

# Messaging Infrastructure is Still in the Dark

## The Observability Illusion Costing Millions

Modern distributed applications heavily rely on robust messaging and streaming platforms like ActiveMQ and Kafka. While powerful, organizations face significant operational hurdles including performance bottlenecks, scalability complexities, reliability risks, and observability gaps.

Despite the widespread adoption of general observability platforms, the specialized nature of messaging infrastructure often leaves critical blind spots, leading to a dangerous “observability illusion.”

This article outlines these challenges, their severe business consequences, comprehensive strategies to overcome them, and explains how meshIQ offers a unique, unified solution to transform operations in these specialized domains from reactive firefighting to proactive engineering, significantly reducing the burden of manual issue investigation and ticket processing by empowering tenants with self-service capabilities.



by **Andrew Mallaband**

Growth Engineering | Enabling Tech Leaders  
& Innovators Around The Globe To Achieve  
Exceptional Results

# The 2:37 AM Wake-Up Call

Raj's phone buzzed aggressively on the nightstand. It was 2:37 AM. He blinked, groggy, and reached for it.

**ActiveMQ Broker 3 memory overload—service degradation.**



Seconds later, another alert.

**Kafka consumer lag exceeding threshold—real-time fraud detection delayed.**



He sat up, heart sinking. Not again.

Raj was the lead platform engineer at a major fintech company. He had lived through enough of these alerts to know what was coming. Just last week, a misconfigured Kafka retention policy had purged transaction logs prematurely, throwing the reconciliation process into chaos. Now, ActiveMQ was choking on a backlog of unprocessed messages, and Kafka was falling behind—again.

He scrambled to log in. The dashboard was a mess: memory exhaustion, disk I/O saturation, spiking consumer lag. Their real-time fraud detection system, which was supposed to block stolen cards instantly, was now minutes behind. Minutes too late.

By the time Raj and his team stabilized the system, the damage was done. Over a quarter-million dollars in fraudulent transactions had slipped through.

When he finally sat back in his chair, exhausted, a message from his VP was waiting:

***“How did this happen? We built this to be resilient.”***

Raj knew the answer. They had the right technologies—ActiveMQ for transactional reliability and Kafka for real-time streaming—but operationally, the complexity of managing these solutions meant they were flying blind. This lack of transparency often forced dependent teams to raise tickets when issues arose, leaving Raj's team constantly reacting to problems they couldn't preempt.



# The Hidden Costs of “Set It and Forget It”

## and the Observability Illusion

When Raj’s team first deployed ActiveMQ and Kafka, they did so with confidence. ActiveMQ handled the company’s core banking transactions with robust JMS guarantees. Kafka powered real-time fraud detection and analytics, streaming millions of events per second. At first, it all seemed to work. But as the business grew, cracks started to show, revealing the substantial underlying costs of managing these “free” open-source tools.

Performance issues cropped up unexpectedly. ActiveMQ brokers would crash under memory pressure due to large message backlogs or improper JVM configuration. Tuning JVM heap sizes helped for a while, but the problems kept resurfacing. Kafka’s much-touted low latency began to degrade, often because disk I/O couldn’t keep up with the pace of incoming messages, leading to message throughput drops. Consumer lag grew—sometimes subtly, sometimes catastrophically

*“We thought scaling was just a matter of adding more brokers,” Raj later reflected.*

*“But traffic skew led to some Kafka partitions getting overloaded, while others sat idle. Consumers started falling behind, and we didn’t know why.”*



**It quickly became clear that scalability wasn’t automatic.** ActiveMQ primarily scales vertically, and its master-slave architecture worked—until a shared SAN failure took down an entire cluster, exposing its limitations for hyper-scale scenarios. Kafka’s partitioning model offered theoretical parallelism, but in practice, poor distribution created hotspots that degraded performance, coupled with the overhead of rebalancing. To compensate, they over-provisioned. It seemed like a safe bet, but the cloud bill ballooned by 40% almost overnight due to resource-intensive scaling.

