



[www.chaossearch.io](http://www.chaossearch.io)

205 Portland Street, Boston  
 MA 02144 (suite 202)  
 Email: [sales@chaossearch.io](mailto:sales@chaossearch.io)

## CHAOSSEARCH

### The company

CHAOSSEARCH is a start-up that has just received its first round of venture capital funding. It is based in Massachusetts.


### What is it?

CHAOSSEARCH is a cloud-native, long-term log and event analytics service running on Amazon S3 that was previously known as Chaos Sumo. It was launched as a beta version in Autumn 2017. A free trial edition is available.

The product uses a serverless architecture and decouples storage from compute, with the former being highly compressed (up to 80%, including relevant indexes).

The company has developed S3 and ElasticSearch API extensions and has built Kibana directly into its console. A more general view of the architecture is illustrated in **Figure 1**, although the TensorFlow support shown will not actually be available until 2019.

“What CHAOSSEARCH brings to the table is an easy, effective way to manage the madness without having to navigate a maze of complex technologies that really only get you part way to solving the issue.”  
 Nomad



The image in this Mutable Quadrant is derived from 13 high level metrics, the more the image covers a section the better. Execution metrics relate to the company, Technology to the product, Creativity to both technical and business innovation and Scale covers the potential business and market impact.

### What does it do?

CHAOSSEARCH uses a patent pending, distributed data fabric that separates storage (S3) from compute. It has been built using containers (Docker) and open source components such as Akka and Kibana. It has developed a read-focused, ElasticSearch compliant API, specifically designed to extend ElasticSearch capabilities to Amazon S3. It also has a patent pending for its Data Edge indexing engine that will automatically compress its results by up to 80 percent.

Data Edge also has some significant advantages other than supporting compression. For example, it supports both relational queries and text search against the same indexed datasets. Further, it includes facilities for live ordering or filtering of a data source, so that result sets can be live rather than static: as the data source changes so does the (virtual) target, which is especially relevant for real-time or streaming use cases.

Other relevant capabilities include support for data filters, which provide the ability to organise your data based on its date, type or source, so that you can search or query the data within one of those contexts. You can also create virtual folders and you can join, aggregate and

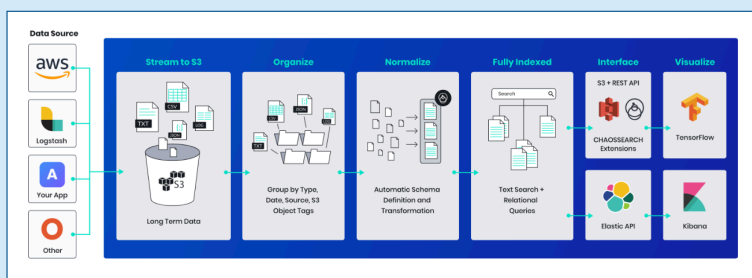


Figure 1 – CHAOSSEARCH conceptual architecture

From a use case perspective, the company targets historical log and event analytics, S3 log data management, AWS application log analytics, data retention to support compliance requirements, and application log analytics for users of SaaS services. In addition, CHAOSSEARCH targets ELK (ElasticSearch, Logstash and Kibana) cost reduction.

Analytics support	★★★★☆
Architecture	★★★★★
Ease of use	★★★★☆
Features	★★★★★

Hybrid data support	★★★★★
Integration	★★★★☆
Performance	★★★★★
Scalability	★★★★★

“Companies are generating petabytes of log/event data and archiving it in S3 storage, where accessing it is slow, complex and cost prohibited. CHAOSSEARCH has solved this problem and built a new and powerful search technology and distributed compute fabric around S3 that opens up fast access to data at scale, at very low cost and high performance.”

Cloud Health Technologies

correlate your data. As noted previously, Kibana, which is an open-source plug-in for ElasticSearch, has been built into CHAOSSEARCH's console (see [Figure 2](#)).

Integration with TensorFlow to provide machine learning and advanced analytics capability, is scheduled for release in 2019.

### Why should you care?

The main problem with traditional approaches to log analytics is that they are expensive. This will often mean that you have to delete data sooner than you would like, so it is only available for analysis for a short period of time rather than the extended duration you would prefer. The reason that these solutions are expensive are several: they do not compress the data well, the indexes that you need to search through the data bloat the amount of storage you require, and there is sometimes a need to over-provision your environment (in order to cater for, say, a distributed denial of service attack).

CHAOSSEARCH is able to reduce the cost of log analytics by automatically indexing the data and compressing the results by as much as 80%. Further, the separation of compute from storage means that the requirements of the former do not impact on the latter. As a result, on a pure cost comparison basis, CHAOSSEARCH significantly undercuts its competition. Its prices are available on its web site for readers to check out. In a quick comparison with another vendor that also publishes its subscription fees online we found that to store 100GB per day, for a year, would save nearly \$21,000 if using CHAOSSEARCH.

Direct costs aren't the only benefit of using CHAOSSEARCH. Apart from its functionality, scalability (elastic by design) and performance, the fact that the offering is serverless is a major advantage, because it means that the service does not need to be provisioned, configured or managed. That said, being realists, we expect that it will be the cost savings that will be of most interest to potential users.

### The Bottom Line

CHAOSSEARCH has entered a busy market in which there are a number of existing vendors. To be successful in such an environment you have to be innovative and since the major issue with log analytics is the cost of storing the data you have to be innovative in a way that allows your offering to reduce those costs. CHAOSSEARCH appears to have done precisely this.

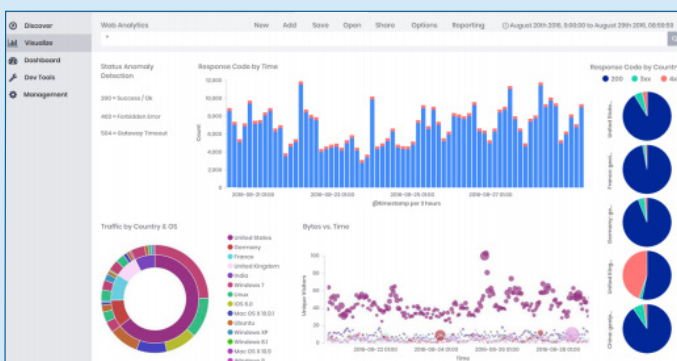


Figure 2 – CHAOSSEARCH Console

FOR FURTHER INFORMATION AND RESEARCH [CLICK HERE](#)