ISSN: 2582-7219

International Journal of Multidisciplinary Research in Science, Engineering and Technology

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Impact Factor: 8.206

Volume 8, Issue 6, June 2025

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)
(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Building a Resilient ODS Integrating Graph Algorithms and Azure for Enterprise Analytics

Sandeep Sarngadharan
Software Architect, SPRI Partners LLC, USA

ABSTRACT: The establishment of a contemporary Operational Data Store (ODS) design and architecture incorporates investment, financial, operational, and market-related datasets across various countries into one central repository at a multinational Private Equity Firm. The architecture incorporates an event-driven and graph-aware design methodology with a selective refresh approach for overcoming challenges related to very large datasets, multiple dependencies between financial indicators, and real-time updating and synching of data. By utilizing various Azure-native services paired with superior graph algorithms, the Platform offers more scalability around data management, better manages intelligent re-calculation capabilities, and provides near- real-time analytics. By offering a consolidated view in terms of corporate reporting this architecture has also significantly improved overall operational efficiency while simultaneously lowering costs and increasing the speed at which decisions can be made. Additionally, this research provides the groundwork for further developing automation and predictive analytic capabilities while emphasizing the value of creating innovative ODS designs that revolutionize the way in which financial/market-related data is managed.

KEYWORDS: Operational Data Store (ODS), Private Equity Firm, Graph-Aware Design, Azure-Native Services, Real-Time Analytics

I. INTRODUCTION

An Operational Data Store (ODS) is a space for aggregating real-time or almost real-time information from multiple operational systems into a single source, enabling organizations to report on and analyze activity across multiple business units. ODS reduces the time needed to collect the most current information about all operational activities, while serving as a connection between transactional and analytical (data warehouse) systems. An ODS acts as a repository for all types of operational data in a standard manner, and allows for the easy creation of operational reports, improves the quality and reliability of data, breaks down data silos, and supports regulatory compliance by providing quick access to accurate data for regulatory purposes. Additionally, the ODS reduces the time it takes for users to identify problems in their integrated operating environments [1].

In addition to the potential benefits of ODS, ODS also has several challenges, including maintaining a stable environment as new data are received, inadequate performance at high loads, the requirement for highly trained individuals to perform ongoing maintenance, and security issues associated with providing an interface to real-time data. When measuring the differences between ODS and data warehouses, the former focuses on accessing current data as needed with little or no changes made to existing information and is therefore best suited for monitoring operations, while the latter stores historical data for in-depth analysis of historical trends and strategic decision-making over a long period of time. Organizations seeking to have instant access to operational data may want to consider an ODS, but must also be prepared to deal with the maintenance, scalability, and security challenges associated with it [2].

An ODS and a data warehouse are different types of systems with different purposes and functions. An ODS is designed to provide a real-time operational reporting capability based on current transactional data from multiple operational systems to help organizations to quickly resolve day-to-day operational requirements. In contrast, a data warehouse provides a means for the storage of historical data, in support of complex analytical queries and long-term strategic planning [3].While the Operational Data Store (ODS) continues to provide updated and newly recorded data, the Data Warehouse is a repository of historical records, both of which have completely different usage applications. The ODS is focused on day-to-day operations and provides quick access to information for trouble-shooting, while the Data Warehouse supports long-range strategic planning and analytics on data trends. Furthermore, the ODS typically have less complex schemas compared to the Data Warehouse, and the Data Warehouse is considered to be the secondary location

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

for the staging of data, necessitating a more extensive set of ETL processes to prepare the data for analysis. The Query Complexity of both solutions is also disparate; the Data Warehouse is considered to have a significant level of complexity for business users, given the number of combinations that are able to be analyzed from extremely large datasets, whereas the ODS can be more easily understood by business users and has fewer combinations that can be requested from relatively small volumes of real-time data. Overall, the approach to business intelligence for the Data Warehouse is historical-based, whereas the approach to operational agility and immediate access is evident in a plethora of additional capabilities within the ODS [3].

