

Assurance reduces time-to-insight, cuts costs, and boosts customer conversion rates with Starburst

The online insurance platform Assurance 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. The company's mission is to make it radically easy for customers to purchase life insurance, mortgages or similar products that are right for them, and data is an essential piece of that strategy.

As customers progress through their buying journey, Assurance has different models running in the background at each step. These models draw on responses, basic customer data, 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%.

Background

At Assurance, data is stored in multiple Postgres databases in AWS, an AWS S3 data lake, and other sources. "We are a very data-driven company," says Principal Data Engineer Shen Wang of Assurance. "We rely on the human touch as well, but we have very powerful metrics that help inform our people at every stage."

Prior to partnering with Starburst and Trino, Assurance struggled to join data from one Postgres DB to another or to 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. The company sells many different insurance products, too, so it was a major challenge to determine which program deserved to be prioritized over another.

"We couldn't scale to the needs of the business," explains Wang. "Each team was fighting to be at the top of the backlog to get their product up and running. 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."



ASSURANCE



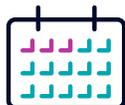
We are a very data-driven company. We rely on the human touch as well, but we have very powerful metrics that help inform our people at every stage.

Shen Wang
Principal Data Engineer
Assurance

Starburst Enterprise & Trino

In search of a more efficient distributed query engine, Assurance tested open-source Presto, but the company experienced limitations in terms of the connector's ability to access certain data types in Postgres. As a web-based company, much of Assurance's data is in JSON, and its predictions are published as JSON objects. Parsing this data was slow and time consuming.

Assurance then deployed Starburst Enterprise (the enterprise edition of open-source project Trino) in one AWS account, using the cloud formation template to set up autoscaling, security groups, various high-performance connectors, and compute clusters. Currently, Assurance has one cluster dedicated to ad hoc queries from the Business Intelligence (BI) layer (Looker and Tableau) and another dedicated to production jobs and data engineering. The impact on Assurance's business has been significant.



Three Months to Three Days

One of the more exciting developments was the accelerated timeline for new projects. "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," Wang explains. "Anything that was low priority dropped from three to four months to three or four days. The more urgent programs went from a couple of weeks to a couple of days."

Plus, these programs put far less stress on Wang and his colleagues, as they no longer required ETL projects.



10% Increase in Conversion Rates

With business units and data scientists leveraging fast access to more data, including disparate JSON blobs, Wang says there has been a significant impact on customer conversions. "With Starburst and Trino our business units and data scientists are generating better insights across our full funnel from marketing all the way to sales and conversions," explains Wang. "This has really helped them affect specific conversion rates at every point. Because we can do it so quickly with Trino and Starburst we were able to get a 10+ basis point increase on conversions in 1/100th the time it would have taken before."



With Starburst and Trino our business units and data scientists are generating better insights across our full funnel from marketing all the way to sales and conversions.

Shen Wang
Principal Data Engineer, Assurance



Real-time Sales Optimization

Assurance serves two customers: potential buyers interested in new products and licensed insurance agents who log on to the platform in hopes of converting these individuals into customers. When the number of potential buyers spikes, Assurance needs to be able to incentivize the appropriate agents to log on and convert those individuals into customers. A failure to do so is essentially lost business. Conversely, wasting the valuable time of these agents might leave them reluctant to log on the next time.

Assurance has the data to optimize these processes, but the company wasn't able to leverage it until recently. As Wang explains: "One of the key benefits of Starburst Enterprise is that we've been able to bring together our shopper and marketing data, along with details on which agents are online, to make predictions that prompt sales managers to react intraday and either incentivize new agents to come online because of an influx of customers or to cool off because of a shortage."

