



Activate the fastest path from data to insight

Financial Services eBook



Contents

01

Enabling transformative insights in Financial Services 3

02

Customer story: FINRA 7

03

Customer story: Global financial institution enables anti-money laundering with Starburst 10

04

A modern approach to securely bring all your data to every critical decision. 12

05

Execute on data anywhere 13

06

Customer story: Assurance 14

07

Customer story: Inter. 16

01

Enabling transformative insights in Financial Services

Intensified competition, market volatility, evolving customer expectations, and an increasingly complex regulatory environment is creating a perfect storm of uncertainty and pressure for financial institutions. As banks, investment firms, insurers, and indexes examine their strategies to continue to maximize return on assets in a compliant and sustainable manner, the role of enterprise and alternative data is becoming increasingly critical to their success. With digital transformation efforts in full flight and cloud adoption on the rise, waiting for data centralization efforts to be completed to start extracting value from data is not an option. Instead, financial institutions need to adopt a framework that enables them to extract critical insights and take action on their troves of data today - regardless of where it's stored. A framework that runs in

parallel and supports their digital transformation strategies to meet customer, regulatory, and competitive pressures.

With all the data available to financial institutions, it provides a massive opportunity for financial services organizations to optimize processes, improve return on assets, manage risks and exposure better, enhance customer experiences, and create new products and services. And with fast and secure access to all that data, build a competitive edge.

By adopting a modern data and analytics approach with secure access to every relevant data set, financial firms can drive monumental change — dramatically improve time to insight to unlock new business models, grow new and existing revenue streams, and deepen customer relationships.



Starburst enables faster, better decisions with every relevant data set.

Starburst is the fastest SQL-based data analytics engine for Financial Services built on open standards. As the single point of access for all data across on-premise, hybrid, cloud, or multi-cloud data stores, Starburst provides the fastest, most efficient analytics engine for your data teams to get from data to insight. Whether it's improving underwriting, recommending the next best product, anti-money laundering, mergers & acquisitions, or simply improving internal analytics to get better insights on return on assets,

Financial services firms can now securely connect data silos across the organization enabling teams to query all relevant data, making it instantly actionable. With Starburst, financial firms can lower the total cost of their infrastructure and analytics investments, prevent vendor lock-in, use their existing BI tools, and adhere to data governance and regulations by enforcing reporting and data access, security, and data sovereignty requirements.

Financial services organizations are leveraging Starburst to drive a broad range of innovative use cases:



Anti money laundering

Fulfill your regulatory obligation to track, monitor, report, and respond to money laundering crimes faster and more accurately. Slow investigations based on incomplete, inconsistent access to data, and long or unreliable queries, increase exposure to regulatory fines due to the inability to detect/prevent fraudulent activity. False positives add an additional layer of overhead, requiring lengthy and resource-intensive manual investigations. By enabling access to all relevant data across your entire data footprint with a lightning-fast query engine, AML teams can get faster insights on petabytes of data where it resides, increasing their confidence and ability to execute their strategy. AML teams can now gain consistent and instant access to analytics across 100+ countries with the necessary data sovereignty, data access, and data governance requirements enforced to combat money laundering crimes while remaining compliant.



Counterparty credit risk & credit valuation adjustment (CVA)

It's never been more important to manage the risk of counterparties failing to make good on contracts. Starburst helps you proactively identify threats, so you can model counterparty exposure to emerging threats, either on a single name or portfolio basis, from geopolitical turmoil to declining industries. Fast, accurate queries from large, diverse sources enable you to view the evolution of a portfolio, identify risk, and modify CVA accordingly.



Customer 360° analytics

In the digital era, customers expect personalized, convenient experiences designed to suit their unique needs – or the needs of their particular business. Digital transactions give you valuable information and plenty of it. With Starburst, you can learn where, when, and how customers engage with your brand, products, or services. Applying analytics to this information gives you the insights you need to improve customer experiences, anticipate life events, improve product recommendations, and drive increased lifetime value.



