



Navigating the Data Landscape: Evaluating Microsoft SQL Server and Presto for Enterprise Needs

Paraskumar Patel

Pittsburgh, PA, USA

Abstract In the realm of data management and analysis, the selection of appropriate technologies is pivotal for businesses aiming to harness the power of data for actionable insights and competitive advantage. This paper presents a comprehensive comparative analysis of two prominent technologies in the field: Microsoft SQL Server and Presto. Microsoft SQL Server, a robust relational database management system (RDBMS), is renowned for its enterprise-grade features, including advanced security measures, sophisticated performance optimization, and seamless integration within the Microsoft ecosystem. In contrast, Presto, an open-source distributed SQL query engine, offers unparalleled flexibility and scalability, excelling in fast, ad-hoc analytics across diverse and distributed data sources. Through an exploration of architectural differences, performance benchmarks, scalability, ecosystem integration, and cost implications, this study elucidates the distinct advantages and limitations of each technology. The findings reveal that while SQL Server is ideally suited for organizations seeking a comprehensive, integrated data management solution within a Microsoft-centric environment, Presto is better aligned with the needs of enterprises prioritizing agile analytics across varied data platforms. This analysis aims to aid decision-makers in aligning their technology choices with organizational needs and strategic objectives, ultimately enabling more informed and effective data management and analysis strategies in the context of modern, data-driven business environments.

Keywords Presto, SQL Server, Architectural Differences, Data Management

1. Introduction

In today's data-driven environment, the ability to efficiently manage, process, and analyze data is crucial for businesses to derive actionable insights and maintain a competitive edge. This has led to the development and adoption of advanced database management systems (DBMS) and data processing engines that cater to various data handling needs. Among these technologies, Microsoft SQL Server and Presto stand out as leading solutions, albeit serving different niches within the broader spectrum of data management and analysis [1].

Microsoft SQL Server, a product of decades of development by Microsoft, is a relational database management system (RDBMS) renowned for its robustness, comprehensive features, and integration with Windows environments. It supports a wide array of transaction processing, business intelligence, and analytics applications in corporate IT environments. SQL Server is designed to manage structured data, with an emphasis on ensuring ACID (Atomicity, Consistency, Isolation, Durability) properties for transaction processing and offering sophisticated tools for query optimization, security, and database administration [1], [2].

On the other hand, Presto, an open-source distributed SQL query engine, is known for its ability to perform interactive analytics on large datasets across various data sources, including traditional RDBMSs, NoSQL databases, and data lakes. Developed initially by Facebook to address the needs of analyzing enormous volumes of data quickly, Presto is designed for high-speed, ad-hoc queries on big data and has been widely adopted for its flexibility, scalability, and the ability to query data where it lives, without requiring data movement.



This paper aims to conduct a comprehensive comparative analysis of Microsoft SQL Server and Presto, highlighting their architectural differences, use cases, performance characteristics, and scalability. By examining these aspects, the paper seeks to provide insights into their suitability for different organizational needs, thus aiding decision-makers in selecting the appropriate technology for their data management and analysis requirements. Through this analysis, we will explore how each technology addresses the challenges of modern data environments and facilitates data-driven decision-making processes.

2. Background

The evolution of data management and analysis technologies has been profoundly influenced by the ever-increasing scale and complexity of data. Organizations today are not only dealing with larger volumes of data but also grappling with the diversity of data types and sources. This section provides a background overview of Microsoft SQL Server and Presto, setting the stage for their comparative analysis.

A. Microsoft SQL Server

Developed and marketed by Microsoft, SQL Server has established itself as a cornerstone in the landscape of relational database management systems [3]. Its inception in the late 1980s, in collaboration with Sybase, marked the beginning of its journey. By the early 1990s, Microsoft took sole ownership and direction of SQL Server's development, significantly evolving its capabilities with each version.

SQL Server's design is anchored in supporting structured query language (SQL), offering a rich set of features for data storage, retrieval, and manipulation. It caters to a wide range of data management tasks, including transaction processing, data warehousing, and business intelligence. Over the years, SQL Server has integrated various tools and services, such as SQL Server Reporting Services (SSRS) and SQL Server Integration Services (SSIS), enhancing its utility in data analytics and integration scenarios.

