G&L

2024-05-22

Workshop

Processing Video CDN Logs at Scale *Cost Effectively*

touchstream Paramount SVTA

Alexander Leschinsky

Co-Founder & CEO

G&L

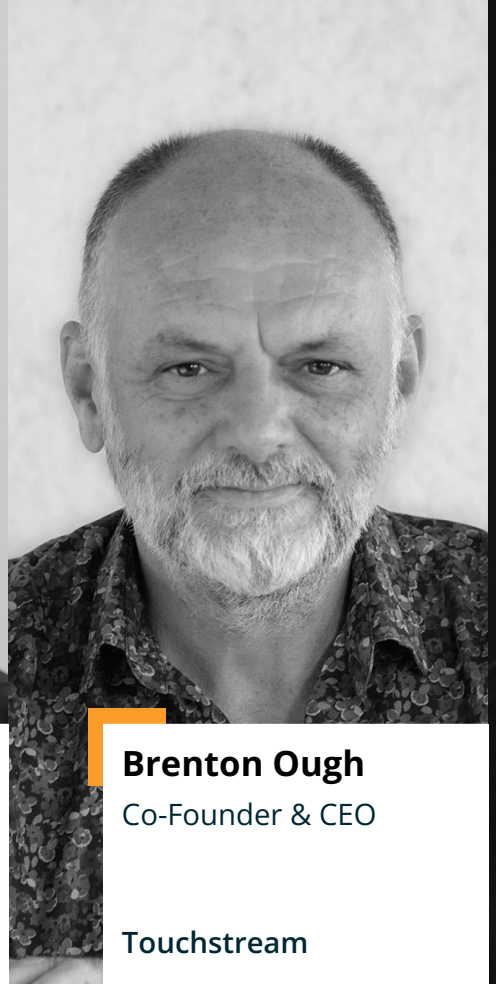

Brenton Ough

Co-Founder & CEO

Touchstream

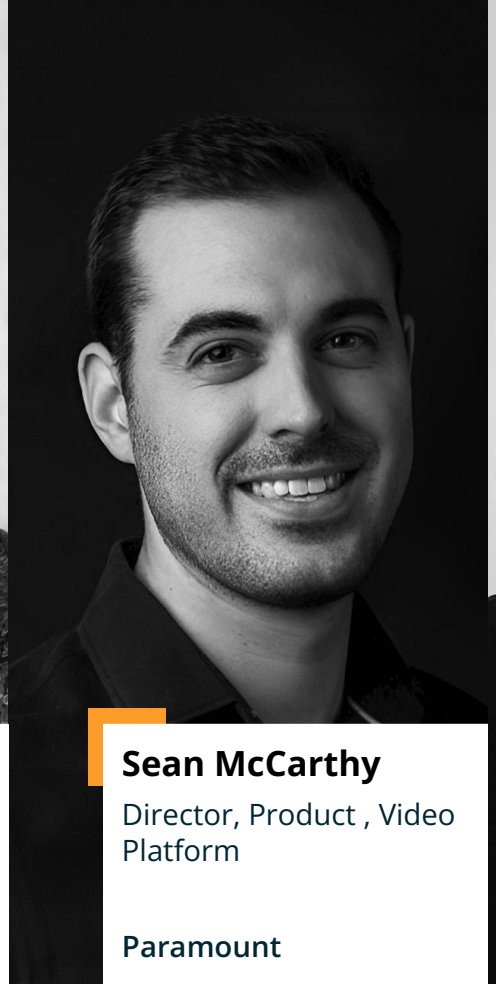

Sean McCarthy

Director, Product , Video Platform

Paramount

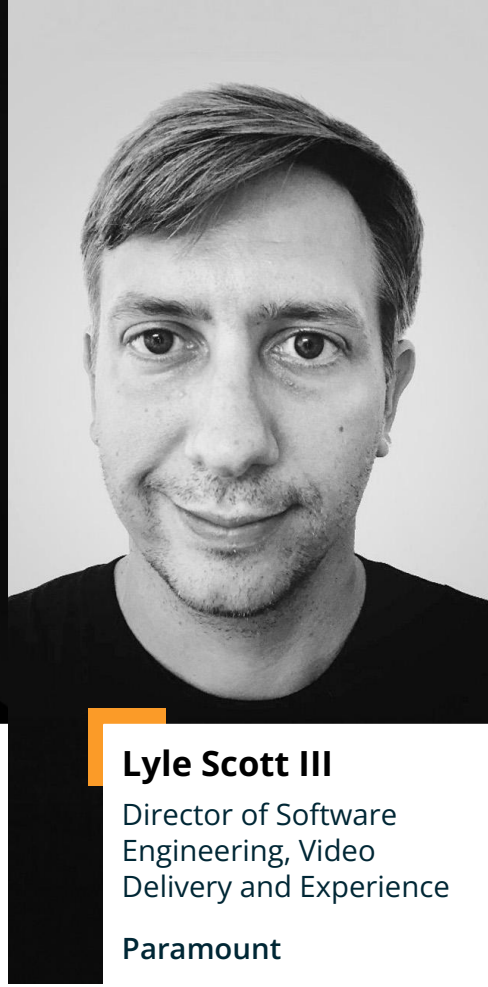

Lyle Scott III

Director of Software Engineering, Video Delivery and Experience

Paramount

What to expect in today’s workshop

Alexander Leschinsky

G&L

Nature of this workshop

- **Processing**
  - **Video CDN Logs**
    - Dealing with different vendors both on the source and tooling side
    - Focus on CDN logs, although many of the principles also apply to other types of logs
  - **at Scale**
    - The larger you are as a broadcaster / OTT platform, the more relevant our finding are for you
  - **Cost Effectively**
    - Log processing can get expensive at the order of magnitude we are looking at
- General
  - No sales pitch
  - This is intended to be a technically detailed Workshop, not a marketing Webinar
  - We are happy to have been endorsed by the SVTA!

Our motivation

Brenton Ough
Sean McCarthy
Lyle Scott
Alexander Leschinsky

G&L

Touchstream motivation