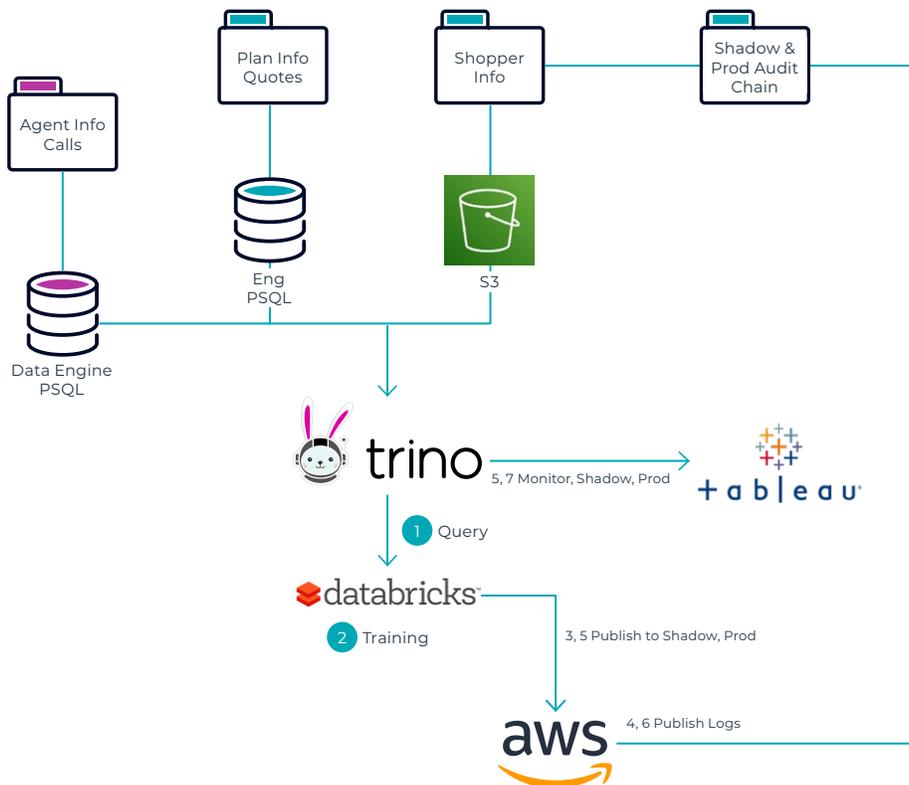
This improved the experience for the agents who drive the sales process, making them more likely to utilize Assurance again.



Driving Higher Quality Sales

Assurance is also able to assess how long a customer remains a customer – the longer, the better. If the customer keeps a policy longer, then Assurance knows that its platform suggested the right product for that individual. By allowing its models access to more data, Starburst is helping to drive better product choices, resulting in customers with higher long-term value.

Assurance Plan Recommendation Model Development Flow



Addressing Acute Business Challenges

The advantages highlighted above relate to Assurance's business as a whole, but the company specifically cites Starburst's now critical role in one of its most important revenue-generating events, the open-enrollment period for Medicare. This seven-week stretch is to Assurance what Valentine's Day is to a florist, and Starburst has delivered tremendous value to the business in three key areas.

1. Optimizing Agent Licensure

To enroll a customer in a Medicare plan, Assurance needs to match them with an agent licensed in that individual's state. Assurance pays the associated licensing fees for its agents, and this cost can total tens of millions of dollars. If the company doesn't have enough agents licensed in a state, it could lose out on revenue. A surplus of agents, on the other hand, amounts to wasted dollars.

The Central Intelligence Unit, a Data Science group within Assurance, needed to develop a model that draws on relevant and real-time data to optimize agent availability and match that to customer demand. Previously, distributed, siloed datasets made this extremely difficult. "Different teams own different systems that produce the relevant data and it's stored in multiple locations and multiple formats," explains Senior Data Scientist Mitchell Posluns. "With Starburst that's not a problem at all. It's the centerpiece that kicks off model development, execution, publishing, and analysis, and it's accelerated by the fact that the Data Science team doesn't have to worry about any of the complexities of collating data or putting it into a consistent format. We can focus on analysis and model formulation."



We estimate that this optimized solution resulted in a reduction of about 40% on that multi-million dollar cost line, which is just huge.

Mitchell Posluns
Senior Data Scientist, Assurance

The result is a more robust model that optimizes agent licensure, ensures sufficient coverage, and allows Assurance to offer more products to more customers, increasing its conversion rates. Plus, the Starburst-backed models have cut costs significantly. "We estimate that this optimized solution resulted in a reduction of about 40% on that multi-million dollar cost line," Posluns says, "which is just huge."

