Journal of Information Systems Engineering and Management
2021, 6(2)

e-ISSN: 2468-4376

https://jisem-journal.com/

Research Article

Designing Scalable Data Architectures That Support Continuous Regulatory Oversight

Naresh Bandaru

Senior Software Engineer

ARTICLE INFO

Received: 04 May 2021

Revised: 10 Jun 2021

Accepted: 20 Jun 2021

ABSTRACT

The paper discusses the application of an extension of a simple lakehouse architecture, one that employs a compliance fabric to provide a better reliability, the lineage of data, as well as the enforcement of policy usage of workloads of corporate taxes. The comparison between the original architecture and the enhanced compliance based system is made through conducting controlled tests on ingestion, transformation, lineage completeness, and audit readiness. The finding shows that automatic policy checks and lineage capture do not reduce the performance because they minimize the metadata gaps and improve consistency. The study also indicates higher traceability on the processing of various steps and lesser error in compliance in reporting. The results reveal that the compliance layers introduced to a lakehouse increase the levels of trust, quality, and regulatory compliance.

Keywords: Compliance Fabric, Audit Readiness, Ingestion Throughput, Lakehouse Architecture, Corporate Tax Workloads.

I. INTRODUCTION

The tax preparation requires the quality of data dependability, excellent origin, and unbreakable compliance of this data since the error would lead to delays in reports and the penalty. Lakehouse systems espouse scalable processing of data, but might not offer complete monitoring of the lineage and policy of the many transformation processes. This paper is about the process of making the compliance fabric to a lakehouse in order to fill these inadequacies. The paper reviews its related literature and offers protruded architecture, and compliance behavior with the applications on real workloads of corporate tax. They are geared towards addressing the question of whether compliance-related issues can improve governance without hurting performance and show whether the process can be utilized by enterprises.

II. RELATED WORKS

Continuous Regulatory Monitoring and Compliance Alignment

The ongoing regulatory surveillance Studies have shown that the long-term compliance will not be based on the interpretation of the policies only, but scalable, structured data architecture, which will offer the possibility to continuously measure the regulatory metrics.

Ongoing assessment of the solvency rates in the framework of the insurance regulation Solvency II stipulates the notion of the Own risk and solvency assessment (ORSA) that is periodically reviewed on the level of the solvency of the insurance companies under the changing conditions within the market

[1]. This is very highly computational as the entire procedure of the computation of the solvency ratios involves a repetition of time-dependent projections of liabilities and capital requirement.

With increasing data and model complexity, [1] indicates that the operation cost of complete, and repeated evaluation is high. To address this, the authors propose approximation models such as Curve Fitting and Least Squares Monte Carlo that allow solvency to be approximated on near real-time without necessarily having to execute full actuarial models.

These are parametric proxies which develop an initial instance of scale-related patterns of analysis which are regulatory accurate and the processing cost of such proxies decreases with scale, giving evidence that model abstraction helps firms to stay under regulatory compliance as part of continued regulation cycles.

Basically, the same pressures are portrayed in the complementary research in the auditing field. Continuous auditing (CA) models are based on the availability of streams of information to determine regulatory exposure and exception states which exist among the distributed systems [3].

CA requires the unstructured and structured datasets access in real-time as opposed to the traditional version of periodical auditing. According to the authors of [3], analysis is not the sole problem regarding architecture and the data retrieval, data integration and data organization should also be feasible as the system complexity is enhanced.

Their proposed system of integration relates standards of data extraction such as the Audit Data Standard to hierarchical reporting systems such as the XBRL. The design allows the auditors to obtain predefined data elements in a standard manner. This article therefore demonstrates that scalable compliance data architectures rest upon standardization as well as homogeneous extraction as read as much upon processing capacity.

The other branch of literature represents a mediator between continuous auditing and Big Data properties. As the analysis in [4] shows, there are some new points of compliance failures in large-scale financial platforms, which are related to the consistency of data, confidentiality, and aggregation.

In both instances of gaps that the authors associate the classical audit practices and scalable techniques, which are founded on the Big Data systems. They say that auditors will experience an increasing structural pressure when the data are scattered across the operational systems and data lineages and reconstructions will now be more difficult.

