



EMERGING TRENDS AND TECHNOLOGICAL ADVANCEMENTS IN DATA LAKES FOR THE FINANCIAL SECTOR: AN IN-DEPTH ANALYSIS OF DATA PROCESSING, ANALYTICS, AND INFRASTRUCTURE INNOVATIONS

HARIHARAN PAPPIL KOTHANDAPANI¹

¹Senior Data Science Analytics Developer at FHLBC.

MS Quantitative Finance @ Washington University in St Louis

[ORCID](#)

Corresponding author: Kothandapani H.P.

© Kothandapani H.,P., Author. Licensed under CC BY-NC-SA 4.0. You may: Share and adapt the material Under these terms:

- Give credit and indicate changes
- Only for non-commercial use
- Distribute adaptations under same license
- No additional restrictions

ABSTRACT The financial sector is increasingly adopting data lakes to manage and analyze vast amounts of structured and unstructured data. This paper explores the latest trends and technological advancements in data lakes, focusing on their application in the financial industry. Key areas of analysis include the integration of artificial intelligence (AI) and machine learning (ML) for predictive modeling, risk assessment, and process automation. Additionally, the paper examines real-time data processing and its importance in enabling immediate decision-making and market responsiveness. The evolution of data lakehouse architecture, which merges data lakes with data warehouses, is discussed as a solution for handling diverse data types and supporting both batch and real-time analytics. On the infrastructure side, the paper explores innovations such as serverless data lakes, which reduce operational complexity and costs while providing scalability. The shift towards hybrid cloud environments is analyzed for its balance of data security and cloud-based scalability. The paper looks at data mesh architectures, which decentralize data management to align with specific business domains, improving scalability and governance. Lastly, the evolution of data lakehouse architecture, which merges data lakes with data warehouses, is discussed as a solution for handling diverse data types and supporting both batch and real-time analytics. This paper aims to provide a clear understanding of how these advancements are shaping the future of data management in the financial sector.

INDEX TERMS artificial intelligence, data lakes, data lakehouse, data mesh, financial industry, hybrid cloud, machine learning

I. INTRODUCTION

The financial industry, characterized by its high frequency of transactions (Pejić Bach et al., 2019), diverse data sources, and stringent regulatory requirements, necessitates a robust data management framework capable of handling vast amounts of structured, semi-structured, and unstructured data. Traditional data management systems, such as data warehouses, have proven inadequate for addressing the growing complexity and scale of financial data, leading to the adoption of data lakes as a more flexible and scalable alternative (Ziegler et al., 2010), (Trelewicz, 2017).

A data lake is a centralized repository that allows for the storage of all types of data in their native formats (Abbasi, 2020). Unlike traditional data warehouses, which require data

to be processed and structured before storage, data lakes enable the storage of raw, unprocessed data, accommodating structured, semi-structured, and unstructured data alike. This characteristic makes data lakes particularly suited for the financial sector, where data can originate from various sources such as transactional systems, market feeds, customer interactions, regulatory filings, and external datasets like social media or news feeds (baccini2013lakes).

The architecture of a data lake is designed to support the ingestion, storage, processing, and analysis of large volumes of data. At the core of a data lake is its storage layer, which is typically based on distributed storage systems. Technologies such as Hadoop Distributed File System (HDFS) or cloud-based storage services like Amazon S3 are commonly used

to provide the necessary scalability and durability for storing petabytes of data. The storage layer in a data lake is highly scalable, allowing financial institutions to store vast amounts of data at a relatively low cost, making it feasible to retain data for extended periods as required by regulatory mandates (Aroraa et al., 2022) (Chessell et al., 2018).

One of the defining characteristics of a data lake is its schema-on-read approach. Unlike data warehouses, which use a schema-on-write method where data must conform to a predefined schema before being stored, data lakes apply the schema only when the data is read or queried. This approach offers significant flexibility, allowing financial institutions to store raw data without the need for extensive preprocessing. As a result, data lakes can accommodate diverse data types, ranging from structured transactional data to unstructured text, images, or audio files. This flexibility is crucial in the financial sector, where the ability to analyze unstructured data, such as customer emails or market sentiment analysis, can provide valuable insights that traditional data warehouses cannot easily support (Zburivsky & Partner, 2021).

The ingestion layer in a data lake architecture is responsible for collecting data from various sources and loading it into the storage layer. In the financial sector, data can be ingested from a wide range of sources, including transactional databases, streaming data from market feeds, batch data from regulatory filings, and unstructured data from customer interactions or external sources. Data ingestion can be performed in real time, near real time, or in batch mode, depending on the requirements of the specific use case. Real-time ingestion is crucial in financial applications like algorithmic trading, where timely access to market data directly influences trading decisions. (Vergadia, 2022).

Data lakes also include a processing layer, which provides the computational resources necessary to transform, analyze, and derive insights from the stored data (Trelewicz, 2017). This layer can leverage a variety of processing engines, such as Apache Spark, Apache Flink, or Presto, to perform distributed data processing tasks. The processing layer supports a wide range of analytics, including batch processing, real-time analytics, machine learning, and ad-hoc querying. In the financial sector, this capability enables institutions to perform complex analyses, such as risk modeling, fraud detection, or customer segmentation, across large datasets with high performance and scalability.

Another key component of a data lake architecture is the catalog and metadata management layer. This layer provides the necessary tools to organize, search, and manage the vast amounts of data stored in the lake. Metadata management is essential in a data lake to ensure that data remains accessible and usable. It includes the creation of a data catalog, which provides a searchable index of all the data stored in the lake, along with metadata describing the data's origin, structure, and any transformations applied to it. In the financial sector, effective metadata management is critical for ensuring data traceability and auditability, which are important for compliance with regulatory requirements (Tanca, 2023).

Data lakes are also characterized by their support for advanced analytics and machine learning. The ability to store and process unstructured data in a data lake opens up opportunities for financial institutions to apply machine learning algorithms to large and diverse datasets. For example, data lakes can be used to train machine learning models on historical transaction data to detect fraudulent activity or to analyze customer behavior and preferences to develop personalized financial products and services. The integration of machine learning with data lakes allows financial institutions to leverage the full breadth of their data assets to gain a competitive edge in the market (Strengtholt, 2023).

Data lakes are often designed to be technology-agnostic, allowing them to work with a wide range of data processing frameworks, analytics platforms, and visualization tools. This interoperability is used in the financial sector, allowing different teams to use the tools and technologies suited to their specific analytical needs. For instance, a data science team might use Apache Spark for machine learning, while a business intelligence team uses Tableau or Power BI for data visualization. The ability of data lakes to support multiple tools and platforms ensures that financial institutions can derive maximum value from their data without being locked into a specific technology stack. As data volumes continue to grow, financial institutions need a data management solution that can scale effectively to accommodate this growth. Data lakes, built on distributed storage systems, can scale horizontally by adding more storage and processing nodes as needed. This scalability ensures that financial institutions can handle the ever-increasing volumes of data without compromising on performance or incurring prohibitive costs. The ability to scale on demand is important in the financial sector, where data spikes can occur due to events like market volatility or regulatory changes. (Pandey & Tripathi, 2016). Traditional data warehouses, with their reliance on expensive, high-performance hardware and their need for extensive data preprocessing, can be costly to maintain and scale. In contrast, data lakes, with their use of commodity hardware and cloud-based storage solutions, offer a more cost-effective option for storing and processing large volumes of data. The cost-effectiveness is important for financial institutions, which must balance data retention for compliance with cost management. Data lakes offer a solution that enables financial institutions to meet regulatory requirements while controlling storage and processing costs.

