

Streaming Analytics vs. Traditional BI in Financial Decision-Making: A Comparative Study

Pavan Kumar Mantha

pavanmantha777@gmail.com

Abstract:

Financial institutions operate under conditions of high transaction volumes, intense regulatory scrutiny, elevated risk exposure, and rapidly changing customer demands. Conventional Business Intelligence (BI) systems — built on batch-based extraction, transformation, and loading (ETL) processes, centralized data warehouses, and standardized dashboards — have been the workhorse of financial decision-making for decades. These systems excel in structured reporting, auditability, and historical analysis, making them essential for regulatory compliance, financial consolidation, and performance management. Nevertheless, the growing speed of digital transactions, online payment mechanisms, and remote interactions with customers has revealed what has limited the power of delayed insights on cases like fraud detection, dynamic credit risk determination, and personalized interactions. Streaming analytics have been developed as a supportive paradigm where data flows are processed continuously as new events take place and it becomes possible to have near real-time insights, alerts and automatic reactions. In comparison to the traditional BI that is about retrospective analysis, streaming analytics is immediate, governed by time, and decision-making. The implementation of streaming analytics in financial services has its potential albeit questioning in its foundational issues of governance, its explanation, cost, the readiness of organizations, as well as its acceptance by the regulatory community. The paper is a full-fledged comparative analysis of streaming architecture and conventional BI with respect to financial decision-making. It explores their architectural roots, latency aspects of their choice and appropriateness in various financial applications. A critical evaluation of the system defines the trade-offs between stability and speed, interpretability and automation, cost effectiveness and complexity of operation. The governance, compliance and explainability are particularly focused on, which are the dimensions that are inadequately explored in the current literature. The paper arrives at the conclusion that streaming analytics is not the replacement of traditional BI but rather its supplement. A financial institution has the most practical way out through a hybrid analytical architecture, which is based on decision criticality and sensitivity to time. This work presents a systematic decision-making model that would be used by the executive team and architects in choosing the most suitable analytical paradigm when making particular financial decisions.

Keywords: Streaming Analytics, Business Intelligence, Financial Decision-Making, Real-Time Analytics, Governance, Decision Latency, Regulatory Compliance.

1. Introduction

1.1 Historical Dominance of Business Intelligence

Financial services institutions, prior to the advent of the internet, have led the data-driven decision-making due to their requirement to handle vast amounts of transactional information which have been keenly controlled by regulatory authorities. [1-3] Since its early phases of consolidating general ledger and financial close, complicated computing of risk-weighted assets, capital adequacy ratios, and other such decision-supportive information, traditional Business Intelligence (BI) platforms have been at the forefront of converting raw operational information into pro-forma information. These systems allowed financial institutions to shed off the system-specific reporting to seek integrated enterprise-wide visibility to facilitate the consistent interpretation of financial performance and risk exposure. The classical BI pipeline, with its batch-based extraction, transformation and loading (ETL) operations serving as a source of data to centralized data warehouses and standardized reporting layers became the industry standard, in large part, because it was so robust and predictable. This architecture guaranteed that the figures that were reported were consistent, balanced, and repeatable by creating controlled data refresh cycles and properly defined transformation logic. These attributes were highly relevant in financial circumstances in which any slight act of inconsistency may result in regulatory discoveries and/or, financial reclaims, and the loss of stakeholder trust. The concentration of data in centralized data warehouses also provided the benefit of having deep historical data that can be used to make longitudinal analysis as well as period-over-period comparisons in financial planning and control. The conventional BI systems got tightly integrated in the overall financial decision-making, such as the evaluation of the quarterly and annual earnings, liquidity, stress testing, and capital adequacy reporting. These systems became so common that regulators and auditors grew accustomed to such determinacy, providing lineages of data, documented control and the capability of replaying reports at the time of a supervisory review of ultimate documentation by the audit team. With time such familiarity solidified the dominance of BI, which establishing a regulatory acceptance was as important as an analytical prowess. Despite being largely retrospective in nature, the fact that these platforms were aligned to the needs of the governance, compliance and strategic planning made them relevant owing to their long term relevance. Consequently, legacy BI has not only found a place as a technology selection, but has become an integral part of financial infrastructure of decision making.

1.2 Rising Demand for Faster Insights

In the past ten years, financial ecosystems are in the midst of a major change as these three factors have precipitated digitization, platformization, and growing interconnectedness of services. The large-scale implementation of digital money, mobile banking apps, and algorithmic trading services and online lending systems has significantly increased the speed of a financial process. The number of transactions that previously used to accrue over hours or days is now being generated on-demand, and at large scales. Because of this, the decision cycle time in financial services has shrunk to milliseconds of time, and the expectation cycle around how quickly insights can be generated and taken to action has been radically changed. This increase has been paralleled by a corresponding rise in real-time risk and opportunity. Fraud efforts are progressively focusing on the instant payment tracks and automated systems and must be identified and acted upon during transaction and not after settlement. The customers that communicate with each other via digital channels anticipate prompt responses, be it in approving a

