



Equifax Simplifies Cloud Operations Management Globally with Cloud Data Platform

OPERATIONAL ANALYTICS AT EQUIFAX

As a global data, analytics and technology company, Equifax plays a critical role in helping employers, government agencies and financial institutions make decisions with greater confidence. Armed with Equifax insights, businesses can provide a seamless and positive experience during consumers' pivotal moments—such as when they're applying for a job or mortgage, financing an education, or buying a car.

Since Equifax processes sensitive financial data, it was critical for its Engineering team to deploy a secure operational analytics solution for internal stakeholders. Each business unit had unique use cases for analytics, such as:



Meeting customer service level agreements



Monitoring and troubleshooting application and infrastructure



Ensuring compliance to internal and regulatory requirements

The common thread across them all? The need for fast, secure, reliable, and diverse access to analytics at scale.

Equifax turned to ChaosSearch® to provide a highly scalable, robust and secure Cloud Data Platform that centralizes data management and simplifies access to analytics. The platform is available as a fully managed SaaS solution connected to data within the company's existing cloud object storage environment. This accessibility of data across cloud providers and environments reduced administrative overhead and data silos, resulting in significant cost savings over the previous deployments. Now, the team can cost-effectively retain important business data for longer periods of time, which helps drive revenue-generating decisions and initiatives.



Industry

Financial Services

Location

Global

Use Cases

- Application and infrastructure performance troubleshooting
- Compliance
- Customer service insights

Impact

- 90% cost savings over legacy stack tools
- No data movement or transformation needed for full text search or data analytics
- High availability and performance across VPC regions and zones

Scale

- 50 TB of data ingested per day

THE CHALLENGE: SCALING ANALYTICS AMIDST STRINGENT OPERATIONAL REQUIREMENTS

Equifax is entrusted with personal information on millions of consumers across 24 countries to assist consumers as they work to achieve financial goals. Equifax has built its business on security, safety, and integrity.

In recent years, as Equifax implemented a best-in-class, multi-cloud environment, a key objective was to facilitate data democratization. In other words, Equifax needed to support diverse data analysis across multiple internal stakeholders—including a variety of departments and roles such as Customer Service, Security, Compliance, and DevOps—spanning numerous regions and availability zones, all while continuing to surpass the highest security standards.

In particular, managing log data at scale was becoming complex and expensive. Business units were spending money on centrally managed logging solutions—which, in many cases, manifested as siloed stacks within individual organizations' available cloud services.

Over time, this resulted in:



Administrative complexity for the Engineering team as they tried to support many siloed deployments.



Operational knowledge gaps and inconsistent analytics processes, tooling, and data sources across stakeholder groups; this prevented Equifax from synthesizing a holistic view across all metrics and from achieving economies of scale.



Fast-growing costs that did not scale in proportion to business growth

ANALYTICS USE CASES, STARTING WITH LOG DATA

The importance of log data to Equifax's business paired with the challenges they already faced managing log data at scale made it a no-brainer to start the data democratization journey there. Numerous business critical functions depended on log analytics at scale.

For example, to continuously improve Equifax products and resolve issues quickly, Equifax Engineering teams need to analyze product data to understand operational insights such as usage, performance and availability. These insights help determine service health and availability, and provide the team with the ability to develop new features quickly or conduct ad-hoc troubleshooting.

Meanwhile, site reliability engineers (SREs) need log analytics for container management and insights, as well as platform health and troubleshooting. Compliance teams regularly need to parse large volumes of reference, transactional, operational and compliance data.

For all of these functions, the time and effort required to analyze this data at scale was quickly becoming unmanageable. Each stakeholder relied on data engineers to move, transform, and cleanse the data before it was available for analytical consumption. This resulted in lagging response times, which became especially problematic when resolving critical platform or service issues. The team needed to achieve a much faster mean time to resolution (MTTR).

“ Taking a piecemeal approach to data analytics got expensive quickly, and it became difficult to pull insights together to paint a broader picture that spanned knowledge silos. That's where I saw the ChaosSearch opportunity. It's a much more ubiquitous platform we can use to access data across disparate systems, at scale.”

