>CASE STUDY_

Sally Beauty transforms their IT and security data management, switching from LogStash and Syslog-ng to a superior cloud-based data engine

HIGHLIGHTS

- Freeing up 75% of security architects' time using Cribl vs managing syslog-ng and Logstash.
- Reduced infrastructure required to run Elastic SIEM by half.
- Saving time by transforming Endgamedata into Elastic Common Schema format

Sally Beauty Holdings is an international specialty retailer and distributor of professional beauty supplies with around $3.9 billion in annual revenues. The company sells and distributes more than 8000 products through more than 5000 stores in the Americas and Europe. To quickly spot and respond to potential security incidents, the security team keeps a close eye on application and device data. With Cribl.Cloud, Sally Beauty's security experts can focus on security instead of data engineering.

To spot unusual activity on endpoints, including point-of-sale (POS) devices, Sally Beauty's security team wanted a cleaner way to bring Endgame endpoint detection and response (EDR) data into Elastic Cloud. "Like many data sources, Endgame doesn't let you pick and choose which components of the data stream you send, and where," says Sheldon Carmichael, Information Security Architect for Sally Beauty. That meant some data sent to Elastic Cloud was redundant. Other data wasn't properly converted to the Elastic Common Schema (ECS), slowing down threat investigation and remediation.

"We wanted a more resilient and flexible data pipeline," Sheldon says. The immediate need was converting the right Endgame data to ECS, and dropping unneeded ields to make room for additional data sources and retain relevant data for longer.

Intercepting the endgame data stream.

Hearing about Cribl.Cloud, Sheldon conducted an extended evaluation, successfully routing and optimizing the Endgame data stream. "Even better," he adds, "I carved up the multiple data types in the Endgame data stream, sending each one via its own route to its own pipeline. That means we can route to destinations other than Elastic as needed. With Cribl, we have full control of what data we send, drop, and enhance — and where we send it to take action."

“I spent a couple of weeks casually working through cribl pipelines to do the most reduction out of the gate. We’re reducing 9.25tb of daily edr data down to a little over 5tb a day – 41% total reduction. Instead of holding data in elastic for seven Days, we’re holding it for 30.”

– Sheldon Carmichael, Information Security Architect

41% data reduction.

By trimming the Endgame data sent to Elastic Cloud, Sally Beauty gets rid of noise — and frees up space in Elastic to keep data for longer. “I spent a couple of weeks casually working through Cribl pipelines to get the most optimization out of the gate,” Sheldon says. “We’re reducing 9.25TB of daily EDR data down to a little over 5TB a day — 41% reduction. Now we retain data for investigations for 45 days instead of seven.”

“We don’t want our security engineers spending hours becoming wizards in Elastic, Logstash, and syslog-ng. With Cribl, they don’t have to, which frees up more time for our real jobs – security. That’s Cribl’s biggest ROI.”

— Sheldon Carmichael, Information Security Architect

Out of the weeds: more time to focus on security.

Cribl frees up compute on the front end and helps the team keep the data for longer retention periods — but the primary win for Sally Beauty is saving time for the security team so they can focus on security — not data engineering. For example, by dropping redundant, blank, or description field clutter, the team can quickly zero in on the fields relevant for investigations. They have the space to bring in backlogged data sources for a more comprehensive view of the company’s security posture. With less data, searches go faster.

Goodnight, syslog-ng server.

Sally Beauty saved even more time by replacing its syslog-ng and Logstash servers with redundant Cribl workers. The cutover was “quick and relatively painless,” according to Carmichael. “Compared to our syslog-ng and Logstash servers, Cribl pipelines are much more straightforward to manage, and they give me both GUI and command line access,” he says. He estimates that his team manages Cribl in one-quarter of the time they spent managing Logstash and syslog. “We don’t want our security engineers spending hours becoming wizards in Elastic, Logstash, and syslog-ng,” says Sheldon. “With Cribl, they don’t have to, which frees up more time for our real jobs — security. That’s Cribl’s biggest ROI.”

The redundant Cribl workers also make their data pipeline more resilient and easier to manage. The team performs rolling upgrades and maintenance on individual workers without interrupting the data flow.