transaction, credit, or a request concerning a service and any delay would have a direct impact on satisfaction and retention. Equally, greater market and information flow require constant frequency of exposure, position, and liquidity parameters surveillance to avoid quick compounding of losses. These forces have revealed structural constraints of batch-based analytics architectures. Classical BI systems, with their focus on the periodic updates of the data and the retrospective mode of analysis, are not able to facilitate the decisions that have to be made in the middle of the events progression. Reports on the same are hours away and can be valuable in reporting and examination of events after the action but can be too late in saving losses or seizing opportunities. As a result, banking institutions have started to review whether or not batch-only analytics is sufficient to support operational and customer-facing decision-making. This growing need to move faster does not only signify a desire to work faster, but a complete change in the mode of making and storing financial value as an operating environment grows more and more in real-time.

1.3 Business Drivers for Real-Time Decisioning




Figure 1: Business Drivers for Real-Time Decisioning

- **Fraud Detection**

One of the most powerful forces that have propelled adoption of real time decisioning in financial services is fraud detection. Current frauds take advantage of immediate payment systems and automatic flow of transactions which do not leave room to intervene in the post-event process. [4,5] Real-time analytics allows institutions the opportunity to evaluate transaction at the point at which it becomes executed, correlate the behavioral signal, historical trends, and contextual attributes in milliseconds. This aspect enables the blocking and challenging of suspicious transactions in real-time and provided it with huge savings in financial losses and downstream recovery expenses. Contrarily, with the delayed identification system using batch analytics, an action is usually undertaken when money has been transferred, or accounts are hacked.

- **Customer Engagement**

Interaction with customers is becoming more based on the capability to react to the behavior of the users in real-time. Real-time personalization enables the financial institutions to offer, message, and engage different representatives according to real-time indicators like browsing, transactions, or channel-related indicators. Quick response, including this offering of an interesting product proposal or a timely service

insurance almost always has a positive impact on the conversion rate and customer contentment. Though helpful in segmenting and strategy, batch-driven insights do not have the urgency needed to affect customer decisions as they are interacting with the company or brand, and thus real-time analytics plays a significant role in facilitating contemporary digital experiences.

- **Risk Mitigation**

Risk mitigation is now not stopped and assessed after some time but made as a continuous monitoring because financial exposures vary swiftly in dynamic markets. Live decision-making assists in continuous pricing and assessment of credit risk, market risk and liquidity positions using real-time transaction and market data. The continuous monitoring of the exposure also helps in early detection of breach of limits, concentration risks or abnormal patterns before transforming these into material losses. With the paradigm shift between the inactive snapshots and the active risk consciousness, the real-time analytics can help an institution to be more responsive and able to react on the threats that are rising instead of responding to them having been fulfilled and the negative outcomes have already become a reality.

2. Literature Survey

2.1 Traditional BI in Financial Services

The conventional Business Intelligence (BI) literature of the financial services sector has continually highlighted the innate basis of structured, dependable, and regulation-acceptable decision-making [6-8]. Financial reporting ecosystems have been dramatically influenced in the last three decades by early paradigm architectures, especially that of Bill Inmon, a leading proponent of the top-down data warehouse model over the enterprise, and the one of Ralph Kimball, a champion of the bottom-up dimensional modelling methodology. Such architectures focus on centralized data integration, data schema stability and historical data maintenance, so that they are especially appropriate to statutory reporting, financial consolidation, and risk disclosures. Empirical research on batch-oriented ETL pipelines also underscores three features: supporting end-of-period reconciliation, controlled data transformations, and reproducible reporting results, which are essential in regulatory frameworks, which include Basel accords, the IFRS, and the SOX. In addition, the conventional BI systems embrace robust metadata management, lineage tracing, and access controls, which enhance the governance and auditability. The literature has however also recognised that such systems are naturally retrospective and that they are geared towards explaining what has occurred instead of responding to what is occurring and hence limits its usefulness in a swift financial environment.

2.2 Emergence of Streaming Analytics

Streaming analytics have their origins in the development of event-based architectures, distributed systems and continuous query processing. The academic research on complex event processing (CEP) and stream processing models shows that unlimited streams of data can be analyzed on a sliding, tumbling and session basis using the windows and allow near-real-time observation of patterns. The literature in the financial services field is focused on high-velocity applications, e.g. money laundering detection, surveillance of a financial market, and algorithmic trading, where aspects of latency can substantially affect outcomes. Case studies across industries have consistently demonstrated the value of streaming analytics through measurable improvements in fraud detection rates, transaction processing

speeds, and response times to anomalous events. Nonetheless, much of this work focuses on measurements of the performance of systems, including latency, scalability, and fault tolerance, without paying much attention to governance issues. Other problems like the lineage of data, explainability of decisions made by a model, the replay of decisions made, and the defensibility of decisions made by regulations are often regarded as secondary concerns or delegated to downstream systems. Subsequently, literature explains streaming analytics as mainly a technology change, as opposed to a decision-governance concept in governed financial settings.

