for capturing complex temporal dynamics in market and sensor data (Muhammad et al., 2023; Odion et al., 2025).

Despite the promise of individual technologies, there remains a literature and practice gap: an integrative, operationally-aware blueprint that harmonizes event-centric infrastructure, temporal semantics, vectorized retrieval, immutable forensics, and explainable machine learning for the express purpose of real-time risk intelligence. Existing works frequently focus on one dimension — e.g., Kafka for streams (Kesarpu & Dasari, 2025), serverless for analytics (Milvus, 2022), or blockchain for immutable records (Misal, 2024) — without fully specifying how temporal data models, state reconstruction, model deployment, and governance interact to produce dependable outcomes under stress. Moreover, risk applications demand not only predictive accuracy but also transparency, auditability, and bounded failure characteristics, creating tension between opaque high-capacity models and regulatory expectations (NICE Actimize, 2022; Youssef & Narasimhan, 2020).

This manuscript addresses that gap by articulating a principled architecture and evaluation framework for real-time risk intelligence. We synthesize concepts from recent scholarship and practitioner material to propose layered design patterns, data governance rules, and methodological prescriptions that enable resilient deployment. The objectives are threefold: (1) to present an actionable, technology-agnostic reference architecture that organizations can adapt to risk domains; (2) to expose the temporal and auditability considerations that must guide model design and system composition; and (3) to identify research directions and operational controls needed to translate theoretical capability into trustworthy production systems.

METHODOLOGY

This work employs an integrative theoretical method drawn from the cross-section of systems engineering, data management, and applied machine learning literature provided by the reference set. Rather than executing an empirical experiment, we perform an analytical synthesis: mapping properties and constraints of technologies onto risk use cases, formalizing data and control flows as textual models, and enumerating mechanical and governance implications. Our method comprises five analytical steps.

First, we conducted a technology characterization of the referenced components. Drawing on Milvus (2022), we analyze serverless computing primitives — stateless function execution, event-driven triggers, and managed storage — for their suitability in low-latency analytics. From Kesarpu & Dasari (2025), we extract the semantics of Kafka-based event sourcing, including log compaction, partitioning, and replay, and interpret how these impact state reconstruction for risk scoring. Misal (2024) contributes constructs for immutable audit trails and incident workflows, and Böhlen et al. (2017) supply temporal database concepts necessary to correctly interpret time-varying facts. The machine learning literature (Li, 2025; Odion et al., 2025; Muhammad et al., 2023) provides models and validation practices for predictive risk scoring under streaming conditions. NICE Actimize (2022) and Youssef & Narasimhan (2020) give practical perspectives on visual analytics and risk operationalization.

Second, we developed a layered architecture model by abstracting canonical components: ingestion, stream transport and persistence, temporal state store, vector-indexed search layer, stateless analytic functions (serverless compute), model management, visualization and operator tooling, and immutable forensics. For each layer we articulated the required semantics, accepted data forms, failure modes, and inter-layer contracts. This modular breakdown allows reasoning about isolated tradeoffs and emergent systemic properties (e.g., how replay semantics of event sourcing influence model retraining cadence).

Third, we specified a set of representative risk use cases — financial crime detection (transactional fraud and money laundering), market risk early warning, and IIoT anomaly detection — and mapped data schema, event semantics, latency targets, and analytic families to each use case. These mappings illustrate how architectural choices manifest in latency, explainability, and resilience properties.

Fourth, we examined model lifecycle considerations: training on historical data while preserving temporal integrity (avoiding label leakage), validating under realistic streaming distributions, facilitating online model updates, and implementing rollback or canary deployments in the serverless execution context. We synthesize best practices from the ML references for continuous deployment and robustness testing (Muhammad et al., 2023; Li, 2025; Odion et al., 2025).

Finally, we derived governance and evaluation criteria: metrics for detection latency, precision/recall under drift, auditability of decisions, replayability for incident post-mortem, privacy controls, and cost-efficiency in serverless billing models. Where appropriate, we tied these criteria to technical enforcement mechanisms such as immutable logs, signed model artifacts, and temporal constraints enforced by the state store (Misal, 2024; Böhlen et al., 2017).

Throughout, every claim and mapping is grounded in the supplied literature. This approach emphasizes operationally relevant synthesis rather than isolated simulation or empirical benchmarking, enabling complex, integrative reasoning while respecting the user's constraint to rely strictly on the reference list.