- Develop better ways to identify root cause of streaming issues
- Going beyond capabilities of standard log analytics tools
  - Data Interchange
  - Data Transformation Pipeline
  - Data Storage
  - Dashboarding
- Making data visualisation intuitive
- Cost-effective approaches

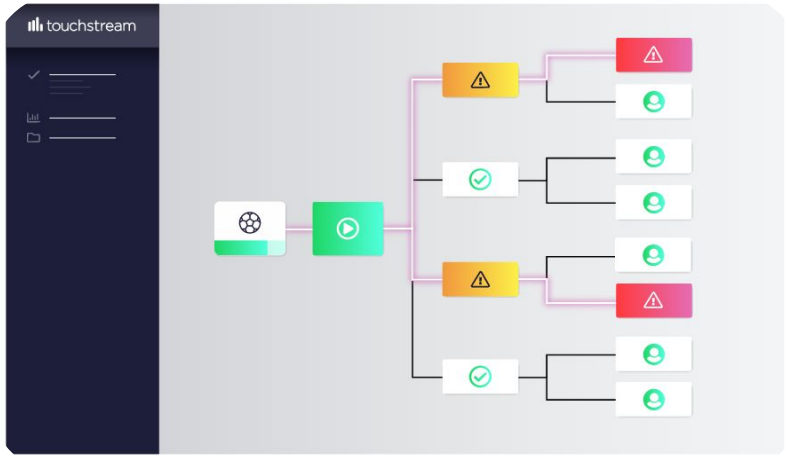

Paramount motivation

- Root cause analysis of stream impacting issues took too long
- Too many operational tools in use
- Standardize black-box CDNs and their unique logging strategies
- Data accessibility issues (need to democratize the data)
- Video stream piracy becoming more of a concern
- Need to perform session debugging
- Business intelligence/Feedback Loops
    - CDN Traffic shaping

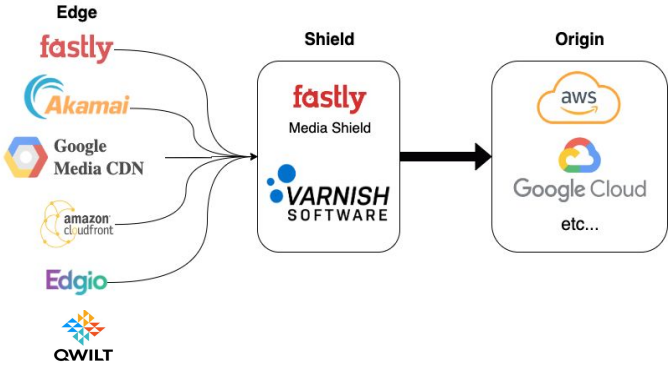

G&L motivation

- Context
  - We design and operate multi-vendor setups in multi-tenant environments
  - Multi-vendor services need data normalization based on raw log data
  - Log processing adds substantial costs if done on scale
  - Aggregated logs only answer the questions you asked when aggregating
  - New questions about the past require long retention periods for raw logs
- Architecture
  - De-couple applications from infrastructure
  - Avoid cloud-provider-specific tools
  - Move applications freely between on-premises and multi-cloud