Key to SQL Server's adoption has been its tight integration with the Windows ecosystem, offering enterprises a cohesive environment for developing and deploying applications. Its support for various data types, comprehensive data management tools, advanced security features, and scalability options has made it a preferred choice for organizations around the globe.

B. Presto

Presto represents a paradigm shift in the field of data processing, emphasizing flexibility and speed in handling big data. Originating at Facebook in 2012, Presto was conceived to address the limitations of existing data analysis tools in coping with the vast and growing datasets within the company. It was open-sourced in 2013, allowing it to benefit from contributions from a broad community of developers and organizations.

As a distributed SQL query engine, Presto is designed to query large datasets distributed across multiple sources swiftly. Unlike traditional RDBMS that require data to be stored within their own systems, Presto can query data where it resides, whether in Hadoop data lakes, traditional databases, or NoSQL systems. This ability to federate queries across different data stores eliminates the need for costly data movement and transformation processes, significantly speeding up data analysis tasks.

Presto's architecture allows it to scale horizontally, adding more nodes to increase query processing power. This scalability, combined with its in-memory processing capabilities, enables Presto to deliver fast query responses, even on petabyte-scale datasets. Its use of ANSI SQL allows data analysts to use familiar syntax and concepts, easing the learning curve and facilitating its adoption.

In comparing Microsoft SQL Server and Presto, this paper will delve into various dimensions, including architectural differences, performance benchmarks, scalability, ecosystem integration, and use case applicability. Understanding the background and evolution of these technologies provides a foundation for appreciating their design principles, strengths, and limitations.

3. Architecture and Design

A. Microsoft SQL Server Architecture

At the core of Microsoft SQL Server's robust data management capabilities lies its sophisticated architecture, designed to facilitate efficient data storage, retrieval, and security. Central to this architecture is the Database



Engine, which orchestrates the processing of queries, transactions, and the management of data stored on disk [4]. This Engine comprises two pivotal components: the Relational Engine, tasked with query parsing and execution plan generation, and the Storage Engine, which manages the physical database files. SQL Server operates on an underlying SQL OS, a specialized subsystem that handles task scheduling, memory management, and input/output operations, ensuring that resources are optimally utilized. Communication between SQL Server and its clients is managed through a versatile Protocols Layer, supporting multiple network protocols to ensure secure and reliable data transmission. Additionally, SQL Server Management Studio (SSMS) offers a comprehensive environment for database administration, providing tools for configuration, monitoring, and managing database infrastructure with ease.

B. Presto Architecture

Presto stands out for its ability to execute fast, distributed SQL queries across large datasets and various data sources, owing to its distinct architecture. The Coordinator Node plays a pivotal role as the management center, responsible for query parsing, planning, and the orchestration of query execution across the worker nodes. These Worker Nodes are tasked with the actual query processing, operating in parallel to ensure efficient data handling and computation. Presto's flexibility and power are further enhanced by its Connector Architecture, which enables it to interface with a myriad of data sources, including traditional RDBMSs, NoSQL databases, and even data lakes, without necessitating data movement. This architecture supports Presto's pipeline processing model, wherein the Coordinator devises a distributed execution plan that divides the query into multiple stages, each executed concurrently by the Worker Nodes. This model allows for the streaming of data between stages, facilitating the processing of extensive datasets with remarkable efficiency [5].

While Microsoft SQL Server's architecture is optimized for transactional integrity, security, and manageability within a centralized data management system, Presto's design prioritizes speed and flexibility for analytics across distributed data sources. SQL Server's architecture supports a wide range of data management and application development scenarios, offering a unified solution for enterprises. In contrast, Presto's architecture, with its distributed execution model and connector-based approach to data sources, is tailored for scenarios requiring rapid, ad-hoc analysis of large volumes of data from diverse sources. Thus, each system's architectural design reflects its target use cases, offering distinct advantages depending on an organization's specific data management and analysis needs.