2. Accelerating Time-to-Insight

The time savings were significant as well. Previously, the Data Science team would have had to collate data, put it into a consistent format, and work through other data-wrangling tasks before even beginning to test its models. "We were able to complete this in a week instead of several months," Posluns adds. "Not only did it accelerate time-to-insight, but it gave us a jump on the licensing process, so we could get more agents licensed in time for the start of the enrollment period."

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. Unexpected events such as TV commercials can lead to spikes in demand. Assurance needs to be able to respond in real-time, reach out to its agents, and bring them online to help shoppers. During one peak period, Assurance's real-time logging tool went down, and the company was effectively flying blind.

Once the panic subsided, the company's Director of Engineering suggested moving streaming events to its data lake, and the Data Science team got to work. Within just three hours, all the key data was streaming into its data lake, which was connected to Tableau via Starburst. "It was not that hard at all, and what resulted was real-time analysis powered by Starburst," says Posluns. "The data sources were abstracted away. We didn't have to worry about sources and formats. We just plugged in to Starburst, built this beautiful and insightful dashboard, shipped it to stakeholders, and got back to business as usual."

What could have taken days or weeks to replace, and likely would have resulted in hundreds of thousands of dollars of lost revenue, took the team a few hours. "We did it quickly enough to present the dashboard at a company all-hands meeting that same afternoon," Posluns recalls. "It was a big win for us in Data Science, and the dashboard is still used daily, with tens of thousands of views on a daily basis."

