CRIBLCON 24

POWERED BY Cribl

Delivering Observability for Highly Available Services with Cribl

JACOB GORNEY
Manager, Site Reliability Engineering, Cribl

JOSH BIGGLEY
Staff Product Manager, Cribl

JOSH BIGGLEY

Staff Product Manager, Cribl

- 26 years in tech.
- 16 years focused on monitoring & observability.
- Experience in manufacturing, healthcare, cloud software, and government operations.

JACOB GORNEY

Manager, Site Reliability Engineering, Cribl

- Leads the cloud infrastructure team, part of Cloud SRE @ Cribl.
- A decade in cloud engineering and scaling distributed services.
- A former customer turned Criblanian.

Agenda

1. **INTRODUCTIONS**
   A little bit about us
2. **THE CHALLENGE**
   What we're solving for
3. **CRIBL.CLOUD PRIMER**
   Our growth trajectory
4. **CRIBL.CLOUD'S OBSERVABILITY**
   About our environment
5. **STREAMING FLEXIBILITY**
   How we leverage Stream
6. **THE EDGE OF DISCOVERY**
   Harvesting data directly from our infrastructure
7. **SEARCH ANYTHING**
   Using Search across all observability data
8. **"GOATFOODING"**
   Shipping better products

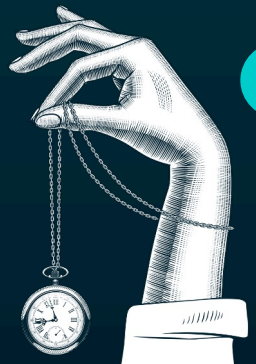

The Challenge

How to scale observability practices for global teams delivering high-growth services.

Observability made easier, not harder.

Cribl.Cloud Primer

Traffic and Growth

DATA PROCESSED

9

Petabytes (Monthly)

DATA OUTBOUND

13

Petabytes (Monthly)

ENVIRONMENTS

3,000+

Active Cribl.Cloud Organizations Today

Thousands of cloud organizations and hundreds of daily sign ups per day across seven globally distributed regions with more locations on the way.

