



Stride, Inc. Streamlines DevOps Analytics with ChaosSearch®

LOG ANALYTICS AT STRIDE

Since 2000, Stride has been a leader in K-12 online education, helping more than two million students reach their full potential by bringing quality hands-on and online curriculum directly into their homes. The company provides both tuition-free and tuition-based online education, as well as college and career prep, and learning solutions for school districts.

Analyzing log data was a critical part of the DevOps team’s troubleshooting process. Stride selected ChaosSearch to leverage its existing cloud object storage infrastructure in Amazon S3, while taking advantage of the Kibana analytics interface the team already knew and loved. Using ChaosSearch, the team tapped into a highly scalable, secure cloud data platform that centralized data management and simplified access to analytics — without the extensive infrastructure maintenance pain of the ELK Stack.

THE CHALLENGE: ANALYZING LOGS IN S3 WITHOUT RESTRICTIONS

Behind the scenes, the Stride DevOps team uses the MuleSoft CloudHub managed service to host hundreds of applications, but capturing logs from each application was restrictive and difficult. The team attempted using Logstash to port logs to their ELK Stack infrastructure (Elasticsearch, Logstash, and Kibana), but found the process costly and difficult to manage. Ballooning retention costs and inconsistent log file configuration issues led to wasted time that could be otherwise allocated to troubleshooting.

To simplify this process, the team looked to take advantage of its existing cloud object storage environment in Amazon S3 to reduce the cost of storing logs, and make them easily accessible for analysis.

“Logstash in ELK doesn’t necessarily work well with S3,” said Mike Burgoyne, senior infrastructure and cloud architect at Stride. “We wanted a way to stream logs to an easier-to-use solution, while using Kibana as our analytics interface.”

Stride

Industry
Education

Location
U.S.

Use Cases

- Application troubleshooting
- Debugging
- Analyzing year-over-year application performance

Impact

- Cost savings over previous Elasticsearch cluster
- Reduction in toil and maintenance for log analytics environment
- Flexibility, high availability and performance

Users

- Developers, testers, operations team

SOLUTION: CHAOSSEARCH CLOUD DATA PLATFORM

Today, the company uses ChaosSearch to send logs directly from CloudHub to S3 without the management complexity of ELK Stack.

“We didn’t want to use CloudWatch and AWS because there were some limitations around querying and other things we needed to do,” said Burgoyne.

“Getting started with ChaosSearch was easy, because we were able to use the existing Kibana interface for querying logs without restrictions.”

Mike Burgoyne

Senior Infrastructure and Cloud Architect
Stride

Today, around 25 developers, testers and operations team members now use ChaosSearch. Primarily, they’re querying logs in response to application issues. For troubleshooting use cases, the team can look for patterns in logs and drill further down into the details with ChaosSearch.

“We use ChaosSearch to help debug critical issues,” said Burgoyne. “Logging is a great reference for when we are having issues in the debugging process.”

THE IMPACT: FLEXIBLE, SCALABLE LOG ANALYTICS AT LOWER COST

The DevOps team has been most satisfied with ChaosSearch’s flexibility and ability to work with existing tools and cloud infrastructure.

“It comes down to how much data we’re ingesting,” said Burgoyne. “Now, we’re able to store a lot more data in S3, and index only what we need to in ChaosSearch, rather than having to ingest everything into ELK.”

The team has newfound flexibility when it comes to log retention, and applying log data to answer certain development or testing questions. For example, the team can now save logs related to application performance testing for longer periods of time in S3, and analyze year-over-year trend results in ChaosSearch.

“The good thing about ChaosSearch is we can store all of our logs in S3, and quickly stand up an index to start searching,” said Burgoyne. “In ELK you can put your logs in an UltraWarm or Warm standby scenario, but you’re still paying for infrastructure and disc cost. S3 is quite a bit cheaper.”

According to Burgoyne, the most useful feature is ChaosSearch’s ability to ingest logs from S3, rather than having to port logs to disc and index them there. And with ChaosSearch, the team does not have to perform updates or infrastructure maintenance tasks, saving valuable time that can be spent elsewhere on more strategic tasks.

Data Management Environment

- Cloud Data Platform: ChaosSearch
- Cloud Object Storage: Amazon S3
- Analytics Interface: Kibana
- Infrastructure-as-a-Service Platform: MuleSoft CloudHub

“ChaosSearch is cheaper than using ELK for logs, and we don’t have to maintain the infrastructure.”

Mike Burgoyne

Senior Infrastructure and Cloud Architect
Stride

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

ABOUT CHAOSSEARCH

ChaosSearch enables organizations to Know Better® by activating the data lake for log analytics, delivering unparalleled operational investigation, visualization, and alerting at scale. ChaosSearch eliminates the architectural bottlenecks that cause today’s complex solutions to fail. The end result: rapid time to insights paired with simultaneous reductions in time, cost, and risk. Cloud-based organizations including Equifax, Blackboard, and Digital River rely on the ChaosSearch Cloud Data Platform to query their cloud object storage directly within their preferred visualization tools.

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