4. Performance and Stability

A. Microsoft SQL Server

When it comes to enterprise-level applications and workloads, Microsoft SQL Server stands out with its robust performance and scalability capabilities. Central to SQL Server's performance optimization are its intelligent query processing, cutting-edge indexing strategies, and pioneering in-memory computing technologies. The Query Optimizer at the heart of SQL Server employs advanced optimization techniques to refine query performance, meticulously analyzing queries to generate the most efficient execution plans. This is done by leveraging detailed statistics and indexes to minimize data scans and enhance data retrieval speeds. SQL Server's approach to query processing is dynamic, with Adaptive Query Processing enabling real-time adjustment of execution strategies to ensure peak performance.

The indexing capabilities of SQL Server are instrumental in boosting data retrieval speeds. SQL Server supports a diverse range of indexes, such as clustered, non-clustered, column store, and spatial indexes, each tailored for specific querying and storage needs [4]. Column store indexes are particularly noteworthy for their role in facilitating high-performance analytics, allowing for real-time operational analytics directly on transactional data [4].

SQL Server also excels in In-Memory OLTP (Online Transaction Processing), which markedly enhances transaction processing speeds and minimizes latency by leveraging in-memory storage rather than traditional disk-based storage. This is complemented by In-Memory Column store Indexes, which enable efficient querying of large datasets, thus supporting real-time analytics on live transactional data [6]. The integration of specialized engines like Apollo column store and Hekaton in-memory engine has been key to supporting hybrid workloads and real-time analytical processing [7].



Scalability in SQL Server is achieved through sophisticated features like partitioning, which simplifies the management of large databases by dividing them into smaller, more manageable segments, and Always On Availability Groups, which offer robust solutions for high availability and disaster recovery. SQL Server can scale both vertically, by augmenting a single server with additional resources, and horizontally, through strategies like data partitioning and federation.

B. Presto

Presto is engineered from the outset for high-velocity, distributed SQL querying across vast datasets. It stands out for its capability to perform queries directly on data in its original location, thereby minimizing unnecessary data movement, and its reliance on in-memory processing for speedy analytics [5].

The distributed query execution architecture of Presto enables it to parallelize queries over numerous servers, breaking down substantial queries into smaller tasks for concurrent execution. This distributed execution framework is crucial for Presto's proficiency in managing petabyte-scale datasets with relative ease [5].

Presto's performance is further enhanced by its utilization of in-memory data processing, which ensures low-latency responses to queries. By maintaining intermediate query results in memory and streaming data between different stages of the query plan, Presto significantly reduces disk I/O, leading to faster query execution times. The connector architecture of Presto is another key feature, enabling it to directly query data from a myriad of sources, including HDFS, S3, relational databases, and NoSQL stores. This capability allows Presto to conduct analytics across diverse data sets without necessitating ETL processes, thereby improving performance by accessing data in its native format.

Presto's scalability is underpinned by its stateless worker nodes, which can be scaled out horizontally to augment query processing capabilities. As the volume of data and complexity of queries increase, additional worker nodes can be seamlessly integrated into the Presto cluster to distribute the workload more evenly, ensuring sustained high performance.

While Microsoft SQL Server and Presto both provide substantial performance and scalability features, their respective approaches mirror their foundational architectures and intended applications. SQL Server is tailored for transactional workloads and intricate querying within a cohesive data management system, offering deep integration with Windows-based environments. Its scalability mechanisms are designed to support both high availability and expansive data warehousing solutions.

Conversely, Presto is distinguished in the analytics sphere, especially for ad-hoc querying across large, distributed data sets. It is optimized for rapidity and efficiency in the concurrent reading and processing of data from various sources. Presto's scalability is inherently adaptive, facilitating dynamic resizing of clusters to meet the demands of fluctuating analytics workloads without detriment to query performance.

In essence, SQL Server is fine-tuned for ensuring transactional integrity and complex query processing within a centralized data ecosystem, whereas Presto is crafted for swift, scalable analytics across decentralized, extensive data architectures.