2.3 Identified Research Gaps

A critical review of literature indicates that there are a number of gaps in the comparative knowledge of traditional existing BI and streaming analytics. To start with, the majority of studies implicitly assume the disruptive nature of streaming analytics as an alternative to previously existing BI systems based on batches after all, instead of exploring the possibility of the two paradigms coexisting and being complementary to one another at various levels of decision making. Second, the emphases on explainability, auditability, and traceability are underexplained, which is the main feature of financial accountability that is difficult to provide in the continuous, event-driven systems. Third, most of the previous studies are disproportionately emphatic on technical performance measures, including processing latency and throughput, and disproportionately lack the analysis of how an outcome of analytics is utilized, managed, and rationalized in financial decision-making processes. This leads to the tendency to abstract away the decision context, namely, who makes decisions, on what boundaries, and with what responsibility. The paper fills such gaps by taking a decision-centric analytical standpoint, in a systematic assessment of traditional BI and streaming analytics not as technologies, but as two-polar financial decision enabling technologies that must have different temporal, regulatory, and governance demands.

3. Research methodology

This study employs a structured literature review methodology to comparatively analyze streaming analytics and traditional Business Intelligence (BI) systems within the context of financial decision-making. Relevant academic literature was identified through searches of Google Scholar, IEEE Xplore, and ACM Digital Library using keywords including "streaming analytics," "real-time analytics," "business intelligence," "financial decision-making," and "event-driven architecture." Sources were selected based on relevance to financial services, analytical architecture, governance, and regulatory compliance. The comparative dimensions evaluated in this study — including latency, decision type, actionability, cost predictability, and architectural complexity — were derived from recurring themes identified across the reviewed literature and validated against documented industry practice.

4. Traditional BI in Financial Services

4.1 Architectural Overview

Financial services Traditional Business Intelligence (BI) systems typically adopt the linear, batch-based architecture that is constructed with [9,10] stability, control, and auditability as the priorities rather than immediacy. Both types of architectural layers have a specific contribution to the consistency of the data and regulatory reliability.




Figure 2: Architectural Overview

- **Source Systems**

The basic operational implementation of classic BI structures is made up of source systems. These are core banking systems, enterprise resource planning (ERP) platform, customer relationship management system, trading and risk platform, and payment processing engine, in financial institutions. These are transaction-based systems, which are not optimized with regard to analytical workloads, but operational throughput. The form of data that is produced in this layer is also highly normalized, application specific, and is often shared across a variety of platforms, and therefore must be integrated downstream before it can be used in enterprise-wide analysis.

- **Batch ETL Processes**

Extract, transform and load (ETL) processes are batch based processes automated to form the backbone of integration between originating systems and analytical repositories. Data will be regularly pulled - usually via extracts on a daily, weekly or monthly basis, purged, standardized, and enhanced based on pre-established business rules. The transformations contain data validation, reconciliation, aggregation, and conformance to enterprise dimensions. Handling of the errors, start-restarting, and controlled execution windows are viable using the batch paradigm, which is required to satisfy financial close schedules and regulatory reporting schedules.

- **Centralized Data Warehouse**

The centralized or the data warehouse serves as the system of record upon which analytical and reporting cause bases are anchored. It contains integrated, time-varying and historical data on business domains which is stored using relational schemas including star or snowflake schemas. With this layer, data quality, consistency, and traceability are highlighted, and it helps with longitudinal analysis, period over period analysis, and regulatory audits. Lineage tracking and metadata management are part and parcel which enable organizations to prove how the figures they are reporting are based on underlying transactions.

- **Dashboards and Reports**

The consumption level of traditional BI structures consists of dashboards and reports. Curated datasets are available to business users, analysts and regulators as a part of standardized reports, scorecards, and visual dashboards. Such outputs are commonly updated according to the range of batch cycles and they are also meant to be periodically, not on-the-fly, operational data-driven. Consequently, this layer is best

suited to the strategic supervision, compliance reporting and management of review, and only responds to the fast changing events minimally.

4.2 Strengths

The conventional Business Intelligence (BI) designs have a number of strengths that precondition their especially strong applicability in the context of financial services where predictability, [11,12] control and accountability are of the utmost importance.

- **Stability**

The traditional BI systems have stability as their key strength factor that is fuelled by well-managed data refresh rates. Limiting variability and operation uncertainty Batch-oriented processing makes sure that data is ingested, transformed, and sanctioned at fixed time intervals. The predictability has enabled financial institutions to match data readiness to business schedules like end of day process, end of month close, and regulatory submission periods. BI platforms also help ensure consistency in the baseline of the analytics and minimize the possibility of the metrics to vary, which makes executive and regulatory decisions consistent and dependable.

- **Auditability**

Persistent, reconciled datasets in central data warehouses are heavily used to help provide a high degree of auditability. Past information is stored in structures that are not easily modified or do not change at all and through this storage; companies can produce reports as they were produced at any given time at any given moment. Embedded in ETL procedures, reconciliation checks provide assurance between reported figures and source systems, and metadata and lineage records are a record of transformation logic. These features allow the auditors and regulators to identify financial results in both directions to their beginning transactions, which is essential to compliance, resolution of dispute, and forensic analysis.