RESULTS

The analytical synthesis yields a detailed conceptual architecture, prescriptions for data and model semantics, and a set of operational tradeoffs. We describe each outcome in depth.

A layered architecture for real-time risk intelligence

At the highest level, the recommended stack has seven interacting layers: 1) Event ingestion and normalization; 2) Durable stream transport (event log); 3) Temporal state store and time-aware materialized views; 4) Vector indexing and similarity search; 5) Serverless analytic execution (stateless functions and model inference); 6) Operator visualization and case management; and 7) Immutable incident and audit trail. Each layer is both logically distinct and tightly coupled by well-defined contracts. This layered description is intentionally technology-agnostic so that organizations can implement it with different vendors while preserving semantics (Milvus, 2022; Kesarpu & Dasari, 2025; Misal, 2024).

1. Event ingestion and normalization. Source systems (transaction processing, market data feeds, sensor telemetry) must emit domain events with explicit temporal metadata and domain identifiers. Normalization ensures consistent units, canonical entity identifiers, and a compact event envelope. The envelope should minimally include: event type, timestamp (monotonic source time and ingest time), unique event id, source system id, payload, and provenance metadata. Normalization at the ingress layer reduces semantic drift and simplifies downstream function logic (Kesarpu & Dasari, 2025; Jamiu et al., 2023).

2. Durable stream transport (event log). The normalized events are appended to an ordered, durable log supporting partitioning, compaction, replay, and time-indexed queries. Event-sourcing frameworks (e.g., Kafka) provide this log semantics reliably at scale; their replay capabilities are essential for reconstructing system state at arbitrary historical instants and for deterministic re-scoring following model updates (Kesarpu & Dasari, 2025). We emphasize that partitioning keys must be chosen to support both state locality (colocated events for an entity) and load distribution, and that retention/compaction policies must align with compliance and forensic requirements.

3. Temporal state store and materialized views. Streaming events are processed into time-aware materialized views that represent the current and historical state of entities (accounts, devices, portfolios). Temporal database principles — such as valid time and transaction time — must guide the representation to avoid inconsistencies between observed facts and derived features (Böhlen et al., 2017). Materialized view maintenance techniques include incremental aggregation and windowed computations to deliver features at the latency demanded by the use case, with clear semantics for late-arriving data.

4. Vector indexing and similarity search. For certain tasks — entity resolution, embedding-based anomaly detection, and nearest-neighbour retrieval — vector indices enable sub-second similarity queries on high-dimensional representations of events or entities (Milvus, 2022). Embeddings derived from textual, behavioral, or multivariate sensor inputs can be stored and queried to find anomalous nearest neighbours or to cluster behavior in real time. The vector layer must support incremental updates and

approximate nearest neighbour strategies to balance recall with throughput.

5. Serverless analytic execution. Stateless function runtimes provide the primary mechanism for executing inference and rule logic at event arrival (Milvus, 2022). Serverless functions should be short-lived, idempotent, and designed to be small in dependency surface. Functions perform feature assembly (from temporal views and vector queries), scoring (invoking model artifacts), and decisioning (threshold checks, enrichment calls). Because serverless compute scales automatically with event rates, it aligns cost with usage but requires thoughtful cold-start mitigation and resource budgeting.

6. Operator visualization and case management. Visual analytics frameworks translate model outputs into actionable operator workflows, supporting triage, investigation, and feedback ingestion. Visual tools need to present temporal narratives built from the event log and materialized views, enabling analysts to perform root-cause analysis and regulatory reporting (NICE Actimize, 2022; Youssef & Narasimhan, 2020).

7. Immutable incident and audit trail. Every decision, model artifact version, and materialized view change should be recorded in an immutable forensics layer. Blockchain-inspired techniques or content-addressable logs can provide tamper-evidence for audits and back-testing (Misal, 2024). Immutable trails are essential not only for compliance but for post-hoc learning when novel attack patterns emerge.