Data democratization refers to the ability to make data accessible to a broader range of users within an organization, enabling them to explore and analyze data without the need for extensive technical expertise. Data lakes facilitate data democratization by providing a centralized repository of raw data that can be accessed and analyzed by different teams within the organization. This approach allows financial institutions to empower their employees to make data-driven decisions, fostering a culture of innovation and agility. For example, a marketing team might use the data in a lake to develop targeted campaigns based on customer behavior,

Aspect	Application	Example
Predictive Modeling	AI and machine learning are used to develop models that predict future trends, market behaviors, and customer actions based on historical data. These predictions are crucial for making informed decisions in areas like investment strategies and credit risk management.	A financial institution uses machine learning to predict stock price movements by analyzing historical trading data, economic indicators, and news sentiment. This model helps portfolio managers make data-driven investment decisions.
Risk Assessment	Machine learning algorithms can analyze vast datasets to identify patterns that indicate potential risks. This enables real-time risk assessment, allowing institutions to respond to emerging threats more swiftly.	A bank employs AI to monitor transactions for unusual patterns that could indicate fraudulent activity. Identifying these risks early, the bank can prevent fraud and minimize financial losses.
Automation	AI-driven automation streamlines the processes of data ingestion, cleansing, and preparation within data lakes. This reduces the manual effort involved and accelerates the overall data processing workflow.	A company automates the data cleaning process in its data lake by using AI to automatically detect and correct inconsistencies in customer records. This ensures that the data used for analysis is accurate and up-to-date, improving the reliability of the insights generated.

Table 1. AI and Machine Learning Integration in Data Lakes for the Financial Sector

Aspect	Application	Example
Immediate Decision-Making	Financial institutions can make decisions based on current data, which is important in fast-moving areas like trading, where timely decisions are key to capitalizing on opportunities or avoiding losses.	A trading firm uses real-time data from stock exchanges and news feeds to adjust its trading strategies on the fly, ensuring that it can respond to market shifts within seconds to optimize returns.
Market Responsiveness	Real-time data processing allows institutions to quickly adapt to new market conditions, customer behaviors, and potential risks. This agility is crucial in a sector that is constantly changing.	A bank leverages real-time analytics to adjust its interest rates for loans and savings accounts based on live economic data, ensuring competitiveness and alignment with market conditions.
Fraud Detection	Analyzing transaction data as it occurs helps in identifying and preventing fraudulent activities in real time, thus reducing the risk of financial loss.	A payment processing company uses real-time data analysis to detect and block fraudulent transactions as they happen, protecting both the company and its customers from potential losses.

Table 2. Real-Time Data Processing in Financial Sector Data Lakes

while a risk management team uses the same data to assess potential risks and develop mitigation strategies.

The use of data lakes in the financial sector also supports the development of more sophisticated data governance frameworks. Data governance refers to the processes and policies that ensure the availability, usability, integrity, and security of data within an organization. In a data lake, data governance is supported by the metadata management layer, which provides tools for tracking data lineage, managing access controls, and ensuring data quality. Effective data governance is essential in the financial sector, where compliance with regulatory requirements and the protection of sensitive customer data are critical. Data lakes provide a platform for implementing robust data governance practices, ensuring that financial institutions can maintain control over their data assets while maximizing their value.

Financial institutions operate in a highly dynamic environment where the ability to process and analyze data in real-time can provide significant competitive advantages. For instance, real-time data analysis is essential for high-frequency trading, where decisions must be made in milliseconds based on the latest market data. Data lakes, when integrated with stream processing frameworks such as Apache Kafka or Apache Flink, can ingest and process real-time data streams alongside historical data, enabling financial institutions to make informed decisions in real-time. This capability is also valuable in other areas, such as fraud detection or customer

service, where timely access to data can improve outcomes.

Data lakes have also enabled the integration of new and emerging data sources in the financial sector. As the volume and variety of data continue to grow, financial institutions are increasingly looking to leverage external data sources, such as social media, news feeds, or alternative data, to gain insights and inform decision-making. Data lakes provide a flexible and scalable platform for integrating these new data sources, allowing financial institutions to combine internal and external data in a single repository. This integration can lead to more comprehensive analyses and better decision-making, as financial institutions can consider a broader range of factors when assessing risks, identifying opportunities, or developing strategies.

The role of data lakes in supporting advanced analytics and artificial intelligence (AI) in the financial sector cannot be overstated. The ability to store and process large amounts of diverse data in a data lake provides the foundation for developing and deploying AI models that can analyze complex datasets and uncover hidden patterns. Financial institutions are increasingly using AI and machine learning to automate processes, enhance customer experiences, and improve risk management. Data lakes enable these AI models by providing the necessary data infrastructure to support their training, deployment, and ongoing refinement. For example, a financial institution might use a data lake to store and analyze transaction data, customer interactions, and market

data to develop an AI-based fraud detection system that can identify suspicious activities in real time.

Moreover, the adoption of cloud-based data lakes has further enhanced their appeal in the financial sector. Cloud-based data lakes offer the benefits of scalability, flexibility, and cost-efficiency, while also providing access to advanced analytics and machine learning tools. Financial institutions can leverage cloud-based data lakes to quickly scale their data storage and processing capabilities without the need for significant upfront investments in hardware and infrastructure. This approach also allows financial institutions to take advantage of the latest advancements in data processing and analytics, as cloud providers continuously update and improve their offerings. The use of cloud-based data lakes can also support global operations, enabling financial institutions to store and analyze data across multiple regions while ensuring compliance with local regulations.

This paper examines the latest trends and technological advancements in data lakes that are relevant to the financial sector. We will discuss how integrating AI and machine learning, adopting real-time data processing, and evolving architectural designs like the data lakehouse are transforming the capabilities of data lakes. Additionally, we will explore infrastructure innovations such as serverless architectures, hybrid cloud environments, and data mesh, which offer new ways to manage and secure financial data effectively.

II. DATA PROCESSING AND ANALYTICS

A. AI AND MACHINE LEARNING INTEGRATION

AI and machine learning have become essential tools in the financial sector in data processing and analytics. Integrating these technologies with data lakes—repositories that store large amounts of raw data—enhances the ability of financial institutions to analyze complex datasets more accurately and efficiently (Ertel, 2018).

AI and machine learning are used in predictive modeling, where they are used to forecast future trends, market behaviors, and customer actions by analyzing both historical and real-time data (Charniak, 1985). These models often utilize algorithms such as linear regression, decision trees, and neural networks, which are trained on extensive datasets to identify patterns and relationships between variables. These predictive models are applied in various areas, including stock price prediction, credit scoring, and economic forecasting. Techniques such as ensemble learning, which combines multiple models to improve overall predictions, and deep learning, which captures complex data patterns, further enhance the accuracy of these predictions (Boden, 1996).

Algorithm 1 Real-Time Data Ingestion and Processing for Immediate Decision-Making Using Data Lakes

Input: Streaming data D , Data lake DL , Decision model M_d , Threshold τ

Output: Insights I

while new data $D_i \in D$ is available **do**

Step 1: Data Ingestion into Data Lake

Ingest data D_i into data lake DL $t_i \leftarrow \text{timestamp}(D_i)$

Step 2: Data Preprocessing in Data Lake

$D'_i \leftarrow \text{clean_and_normalize}(DL[D_i])$ // Clean and normalize data within the data lake

$D''_i \leftarrow \text{transform}(D'_i)$ // Apply necessary transformations in data lake

Step 3: Feature Extraction

$F_i \leftarrow \text{extract_features}(D''_i)$ // Extract features from preprocessed data in data lake

Step 4: Decision-Making Analysis

$P_d(D_i) \leftarrow M_d(F_i)$ // Apply decision-making model M_d stored in data lake

Step 5: Immediate Decision Making

if $P_d(D_i) > \tau$ **then**

$I_i \leftarrow \text{generate_insight}(P_d(D_i))$ // Generate actionable insight

$\text{update_dashboard}(I_i)$ // Update decision-maker's dashboard with I_i

end

Step 6: Logging and Monitoring in Data Lake

$\text{log}(DL, t_i, D_i, P_d(D_i), I_i)$ // Log data, predictions, and insights in data lake

$\text{monitor_model_performance}(DL[M_d], P_d(D_i))$ // Monitor and update model performance in data lake

end

In risk assessment, machine learning algorithms analyze vast datasets to identify patterns that indicate potential risks. Unsupervised learning techniques like clustering and anomaly detection are used to uncover hidden relationships and detect anomalies within data, which may signal fraud or financial irregularities. For example, clustering algorithms can group similar transactions or accounts, making it easier to spot unusual patterns. Anomaly detection is useful for real-time risk assessment, enabling institutions to respond quickly to emerging threats. Additionally, reinforcement learning is increasingly used to develop adaptive risk management systems that can adjust dynamically to changing financial environments.