Significantly, an enterprise should consider all of the security issues associated with implementing an ODS due to the inherently sensitive data it holds, primarily because the majority of data held by the ODS relates to private customer and financial records. The primary security issues associated with the implementation of an ODS are: (I) there is a risk of exposing sensitive data due to unauthorized access; (II) vulnerabilities in the integration of varying security measures from various source systems; (III) real-time issues associated with implementing security controls that do not interfere with the performance capabilities of the ODS; and (IV) the challenge for an enterprise to determine appropriate access restrictions to avoid insider threats. Also, compliance with respect to the Data Protection Regulations (e.g. GDPR, HIPAA) will become difficult for an enterprise to maintain due to the integrated nature of the data within the ODS. As the volumes of data stored in an ODS increase, the difficulty associated with monitoring the security of an ODS will increase, thus creating risks associated with the performance of the system and/or with the potential for undetected security breaches. To mitigate the risks associated with storing sensitive data in an ODS, an enterprise should implement a multi-faceted security framework that includes encryption, access restrictions, constant monitoring, and routine auditing of the ODS [4].

A top private equity company with assets exceeding $75 million has greatly improved the company's data management abilities through creating an advanced Operational Data Store (ODS). As the primary cloud architect for this project, I integrated multiple data sources from investment, financial, operational, and market into the new ODS so that the company now has access to real-time analytical capabilities. This increased visibility provides the company with improved capabilities to analyze potential transactions to assist with portfolio management and measuring overall portfolio risk. There were over 50 million data points processed through the new ODS, including approximately 20 million cash flow records, giving the firm insights and flexibility they had never before experienced. Prior to this change, the company relied on a third-party portfolio tracking solution that restricted firms' ability to access their investments and created a lack of transparency between the firm's platform. In this setup, the company could not easily pull data from its various platforms, creating problems with disjointed analysis. In addition, the existing Key Performance Indicators (KPIs) were subject to slow recalculation due to increased volume, and scalability was a challenge due to increasing volumes of data. The new ODS architecture allows the firm to efficiently extract, synchronize, and convert vast volumes of data at enterprise level speeds, allowing for enhanced data-driven decision-making. The IRR (Internal Rate of Return) is a key indicator for measuring the performance of private equity funds and portfolios and gives investors an investor's yearly return based on time of cash flows that were returned to the investors. Higher IRR means capital was deployed most efficiently by the portfolio managers and therefore, provides a better return to the investors.

MOIC or Multiple on Invested Capital represents the total/value/cash return of the fund by giving a complete/aggregated view of both realized/unrealized return. Therefore, it can be seen as looking at value creation without the value of time aspect. Time-weighted return (TWR) calculates the compound growth rate of the portfolio based on the investment cash flows into/out of the fund and is primarily used to compare investment manager performance across numerous different funds. Distributions to Paid In (DPI) express the cash that has been returned to investors versus their total invested capital and only represents realized returns. Collectively, these three metrics provide investors and fund managers an overall view of both investment performance, risk and efficiency, helping both types of entities make better investment decisions [5].

Generally, when people evaluate investment funds, MOIC has been used more frequently than IRR because of wealth creation being more easily assessed/compared with the former than the latter. This trend becomes even more apparent with long-term investments, early-stage portfolios and due to uncertainty surrounding the timing of exits. As a result, MOIC represents a straightforward metric that indicates how much total value has been returned per dollar invested, which is particularly relevant for early-stage venture capital funds, as well as long-term/illiquid assets. Because IRR places more emphasis on the timing of cash flow, it is better suited for providing an assessment of the timing and speed

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

at which an investment provides return and therefore, IRR is used to evaluate how effectively an investment can return value. Generally, both metrics are combined to evaluate the performance of the investor and the funds they manage by taking into consideration both time and total value generated [6].