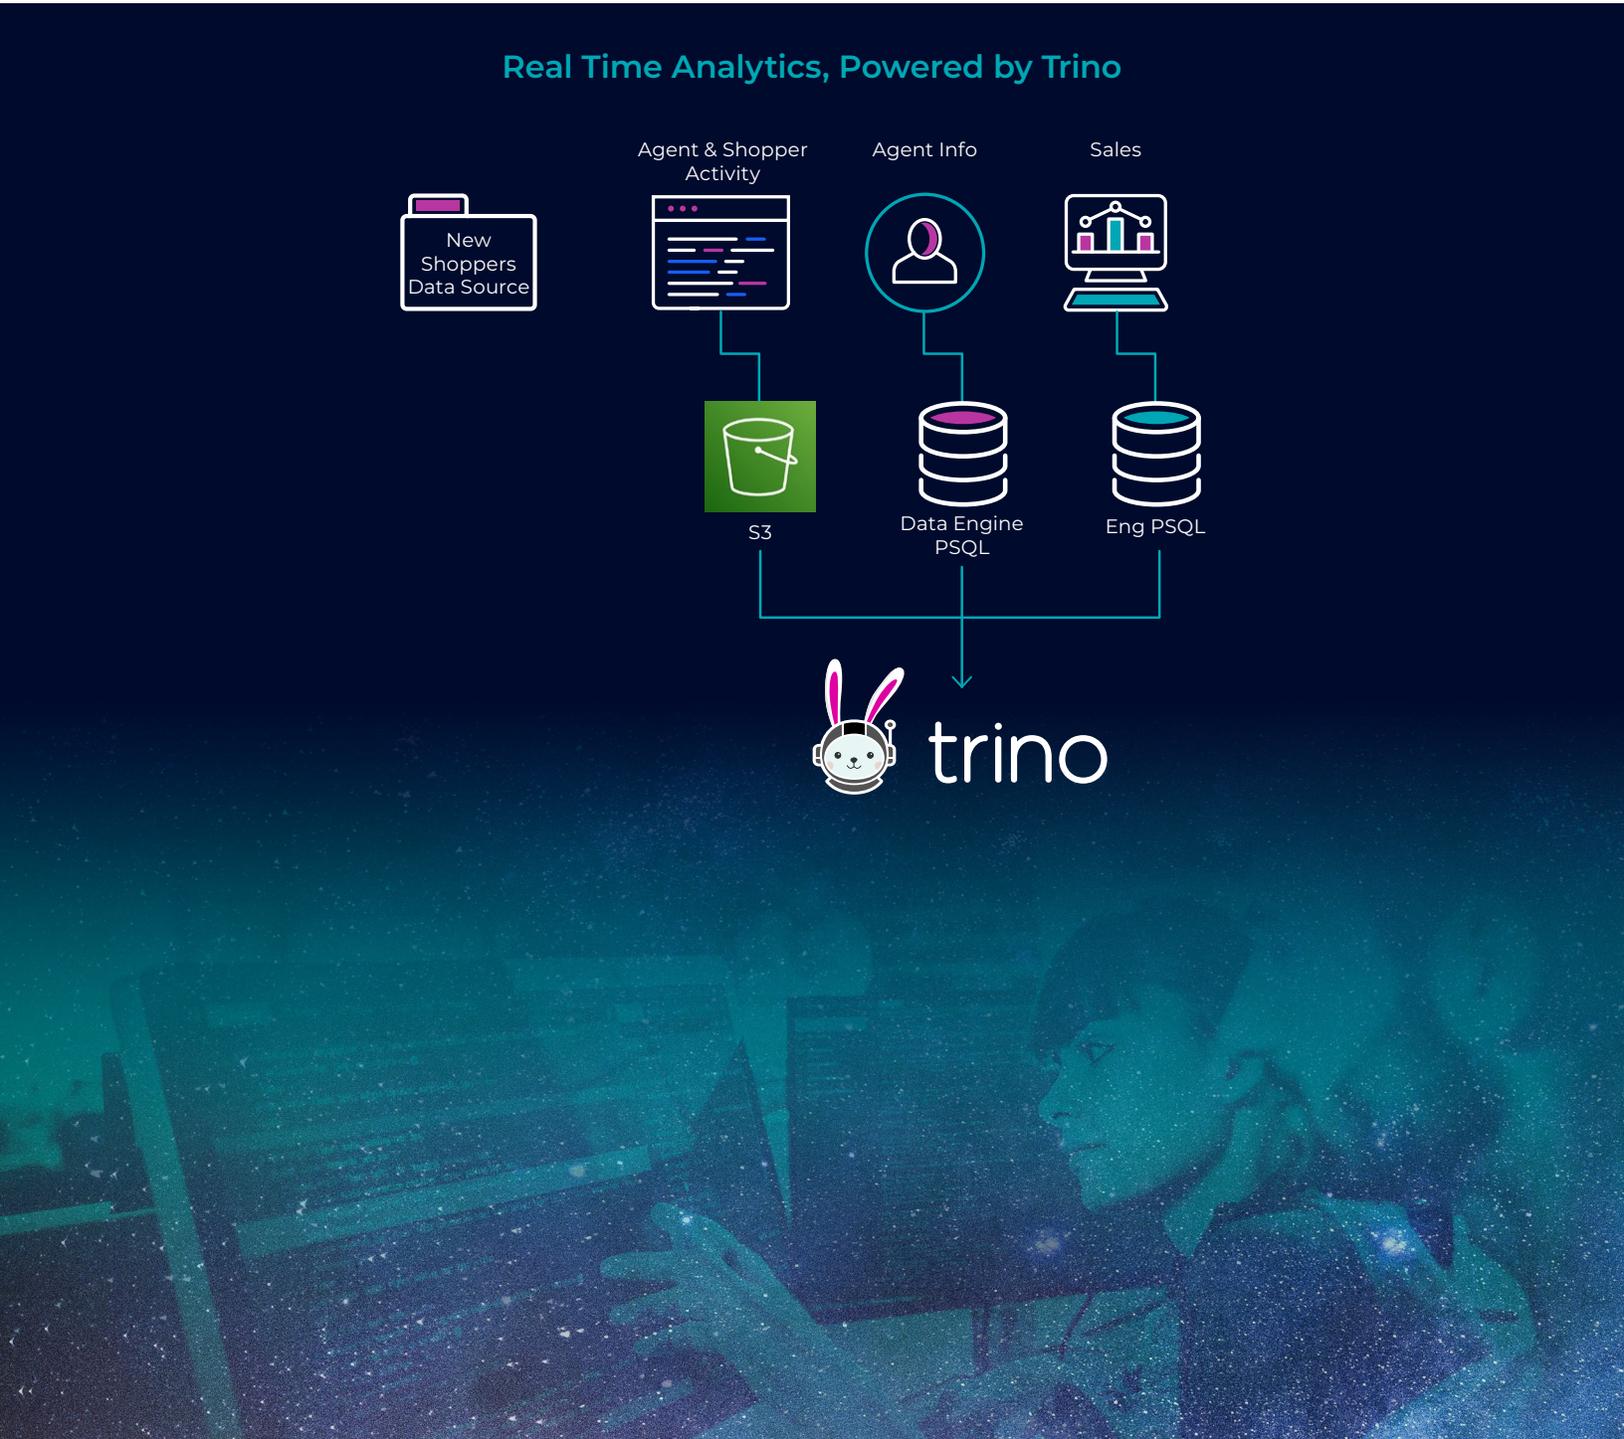
3. Enhancing Recommendation Engines

Assurance is always looking to improve its ability to suggest the right product for the right customer. To do so, the company's recommendation engines have to draw from proprietary Assurance data, third-party sources such as credit agencies, and more. "We don't expect to get everything right on the first try, so we need a really rapid development process to very quickly design, implement, evaluate, and improve our model," explains Posluns. "This is a really important business problem to unlock and it's where Starburst comes in. Starburst is a critical piece of this rapid development model."

With Starburst, Assurance is able to immediately ingest the data needed for its initial model training, iterate, deploy, and monitor the model's performance in a shadow environment. "Then when we're ready to go," Posluns says, "we can deploy the model into production and actually start using it as a shopper goes through their experience."

More robust, reactive modeling leads to better recommendations, which turns into improved customer retention rates and increased revenue.

Real Time Analytics, Powered by Trino



The Starburst Enterprise and Trino Difference



Performance & Cost-based Optimization

Switching to Starburst Enterprise accelerated queries by 20%-30%, reducing the company's compute costs. "We saw the speed-up because queries we set up were re-ordered to have a better join options setup," Wang says. "That was amazing."



Global Security & Role-Based Access

Control (RBAC)

Assurance has access to Personally Identifiable Information (PII) and Personal Health Information (PHI), so strong security is absolutely critical. "We are rightfully concerned about the security of our data, so we were thrilled to bring in Apache Ranger with RBAC to limit users to specific rows, columns, tables, and schema," Wang notes. "This segmentation of access is critical. We are able to lock down who has access to what data."



Enterprise Support from the Trino Experts