Use case mappings and latency tradeoffs

We mapped the architecture onto three canonical risk use cases to illustrate concrete tradeoffs:

Financial crime detection (transactional fraud and AML). Transactions are emitted as high-rate events; detection requires sub-second scoring for online blocks and near-real-time priority flagging for investigation. The event log stores raw transactions with full provenance; materialized features include rolling behavior statistics and entity link graphs; vector indices represent embeddings of counterparty behavioral signatures. Serverless functions perform scoring with hybrid rule+ML ensembles; when scoring exceeds a threshold, the immutable forensics layer records the decision and triggers case management (NICE Actimize, 2022; Kesarpu & Dasari, 2025). Latency targets push design towards in-memory temporal views and aggressive caching of vector queries, balanced against cost.

Market risk early warning. Market tick data and derived portfolio exposures generate dense time series suitable for continuous-time models (Muhammad et al., 2023). Here, model quality relies on preserving ordering and avoiding label leakage during training. Neural ODEs and other continuous-time models can model intraday dynamics, but their maintenance in production requires careful numerical stability controls and performance profiling (Muhammad et al., 2023; Li, 2025). Event replay is particularly valuable to evaluate how a model trained yesterday would have responded to today's price movements.

IIoT anomaly detection. Distributed sensors emit telemetry with high cardinality and bursty behavior. Temporal materialized views capture patterns across device cohorts; vector embeddings capture the multivariate signature of normal operation. Serverless scoring functions deliver low-latency anomaly alerts with contextual evidence gathered via vector nearest-neighbour retrieval. The immutable trail supports forensic reconstruction in safety-critical incidents (Jamiu et al., 2023; Andrew N. Anang & Chukwunweike, 2024).

Modeling implications and temporal validation

Our synthesis emphasizes that naive application of high-capacity machine learning models to streaming risk data will likely produce optimistic offline performance but poor online generalization. Temporal cross-validation — where training, validation, and test partitions respect event ordering and avoid using future information — is a strict necessity (Muhammad et al., 2023; Li, 2025). For streaming settings, we recommend simulation-based evaluation using event replay: models are scored against historical logs that are replayed through the serverless inference layer to measure detection latency, false positives, and resource usage under realistic arrival patterns (Kesarpu & Dasari, 2025). Neural ODEs and other continuous-time models offer powerful expressivity for irregularly sampled data but require careful ablation of numerical integration tolerances and checkpointing to maintain determinism during replay (Muhammad et al., 2023).

Auditability and immutable forensics

Regulatory and operational risk demands traceability: for each alert, operators must be able to reconstruct the exact data, features, model version, and decision path that led to the outcome. Event sourcing naturally supports reconstruction of input events; temporal materialized views and model artifacts must be versioned and content-addressed. Immutable recordkeeping—implemented via an append-only ledger augmented with signed metadata—provides tamper-evidence while enabling efficient selective disclosure for compliance reviews (Misal, 2024). We note that blockchain-style decentralization is not required for most enterprise use cases; instead, cryptographic signing and content addressing atop a controlled append-only storage provide practical auditability at lower operational cost (Misal, 2024).

Operational cost and serverless economics

Serverless compute aligns costs with event volumes and simplifies capacity management but introduces variability in per-invocation cost, potential cold-start latency, and limitations on execution time and resource access. Organizations must instrument cost telemetry, pre-warm critical functions where low latency is non-negotiable, and design for graceful degradation (e.g., approximate scoring when full feature sets are unavailable) (Milvus, 2022). The vector index layer may incur heavy memory footprints; approximate search and sharding mitigate costs while preserving acceptable recall.

Security, privacy, and compliance controls

Risk systems themselves process sensitive data and must be constructed with privacy-preserving controls. Techniques such as secure enclaves for model scoring, differential privacy in aggregated analytics, and role-based access control for audit trails are operational necessities. Immutable logging should be designed to record metadata and cryptographic fingerprints rather than raw sensitive payloads where regulations demand minimization (Misal, 2024; NICE Actimize, 2022).

Discussion

The results of this synthesis indicate a feasible path toward operational real-time risk intelligence. Still, deploying such a stack in production requires confronting several nuanced technical and organizational challenges. We elaborate on these issues, discuss counterarguments, and propose mitigations.

