



THE CHAOSSEARCH DATA LAKE PLATFORM

Scalable Log Analytics for Application Troubleshooting and DevOps Efficiency

CHALLENGE

APM and Observability Tools: Incomplete Log Coverage and Cost-Prohibitive at Scale

Application performance monitoring (APM) and observability tools ingest logs, traces, and metrics and present a near real-time view of application behavior in feature-rich dashboards. DevOps and software development teams rely on them to monitor and troubleshoot application performance. However, the rapid adoption of containers and microservices architecture—and the corresponding explosion in the volume of log data—has exposed the shortcomings of these ubiquitous tools for log analytics at scale:

Limited Data Retention

APM/Observability tools routinely cap log data which inhibits historical and long-term analysis. APM/observability tools are best for short-term (<30 days) of operational data, not large-scale logs. This creates analytical blind spots.

Data Movement and Multiple Data Copies Dramatically Increases Costs

You could be paying for the same data multiple times. You may be paying to store your data, send it to your APM/Observability tool (egress), and again for data ingest. Plus, data movement can also violate compliance mandates.

High Storage Costs

APM/Observability tools add metadata to aid in analysis, but this inflates data size. And parsing logs into fields also increases data size. This increases storage costs or forces log retention trade-offs.

Not Optimized for Logs

The database backend for many APM/Observability solutions (typically Lucene) is optimized for telemetry data such as metrics and events. High log volume and searches across large data sets often cause system bottlenecks and slow query response.

A NEW APPROACH TO LOG ANALYTICS

Imagine sending all your data to your cloud environment in its native format—no parsing or schema changes. ChaosSearch indexes all data as-is, without transformation, while auto-detecting native schemas. The ChaosSearch Data Lake Platform enables the search and analysis of cloud data with unmatched scale, resiliency, and cost savings.



SOLUTION

Scalable Log Analytics for Application Troubleshooting and DevOps Efficiency

ChaosSearch empowers customers to Know Better™, activating the data lake for analytics. Unlike APM/Observability solutions, ChaosSearch indexes all log data in your cloud object storage, as-is, without the need for any data transformation or data movement. Our revolutionary approach supports the cloud environments' scalability and economic advantages so organizations can retain more log data to fine-tune container and application troubleshooting and optimize the user experience.

CHAOSSEARCH BENEFITS FOR DEVOPS, ARCHITECTS, AND APP TROUBLESHOOTING TEAMS

Unlimited Data Retention

ChaosSearch treats your cloud object storage (AWS/S3, GCP/CS) as first-class citizens allowing unlimited retention and scalability at much lower cost. Troubleshooting teams have better log coverage to identify performance issues and optimize the user experience.

No Data Movement

ChaosSearch indexes the data in your cloud environment while auto-detecting native data schemas—no complex pipelines, no data transformation, or parsing on ingest. You also retain complete control over your data with full RBAC for security and compliance.

Cost-Performance at Scale

With a virtually unlimited total capacity, you can maintain long-term data retention, avoiding the painful “cost vs. retention” tradeoffs prevalent with APM/Observability tools. Our data ingest rate leverages the full power of the cloud without introducing any latency. Process tens of terabytes per day!

High Performance Queries at Scale

The patented Chaos Cache™ is a non-memory caching system that returns queries in seconds at Petabyte scale. Chaos Cache allows for infinite storage for queries and results.

“ChaosSearch has been a right-fit solution for us. We have a reliable platform that “just works” without our intervention or care—with incredible support to back it. This is one of the few products that simultaneously delivers real cost savings with real increases in productivity.”

Jason Standiford, VP Engineering

REVINATE

ABOUT CHAOSSEARCH

ChaosSearch empowers data-driven businesses like Blackboard, Equifax, and Klarna to Know Better™, delivering data insights at scale while fulfilling the true promise of data lake economics. The ChaosSearch Data Lake Platform indexes a customer's cloud data, rendering it fully searchable and enabling data analytics at scale with massive reductions of time, cost and complexity. The Boston-based company raised \$40M Series B in December 2020 and is hiring to support its hyper growth.

For more information, visit ChaosSearch.io or follow us on [@ChaosSearch](https://twitter.com/ChaosSearch) and [LinkedIn](https://www.linkedin.com/company/chaossearch).