Currently, a modern ODS integrates data related to the market, financials, operations and investments, allowing businesses and investors alike to assess transactions and portfolios in real-time and analyze risks as they happen. By eliminating data silos and centralizing information from many sources, an ODS improves decision-making efficiency by providing users immediate access to the latest market information. An ODS is a great tool for all leaders because it helps them respond quickly to the ever-changing marketplace, as well as identifies new risks and better manages those risks. Since an ODS provides self-service access to critical information for both tactical and strategic decisions, it eliminates the need for complex batch processing and long-term reporting cycles, which helps facilitate more agile and quicker portfolio management, transactional assessments and risk management processes [7].

With the implementation of modern ODS analytics solutions (such as), the utilization of ODS for investment, financial and operational purposes has also become increasingly common. Solutions are customized for investment management, financial institutions (banks) and service providers alike and allow users to build real-time dashboards and stream live data analysis by utilizing ODS. Examples of ODS-analytics solutions include: Tinybird for high availability/SQL access; ClickHouse Cloud for time-series analysis; Apache Druid for high-high-concurrency analytics; Apache Pinot for both real-time and historical data; Materialize for use with KAFKA; Apache Flink for complex event processing; and Amazon Kinesis for AWS-based environments. Therefore, these types of real-time ODS technologies improve the dashboards used to assess, monitor and evaluate risk, manage the portfolio and produce operational reports through the ability to access all of the above information in a unified/real-time environment.

II. RELATED WORK

There is a growing appreciation for Operational Data Stores (ODS) regarding their contributions to real-time analytics as illustrated in some scholarly writings and publications. A few key resources are Atlan's architectural guide and TDWI's study, which define ODS as a central hub to combine data from multiple sources to enhance operational decision-making and supports the ability to continuously update data while providing a transitional layer for further transfer of that data between an ODS and a warehouse. An ODS functions as a repository containing real-time data and data transformation for tactical decision making in the operational realm, allowing for continued information exchange between operational teams and information technology. Access to data transformation will usually only be available in operational environments through granular inquiries along with the ability to perform regular updates of data [8][9].

However, implementing ODS presents challenges such as significant complexity in integrating data, scalability issues, and the maintenance of data consistency and quality, largely due to the frequency by which ODS updates data. The key advantages of employing an ODS include providing real-time visibility to support operations in industries such as retail and banking, along with the ability to perform unified analytics across multiple analytic platforms, along with ensuring high-quality data prior to entering a warehouse. The downside of ODS is that ODS do not support complicated trend analyses or accommodate historical data; consequently, ODS may require high maintenance and require extensive resources for quasi-real-time processing and may have issues integrating with legacy systems. Overall, while ODS may provide a strong solution for real-time analytics, careful planning must be considered when designing an ODS concerning issues related to integration, scalability, and maintenance. [10][11]

Operational data stores in finance have been the focus of multiple academic and technical studies covering design, integration, security, and Real Time Analytics. In 2025 the first comprehensive study of Automated Financial Operations at Amazon provides an overview of how AWS databases can build a scalable, event-driven operational data store (ODS). This work highlights some key characteristics of successful scalable event-driven systems for managing financial transactional data: scalable solutions, data quality and reconciliation.

Li et al. (2022) suggest a dynamic solution to store cloud-based transactional data utilizing attributes based on fuzzy logic and rough sets, to increase data security and reduce processing delays. Best practices for building scalable data warehouses for financial analytics are discussed by Edapurath (2024), highlighting the challenges of integration, performance, security, and regulatory compliance, which are applicable to ODS. Using a multidimensional view of the

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)
(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

financial industry for data integration and analytics, Ferreira (2017), provides insights into building financial data warehouses that can effectively manage data from multiple sources. While all of the above-mentioned articles focus on challenges associated with maintaining data quality, processing transactions in real-time and regulatory compliance, they all share a common theme of continuing to build scalable, secure and integrated architectures for managing financial operational data stores [12][13][14].

