

LADDERLIFE.COM

Ladder Puts Data in Motion to Underwrite Policies in Minutes



Headquarters

Palo Alto, CA

Industry

Life Insurance

Challenge

Apply machine learning to accelerate the life insurance underwriting process, enabling customers to apply for and activate policies within minutes if approved

Solution

Use Confluent to deploy a reliable, scalable, and low-maintenance event-driven architecture that streams data to machine learning models in near real time

Results

- Reliability maximized
- Focus on adding value sharpened
- Rapid ROI achieved
- Transition completed in minutes

Obtaining life insurance was—and in many cases still is—a weeks-long, arduous process involving in-person medical exams, paper forms, phone interviews, and brokers. Ladder has eliminated all of that with an online, direct-to-consumer, full-stack approach to life insurance that is powered by data and AI. In just five minutes, eligible customers can apply, get approved, and activate their policy for immediate coverage.

To improve the customer experience, make life insurance easier to get, and streamline the underwriting process, Ladder relies on a continuous flow of data from third-party providers to its AI underwriting engine. From the start, Ladder designed its data architecture around Apache Kafka® for these essential data flows. More recently, however, the company's explosive growth—more than quadrupling year over year—began to put a strain on this architecture and the team that supports it. To improve scalability and reliability while reducing administrative overhead, Ladder transitioned from self-managed Kafka to Confluent.

"The strength of the engineering team at Ladder really lies in making a great customer experience and automating business processes, rather than running a Kafka cluster. Putting our data in motion with Confluent allows us to focus on differentiating Ladder with automated underwriting built around machine learning models fed by real-time data."

— NICK HANSEN, SOFTWARE ENGINEER AND PLATFORM TEAM LEAD, LADDER

Now that Ladder has transitioned its event-driven architecture to Confluent, the company is positioned to continue scaling as it meets the rising demand for its approach to offering quality life insurance online. "Ladder's continued growth has been very exciting, and Confluent is going to provide a strong foundation for that growth," says Hansen. "Over time, there may be aspects of our infrastructure that we'll need to rethink or replace as we scale, but we know we're set with Confluent for a very long time."

Technical Solution

When a customer applies for term life insurance at Ladder, the company's automated systems contact third-party data providers and, with the customer's permission, access health records, driving records, and other information. That data is fed into machine learning models, which underwrite each policy using far more data points than a human underwriter could reasonably take into account. Because Ladder's business model requires all of this to be completed in minutes, Ladder engineers made data streaming with Kafka a core part of their data platform before the company launched its first product.

It was the company's growth surge that highlighted some of the limitations of Ladder's initial setup, which was a self-managed Kafka cluster deployed on Kubernetes running on AWS. Those limitations included reliability issues and administrative overhead, which would only grow more acute as the company scaled up.

"We've seen significant growth, starting with our launch and continuing on until now, and with that growth it's become more difficult and time-consuming to manage our own Kafka cluster, particularly since we don't have a dedicated DevOps team," says Hansen. "So we evaluated the fully managed options available, and Confluent was the natural choice for us."

Hansen and his team worked with Confluent engineers to plan the move from their self-managed cluster to Confluent using [Replicator](#). "We moved the producers first and let the consumers drain the existing topics, then we moved the consumers over," says Hansen. "It went amazingly well, and we were up and running on our production Confluent cluster in minutes."

To meet regulatory guidelines on protecting personally identifiable information, Ladder decided to encrypt all data flowing through Confluent, which they did by writing serializer and deserializer classes that they plugged into their existing producers and consumers. Now, all data flowing through Confluent—including customer information from third-party providers, user website interactions, unit test results, and more—is fully encrypted in transit.

Hansen's team uses [Confluent Control Center](#) to monitor their production cluster and track key performance indicators. "With Control Center, we can visualize what's going on and quickly diagnose issues," Hansen notes. "For example, we keep an eye on consumer lag metrics, because those are vital for our ability to underwrite and activate a policy in minutes. And, if we start to see lag we can drill down to find optimizations or improvements—sometimes as simple as adding more consumers—that can bring the lag back down."

Having Confluent as a common middle layer between producer and consumer services has greatly simplified integrations. "We have a lot of data sources and a lot of sinks, and we're always adding more," Hansen explains. "Without Confluent, that becomes an N-squared problem to integrate each source with each sink; with Confluent we just need N connectors and we're done."

Business Results

- **Reliability maximized.** "Since we moved to Confluent, we've had no service disruptions related to Kafka," says Hansen. "That reliability helped us achieve five nines last year, which builds trust in engineering within Ladder, and trust in our brand among our customers."
- **Focus on adding value sharpened.** "Because we're no longer spending so much time managing our Kafka cluster, we can focus our engineering resources on what's important to us as a business: improving the customer experience and improving our underwriting, which drives down costs for our customers," says Hansen.
- **Rapid ROI achieved.** "We performed a cost-benefit analysis that took into account how much we were spending to manage Kafka ourselves and how much we lost when it was down," says Hansen. "It was clear from that analysis that the move to Confluent had paid off quickly."
- **Transition completed in minutes.** "When we moved from self-managed Kafka to Confluent, we scheduled the switchover at night to be prepared for any issues," says Hansen. "It turned out not to be necessary; the whole process took all of five minutes, with no downtime for the user."

With a scalable, low-maintenance data streaming architecture in place, Ladder is now set to continue its robust growth. "Confluent is the underlying infrastructure that lets us orchestrate our business logic," says Hansen. "To be honest, our team mostly takes Kafka for granted these days. We're not worried about outages, and we know we have an architecture that is easy to work with and will scale as Ladder continues to grow."

"Confluent is the underlying infrastructure that lets us orchestrate our business logic. To be honest, our team mostly takes Kafka for granted these days. We're not worried about outages, and we know we have an architecture that is easy to work with and will scale as Ladder continues to grow." – NICK HANSEN, LADDER

Learn More About Ladder

ladderlife.com