Then came the data integrity issues—small misconfigurations with massive consequences. A single Kafka topic with cleanup.policy=compact mistakenly applied wiped older transaction records. Incorrect producer configurations (e.g., acks=0) led to message loss. An ActiveMQ broker crash resulted in the loss of in-flight messages that hadn’t yet been persisted. These weren’t one-off incidents; they were signs of a fragile system, leading to eroded data integrity, potential non-compliance, and severe reputational damage.

Even their existing general-purpose observability tools failed them. While they had invested in comprehensive commercial platforms and OpenTelemetry, these tools, while excellent for infrastructure and application monitoring, lacked the deep, specialized context needed for messaging and streaming. Dashboards showed broker status and system health, but lacked the granular insights into queue depths, message flow patterns, consumer group rebalances, or the nuanced interplay of JMS settings that are unique to these domains.

**The business impact of technical issues wasn’t truly visible. Meanwhile, the team was drowning in alerts—hundreds of them each day.** Most were noise, contributing to alert fatigue. The ones that mattered were often buried too deep to catch in time, resulting in longer Mean Time To Recovery (MTTR).

This lack of actionable, specialized insights meant that application teams, facing issues with their services, had no choice but to raise support tickets, forcing Raj’s team into constant reactive investigations, spending valuable time deciphering generic logs and metrics. This constant influx of tickets represented significant toil for the platform engineering team, diverting their attention from strategic initiatives to reactive firefighting and consuming expensive, highly skilled resources.

# From Firefighting to Forward-Thinking

Things came to a head when the CFO cornered Raj in a meeting and asked:

*“Why do we keep losing money on preventable issues? And why is our real-time data always behind?”*

Raj didn't have a good answer. But he knew they couldn't keep going like this.

He decided it was time to stop reacting and start engineering for reliability.

The first step was tuning their systems like they mattered—because they did. On ActiveMQ, they disabled unnecessary message IDs/timestamps to reduce overhead, optimized message processing efficiency with prefetch limits and asynchronous dispatch, and managed persistence effectively.

Kafka was similarly overhauled: they optimized broker configurations (partition count, segment size), tuned producer settings (message/batch size, compression, acks levels), fine-tuned consumer settings (fetch sizes, consumer instances), and ensured optimal hardware and network infrastructure.

*“We stopped treating it like a black box,”* Raj said.

*“We started tuning it like the mission-critical system it was.”*

Next, they tackled availability and scalability. ActiveMQ was upgraded leveraging its built-in replication and replicated message store, eliminating their dependency on a fragile shared SAN and enabling robust failover. Kafka was hardened with rack-aware replication and a strict `min.insync.replicas=2` policy, ensuring durability even in the event of a broker failure, along with effective consumer group management.

But perhaps most importantly, they shifted how they measured success. Consumer lag was no longer just a metric—it became a critical Key Performance Indicator (KPI). They rewired their alerts to focus on early indicators like memory saturation and I/O bottlenecks, giving them the chance to act before systems failed.

They also explored automation and self-service capabilities to streamline routine tasks, provision resources, and empower tenants and stakeholders with direct, actionable insights into the health and performance of their dependent services, aiming to reduce the flood of reactive tickets and free up precious platform engineering time.

The difference was immediate. But while the systems ran better, managing them still required significant time and attention.

**That's when they discovered meshIQ.**

# Enter meshIQ

## Operational Intelligence for Specialized Challenges

Raj was skeptical at first. They'd tried plenty of “single pane of glass” tools before—most were underwhelming. But meshIQ was different. It wasn't another general observability tool; it was purpose-built for the intricacies of messaging and streaming. **For the first time, they had unified, real-time observability across both Kafka and ActiveMQ from a single, coherent dashboard.** meshIQ didn't just show generic system metrics—it provided deep, specialized context unique to messaging environments. They gained unparalleled visibility into queue depths, message delivery rates, consumer group health, and the granular details of message flow that their existing observability tools simply couldn't provide.

Its AI/ML-Driven Insights leveraged machine learning to predict bottlenecks and outages, enabling proactive issue identification and resolution across complex topologies, highlighting anomalies specific to messaging patterns. This directly addressed their “observability blind spot” by providing end-to-end data flow visibility and translating technical performance into clear business outcomes.