5. Ecosystem and Integration

A. Microsoft SQL Server

The Microsoft SQL Server ecosystem is a comprehensive and interconnected framework, showcasing its long history in the data management sector and its seamless integration with the Windows operating system alongside Microsoft's extensive collection of enterprise applications. At the heart of this ecosystem lies SQL Server Integration Services (SSIS), a robust platform designed for constructing high-level data integration and transformation solutions. SSIS plays a crucial role in data warehousing, facilitating the extraction, transformation, and loading (ETL) of data from a multitude of sources into SQL Server, thereby enabling efficient data management and analysis. Complementing SSIS are SQL Server Reporting Services (SSRS) and SQL Server Analysis Services (SSAS), which are indispensable for generating reports and analyzing data. These services empower businesses to make informed decisions based on comprehensive data insights, offering advanced interactive reporting, dashboards, and data mining capabilities.

Expanding its horizons into the cloud, SQL Server embraces Azure SQL Database, a cloud-centric adaptation that brings forth scalability, high availability, and enhanced security. This integration not only propels SQL



Server into the realm of cloud-native applications but also ensures its relevance in hybrid and cloud environments by fostering seamless connections with other Azure services. Moreover, while Microsoft Power BI is not exclusive to SQL Server, its profound integration with the server amplifies the analytical and reporting functionalities, providing users with powerful business intelligence tools and interactive visualizations.

B. Presto

In contrast to the tightly-knit structure of the SQL Server ecosystem, the Presto ecosystem is celebrated for its open and versatile nature, designed to thrive in a heterogeneous data landscape. The foundation of Presto's flexibility is its array of connectors, enabling queries across diverse data sources such as Hadoop HDFS, Amazon S3, MySQL, PostgreSQL, Cassandra, and Kafka. This versatility positions Presto as a central figure in varied data architectures, fostering an environment of comprehensive data accessibility and analysis. The governance of Presto under the Linux Foundation through the Presto Foundation underscores the project's commitment to community growth and development, ensuring its stature as a premier solution for rapid and scalable data analytics [5].

Presto's compatibility with leading data visualization and business intelligence tools, including Tableau, Looker, and Superset, enhances its utility by facilitating the creation of engaging, interactive dashboards and reports directly from Presto-queried data. Additionally, Presto's availability as a managed service on cloud platforms like AWS (Amazon Athena) and Google Cloud extends its reach, offering scalable and serverless data querying capabilities that simplify leveraging Presto in cloud-based environments [5].

When contrasting the ecosystems of Microsoft SQL Server and Presto, it's evident that each caters to distinct aspects of data management and analysis. SQL Server's ecosystem is deeply integrated with Microsoft's suite of products and services, providing a unified platform for database management, business intelligence, and analytics. This integration is particularly advantageous for organizations that are heavily invested in Microsoft technologies, facilitating the development and management of data-centric applications within Windows and Azure frameworks.

On the other hand, Presto's ecosystem prioritizes openness and interoperability, emphasizing its utility as a versatile query engine that connects with multiple data sources. Its extensive support for various connectors and BI tools ensures adaptability to any data architecture, whether it involves on-premise systems or cloud-native solutions. Thus, while SQL Server offers a holistic, integrated solution for data management and analytics, Presto emerges as a flexible and agile tool for querying and analyzing data across a wide array of distributed data stores, complementing existing data management frameworks.

6. Cost and Open Source Aspects

A. Microsoft SQL Server Cost Structure

Microsoft SQL Server operates on a commercial licensing model, with costs varying based on several factors including the edition chosen (e.g., Enterprise, Standard, Web, Express), the number of cores, and whether it's being run on-premises or in the cloud. The Enterprise edition, offering the full suite of features including advanced data management, business intelligence, and analytics capabilities, is at the higher end of the pricing spectrum, reflecting its comprehensive offerings tailored for large-scale, critical operations.

For smaller businesses or less critical applications, the Standard edition provides a more cost-effective option, albeit with fewer features. SQL Server also offers a free Express edition, designed for desktop, small server, and development purposes, which includes the core database features but limits database size and the number of cores.

Additionally, running SQL Server on Azure as a managed service (Azure SQL Database or Azure SQL Managed Instance) involves a pay-as-you-go pricing model, offering scalability and flexibility without upfront hardware investments, but with operational costs based on resource consumption.