Data sources

Sean McCarthy
Lyle Scott

Paramount

G&L

What data are we working with?

- Classic CDN logs
- CDN logs enriched with CMCD
- Traditional client analytics metrics
- Client events
- Additional access logs (next slide)

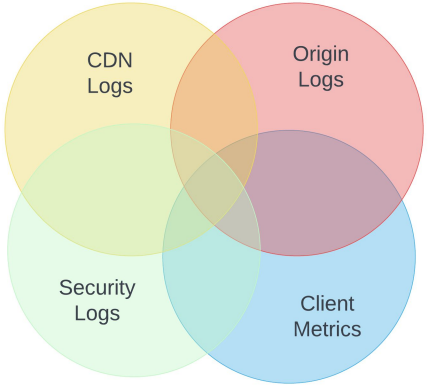

CMCD Device Coverage (Paramount)

Dash.js	Shaka.JS	HLS.js	iOS	Exoplayer	Roku
Query Strings	Query Strings	Query Strings	TBD	Query Strings (To be released)	Headers (default)

End-to-End Monitoring

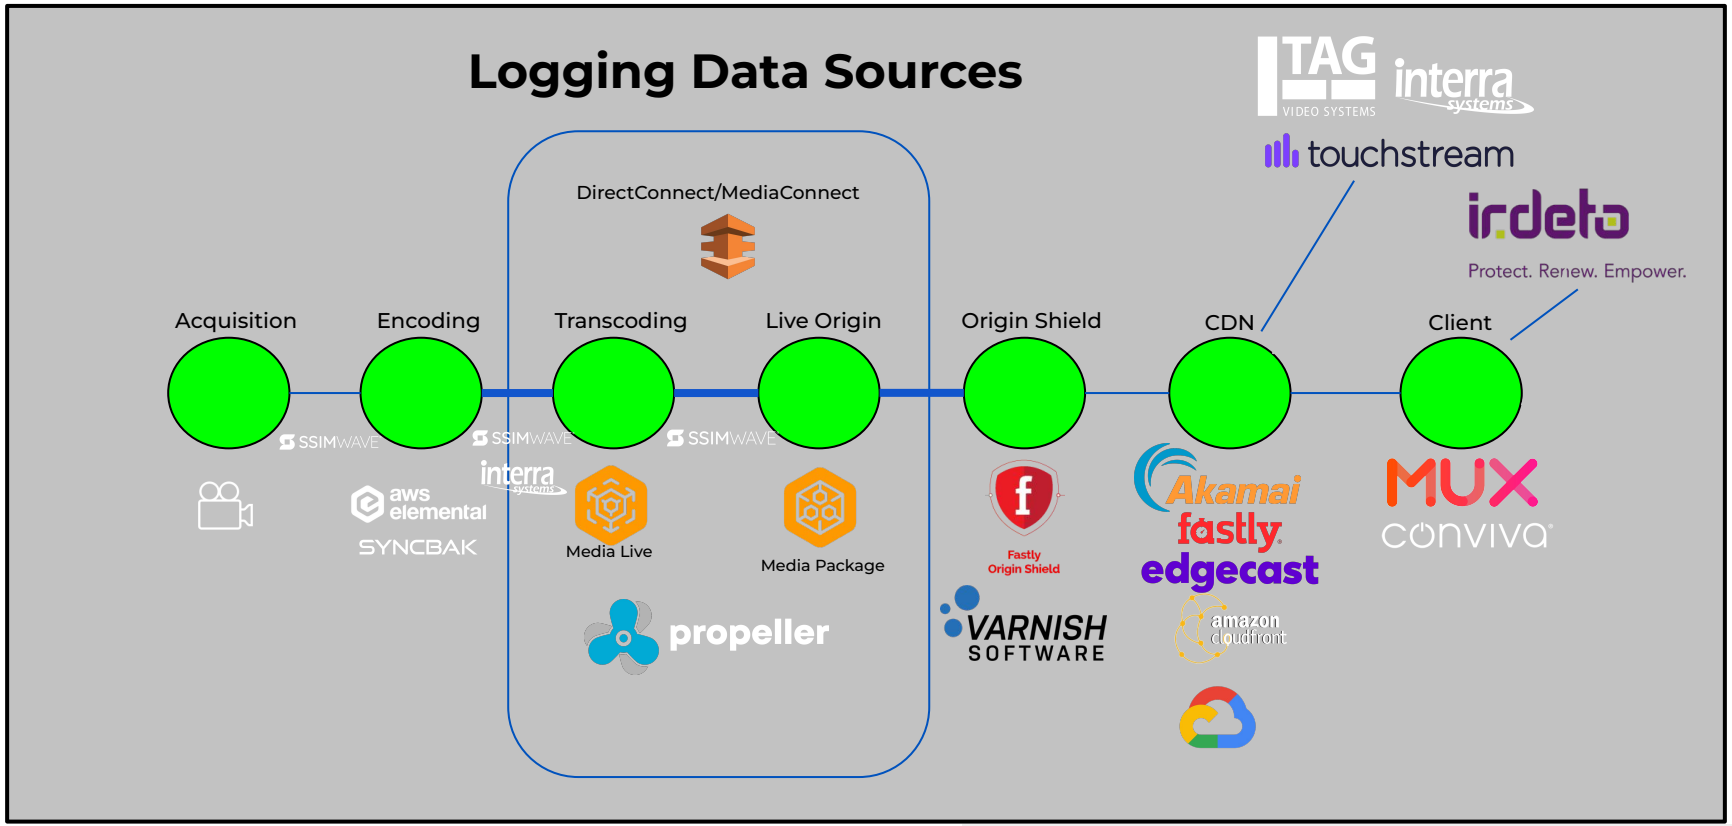

Challenges of high cardinality logs & metrics

Sean McCarthy
Lyle Scott

Paramount

G&L

Normalization

- Normalize CDN metrics across vendors
  - Consistent and comparable data
- Work with CDN vendors to understand logging nuances
  - Fastly: time to last byte documentation very clear
  - Akamai: means different things
  - Cloudfront: numbers are completely different, they might only send if an object is already in memory, vs in disk, then in two different fields, time consuming
- Paramount's field mappings (QR Code)

Fastly	Cloudfront	Edgio	Akamai
503 at timeout value; or enable 'streaming miss' feature and you'll see a 200 with "complete' == false	Cache status == "miss" , Http status == "4xx" or "5xx", Value of the 'x-edge-result-type' field is "Error", value of 'x-edge-response-result-type' is **NOT** "Error", TTLB can be > timeout value,	"/con-err" as cache status, with varying http status. This is a special field Edgecast needs to enable.	