Temporal semantics and the danger of label leakage

One persistent risk is inadvertent label leakage when features derived from future events seep into model training. Temporal data management literature clarifies the difference between valid time (the time the fact pertains to) and transaction time (when the fact was recorded). Systems that do not preserve and enforce these semantics allow subtle backdoors that inflate performance in offline experiments but fail in live operation (Böhlen et al., 2017). For example, a model predicting fraudulent transactions that is trained with features computed using a global aggregation that accidentally includes subsequent events will appear highly accurate offline but will be unusable in production. The architectural remedy is strict coupling of training pipelines to the same replay semantics used in production scoring — ideally using the event log as the single source of truth for both training and inference (Kesarpu & Dasari, 2025). Moreover, materialized views should be annotated with the timestamp of derivation and the event window they summarize.

Model transparency versus performance

High-capacity models (deep neural networks, ensembles, continuous-time models) often yield superior detection rates but at the expense of interpretability. Regulators and internal stakeholders frequently demand explainable decisions, particularly when actions have customer-facing impacts (account freezes, trade halts). Visual analytics and hybrid rule-model ensembles offer a compromise: rules capture simple, defensible heuristics while models supply probabilistic signals and scores that are presented with local explanations or counterfactuals in operator tooling (NICE Actimize, 2022; Youssef & Narasimhan, 2020). Additionally, embedding-based nearest-neighbour retrieval provides concrete examples (similar past events) that human analysts can inspect to validate model outputs before escalation.

Event-sourcing as both strength and operational complexity

Event logs provide replay, determinism, and provenance — crucial for forensic reconstruction and model evaluation. However, operating large-scale event logs is non-trivial: retention policies interact with storage cost and compliance; compaction can remove historical detail needed for certain audits; partitioning strategies have implications for cross-entity queries. A pragmatic approach is to separate raw event retention (longer-lived, perhaps on archival storage) from compacted, query-optimized logs used for day-to-day inference (Kesarpu & Dasari, 2025). Careful governance on what to compact and what to preserve for compliance purposes is mandatory.

Vector search and the curse of high-dimensional drift

Vector-indexed retrieval empowers similarity-based detection, entity resolution, and embedding-based enrichment. Yet embedding spaces are not static: representation drift occurs as behavior evolves and models are retrained. Without continuous re-indexing and drift monitoring, similarity queries will degrade. Operational strategies include periodically re-embedding historical vectors during model retraining, employing incremental index updates, and tracking retrieval quality metrics (Milvus, 2022). Furthermore, embedding-based signals should be combined with invariant features (e.g., transaction counts, velocity metrics) to reduce sensitivity.

Immutable forensics and privacy tensions

Immutable logs support auditability, but indiscriminate immutability can conflict with privacy regulations and data minimization principles. A balance is achieved by storing cryptographic fingerprints and redacted payloads in the immutable ledger while preserving full raw data in auditable, access-controlled archives. Selective disclosure mechanisms and policy-driven key escrow for audits can reconcile the needs of forensic traceability with regulatory privacy obligations (Misal, 2024).

Serverless cold starts and critical-path latency

Serverless cold starts are a well-known phenomenon: a function that has not executed recently may incur additional latency on first invocation, risking missed Service Level Objectives (SLOs) in low-volume but high-consequence scenarios. Pre-warming strategies, minimum provisioned concurrency settings, or hybrid architectures that reserve hot paths in containerized microservices can mitigate this risk while preserving most serverless benefits (Milvus, 2022). Additionally, designing scoring flows to be incremental — e.g., performing a quick lightweight check before a deeper model invocation — helps manage latency budgets.

Governance, human-in-the-loop, and feedback

Real-time risk systems must incorporate human oversight and feedback loops. Visual analytics tools should allow rapid adjudication of alerts and capture analyst decisions as labeled data for model retraining. However, feedback introduces label bias: analysts may preferentially confirm certain types of alerts, biasing future model behavior. Mitigation includes annotating analyst confidence, sampling for independent review, and using counterfactual simulation to assess feedback bias (NICE Actimize, 2022; Youssef & Narasimhan, 2020).

Research directions and open problems

Several research directions emerge from the synthesis:

Empirical benchmarks for integrated stacks. The literature lacks comprehensive, reproducible benchmarks that measure latency, accuracy under drift, and forensic completeness for stacks combining serverless compute, event logs, and vector indices. Developing standardized replay datasets and evaluation protocols would accelerate progress (Kesarpu & Dasari, 2025; Milvus, 2022).