- **Regulatory Acceptance**

The long history of usage and full comprehension of the data lineage mechanisms offer traditional BI platforms ubiquitous regulatory acceptance. Supervisory bodies are used to batch form of reporting, data that have been defined in standard format as well as the control point that are documented. This clear distinction between the data ingestion, transformation, and consumption layers into data enable the transparency and validation levels at regulatory audits. Consequently, reports generated in BI tend to be viewed as official accounts of the financial status of an institution and therefore, classic BI is one of the foundations of statutory reporting and risk governance models.

4.3 Limitations

Regardless of their stability and governance advantages, traditional Business Intelligence (BI) architectural solutions have a series of inherent limitations that restrict their performance in contemporary, rapidly evolving financial contexts. Among the most significant disadvantages is high latency: data is typically made available only after several hours or even days, following batch-based extraction, transformation, and loading processes. While this is acceptable for statutory reporting and historical analysis, the latency introduces a time gap between actual events and their analytical interpretation. Delays in visibility can result in missed opportunities in areas such as payment processing, market risk monitoring, and liquidity management, reducing responsiveness to emerging risks. Closely related to latency is the reactive nature of traditional BI-generated insights. Since only

snapshots of historical data are analyzed, BI platforms are more descriptive than predictive or proactive — informing decision-makers of lagging indicators, trend summaries, and post-hoc variance analyses rather than enabling real-time intervention. Rather than preventing adverse events, corrective actions are often taken only after losses from fraud, limit violations, or service failures have already materialized, which decreases the overall efficiency of decision-making. A further weakness is the limited support for operational decision automation. Traditional BI deliverables are generally human-facing — delivered via dashboards and reports — rather than integrated with transactional or control systems. The separation of analytical and operational systems makes it difficult to trigger automated responses to analytical findings. As a result, despite the detection of risks or anomalies, remediation is commonly still performed manually. This human-in-the-loop dependency introduces delays, limits scalability, and creates inconsistency, particularly in high-volume environments. Collectively, these constraints underscore the growing misalignment between traditional BI architectures and the real-time, automated decision-making needs of modern financial services.

5. Streaming Analytics: Concepts and Capabilities

5.1 Event-Driven Processing

Event-driven processing is the conceptual basis of streaming analytics because it treats the incoming data point as an independent event that is processed immediately it is created. [13,14] Streaming architectures handle incoming events in real-time, whereas batch-based systems process historical data in batch, and thus do not store repeatedly encountered state or timely information, unlike their counterparts. The system can preserve aggregates, counts, and correlations sliding, tumbling, or session-based windows with the help of state management. Consequently, decisions are not determined only at specific points of reporting but they are analyzed at all times. This paradigm is very much aligned with the working realities of the contemporary financial systems, where transactions and market updates together with and in regards to customers are incessant in occurring and they need to be analytically interpreted immediately.

5.2 Continuous Computation and Alerting

Streaming analytics substitutes non-continuous, periodic computation in favor of continuous, incremental computation of metrics and indicators. Analytical results are constantly updated in real time as new events come without necessarily having to re-compute complete datasets. This is a continuously computed model that facilitates low-latency anomaly, threshold excess detection and complex event pattern detection. Rather than using dashboards which are periodically updated, streaming systems include alerting and action capabilities as part of analytics. In case of predefined conditions, the alerts may be issued to the stakeholders or automated activities may be taken in downstream systems. This close connection between analytics and action allows serving a response more quickly and involves taking a proactive role in the areas of financial situations that are at risk or time sensitive.

5.3 Key Capabilities

The most important feature of streaming analytics is that it provides close to real-time insights and, as a result, organizations can monitor and analyze what happens within several seconds or milliseconds of their happening. The immediacy facilitates timely identification of the such risks and opportunities

before the other processing batches are done. There is a second major capability, that is, automated response when the result of an analysis is an automatic state reaction to the system, e.g., transaction blocking, limit control, or escalation of a workflow, which causes less interaction with a human operator. Lastly, streaming analytics is superior in patterns of time, sequences, trends, and correlations that are dynamic, or developing over time instead of showing a snapshot. Streaming systems enable the development of an action-oriented and dynamic analytical layer to supplement traditional BI by analyzing event streams in entirety within the current financial decision-making landscape.

6. Decision Latency and Business Impact

6.1 Comparative Dimensions

Dimension	Traditional BI	Streaming Analytics
Latency	High	Low
Decision Type	Strategic/Tactical	Operational
Actionability	Low	High
Cost Predictability	High	Low
Architectural Complexity	Moderate	High

Table 1: Comparative Dimensions




Figure 3: Graph representing Comparative Dimensions

Note: Values shown are illustrative relative weightings derived from literature synthesis and are not empirically measured scores

- **Latency**

Latency is the elapsed time between data production and its availability as actionable insight. This dimension results in the lower score of traditional BI systems due to their processing that is based on the batch which implies the existence of delays between production and analysis availability of the data. [15,16] Such delays are tolerable when used in retrospective reporting and frequency reviews but constrain responsiveness in time-sensitive situations. By comparison, streaming analytics supports the

high level of alignment with low-latency needs through processing events in real time, delivering insights in seconds or milliseconds. This is useful especially in the context of operational monitoring, fraud prevention and real-time risk reduction.