Fraud detection

Protect, detect, and respond is the name of the game when it comes to combating fraudulent activity. To minimize financial losses, financial firms can now enhance their measures to identify and stop theft across account takeovers, credit theft, fraudulent credit, and claims applications. With more data from more sources, fraud identification and investigation becomes a more precise art.



Financial and investment modeling

Asset managers are utilizing new data sets to develop new strategies for investing. By developing models around social media, investors can gain specific insight into sentiment and develop trading signals. Whole new categories of investing emerge from leveraging a single access point for analytics. Analysts are also empowered to model stress tests with a complete picture, as Starburst federates queries across any data source.



Optimize KYC for underwriting

Equip underwriting teams to access all relevant data across any data store for risk assessment. By building a complete KYC profile, underwriting teams can quickly and reliably process the data to enable real-time decision-making, providing financial clients with immediate outcomes and firms with greater confidence in acceptable exposures, risks, and return on assets potential.



Regulatory reporting

Financial institutions must obtain a holistic view of their data landscape to ensure that no critical detail is left out when meeting regulatory reporting requirements. However, accomplishing this is no easy task. Firms need armies of analysts dedicated to data management, analytics, and reporting to meet the requirements. Whether for trade activity reporting or stress testing for market liquidity, financial services firms can gain a reliable path to ensuring all relevant information, regardless of where it is stored, is easily accessible and quickly queried for timely and accurate regulatory reporting needs. To streamline future needs, easily create data products of data sets that are required for each reporting cadence.

02

FINRA®

Every financial services company is a data company. That's because insights from analytics have become an essential driver of innovation, optimization, and competitive edge. The financial services industry has embraced the opportunity. Across the globe, retail banks, investment banks, brokers, exchanges, wealth managers, insurers, fintech companies – even central banks and regulators – are using data analytics to radically transform their organizations. And the gains they're seeing are remarkable.

Analyzing 80 billion trade records daily using AWS and Starburst Enterprise

The Financial Industry Regulatory Authority (FINRA) is a not-for-profit organization authorized by the U.S. Congress to regulate a critical part of the securities industry – brokerage firms doing business with the public. FINRA carries out this mission by analyzing billions of daily trading events from financial institutions to detect fraud, insider trading, and abuse.