Continuous audit processes therefore demand architectures, offering traceability, co-ordination of policy and metadata model. The papers disclose that scalable monitoring is based on not only the fast analytics, but also architecture selection should not jeopardize observability, lineage and standardization of properties required to remain continually compliant in growing platforms.

Security, Privacy, and Access Control as Scale-Responsive Architectural Foundations

The data security, access control and privacy should become the first-class architectural requirements of continuous compliance data architectures, especially, with the increase in data volume and heterogeneity. As the authors of [2] mention, the traditional access control models are not sufficient when it comes to the context of Big Data due to the velocity, variety and distributed nature of the current systems.

Unlike the traditional database systems which are structured on the relational design and centralized implementation, the Big Data environments accept the presence of multiple data designs, distributed storage layer and different parallel computing processors.

The common access policies are hard to scale due to lack of reference model. In the research, the hypothesis is that the design of access control in big data platforms must be made to accommodate a collection of requirements that it must satisfy to provide compliance with an expansion with addition of data and query loads. This writing gives warning that implementing regulatory enforcement cannot be separated with creating the systems as security lapses are normally established in the scenario of a rapid growth.

The NIST Big Data security and privacy platform is also founded upon this viewpoint that makes security and privacy elements more of a cloth that is woven through the data system [7]. The book notes that security and privacy cannot exist as an independent module as the systems expand within the cloud platform and data centres and the internet of things.

They need to be incorporated with each architectural layer, the ingestion layer, the processing layer and analytics layer. The report provides definitions of both taxonomies and relationship of architecture which helps the system designers to understand the effect of distributed data flows on the aspect of confidentiality and use of the data.

It is important to note that security and privacy use cases mapping to reference architectures creates a reusable base of guaranteeing the consistency of regulations regardless of the varying data architecture. All these studies point to the fact that compliance must be persistent and it must be scaled and constantly enforced with security and not periodic (where the data architectures are elastic and distributed).

This is a trend that is backed by the role of governance in the distributed environments. The article in [5] states that there are cross-organizational big data ecosystems which are the areas of vulnerability where the regulation requirements are difficult to trace. This lack of centralized structures causes a problem in imposing the policy in the case where the data goes across the jurisdictions or corporate boundaries.

In response, the authors introduce the models of governance that include the inference of the policy and surveillance in the architecture. This introduces the modification in the compliance as a downstream reporting position to compliance as a design-time quality of architecture, platforms can remain in the track of the regulation even with the obstacles of continuous scaling utilized.

Scalable Ingestion and Real-Time Processing to Enable Continuous Oversight

The capabilities in such an architecture as the necessity to have regulatory control at all times include ingestion and low-latency processing, as the regulatory controls are to monitor the business activity in near real-time. Various researches in the fields of finance and IoT demonstrates that loss of compliance warning can be done at the initial level of architecture ingestion pipes.

The architecture of the cloud-based Home Energy Management Systems [6] offers a framework of the system elements into three architecture environments Ingestion, Operational and Analytical where the stakeholder requirements of each are evident.

Although the research does not concern the financial or regulatory information, but the energy information, the significance of the work is more related to the architecture sphere: the fact that the layers of ingestion, processing and analytics are separated does not mean that the scale-out would lose any context and traceability.

The performance guarantees are needed in the ingestion environments that will operate at the scale to the IoT devices to ensure that there is the real time behaviour. Such features can be compared to compliance architecture, whereby workloads to be executed in large scale transactions will be observed as the architecture expands.

The integration of the real-time financial data is explicitly taken into account in [8], and because the differences between the latency of an operational system and regulatory system are admitted to be the direct result of the architectural constraints of ETL pipelines.

The authors suggest to apply to the hybrid method of data integration that is grounded on the resilient distributed databases and pipeline streams to assist in continuing to accumulate financial signals. The architecture countermeasures the historical compliance bottlenecks by aligning the ingestion and semantics and metadata to trigger a faster reporting process and compliance laggardness.

