

TOOLSTATION.COM

Toolstation's Approach to Modern Data Streaming in the Cloud



Headquarters

United Kingdom

Industry

Retail

Challenge

Limitations with click and collect (buy online, pick up in store) due to a legacy batch processing paradigm, during a pandemic, increased the pressure on a small engineering team to shift to an event-driven architecture powered by Confluent, in the cloud.

Solution

Confluent offered both the visionary qualities and the crucial practical features to help Toolstation eliminate the limitations of legacy batch processing, contributing to immediate revenue growth—as well as multicloud capability with GCP integration built in.

Results

- Ability to quickly implement and scale and move into production in six weeks with the help of Confluent's Professional Services
- An end to the risks and the runtime issues of batch processing
- The advantage of multicloud capability, and integration with GCP
- A cost-effective solution to maintain a healthy bottom line

Purchased by British builders' merchant and home improvement retailer Travis Perkins in 2014, Toolstation is a omnichannel retailer of building tools and materials with more than 500 branches in the U.K.—plus over 100 in the Netherlands, Belgium, and France. As such, Toolstation's IT team is something of an anomaly: a small, agile team ensconced in a bigger, more traditional organization.

Even as the IT team has grown and become more structured over the last few years, Head of Architecture Stuart McGrogan says, "We're still architects with a small 'a.' We're pragmatic and delivery-focused—not so much a blue-sky-thinking sort of enterprise."

To support a hands-on style of progress, achieve rapid, steady growth, and to address limitations with the click-and-collect shopping in their stores (60 click-and-collect orders per branch/store per day), the IT team aimed to replace the legacy batch processing infrastructure, solve for technical debt, and remove silos between departments. There's an ambitious roadmap in place to become fully cloud native, with a modern tech stack, by 2023. An enormous factor in that evolution is the ability to put data in motion with Confluent's data streaming platform so the organization can take full advantage of an event-driven architecture.

"When it comes to batch processing versus stream processing, both of them work. But one is a much more efficient way of doing it. A lot of the cost savings over the last twelve months has come from moving into a more efficient way of developing applications with stream processing."

— Rob Contreras, Technical Architect, Toolstation



The Challenge

The balance of excellent CX and a healthy bottom line

"Everything we do as an organization is done from either the CX point of view—to improve the customer experience—or for profit," says McGrogan. "Our mantra across the business is to prioritize things that will generate revenue. It's a blessing and a curse, because sometimes things that need to be done don't seem as important because they won't make a difference to the bottom line."

Case in point: The legacy infrastructure Toolstation has been operating on for monitoring their stock levels was essentially one big MySQL database instance with a host of polling processes, which meant that nothing was real time. In addition, batch jobs ran on a periodic basis (every minute or every hour) and there was a lot of integration work needed between systems. This meant that customers were occasionally buying the same item online and at their stores at the same time, due to out-of-sync data between systems. Further, stores were limited to only being able to view 60 click-and-collect customer orders at any given time, meaning orders were getting queued up overnight. The IT team wanted to move away from that paradigm into using

data streaming, with an event-driven architecture powered by Confluent and built on top of Apache Kafka®, in the cloud. But the cloud presents a problem. With it, you get all the CX benefits of an always-on infrastructure, reliability, and resilience, but the risk, McGrogan says, is "a bill that reflects it. For this reason, consumption-based models have been tricky for us."

"The cloud is not cheaper than it would be for us to run a traditional datacenter," he continues. But using the right combination of cloud-based components helps: "With Confluent, we get consistent pricing that makes it easier—which is important, because we have a budget, and we're a low-margin business."

The Solution

Migrate first, modernize second

Before the team decided on Confluent, the initial question was: Which approach to take? Rob Contreras, technical architect at Toolstation, says: "We looked at various different (integration and middleware) solutions and arrived at Confluent because of the ability to use stream processing and connectors as a methodology to try to move away from the MySQL database, bit by bit."

They considered building their own architectural solution for event streaming, but ultimately decided they'd much rather consume an existing managed service for Kafka. McGrogan says: "Rather than building our own complex infrastructure, we looked at the best option, and Confluent came up, obviously—you invented Kafka. So we just went with that."