AI-driven automation streamlines data ingestion, cleansing, and preparation processes within data lakes. Traditional data processing workflows, which are often manual and time-consuming, are made more efficient through AI. Natural

language processing (NLP) techniques, for instance, can automatically categorize and extract information from unstructured data sources such as financial news, reports, and social media. This data is then ingested into data lakes for further analysis. AI-powered data cleansing tools also use machine learning algorithms to detect and correct inconsistencies and other data quality issues, ensuring that the datasets used for analytics are accurate and reliable. This automation reduces the manual effort required and speeds up the overall data processing workflow, allowing financial institutions to derive insights more quickly and at a lower operational cost (Charniak, 1985).

The integration of AI and machine learning into data lakes enables financial institutions to significantly enhance their analytical capabilities.

B. REAL-TIME DATA PROCESSING

Real-time data processing is vital in the financial sector, where the ability to act on current information is crucial. Implementing real-time data streaming and analytics within data lakes allows financial institutions to make immediate decisions and improves their responsiveness to rapidly changing market conditions. This capability is important in areas such as trading, risk management, and fraud detection, where the timeliness of data processing can have significant financial implications.

Real-time data processing enables financial institutions to make decisions based on the most current data available. This is especially critical in fast-moving areas like trading, where milliseconds can mean the difference between profit and loss. Using real-time data streams, traders can monitor market movements, execute trades, and adjust strategies instantly, thus taking advantage of opportunities or avoiding potential losses. Technologies like distributed computing frameworks such as Apache Kafka for data streaming and Apache Flink or Apache Spark for real-time analytics support the processing of high-velocity data feeds, ensuring that financial institutions operate with the most up-to-date information.

Market responsiveness is another key benefit of real-time data processing in the financial sector. The ability to quickly adapt to new market conditions, customer behaviors, and potential risks is essential in an industry characterized by constant change and volatility. Real-time analytics enable institutions to continuously monitor market indicators, economic data, and other relevant factors, allowing them to swiftly adjust their strategies and operations in response to emerging trends. In algorithmic trading, for instance, real-time data feeds are used to automatically trigger buy or sell orders based on pre-defined criteria, ensuring that trading strategies remain aligned with current market conditions. Real-time customer data analysis also allows financial institutions to personalize services and offers, enhancing customer engagement and satisfaction.

Algorithm 2 Adaptive Market Response via Real-Time Analytics Using Data Lakes

Input: Market data M , Behavior data B , Data lake DL , Strategy model S , Thresholds θ_M, θ_B

Output: Strategy adjustments S_t

while new data $M_t \in M$ and $B_t \in B$ is received **do**

Step 1: Data Ingestion into Data Lake

Ingest market data M_t and customer behavior data B_t into data lake DL $t \leftarrow$ current_timestamp()

Step 2: Data Preprocessing and Storage

$M'_t \leftarrow$ clean_and_normalize($DL[M_t]$) $B'_t \leftarrow$ clean_and_normalize($DL[B_t]$)

Step 3: Change and Anomaly Detection

$\Delta M_t \leftarrow$ detect_change(M'_t) $\Delta B_t \leftarrow$ detect_anomaly(B'_t)

if $\Delta M_t > \theta_M$ **then**

$F_M \leftarrow$ extract_features(ΔM_t) $A_M \leftarrow$ analyze_market($DL[S], F_M$) // Analyze market changes using strategy model from data lake

end

if $\Delta B_t > \theta_B$ **then**

$F_B \leftarrow$ extract_features(ΔB_t) $A_B \leftarrow$ analyze_behavior($DL[S], F_B$) // Analyze behavior anomalies using strategy model from data lake

end

Step 4: Strategy Adjustment and Deployment

if $\Delta M_t > \theta_M$ or $\Delta B_t > \theta_B$ **then**

$S_t \leftarrow$ update_strategy($DL[S], A_M, A_B$) // Update strategy in data lake
deploy(S_t) // Deploy updated strategy

end

Step 5: Logging and Monitoring in Data Lake

log($DL, t, M_t, \Delta M_t, B_t, \Delta B_t, S_t$)
monitor_performance($DL[S], A_M, A_B$)
// Monitor and update strategy model performance

end

Fraud detection is a critical application of real-time data processing in the financial sector. Through analyzing transaction data as it occurs, financial institutions can identify and prevent fraudulent activities in real-time, reducing the risk of financial loss. Machine learning models, specially those based on supervised learning techniques like logistic regression and support vector machines, are employed to detect patterns indicative of fraudulent behavior. These models are trained on historical transaction data labeled as legitimate or fraudulent, enabling them to identify suspicious activities with a high degree of accuracy. Unsupervised learning techniques such as clustering and anomaly detection are also used to identify outliers and irregularities in transaction data that may not match known fraud patterns but still require investigation. The real-time aspect of this analysis is crucial, as it enables institutions to take immediate action, such as

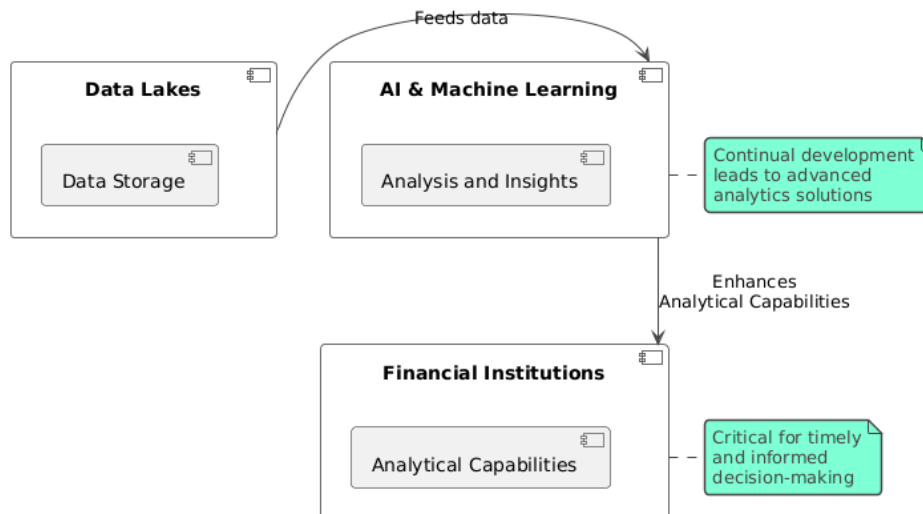


Figure 1. The diagram describes the flow of data from data lakes to AI and machine learning systems, which in turn enhance the analytical capabilities of financial institutions. It highlights the process and impact of the integration, suitable for conveying how technological advancements bolster financial operations.

blocking a transaction or alerting the customer, minimizing potential losses.