Papers released recently in the industry discuss the evolution of and the use of Operational Data Store (ODS) architecture in modern data environments. These publications also compare ODS and Data Warehouses. RisingWave (2024) has reported that ODS offer real-time operational visibility of your business, allowing for rapid access to data, whereas Data Warehouses support historical analytics and provide insights into how your company operated over time. GigaSpaces (2024) examines next-generation ODS architectures to address some of the latency and concurrency challenges that traditional architectures face. They are focusing on multi-region designs using a data fabric, and on real-time data synchronization so that customers can expand their capabilities in a hybrid cloud architecture and have better scalability and availability.

Materialize's (1998, updates 2023) most recent publication emphasizes that ODS was designed to process large volumes of operational data in real time, while traditional Data Lakes and Data Warehouses support the needs of modern Data Applications with batch data processing. The above papers demonstrate the transition of ODS systems in response to the demands of global organizations, an increase in demand for real-time analytic capabilities and the use of Hybrid Cloud environments, while at the same time continuing to invest in improving performance, scalability, and integration [15][16].

Next-generation designs of ODS architecture can differ significantly from traditional designs, with respect to performance, scalability, real-time capability, and integration flexibility. Traditional ODS architecture relies heavily on either relational databases or NoSQL databases based on disks, making them unsuitable for real-time applications due to latency and concurrency. Additionally, traditional ODS architecture focuses on the use of operational reporting, making it more difficult to support today's digital applications. The manual processes required to perform schema mapping and integrate data from disparate systems create time-consuming bottlenecks. Due to the integral relationship between ODS architecture and their System of Record, traditional ODS architecture suffers from outages and subsequently fails to provide a high level of availability as described in (Table 1) below.

Feature	Traditional ODS	Next-Generation ODS
Architecture	Relational/Disk-based NoSQL	Event-driven, In-memory, Distributed
Data Freshness	Batch, Periodic Updates	Real-time, Continuous Updates
Use Cases	Operational Reporting	Real-time API, Advanced Analytics
Scalability	Limited	High, Multi-region Support
Cost	Lower Initial, Higher OpEx	Lower Long-term, Optimized

Table 1: Traditional vs next-generation Operational Data Stores (ODS)

In contrast, future-generation operational data store (ODS) technology employs distributed, in-memory processing/storage systems with high-performance and high-concurrency capabilities. Future ODS architecture is built using event-driven, real-time processing platforms that are capable of processing many millions of events per second and providing very fast response times, thus enabling rapid reaction to business events. ODS architecture leverages automatic schema discovery coupled with a unified API, thus simplifying integration and accelerating implementation timelines. ODS architecture also decouples the API from transactional systems to maintain application continuum during transactional system outages, supports adaptive tiered storage and facilitates seamless global multi-region data synchronization to maximize performance and minimize cost for hybrid cloud global environments. Finally, by leveraging both historical and real-time data, future ODS will improve analytical models, as well as enable real-time analytics and predictive modeling capabilities, thus providing improved flexibility, reliability, and scalability to achieve business digital transformation initiatives.

Future operational data stores (ODS) have been designed specifically for real-time analytics and scalable architectures. Models developed for next-generation ODS include the person-centric OMOP Common Data Model for standardizing health-related information that can also be applied to finance, and the Microsoft Common Data Model to provide common schemas for the way organizations interact with data. The CDISC Study Data Tabulation Model is used in clinical research or clinical study data sharing to achieve consistent formatting of shared data, and both the Sentinel and PCORnet CDMs support federated analytic processing and the harmonization of data for large-scale analyses. Additionally, the use of graph models continues to increase in popularity for the modeling of complex relationships and to provide more sophisticated query capabilities. Collectively, these models expand the standardization, integration, and analytic capabilities within next-generation ODS environments [17].

III. SYSTEM ARCHITECTURE

For a global private equity firm, the architecture we created allows them to analyze and integrate large amounts of real-time financial data from all of their portfolio companies into one unified view. By being able to merge and analyze multiple types of financial data (investment data, market data, etc.) simultaneously, they were able to perform real-time analysis and generate comprehensive reports. Among the many technical barriers faced in delivering this architecture were the following:

1. No native means of tracking change within the architecture itself
2. The need to intelligently re-calculate interdependent financial metrics
3. Extreme amounts of data
4. The need to control costs and latency for data refreshes.