- **Decision Type**

The decision type dimension reflects the primary category of decisions supported by each analytics paradigm. Traditional BI is strongly aligned with strategic and tactical decisions, where aggregated historical data and trend analysis support long-term planning, performance management, and regulatory reporting. Streaming analytics, by contrast, excels at operational decision-making, enabling immediate responses to events as they unfold. While streaming analytics is less suited to enterprise-wide strategic assessments that require deep historical context, its strength lies in execution-level, time-sensitive decisions — precisely the area where traditional BI is weakest.

- **Actionability**

Actionability is a measure of how the analysis outcomes directly result in a decision or intervention. Conventional BI systems are mainly person-centered, where an analyst and a manager are required in order to read reports and take action, which restricts expediency and scale. Streaming analytics, in its turn, is intended to facilitate automated or semi-automated reactions. The results of the analysis can be used to initiate alert, workflow escalation or system level controls and thus the interventions can be much faster and consistent. It is this increased degree of automation that justifies the greater conformity of streaming analytics to action-oriented applications.

- **Cost Predictability**

Cost predictability indicates the continuity and visibility of the operational costs of analytics platform. Conventional BI setups usually use infrastructure that is fixed and scheduled workloads that lead to fairly predictable costs that are in line with budgeting and financial planning periods. Streaming analytics platforms however, are mostly based on elastic compute scales and real time processing which results in unpredicted cost profiles due to event volume and processing intensity. This elasticity provides scalability but minimizes cost predictability especially in high throughput financial cases.

- **Architectural Complexity**

The design, integration and operational sophistication needed to deploy and support analytics systems is captured by the complexity of architecture. The complexity of traditional BI architecture is middle-range and the practice is well established and standardized, which leads to moderate complexity and mature operation practice. The streaming analytics systems include distributed event ingestion, stateful processing, fault tolerance, and low-latency guarantees, which add complexity to the design and operation aspect. The increased complexity of streaming architectures needs specialized skills and governance infrastructure, which can augment the barriers to adoption over their performance benefits.

7. Use Case Comparison in Financial Services

7.1 BI-Dominant Use Cases




Figure 4: BI-Dominant Use Cases

- **Regulatory Reporting**

One of the best BI-dominant use cases would be regulatory reporting because it has very high standards of accuracy, consistency, and auditability [17,18]. The audits and supervisory reviews require that financial regulators implement well-identified data lineage, reconciliation with the source systems and repeatability of report generation. The conventional BI systems with their batch processing models and central data warehouses are very appropriate in the creation of standard disclosures like capital sufficiency, liquidity measures and regular financial statements. BI is the ideal paradigm to use when compliance-driven reporting is required due to its focus on historical correctness rather than instantaneous nature.

- **Financial Performance Analysis**

The analysis of financial performance depends more on the aggregated historical data to compare the profitability, cost efficiency, and the business unit performance of a given time frame. Classical BI systems also allow time geographic product- and customer group-multiplexed multidimensional analysis and facilitate budgeting, forecasting, and management review. The latency caused by the processing in batches does not have a significant impact on the quality of the decisions since they are usually analyzed on a monthly or quarterly basis. The metrics and definitions also give consistency in the BI tools where the performance levels are comparable over different reporting periods.

- **Long-Term Trend Analysis**

Trend long term analysis puts emphasis on structural trends and directional changes over a long period of time. The same can be well done on the traditional BI by keeping repositories of deep historical data and consistent dimensional frameworks which facilitate longitudinal examination. Analysts do not have to worry about temporary fluctuations because they are able to view or analyze the multi-year revenue growth, risk exposure, customer behavior, or market performance. This can create informed strategies and policies, which makes BI relevant to the information that can be used to make long-term financial decisions.

7.2 Streaming-Dominant Use Cases

- **Fraud Detection**

The best example of streaming-dominant use case is fraud detection in which the value of analytics depends heavily on the speed of response. The streaming analytics systems drill down to transactions and with the help of short time windows, correlate the events to detect any fraud like patterns. A real-time detection will allow organizations to stop potential transactions or proceed with investigations on cases, before financial depletion takes place. Streaming analytics is much more efficient than BI based on flat rates to fight real-time fraud simply because of the capacity to process high-velocity streams of events.

- **Real-Time Credit Approvals**

The real-time credit approval processes demand real-time evaluation of risk when interacting with the customers. The streaming analytics supports this need by incorporating real-time data on transactions, behavioral clues, and risk models to assess creditworthiness in a few seconds. Rapid approval and speed of customer journey can be automated and executed, with the risk controls in place. Such real-time decisions cannot be made with the help of traditional BI, with its slow insights.