Algorithm 3 Real-Time Fraud Detection in Financial Transactions using Data Lakes

Input: Transaction data T , Data lake DL , Fraud detection model M_f , Threshold θ_f

Output: Fraud alerts A

while new transaction $T_i \in T$ is received **do**

Step 1: Data Ingestion into Data Lake

Ingest transaction data T_i into data lake DL $t_i \leftarrow \text{timestamp}(T_i)$

Step 2: Data Preprocessing in Data Lake

$T'_i \leftarrow \text{clean_and_normalize}(DL[T_i])$ // Clean and normalize data in data lake

Step 3: Feature Extraction

$F_i \leftarrow \text{extract_features}(T'_i)$ // Extract relevant features from preprocessed data

Step 4: Fraud Probability Estimation

$P_f(T_i) \leftarrow M_f(F_i)$ // Apply fraud detection model from data lake

Step 5: Fraud Detection

if $P_f(T_i) > \theta_f$ **then**
 | $A_i \leftarrow \text{raise_alert}(T_i, P_f(T_i))$

end

Step 6: Logging and Monitoring in Data Lake

$\text{log}(DL, t_i, T_i, P_f(T_i), A_i)$ // Log data and results in data lake
 $\text{update_model_performance}(DL, M_f, P_f(T_i))$
 // Update model performance metrics stored in data lake

end

for immediate decision-making using data lakes is designed to manage the continuous influx of streaming data from multiple financial sources, ensuring decisions are based on the most current information. Data D_i is ingested into a data lake DL , where each data point is timestamped to maintain sequence and context. The data lake acts as a scalable and efficient central repository for data storage and management. The ingested data undergoes preprocessing, including cleaning and normalization, resulting in a refined dataset D'_i . Further transformations convert D'_i into D''_i , aligning the data with the decision-making model's requirements. The preprocessing and transformation steps leverage the data lake's capabilities to handle large data volumes and apply consistent, complex processing operations, ensuring the data is optimally formatted for analysis.

Feature extraction follows, where relevant features F_i are derived from the transformed data D''_i . These features represent key attributes necessary for informed decision-making and are fed into a pre-trained decision-making model M_d , stored within the data lake. The model computes a decision probability $P_d(D_i)$ for each data point D_i , given by $P_d(D_i) = M_d(F_i)$. This probability indicates the likelihood that an event or outcome is significant enough for an immediate response. Maintaining the model within the data lake ensures the decision-making process is efficient and scalable, allowing dynamic updates as new data or changing conditions arise.

The algorithm concludes with decision-making and logging. The computed decision probability $P_d(D_i)$ is compared against a predefined threshold τ , determining whether an actionable insight I_i should be generated. If $P_d(D_i) > \tau$, an insight I_i is produced and updated on a decision-maker's dashboard for real-time response. All relevant data, including the input, computed probability, and generated insights, are logged back into the data lake for future reference and

The algorithm for real-time data ingestion and processing

auditing. This logging supports continuous monitoring and performance evaluation of the decision-making model M_d , ensuring accuracy and reliability over time. The data lake serves as the backbone of this operation, providing a unified platform for data storage, processing, model management, and decision support, enabling financial institutions to react swiftly to real-time developments.

The algorithm for adaptive market response via real-time analytics using data lakes aims to leverage the data lake's capabilities for the continuous ingestion, analysis, and strategic adjustment of market and customer behavior data. Streaming market data M and customer behavior data B are ingested into the data lake DL . The data undergoes cleaning and normalization within the data lake, producing preprocessed versions M'_t and B'_t . Significant changes in market data ΔM_t and anomalies in customer behavior ΔB_t are detected using sophisticated algorithms. These detection processes utilize the data lake's computational power and storage capabilities to manage large data volumes efficiently.

Once changes and anomalies are detected, the algorithm extracts features F_M and F_B from the change data ΔM_t and anomaly data ΔB_t , respectively. These features are analyzed using a pre-trained strategy model S , stored within the data lake. The analysis produces insights A_M and A_B , which inform strategic adjustments. If either ΔM_t exceeds the threshold θ_M or ΔB_t exceeds θ_B , the strategy S_t is updated in the data lake to respond to the detected changes and anomalies. The updated strategy is then deployed in real-time to adapt to market conditions and customer behaviors.

The final steps involve logging and monitoring within the data lake. All relevant data, including the ingested data, detected changes, anomalies, and adjusted strategies, are logged for future reference and auditing. The performance of the strategy model S is continuously monitored and updated within the data lake, ensuring it remains effective as market conditions evolve. The data lake provides a centralized platform for managing the entire process, from data ingestion and preprocessing to feature extraction, model analysis, strategy adjustment, and performance monitoring. This integration enables financial institutions to dynamically adapt their strategies in real-time, responding effectively to changing market conditions and customer behaviors (macey202197).

The algorithm for real-time fraud detection in financial transactions using data lakes leverages the data lake's capabilities to continuously monitor and analyze transaction data, detecting potential fraud through advanced machine learning models and predefined rules. Transaction data T is ingested into the data lake DL , where each transaction T_i is timestamped. The data undergoes preprocessing, including cleaning and normalization, resulting in a refined dataset T'_i . Features F_i are extracted from T'_i , capturing key characteristics and patterns indicative of fraudulent behavior. These features are fed into a pre-trained fraud detection model M_f , stored within the data lake.

The model computes a fraud probability $P_f(T_i)$ for each transaction T_i , given by $P_f(T_i) = M_f(F_i)$. This proba-

bility is compared against a predefined threshold θ_f , and if $P_f(T_i) > \theta_f$, the transaction is flagged as potentially fraudulent, and an alert A_i is generated. The data lake's centralized storage and processing capabilities enable efficient and scalable analysis, ensuring timely detection of fraudulent activities. All relevant data, including transaction details, computed probabilities, and alerts, are logged into the data lake for auditing and future analysis.

Continuous monitoring and performance evaluation of the fraud detection model M_f are integral to the algorithm, facilitated by the data lake's infrastructure. The model's performance is updated based on feedback from identified fraud cases and false positives, ensuring its accuracy and reliability over time. The data lake serves as the backbone of the fraud detection process, providing a unified platform for data storage, processing, model management, and decision support. This approach enables financial institutions to effectively identify and respond to fraudulent activities in real-time, minimizing financial losses and enhancing security measures. Real-time data processing in data lakes provides financial institutions with a competitive edge by enabling them to act on fresh data insights as they emerge. This capability is not only about speed but also about processing and analyzing vast amounts of data in real-time, made possible by the scalability and flexibility of data lakes. Integrating real-time data processing with advanced analytics, financial institutions can stay ahead of the competition, improve operational efficiency, and enhance their risk management capabilities. As the financial sector continues to evolve, the demand for real-time data processing and analytics is expected to grow, driving further innovations in this area and enabling institutions to better navigate the complexities of the modern financial sector.

C. EDGE COMPUTING INTEGRATION

Edge computing is increasingly being integrated with data lakes in the financial sector to enhance real-time data processing and analytics capabilities. Edge computing involves processing data closer to its source—at the network's edge—rather than relying on centralized data centers. This approach reduces latency, improves data privacy, and enhances the overall efficiency of data processing, making it highly valuable for financial applications that require rapid analysis and response.

One of the primary benefits of edge computing integration is the reduction of latency in data processing. In traditional cloud-based architectures, data must be transmitted from the source to a centralized server for processing and then back to the endpoint, which can introduce significant delays in time-sensitive financial applications such as real-time trading or fraud detection. Through processing data at the edge, near the point of generation, these delays are minimized, allowing for faster analysis and decision-making. For example, in high-frequency trading, where milliseconds can determine the profitability of a trade, the ability to process and act on data in real-time is crucial. Edge computing enables financial institutions to execute trades based on the most

Aspect	Application	Example
Reduced Latency	By processing data at the edge—near its point of generation—institutions can reduce the time it takes to analyze and act on data. This is especially valuable in applications like real-time trading and fraud detection.	A financial trading platform uses edge computing to process and analyze trading data locally, enabling real-time decision-making that is crucial for high-frequency trading strategies.
Enhanced Data Privacy	Processing data at the edge minimizes the need to transmit sensitive information over networks, improving data privacy and security.	A bank implements edge computing in its ATMs, where transaction data is processed locally to reduce the amount of sensitive data sent over networks, thereby enhancing privacy and security.
Better Customer Experience	Faster data processing at the edge can improve service responsiveness, leading to a better customer experience in applications like mobile banking and automated teller machines (ATMs).	A mobile banking app utilizes edge computing to provide real-time balance updates and transaction processing, ensuring that users experience minimal delays, which enhances overall satisfaction.

