MEMSQL

RTInsights.com
Accelerate Your Business with Real-Time Insights

SPECIAL REPORT

BENEFITS OF REAL-TIME TRANSACTION PROCESSING IN BANKING AND FINANCIAL SERVICES

Financial services institutions (FSIs) are under increasing pressure to keep pace with a number of significant industry changes, including expectations of higher-touch and more personalized customer experiences, a growing number of privacy and data protection regulations, increased cyber threats, and a rise in fraud. This is all in addition to the constant presence of tight competition, and the still-ongoing need for digital transformation.

A common thread across all operational areas is the need for faster and more intelligent decision-making. Examples where such capabilities are required include:

- **Fraud monitoring** that is personalized to the individual customer's transactions, protecting the business and preventing customer disappointment
- **Dynamic portfolio analysis** and optimization of real-time trading decisions to manage risk and overcome the reliance on stale notions to optimize assets in dynamic markets, especially to meet long-term strategic asset allocations
- **Cyber threat protection** that identifies previously unknown attacks and mitigates against threats in real time as information from a broad array of sources is assimilated
- **Regulatory compliance awareness and enforcement,** where user access patterns that do not appear to meet regulatory requirements are identified to ensure data and privacy are not compromised

In all of these areas, there is considerable room to provide services that deliver a competitive advantage. An organization that can make use of data and insight at scale in real time (or near real-time) stands to differentiate itself from competitors, and offer superior products and customer experiences.

Unfortunately, current databases and associated data services often break down under the strain of real-time data ingestion, analytics, and transactions. What is needed to meet these new customer and business demands is instantaneous analysis leading to actionable insights and the ability to automate responses.

Different Approaches to Consider

One approach to enabling real-time capabilities such as these would be to simply add a system to perform transactional analysis separate from existing database systems and data warehouses. The existing systems could still be used to perform historical monitoring, while a transactional system could accelerate analysis and transactions.

Several problems result from such an approach, however. Having operational and analytical data in different places introduces complexity (and oftentimes with it, maintenance overhead and system fragility) and slows data processing significantly. Most analytics systems do not perform well with the operational demands of streaming data. Also, having operational and analytics data in two places can introduce data accuracy problems and make regulatory auditing more difficult and time-consuming.

What's needed to enable smarter, real-time decisions is a translytic data infrastructure, which is a hybrid approach that runs both transaction and analytical workloads at scale on an integrated SQL database.

What's needed to enable smarter, real-time decisions and to reduce the aforementioned problems is a translytic data infrastructure, which is a hybrid approach that runs both transaction and analytical workloads at scale on an integrated SQL database.

The reason translytic processing is needed is because today's systems do batch ingesting, transaction processing, ETL, data warehouse arrival, and query processing. Translytical changes this batch process and manual ETL to streaming ingest with built-in ETL that includes transaction processing (the part needed for analytics, not necessarily core customer transaction processing) and data ready for apps and analytics in real time or near real-time.

Benefits of Translytical Databases for Modern Workloads

The increasing demand to support operational workloads that incorporate real-time analysis—such as recommendations, targeting, and fraud—is leading to the adoption of hybrid transactional analytic database systems. Industry analysts have keyed in on this trend, dubbing it by monikers including hybrid transaction/analytical processing (HTAP), to hybrid operational and analytics processing (HOAP), and translytical (transactional and analytical) data platforms.

A scale-out, in-memory architecture capable of processing transactional data with a memory and disk-based engine for analytics are key technology enablers of such hybrid transactional-analytic databases. A hybrid approach delivers the performance needed to process and analyze streaming data. An added benefit is the simplicity of the architecture. Only one system needs to be maintained, with no data movement between separate systems.

Many benefits can be realized using translytic processing. They include:

Improved customer experience. FSIs strive to provide clients with personalized, high-quality experiences. They also seek to improve or complement the experience already delivered to the most profitable customers. This may manifest itself in the ability to intelligently respond to a customer interaction or planned action for improved portfolio management and real-time information for stock trading.