B. Presto: Open-Source and Cost Implications

Presto is an open-source distributed SQL query engine, which means it is free to download and use. This open-source nature allows organizations to deploy and scale Presto across their data infrastructure without the licensing costs associated with commercial software. The main costs associated with Presto would come from



the computing resources (servers or cloud services) needed to run the software and the operational expenses related to managing those resources and the Presto environment [8].

For cloud-native organizations, managed Presto services offered by cloud providers (e.g., Amazon Athena) operate on a pay-per-query pricing model, where costs are incurred based on the amount of data scanned by queries. This model can offer cost savings for businesses with varying workloads, as they pay only for what they use.

The fundamental difference in the cost structure between Microsoft SQL Server and Presto stems from their commercial vs. open-source nature. SQL Server's commercial licensing can represent a significant investment, particularly for large-scale deployments and when opting for the more feature-rich editions. However, this cost also includes support and maintenance services provided by Microsoft, which can be critical for enterprises requiring reliability and accountability.

Presto's open-source model offers flexibility and potential cost savings, particularly appealing to businesses looking to leverage modern, distributed data architectures [8]. However, it's important to consider the total cost of ownership, including the need for skilled personnel to install, manage, and optimize the Presto environment, as well as the underlying hardware or cloud services costs.

In summary, the choice between SQL Server and Presto from a cost perspective will depend on the specific needs, scale, and existing infrastructure of an organization, as well as the value placed on commercial support and maintenance versus the flexibility and cost-efficiency of an open-source solution.

7. Advantages and Limitations

Microsoft SQL Server

Microsoft SQL Server emerges as a comprehensive solution, offering a robust suite of tools that cater to various aspects of data management, business intelligence, and analytics. This versatility positions SQL Server as a one-stop solution for organizations seeking to address a broad spectrum of data needs. It is particularly renowned for its enterprise-grade security features, including advanced data encryption, comprehensive auditing capabilities, and a suite of compliance tools, all designed to ensure that data is managed securely and efficiently. In terms of performance and scalability, SQL Server stands out with its adept performance tuning, sophisticated indexing strategies, and in-memory computing capabilities. These features collectively ensure optimal performance for both transaction processing and analytics applications, while also offering scalable solutions that accommodate organizational growth, whether through vertical or horizontal expansion. Furthermore, SQL Server's seamless integration with the broader Microsoft ecosystem—encompassing Azure, Power BI, and various enterprise applications—enhances productivity and provides a user experience that is both cohesive and efficient.

However, Microsoft SQL Server is not without its limitations. The licensing and operational costs associated with SQL Server can be prohibitive, especially for large-scale deployments or for enterprises that require access to its advanced features. This cost barrier can make SQL Server less accessible for small to medium-sized businesses. Additionally, while SQL Server supports Linux, its deep integration with Windows could limit its utility in environments that primarily utilize other operating systems, potentially restricting its applicability.

Presto

On the other hand, Presto is celebrated for its flexibility and open-source nature. Its flexible architecture enables users to query data from multiple sources in situ, facilitating analytics across disparate data stores without the necessity of moving data. Designed for distributed computing, Presto excels in scalability, allowing for the efficient horizontal addition of more nodes to manage increased query loads. The open-source model of Presto fosters a vibrant community that contributes to its continual improvement, driving innovation and the rapid introduction of new features and connectors. Additionally, Presto's cost-efficiency is notable, offering organizations a more economical option by eliminating licensing fees and allowing deployment on commodity hardware or in cloud environments [5].

Nevertheless, Presto faces its own set of challenges. The operational complexity associated with deploying and managing a Presto cluster, particularly at scale, demands expertise in distributed systems and may increase the operational overhead. Presto's focus on analytical queries means it lacks support for transactional operations or updates to data, necessitating the use of additional systems for transactional data management. Moreover, Presto's



in-memory processing, while efficient, can be resource-intensive for large queries, necessitating careful resource management and potentially leading to increased costs in cloud deployments.