- **Best-Time-to-Contact Optimization**

Best-time-to-contact optimization is trying to identify the most effective time to approach customers depending on their past actions and a situational indicator. The analytics of streaming continuously analyzes the interaction between the customer, street channel activity, and patterns of response to determine the best engagement periods. The personalization in real time allows the financial institutions to maximize the numbers of conversions and customer satisfaction. It was the dynamic aspect of this use case where continual analysis and expediting is required which makes streaming analytics the default paradigm over the unchanging, past-focused BI methods.

8. Governance, Compliance, and Explainability Considerations

8.1 Audit Trails

Financial analytics presuppose auditability as a basic requirement, and the traditional BI systems can easily deliver robust audit trails as a result of deterministic and batch-based processing. [19,20] The reports created out of centralized data warehouses are completely replayable since the data sets are persistent, reconciled and versioned by reporting period. This enables regulators and auditors to easily recreate the results and trace the figures that have been reported to the origin transactions with a lot of confidence. However, by contrast, streaming analytics can process irrevocable streams of events and have continuously changing state, an attribute that does not necessarily give those systems replayability. Streaming architectures should strategically adopt event retention, immutable logs and state versioning in order to reach similar auditability. The absence of such design considerations makes it hard to recreate historic decision conditions to undermine the defensibility of a regulation.

8.2 Explainability Challenges

Explainability becomes a major issue to real-time analytics especially when the actions are informed by automated rules or machine learning tasks. In the traditional BI, the findings are made using a set of unchanging data as well as documented modifications, and it is simpler to justify outcomes when conducting an audit or regulatory review. Streaming analytics, however, analyzes decisions in dynamic

settings whereby inputs, state, and thresholds can change at any point. To recreate the rationale behind a certain step at a certain point, it not only is essential to record the event that made this action happen, but also the active versions of the rules, the models and the data that existed in the situation at that point. This lack of contextual snapshot makes post hoc difficult to explain, and makes the decision-making more regulatory, particularly in customer-impacting decisions like transaction blocking or credit denial.

8.3 False Positives and Risk

False positives can really happen with streaming analytics systems since it is very possible that the decision is taken without much contextual information at the time of the assessment. Incomplete historical visibility and short time windows may incorrectly classify benign irregularities as risks; hence, resulting in unwarranted alerts or automated interventions. The consequences of such false positives in the financial services can work against customer experience, operational workload, and trust towards automated systems. Therefore, threshold design, adaptive tuning, and multi-stage decision logic constitute the necessary aspects to achieve a balance between the responsiveness and the accuracy. Unlike the prior checks, more escalation mechanisms and feedback loops can be introduced to counter the risk that would allow high impact decisions to be reviewed and corroborated before enforcement.

9. Cost, Complexity, and Organizational Readiness

9.1 Cost Factors



Figure 5: Cost Factors

- **Infrastructure Elasticity**

Streaming analytics systems are usually based on the elastic scalable architecture that is able to support the volumes of events and maximum processing loads. Although this elasticity can perform highly and with elasticity, it leads to variation in costs, since the usage of compute and storage is affected by data velocity and data retention needs. The streaming systems are not as predictable as traditional BI systems in terms of the expected cost spikes during higher transactional periods. The necessary steps to avoid uncontrolled expenditure are therefore the effective capacity planning and workload throttling, and cost governance mechanisms.

- **Operational Monitoring**

The dynamics of the streaming analytics requires advanced operational monitoring to develop reliability and performances. Ingestion lag, processing latency, state growth, and failure recovery are some of the system health indicators that should be monitored by organizations. Such overhead of monitoring costs

operations and demands investment in observability at the tool level and specialized support teams. Unlike conventional BI, when the failure may be ameliorated, under expected timeframes, a streaming system needs uninterrupted patrol in order to prevent possible breakdown of service or erroneous decision.

- **Skilled Personnel Availability**

Development and sustenance of streaming analytics overcomes need technical skills within fields of distributed system, event processing and real time data structure. This resourcefulness of such competent employees directly affects cost and viability of projects. The lack of talent may result in higher costs of hiring and adoptions, as well as inadequate expertise might result in weak implementations. Because of this, workforce preparedness and investment on training should be the parts that organizations should consider as part of the streaming analytics costs frameworks.

9.2 Organizational Maturity

The success of streaming analytics is closely related to the possibility of an organization to take prompt actions in response to analytical insights. Real-time detection itself is not something that creates value unless operational processes, decision-making processes, and governance structures are configured towards responding quickly. Companies that have a slow approval process or ownership confusion might not translate the insights into actions in a timely manner and therefore, the alarm bells might keep piling up. Real-time analytics in this situation can inundate the users with noise and not help them make better decisions. Thus, the attainment of the complete advantages of streaming analytics requires the presence of organizational maturity, in terms of accountability clarity, automated processes, and decision authority at the correct level.