Portfolio managers that work with hedge funds, sovereign wealth funds, endowments, bond management institutions, or investment banks manage up to tens of thousands of positions. Their ability to meet increasing customer demands is dependent on increased use of analytical applications that can provide intelligent, personalized responses to live market events. Those applications need response times to drop to sub-second levels.

Analytics applications built on translytic processing allow customers to investigate investment alternatives with the latest market data seamlessly, without troublesome pauses. This allows them to make better investment decisions.

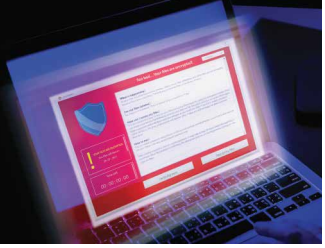

Translytic processing allows an organization to apply sophisticated threat detection queries on live data to identify threats in real time. This can mean the difference between incurring a major, publicly humiliating loss versus blocking fraudulent transactions before they happen.

Real-time fraud detection. Historically, fraud detection has been done using common batch processing runs at the end of the day, or based on simple historical models and rules. This process is inadequate to detect more sophisticated or fast-changing fraud events that occur throughout the day, or to adapt to the unique buying habits of individual customers.

Translytic processing allows an organization to apply sophisticated threat detection queries on live data to identify threats in real time. This can mean the difference between incurring a major, publicly humiliating loss versus blocking fraudulent transactions before they happen.

Increasingly, fraud detection is employing machine learning (ML) and artificial intelligence (AI) to spot anomalies and outliers. Most approaches require that any data analyzed must be run on a separate analysis system dedicated to ML or AI due to their compute- and data-intensive nature. Translytic processing allows the event data and the model scoring to remain in place and run instantly.

Strengthened risk management. Financial services organizations face ever-increasing regulations. Beyond Basel III (the international regulatory framework for banks), organizations have to comply with new and updated versions of other regulations. The EU General Data Protection Regulation (GDPR), Payment Services Directive 2 (PSD2), Solvency II, and Markets in Financial Instruments Directive II (MiFID II) all went into effect in 2018. Additional regulations are emerging all the time at the country level around the globe, as well as at the state level in the United States. For example, there is the California Consumer Privacy Act of 2018 (CCPA), a bill that enhances privacy rights and consumer protections for residents of the state of California.

To comply with these edicts, institutions must track risk measures for their own portfolios, trading departments, funds they offer, and individual client portfolios—all in real time. The right solution will offer a series of analytic capabilities that improve the service or reduce exposure by performing pre-trade analytics on real-time event data, or predicting when to optimize the rebalancing of a portfolio.

Creation of a "data culture." FSIs are under increasing pressure to act both quickly and correctly. They must take all available information into account and follow all applicable regulations in their markets. This can only be done with a widespread data culture. Currently, in most organizations, this is not the reality. FSI companies are limited by data availability, query speed, and concurrency restrictions. Translytic processing can help address all of these issues and can support a wide range of uses simultaneously.

What's Needed?

The right scalable translytic architecture makes all the needed data available in one place for line-of-business analytics, executive insights, reporting, compliance, governance, and operationalizing machine learning and artificial intelligence.

Data silos create barriers, complexity, and latency, and can create serious business problems. A translytic solution replaces silos to simplify data architectures and improve auditing and traceability of actions. In addition, making it easier to combine, analyze, and act on all of the data in your enterprise unlocks new capabilities. For instance, information used to identify a real-time cyber attack also might provide insights into the use of data by authorized individuals in ways that could expose a business to regulatory liabilities.

