Cribl

>SOLUTION BRIEF_

Fast Growing US City Keeps a Close Eye on User Experiences and Security by Routing Log Data to Different Analytics Tools

For one of the Fastest Growing US Cities - and Capital of its state — data is a precious resource. Moving data from more sources to more destinations falls to the City's centralized IT department.

"Our applications and infrastructure produce 2TB of log and event data daily that various City departments use to spot cyberthreats and performance issues," says the Senior Enterprise Architect for the city. Previously, all log and event data—from VPNs, firewalls, CrowdStrike, Microsoft 365, Windows, and more—went to a single destination, Splunk." But as data volumes grew, the team needed more choice and control over which data was routed where. "Whenever a department requested a new destination, like Azure Sentinel, admins had to drop whatever they were doing to reconfigure Splunk Universal Forwarders and syslog-ng servers," the Senior Enterprise Architect says. "And data volume was approaching our Splunk license limits." The tipping point came as the team started planning a migration to Splunk Cloud. "If we didn't get data from our dozens of syslog-ng servers into a data pipeline, the migration would be a nightmare. We could no longer afford to not have an observability pipeline."

HIGHLIGHTS

- City teams send log and event data to more destinations to strengthen cybersecurity and improve citizen experiences with city services.
- Selectively routing data to Azure Sentinel, Amazon S3 or Splunk helps the city better manage license and storage costs.
- With the observability pipeline in place, the team was able to confidently migrate to Splunk Cloud in half the time.

Freedom to Route Data Anywhere

The City consulted with a leading analyst firm to recommend an observability pipeline. On the shortlist was Cribl Stream, which team members had seen years earlier in a Splunk User Forum. The team selected Cribl Stream after also considering Splunk, Edge Delta, Calyptia, Fluentd, and Kafka.

> “Adding a new destination in Cribl typically takes just minutes. We can make forwarding decisions without bugging syslog admins and interrupting their other work.”

> “Splunk is one of many destinations, so it doesn’t make sense for it to own all of our data,” Senior Enterprise Architect says. “Edge Delta wasn’t proven, and the others aren’t supported. Only Cribl met our all our requirements as a cost-conscious government entity running critical online services.”

The team decided on a hybrid deployment, combining the convenience of Cribl.Cloud (software-as-a-service) with the ability to pull in on-prem data and utilize on-prem storage to comply with regulatory requirements. The City purchased Cribl through public-sector reseller Freeit Data Solutions, tapping Cribl Professional Services for deployment. The team also takes advantage of Cribl’s no-cost learning resources, like documentation, Cribl University, labs, sandboxes, and the Cribl Curious forum.

You Want Data Sent Where? Sure.

Now the team can quickly fulfill requests from departments to deliver log and event data to new applications like Azure Sentinel, or new storage locations like Amazon S3. What used to take days to fulfill new data onboarding requests, now takes a few hours and reduces the back and forth that historically slowed down the overall process.

> “Adding a new destination in Cribl typically takes just a few minutes,” says the Senior Enterprise Architect. “Now we can make forwarding decisions without bugging admins and interrupting their other work.”

The team now routes some data to multiple locations. For example, firewall and VPN logs are sent to both Splunk Enterprise Security and Azure Sentinel. On the other hand, Microsoft O365 telemetry is sent to Sentinel alone, conserving the Splunk license.

> “With Cribl.Cloud we don’t have to worry about under – or overestimating licensing needs because we pay only for what we use.”

Tagging: Data Governance Superpower

The power of tagging came as a welcome surprise after Stream was already deployed. The team tags data in Cribl Stream to control what data is sent where, and for access control. Filters on the Splunk side exclude tagged data from query results unless the requester is authorized.

> “When we didn’t know about tagging we didn’t miss it,” said the Senior Enterprise Architect. “But now that we have tagging in Cribl we’d never give it up.”

> "Before we had Cribl Stream, troubleshooting meant a lengthy back-and-forth with the application vendor to get access to databases and find buried log information. Now, with direct visibility into log data, application owners can spot performance issues sooner, taking action when it counts."

Real Value to Residents and the Workforce

Just months into the deployment, Cribl Stream is already helping City departments meet their mission. One example is faster identification of issues with the computer-aided dispatching (CAD) system used for 9-1-1 calls. Avoiding a 5-10 second delay in dispatching first responders — say, because the SQL server is churning — can affect outcomes.

> **"Before we had Cribl Stream, troubleshooting the CAD system meant a lengthy back-and-forth with the application owner to get access to databases and find buried log information," says the Senior Enterprise Architect. "Now, with direct visibility into log data, application owners can quickly identify the source of performance issues to resolve them sooner."**

Teed Up: Simpler Migration to Splunk Cloud

Cribl Stream will pay more dividends when CTM migrates to Splunk Cloud.

> **"Now that we have Cribl we won't have to pester two dozen stakeholders to reconfigure their Universal Forwarders and syslog-ng servers," the Senior Enterprise Architect says.**

It's estimated that Cribl will cut migration time in half and save approximately four hours of work for each of dozens of server admins.

New ideas just keep coming.

> **"With a robust processing engine and data pipeline we can coach departments on what data they have and how they can use it to deliver better experiences.," the Senior Enterprise Architect says.**

One idea: improve the work-from-home experience by using Cribl Edge to collect telemetry from end-user devices. (Installing Splunk Universal Forwarders on 15,000 endpoints isn't practical.) The Team also wants to share log and event data with the City's Data Science team to give them new sources of insights for improving resident services.

Summing up, the Senior Enterprise Architect says, "The biggest value of Cribl Stream for our city is opening people's eyes to what data we have. Log data isn't boring once you understand what you can do with it. Personally, it's been very satisfying to optimize a very complex process to make government more efficient."

TL;DR

- Before Stream, adding a new data source to Splunk required hours of work to configure Universal Forwarders and syslog servers.
- With Cribl, adding new destinations takes just minutes.
- Tagging data in Cribl simplifies access control.
- Bottom line: City teams get the data they need sooner—and the existing Splunk license goes further.

Get Cribl Stream, and take control of your data.