To overcome these challenges, we designed a graph-aware, event-driven architecture that incorporates a selective refresh approach based on a Directed Acyclic Graph (DAG) that allows for efficient re-calculation of the financial interactions using algorithms such as topological sorting, breadth-first search, etc., while selectively refreshing only those data points that were changed. Additionally, this architecture takes advantage of the Azure ecosystem for data ingestion, orchestration, event handling, and distributed processing. As a result, we were able to resolve a number of enterprise-level issues—including the need for real-time synchronization of financial data across disparate platforms—and also to provide a consistent API throttling framework to ensure the reliability of the API.

The implementation of this architecture provides not only improved data accuracy but also increased accuracy, greater data freshness, reduced processing times, and reduced operating costs. Additionally, it serves as the primary data source for artificial intelligence and machine learning applications and provides a framework for making better investment decisions. Thus, this architecture demonstrates that the use of an event-driven design that incorporates graph structures enables organizations to overcome scalability and performance issues during the development and deployment process. The architecture is built on a modern event-driven model that leverages Azure for enhanced performance and scalability in processing and analyzing large volumes of data across numerous operational contexts. As illustrated in Figure 1 below.

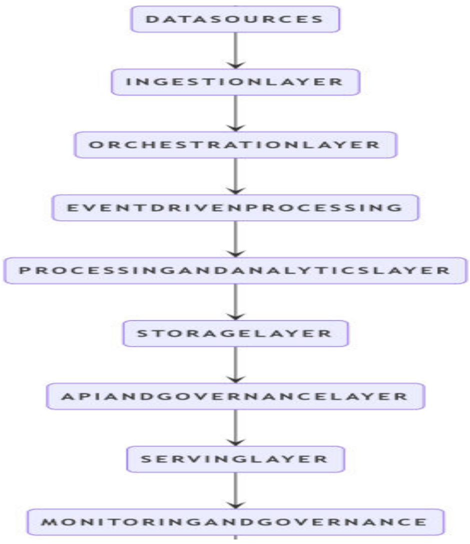

Figure 1: Operational Data Store (ODS) Architecture

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)
(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

1. **Sources of Data:** Internal systems, market feeds/FX, Salesforce, Workday, CapIQ, and portfolio monitoring.
2. **Ingestion Layer:** Azure Functions allow batch and real-time ingestion of data from multiple sources.
3. **Orchestration Layer:** Azure Data Factory is used to manage pipeline executions and scheduling of processing.
4. **Event-Driven Processing:** Azure Event Grid and Service Bus provide near real-time processing by propagating events and triggering changes in near real-time.
5. **Graph Engine:** It models cash flows, valuations, hierarchies, metrics, operational/financial relationships, etc. in a Directed Acyclic Graph (DAG).
6. **Recalculation Engine:** To ensure the accuracy and efficiency of recalculation engines, the only affected metrics are recalculated using a topological sort and BFS algorithm.
7. **ETL & Analytics Layer:** Using Azure Fabric (Spark), it operates a Distributed ETL and Graph- based Recalculation Algorithm.
8. **ODS Storage Layer:** Azure SQL acts as the ODS (Operational Data Store) and is designed specifically for Analytics and Near Real-time access.
9. **API Throttling:** API rate limiting helps to control marketplace requests from Market Data suppliers for them to avoid hitting usage quotas.
10. **Serving Layer:** The ODS will serve near real-time dash boards, reporting and Analytical capabilities for the ODS findings through both BI tools and APIs
11. **Monitoring & Governance:** Azure Monitor and Azure Cost Management provide oversight over the costs associated with using the platform and the health of the systems.