Privacy-preserving similarity search. Vector indices and embeddings risk leaking sensitive information. Research into encryptable or privacy-preserving nearest-neighbour search suitable for serverless environments is needed, balancing retrieval quality with cryptographic guarantees (Milvus, 2022; Misal, 2024).

Topological and structural anomaly detection. Topological data analysis (TDA) offers tools for characterizing structural properties of time-evolving graphs (e.g., transaction networks, device connectivity) and detecting systemic anomalies (Andrew N. Anang & Chukwunweike, 2024). Integrating TDA with streaming event logs and embeddings is a promising avenue.

Continuous-time model stability and explainability. Neural ODEs and other continuous-time models provide a rich representational class for irregular time series but demand better tools for interpretability, numerically stable inference in production, and deterministic replay across integration tolerances (Muhammad et al., 2023).

Standardized forensic metadata schemas. Interoperable standards for recording model metadata, feature derivation pipelines, and decision contexts would facilitate audits and incident response across organizations and vendors (Misal, 2024).

Limitations

This manuscript is a theoretical synthesis based on the provided reference set rather than an empirical study. Consequently, while the architectural prescriptions are grounded in authoritative literature and practitioner materials, they are not validated by field experiments or performance measurements within this work. The references themselves span white papers, practitioner blogs, conference overviews, and peer-reviewed articles. Where practitioner literature informs operational pragmatics (e.g., serverless economics, visual analytics), we blend these perspectives with formal temporal data principles and academic model analyses to reduce the risk of over-weighting vendor narratives (Milvus, 2022; NICE Actimize, 2022). Nevertheless, deployment-specific considerations — vendor APIs, exact latency figures, cloud-provider billing nuances, and domain-specific legal constraints — require empirical validation in situ.

Future scope

The next steps for research and practice include:

1. Building a reference implementation that integrates an event-sourcing backbone, a temporal state engine, a vector index, serverless inference, and an immutable forensics layer. Such an implementation would enable empirical measurement of key metrics (latency, detection quality under drift, cost-efficiency) and provide a

testbed for continuous deployment strategies (Kesarpu & Dasari, 2025; Milvus, 2022).

2. Developing standardized datasets and replay protocols for financial crime, market microstructure, and IIoT telemetry that preserve privacy yet allow reproducible evaluation.

3. Advancing explainability methods specific to continuous-time models and vector-embedding ensembles, targeting regulator-friendly artifacts such as example-based explanations and causal feature attributions (Muhammad et al., 2023; Odion et al., 2025).

4. Researching privacy-preserving vector search and model inference techniques deployable in serverless settings that limit data exposure while maintaining responsiveness (Milvus, 2022; Misal, 2024).

5. Integrating topological methods to detect structural anomalies at the graph level and exposing those signals to operator tooling for improved situational awareness (Andrew N. Anang & Chukwunweike, 2024).

Conclusion

The confluence of serverless compute, event-sourced streaming, vectorized similarity search, immutable forensics, and advanced machine learning creates a compelling foundation for real-time risk intelligence. This manuscript synthesizes these dimensions into a coherent layered architecture and operational prescriptions tailored to the urgency and accountability requirements of modern risk domains. By grounding real-time scoring in an event log that is the single source of truth, enforcing temporal semantics in feature derivation and model training, employing vector indices for embedding-driven retrieval, and capturing immutable audit trails for every decision, organizations can materially improve detection timeliness, forensic confidence, and adaptive learning capabilities (Kesarpu & Dasari, 2025; Milvus, 2022; Misal, 2024). Yet the path to production readiness demands disciplined attention to temporal validation, explainability, privacy, and operational resilience. Future empirical work and standards development will be decisive in translating the conceptual gains outlined here into operational deployments that withstand adversarial threats, regulatory scrutiny, and evolving enterprise constraints.