Translytic platforms for FSIs must offer enterprise-class components, tools, and features. Critical elements include:

- A core data platform that supports high availability, security, performance, scalability, and management ease
- Data integrity and consistency to deliver trusted data
- Native tools or integration with third-party vendors' tools to support data management functions
- Support for concurrent queries, transactions, analytics, data ingest, and inflow/outflows
- Deployment options for on-premises, cloud, or a hybrid approach
- Access to data using standard connectivity including SQL, ODBC/JDBC, XML, or REST

Such a translytic platform can support fast insights, machine learning, streaming analytics, and high-performance transactional processing. FSI organizations are no strangers to real-time applications. However, the differentiator with using a translytic database is that it significantly reduces data silos and duplication in addition to simplifying application design, coding, and customization, while improving throughput and responsiveness.

A translytic database can be used as a point solution within an existing database infrastructure. For instance, analytics responsiveness and concurrency requirements are often a pain point, within a division of a company, or across the organization as a whole. Using a translytic database alongside existing operational database or data warehousing solutions can improve responsiveness and concurrency while significantly reducing data duplication when offloading existing, often overloaded, database systems.

An ideal solution is to perform all workloads within a single database. This allows an organization to store and process data in a single, integrated platform, enabling synergies across operational areas. Additionally, the right platform can help deliver real-time insights for stock trading, fraud detection, and machine learning analysis. These applications require high availability to data 24x7, and access to the data must be very low latency. Only translytic platforms meet these requirements.

Moreover, translytic databases avoid errors and problems that occur in other analysis systems where data must be extracted, transformed, and loaded (ETL) before it is analyzed. Translytic databases overcome these problems by delivering a real-time, trusted view of critical business data. This ensures information is accurate and helps guarantee consistency across the organization.

Machine learning models that run on translytic databases do not require that data be moved to external machine learning platforms. This saves time, improves results, and cuts costs.

New types of analysis are made easier with translytic databases. For example, many FSIs employ sophisticated predictive and prescriptive analytics, complemented by machine learning. These efforts help companies use prior decisions and actions to enhance what was previously done to gain more informed predictions about customers, business processes, and operations. Machine learning models that run on translytic databases do not require that data be moved to external machine learning platforms. This saves time, improves results, and cuts costs.

Scalability Is the Critical Element

Another issue to consider when selecting a translytic database is the ability to run transactional and analytic workloads at the scale required with today's FSI datasets.

Legacy relational databases can't be scaled easily or cost effectively. The challenge is, while the database processes transactions, it can't be used to process queries at the same time. If the data is offloaded to a data warehouse, that process is batch-based, which means nightly (or at best, hourly) runs. A complex extract, transform, and load process must then be run to ready data for analytics; and analytics, even against this stale data, are often not available in real time. Attempts to meet the need for 24x7 data collection and monitoring can generate very large tables, or many smaller ones, either of which results in slow-performing queries.

A common trade-off is to break up the data into separate silos or batch the data into low-cost data lakes—which requires expert tuning and caching. This leads to query delays, inaccurate results that can ultimately cause customer attrition—if not regulatory problems—and increased costs and complexity for performance workarounds.

NoSQL databases overcome the scalability barrier, but are designed for unstructured or semi-structured data. They can't natively maintain consistency and proper join patterns for a structured data store, as required for efficient analytics. In addition, the specialized skills required to work with these tools restricts access from many line-of-business users. This makes NoSQL less than ideal to fully address the range of problems of legacy operational data for real-time or near real-time decision-making at scale.

The only way to deliver true real-time decision-making at organizational scale is to combine ingest, transactions, and analysis in a single, continually updated, fast database. This is where the newest generation of database technology can help.

Specifically, scalability is addressed using translytic databases that are fully distributed and can run as a single database across a pool of machines while maintaining consistency and stability. This allows an organization to easily scale out. This new generation of databases can speed up transaction processing, avoiding batching on ingest. They also can speed up analytics processing, making operational analytics at scale possible. And they can combine both processes and enable true, end-to-end, real-time decision-making.

Selecting the Right Tools for Modern FSI Applications

Translytic processing can be the underlying technology that supports new applications, including real-time dashboards for portfolio analysis, portfolio management analytics for reduced risk and optimization, and fraud prevention.

Financial markets move quickly. Data access and insights in real time gives banks, asset managers, analysts, and advisors a competitive advantage. To achieve that advantage, FSIs need the ability to ingest, query, and serve data to thousands of simultaneous users.