The data engineering team at Assurance is in high demand, so being able to engage with an extremely qualified team of Trino experts when a new feature

needs to be set up or hurdles arise saves them valuable time. "I'm very impressed with the support team and their response time," Wang stresses. "They've been really helpful in setting up our security and helping with anything that goes wrong."



Data Scientists as Data Engineers

Now that Assurance's teams have fast, federated access to more data sources, Wang and his data engineering team are seeing more creative use cases relying on more and more complex queries. The group is exploring the possibility of using Starburst for specific ETL jobs, shifting to a multiple account world with AWS, and potentially deploying a Starburst-to-Starburst connector.

Overall, the platform is empowering data scientists. "We're trying to push the envelope and allow our data scientists to be more like data engineers," Wang adds. "Starburst and Trino have really helped them get more comfortable with accessing their data because it's in a SQL language they're familiar with."



Starburst has really enabled us to accelerate time-to-insight, improve our conversion rates, and enable more robust modeling, which have all led to better business outcomes.

Mitchell Posluns
Senior Data Scientist, Assurance

CONCLUSION

Most modern organizations describe themselves as being data driven, but at Assurance this isn't just marketing talk. The company's current CEO was its former Chief Data Scientist, and providing faster access to more data relating to different parts of the business and sales process drives measurable results. In providing high-performance, distributed access to more data in an enterprise-ready platform, Starburst has proven to be a valued partner, according to Wang: "Every business team relies heavily on the data that is accessed through Starburst to help our customers get the best insurance policies possible and to help our agents be more successful."

