


CRIBLCON²⁴

POWERED BY  Cribl

s/Chaos/Control/g

Modernizing the data
pipeline with Cribl.

Aaron Wilson
SRE Manager, iHerb





AARON WILSON

SRE Manager, iHerb

Really just an engineer pretending to be a manager. Over 25 years in the industry, starting back in physical hardware and datacenter days, to today's modern DevOps, Cloud, and virtual server environments.




JON RUST

Solution Engineer, Cribl

30 years of log wrangling.
Former Splunk Arch, and ISP owner
"Logs flow through my veins, and I dream in regex."

Agenda

- 1 **The many layers of logging**
- 2 **Simplifying the layers**
- 3 **Making the layers even better**
- 4 **Takeaways and Tips**



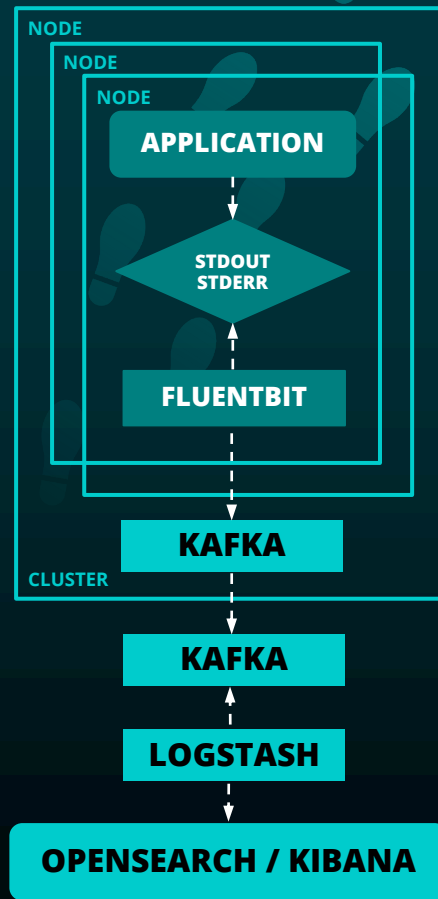
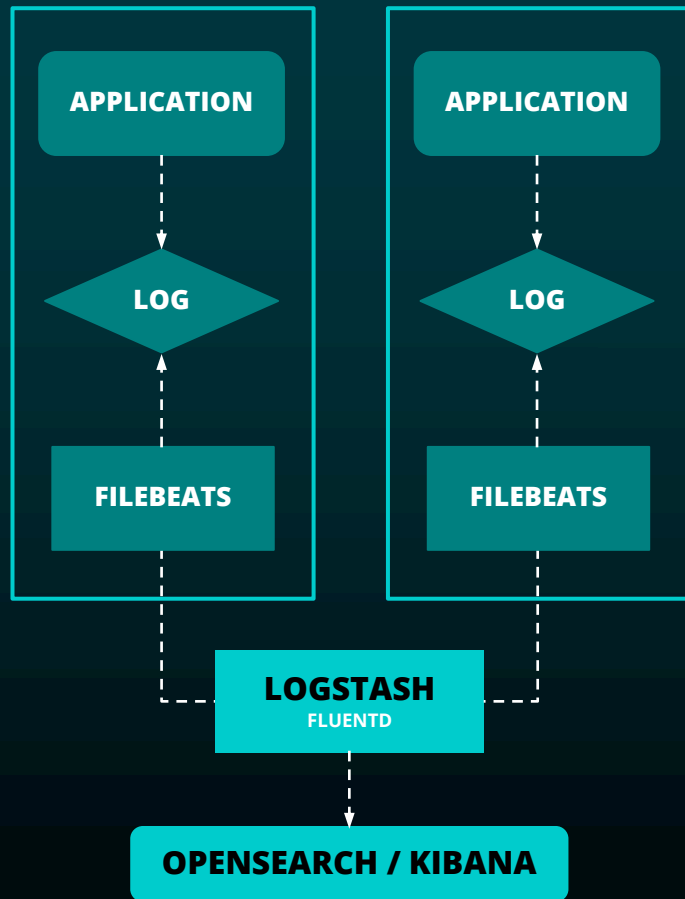


"Logging has
layers, onions
have layers (Sic)."

Trademarked green ogre.



Logging Layers



It began as a solution to our logging problems.



VM / Physical hosts:

- A myriad of logging locations
- No central management of the config files
- Manual changes anytime something new was added or needed

K8s:

- Daemonset fluentbit for node level container log pickup
- Deployment fluentd for regex basic log manipulation
- Central Kafka for temporary log storage
- More fluentd with major regex log manipulation



"Not everybody
likes onions."