Tastes like middleware.

Besides moving Endgame data into Elastic, Sally Beauty also uses Cribl to pull IT and Security data from sources that don’t have the capability to push, like their Network Detection and Response (NDR) solution and Microsoft Office365. “Before Cribl, to pull data from our NDR solution and O365, we had to run python scripts on separate servers acting as middleware to pull the data we wanted. That just adds to the systems and resources to maintain, along with everything else we have to manage manually vs in an automated fashion.” says Sheldon. “Cribl allowed us to collapse multiple functions and resources into a single solution with enterprise-grade support,” said Sheldon. This resilience, time savings, and flexibility is what ultimately persuaded Sheldon’s leadership to invest in Cribl.

Looking to the future: using built-in detections.

By trimming the Endgame data sent to Elastic Cloud, Sally Beauty gets rid of noise — and frees up space in Elastic to keep data for longer. "I spent a couple of weeks casually working through Cribl pipelines to get the most optimization out of the gate," Sheldon says. "We're reducing 9.25TB of daily EDR data down to a little over 5TB a day — 41% reduction. Now we retain data for investigations for 45 days instead of seven."

Now Sheldon is planning to use Cribl to modify the Endgame dataset to emulate Elastic Agent data streams. "By doing this we can take advantage of Elastic's hundreds of built-in detections, and add custom exclusions that persist when Elastic updates its rules, without having to modify each rule every time there's an upgrade," he says. The potential time savings are huge. Today the team has to clone the built-in detection rules, modify the observed indices in the rule to include Endgame, and add custom exceptions to tune the rules to Sally Beauty's environment.

By using Elastic's built-in detections and the modified data stream from Cribl, Sheldon expects to save significant time supporting upgrades, and to strengthen the company's security footprint by using more properly tuned and updated detection rules. "A future version of Elastic might give me that capability," he says. "But the fact that I can do it today with Cribl is pretty awesome, and allows me to apply the same functions to any other observability data we bring into our pipeline."

Next up for Sally Beauty? Using Cribl to enrich the Endgame stream — for example, by mapping IP addresses to the device's location (geo-IP) and incorporating threat intelligence. Wrapping up, Sheldon says, "We've seen big benefits from Cribl, and I'm a fan. If anyone reading this comes to a Cribl meetup, look me up if you want to talk nerdy."

> **"Our CISO knows the value of NDR and O365 observability data for security, and he also understood it was a huge pain for our team to search and parse the data. With Cribl, the world is our oyster when it comes to search."**
>
> — Sheldon Carmichael, Information Security Architect

TL;DR

- Sally Beauty wanted control over which Endgame data it shipped to Elastic Cloud.
- Cribl.Cloud parses Endgame data, standardizes it to match the ECS, and reduces it before sending it to Elastic.
- Cribl reduced daily EDR data by 41% — from 9.25TB to 5TB.
- The company can now retain EDR data for 45 days instead of 7.
- The team replaced syslog-ng and Logstash with Cribl, estimating it takes ¼ of the time to manage Cribl versus the legacy tools.
- Eliminated the need to write and manage python scripts to collect "pull" data sources like NDR and O365 logs.
- Not having to learn the ins and outs of logging software frees up more time for engineers to work on security.
- Selecting Cribl.Cloud gave the team fast time to value and control over cloud latency andegress costs.

ABOUT CRIBL

Cribl, the Data Engine for IT and Security, empowers organizations to transform their data strategy. Customers use Cribl's vendor-agnostic solutions to analyze, collect, process, and route all IT and security data from any source or in any destination, delivering the choice, control, and flexibility required to adapt to their ever-changing needs. Cribl's product suite, which is used by Fortune 1000 companies globally, is purpose-built for IT and Security, including Cribl Stream, the industry's leading observability pipeline, Cribl Edge, an intelligent vendor-neutral agent, Cribl Search, the industry's first search-in-place solution, and Cribl Lake, a turnkey data lake. Founded in 2018, Cribl is a remote-first workforce with an office in San Francisco, CA.

Learn more: www.cribl.io | Try now: Cribl sandboxes | Join us: Slack community | Follow us: LinkedIn and Twitter

CS-0014-EN-3-0524