Crucially, meshIQ enabled self-service by empowering tenants and stakeholders with direct access to these vital, specialized insights into their specific services. This dramatically reduced the need for them to raise tickets, enabling teams to self-diagnose and resolve many issues without platform team intervention, saving Raj's team countless hours of investigation and reducing toil. Importantly, in the transformed state, self-service could also include optional approval steps for configuration changes—ensuring that while teams move faster, risk is still governed and mitigated through controlled workflows.

*“It wasn't just visibility,” Raj explained. “It was clarity. We could see lag, disk utilization, throughput, everything—and understand what it meant for the business, and our application teams could see it too, with the specific messaging context they needed. This meant fewer tickets for us to triage, and more time for strategic work.”*

meshIQ also brought streamlined management and automation into the fold. Kafka partitions were rebalanced dynamically, eliminating the manual sharding work that had cost the team hours each week.

Developers could now provision queues and topics on their own using a single API and centralized access control, unblocking delivery cycles and removing ops bottlenecks. This self-service capability extended beyond just provisioning, offering stakeholders direct views and simplified control over their messaging resources, with the ability to define policies that controlled the scope of what they could do within the services they were responsible for. This proper partitioning of insights and control drove much greater efficiencies, acknowledging that platform engineering resources are very expensive and often have many calls on their time. meshIQ also provided comprehensive audit trails for all modifications.

The expert support experience was another game changer. Instead of Googling obscure forum posts at 3 AM, Raj's team had access to 24x7 technical support from ActiveMQ and Kafka experts. They shared tuning templates, best practices, and insights grounded in real-world experience.

The results spoke for themselves. Within six months, they had experienced zero message loss incidents. Fraud detection latency dropped from minutes to milliseconds. Their cloud bill was 30% lower, thanks to right-sizing and smarter scaling decisions, leading to reported TCO reductions of up to 50% for Kafka deployments compared to traditional vendors.

meshIQ even facilitated seamless, zero-downtime migrations when needed. And perhaps most significantly, the volume of reactive support tickets from dependent teams plummeted, freeing Raj's team to focus on strategic initiatives rather than constant firefighting, proving the immense value of specialized self-service capabilities.

His VP, once skeptical, sent a short message after the latest quarterly report:

*“Whatever you did—keep doing it.”*



# Epilogue: The Choice Is Yours

Raj's story isn't unique. It's playing out in quiet, chaotic ways across every industry that relies on messaging infrastructure. The perception that general observability tools can fully address specialized messaging challenges is a dangerous one, often leading to a false sense of security. The burden of reactive ticket management due to a lack of delegated insights disproportionately impacts highly valuable platform engineering teams.

Maybe your system is already showing signs—rising lag, growing cloud bills, inconsistent throughput, or a gradual erosion of data integrity. Or maybe you've already had your 2:37 AM moment, and your platform team is overwhelmed by support tickets because dependent teams lack the specialized insights and self-service capabilities to help themselves.



Before you dismiss this as “just another story,” ask yourself:

- 01 How confident am I that my messaging layer is truly resilient?
- 02 Could my teams resolve more issues themselves if they had the right specialized insights?
- 03 How much engineering time am I losing each month to reactive firefighting?
- 04 Are my scaling decisions data-driven, or based on educated guesswork?
- 05 If a critical broker failed tonight, how quickly would we really recover—and at what cost?



You don't have to keep firefighting.

## meshIQ offers a new path:



### UNIFIED, SPECIALIZED OBSERVABILITY

across Kafka, ActiveMQ, and more, with AI/ML-driven insights for proactive issue detection.



### EMPOWERING SELF-SERVICE

for tenants/stakeholders with granular insights into their services and controlled provisioning capabilities, significantly reducing toil for platform teams and freeing up expensive engineering resources.



### OPTIONAL APPROVAL WORKFLOWS

for configuration changes, balancing speed with governance.



### PROACTIVE AUTOMATION

that removes manual toil, streamlines provisioning, and simplifies scaling guesswork.



### ENTERPRISE-GRADE RELIABILITY

backed by deep expertise and proven cost savings.

You can keep hoping your next outage isn't the big one—or **you can take control and engineer for resilience with tools built for the complexities you face.**

**Because in modern operations, resilience isn't built on luck. It's engineered.**

TALK TO AN EXPERT

BOOK A DEMO