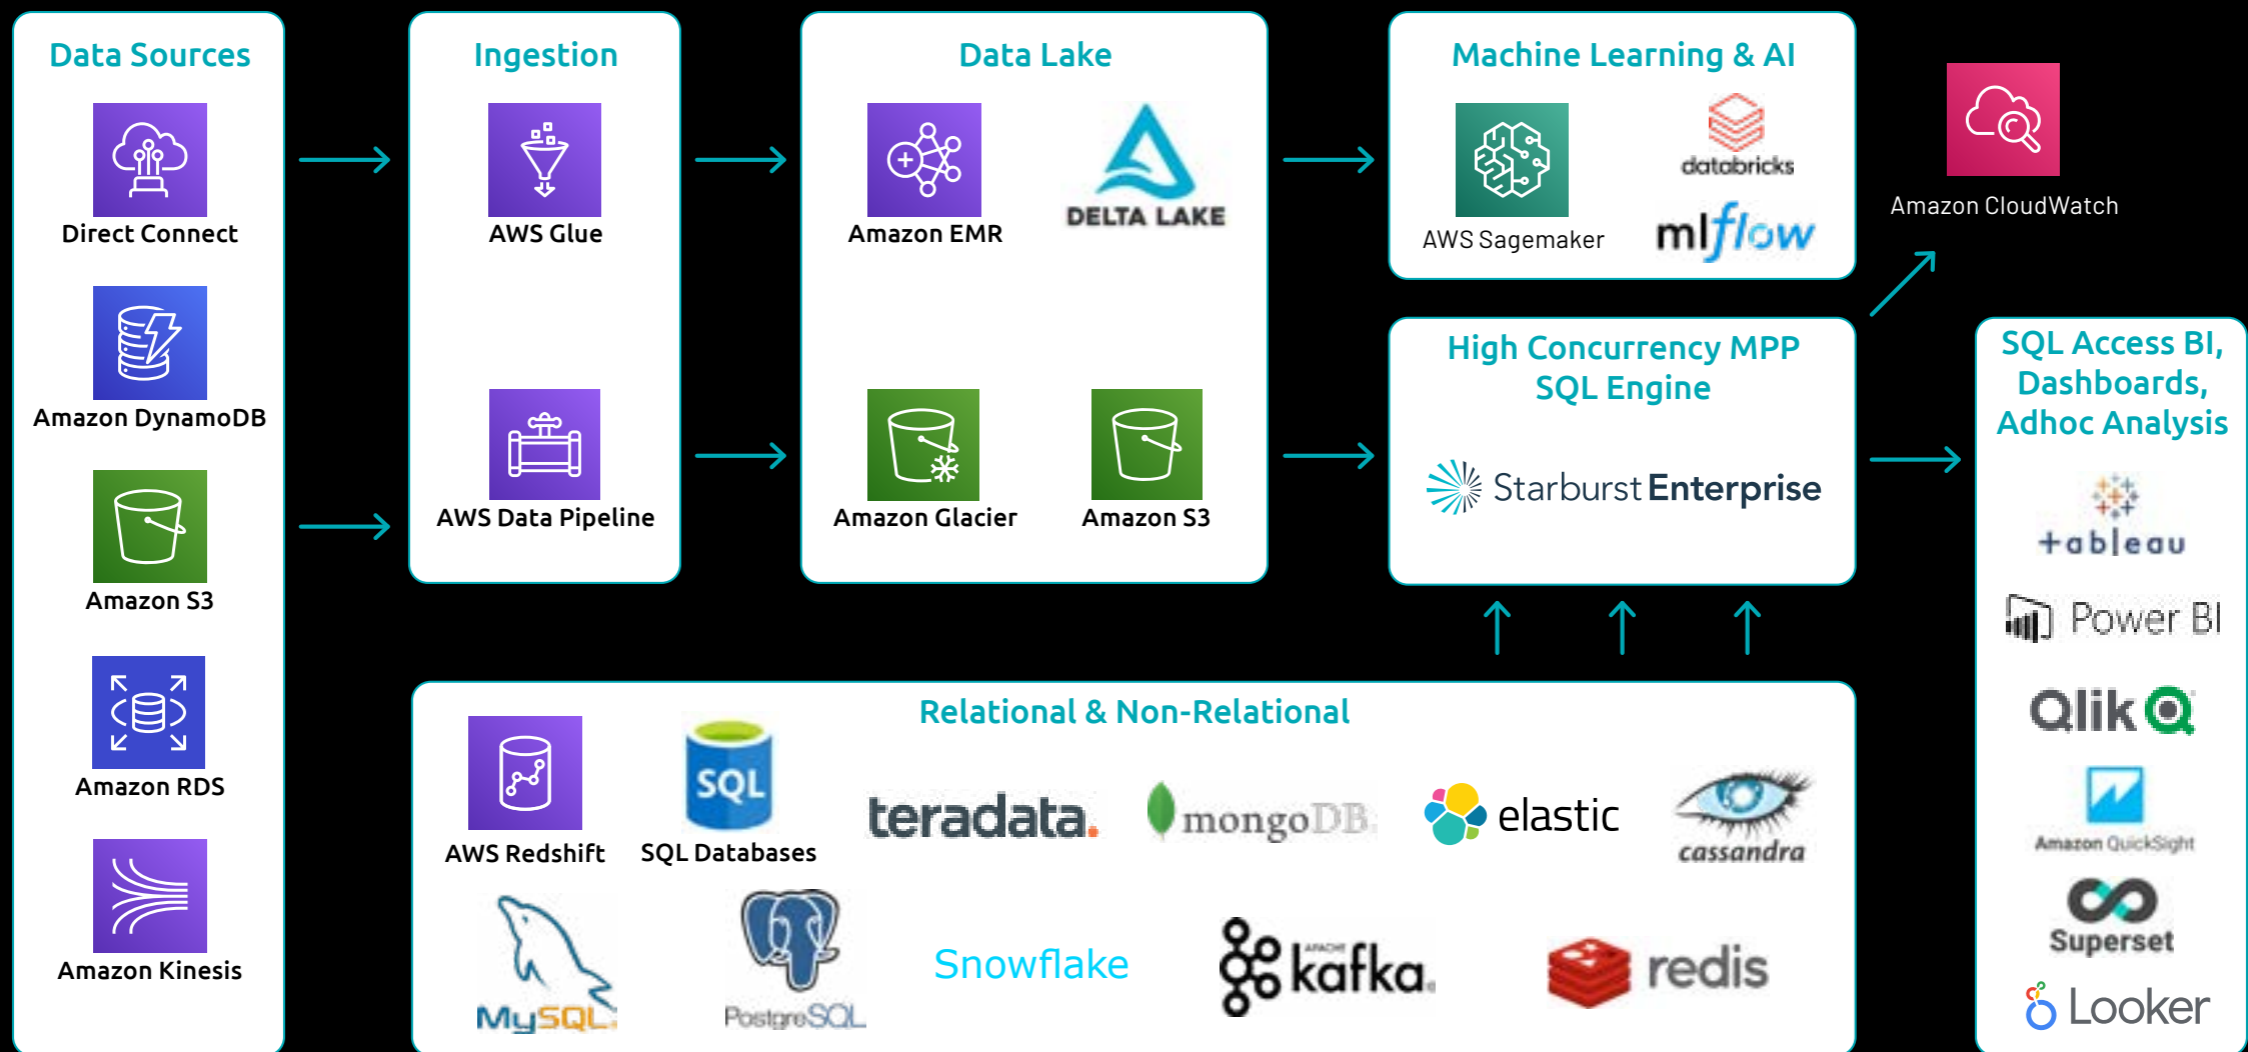
FINRA migrated its legacy data warehousing systems to an Amazon Web Services (AWS) data lake to address the challenges of massive data growth and increasing demand for efficient computing. When redesigning its data platform, FINRA chose to separate compute and storage and query its multi-PB AWS data lake using Starburst Enterprise, the world's fastest distributed SQL query engine.