Cribl.Cloud's Observability

High Level Monitoring System Design

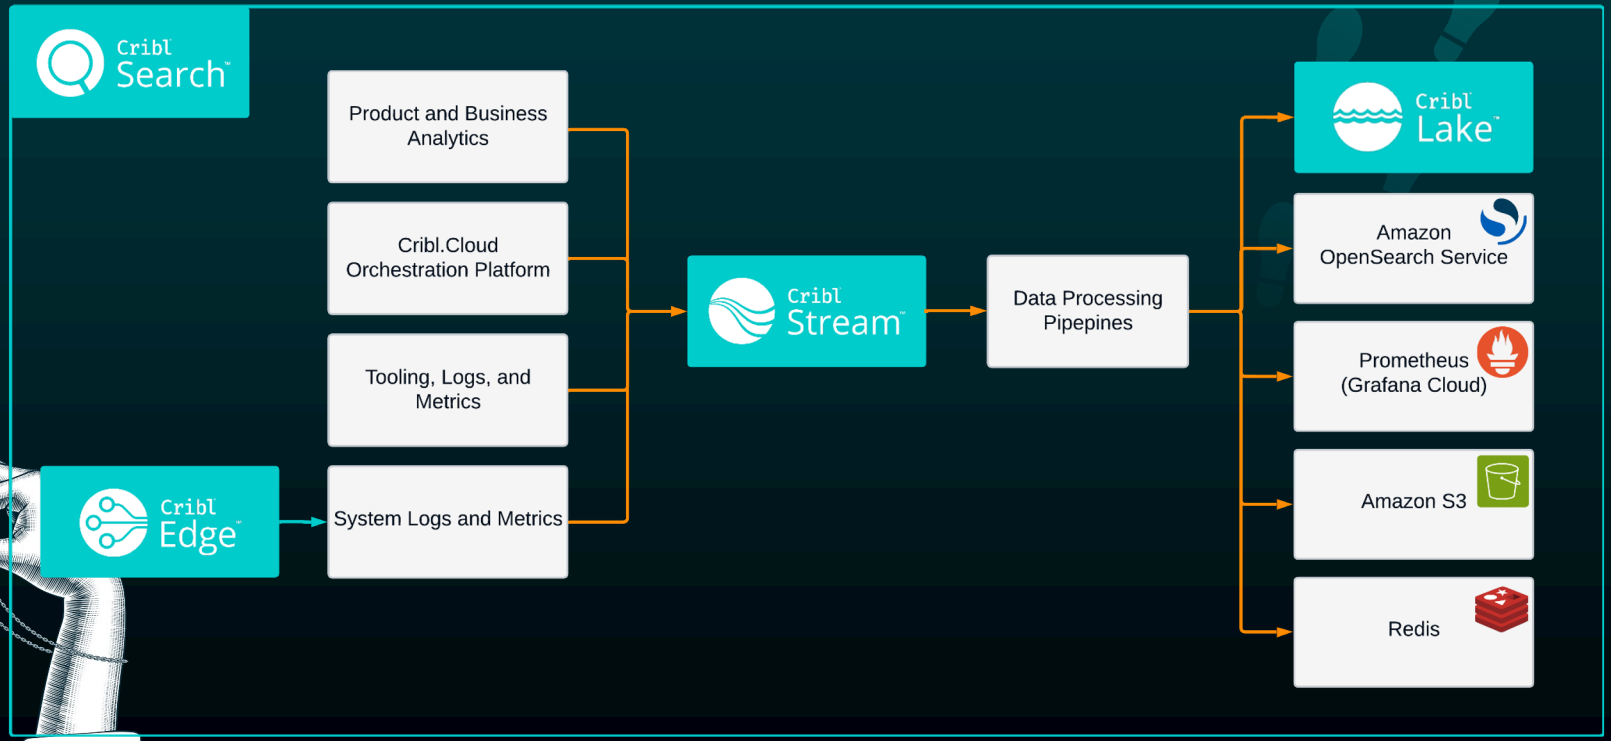

Cribl.Cloud's Observability

High Level Monitoring System Design

Cribl Stream & Edge

Edge runs on all cloud infrastructure.

Cribl Stream at the center; responsible for processing metrics and log data.

Destinations

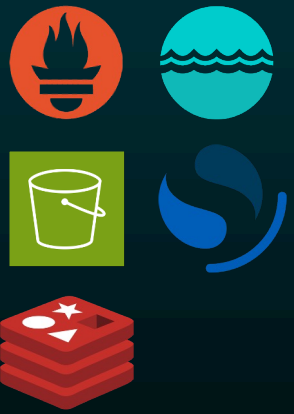

Cribl Search & Lake

Events and metrics stored in Cribl Lake, S3, OpenSearch, Grafana Cloud (Prometheus), and Redis.

Cribl Search as our "pane of glass" over all data stores.