10. A Decision Framework for Choosing Between BI and Streaming

10.1 Decision Matrix

The decision matrix offers a systematic method of highlighting which of the traditional BI or streaming analytics is more applicable to a particular financial use case. One of the main requirements is time sensitivity: low-urgency decisions (e.g., regulatory disclosures or regular performance reviews) are well supported with BI systems, but high-urgency ones (e.g., fraud prevention or transaction risk scoring) are better supported with streaming analytics. The two paradigms also vary under the impact of regulation. Use cases that receive high regulatory scrutiny often favor BI because of its authority nature of processing, fixable datasets and the already established auditability. Streaming analytics is more suitable in moderate regulatory impact and when quick action is important and when one does not need a fully reconciled past perspective. Lastly, the event volume is important; BI is useful in medium data volumes, which are batch-processed, but streaming architecture is configured to support extremely high rates of events (and has continuous event ingestion and processing). Collectively, the criteria assist companies in ensuring that the decisions made by analytics are in line with decision needs and not the adoption of technologies separately.

10.2 Hybrid Architecture Recommendation

Instead of considering BI and streaming analytics to be mutually exclusive, a hybrid architectural design can be used as a more balanced and resilient system for financial services. This layered approach mirrors

the Lambda Architecture proposed by Marz and Warren (2015), which similarly separates high-speed real-time processing from batch-oriented historical analysis, unifying both layers into a coherent analytical system. Streaming analytics occupies the topmost position in this stratified model, performing the role of front-line to identify the high-velocity event, make decisions in real time, and automatically respond to them. Conventional BI is downstream in that it combines streaming results with more extensive enterprise information to offer oversight, reconciliation, and historical examination. This division of concerns enables organizations to achieve both responsiveness and governance simultaneously. Streaming systems provide speed and automation, whereas the BI platforms can provide explainability, regulatory adoption, and accountability in the long run. An effective combination of both the paradigms in a coherent system of making decisions can allow a financial institution to sustain different levels of decisions and maximize the performance at operational, tactical, and strategic levels.

11. Challenges and Future Outlook

11.1 Overuse of Streaming

Among the new pitfalls in modern analytics usage, there was a tendency of excesses in the application of streaming technologies in situations when real-time processing is not going to significantly enhance the quality of a decision. Not every financial decision is time-constrained, and streaming analytics are not a solution that should be applied blindly, as it can impose unwarranted architectural complexity and cost. Not only continuous processing is expensive in elastic infrastructure resources, but it must be monitored continuously, and specialized expertise is required, which may be offset by low-urgency or retrospective use cases. In scenarios where real-time insights cannot be acted on directly, streaming outputs have a risk of being treated as redundant not value-generating inputs. This means that, organizations have to be very disciplined in matching streaming adoption to clearly articulated requirement decisions to ensure that streaming implementations are not carried out in a technology-based and not a value-based fashion.

11.2 Convergence of Analytics Platforms

The convergence of batch and streaming analytics into hybrid platforms which enable the various modes of processing by a shared governance structure is a prominent trend in the industry. Architectures in which historical and real-time data coexist are being increasingly provided by vendors and open-source ecosystems, to facilitate ease when moving to, and between, batch analytics, micro-batching and actual streaming. This intersection minimizes data siloing, eases architecture, and enables organizations to use harmonized security, metadata administration, and lineage over analytics loads. The result of unifying processing paradigms is that financial institutions can eliminate redundant logic and infrastructure as well as enable more diverse decision horizons expressed using a single analytical ecosystem.

11.3 Regulatory Scrutiny

Regulatory scrutiny is likely to increase as financial institutions add more and more automation in real-time and other algorithms-driven decision-making processes. Any actions performed by machines which directly involve consumers (e.g., blocking of transactions, decision-making, placing restrictions on accounts) are concerned with the issue of fairness, transparency, and responsibility. Regulators will tend to require more explainability, written control procedures and being able to recreate decision situations in retrospect. Such a change will necessitate organizations to weave compliance concerns into straight

into streaming systems, such as versioned regulations, model provenance, and audit trails of decisions. The future success of streaming analytics will then rely on more than the technical performance but also the possibility to prove responsible and controlled automatization.

12. Conclusion

Streaming analytics is another crucial transformation of the analytical capacity offered by financial organizations that allows processing the high-velocity data streams and making automated and immediate decisions in time-sensitive situations. With more and more financial ecosystems growing digital and connected, the pressure to gain quick access to the transactions, customer activity, and market trends keeps growing. Streaming analytics is a direct response to this requirement as it allows evaluating what occurs all the time and responding to it in real-time, increasing the sensitivity of such areas as fraud detection, providing credit, and communicating with customers. Its strength is operational agility- it helps institutions to do what they can at the time when the decisions can get the most effect. However, this work still proves that classical Business Intelligence (BI) is a vital component of financial analytics. BI platforms offer the governance, the stability and the auditability needed in regulatory reporting, financial consolidation and long-term strategic analysis. Their batch-based designs provide reconciled data, predictable results, and documented lineage of data, which are essential to the preservation of regulatory confidence and organizational responsibility. Traditional BI is still the best paradigm to use when decisions should be made based on accuracy, explainability and consistency across history rather than on the present. One of the main contributions that this work has made is its ability to demonstrate that streaming analytics, in fact, should not be considered an alternative to BI, but instead, it is an additional ability that operates in an extended decision-horizon. The analysis of analytics paradigms in terms of a decision-centric perspective, the research indicates that the applicability of the BI or a streaming analytics is deeply rooted in the nature of the decision context in question, such as time sensitivity, regulatory influence, and actionability. Any efforts to consolidate all decisions into some real-time model endanger to add cost and complexity without adding value in similar amounts and sole BI dependency prevents responsiveness in high-rates change scenarios. The conclusion highly suggests implementation of hybrid approach to analytics whereby streaming analytics will deal with immediate detection and automated response, whereas BI will be used to provide oversight, reconciliation and strategic insight. A layered approach to this kind of service allows financial institutions to trade off speed against control, innovation against compliance, and automation against trust. With convergence in analytics platforms and heightened regulatory expectations regarding explainability, organizations that adopt both paradigms deliberately, as opposed to following one at the cost of the other, will be well-placed to overcome the challenges of an ever-more real-time of financial existence.