Challenge: siloed data warehouse platforms, costly to scale and analyze

Until a few years ago, FINRA ran its data warehousing infrastructure on-premises. Organizational barriers and scalability limitations forced FINRA to create separate analytic silos, each handling a subset of the entire dataset. The resulting data fragmentation made analytics difficult. The growing data volume and analytical needs also started to exceed the capacity of its legacy systems. Scaling was expensive and difficult.

Solutions were sized to handle peak capacity, which became very costly. Expanding them was only possible through long procurement processes, and data had to be moved constantly, so it took far too long to perform analytics, slowing time to insight. To address its mounting data storage and processing challenges, FINRA decided to completely rethink its data platform. In 2014, the organization moved from on-prem to an AWS data lake model.

Example Architecture



Solution: Starburst

One remaining challenge was to select an interactive SQL engine that would match the performance of FINRA's legacy MPP SQL systems. For ad-hoc analytics, the query SLAs are measured in seconds. FINRA chose Trino because it was the only SQL engine able to operate at a petabyte scale in the cloud and execute concurrent queries interactively against data stored on S3. Starburst Enterprises' proven integration with AWS was another essential feature.

“Starburst is very data-lake friendly. It was as if it was built for that model. That was a key differentiator for us. We were very invested in the data lake.”

- Ivan Black Director, FINRA

Today, FINRA uses Starburst Enterprise for ad-hoc data profiling, BI, and reporting. Teams of data analysts and scientists execute multiple concurrent SQL queries via JDBC and ODBC clients. Trino then authenticates requests with Active Directory using LDAP and authorizes them via Hive Metastore table permission checks. Finally, during query execution, Trino reads the ORC table data directly off S3. FINRA has also built several interactive web applications which leverage Starburst as their backend SQL query engine to access data in the AWS Data Lake.