Cribl.Cloud's Observability

By the Numbers

EDGE NODES

6,400+

Running Nodes

EVENTS

310M

Processed Per Day

PROMETHEUS

650K

Active Metric Series

SEARCH JOBS

500+

Completed Monthly

SCANNED EVENTS

250M

Searched Monthly

DATA STORAGE

~50TB

Object Storage Usage

Workers

6

Uptime

Max: main-default: 10.2 weeks

Min: main-default: 10.2 weeks

CPU counts

Max: main-default 32

Min: main-default 32

Memory

Max: main-default 61.7 GiB

Min: main-default 61.7 GiB

Version

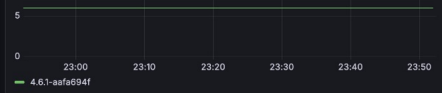

CPU %

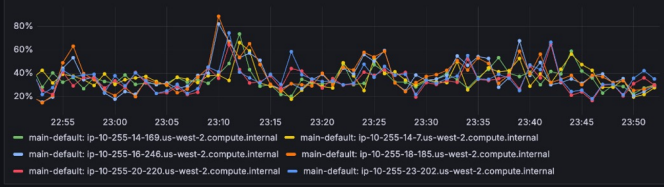

Memory Used %

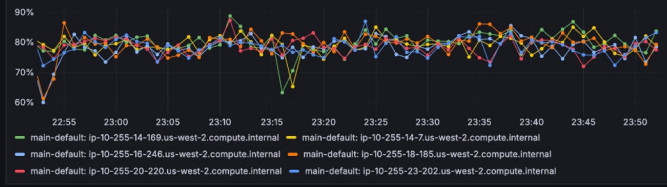

Disk Read/Writes

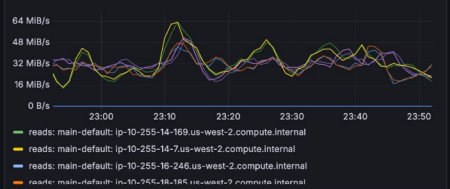

Open FDs

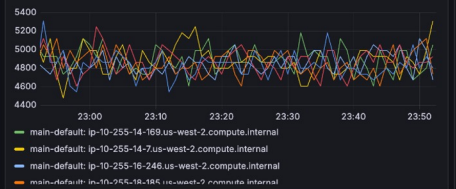

filesystem used

Network Receive/Transmit

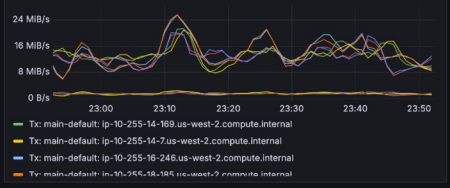

TCP sockets

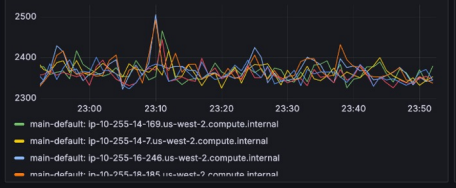

Healthy instances

TCP/UDP Errors

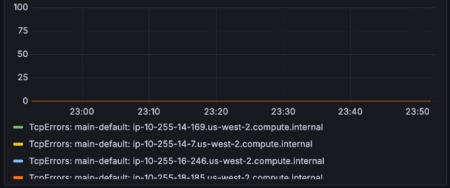

Story: Streaming Flexibility

Story: Streaming Flexibility

How Cribl.Cloud Leverages Stream

- Cribl Stream is the heart of our operation.
  - `Logs, Metrics, API collectors, etc.`
- Cribl is a consumer of Cribl.Cloud
- Quick decision making on the fly:
  - `S3 -> Cribl Lake`
  - `Metric aggregations for Grafana Cloud (Prometheus)`
  - `OpenSearch adoption for indexed events`

Cribl Stream gives us the ability to choose what to do with our data when it matters most.

Stream

Home Manage Monitoring Settings

Search Stream...

Overview Data System Reports Flows Logs Notifications

Live All Groups All Workers 15min

8	23	271	296	568.48GB	418GB / 690GB
Groups	Workers	Processes	CPUs	Memory	Storage

Events In and Out

Live

THRUPUT IN (AVG)	EVENTS IN	THRUPUT OUT (AVG)	EVENTS OUT
18.23kEPS	15.32m	28.90kEPS	24.28m

35.67k
26.75k
17.84k
8.92k
0.00

19:40:00 19:42:00 19:44:00 19:46:00 19:48:00 19:50:00

Bytes In and Out

Live

THRUPUT IN (AVG)	BYTES IN	THRUPUT OUT (AVG)	BYTES OUT
4.15MBPS	3.41GB	20.64MBPS	16.93GB

24.83MB
18.62MB
12.41MB
6.21MB
0.00B

19:40:00 19:42:00 19:44:00 19:46:00 19:48:00 19:50:00

Free Memory

26.75GB
20.06GB
13.37GB
6.69GB
0.00B

19:40:00 19:41:00 19:42:00 19:43:00 19:44:00 19:45:00 19:46:00 19:47:00 19:48:00 19:49:00 19:50:00 19:51:00

CPU Load Average (1min)

View CPU load by node

3.12
2.34
1.56
0.78
0.00

19:40:00 19:41:00 19:42:00 19:43:00 19:44:00 19:45:00 19:46:00 19:47:00 19:48:00 19:49:00 19:50:00 19:51:00

Story: The Edge of Discovery

Story: The Edge of Discovery

How Cribl.Cloud Leverages Edge

- The Cloud SRE team pushed the boundaries of Edge node scaling with 4,000+ Edge nodes deployed and reporting to a single Stream Leader.
- Cribl Edge is used for:
  - `OS-level telemetry and logs`
  - `Product logs and metrics`
  - `Security audit events`
  - `System state, health, updates, etc.`
- Edge teleport used for advanced troubleshooting, eliminating the need for shell access.