These capabilities allow firms to support their mission and their customers better. Translytic processing can be the underlying technology that supports new applications, including real-time dashboards for portfolio analysis, portfolio management analytics for reduced risk and optimization, and fraud prevention. These applications put sophisticated decision-making power in the hands of FSI staff and investors.

Selecting the right technology partner is critical. Whatever solution is chosen must easily scale, minimize latency, and work with existing data. It also must offer enterprise performance, management tools, and security.

These are all areas where MemSQL excels. MemSQL is a distributed, highly scalable, translytical SQL database that can run both in the cloud and on premise. It delivers maximum performance for transactional and analytical workloads with familiar relational data structures.

MemSQL ingests data continuously to perform operational analytics for the front lines of your business. Ingest millions of events per day with ACID transactions while simultaneously analyzing billions of rows of data in relational SQL, JSON, geospatial, and full-text search formats.

Organizations that use MemSQL get high ingestion performance at scale and eliminate the need for costly data integration tools through the use of built-in batch and real-time data pipelines. MemSQL lets organizations achieve ultra-fast query responses across both live and historical data using familiar ANSI SQL.

Its highly scalable, distributed system balances data and queries across a cluster of cloud instances or commodity hardware for maximum performance, concurrency, and availability. Infrastructure elements that enable these characteristics include:

- **Distributed storage** that allows for the processing of data across clusters of machines for maximum durability, resilience, and performance. Tiered, distributed storage helps an application efficiently leverage system memory.
- **Massively parallel architecture** that provides a robust, parallel execution engine for reading and writing queries, delivering ultra-fast performance.
- **Big data capacity** that allows for the efficient storage of petabytes of data on low-cost disk and cloud storage (for data science or archive requirements) while maintaining nearly instantaneous retrieval for fast deep analysis.

MemSQL meets or exceeds data security requirements mandated by the plethora of global and local regulations. This is done without compromising database performance. Administration tools allow for easy management of how users and roles access data to support workloads in complex organizations and regulated environments.

Combined, these capabilities allow MemSQL to deliver:

- Extreme performance
- Massive scalability
- An easy-to-use SQL architecture
- The ability to run natively on any cloud infrastructure or commodity hardware
- Reasonable price, with strong price-performance and total cost of ownership (TCO)

With such capabilities, FSIs can use translytic processing to improve operations and offer enhanced services to differentiate their offerings. In a highly competitive marketplace, this can help meet the demands of existing customers and incite others to become customers. This is why half of the top 10 banks in North America already use MemSQL to power some of their most important decisions and most critical applications.

Adopting any new technology in an FSI is not taken lightly. In many cases, the need for a solution like MemSQL arises when something breaks in an organization's architecture, or when a required service or demand cannot be met with existing infrastructure. In such cases, MemSQL is brought in to address that specific problem. Once the capabilities and nondisruptive nature of the solution are realized, its use is expanded to other areas.

To learn more about translytic processing, visit https://msql.co/finserv.

RTInsights.com
Accelerate Your Business with Real-Time Insights

RTInsights is an independent, expert-driven web resource for senior business and IT enterprise professionals in vertical industries. We help our readers understand how they can transform their businesses to higher-value outcomes and new business models with IoT real-time analytics. We provide clarity and direction amid the often confusing array of approaches and vendor solutions. We provide our partners with a unique combination of services and deep domain expertise to improve their product marketing, lead generation, and thought leadership activity.

MEMSQL

MemSQL is The No Limits Database™, powering modern applications and analytical systems with a cloud-native, massively scalable architecture for maximum ingest and query performance with the highest concurrency. MemSQL envisions a world where every business can make decisions in real time and every experience is optimized through data. Global enterprises use the MemSQL distributed database to easily ingest, process, analyze, and act on data in order to thrive in today's insight-driven economy. MemSQL is optimized to run on any public cloud or on-premise with commodity hardware.