Capacity and resilience became imperative during the pandemic, when e-commerce became the primary driver for the business. This spurred a "lift and shift" of their legacy tech stack to the cloud within six weeks, in which Confluent and professional services played a key role. The first data streaming project was around getting a real-time view of their stock levels by moving from inefficient polling API processes (from each of the 2,000+

cash registers within their stores) to real-time event streaming, and removing the limitation for click-and-collect orders. Confluent gives Toolstation the ability to push new orders to stores in real time, which showed an immediate impact on revenue. The second use case went live in the warehouse environment in 2021, with Confluent connecting two different third-party systems for a new voice picking system. The third use case is around push notifications where now, anytime a customer places, tracks, or picks up an order, it all goes through Confluent. Toolstation has also migrated to a new search engine, so that when product information is updated or stock levels or item prices are updated, all this goes through Confluent and then gets pushed to the search engine.

While the team is still in the process of migrating away from MySQL, progress has been excellent so far. McGrogan says, "A lot of our workloads are moving from batch-based processing to a stream of data coming through from these events, and moving this to datastores that are constantly being updated." Toolstation currently has around 10 apps running on Confluent, with a few more projects about to launch.

"We looked at building our own architecture, but we'd much rather consume services than build our own complex Kafka infrastructure. And you [Confluent] invented Kafka, so obviously we just went with that."

— Stuart McGrogan, Head of Architecture, Toolstation

The Results

Here's how Toolstation has succeeded on the path to modernization:

De-risk the legacy database

Previously, everything in their service-oriented architecture was batch-driven, and that created a lot of technical debt. "The problem with legacy is, you never maintain it. It just gets put in a box and left alone. The only time you look at it is when it breaks," says McGrogan. "Because it is a service-based architecture, the database is one of our biggest risks. It's there, everything runs from it, so Confluent gave us a way of de-risking it."

The advantage of multicloud capability

Toolstation moved to the cloud about two years ago and is currently focused on becoming as fully managed as possible when it comes to local services: "I definitely prefer the approach of not tying ourselves into any specific cloud. That's the reason Confluent was appealing to us," says Contreras.

Specific integration with GCP

While a multicloud approach was important, the ability to use Google Cloud Platform (GCP) right away was critical. McGrogan says: "We wanted Kafka, and a hosted version of that made sense. And it works in GCP. Although we want to be cloud-agnostic, we need things to work with GCP."

A more cost-effective, efficient, flexible solution

As a fast-paced retail organization, Toolstation has experienced consistent 25-30% year-over-year growth—even higher during COVID-19, in which Confluent played a key role. Although growth has started to level out again now, the team will always choose a platform that provides a path to profit without enormous operating costs. McGrogan says: "Value for money is a key driver for us. We probably have an IT budget that's 50% of what most companies would be spending on IT for a comparable size business. We don't do vanity projects. There's no IT-driven projects. The things we do, we do for a reason—because the business needs it. IT is an integral part of the business."

"Confluent is one of those foundational architectural components that, now that we know, will make up any kind of solution that we put together."

— Stuart McGrogan, Head of Architecture, Toolstation

The Future

What's next for Toolstation?

By the start of 2023, the Toolstation team plans to be fully cloud native, but the approach has been "migrate first, modernize second." "Now," says McGrogan, "we're in a position where we're starting to put modernization in place. That's the benefit of being pragmatic. If we had been waylaid with architectural dreams, we'd have modernized a lot of the components. But the way you do things quickly is to be pragmatic, not an architectural purist. That's what we did, and it was really successful. We have big projects coming up this year, which would be massive to the business, in which Confluent is going to be an integral part."

As Toolstation continues to mature in its modernization efforts, and hopefully experience continued revenue growth year over year, data streaming will be an invaluable element of the trajectory, enabling the IT team to grow out a multicloud approach while continuing to innovate on applications.

Learn More About Toolstation

<https://www.toolstation.com/>

Learn More About Confluent Cloud

<https://www.confluent.io/confluent-cloud/>