Cribl Edge reduces our dependence on agents and allows us to make the same quick decisions we make in Stream at the service's edge.


Edge
Home
Manage
Settings
Search Edge...
Fleet prod-workers
Overview
Explore
Collect
Health
Target Software Version 4.6.1
f514f42
Commit
Deploy
Monitor
Map View
List View
Add/Update Edge Node
Live
5min
2003
Edge Nodes
3.03m
Events In
909.29MB
Bytes In
Sources
File Monitor 7
Exec 4
Others 11
Destinations
Cribl HTTP 2
Cribl TCP 1
Others 2
Events In and Out
THRUPUT IN 12.62K EPS (AVG)
EVENTS IN 3.03M
THRUPUT OUT 7.41K EPS (AVG)
EVENTS OUT 1.78M
13.89k
10.27k
6.85k
3.42k
0.00
20:50:00
20:50:20
20:50:40
20:51:00
20:51:20
20:51:40
20:52:00
20:52:20
20:52:40
Bytes In and Out
THRUPUT IN 3.79 MBPS (AVG)
BYTES IN 909.45 MB
THRUPUT OUT 11.41 MBPS (AVG)
BYTES OUT 2.67 GB
11.56MB
8.67MB
5.78MB
2.89MB
0.00B

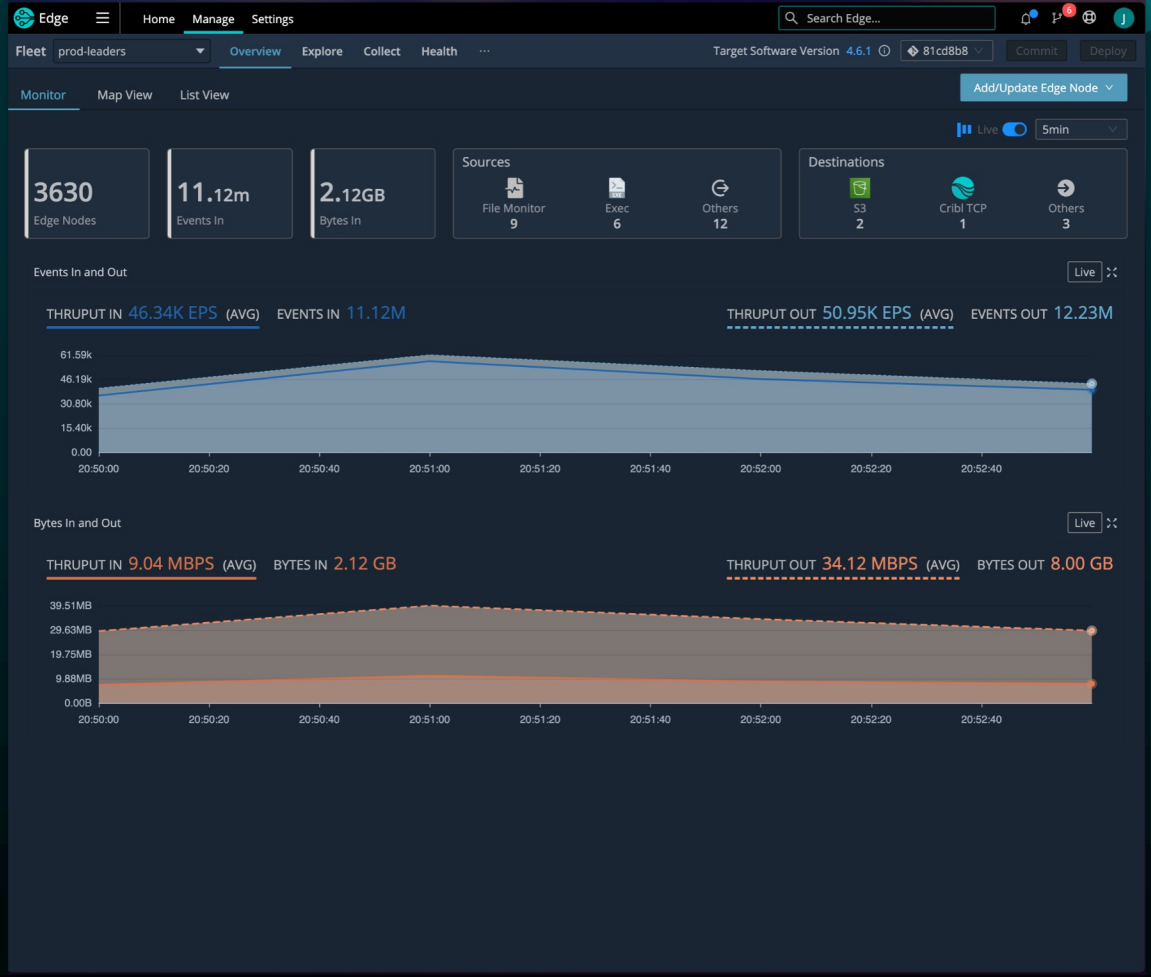
Edge
Home
Manage
Settings
Search Edge...
Fleet prod-leaders
Overview
Explore
Collect
Health
Target Software Version 4.6.1
81cd8b8
Commit
Deploy
Monitor
Map View
List View
Add/Update Edge Node
Live
5min
3630
Edge Nodes
11.12m
Events In
2.12GB
Bytes In
Sources
File Monitor 9
Exec 6
Others 12
Destinations
S3 2
Cribl TCP 1
Others 3
Events In and Out
THRUPUT IN 46.34K EPS (AVG)
EVENTS IN 11.12M
THRUPUT OUT 50.95K EPS (AVG)
EVENTS OUT 12.23M
61.59k
46.19k
30.80k
15.40k
0.00
20:50:00
20:50:20
20:50:40
20:51:00
20:51:20
20:51:40
20:52:00
20:52:20
20:52:40
Bytes In and Out
THRUPUT IN 9.04 MBPS (AVG)
BYTES IN 2.12 GB
THRUPUT OUT 34.12 MBPS (AVG)
BYTES OUT 8.00 GB
39.51MB
29.63MB
19.75MB
9.88MB
0.00B

