

A Technology Provider Scaled Up to 75 Terabytes a Day Without Toil

Key Takeaways



Increased team productivity 2x



Reduced logging costs by 90%



Improved margins from 20% to 90%

The Challenges

Before engaging Era Software, the customer, a technology company, experienced challenges to growing their business due to the high cost of running Elasticsearch clusters at scale. With many disks to track and outages to worry about, they were constantly stressed about maintaining their Elasticsearch system. For their whole process of ingesting, storing, and querying log data, they required about 10,000 CPUs, using a combination of expensive solid-state drives and spinning disk drives (cheaper, but much slower performance).

Their systems were bombarded with around 20,000 requests/sec, scaling up to 50,000 requests/sec at times. The data payload would also vary. During an outage, log volumes would increase very quickly. The customer realized they had to transition to a modern architecture optimized for log management to support their use case and sustain their growth.

The Solution

In early 2021, the customer began ingesting 25% of their data into EraSearch. Observing their environment, Era Software's team was quickly able to optimize the CPUs for ingesting and indexing data. As the customer migrated their data to EraSearch, Era Software's customer success team worked with their engineering team to optimize for cost and performance.

Using EraSearch, the combination of high ingestion and fast queries and decrease in disk utilization allows the customer to use slow-performing disks. For a much smaller footprint, they're getting high performance over their entire range of data. Because EraSearch doesn't require hot replicas, it's cost-effective for the customer to run all their data on hot nodes, as well as to keep a full copy in a cold object store. Now, they don't have to worry about what's hot and what's not.

Their journey scaling to 75 terabytes/day

With fewer disks to monitor, their engineers can fine-tune the disks to optimize performance and cost. For their use case, we discovered that the less expensive disk delivers sufficient performance. Our customer success engineering team also helped them identify disks that weren't performing well so they could swap them out to improve performance. And in their new cloud environment where they're running sixteen c5a.16xlarge EC2 instances, they can adjust disk performance dynamically.

Learn More

TRY ONLINE
cloud.era.co/signup

LEARN MORE
era.co

The customer's ability to scale to 75 TB/day and more is enabled by EraSearch's innovative architecture and these design principles:

- **Zero-schema design** enables ultra-fast data ingestion and searchability
- **Decoupled storage and compute** allows users to pay for only infrastructure resources they consume without dedicated processing and storage resources
- **Object Storage** provides low-cost storage for cold data

EraSearch's multi-tier design allows scaling of each tier independently. The indexing tier is the most CPU intensive. This tier takes in the 20,000 requests/sec, with log data payloads of various sizes, and indexes the data quickly before sending the data along to be stored and queried immediately. By optimizing this tier, the customer is able to use 10x fewer CPUs (from 10,000 to 1000) to handle both ingestion and querying of 75 terabytes/day — a significant improvement.

The Results

The customer's transition to EraSearch has led to increased engineering productivity and improved margins. By reducing operations toil with EraSearch, the customer's business reaps the benefits of:

- **Team Productivity.** They were able to reduce the number of engineers dedicated to maintaining observability from 12 to 2.
- **Improved margins.** Reducing logging costs by 90% helped improve their margins from 20% to 90%.
- **A new business model.** Their engineers now have more time to innovate.

The technology provider's success is evident from its reduction of operations toil and its ability to scale to 75 terabytes a day.

Summary

Before transitioning to EraSearch, the customer faced many challenges growing their business due to the high cost of operating Elasticsearch clusters at scale. In order to sustain their growth, they knew they needed to transition to a modern architecture for log management quickly. With EraSearch, the customer has reduced operations toil, resulting in increased productivity and improved margins. Now, their business and technical decision makers can focus on introducing a new business model to generate new revenue streams.