Outcomes: Faster insights at lower cost

Leveraging both Starburst and AWS cloud storage eliminated the need to invest in expensive proprietary big data appliances to support ever-increasing volumes of data. Starburst Enterprise enabled FINRA to analyze its data interactively in an ad-hoc manner without the data copying and loading required in the past. The migration from legacy data warehousing systems was seamless to end users, and researching market manipulation and investigating potential fraud is now faster than before. Overall, Starburst has given FINRA a scalable, cost-effective way to analyze its constantly growing volumes of data.

“FINRA monitors market data for trading fraud. Starburst separates compute and storage, making it possible to scale economically and analyze 25PB of data—100B rows of new data per day from 25+ sources.”

- Ivan Black Director, FINRA

03

Global financial institution enables anti-money laundering with Starburst

One of the world's leading investment banking and financial services companies is on a mission to help make financial lives better with modern financial solutions to contribute to a growing economy. The multinational investment bank serves millions of people and institutions in over 100 countries. Banks are typically associated

with their services, such as checking and savings services, loans and financing, investment management, etc. However, as worldwide banks mediate millions of transactions per day, like every other financial firm, this institution is at a high risk of financial crimes.

Challenges: Complex regulatory environment. Data sovereignty regulations. Lengthy AML investigations.

AML regulations are constantly changing to keep up with money laundering trends. Financial institutions must stay on top of the rapidly changing regulatory landscape to monitor increasingly sophisticated money laundering techniques. On the other hand, data sovereignty regulations present a barrier to monitoring illegal activity. It is becoming commonplace that each country has its own data privacy regulations. In certain countries, regulations prevent banks and institutions from taking data out completely. Identifying a bad actor or institution without proper access to account data

becomes extremely difficult. Furthermore, with millions of transactions to monitor daily worldwide and a wide range of products and services, detecting illegal activity becomes increasingly complex. For this financial institution, the amount of data that needs to be collected and stored is enormous. With their data archived in multiple regions around the world, data duplication was costly, access to data was challenging and time-consuming, and as a result, the bank wasn't able to meet the time-based demands required by government regulators, such as the Office of the Comptroller of the Currency (OCC).

Outcome: Money laundering detection and prevention is accelerated by the instant retrieval of suspicious activity data.

With Starburst, the multinational bank improved the detection of money laundering, reduced the number of false positives, delivered timely and actionable insights to AML investigators, stayed ahead of ever-changing regulations, and reduced costs.

Solution: Enable global access without data duplication or movement

With Starburst, the bank can access data from anywhere using Starburst without needing to duplicate or move it. Not only does this solve the challenges around timely access to data, but it also solves the problem of data access from countries with strict data sovereignty rules. "Many of the challenges will no longer exist as we are not moving data out of another country or reading any personal information," says the Head of Anti-Money Laundering at the multinational bank.