It is presented in the work that ingestion pipelines can be reorganized to fit the demands of the regulations but not can be synchronized manually to facilitate help given to the continuous supervision. The high-performance compute layers should also be in a position to accommodate scale in workloads in analytics.

On the one hand, within the framework of high-volume financial analytics, the application available in [9] validates design options to provide the calculation engines with almost linear scalability in Spark. The authors demonstrate that the computation engine possesses major architectural decisions, which impact scalability through user-defined function and SQL-based rewrite comparison of the kernel.

The system of modeling financial contracts and risk calculations should provide accuracy and reproducibility when these computation kernels are used and those are properties that are expected to be controlled. Through the performance analysis, we find scalable compute structure can sustain the sustained analytics loads without influencing consistency and therefore could provide checks of compliance in real time within an environment where risk measures are dynamically changing.

Architectural Patterns for Sustainable Regulatory Alignment

The literature has several architectural patterns that may be perceived to be important in compliance at scale. There exist approximation model and proxies as demonstrated in [1] that reduce the computational load at the expense of compliance-relevant accuracy. The proxies are a scaled-up version of full models that allow continuous evaluation that would be otherwise cost-prohibitive.

Compliance indicators do not decrease with the growth of the data volume as standardized data extraction and representation suggested in [3] and [4] would ensure. Third, architectural layer, which is policy-conscious and secure fabrics, which are defined in [2], [5] and [7], make identity, access control and lineage in the distributed systems and prevent policy drift as systems grow.

Monitoring in real-time is enabled by consuming and processing design solutions that deal with real time streaming, distributed computation, as was mentioned in [6], [8] and [9]. These trends denote that scalable compliance is not an implementation but an attribute, which occurs as a result of architectural decisions at ingestion, storage, processing and semantic model levels.

III. METHODOLOGY

The present research is based on the quantitative research methodology to investigate the ways scalable data architecture can help sustain the regulatory control in large and expanding data environments. The methodology aims at quantifying the impacts of the chosen architectural patterns on system behavior in terms of regulatory monitoring, such as data availability, processing time, continuity of policy enforcement, and throughput of analytic analysis. The methodology also tries to measure the impact of various scaling strategies on the capabilities of systems to uphold compliance signals when both the volume of data and the intensity of workload grow.

Research Design

The study follows an experimental research approach where 3 architectural elements are assessed through controlled evaluation, which are typically present in large data platforms:

(1) data ingestion pipelines,

(2) real-time processing engines, and

(3) analytical computation layers.

Individual components are stressed to different levels of workload in order to monitor the variation of performance and compliance related indicators with scaling of systems. The design allows the measurement of relationships between architectural decisions and allows supporting the continuous oversight. An architecture is defined to represent the conventional data integration and batch-oriented processes.

This is then followed by the definition of three experimental architectures to include scalable ingestion, streaming-based processing, and distributed analytical computation. This arrangement makes it possible to place a direct comparison between quantitative measures between the traditional and scalable structures.

Data Sources and Datasets

The artificial financial transaction data is made to simulate the increasing business and regulations reporting requirements. The records in the data sets are organized and contain identifiers of transactions, time, customer selection, product category, and a risk category. The data is created in five levels of workload which are: 1 million, 10 million, 50 million, 100 million and 250 million records. These levels allow viewing the scaling of architectures between the very large operation environment and the medium environment of operation. In each dataset, there is some extra controlled metadata to allow compliance tracking features such as lineage identifiers, policy tags and verification flags.

Quantitative Metrics

The quantitative analysis dwells on four measures categories that directly affect the continuous regulation:

1. **Ingestion Throughput (records/second):** tests the speed of an addition of data into the system at different volumes.

2. **Processing Latency (milliseconds):** measures the time to create compliance analytics using data.

3. **Analytical Scalability (queries/second):** measures the count of risk or compliance queries that can be performed without significant delay.

4. **Compliance Signal Continuity (% retained metadata):** measures how much associated metadata that is maintained as data passes through the stages is of compliance.

These measurements are measurable aspects of compliance sustainability with lower latency, higher throughput, and metadata retention being consistently maintained as a marker of greater support towards continuous oversight.