Table 3. Edge Computing Integration in Data Lakes for the Financial Sector

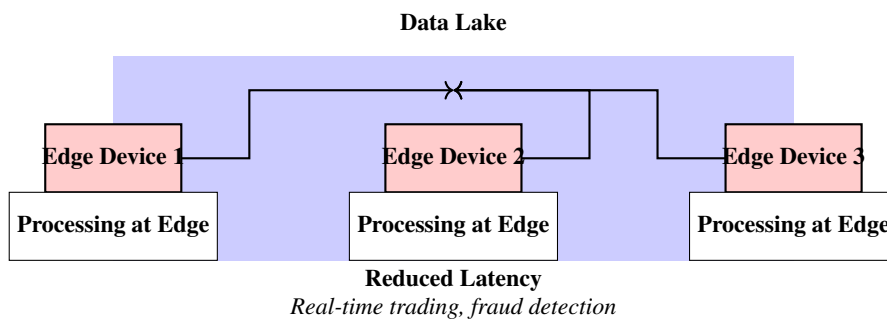


Figure 2. Integration of Edge Computing with Data Lakes in the Financial Sector. Edge devices process data closer to its source, reducing latency and enabling real-time analytics for applications like trading and fraud detection.

current market data, thereby enhancing their competitiveness and operational efficiency.

Enhanced data privacy is another significant advantage of edge computing integration. In the financial sector, sensitive data such as transaction records, customer information, and financial statements must be handled with the utmost care to prevent unauthorized access and data breaches. Processing data locally at the edge, the need to transmit sensitive information over potentially insecure networks is reduced, which helps to safeguard data privacy and security. This localized processing also aligns with regulatory requirements in many jurisdictions that mandate strict controls over the transmission and storage of personal and financial data. Furthermore, by limiting the amount of data that needs to be transferred to centralized servers or the cloud, financial institutions can reduce their exposure to potential cybersecurity threats.

Edge computing also contributes to a better customer experience by improving the responsiveness of financial services. For instance, in mobile banking applications or automated teller machines (ATMs), edge computing can process customer requests and transactions on-site, reducing the time it takes to complete operations. This faster processing translates into a more seamless and efficient user experience, which is critical in an industry where customer satisfaction is closely tied to service speed and reliability. Moreover, edge computing supports the implementation of advanced features such as personalized financial advice or instant transaction notifications, which require real-time data analysis to deliver

relevant and timely information to customers.

The integration of edge computing with data lakes allows financial institutions to extend their real-time data processing capabilities while maintaining strong data privacy standards. Through processing data closer to where it is generated, edge computing reduces latency and enhances the speed of decision-making, which is essential in the fast-paced financial sector. At the same time, by keeping sensitive data local, edge computing improves data privacy and security, helping institutions comply with regulatory requirements and protect customer information. The adoption of edge computing is likely to increase, providing institutions with a more efficient and secure framework for handling real-time data analytics. This integration represents a significant step forward in the ability of financial institutions to leverage data effectively while meeting the demands of modern digital finance.

III. INFRASTRUCTURE AND ARCHITECTURAL INNOVATIONS

A. SERVERLESS DATA LAKES

Serverless architecture is gaining traction in the financial sector as a means to enhance data management efficiency. This architectural model allows financial institutions to simplify operations, reduce costs, and maintain flexibility and scalability without the burden of managing underlying server infrastructure. In a serverless environment, the cloud service provider manages the server infrastructure, automatically allocating resources as needed to handle data processing tasks.

This approach is well-suited to data lakes, where the volume and complexity of data can vary significantly.

The adoption of serverless data lakes simplifies data management by abstracting the complexity of infrastructure maintenance (John & Misra, 2017). Traditionally, managing a data lake requires significant oversight of the hardware and software stack, including server provisioning, scaling, patching, and maintenance. These tasks are resource-intensive and divert IT teams from focusing on more strategic activities such as data analysis, security, and the development of new data-driven products. In a serverless data lake, the cloud provider handles all these infrastructure concerns, freeing IT personnel to concentrate on extracting value from data rather than managing the environment in which it resides. This shift allows for a more agile and responsive IT organization, capable of rapidly adapting to changing business needs and technological advancements.

Cost efficiency is a significant advantage of serverless data lakes in the context of financial institutions that often face tight budget constraints. In traditional data management systems, institutions must provision enough server capacity to handle peak loads, leading to over-provisioning and underutilization during periods of lower demand. This results in wasted resources and higher operational costs. Serverless architecture, on the other hand, operates on a pay-as-you-go model, where institutions are billed only for the computing resources they actually consume. This approach aligns costs directly with usage, reducing the financial burden associated with idle resources and providing a more predictable cost structure. For financial institutions, this model can lead to substantial cost savings in environments with variable workloads where data processing demands can fluctuate significantly.

Scalability is offering financial institutions the ability to automatically adjust to changing data volumes and processing needs. As data volumes grow or processing demands increase, the serverless architecture seamlessly scales the underlying resources to meet these requirements without the need for manual intervention. This scalability is crucial for financial institutions that deal with large, dynamic datasets and require real-time analytics capabilities. For example, during periods of high market volatility, trading systems may need to process significantly more data than during stable periods. A serverless data lake can automatically scale to accommodate this surge, ensuring that analytical performance remains consistent regardless of the workload. This flexibility allows financial institutions to handle varying workloads efficiently, avoiding the risks associated with over-provisioning or under-provisioning resources.

The implementation of serverless data lakes represents a streamlined approach to data management (Gorelik, 2019), enabling financial institutions to handle their data more effectively while reducing operational overhead. Because of the eliminating the need to manage infrastructure, institutions can focus on higher-value tasks such as data analysis, model development, and decision-making. The pay-as-you-go cost

model offers financial efficiency, ensuring that institutions only pay for the resources they use, which is beneficial in managing unpredictable or fluctuating workloads. Finally, the ability to scale automatically ensures that institutions can maintain performance and reliability even as data demands change. As the financial sector continues to evolve, the adoption of serverless data lakes is likely to increase, providing a more agile, cost-effective, and scalable solution for data management in an increasingly data-driven industry.

B. HYBRID CLOUD ENVIRONMENTS

Hybrid cloud environments are increasingly being adopted by financial institutions as a strategic solution to balance the need for data security with the scalability and cost-efficiency benefits of cloud computing. With combining on-premises infrastructure with cloud-based solutions, these environments offer a tailored approach that can meet the diverse and stringent requirements of the financial sector.

Enhanced security is one of the primary advantages of hybrid cloud environments. Financial institutions often deal with highly sensitive data, such as personal customer information, transaction records, and confidential financial reports, which are subject to strict regulatory requirements. Storing this critical data on-premises, institutions retain full control over their most sensitive information, ensuring compliance with regulatory mandates and reducing the risk of unauthorized access. On-premises storage also allows for the implementation of robust security measures that are specifically tailored to the institution's needs, including advanced encryption, multi-factor authentication, and regular security audits.

Scalable solutions are a key benefit of leveraging cloud services in a hybrid environment. Non-sensitive data, along with less critical applications, can be hosted in the cloud, providing financial institutions with the ability to quickly scale their operations as data volumes increase or as new applications are deployed. This flexibility is advantageous in scenarios where data processing demands can vary significantly, such as during peak trading periods or in response to market events. Cloud platforms offer virtually unlimited scalability, allowing institutions to access additional computational resources on-demand without the need to invest in and maintain additional on-premises infrastructure.

