

WWW.CERVED.COM

Cerved Expedites Strategic Decision-Making with Real-Time Data Access from Confluent



Headquarters

Milan, Italy

Industry

Financial Services

Challenge

Distributed databases using different technologies were causing slow query times and high latency as the system struggled to locate and process data at speed—which negatively impacted service delivery.

Solution

Cerved partnered with Confluent to provide instant access to business-critical data to empower strategic decisions and support.

Results

- Near-real time data propagation (event-driven architecture)
- Increased availability with Confluent’s managed service that’s helped build resilient systems with a 99.99% SLA
- Increased security with automatic monitoring and alarming
- Increased scalability with a higher number of projects and automated expansion of platform capacity
- Cost savings thanks to reduced infrastructure costs and governance, with 24/7 dedicated support from Confluent Cloud
- Disaster recovery that’s automatically updated with Cluster Linking

Technical Solution

- Advanced Stream Governance
- Schema Registry
- Connectors
 - Amazon S3 Sink Connector
 - Databricks Delta Lake Sink Connector

Cerved is an IT company and leading provider of credit information, market advice, and risk intelligence serving a variety of sectors. Over 30,000 businesses and 95% of Italian banks use Cerved to identify new strategies, protect themselves from risk, improve performance, and grow sustainably.

Data is the driving force behind Cerved's continued success. But when that data is split across multiple platforms that don't communicate, accessing key insights is time-consuming and wastes resources that could be better used elsewhere.

To provide a faster and more effective B2B service, Cerved needed a streaming platform that could handle huge amounts of data and form the basis for its new, connected ecosystem.

In 2017, Confluent provided Cerved with an on-premises platform to manage faster access to business-critical data at the touch of a button. The information services provided could now deliver strategic solutions and risk management support to the business in near real time.

Solving the primary challenge: difficulty decoding data

Cerved was looking to build a unique data ecosystem that could access public, Chamber of Commerce, web, and social data—alongside data from business partners like Experian and Assilea. These datasets provide their customers with critical information and insights to help them improve decision-making and gain a competitive advantage over their competitors.

“To give our customers the most relevant, up-to-date, and reliable advice, we need continual access to data. But as all of our data runs through different technologies, and we’ve got lots of applications communicating with each other simultaneously, it was a big issue to manage the entire infrastructure.”

— PASQUALE FOSSO, HEAD OF DATA ARCHITECTURE AT CERVED

But running one of the biggest data ecosystems in Italy isn't easy. Cerved has hundreds of databases, with millions of potential data points that teams need to locate, access, and transmit. Data has to be available on demand to scan, capture, identify, and interrogate every piece of information needed to paint a clear picture of financial risks and market opportunities.

Event-driven architecture

In a big push for digitization across the company, Cerved needed a better way to facilitate data access between its disparate databases and systems.

“Managing an event-driven architecture is always complex and can be pretty debilitating when it impacts service delivery. Thankfully, with Confluent, we have found the best solution for our event-driven approach even when accessing a large volume of data, which helps us deliver insights much faster,” said Fosso.

Cerved quickly realized many benefits using Confluent Platform, but managing on-premises clusters internally was too resource intensive. To accelerate new strategic projects and reduce operational costs, the company adopted [Confluent Cloud](#).

The migration to Confluent Cloud

The migration project started in the latter part of 2022, with the entire production environment deployed within six months. The process was managed by Cerved’s Data Platform team, which involved 12 dev teams across the Italian territory.

Overall, 200 applications, 789 topics, and 2,452 partitions with 1,107 schemas were migrated. In total there are now 40,000 connected clients with 2TB of storage.

The applications are based on languages such as Scala, Python, Java, and .NET with development/connection frameworks including Spring Boot, Kafka Streams, Spark, and Apache Kafka® connectors.

“Availability of data is crucial, and having a fully managed platform for the data distribution lets us provide solutions to our customers at speed,” explained Fosso.

Enjoying the results of migrating to Confluent Cloud

With a unique and consolidated data ecosystem using Confluent Cloud, Cerved has seen some great results, including:

Around-the-clock availability

With Confluent’s SaaS, Cerved has a network of resilient systems that are always on and continually adhere to the [99.99% SLA](#).

Strengthened security

Cerved uses Confluent Cloud to transfer large volumes of data in near real time and in a secure platform that implements encryption at rest on both the stored and in transit data for all the consumers of the platform.

Increased scalability

Thanks to the flexibility of Confluent’s managed service, Cerved is able to increase the number of new projects, without having to worry about infrastructure sizing, because the managed platform automatically adjusts any necessary expansions of the clusters.

With Confluent, this process is quick and simple, letting teams focus on what they do best: providing strategic B2B services to their customers.

Cost-savings thanks to fully managed support

With a fully managed Kafka service from Confluent Cloud, Cerved has entirely offloaded the operational burden which helps them reduce infrastructure and management expenses.

“With Confluent at the helm, we’ve cut back significantly on operating expenses, as we no longer need to waste time chasing data, managing capacity, or maintaining our growing Kafka deployment,” said Fosso.

Disaster recovery

Implementing a disaster recovery strategy with [Cluster Linking](#) increases the availability and reliability of mission-critical applications, helping to minimize data loss and downtime in the event of an unexpected disaster.

What’s next for Cerved?

Cerved can now access critical information across its ecosystem in a matter of seconds, which helps its growing team of consultants turn raw data into strategic insights and innovative ideas faster than ever before.

But what does the future hold for the business intelligence giant?

“In the coming months, we’re looking into the possibility of broadening our use cases of Confluent, as it’s become one of the most important pillars of our data mesh architecture,” said Fosso.

Learn More About Cerved: <https://www.cerved.com/>