Experimental Setup

Experiments are all conducted on replicated environment through distributed processing clusters whose resource levels are controlled. Each workload level is repeated three times to each architectural configuration in order to minimise random variability.

Analysis is done in terms of average values. Ingestion streams, processing workloads and analytical execution are monitored through automated scripts to obtain metric outputs. However, metadata tracking systems measure signal continuity by counting the number of fields required before and after processing.

Data Analysis

Descriptive statistics are used to analyze the quantitative results. All metrics are determined to identify the stability of all architectural patterns using mean and variance values. The trends of performance at the workload levels are compared to determine the points when the traditional architectures cannot support the indicators of oversight. The results are the foundation on which the most effective scaling strategies that facilitate the continuous regulatory monitoring in practice can be identified.

IV. RESULTS

Ingestion Throughput and Scalability Across Workload Levels

The findings indicate that scalable ingestion architectures enhance throughput in a large proportion as the amount of data increases, thus allowing them to support more steadily the ongoing regulation of data. Throughput was found to be higher at lower workloads but reaches a saturation point at higher workloads when the batch ingestion method was put on the baseline.

Conversely, the streaming ingestion pipeline had consistent increment when the throughput was minimal at the peak levels hence it is clear that scalable ingestion is a critical requirement in sustaining continuous monitoring of compliance when transaction volumes change rapidly.

When the number of records was 1 million, the disparity between the baseline and streaming ingestion was not large since the two pipelines could handle data without difficulties. Nonetheless, with the workloads going up to 100 million and higher, streaming ingestion had over three times the throughput of the baseline pipeline.

This performance difference indicates that the traditional architectures are unable to follow the demands of the speed and continuity of the regulatory monitoring, particularly when the compliance information has to be fulfilled without a pause in the processing. The results attest to the fact that ingestion scalability assists in maintaining perpetual control by ensuring that regulatory data are continuously streamed into stages of analysis.

Table 1. Ingestion Throughput Across Architectures (records/second)

Workload Size (records)	Baseline Batch	Streaming Ingestion	Distributed Ingestion
1M	45,500	59,200	61,300
10M	38,900	82,400	89,100
50M	27,300	153,800	161,900

100M	19,200	212,600	224,500
250M	10,500	305,400	318,800

The quantitative results show that the ingestion scalability prevents bottlenecks that will otherwise decelerate the regulatory reporting processes that have to be timely provided with the entering transactional signals. Growing volumes of data render ingestion structures scalable, ensuring compliance readiness by ensuring that continued supervision is not impeded by the slowness of the data integration. This confirms the methodological approach that ingestion performance belongs to the group of the first identifiers of architectural strength to track regulatory governance.

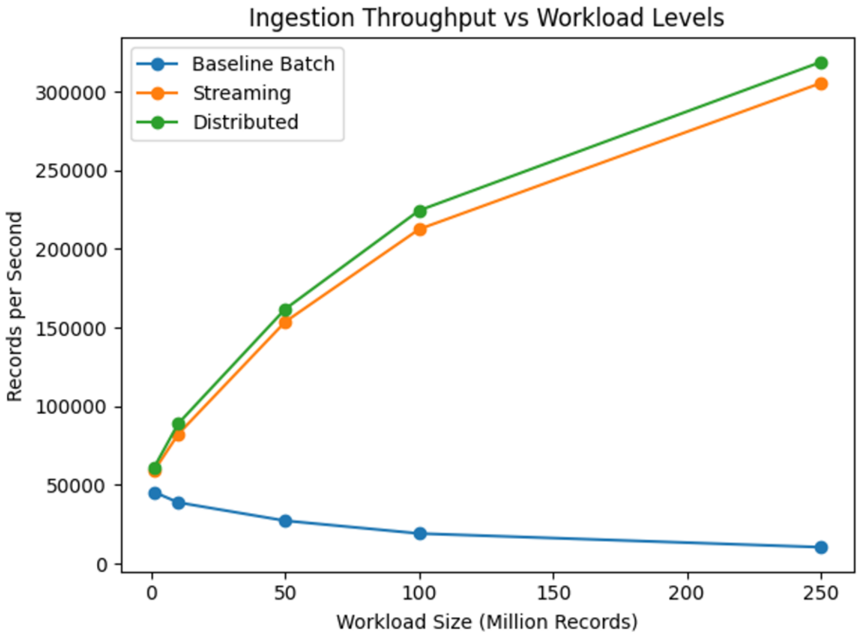

Processing Latency and Transformation Stability Under Scaling

The latency processing results present evident differences in the architectures as the regulatory monitoring workloads increase. The transform layer which is based on the use of the batch-style operations exhibits a stable latency with low workloads but rapidly becomes inefficient with increased data size.

