✳ CONFLUENT

Alight Solutions Increases the Pace of Digital Deliveries with a Unified Data Platform

alight

✳ **Confluent Platform**   ▦ **Schema Registry**

Headquarters
Lincolnshire, IL, USA

Industry
Business Services

Challenge
Integrate data from numerous back-end systems into a forward cache to improve performance, simplify consumer-facing application development, and reduce demand on costly mainframes.

Solution
Use Confluent Platform and Apache Kafka to implement scalable, secure streaming data pipelines that refresh data in the cache in near real-time near real-time data streams of plant data.

Results
- Mainframe costs reduced
- Delivery of new solutions accelerated
- Platform established for future expansion
- Facilitate Innovation

Alight Solutions, a leader in technology-enabled health, wealth and human capital management solutions, recently embarked on an initiative to align the company's internal organization with its next-generation digital strategy. The strategy was adopted to reduce barriers to enterprise innovation and accelerate the pace at which Alight brings new digital products to the marketplace. As part of this initiative, Alight developed a low-latency, single view of the customer application using Confluent Platform. This solution is powered in part by Apache Kafka® with data from a wide variety of back-end systems, including mainframe and Java applications as well as cloud-based human capital management software.

"We partnered with Confluent to build what we've named our 'Unified Data Platform,' which is a forward cache of all of the various data sources that we use to service our customers," explains Chris Roberts, Vice President Enterprise Architecture at Alight. "As the primary source of individual consumer data for our customer-facing solutions, the Unified Data Platform is a foundational component to our future state architecture – one that is lowering costs, reducing complexity, increasing the speed at which we can deliver products, and creating new opportunities for innovation."

Challenges

Prior to developing the Unified Data Platform, Alight integrated back-end information within a web portal or a mobile app, for example. This approach had several drawbacks, particularly in the areas of performance, cost and time-to-market for new solutions. "When a user logged in to check their benefits, for example, our applications were reaching out to four or five different systems to get the data they needed," Roberts recalls. "The more systems we were integrating with at the time of the consumer interaction, the slower the system got. Aside from the performance issues, we also saw slowdowns in our pace of development due to the complexity of integrating with all these disparate systems. Lastly, because our mainframe systems represent a significant component of our budget, we were looking to reduce

the demand that we put into our z/OS solution by having it service a request every time a user logs in."

When Alight decided to implement a forward cache to address these issues, Roberts and his team focused first on identifying the right database for the cache and then pivoted to the next challenge – finding a scalable, secure way to get information into the cache when data changes. "Including our clients' employees and their families, we have about 70 million unique records on file throughout our various systems. So, data privacy and security are first and foremost in our minds. We needed to put the right data protections in place, including encryption of data at rest and in transit," says Roberts. "In addition, we needed to handle large-scale data changes based upon near-time or real-time events in the system so that we could refresh the cache on-demand with subsecond response times."

The team considered a range of technology options, including traditional queuing models and direct HTTPS calls, but none of these fully met the requirements the team had defined in terms of delivery assurance, performance or required skills. After further research, discussions with partners, vendors and others in the industry; and some proof-of-concept (POC) projects the team was ready to move forward with their choice: Apache Kafka. "The POCs helped us to validate that a Kafka solution could scale and get us to the right place," says Roberts. "At that point, we decided that if we're going to put this in as a key customer-facing component, then we had to ensure that we had the right support model in place from a data security standpoint as well as an active product support standpoint."

Solution

Alight used Confluent Platform, based on Apache Kafka, to build the streaming data pipelines that move data between systems and applications in the company's Unified Data Platform.

After taking time to design the necessary data models and schema, the Alight team worked with Confluent engineers to finalize their designs. "One of the things we did right on this project was work with Confluent early on," says Roberts. "We have folks here who have a really good head for data models and schema design, but it was a real positive to bring Confluent consultants in to validate what we had done and help us determine the right order for the work we would be doing."

Before jumping in with the implementation work, Alight development and operations teams attended onsite training provided by Confluent to accelerate ramp-up. "Another good decision we made was to schedule training early in the process so that everyone up and down the organization – from architect to developer – had a good sense for what they would be delivering," says Roberts. "We had four or five days of training for developers and for operations teams as well. That was instrumental in getting things off the ground in the right way, because after that no one was struggling or guessing what the best approach to take was."

Following the training, the team's focus shifted to getting Kafka up and running in the corporate data center, to the initial development of consumers and producers, and to set up the processes that govern message extraction and publishing.

Initially, development on the platform was limited to a small group of about six engineers, but it soon scaled out to a group ten times larger. "Having scaled out the system quite a bit, the rest of our platform organizations are now building the broader set of APIs based upon our forward caching model, and really deepening the information that's in that cache," says Roberts.

Part of that effort includes extending the platform to cloud services, such as the Salesforce.com Service Cloud. "To sync up information with the Salesforce.com Service Cloud, which is now a consumer of our back-end information, we are using Kafka to publish change events," Roberts explains.

Looking ahead, Alight has plans to expand its use of Confluent Platform beyond the forward cache. "This year and next we will be exploring an increase in event-driven business models, in which we will, for example, publish events to customers for mobile push messaging or refresh third-party partners with information as necessary," says Roberts. "As we pursue these goals, there are likely to be use cases for the KSQL, the streaming SQL engine for Kafka, connector management, and the Confluent Control Center."

The Unified Data Platform we built with Confluent has gotten our organization to think differently about how we deliver solutions – with a variety of layers and open APIs – that are more advanced than what we had been delivering. That shift in organizational thinking is what I'm most proud of.

—
Chris Roberts, Vice President Enterprise Architecture, Alight Solutions

Results

Mainframe costs reduced. "Our primary outsourcing delivery platform has a back-end COBOL-based application that runs on z/OS," says Roberts. "Our mainframe based environment is significant and we have to scale it up to meet peak demand during our peak season. The UDP platform enabled us to lower costs by offloading work to the forward cache and reducing demand on our mainframe systems."

Delivery of new solutions accelerated. "In talking with our business leaders about what we've built, one of the biggest potentials is in our ability to accelerate delivery of customer-facing solutions, which improves agility," Roberts says. "We're no longer seeing the same kind of scalability issues we were seeing previously and integration complexity has come down. System integrations can still be complex, but now they are centralized and done in a well-defined way, rather than project-by-project and consumer-by-consumer."

Platform established for future expansion. "Our teams like working on the platform; they know how to work with Kafka and we're integrating newer frameworks and engines on the back end," says Roberts. "In our environment, we already make use of Kafka and MongoDB, but we're introducing Docker, Spring Boot, Angular, and more, which our developers are definitely enjoying. Developer morale is high on the consumer side as well because they can use a single, well-defined API to instead of several different APIs for different systems."

Platform for Innovation. "With secure, consistent access to consumer data, we can now free our solution designers and developers to try out new delivery channels and ideas, such as incorporating consumer data based responses into Lisa, our chatbot platform. Without this level of consistent, constantly updated view, we would be unable to make these kinds of investments easily."

One of the things we did right on this project was work with Confluent early on.

—
Chris Roberts, Vice President Enterprise Architecture, Alight Solutions

Learn More About Alight

go.alight.com