Miniature beast of burden



A hand holding a magnifying glass over a bookshelf in a library. The background is a dark, atmospheric illustration of a library with tall bookshelves and a patterned rug. The text is centered in a large, white, serif font.

Logging could be
like Parfait
Everybody likes
Parfait

Everybody likes simplified layers

VM / Physical / K8s



Edge on all sources

- Centralized config point
 - Easy to adjust
- Simple log manipulation easy to handle



Stream for heavy lift on log manipulations

- HPA deployments on a dedicated "monitoring" cluster
- More intense log manipulation

24

Fleets

2418

Nodes

19

Config versions

24

Routes

293

Pipelines

375

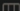

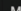

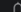
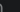
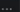


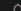



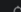

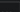

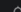
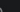
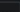


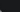




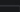

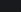
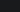
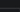
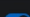


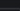
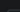



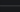

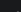

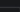



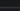
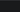

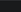
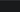
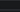
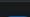
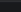
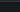
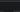
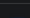
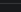
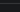
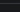
Sources

143

Destinations

Start typing to find Fleets

Add Fleet

 Name	Description	Group Type	UI Access 	Nodes	Active Rules	Details	Deployed version	Target Software Version 	Msg	Actions
default_fleet	Default Fleet	Hybrid		0	1	View	d26afc1	None	 	Add Subfleet Commit Deploy 
 preprod		Hybrid		0	0	View	8d46dd8	None	 	Add Subfleet Commit Deploy 
monitoring-edge-preprod		Hybrid		19	1	View	d26afc1	None	 	Add Subfleet Commit Deploy 
warehouse-edge-preprod		Hybrid		297	1	View	144fc36	None	 	Add Subfleet Commit Deploy 
utilities-edge-preprod		Hybrid		171	1	View	846476f	4.6.0	 	Add Subfleet Commit Deploy 
central-edge-preprod		Hybrid		222	1	View	48b29c0	None	 	Add Subfleet Commit Deploy 
catalog-edge-preprod		Hybrid		121	1	View	939b596	None	 	Add Subfleet Commit Deploy 
checkout-edge-preprod		Hybrid		67	1	View	7385706	None	 	Add Subfleet Commit Deploy 
 preprod-vm		Hybrid		0	0	View	d26afc1	None	 	Add Subfleet Commit Deploy 
linux-edge-preprod		Hybrid		0	1	View	d26afc1	None	 	Add Subfleet Commit Deploy 
windows-edge-preprod		Hybrid		26	1	View	0eac993	None	 	Add Subfleet Commit Deploy 
 prod		Hybrid		0	0	View	d26afc1	None	 	Add Subfleet Commit Deploy 
catalog-edge-prod		Hybrid		653	1	View	08975fc	None	 	Add Subfleet Commit Deploy 
checkout-edge-prod		Hybrid		267	1	View	22a8ccc	None	 	Add Subfleet Commit Deploy 



The sweet parfait toppings!



Long-term Storage

S3 with easy
replay/rehydration
Removed cold storage
Reduced hot/warm storage
to 2 weeks



Tracking

Easy, visible log path
tracking and errors to
correct "missing" logs



Routing

Effortless routing to
multiple endpoints and
pipelines



But even Parfaits
can get better!

Always move forward.

Cribl[®].Cloud

Cribl.Cloud

- Moved our segregated preprod and prod leaders to the cloud
- On prem cluster cost savings – Basically took up an entire 4XL node in each env (leader plus all the daemonsets needed)
- Performance gains – Really needed an even bigger node, as performance was laggy with all the edge nodes



Cribl Search

- Step 1 will be to do away with log rehydration when requested. Users will have access to 1 year of history via Search
- Step 2 will be to do away with OpenSearch entirely, relying on Search



Cribl Lake

- Looking to replace our current data engineering pipeline of kafka connectors, mirror makers with Lake

The gumdrop buttons



A hand holding a magnifying glass over a server room floor. The background is a dark teal server room with rows of server racks and a door labeled 'CONE' in the distance.

Key takeaways

ROI

- **Time savings:** teams can focus on other projects rather than toil away with logging and telemetry data
- **Stability** makes the day-to-day usually negligible
- **Less stress:** quick changes, routing, and pipelines rules make work quick and painless