Story: Search Anything

Story: Search Anything

The Best Compliment to Existing Tools and Data

- All of our data is connected to Cribl Search and is instantly available.
- Complements other tools in our arsenal.
- Available to our entire Product Engineering team.
  - `Tracks the impact of bugs and performance issues.`
  - `Encourages internal bug finding and troubleshooting.`
  - `Brings product engineering closer to the products and features our customers use every day.`

The perfect tool that brings everything together and unlocks additional value out of the tools we already use.

Story: Search Anything

Indexed vs Object Store

- Cribl.Cloud primarily stores data in three places:
  - `Cribl Lake`
  - `Amazon S3`
  - `Amazon OpenSearch Service`
- Amazon OpenSearch used for indexed data.
  - `OpenSearch support in Cribl Search was a joint effort between SRE and Search engineering.`
- The rest lives in Cribl Lake and Amazon S3.

Searches Over 24h

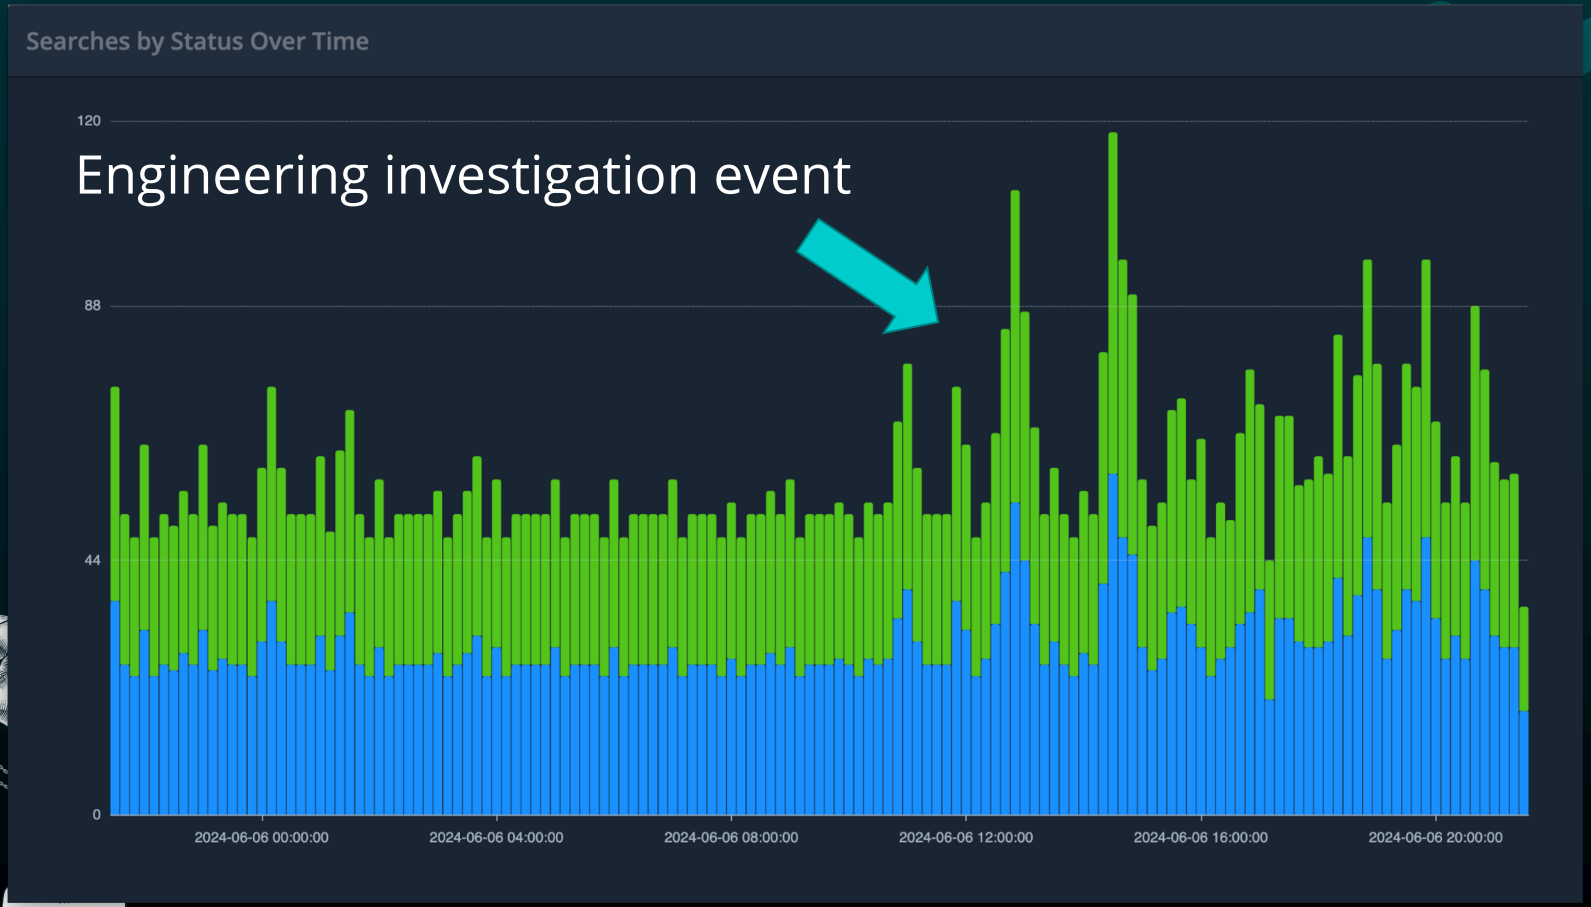

“Goatfooding”

Goatfooding

Internal Product Improvement

- Cribl products are used for all production cloud services.
- Always tackle problems with a “Cribl first” approach.
- Drive product features and bugs, from the inside out.

And push boundaries while doing it!

If we don’t use Cribl products to solve our problems, how can we be sure we solve the challenges our customers face?

Goatfooding

Driving Product Quality

- Cribl.Cloud's observability tooling is used by the entire engineering organization.
  - `We also run pre-release builds!`
- Those internal builds are monitored and used extensively.
  - `Immediate feedback from internal stakeholders.`
  - `Bugs filed and addressed before reaching customers.`
  - `Deep understanding and feedback of new product features and capabilities.`

We ship higher quality software by introducing pre-release versions of our products in our production environments.

Lessons Learned & Good Practices

Recap

Major Highlights

Cribl uses the full suite of products and services in production to support Cribl.Cloud.

Internal customers are customers! We leverage every opportunity we can get.

Stream, Edge, Search, and Lake allows us to grow with our service and make data decisions at supersonic speed.

We build our engineering culture around our customers with an emphasis on quality. We drive quality through adoption.


Thank you!
Come visit us in the CooLAB!