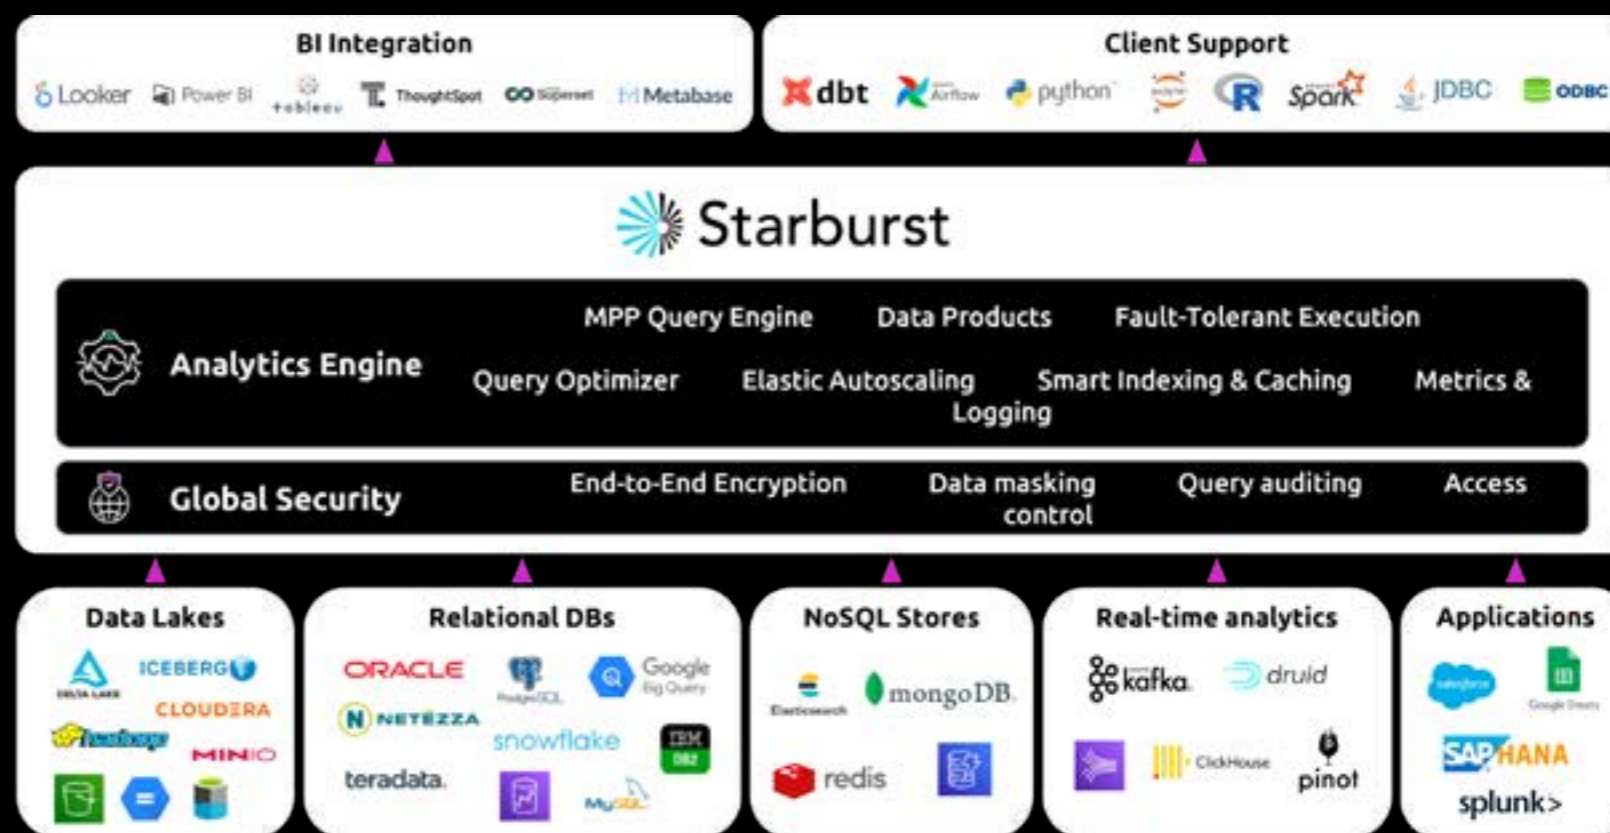
- **Enhanced money laundering detection model** - Each time a new financial product is created, the team identifies the AML risks associated with the product and creates a machine learning model to prevent and detect illegal transactions. Starburst allows AML data scientists to rapidly test new models, update existing models, and profile new data. The more intelligent models increase the efficiency of the alert investigation process. In the future, the machine learning models will access all historical and present data sets to constantly train and improve the model in real time.
- **Timely and actionable insights** - When investigators needed more information on a case, they had to ask the monitoring team to copy and send extracts of information from multiple physical warehouses. Once the data was available, access was further delayed due to the poor performance of Apache Hive. With the built-in infrastructure of Starburst, the investigations team can access AML data in near real-time and across multiple data sources.
- **Data governance to lower compliance risk** - With government regulations for AML constantly evolving with modern money laundering trends, Starburst's federated computational governance ensures continual compliance and data security.
- **Reduced costs** - Less data duplication and quicker time-to-insight results in cost savings. The cost of investigations is also lowered by putting more complete data in the hands of investigators, which reduces the rates of false positives.
- **Reduced risk of fraud and regulatory fines** - Access to big data and big data analytics (in partnership with Tableau) helps the financial institution address fraud more effectively to better monitor transactions and make decisions effectively. This reduces the risk of fraud and the risk of penalty for AML violations.