When comparing Microsoft SQL Server and Presto, it's evident that each offers unique advantages tailored to specific use cases. SQL Server is particularly well-suited to environments requiring comprehensive data management capabilities, advanced security measures, and tight integration with Microsoft products. Its main drawbacks are related to cost and a preference for Windows environments. Presto, conversely, is ideal for organizations seeking flexibility, scalability, and cost-effectiveness, especially for fast analytics across diverse data sources. However, its operational complexity and limited transactional support may necessitate additional solutions for complete organizational data management.

In conclusion, the decision between Microsoft SQL Server and Presto should be based on an organization's specific needs, existing IT infrastructure, and strategic objectives. SQL Server is likely the better option for entities looking for a comprehensive, enterprise-grade data management solution within a Microsoft-centric environment. Presto, however, offers a dynamic and scalable approach for conducting agile, big data analytics across a variety of data platforms. Microsoft SQL Server emerges as a comprehensive solution, offering a robust suite of tools that cater to various aspects of data management, business intelligence, and analytics. This versatility positions SQL Server as a one-stop solution for organizations seeking to address a broad spectrum of data needs. It is particularly renowned for its enterprise-grade security features, including advanced data encryption, comprehensive auditing capabilities, and a suite of compliance tools, all designed to ensure that data is managed securely and efficiently. In terms of performance and scalability, SQL Server stands out with its adept performance tuning, sophisticated indexing strategies, and in-memory computing capabilities. These features collectively ensure optimal performance for both transaction processing and analytics applications, while also offering scalable solutions that accommodate organizational growth, whether through vertical or horizontal expansion. Furthermore, SQL Server's seamless integration with the broader Microsoft ecosystem—encompassing Azure, Power BI, and various enterprise applications—enhances productivity and provides a user experience that is both cohesive and efficient.

However, Microsoft SQL Server is not without its limitations. The licensing and operational costs associated with SQL Server can be prohibitive, especially for large-scale deployments or for enterprises that require access to its advanced features. This cost barrier can make SQL Server less accessible for small to medium-sized businesses. Additionally, while SQL Server supports Linux, its deep integration with Windows could limit its utility in environments that primarily utilize other operating systems, potentially restricting its applicability.

B. Presto

On the other hand, Presto is celebrated for its flexibility and open-source nature. Its flexible architecture enables users to query data from multiple sources in situ, facilitating analytics across disparate data stores without the necessity of moving data. Designed for distributed computing, Presto excels in scalability, allowing for the efficient horizontal addition of more nodes to manage increased query loads. The open-source model of Presto fosters a vibrant community that contributes to its continual improvement, driving innovation and the rapid introduction of new features and connectors. Additionally, Presto's cost-efficiency is notable, offering organizations a more economical option by eliminating licensing fees and allowing deployment on commodity hardware or in cloud environments [5].

Nevertheless, Presto faces its own set of challenges. The operational complexity associated with deploying and managing a Presto cluster, particularly at scale, demands expertise in distributed systems and may increase the operational overhead. Presto's focus on analytical queries means it lacks support for transactional operations or updates to data, necessitating the use of additional systems for transactional data management. Moreover, Presto's in-memory processing, while efficient, can be resource-intensive for large queries, necessitating careful resource management and potentially leading to increased costs in cloud deployments.

When comparing Microsoft SQL Server and Presto, it's evident that each offers unique advantages tailored to specific use cases. SQL Server is particularly well-suited to environments requiring comprehensive data management capabilities, advanced security measures, and tight integration with Microsoft products. Its main drawbacks are related to cost and a preference for Windows environments. Presto, conversely, is ideal for



organizations seeking flexibility, scalability, and cost-effectiveness, especially for fast analytics across diverse data sources. However, its operational complexity and limited transactional support may necessitate additional solutions for complete organizational data management.

In conclusion, the decision between Microsoft SQL Server and Presto should be based on an organization's specific needs, existing IT infrastructure, and strategic objectives. SQL Server is likely the better option for entities looking for a comprehensive, enterprise-grade data management solution within a Microsoft-centric environment. Presto, however, offers a dynamic and scalable approach for conducting agile, big data analytics across a variety of data platforms.

