SOLUTION BRIEF

How Cribl Stream Optimizes Your Splunk Cloud Migration

THE CHALLENGE

Splunk's workload based pricing masks the true cost of unlimited data ingestion by penalizing searches and driving up storage costs over time.

THE SOLUTION

Cribl Stream optimizes data on its way into Splunk, delivering optimal search performance and low cost options for long-term storage.

THE BENEFITS

- Classify and optimize data on the way into Splunk, ensuring you're not spending workload cycles processing data at its most expensive point.

- Optimize Splunk search performance by removing redundant data before it is indexed.

- Take advantage of low cost object storage and replay data on your terms.

SOLUTION BRIEF
—

How Cribl Stream Optimizes Your Splunk Cloud Migration

Shifting Splunk's core pricing metric from daily ingest to workload sounds like it solves the Splunk tax. After all, everyone wants to put more data into Splunk, and removing ingest constraints appears to solve that pain for its customers. However, the shift to workload pricing comes with its own challenges - challenges an observability pipeline is uniquely positioned to solve.

As the leading observability pipeline, Cribl Stream provides distinct benefits to users on Splunk's workload pricing:

- *Optimizing getting data in*
- *Faster, more relevant searches*
- *Managing storage growth and costs*

Optimize Getting Data In (GDI)

Users of Splunk's workload pricing are told that the amount and types of data ingested into Splunk no longer matter because the metric has shifted. However, this is a misconception that drives up your costs. Ingesting data is a workload, and you will be billed accordingly. This is especially true for unclassified data lacking defined index sources and source types, as well as unstructured data, like XML payloads and large events. Both of these data types consume large amounts of CPU and memory during ingest and indexing.

With Cribl Stream, you can classify and optimize data on the way into Splunk, ensuring you're not wasting precious workload cycles on data processing at its most expensive point.

Getting data in is only part of the equation. You only put data into Splunk to do something with it. That's where better search performance becomes critical.

Faster, More Relevant Searches

Shifting from ingest pricing to workload pricing untethers teams from immediate concerns about how much data they are bringing into their Splunk environment. However, unbridled ingest has a downstream effect on search performance and relevance. Bringing in low value data just because you can impacts indexing, resulting in slower performing queries that drive up workloads.

By putting you in control over what data goes into Splunk, you can:

- *Route low value data to the right storage tier and replay it later*
- *Filter out duplicate or null values, or even convert data to more efficient formats*
- *Redact sensitive data before it lands in long-term storage*

Each of these factors lowers your overall workload costs because you're creating a better data product for Splunk to work with. Of course, that data product needs somewhere to live. This is where managing your storage becomes essential in workload pricing.

Manage Storage Growth and Costs

The Splunk workload pricing model licenses storage in 500GB blocks, and these costs can increase substantially if your data retention durations are longer than ninety days. It's easy to exceed your storage allocation since you're no longer thinking in terms of ingested data.

Cribl Stream helps you manage storage costs in two ways. First, Stream lets you reduce the amount of data you're storing by removing redundant fields and optimizing data for your needs, resulting in lower overall storage costs. Next, Stream allows you to archive data to low-cost object storage that is more cost effective than storage offered by Splunk. If you later decide you want to use archived data, Stream allows you to replay data and deliver it to multiple destinations.