

Is Your Legacy ELK Stack Falling Over Under The Weight Of Expanding Data?

These days, rapidly scaling log and event data is shining a harsh light on the complexity of maintaining a legacy ELK Stack.



To keep up with the flood of log and event data, operators must continue to add additional Elasticsearch servers and expensive disk to each system.

In fact, most businesses running an ELK Stack are forced to archive or delete log and event data after a few days or weeks because it's just too complex and cost-prohibitive to retain. With log and event storage – the indices can be larger than the source data itself. The growing flood of log and event data has companies scrambling to add additional Elasticsearch servers and expensive disks to each system to continue to scale.

Offloading indices to S3 only slows your time to answers

To avoid building a complicated Hot/Warm Elasticsearch cluster, many operators leverage the built-in AWS plugins for Elasticsearch. They snapshot their indexes directly to Amazon S3 to age out old data from the cluster. However, indices sent to Amazon S3 first need to be restored to a running, hot Elasticsearch cluster before you can ask questions. If your existing Elasticsearch cluster does not have the available disk space to restore to, you would need to provision additional servers to complete the restore operation. This is a time consuming, complex, and expensive process that drastically increases the time required to gain value from your data.

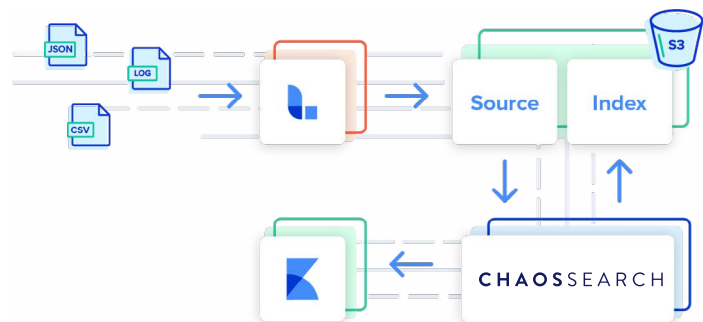


Store everything. Ask anything.

Introducing CHAOSSEARCH. The New Standard in Log Search Analytics

Leverage the low cost of Amazon S3 for long term log retention

Don't wait for
answers –
store everything
& ask anything!



With **CHAOSSEARCH**, use your Amazon S3 account to store and index all your log and event data without ever having to move the data or process into a separate database. **CHAOSSEARCH** stores all of the indexes in your Amazon S3 account in a highly compressed state, yet still fully searchable and queryable. Leave all your data fully indexed within your Amazon S3 bucket, and get quick answers to your questions.

Don't spend time and energy building Elasticsearch clusters to support restoring your Lucene indexes to ask questions. Just leverage the **CHAOSSEARCH** platform to search, query, and visualize your data, all instantly, all without ever having to move your data.

ELB (1.35 billion)	CHAOSSEARCH	elasticsearch
Index - Time	1.6 hours	4.2 days
Index - Size	150 GB	2.1 TB
Index - Cost	c4.2xlarge (<\$1.00)	c4.2xlarge (\$38.00)

Start your free trial of CHAOSSEARCH today!