8. Conclusion

In this comprehensive comparative analysis, we have delved into the intricacies of Microsoft SQL Server and Presto, highlighting their architectural differences, performance characteristics, scalability, ecosystem integration, and cost considerations. Through meticulous examination, the paper elucidates the distinct advantages and limitations inherent to each technology, providing valuable insights for organizations navigating the complex landscape of modern data management and analysis.

Microsoft SQL Server emerges as a robust, enterprise-grade RDBMS, offering a comprehensive suite of features for transaction processing, business intelligence, and analytics. Its integration within the Microsoft ecosystem, advanced security measures, and sophisticated performance and scalability mechanisms position it as a compelling option for organizations deeply entrenched in Windows-based environments or those requiring a unified, feature-rich data management solution.

Conversely, Presto is celebrated for its exceptional flexibility, open-source nature, and remarkable ability to perform fast, distributed SQL queries across heterogeneous data sources. This makes Presto an ideal choice for organizations prioritizing agile, cost-effective analytics over a wide array of distributed data architectures. Despite its advantages, Presto's operational complexity and focus on analytical rather than transactional processing highlight the necessity for complementary systems to address a broader range of data management requirements.

The decision between Microsoft SQL Server and Presto should, therefore, be predicated on a thorough understanding of an organization's specific data management needs, existing infrastructure, and strategic goals. For entities seeking a comprehensive, integrated solution within a predominantly Microsoft-centric environment, SQL Server offers unmatched depth and breadth of functionality. In contrast, Presto stands out for enterprises aiming to leverage flexible, scalable analytics across diverse data platforms, especially when cost-efficiency and open-source adaptability are paramount.

In conclusion, as data continues to play a pivotal role in shaping business strategies and operational efficiencies, the choice of a data management and analysis solution becomes crucial. This paper has endeavored to provide a detailed comparison of Microsoft SQL Server and Presto, aiding decision-makers in selecting the technology that best aligns with their organizational needs and strategic objectives. In the rapidly evolving data landscape, understanding the strengths and limitations of these technologies will be key to leveraging data for competitive advantage, fostering innovation, and driving data-driven decision-making processes.

References

- [1]. W. N. Suliyanti, "Studi Literatur Basis Data SQL dan NoSQL," KILAT, vol. 8, no. 1, pp. 48–51, May 2019, doi: 10.33322/KILAT.V8I1.460.
- [2]. E. C. Foster and S. V. Godbole, "Overview of SQL," Database Systems, pp. 171–175, 2014, doi: 10.1007/978-1-4842-0877-9_10.
- [3]. S. Binani, A. Gutti, and S. Upadhyay, "SQL vs. NoSQL vs. NewSQL- A Comparative Study," Communications on Applied Electronics (CAE), vol. 6, no. 1, 2016, [Online]. Available: www.caeaccess.org
- [4]. P. A. Larson, E. N. Hanson, and M. Zwilling, "Evolving the architecture of SQL Server for modern hardware trends," Proc Int Conf Data Eng, vol. 2015-May, pp. 1239–1245, May 2015, doi: 10.1109/ICDE.2015.7113371.



- [5]. R. Sethi et al., “Presto: SQL on everything,” Proc Int Conf Data Eng, vol. 2019-April, pp. 1802–1813, Apr. 2019, doi: 10.1109/ICDE.2019.00196.
- [6]. A. T. Kabakus and R. Kara, “A performance evaluation of in-memory databases,” Journal of King Saud University - Computer and Information Sciences, vol. 29, no. 4, pp. 520–525, Oct. 2017, doi: 10.1016/J.JKSUCI.2016.06.007.
- [7]. P. Å. Larson, A. Birka, E. N. Hanson, W. Huang, M. Nowakiewicz, and V. Papadimos, “Real-time analytical processing with SQL server,” Proceedings of the VLDB Endowment, vol. 8, no. 12, pp. 1740–1751, Aug. 2015, doi: 10.14778/2824032.2824071.
- [8]. E. G., “Microsoft SQL Server vs Presto | What are the differences?” [Online]. Available: <https://stackshare.io/stackups/microsoft-sql-server-vs-presto>