04

A modern approach to securely bring all your data to every critical decision

Various departments across your organization have different data needs. These needs are served with different dashboards, tools, databases, and more. A modern strategy works to enable users across your organization with tools that create efficiencies. Data professionals are empowered to leverage their preferred analytical tools and data sources, thus eliminating unnecessary data silos and duplication. As you work on creating a single source of truth via data migration efforts, Starburst creates a single point of access for all your data so you can immediately extract valuable insights to increase return on assets, mitigate risks, and adhere to regulatory requirements.

Starburst is the analytics engine that fits into any environment



05

Execute on data anywhere

With Starburst, data teams can unlock data silos across business units, clouds, countries, and regions by gaining a secure single point of access across all your data sources. With over 50 enterprise data connectors from on-premise or the cloud, data lakes to data warehouses, and more, Starburst provides you with a holistic view of all your data. Starburst can connect with the most popular data sources, including MySQL, Hive, PostgreSQL, Snowflake, Delta Lake, and many more. Starburst's enhanced connectors drastically improve overall performance and security with parallelism, pushdown, table statistics, and fine-grained access control, while fault-tolerant execution ensures the reliability that every query is executed completely.

Data teams can also introduce massive productivity gains by enabling data producers and consumers to create, publish, discover, and manage curated datasets in the form of data products. Lastly, having data stores on multiple clouds or regions/countries is no longer a barrier to creating value. Starburst enables financial services firms to link catalogs and data sources to other catalogs and sources without needing to copy or move data. Starburst becomes a gateway for unlocking data access across geographies while ensuring access controls are in place and data residency requirements are honored - a global approach to common initiatives like KYC, Fraud detection, AML, and reporting just got significantly easier.

To learn more, please visit starburst.io

06



Assurance reduces time-to-insight, cuts costs, and boosts customer conversion rate.

Assurance, a subsidiary of Prudential Financial, is an online insurance platform that operates between consumers and the hundreds of thousands of insurance and financial services products on the market, leveraging data science and a human touch to find the ideal match for each individual.

Assurance has different models running in the background at each step of the customer journey that draws on responses and basic customer data,

and insurance policy offerings and other data to generate the best possible recommendations. By deploying Starburst, Assurance has reduced time-to-insight from weeks to hours, provided Data Science and BI teams with easy access to more data sources, strengthened security, increased its overall conversion rate by 10%, and reduced one recurring cost by 40%.

Challenge: Joining data from multiple sources was time-consuming, manual, and unscalable

At Assurance, data is stored in multiple Postgres databases in AWS, an AWS S3 data lake, and other sources. Before partnering with Starburst and Trino, Assurance struggled to join data from one Postgres DB to another or its data lake. Typically, an engineer would have to copy data across manually for data scientists or other teams to run dashboards or create reports from two or more separate datasets. Data engineers are a shared resource across the company, so when a business group wanted access to a new data source, they had to get in line.

“We couldn’t scale to the needs of the business. For anything that wasn’t a top priority for the company, that elongated time-to-insight from weeks to months, and it really limited our ability to connect disparate parts of the business.”

- Shen Wang, Principal Data Engineer, Assurance

Solution: In search of a more efficient distributed query engine, Assurance tested open-source Trino, but the company experienced limitations with the connector's ability to access certain data types in Postgres.

Assurance then deployed Starburst in one AWS account, using the cloud formation template to set up autoscaling, security groups, high-performance connectors, and compute clusters. Currently, Assurance has one cluster dedicated to ad hoc queries from the Business Intelligence layer (Looker and Tableau) and another dedicated to production jobs and data engineering.

Outcomes: Faster insights and improved customer conversion rates

“When we introduced Starburst and Trino, we were able to instantly bring together all these different data sources and federate our queries, so we were able to reduce our time-to-insight.”

