

Consolidate Kafka Infrastructure and Accelerate Migrations While Cutting Costs by 20-40%

Whether you're migrating to the cloud, consolidating sprawled clusters, or rightsizing your infrastructure, success requires visibility you don't have today: what you're running, who's using it, what it costs, and how to act on it without disruption.

Kafka sprawl has become the hidden tax on every data initiative.

What started as a few clusters has grown into a fragmented estate spread across OSS, MSK, and Confluent, with no visibility into what's connected, who owns what, or what it actually costs. We see this repeatedly across enterprises of all sizes: organizations paying for 8 clusters when they need 3, negotiating vendor contracts blind, and running migration projects that stall because no one can map the dependencies. The result is budgets that balloon, timelines that slip, and consolidation initiatives that never get off the ground.

01

Kafka costs grow without accountability or clear ROI

Infrastructure spend climbs year-over-year with no way to tie costs back to business value or justify the investment.

No chargeback model means platform teams can't attribute costs to business units or hold teams accountable for usage.

Engineers field constant questions about resource usage with no data to provide answers.

02

You know there's waste, but you can't safely act on it

Consolidation could cut costs by 20-40%, but the risk of destabilizing production makes it a non-starter.

Team isolation requires dedicated clusters, multiplying infrastructure costs even when utilization is low.

Cleaning up idle topics or reducing partitions is impossible without knowing who's consuming what.

03

Migration and consolidation projects stall or fail

Migration budgets balloon 30-50% as hidden dependencies surface, delaying time-to-value on strategic data initiatives.

30% of engineering capacity consumed by manual topology mapping and cross-team coordination.

Teams discover critical blockers during cutover, forcing rollbacks and restarts.

With Conduktor

Usage analytics and chargeback by team create accountability. Real consumption data gives you leverage in vendor negotiations.

With Conduktor

Virtual clusters enable logical isolation on shared infrastructure. Topic concentration reduces partition counts by 90%+.

With Conduktor

Automatic resource relationship discovery and transparent traffic routing enable migrations without app code changes.

IMPACT WITH CONDUKTOR

REDUCE

20-40%

Infrastructure Costs
Through consolidation & waste elimination

ACHIEVE

Zero

Downtime Cutover
With transparent traffic routing

CUT BY

50%

Migration Timeline
vs. manual coordination approach

OVER

\$500k

Year 1 Value
From cost savings, faster migration & efficiency gains

Leading European Airline

Consolidated 25 Clusters Across 3 Data Centers and Migrated to Cloud in 9 Months

A major European airline serving 100,000+ passengers daily faced a common challenge: Kafka sprawl across 25 clusters in three data centers with 170 dependent applications and no clear view of dependencies or ownership. Manual consolidation planning would have taken months and coordinating app teams for migration seemed impossible without disrupting mission-critical flight operations.

| | | | |
|-----------------------|-------------------|----------------|----------|
| 25 | 170 | 9 months | Zero |
| Clusters consolidated | Apps transitioned | Total timeline | Downtime |

Using Conduktor, the platform team mapped application dependencies automatically, identified consolidation opportunities, and sequenced migrations by complexity. Gateway enabled transparent traffic routing with no app code changes or team coordination required.

"The tools Conduktor provides made our migration seamless and the level of collaboration and support has been outstanding."

– Software Engineer, Leading European Airline

How We Do It

Conduktor provides the visibility to understand your Kafka estate and the control to consolidate and migrate without disruption.

Usage Analytics & Chargeback

- Track storage and partition usage with cost attribution by team to identify empty, stale, and underutilized topics.
- Understand where demand is coming from and drive accountability.

Virtual Clusters & Topic Concentration

- Deploy isolated logical clusters on shared infrastructure and consolidate low-volume topics to reduce partition counts by 90%+.
- Lower storage costs and broker CPU while maintaining logical separation.

Cross-Provider Resource Discovery

- Unified visibility across OSS Kafka, MSK, Confluent, and Aiven automatically indexes topics, consumer groups, connectors, and schemas.
- Plan consolidation and migration with complete data, not spreadsheets.

Transparent Traffic Routing

- Gateway sits between apps and Kafka, enabling cluster switches and consolidation without application changes.
- Seamlessly migrate applications, cluster-by-cluster.



3

Ready to consolidate and
optimize your Kafka estate?

Watch a 5 minute demo to see how
Conduktor reduces infrastructure
costs and accelerates migration.



• See it in action now