



THE CHAOSSEARCH DATA LAKE PLATFORM

Activate your Amazon S3 Data Lake

CHALLENGE

Log Volume and Velocity Limit Insight at Scale

As log volume grows, existing log analytics solutions become difficult to manage and expensive to scale. Organizations must dedicate full time employees (FTE) to manage complex pipelines, transform data for analysis, and manually configure and manage the complex compute and storage clusters that comprise traditional log analytics architectures. To reduce costs and administrative toil, customers limit data ingest and shorten retention, limiting their ability to extract timely and meaningful insights.

SOLUTION

Scalable Log Analytics for Efficient Cloud Operations

ChaosSearch empowers you to Know Better™, activating the data lake for analytics. The ChaosSearch Data Lake Platform connects to and indexes data within your Amazon S3, enabling search, SQL and machine learning workloads with infinite scale, lower cost, no data movement, and faster time to insights. Whereas other solutions require complex data pipelines consisting of parsing or schema changes, ChaosSearch indexes all data as-is, without transformation, while optimizing for both data size and performance. Through open APIs, users continue working in the tools they know and trust, such as Kibana, Elastic, Looker, and Tableau. ChaosSearch eliminates complexity, unlocks data retention limits, and reduces costs up to 80%, delivered in a fully managed service

Cost vs. Data Retention Trade-Offs

Organizations are often forced to make trade-offs in log retention due to the high cost and unreliability of storing log data at scale. They face the choice of throttling log ingest or shortening retention. In either case, you are constantly reducing visibility and limiting insights to save money.

Containers and Microservices Architectures Complicate Troubleshooting

With the growth of containers and microservices architecture, trimming the number of logs analyzed to reduce costs results in analytical blindspots that exacerbate performance and availability issues. Architects and engineers can't create a complete view of all the data they need to troubleshoot and optimize cloud services and infrastructure.

Management Overhead & Complexity

An Elasticsearch cluster—whether managed in the cloud or on-premises— becomes brittle at scale and requires active, hands-on administration. Compute and storage are tightly coupled resulting in cost spikes or failures when log volumes inevitably rise.

Complex Data Pipelines and Movement

Legacy analytics solutions require complex data pipelines consisting of parsing or schema changes. Organizations must dedicate full time employees (FTE) to create and manage pipelines and transform data for analysis. Data is often duplicated, which adds cost and increases risk of security and compliance violations.



ACTIVATING YOUR AMAZON S3 DATA LAKE

Imagine sending all your data to Amazon S3 in its native format—no parsing or schema changes required. ChaosSearch indexes all data as-is, without transformation, while auto-detecting native schemas. Architects, developers, and engineers can improve log coverage and ensure the stability and agility of cloud applications, infrastructure, and services.



Eliminate the Cost vs. Retention Trade-Off

Without the need to sacrifice data ingest or retention, ChaosSearch lets you keep more log data from multiple systems in a single repository—your Amazon S3. CloudOps, DevOps, and SecOps teams work with the same data to collaborate more effectively, isolate issues faster, and identify trends over time that can impact the business.

Optimize Microservices Architecture

Architects and engineers get full visibility across containers and microservices for better data analysis and troubleshooting. ChaosSearch transforms your Amazon S3 into a hot analytics platform with no data movement or ETL required. You'll be able to keep more log data from multiple systems in a single repository. All your data is in your Amazon S3 for as long as you need—at a fraction of the cost of existing solutions.

Put an End to Management Headaches

ChaosSearch is a fully managed service that eliminates the need to add, deploy, and manage hardware and software. ChaosSearch automatically scales on the fly so you can build your cloud applications and services without worrying about managing your logs. And it's easy to get started. Just land your logs in Amazon S3, connect ChaosSearch, and leverage our published Elasticsearch API/Kibana UI. You eliminate management overhead and reduce costs without making your users change their behavior or the tools they use.

Eliminate Data Movement

Whereas other solutions require complex data pipelines consisting of parsing or schema changes, ChaosSearch indexes all data as-is, without transformation, while optimizing for both data size and performance. All the data is read from and kept in customers' buckets and under their control. By combining Amazon S3's advanced Identify and Access Management (IAM) functionality and roles with a Role Based Access Control (RBAC) overlay, ChaosSearch customers enjoy exceptional security and granular data access control.

“Until now, the data has been a trickle. Now with ChaosSearch, we're ready to turn on the firehouse and help our customers, for efficacy and driving growth.”

Mark Hill, Director
IT Operations

Digital River®

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, find us on the AWS Marketplace, or follow us on Twitter @ChaosSearch and LinkedIn.

info@chaossearch.com | www.chaossearch.io