Cost optimization is also significant advantage of hybrid cloud environments. Financial institutions can optimize their IT spending by strategically distributing workloads between on-premises and cloud resources. Critical workloads that require high security and low latency can be run on-premises, where institutions have already made significant capital investments. Meanwhile, less critical tasks and non-sensitive data processing can be offloaded to the cloud, taking advantage of the cloud's pay-as-you-go pricing model. This approach allows institutions to reduce capital expenditures on infrastructure while benefiting from the operational flexibility and lower costs associated with cloud services.

Aspect	Application
Simplified Management	Serverless data lakes eliminate the need for managing underlying infrastructure, allowing IT teams to focus on data analysis and strategy instead of maintenance tasks.
Cost Efficiency	The serverless model operates on a pay-as-you-go basis, meaning institutions only pay for the computing resources they use, which can lead to significant cost savings.
Scalability	Serverless data lakes automatically scale to meet changing data volumes and processing needs, ensuring that financial institutions can handle varying workloads without over-provisioning resources.

Table 4. Benefits of Serverless Data Lakes in the Financial Sector

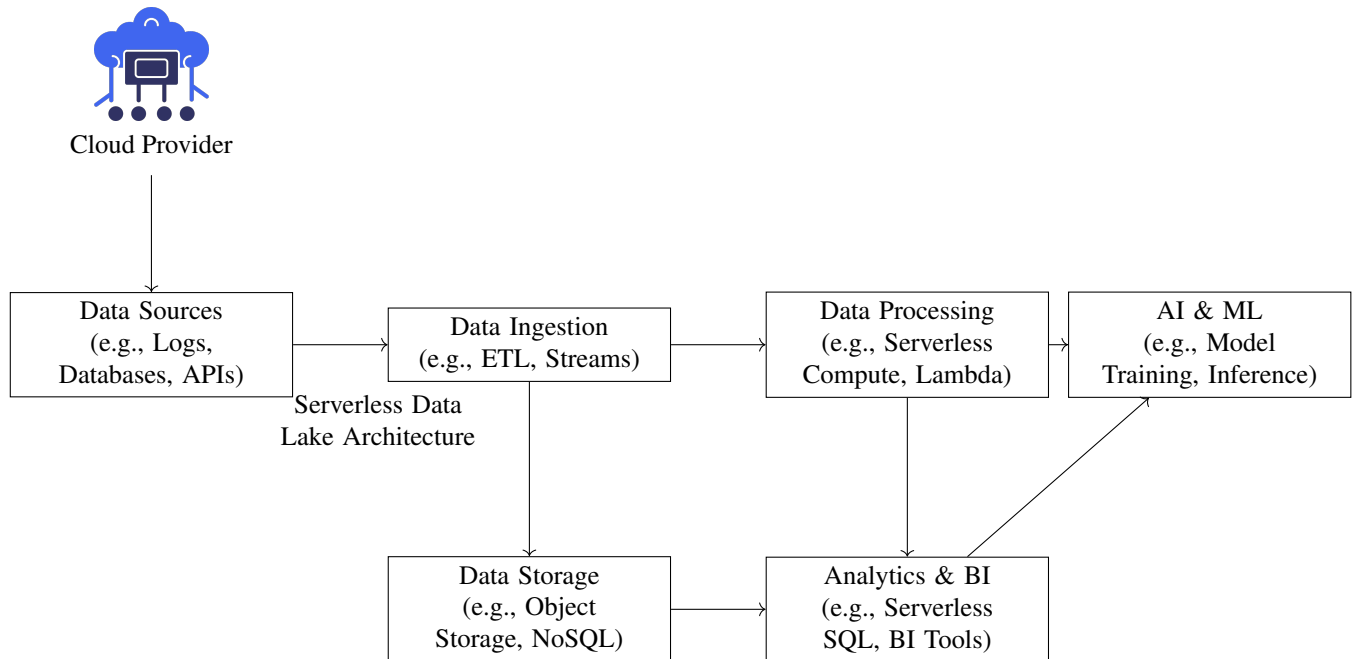


Figure 3. Architecture of a Serverless Data Lake

Aspect	Application
Enhanced Security	Sensitive data can be stored on-premises, ensuring that institutions maintain control over their most critical information while complying with regulatory requirements.
Scalable Solutions	Non-sensitive data and less critical applications can be hosted in the cloud, providing institutions with the ability to scale their operations efficiently as their data needs grow.
Cost Optimization	By using a hybrid approach, institutions can optimize their IT spending, utilizing on-premises resources for critical workloads while leveraging the cost efficiencies of cloud-based services for other tasks.

Table 5. Benefits of Hybrid Cloud Environments in the Financial Sector

Hybrid cloud environments offer a flexible and strategic solution for managing financial data, enabling institutions to balance the need for stringent data security with the demand for scalable and cost-effective IT resources. Financial institutions can create a hybrid architecture that supports their operational goals while maintaining compliance with regulatory standards. As the financial sector continues to evolve, the adoption of hybrid cloud environments is likely to expand, providing a robust framework for managing the increasingly complex demands of modern financial data processing and storage.

C. DATA MESH ARCHITECTURES

Data mesh architecture is an emerging approach to data management that seeks to overcome the limitations of centralized data architectures by decentralizing data ownership and aligning it with specific business domains. This approach is well-suited to the financial sector, where different departments, such as trading, risk management, and customer service, have distinct data requirements and operational objectives.

Decentralized data management is a core principle of data mesh architecture. In this model, each business domain within a financial institution is responsible for managing its own data. This decentralized approach allows each domain to

Aspect	Application
Decentralized Data Management	Each business domain, such as trading or customer service, manages its own data, allowing for more responsive and specialized data handling.
Scalability	Data mesh architectures are scalable by design, as each domain can independently manage its data without creating bottlenecks associated with centralized data management.
Improved Governance	By aligning data ownership with business domains, data mesh architectures enhance data governance, ensuring that each domain is accountable for the quality, security, and compliance of its data.

Table 6. Benefits of Data Mesh Architectures in the Financial Sector

handle its data in a way that best supports its specific needs and goals. For instance, the trading department may prioritize real-time data processing capabilities to respond quickly to market changes, while the customer service department may focus on ensuring that customer data is accurate and up-to-date to improve service quality.

Scalability is inherently built into data mesh architectures, as each domain operates independently, managing its own data and scaling its resources as needed. This independence eliminates the bottlenecks often associated with centralized data management systems, where data flows through a single, centralized pipeline. In a data mesh, as data volumes grow or as new data sources are integrated, each domain can scale its data processing infrastructure without affecting other parts of the organization. This decentralized scalability is beneficial in the financial sector, where different departments may have varying data throughput and storage requirements, depending on their operational focus and the volatility of the market.

Improved governance is another significant benefit of data mesh architecture. Aligning data ownership with specific business domains, data mesh ensures that those who are most familiar with the data are responsible for its quality, security, and compliance. This alignment enhances data governance by making data management more accountable and transparent. Each domain is tasked with ensuring that its data meets regulatory standards and internal policies, reducing the risk of compliance issues and improving overall data integrity. This domain-specific governance model also allows for more granular control over data access, ensuring that sensitive information is protected according to the unique requirements of each business area.

Data mesh architecture offers financial institutions a way to align their data management practices with their organizational structure, improving both scalability and data governance. Data mesh enables each business domain to operate more autonomously and efficiently, supporting the diverse analytical and operational needs of the institution. As financial institutions continue to seek ways to manage growing data volumes and complexity, data mesh architecture is likely to become an increasingly important tool in achieving flexible, scalable, and well-governed data management practices.

D. DATA LAKEHOUSE ARCHITECTURE

Data lakehouse architecture represents an innovative evolution in data management, combining the strengths of both data lakes and data warehouses. This hybrid approach ad-

dresses the limitations inherent in each system by creating a unified environment capable of handling diverse data types and analytical workloads, making it advantageous for financial institutions that deal with complex and varied datasets.