REFERENCES:

- [1] Kimball, R., & Ross, M. (2013). The data warehouse toolkit: The definitive guide to dimensional modeling. John Wiley & Sons.
- [2] Inmon, W. H., Strauss, D., & Neushloss, G. (2010). DW 2.0: The architecture for the next generation of data warehousing. Elsevier.
- [3] Etzion, O., & Niblett, P. (2010). Event processing in action. Manning Publications Co..

- [4] Akidau, T., Bradshaw, R., Chambers, C., Chernyak, S., Fernández-Moctezuma, R. J., Lax, R., ... & Whittle, S. (2015). The dataflow model: a practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing. *Proceedings of the VLDB Endowment*, 8(12), 1792-1803.
- [5] Chandrasekaran, S., Cooper, O., Deshpande, A., Franklin, M. J., Hellerstein, J. M., Hong, W., ... & Shah, M. A. (2003, June). TelegraphCQ: continuous dataflow processing. In *Proceedings of the 2003 ACM SIGMOD international conference on Management of data* (pp. 668-668).
- [6] Moss, L. T., & Atre, S. (2003). *Business Intelligence Roadmap: The Complete Project Lifecycle for Decision-Support Applications*. Addison-Wesley Professional.
- [7] Supervision, B. (2011). *Basel committee on banking supervision. Principles for Sound Liquidity Risk Management and Supervision* (September 2008).
- [8] Inmon, W. H. (2005). *Building the data warehouse*. John Wiley & sons.
- [9] Alfian, G., Ijaz, M. F., Syafrudin, M., Syaekhoni, M. A., Fitriyani, N. L., & Rhee, J. (2019). Customer behavior analysis using real-time data processing: A case study of digital signage-based online stores. *Asia Pacific Journal of Marketing and Logistics*, 31(1), 265-290.
- [10] Allen, S. (2012). *Financial risk management: A practitioner's guide to managing market and credit risk*. John Wiley & Sons.
- [11] Carbone, P., Fragkoulis, M., Kalavri, V., & Katsifodimos, A. (2020, June). Beyond analytics: The evolution of stream processing systems. In *Proceedings of the 2020 ACM SIGMOD international conference on Management of data* (pp. 2651-2658).
- [12] Dayarathna, M., & Perera, S. (2018). Recent advancements in event processing. *ACM Computing Surveys (CSUR)*, 51(2), 1-36.
- [13] Kolajo, T., Daramola, O., & Adebisi, A. (2019). Big data stream analysis: a systematic literature review. *Journal of Big Data*, 6(1), 1-30.
- [14] Chen, H., Chiang, R. H., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS quarterly*, 1165-1188.
- [15] Ahmad, A. (2015). Business intelligence for sustainable competitive advantage. In *Sustaining competitive advantage via business intelligence, knowledge management, and system dynamics* (pp. 3-220). Emerald Group Publishing Limited.
- [16] Al-Aqrabi, H., Liu, L., Hill, R., Ding, Z., & Antonopoulos, N. (2013, March). Business intelligence security on the clouds: Challenges, solutions and future directions. In *2013 IEEE Seventh International Symposium on Service-Oriented System Engineering* (pp. 137-144). IEEE.
- [17] Stonebraker, M., Çetintemel, U., & Zdonik, S. (2005). The 8 requirements of real-time stream processing. *ACM SIGMOD Record*, 34(4), 42-47.
- [18] Basel Committee on Banking Supervision. (2013). *Principles for effective risk data aggregation and risk reporting (BCBS 239)*. Bank for International Settlements.
- [19] Friedman, E., & Tzoumas, K. (2016). *Introduction to Apache Flink: stream processing for real time and beyond*. " O'Reilly Media, Inc."
- [20] Holzinger, A. (2018, August). From machine learning to explainable AI. In *2018 world symposium on digital intelligence for systems and machines (DISA)* (pp. 55-66). IEEE.
- [21] Marz, N., & Warren, J. (2015). *Big Data: Principles and best practices of scalable real-time data systems*. Manning Publications.