elastic

CUSTOMER ANONYMIZED

CUSTOMER STORY: SEARCH

Financial services · CRM search · Hadoop batch → Elastic

Real-time CRM search at 20,000 events per second: A leading Indian private bank leaves Hadoop batch behind for Elasticsearch

One of India's top four private banks kept Hadoop for batch analytics and moved real-time CRM and transaction search onto three workload-isolated Elasticsearch clusters, with Kibana as one front end for business and IT.

20,000/sec
Real-time CRM search, not batch

750 GB/day
Application logs searchable as they land

3 clusters
Workload-isolated, beside Hadoop

1 platform
One Kibana layer for business and IT teams

The challenge

- ✕ Hadoop batch could not answer a relationship manager's customer search the moment it was made.
- ✕ Application logs were not surfaced as issues happened, leaving no real-time view of service health.
- ✕ CRM data sat in Hadoop and application data in separate databases, so teams saw different pictures.
- ✕ Customer-facing systems had no disaster recovery, and tuning outpaced what the team could handle.

The solution

- ✓ Three purpose-built Elasticsearch clusters, each sized for one workload, run beside the Hadoop lake.
- ✓ The search cluster feeds the CRM directly at 20,000 events per second; transaction search runs at 12,000.
- ✓ The log cluster indexes 750 GB a day of microservice logs as they land.
- ✓ Kibana is the shared front end for business and IT, delivered with managed-service partner Ashnik.

How it works

1. Collect
Fluentd gathers logs from Kubernetes

2. Ingest
Logstash and custom apps feed Elasticsearch

3. Isolate
Search, transaction, and log clusters

4. Explore
Kibana serves relationship managers and IT

Before

- Customer360 lookups ran on Hadoop batch, too slow for live use
- Application logs not surfaced as issues happened
- Business and IT data in separate silos, no shared view

After

- → Real-time profile and transaction search at 20,000 events per second
- → 750 GB a day ingested and searchable as it lands
- → One Kibana layer, one source of truth for both teams

> "Log ingest and RM search pull a cluster in opposite directions. One needs sustained indexing; the other needs **predictable query latency at 20,000 events per second**."
>
> — Santanu Patel, Technical Consultant, Ashnik

Why it matters beyond one cluster

The right architecture gives every tool the job it was built for: Hadoop keeps the batch analytics, and Elasticsearch handles real-time search. Disaster recovery is the next priority, and workload isolation lets each cluster be protected on its own terms. New workloads land on an appropriately sized cluster without a platform-wide redesign.

See how Elasticsearch runs real-time search at scale alongside your data lake or start a free trial →

 © 2026 elasticsearch B.V. All rights reserved.