Best practices

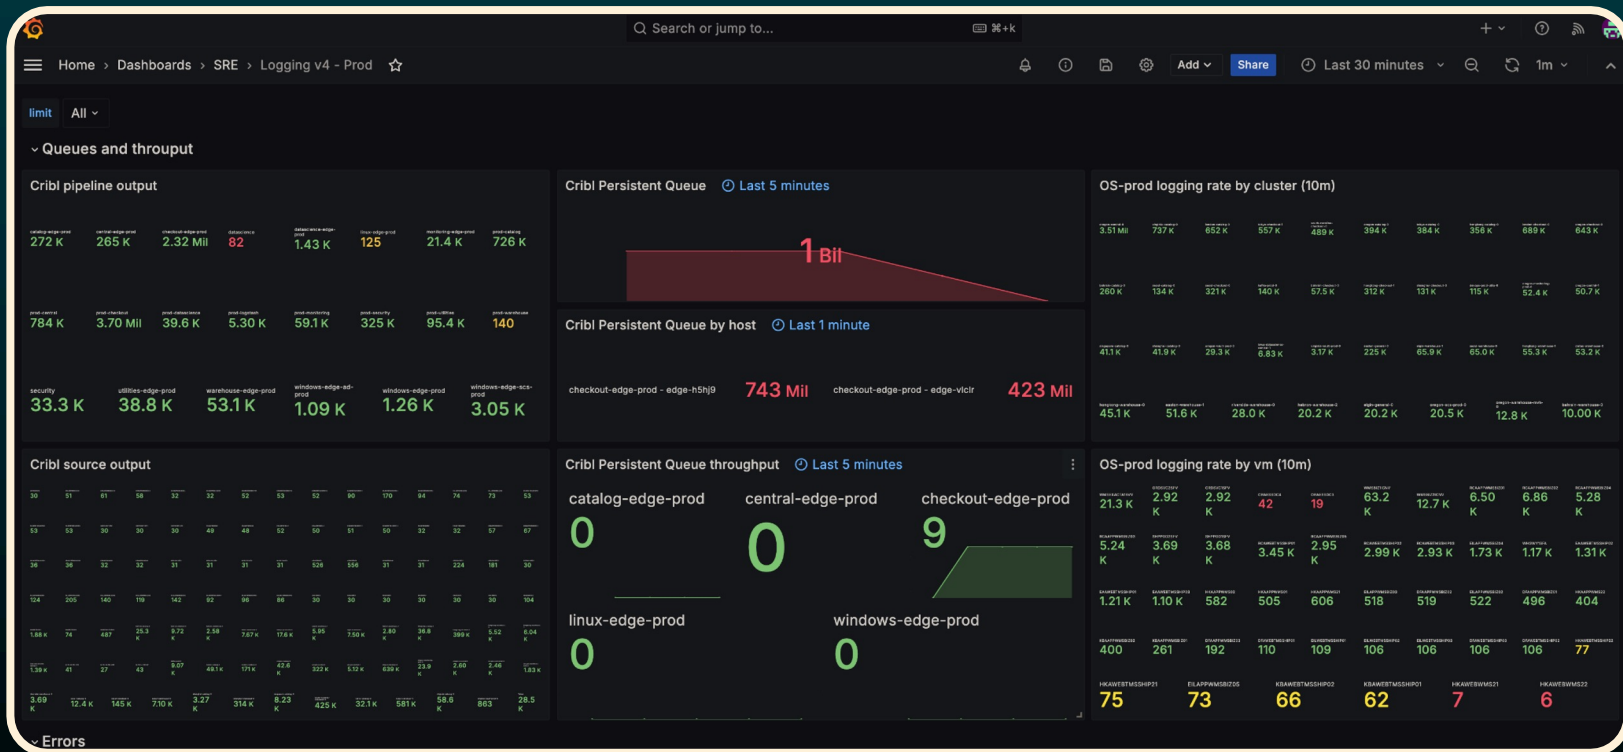
- Using a CD solution like Flux to manage the Edge and Stream workers
- Easy gitops based yaml updates to release new version, tune containers, and push entire new clusters
- Most of our Devops tooling is managed this way

A hand holding a magnifying glass over a server room floor. The background is a dark teal server room with rows of server racks. A door with a 'CLOSE' sign is visible on the left. The floor has a checkered pattern.


...and tips / tricks:

- Pass data streams between different groups / streams with CribITCP. No extra ingest fees
- Use tags
- Naming conventions
- Use multiple layers of Persistent Queues, and tune them. Smaller at Edge, larger at Stream
- Packs provide another level of hierarchy, and easy portability

Grafana metrics.



Edge PQ setup:

 Destination > Cribl TCP
CriblTCP-catalog

Configure

Status

Charts

Live Data

Logs

Test

Notifications

General Settings

Persistent Queue Settings

TLS Settings (Client Side)

Timeout Settings

Processing Settings

Post-Processing

Advanced Settings

Max file size ⓘ

100 MB

Max queue size ⓘ

10GB

Queue file path ⓘ

\$CRIBL_HOME/state/queues

Compression ⓘ

Gzip

Queue-full behavior ⓘ

Block

Strict ordering ⓘ

No

Drain rate limit (EPS) ⓘ

0

Clear Persistent Queue

Q&A

We're goating
you to ask us
anything...





Thank you!