Jeff Kinsharf, SVP

Engineering Services and SRE
Equifax

Users

- Site Reliability Engineering
- Engineering Teams
- Customer Service Operations
- Compliance

Data Management Environment

- Cloud Data Platform: ChaosSearch
- Cloud Object Storage: Amazon S3 and Google Cloud Storage
- Analytic Tool(s): Kibana



SOLUTION: THE CHAOSSEARCH® CLOUD DATA PLATFORM

When evaluating analytics solutions, the Equifax Engineering team knew they needed consistency across all stakeholder groups, while maintaining the security and compliance measures required for operating at global scale. They sought a cloud-native platform that could provide high availability across regions and zones around the globe. And each Equifax business unit needed the capability to create their own data lake by storing logs in cloud object storage (such as Google Cloud Storage).

“*With ChaosSearch, we no longer have to move or transform our data. That means we don't have to think about how we're going to use the data before we access it. It's just there for us to query. That helps us deliver new opportunities to the business, or discover usage patterns in our systems.*”

Jeff Kinshurf, SVP

Engineering Services and SRE
Equifax

The ChaosSearch Cloud Data Platform—which activates the data lake for analytics by indexing cloud data—uniquely aligned with Equifax's needs. ChaosSearch was purpose-built for cost-effective, highly scalable analytics encompassing full text search, SQL and machine learning capabilities in one unified offering. The patented ChaosSearch technology instantly transforms cloud object storage into a hot, analytical data lake.

Equifax decided to move forward with ChaosSearch, starting with log analytics as the first use case, to resolve the fragility and untenable maintenance costs of its current providers at scale.

Using ChaosSearch, each business unit within Equifax can now ingest, index and store logs and other business-critical data in a regional cloud object storage aggregation bucket. Each Equifax business unit owns the account or project associated with its cloud object storage bucket. Read-only access is given to the ChaosSearch service to ensure security and compliance, while providing fast, highly available analytics capabilities to key stakeholders.

The ChaosSearch platform gives users an intuitive web user interface, making it much easier to analyze data by creating indexes or object groups. Using its embedded Kibana interface, users can query this data, or create visualizations and dashboards to view trends over time. This removes the need to rely on a data engineer to model or cleanse the data before consumption, speeding up the time to insight.

By using role-based access control (RBAC) groups, the Equifax team can ensure that analytics access is limited to only users with permission to view specific log data and related Kibana artifacts.



THE IMPACT: SECURE DATA ANALYTICS WITH ECONOMIES OF SCALE

With ChaosSearch, Equifax finally has a single access pane for analytics across cloud providers and environments. The team is on track to ingest 50 TB per day across both Amazon and Google cloud storage platforms in 2022. Regardless of the cloud service provider each region chooses to use, ChaosSearch gives users real-time access to business-critical data at scale, without having to compromise on data retention time frames.

Using ChaosSearch, Equifax has minimized hardware maintenance and simplified the management and administration previously associated with log analytics. What's more, reduced complexity within the analytics process allows teams to innovate faster and rely on long-term data. This helps detect trends in application and infrastructure performance, as well as avoid the risks of persistent security issues. So far, the migration to ChaosSearch has amounted to a 90% cost reduction vs. Equifax's previous providers.

In the future, the Equifax team plans to extend its use of ChaosSearch to other data types beyond logs. For example, the team plans to parse CI/CD platform data using BI tools to understand the usage patterns of its development team. While operational logging data is used today to determine the health of the platform, the team sees many other business analytics use cases for ChaosSearch. BI tools can be leveraged with data accessed via ChaosSearch to gain insights on the adoption of new service offerings.

“ChaosSearch gives us the ability to simplify our data analytics workflows, while still remaining an incredibly powerful solution.”
We can now take advantage of the scale of a data lake platform, without sacrificing cost and performance.”

Jeff Kinsherf, SVP
Engineering Services and SRE
Equifax

To begin your free trial of ChaosSearch, visit chaossearch.io

ABOUT CHAOSSEARCH

ChaosSearch helps modern organizations Know Better® by activating the data lake for analytics. The ChaosSearch Cloud Data Platform indexes customers' cloud data, rendering it fully searchable and enabling analytics at scale with massive reductions of time, cost and complexity. ChaosSearch was purpose-built for cost-effective, highly scalable analytics encompassing full text search, SQL and machine learning capabilities in one unified offering. The patented ChaosSearch technology instantly transforms your cloud object storage (Amazon S3, Google Cloud Storage) into a hot, analytical data lake.

For more information, visit ChaosSearch.io, find us on the AWS Marketplace, or follow us on Twitter @ChaosSearch and LinkedIn