The architecture of a data lakehouse enables unified data management by integrating the flexible storage capabilities of data lakes with the structured, query-optimized environment of data warehouses. Data lakes traditionally excel at storing large volumes of raw, unstructured data, such as text files, logs, and multimedia content, while data warehouses are optimized for structured data that is organized into relational tables, supporting complex queries and business intelligence operations. In a data lakehouse, these two paradigms converge, allowing for the storage, processing, and querying of both structured and unstructured data in a single platform. This unification simplifies data management by eliminating the need for separate systems to handle different data types, thereby reducing the complexity and overhead associated with maintaining multiple data storage solutions.

The versatility of data lakehouse architecture extends to its support for a wide range of analytical processes. Blending the flexible, schema-on-read capabilities of data lakes with the robust, schema-on-write model of data warehouses, the lakehouse supports everything from traditional business intelligence (BI) and reporting to more advanced analytics, including machine learning and real-time data processing. For example, a financial institution can use the same data lakehouse platform to run SQL queries for generating financial reports, execute complex machine learning algorithms to predict market trends, and process streaming data for real-time fraud detection. This versatility not only enhances the analytical capabilities of financial institutions but also ensures that they can leverage a broader spectrum of data to drive decision-making.

From a cost efficiency standpoint, the data lakehouse model offers significant advantages. Integrating the storage and processing capabilities of data lakes and data warehouses into a single architecture, the need for data duplication is greatly reduced. Traditionally, data needed to be copied from a data lake to a data warehouse for analysis, leading to increased storage costs and potential issues with data consistency. In a lakehouse, the data remains in one place, accessible for both batch processing and real-time analytics, thus streamlining operations and reducing costs associated with data movement and storage. Moreover, the

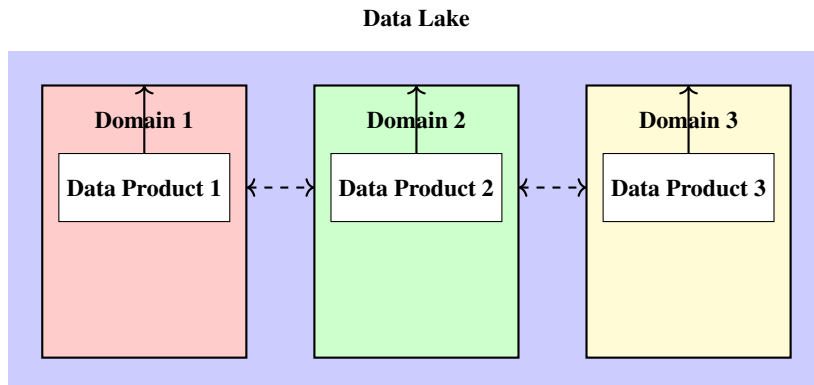


Figure 4. Data Mesh Architecture in Data Lakes. The architecture is composed of multiple domains within a data lake, each containing data products that are interconnected.

Aspect	Application	Example
Unified Data Management	The data lakehouse architecture enables institutions to store and process all types of data in a single platform, simplifying data management and reducing the need for multiple storage solutions.	A financial institution consolidates its customer transaction records (structured data) and social media sentiment analysis (unstructured data) into a single data lakehouse platform, making it easier to perform comprehensive customer analysis.
Versatile Analytics	By merging data lakes' flexibility with data warehouses' structure, the lakehouse model supports a wide range of analytical activities, from traditional business intelligence to advanced machine learning.	A bank uses its data lakehouse to perform both traditional reporting on quarterly performance metrics and predictive analytics to forecast future loan default risks using machine learning models.
Cost Efficiency	Integrating the two architectures reduces data duplication and streamlines operations, leading to lower costs for data storage and processing.	A financial services company saves costs by reducing the need for separate data warehouses and lakes, instead leveraging a lakehouse architecture that handles both structured transaction data and unstructured customer feedback, reducing overall infrastructure expenses.

Table 7. Applications of Data Lakehouse Architecture in the Financial Sector

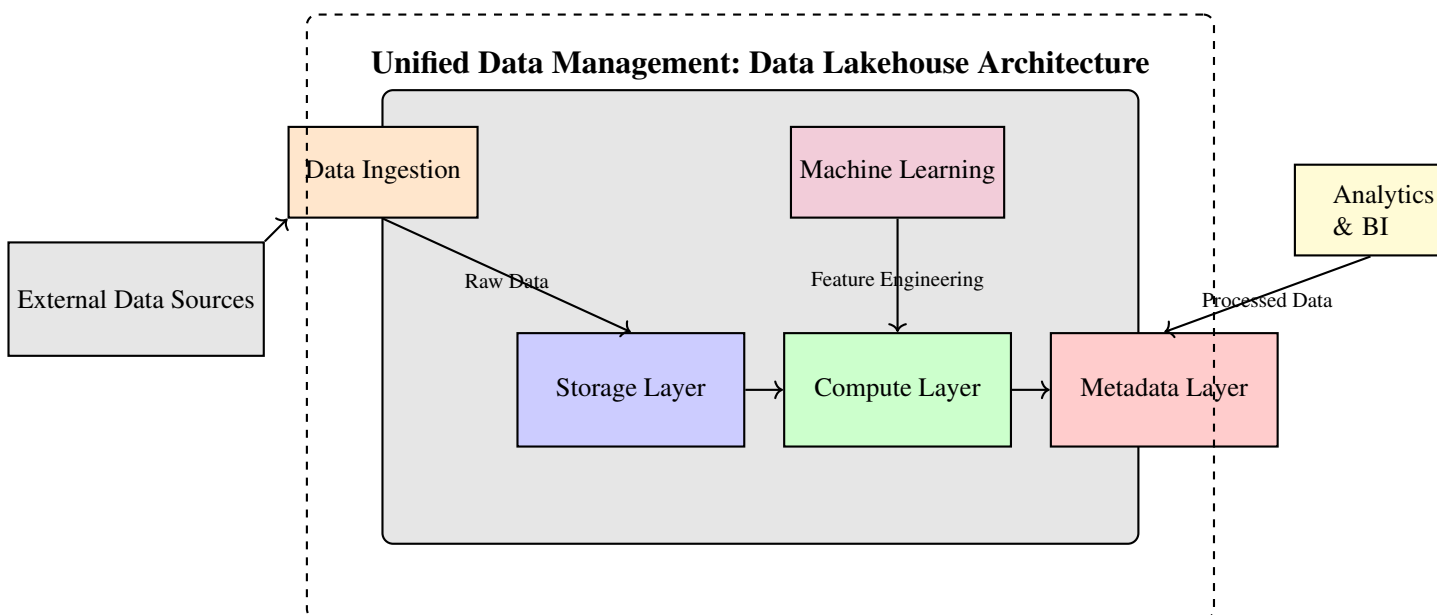


Figure 5. Unified Data Management: The data lakehouse architecture enables institutions to store and process all types of data in a single platform, simplifying data management and reducing the need for multiple storage solutions.

lakehouse architecture often leverages scalable cloud-based infrastructure, which allows financial institutions to manage large datasets more cost-effectively, scaling resources up or down based on demand.

The adoption of data lakehouse architecture provides financial institutions with a powerful platform capable of handling the complexities of modern data analytics. Supporting both traditional and cutting-edge analytical workloads, lakehouses enable institutions to extract more value from their data assets while maintaining operational efficiency. As data continues to grow in volume and complexity, the lakehouse model is likely to become an increasingly important tool for financial institutions, enabling them to navigate the challenges of data management and analytics with greater agility and effectiveness.

IV. CONCLUSION

The financial sector is increasingly dependent on advanced data management strategies to handle the growing volume and complexity of data. Data lakes, supported by emerging technologies and infrastructure innovations, offer a versatile and scalable solution for managing this data. The integration of AI, machine learning, and real-time processing capabilities into data lakes enhances their ability to provide valuable insights and support critical financial operations.

