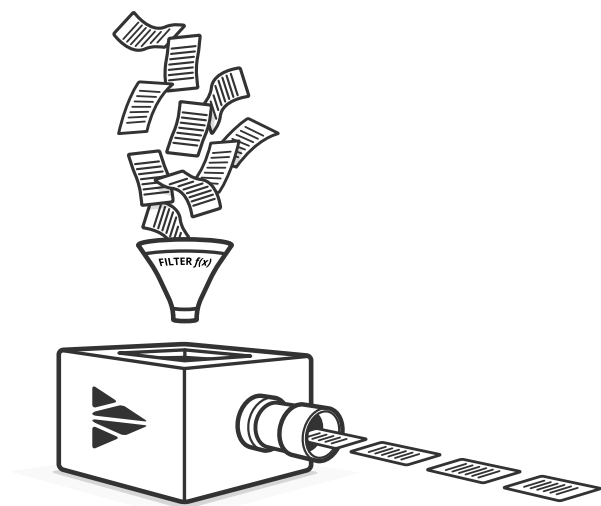Cribl

Cribl Stream – Data Reduction

Reducing the volume of your IT and security data can help you control license compliance, reduce your infrastructure budget, and improve the performance of your analytics tools. Cribl Stream helps you get control over increasing data volumes by letting you trim unused data. Data that has little to no analytical value can account for as much as 50% of your data. You should never pay to ingest, analyze, or store data that doesn't help you solve your observability or security challenges. Cribl Stream helps you tune out the noise and focus on the signal.

THE CHALLENGE

Log data is growing at 28% CAGR,. You may be anxious about exceeding daily ingest limits, and employ workarounds to stay under the limit. Long-term retention of data in analytics tools is expensive and can slow query performance.

The challenge.

Logging systems and other observability tools are at capacity. IT and Security teams are facing more challenges than ever with dramatically increasing volumes of machine-generated data already growing at 28% CAGR, and the incredible number of shapes and forms that data often takes. Most traditional data engine solutions focus on structuring data into columns and rows and putting them into tables, which requires a massive investment in labor in order to realize value from these solutions. Cribl approaches IT and Security data from first principles thinking, and is now unveiling its full suite of products built upon a powerful data engine that is designed especially for voluminous and messy IT and Security data.

THE SOLUTION

Eliminate up to 50% of ingested log volume to optimize licenses, control costs, and improve system performance. Eliminate duplicate fields, null values, and data that has little analytical value. Filter events for dynamic sampling, or aggregate log data into metrics. All while preserving the crucial information required by your analytics platforms. Keep a full-fidelity copy in low-cost storage and replay it if needed.

The Cribl solution.

Using Cribl Stream or Edge can reduce up to 50% of ingested log volume to help control costs and improve system performance. Easily eliminate duplicate fields, null values, and any elements that provide little analytical value using dynamic sampling. From the same interface, Cribl customers can filter and screen events or aggregate log data into metrics for volume reduction at scale – all while keeping a full-fidelity copy in low-cost storage to replay if needed.

THE BENEFITS

- **Eliminate license anxiety** – control your log volume to optimize or reduce ingest.
- **Improve performance** – declutter data in your analytics tools for better query and search performance.
- **Reduce infrastructure** – keep less data in logging tools and retain it in object storage for compliance.

Facets of data reduction.

Eliminate noise and analyze signal.

Over time, IT and security data can become bloated. The same timestamp is recorded 4 different times in 3 different formats, or a paragraph of legal information is appended to every log – adding zero usable information. Stream helps you cut out the noise and lets you focus on data with high analytical value. Sometimes you've collected verbose log data when all you need is a subset of metrics. For these rare use cases, Stream lets you aggregate logs into metrics, which can trim data volume by a factor of 1,000 to 1. You can also capture data samples to see the impact of your reduction efforts, and more strategically sample events with redundant or uninteresting information. Cut the noise, and route only valuable data, to eliminate license anxieties and slash infrastructure costs.

Improve performance of your analytics systems.

Analytics tools are often called on to perform "needle-in-a-haystack" style searches or perform correlations of data to look for patterns. The more data these tools have to sift through and analyze, the poorer the performance. Stream can help your logging tools search through a subset of data with the highest analytical value. When these tools aren't bogged down by all of the clutter of low-value data, they perform much better and require a smaller infrastructure footprint. This can dramatically reduce compute expenditures.

Separate your systems of analysis and retention.

Analytics tools are great for analyzing data. They don't measure up nearly as well as a destination for long-term retention. Separating the systems where you analyze data in real time from where you store it long-term allows you to keep more data, for longer periods of time, and for a lot less money.

Routing full-fidelity data to lower-cost object storage can be a significantly better choice for long-term retention. Because Stream can replay this data to any tool at a later point in time, you can reduce the retention period for data in your analytics tools. Drop data after you've made your initial analysis, for better performance and a leaner infrastructure. subset of data with the highest analytical value. When these tools aren't bogged down by all of the clutter of low-value data, they perform much better and require a smaller infrastructure footprint. This can dramatically reduce compute expenditures.

> Cribl Stream helps you cut out the noise and lets you focus on data with high analytical value.

Summary.

Using Stream to implement an observability pipeline helps you unlock the value of all of your IT and security data. Stream helps you eliminate any data that doesn't provide useful information about the well-being of your applications and systems. Cut up to 50% of your data volume to eliminate license anxiety, improve the performance of your analytics tools, and reduce the size and cost of your infrastructure.

>SOLUTION BRIEF: "CRIBL STREAM – DATA REDUCTION"_