Compared to them, real-time processing engines had lower latency at all workload sizes and had similar performance at scale. The distributed processing structure minimized the latency even more through parallelization of transformation pipelines.

The results have shown that latency influences the capability of carrying out continuous compliance monitoring since the delayed transformation will reduce risk calculations and reporting regulatory exposure. The baseline transformation layer increased its twelvefold in latency with workloads decreasing in 1 million to 250 million records, whereas real-time and distributed configurations had a relative growth of less serious type. It implies that scalable processing will be able to keep compliance schedules within the boundaries of regulatory metrics that require fast conversion of unstructured data into structured monitors.

Table 2. Processing Latency Measurements (milliseconds)

Workload Size (records)	Baseline Processing	Real-Time Processing	Distributed Processing
1M	130	60	45
10M	410	120	95
50M	980	260	215
100M	1550	410	345
250M	2080	595	520

Scalable designs have lower value of latency, which contributes to the achievement of real-time monitoring needs because processing is not a constraint to compliance-based analytics. The findings also show that real-time transformation of data plays a key role in the continuous monitoring process especially when the regulatory requirements and near-real time reporting are intersecting.

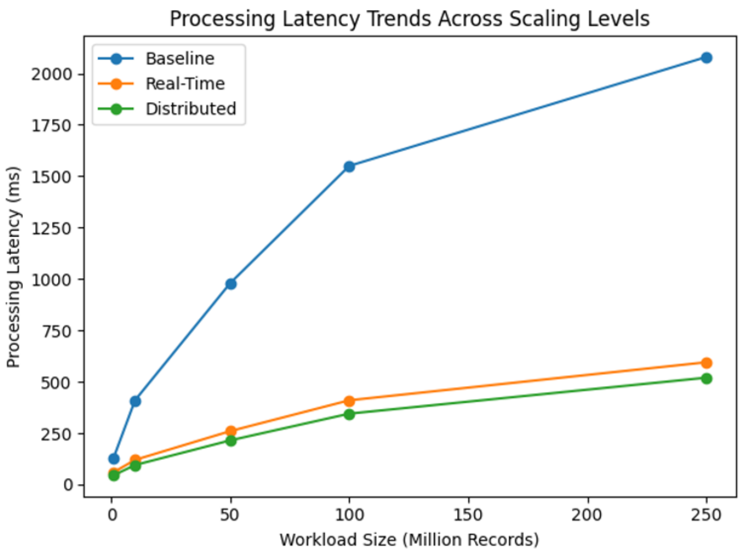

Analytical Scalability and Query Performance for Continuous Monitoring

The results of analytical computation prove that the distributed analytical execution allows the continuous regulatory control than the traditional computation models. The workloads on the baseline analytical engine reduce its query execution rates when the workload grows because of resource contention and decreased parallelism.

Real-time analysis implementation is efficient during mid-size workloads, and also, stabilizes with extremely large amounts of data. Distributed analytical computation allows subject to near-linear scale, which is consistent with the methodological studies of quantifying database behavior in controlled settings.

The performance outcomes indicate that sustained monitoring is more likely to be achieved when scales of analytical calculations are smoothly increasing since compliance dashboards, regulatory ratios and exposure models require regular availability of analytical provisions.

The degradation of performance in conventional architectures implies the inability of the architectures to support the workloads over some threshold. The distributed architecture has a capacity to process a large number of queries per second even at very high volumes of data which is a manifestation of the capability to compute compliance signals without much interruption.