The implemented solution removed the need for manual reconciliations and ensured near real-time updates across systems by increasing data precision and improving operational capacity through almost instantaneous validation of information. The company developed a standardized auto-throttling approach that allowed for the efficient use of market data provider's Application Programming Interface (API) requests while also preventing service interruptions or request counter depletion. By employing this method, the organization has increased compliance with quota and increased reliability. The selective auto-throttling framework dynamically adjusts the APIs request pattern based upon actual usage (against Quota limits) allowing for greater efficiency and reliability and still follows the standard governing API access. These systems not only meet immediate OA (operational accuracy) needs but provide a framework for continued growth with a focus on Consumer and Employer Data Management (CDM & OMD).

The new ODS architecture has also significantly improved operations, with end-to-end data processing time decreased through almost real-time access of data for investment team members' use in the assessment of a portfolio and for transaction processing functions (TPF). Also, increased Accuracy through the elimination of manual reconciliation errors results in all data reflected across all systems immediately in the same way. Increased freshness (90%) of data provides near immediacy in updating risk dashboards and portfolios after market events have occurred.

Cost-wise, the implemented systems yielded substantial financial savings of over 50% on cloud compute and storage costs through the implementation of the selective recalculation methodology along with the resultant decrease in operational overhead due to the elimination of the need for manual processes. Also, the selective refresh model of the ODS has dramatically reduced the burden placed on upstream systems from licensing and infrastructure needs. In a further example of how the ODS is being used, it currently serves as a major data source for building predictive models using Artificial Intelligence (AI) and Machine Learning (ML) Models used for creating risk assessments ensuring the most up-to-date data utilized to build the ai models is current for the existing investments.

Increased accuracy in risk scoring and improved access to reliable time-series historical data for data scientists. From a strategic perspective, the ODS has become the only source of truth for enterprise-level performance reporting, resulting in increased transparency and reliability of financial reporting, faster investment decision making by providing integrated financial metrics to the team, and as a basis for automation and data-driven strategies of the future and successful companies within a competitive investment marketplace. The proposed enhancements to the ods will allow users to advance analytics and predictive modeling using complex artificial intelligence and machine learning models for improved risk assessment and investment forecasting capability.

It is recommended that data governance and compliance processes are automated to ensure compliance with regulations that include automated audit trails and real-time monitoring of data quality. Increased user self-service capabilities for

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

analytics and reporting visualization should allow users to create custom reports without technical knowledge. It is recommended to expand data integration by adding external data sources that will enhance decision-making capability. Additionally, continuous optimization of cloud-native technologies is critical to scale and perform to meet growing data volumes and user demand. Overall, improvements to functionality and user Empowerment will result from enhancements. There are several key measures that should be evaluated to determine the effectiveness and impact of an Operational Data Store (ODS) along with the real-time analytical platform being used. To evaluate the accuracy and quality of the data produced by the platform, some of the metrics commonly utilized are data quality score which evaluates correctness, consistency, completeness of data and two additional measures known as synchronization accuracy and the data validation (success rate).

Operational efficiency levels are evaluated based on time-to-insight, data pipeline latency, query performance and dashboard load time. Reliability of systems is based on system availability, Mean Time to Repair (MTTR) and error rates. Economic effectiveness is based on cost per data job and data processing throughput. The User Experience is determined based on metrics related to User Engagement and Stakeholder Satisfaction. The performance of AI/ML models is evaluated based on Model Accuracy, Time-to-Model Deployment and Frequency of Model Retraining. Finally, Security and Compliance are monitored through compliance with security standards, frequency of security incidents, and response time to risks. The KPIs provide a holistic view of the performance of the ODS platform to stakeholders, thus allowing the opportunity to continue to improve the ODS and demonstrate the value that the ODS provides.