The study thoroughly examines the emerging trends and technological advancements in data lakes within the financial sector, providing a detailed analysis of how these innovations are reshaping data management, processing, and analytics. Key among these trends is the integration of artificial intelligence (AI) and machine learning (ML), which are increasingly critical for enhancing predictive modeling, risk assessment, and process automation. These technologies allow financial institutions to analyze large datasets with greater accuracy, enabling more informed decision-making and efficient operations. The ability to predict future market trends and customer behavior through AI-driven models is crucial, as it supports strategic planning and risk management. Additionally, AI and ML facilitate the automation of data processes within data lakes (Simon, 2021), streamlining workflows and reducing manual effort.

Real-time data processing is another significant development that the study highlights as essential for the financial sector. In an industry where timely decision-making can be the difference between profit and loss, the ability to process and analyze data as it is generated is crucial. The integration of real-time data streaming and analytics in data lakes allows financial institutions to respond quickly to market changes, customer behaviors, and emerging risks. This capability not only improves market responsiveness but also enhances the ability to detect and prevent fraudulent activities by analyzing transaction data in real-time. The study emphasizes that real-time data processing is becoming a standard requirement for financial institutions seeking to maintain a competitive edge in a fast-paced market.

The paper also explores the evolution of data lakehouse

architecture, which combines the benefits of data lakes and data warehouses into a unified platform. This hybrid approach allows financial institutions to manage both structured and unstructured data in a single environment, supporting a broad range of analytical activities. The data lakehouse model simplifies data management by reducing the need for multiple storage solutions and eliminating data silos. This versatility enables institutions to handle diverse data types and perform various analytics tasks, from traditional business intelligence to complex machine learning. The study points out that the lakehouse architecture not only enhances analytical capabilities but also reduces operational costs by streamlining data storage and processing.

In addition to these advancements, the integration of edge computing with data lakes is identified as a growing trend that offers significant benefits for the financial sector. Edge computing involves processing data closer to its source, which reduces latency and improves the speed of real-time analytics. This is important for applications such as real-time trading and fraud detection, where quick analysis and response are critical. Edge computing also enhances data privacy by minimizing the need to transmit sensitive information across networks, thereby reducing the risk of data breaches. The study highlights that the integration of edge computing with data lakes can lead to improved customer experiences in areas such as mobile banking and automated teller machines (ATMs), where faster data processing translates to quicker service and greater customer satisfaction.

On the infrastructure side, the paper discusses the adoption of serverless data lakes as a significant innovation in data management for the financial sector. Serverless architecture eliminates the need for managing underlying infrastructure, allowing IT teams to focus on more strategic tasks such as data analysis and system optimization. The serverless model is cost-effective, operating on a pay-as-you-go basis that ensures institutions only pay for the computing resources they actually use. This flexibility is especially beneficial in an industry where data processing demands can fluctuate significantly. Serverless data lakes also offer scalability, automatically adjusting to accommodate changing data volumes and processing requirements. The study suggests that serverless data lakes provide financial institutions with a more efficient and cost-effective approach to data management, reducing operational overhead while maintaining the ability to scale as needed.

Hybrid cloud environments are another area of focus, offering a solution that balances the need for data security with the scalability and cost advantages of cloud computing. The study notes that hybrid cloud environments allow financial institutions to store sensitive data on-premises, ensuring compliance with regulatory requirements and maintaining control over critical information. At the same time, less sensitive data and applications can be hosted in the cloud, providing the flexibility to scale operations as data needs grow. This approach optimizes IT spending by leveraging on-premises resources for high-priority tasks while utilizing the

cost efficiencies of cloud-based services for other workloads. The study concludes that hybrid cloud environments provide a balanced and flexible solution for managing financial data, addressing the sector's dual need for security and scalability.

Finally, the study explores the concept of data mesh architecture, a relatively new approach to data management that decentralizes data ownership and aligns it with specific business domains. This architecture is well-suited where different departments often have unique data requirements. Data mesh allows each business domain, such as trading or customer service, to manage its own data independently, leading to more responsive and specialized data handling. This decentralization also improves scalability, as each domain can manage its data without the bottlenecks associated with centralized data management systems. Additionally, data mesh enhances data governance by aligning data ownership with specific business areas, ensuring that each domain is responsible for the quality, security, and compliance of its data. The study suggests that data mesh architecture could be a transformative approach for financial institutions, offering a way to align data management practices with organizational structures and improve both scalability and governance.

VECTORAL PUBLICATION PRINCIPLES

Authors should consider the following points:

- 1) To be considered for publication, technical papers must contribute to the advancement of knowledge in their field and acknowledge relevant existing research.
- 2) The length of a submitted paper should be proportionate to the significance or complexity of the research. For instance, a straightforward extension of previously published work may not warrant publication or could be adequately presented in a concise format.
- 3) Authors must demonstrate the scientific and technical value of their work to both peer reviewers and editors. The burden of proof is higher when presenting extraordinary or unexpected findings.
- 4) To facilitate scientific progress through replication, papers submitted for publication must provide sufficient information to enable readers to conduct similar experiments or calculations and reproduce the reported results. While not every detail needs to be disclosed, a paper must contain new, usable, and thoroughly described information.
- 5) Papers that discuss ongoing research or announce the most recent technical achievements may be suitable for presentation at a professional conference but may not be appropriate for publication.

References

Abbasi, A. (2020). *Aws certified data analytics study guide: Specialty (das-c01) exam*. John Wiley & Sons.

Arora, G., Lele, C., & Jindal, M. (2022). *Data analytics: Principles, tools, and practices: A complete guide for advanced data analytics using the latest trends,*

tools, and technologies (english edition). BPB Publications.

- Boden, M. A. (1996). *Artificial intelligence*. Elsevier.
- Charniak, E. (1985). *Introduction to artificial intelligence*. Pearson Education India.
- Chessell, M., Scheepers, F., Strelchuk, M., van der Starre, R., Dobrin, S., Hernandez, D., et al. (2018). *The journey continues: From data lake to data-driven organization*. IBM Redbooks.
- Ertel, W. (2018). *Introduction to artificial intelligence*. Springer.
- Gorelik, A. (2019). *The enterprise big data lake: Delivering the promise of big data and data science*. O'Reilly Media.
- John, T., & Misra, P. (2017). *Data lake for enterprises*. Packt Publishing Ltd.
- Pandey, D., & Tripathi, S. (2016). Data-lake: Requirement to deployment. *Advances in Computing, Control and Communication Technology, 1*, 200.
- Pejić Bach, M., Krstić, Ž., Seljan, S., & Turulja, L. (2019). Text mining for big data analysis in financial sector: A literature review. *Sustainability, 11*(5), 1277.
- Simon, A. R. (2021). *Data lakes for dummies*. John Wiley & Sons.
- Strengholt, P. (2023). *Data management at scale*. O'Reilly Media, Inc.
- Tanca, L. (2023). Enabling real-world medicine with data lake federation: A research perspective. *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB Workshops, Poly 2022 and DMAH 2022, Virtual Event, September 9, 2022, Revised Selected Papers, 13814*, 39.
- Trelewicz, J. Q. (2017). Big data and big money: The role of data in the financial sector. *IT professional, 19*(3), 8–10.
- Vergadia, P. (2022). *Visualizing google cloud: 101 illustrated references for cloud engineers and architects*. John Wiley & Sons.
- Zburivsky, D., & Partner, L. (2021). *Designing cloud data platforms*. Simon; Schuster.
- Ziegler, H., Jenny, M., Gruse, T., & Keim, D. A. (2010). Visual market sector analysis for financial time series data. *2010 IEEE Symposium on Visual Analytics Science and Technology*, 83–90.

...