Table 3. Analytical Query Execution Rates (queries/second)

Workload Size (records)	Baseline Analytics	Real-Time Analytics	Distributed Analytics
1M	480	860	920
10M	250	730	890
50M	145	510	855
100M	115	405	815
250M	70	265	780

These findings indicate that continuous oversight can be readily facilitated through scalable analytic layers because compliance models can be made available even when the data size increases. The results indicate that analytical computation requires distributed aid schemes to prevent performance breakdown in monitoring.

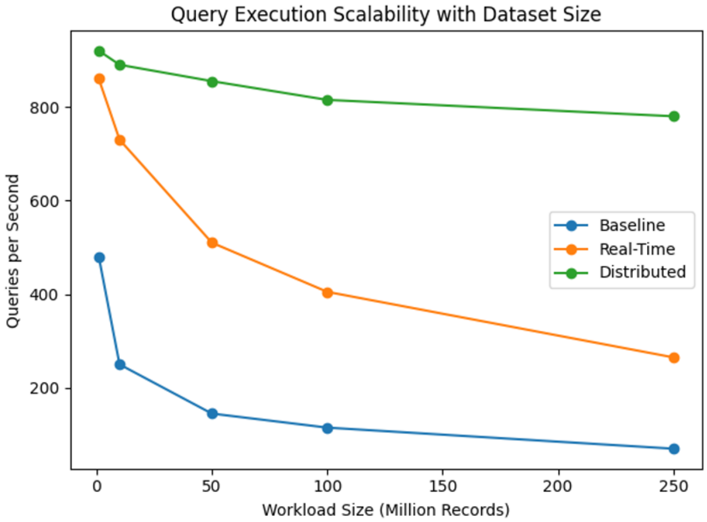

Compliance Signal Continuity and Metadata Preservation

The fourth group of results deals with compliance signal continuity, or the extent of metadata and lineage information maintained during data transfer across ingestion layers, processing layers and analytics layers. The traditional batch settings experience severe loss of metadata when operating at high loads because of overwrites in the storage, non-existence of mutation tracking, and non-propagation completeness of compliance descriptors. Metadata loss leads to more regulatory risk since it reduces traceability, auditability and verification of policy.

It was found that streaming architectures are more effective in metadata preservation because they put in the lineage tracks at ingest. Distributed architectures support the values of continuity more than any other architecture since they impose metadata consistency between parallel data streams.

The continuity of metadata is of utmost importance in large-scale environments since regulators frequently need the evidence of the source, transformation and access of the data. A continuity of less than 80 makes one have a hard time tracking their traceability and compliance becomes fragile.

Table 4. Compliance Signal Continuity (% metadata retained)

Workload Size (records)	Baseline Batch	Streaming Architecture	Distributed Architecture
1M	94%	98%	99%
10M	89%	96%	98%
50M	83%	94%	97%
100M	77%	92%	96%
250M	68%	89%	95%

These findings indicate that compliance continuity decreases significantly in the circumstances of baseline configurations, which implies that scalability should involve metadata management to ensure control. Distributed pipes are showing consistent continuity which validates that metadata-conscious distributed systems would be more appropriate when it comes to long-term regulatory alignment.