The design of a data platform is based primarily on kpis that are structured around reliability, efficiency, quality of data and business impact. Among these KPIs is a KPI known as data quality score which is designed to ensure that analytics developed using this platform are based upon reliable quality data as measured by consistency, correctness and completeness. A second KPI that reflects the efficiency of moving data from a data source to its destination is data pipeline latency, and as a result, it impacts the speed with which an organization can achieve Insight and make decisions using that insight. system uptime/availability indicates the amount of time when the data platform is operational, which is critical to maintain the ongoing operations of a business. Query performance measures how quickly an organization can complete a data query and is therefore important for user satisfaction and productivity.

Another KPI is dashboard load time, by monitoring how quickly a user can view a dashboard and therefore make a quick decision or respond to information requests. Data processing throughput measures how rapidly data is processed or handled. Error Rate measures the amount of errors that occur when processing data (indicating that there may be Reliability Issues). Cost per data task allows for monitoring the amount paid for processing data which in turn provides insight into the financial efficiency of working with data. Time-to-market for new data pipelines or new datasets provides an indication of how quickly an organization can implement a new source of data, which in turn increases the agility and scalability of an organization. Finally, stakeholder satisfaction is an indication of how users have rated the usability and reliability of the ods and what they perceive to be the true value of the ods for their organizations. When combined, these KPIs provide a clear overview of the ODS platform's performance to ensure that the ODS platform can continue to meet the Business Needs for reliability, speed, quality, and cost-effectiveness.

Estimated KPIs (based upon annual data) for large data platforms are as follows: data quality ranges from 90%-97%; the latency of the data pipeline improved from 15 minutes to 5 minutes; system uptime/availability is estimated between 98% - 99.5%; query performance dropped from 300ms to 150ms; dashboard loading time decreased from 5 seconds to 2 seconds; throughput of data processing has doubled from 10 million records/hour to 20 million records/hour; error rate decreased from 2% to0.5%; cost per task decreased from $0.15 to $0.09; and stakeholder satisfaction improved from 3.8 to 4.5. Graphically representing this data in either line or bar charts will provide trends within the data as described within Figure 2 below:

© 2025 IJMRSET | Volume 8, Issue 3, March 2025| DOI:10.15680/IJMRSET.2025.0803348

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

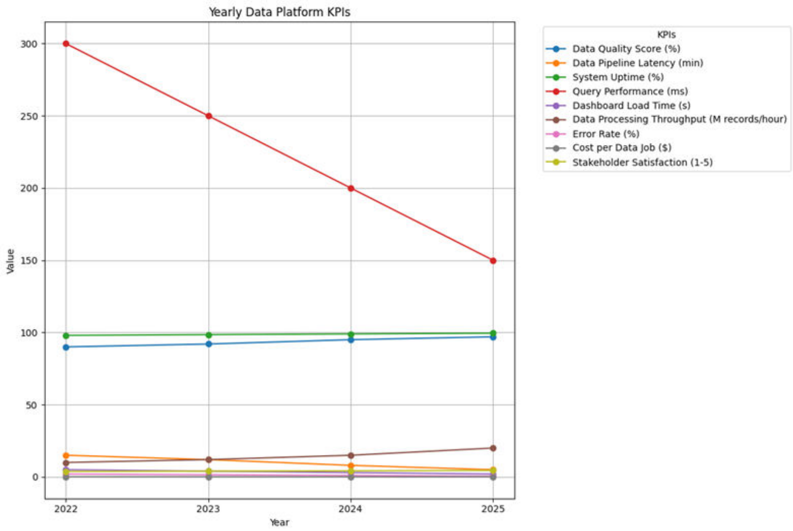

Figure 2: Key Data Platform KPIs

IV. CONCLUSION

The ODS is designed to optimize the ways in which businesses can manage their data by providing real-time analysis, intelligent recalculation of impacted data, and scalable synchronization of systems. This architecture has the potential to provide business with many operational, financial and strategic benefits such as faster time to process, lower costs, improved data accuracy and increased capability for making better decisions. Event driven processing and graph algorithms provide a means of ensuring that only the impacted data points will trigger a recalculation of data; thus, increasing both the efficiency and reliability of the architecture. Also, the development of a single source of truth for enterprise reporting has equipped the investment team with real-time information and has built a strong foundation for future automated processes and predictive analytical capabilities. Future developments for the ODS include more in-depth advanced analytics with the implementation of additional AII/ML models, automation of data governance processes, enhanced self-service analytics and visualizations, enhanced insertion of external data sources, and continued optimization of performance. The future enhancements to the ODS will keep businesses on the cutting edge and will provide organizations with innovative, efficient and strategically successful solutions in the ever-evolving landscape of the investment industry.