- Shen Wang, Principal Data Engineer, Assurance

Leveraging Starburst accelerated the timeline for new projects from months or weeks to days. Trino's ability to function as a single point of access for all data also allowed the team to respond quickly to a major issue during the peak of the enrollment period. With Starburst and Trino, business units and data scientists are generating better insights across the full funnel, resulting in a 10+ basis point increase in conversions in 1/100th the time it would have taken before.

Banco Inter unlocks the value of data with revenue-generating insights and reduced infrastructure costs.

Banco Inter (Inter) is a financial technology company based in Brazil. As Brazil's first digital bank, the company has a history of paving the way for digital innovation in the banking industry. After evolving the business to become a provider of financial and non-financial products and services through Inter's Super App, the Brazilian bank transitioned to the data lake architecture to unlock the power of big data.



Challenge: Without a scalable and performant query solution, Inter struggled to generate quick insights that were needed to shape the future of the digital marketplace.

Inter's data users were held "hostage" by Inter's analytics processes, completely dependent on data teams to serve them the data. Each request for data sets required hours, and sometimes days, to move the data to another zone in the data lake before the data team could enable a connection to analytics and visualization tools for team members to view it.

Solution: Starburst & Trino

Starburst's federated semantic layer simplifies data integration and access while optimizing query performance across distributed data assets. Now, data consumers can rapidly deliver revenue-generating insights on consumer behavior. With access to a complete set of data from sources such as the Super App, NASDAQ, Brazil's stock market (B3), and the central bank, Inter can create personalized recommendations for the right product to be provided to the right customer at the right time. "We have upselling and cross-selling campaigns that we had to stop before because we didn't have the query performance or data access. Now that we have Starburst, we've unlocked projects that generate sales from personalized recommendations to our products and services. This is the path to monetization," shares Pedrosa.

Outcomes: Faster insights, lower infrastructure costs, and greater return on assets

Now, the user can easily access the platform and explore the data they need; data scientists can easily connect to Starburst and explore data through their notebooks. Starburst's plug-and-play connectors speed up business intelligence (BI) applications and data visualization, enabling the business user to use Trino as a data serving layer without requiring engineering teams to act. Since its deployment, Starburst has supported Inter's growth in several ways:

- Revenue Generation - Data consumers have a single point of access to all of the company and customer's data to understand consumer behavior, resulting in upselling and cross-sell opportunities. Greater insights into the performance of Inter's products and services also help Inter enhance its marketplace offerings, thereby increasing profitability.
- Increased Analytics Access - Starburst democratizes data for Inter by allowing thousands of data consumers to concurrently access huge data sets at high query volume with low latency. Before, only a few team members could access data. Within one month of deployment, over 120 users were empowered to explore data to prove their hypotheses on their preferred analytic and reporting tools.
- Faster Time-To-Insight- With 50+ out-of-the-box connectors, Inter can provide immediate access for near real-time data exploration.
- Potentially Hundreds of Thousands in Cost Savings Yearly - Inter has reduced infrastructure costs dramatically by moving from Amazon Redshift and Amazon Athena to S3 and Starburst. In addition to data warehousing costs exceeding \$100,000 per month, the cost per query on Amazon Athena ranged from \$40 to \$ 1,000.

"Every new process, every new idea, and every new initiative can and will be based on data. Starburst empowers our decision-makers to have the best information on a complete set of data. It is a shifting point that will impact the future of the company."
– André Gortari, Data Engineering Manager at Banco Inter

About Starburst

Starburst is the analytics engine for all your data. We provide the fastest, most efficient analytics engine for your data warehouse, data lake, or data mesh. We unlock the value of distributed data by making it fast and easy to access, no matter where it lives. Starburst queries data across any database, making it instantly actionable for data-driven organizations. With Starburst, teams can lower the total cost of their infrastructure and analytics investments, prevent vendor lock-in, and use the existing tools that work for their business. Trusted by companies like Apache Corporation, Comcast, Doordash, FINRA, Marks and Spencer, and VMware, Starburst helps companies make better decisions faster on all data.

[Starburst.io](https://starburst.io)