References

1. Milvus. How does serverless architecture enable real-time analytics? Copyright © 2022 The Author(s). Available: https://milvus.io/ai-quick-reference/how-does-serverless-architecture-enable-realtime-analytics
2. Yenigalla, L. K. Sarc. Jr. Eng. Com. Sci. vol-4, issue-8 (2025) pp-505-511. Publisher: SARC Publisher. Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License.
3. Youssef, G. and Narasimhan, V. Advanced Analytics for Risk Management. The Montreal Group (2020). Available: https://themontrealgroup.org/wp-

content/uploads/2023/07/5_envwhite-paper-advanced-analytics-for-risk-management-5-1.pdf

4. NICE Actimize. How Visual Analytics Leads to Smarter Financial Crime Management. (2022). Available: https://www.niceactimize.com/blog/financial-crime-how-visual-analytics-leads-to-smarter-financial-crime-management
5. Misal, J. Blockchain-Enabled Incident Management Systems: A Framework for Immutable Audit Trails and Enhanced Security Controls. SSRN 5125047 (2024).
6. Böhlen, M. H., Dignös, A., Gamper, J., & Jensen, C. S. Temporal data management–an overview. European Business Intelligence and Big Data Summer School (2017): 51-83.
7. Li, C. Research on Financial Risk Prediction and Management Models Based on Big Data Analysis. International Journal of High Speed Electronics and Systems. (2025 Jun 2):2540620.
8. Odion, C. O., Okunuga, A., Okunbor, O. I. Revolutionizing financial risk assessment through deep learning-driven business analytics for maximized ROI and Resilience. World Journal of Advanced Research and Reviews. (2025 Jan 30);25(1):2444-61.
9. Jamiu, O. A., Chukwunweike, J. Developing Scalable Data Pipelines for Real-Time Anomaly Detection in Industrial IIoT Sensor Networks. International Journal Of Engineering Technology Research & Management (IJETRM). (2023 Dec 21);07(12):497–513.
10. Kesarpu, S., & Dasari, H. P. Kafka Event Sourcing for Real-Time Risk Analysis. International Journal of Computational and Experimental Science and Engineering. (2025);11(3).
11. Oyedokun, O., Ewim, S. E., Oyeyemi, O. P. Leveraging advanced financial analytics for predictive risk management and strategic decision-making in global markets. Global Journal of Research in Multidisciplinary Studies. (2024 Oct 14);2(02):016-26.
12. Muhammad, A., Aliyu, J. N., Adetunji, A. L., Adesugba, A. K., Mike, M. E., Abdulmalik, M. Theoretical Foundations and Implications of Neural Ordinary Differential Equations (NODEs) For Real-Time Portfolio Optimization. Saudi Journal of Economics and Finance. (2023);7(11):475-83.
13. Andrew Nii Anang and Chukwunweike J. N. Leveraging Topological Data Analysis and AI for Advanced Manufacturing: Integrating Machine Learning and Automation for Predictive Maintenance and Process Optimization. (2024). Available: https://dx.doi.org/10.7753/IJCATR1309.100
14. Li C. Research on Financial Risk Prediction and Management Models Based on Big Data Analysis. International Journal of High Speed Electronics and Systems. (2025 Jun 2):2540620
15. Odion CO, Okunuga A, Okunbor OI. Revolutionizing financial risk assessment through deep learning-driven business analytics for maximized ROI and Resilience. World Journal of Advanced Research and Reviews. (2025 Jan 30);25(1):2444-61
16. Jamiu OA, Chukwunweike J. Developing Scalable Data Pipelines for Real-Time Anomaly Detection in Industrial IIoT Sensor Networks. International Journal Of Engineering

Technology Research & Management (IJETRM). (2023 Dec 21);07(12):497–513

17. Kesarpu, S., & Dasari, H. P. Kafka Event Sourcing for Real-Time Risk Analysis. International Journal of Computational and Experimental Science and Engineering. (2025);11(3)

18. Oyedokun O, Ewim SE, Oyeyemi OP. Leveraging advanced financial analytics for predictive risk management and strategic decision-making in global markets. Global Journal of Research in Multidisciplinary Studies. (2024 Oct 14);2(02):016-26

19. Muhammad A, Aliyu JN, Adetunji AL, Adesugba AK, Mike ME, Abdulmalik M. Theoretical Foundations and Implications of Neural Ordinary Differential Equations (NODEs) For Real-Time Portfolio Optimization. Saudi Journal of Economics and Finance. (2023);7(11):475-83

20. Andrew Nii Anang and Chukwunweike JN. Leveraging Topological Data Analysis and AI for Advanced Manufacturing: Integrating Machine Learning and Automation for Predictive Maintenance and Process Optimization. (2024). Available: https://dx.doi.org/10.7753/IJCATR1309.100