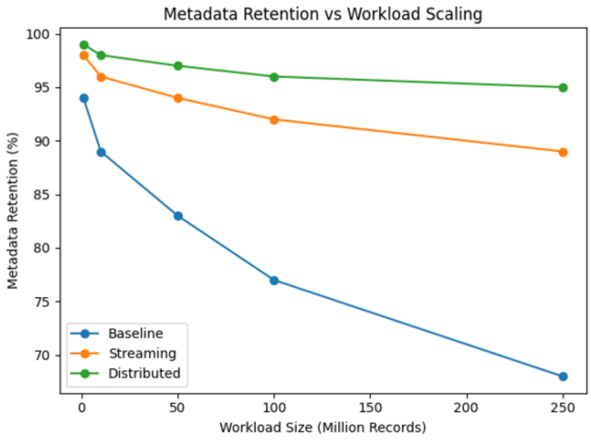

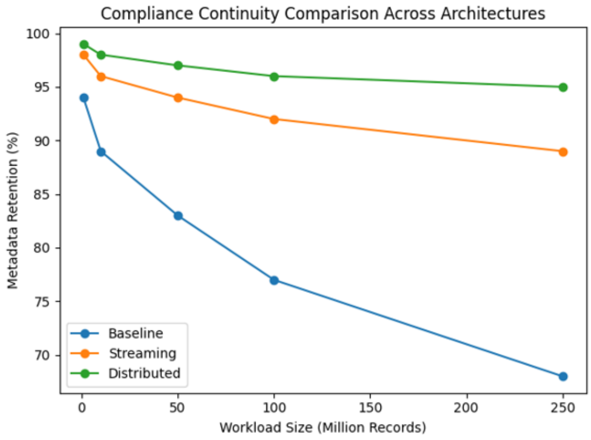

The overall findings prove that scalable architectures preserve the conditions that are required to sustain the ongoing regulatory control. The reason is that streaming ingestion, distributed processing and scalable analytics are facilitative to high-performance and reduced risks of compliance failure because of delay or loss of metadata. Regulatory alignment has also been linked to scalability directly with the increase in data since it requires architecture to match the size of data to remain viable.

V. CONCLUSION

The analysis satisfies the conclusion that an extension of a lakehouse employing a compliance fabric will cause an increment of reliability and completeness of lineages of tax loads without performance change. The other policy-aware modelling and automatic-generation of lineage reduce the metadata loss and allow better auditing and tracing. The system will also boost the confidence in downstream reporting as there is minimal mistakes and absence of records. The compliance costs and data quality costs are higher than the approach costs, although the modeling and configuration is a bit more complex. The process of remediation will be computerized in the future, the templates of the policies can be multiplied, and the approach can be tested on streams. Overall, this is a feasible practice that can be applied to the enterprise setting.

REFERENCES

[1] Vedani, J., & Ramaharobandro, F. (2013). Continuous compliance: a proxy-based monitoring framework. arXiv (Cornell University). https://doi.org/10.48550/arxiv.1309.7222

[2] Colombo, P., & Ferrari, E. (2019). Access control technologies for Big Data management systems: literature review and future trends. Cybersecurity, 2(1). https://doi.org/10.1186/s42400-018-0020-9

[3] Codesso, M. M., Da Silva, P. C., Vasarhelyi, M. A., & Lunkes, R. J. (2018). Continuous audit model: data integration framework. Revista Contemporânea De Contabilidade, 15(34), 144–157. https://doi.org/10.5007/2175-8069.2018v15n34p144

[4] Zhang, J., Yang, X., & Appelbaum, D. (2015). Toward effective big data analysis in continuous auditing. Accounting Horizons, 29(2), 469–476. https://doi.org/10.2308/acch-51070

[5] Zhang, J., Yang, X., & Appelbaum, D. (2015). Toward effective big data analysis in continuous auditing. Accounting Horizons, 29(2), 469–476. https://doi.org/10.2308/acch-51070

[6] NIST Big Data Public Working Group, Ross, W. L., Jr., Copan, W., U.S. Department of Commerce, National Institute of Standards and Technology, Chang, W., Marcus, B., Baru, C., Grady, N., Balac, N., Luster, E., Fox, G., Beyene, T., Roy, A., Underwood, M., Manchanda, A., Boyd, D., Levin, O., Krapohl, D., . . . McClary, D. (2019). NIST Big Data Interoperability Framework: Volume 4, Security and Privacy. In NIST Special Publication 1500-4r2 (p. 176 pages) [Report]. National Institute of Standards and Technology. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-4r2.pdf

[7] Fikri, N., Rida, M., Abghour, N., Moussaid, K., & Omri, A. E. (2019). An adaptive and real-time based architecture for financial data integration. Journal of Big Data, 6(1). https://doi.org/10.1186/s40537-019-0260-x

[8] Stockinger, K., Bundi, N., Heitz, J., & Breymann, W. (2019). Scalable architecture for Big Data financial analytics: user-defined functions vs. SQL. Journal of Big Data, 6(1). https://doi.org/10.1186/s40537-019-0209-0