REFERENCES

1. “What is an operational data store?”, Ben Lutkevich, Feb 23, 2022, https://www.techtarget.com/searchoracle/definition/operational-data-store.
2. “Operational Data Store vs. Data Warehouse”, September 2, 2024, https://www.trianz.com/insights/operational-data-store-vs-data-warehouse.
3. “ODS vs Data Warehouse: Unveiling the Key Differences”, May 11, 2024, https://risingwave.com/blog/ods-vs-data-warehouse-unveiling-the-key-differences/.
4. “Top 6 Data Security Challenges and How to Address Them”, November 16, 2021, https://www.fortra.com/blog/top-6-data-security-challenges-and-how-address-them.
5. “Key Metrics for Venture Capital Performance: IRR, MOIC, DPI, TVPI, Markup Rate, and Write-Off Rate”, 19 January 2025, https://www.francescatabor.com/articles/2025/1/19/ understanding-key-venture-capital-metrics-markup-rate-moic-and-gross-irr-1.
6. “MOIC vs IRR”, Matt Shibata, December 30, 2022, https://thoughtfulfinance.com/moic-vs-irr/.
7. “Do You Need An Operational Data Store?”, Gaine Technology, October 11, 2023, https://www.gaine.com/blog/do-you-need-an-operational-data-store.
8. “Boost Business Analytics with an Operational Data Store”, Edward Huskin, 2021, https://tdwi.org/articles/2021/04/16/arch-all-boost-business-analysitcs-with-operational-data-store.aspx.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

9. "What is an Operational Data Store: A Complete Guide for 2024", December 28th, 2023, https://atlan.com/what-is/operational-data-store/.
10. "A dimensional warehouse for integrating operational data from clinical trials", Michael A Farnum, Lalit Mohanty, Mathangi Ashok, Paul Konstant, Joseph Ciervo, Victor S Lobanov, Dimitris K Agrafiotis, 2019 Apr 3, https://doi.org/10.1093/database/baz039.
11. "Real-Time Analytics: Concepts, Architectures, and ML/AI Considerations", WEISI CHEN, ZORAN MILOSEVIC, FETHI A. RABHI, ANDREW BERRY, 12 July 2023, https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10183999.
12. "Research on dynamic and secure storage of financial data based on cloud platform", Jingyi Li, January 7, 2022, https://doi.org/10.3233/WEB-210472.
13. "Building Scalable Data Warehouses for Financial Analytics in Large Enterprises", Naveen Edapurath Vijayan, 2024, https://philarchive.org/archive/NAVBSD-2.
14. "Building an Effective Data Warehousing for Financial Sector", José Ferreira, Fernando Almeida, José Monteiro, Automatic Control and Information Sciences, 2017, Vol. 3, No. 1, 16-25, DOI:10.12691/acis-3-1-4.
15. "What Happened to the Operational Data Store?", Arjun Narayan, August 21, 2024, https://materialize.com/blog/operational-data-store/.
16. "How a Next Generation Operational Data Store (ODS) Drives Digital Transformation", April 17, 2023, https://www.gigaspaces.com/blog/operational-data-store-ods/.
17. "Evaluating common data models for use with a longitudinal community registry", Maryam Garza, Guilherme Del Fiol, Jessica Tenenbaum, Anita Walden, Meredith Zozus, 2016 Oct 29, https://doi.org/10.1016/j.jbi.2016.10.016.

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH IN SCIENCE, ENGINEERING AND TECHNOLOGY