axual

Data Governance & the Road to Compliance with Apache Kafka

How to create secure, controlled data access for all

Table of contents

DATA COMPLIANCE

Data is everybody's business

Digital transformation has made data "the new oil" of the 21st century economy. And the data that fuels us needs to be **secure**, **real-time** and **compliant**. That means every one of us — from smartphone owners to the Pentagon — must make data governance and compliance our business.

Protecting data means complying with legislation to **protect sensitive digital assets** by storing and managing them to safeguard against loss, theft, corruption and ransom. It makes sure that user access is tightly controlled to detect and thwart unauthorised attempts at entry.

Compliance starts with governance. It's your baseline. If you don't have the means to organise and control your data, you can never move on to the next step: complying with all applicable regulations.

WHY IS IT SO IMPORTANT?

Compliance is mandatory

Inadequate data protection leaves you vulnerable on three fronts: client litigation in the event of a breach, public prosecution for violating regulations, and ransomware attacks.

If the integrity of your data is at risk, it can also damage your reputation and even close your business. But good data management and proactive security measures aren't just about defence. They are also essential to secure trust between the security provider and the data's owner.

FINANCE

In the financial sector, simple compliance is rarely enough. Watchdogs keep a close eye on every move and consumers demand 100% security and reliability. Financial service providers can't afford to see their data protection compromised. Data integrity and data security strategies generally need to exceed the minimum requirements of compulsory compliance.

ENERGY

Companies in the energy sector are also legally compelled to comply with data protection laws, like GDPR and its dizzying array of rules for data storage. The transition to renewable energy demands insights into market and customer behaviour — insights that feed on data. Smart solutions allow you to safely handle and access ever-increasing data flows.

APACHE KAFKA

The gold standard in data streaming

Apache Kafka is an open-source distributed event streaming platform. It provides high-performance data pipelines, streaming analytics, data integration, and mission-critical applications.

With more and more companies using Kafka, particularly in the Finance and Energy sectors, it has established itself as the **de facto standard for data streaming today**.

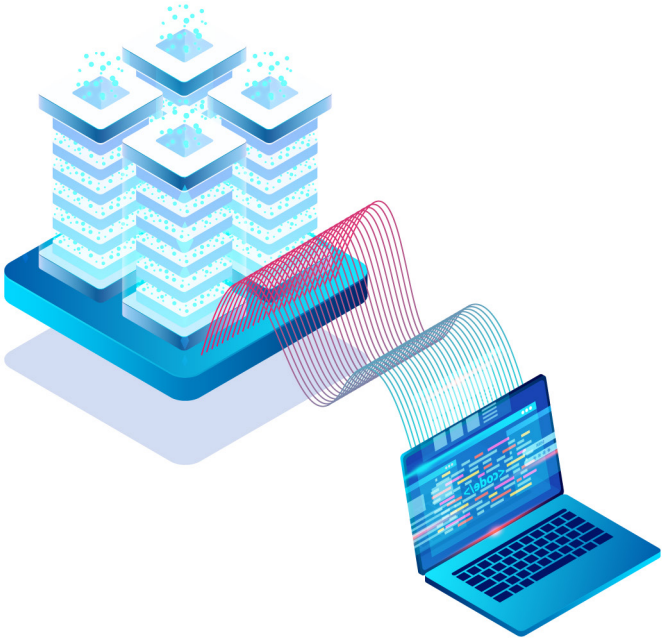

More than **80%** of **Fortune 100 companies** trust Kafka

70% in **Banking & Finance**

Rabobank Goldman Sachs BARCLAYS

60% in **Energy & Utilities**

Eneco

100% in **IT & Services**

ORACLE

cisco

salesforce

CHALLENGES

Challenges of using Kafka

When it comes to compliance, there are pain points any organization working with Kafka needs to overcome.

Who owns it?
Kafka Broker configurations all have their defaults, but some of those defaults don't really work in your favor. Auto topic creation is one of those examples. Nonexisting topics are automatically created without any check on completeness or correctness. A simple typo in the producer configuration can create a real mess. How can you clean it up? Does the topic hold sensitive data? Who can the consumer contact with questions?

Where did my topic go?
As easy it is to create a topic, it is just as simple to delete one. That is not such great news if you are relying on the topic data for your consumer application.

Which applications? Which permissions? Which topics?
By default, authorizing applications to access specific topics is not configured on Kafka. It means anyone can produce and consume to any topic. Depending on your needs, you might opt for Mutual SSL to secure your topics. But then you need to administer access by applications and principals.

This can quickly become an administrative nightmare for your platform team, who are forced to rely on administration scripts or poorly supported, unreliable open source tools — if there are any.

AXUAL'S SELF-SERVICE

Why self-service makes sense

Your DevOps teams need to focus on delivering functionality. They can't get their job done if they are continually occupied by compliance issues.

Axual's Self-Service provides a secure and effective solution for data governance. At the same time, it resolves many other issues by limiting the possibility of non-compliance.

Using a self-service interface is an efficient way to enforce a governance structure. It allows DevOps teams to freely interact with Kafka, while conforming to your own predetermined rules and practice.

By allowing developers to create and manage topics themselves, they can do their best work without relying on a central team — ensuring autonomy, efficiency and quick feedback.

GOVERNING DATA

Governance to the rescue

What organizations need is a structure of data governance that allows for **secure, controlled access** for Kafka topic administration — without losing essential **business agility**.

OUR SOLUTION

How we do it

Stream data and application ownership are key.
No Kafka topic is created before an owning team and metadata has been defined. Each team has the responsibility to selectively allow produce/consume access to the topic. A one-stop, easy-to-use, self-service interface maintains the big picture within your organization.

Security is at the heart of our platform.
Every single application must use SSL certificates in the connection to Apache Kafka. To prevent cross-environment data mangling, exactly one certificate is used per environment, which is validated to make sure only trusted connections are allowed.

Self-service offers agility.
Our Stream Team guarantees the streaming platform SLA for you. That means your DevOps teams can focus on configuring and controlling their topics and applications. No need for a central gatekeeper as the responsibility lies with the people who know most about the data.

Environment-specific rules allow precise customization.
Our streaming platform enables more fine-grained governance through environment-specific rules. On non-production environments, where no real customer data exists, the rules can be a bit more relaxed as compared to environments with sensitive data.

USE CASE

How Rabobank achieves Governance with Kafka

Challenges:

- Anyone can create and delete (oops!) a topic
- Anyone can produce and consume from any topic
- Extremely hard to continuously store secure data against threats and therefore to comply with e.g. GDPR.

Safely scale Kafka with Axual:

- As Kafka topics are being created by teams, they automatically become the owner
- Teams are in charge to allow produce/consume requests from any application to/from their topics
- Every single application uses secure and authenticated connections to Apache Kafka
- Teams are able to safely test different topic configurations or schema versions in a non-production environment, before going to production.

Business impact & ROI:

- Teams move faster due to self-service capabilities – no need to rely on a central data team
- Better sleep knowing your teams are compliant to GDPR and other regulations
- Saved €400,000 in maintenance costs – the maintenance team shrunk from 6 FTE to 1 FTE

About Rabobank:

Country: The Netherlands

Industry: Financial services

Size: 48,000 employees

Customers: Serving 9.5 million customers worldwide

Finances: €681 billion in assets

IT Cluster:

120+ DevOps Teams

400+ Registered Applications

180+ Data Streams in production

0.5 – 1B Messages sent per day

axual

Experience it yourself

Discover the all-in-one Kafka platform. Get started today with a free SaaS trial or request a personal demo.

REQUEST A DEMO