HTTP status == 200, custom field to include "client_err_abort" server error code

Distributed Request Tracing

Edge-to-Shield
- Log Edge UUID
- Add UUID to shield request header
- Log this value as shield request ID

Shield-to-Origin
- Return origin request UUID in response header
- Log as “origin_request_id” in shield log line

Mid-tier “intra-CDN” logs
- Breadcrumbs?
- Full log lines?
- Nothing?

Paramount

MediaShield Logging Logic

```
# Set x-edge-request-id
  if (fastly.ff.visits_this_service == 0 && !req.http.x-edge-request-id) {
    if (req.http.X-EC-Uuid){
      # Use the Edgecast Uuid if it is the upstream CDN
      set req.http.x-edge-request-id = "ec-" + req.http.X-EC-Uuid;
    } elsif (req.http.X-Amz-Cf-Id){
      # Use the CloudFront Uuid if it is the upstream CDN
      set req.http.x-edge-request-id = "cf-" + req.http.X-Amz-Cf-Id;
    } elsif (req.http.ak-request-id){
      # Use the Akamai Uuid if it is the upstream CDN
      set req.http.x-edge-request-id = "ak-" + req.http.ak-request-id;
    } else {
      # Generate one if there is no upstream CDN, direct request to Fastly
      set req.http.x-edge-request-id = "fa-" + uuid.version4();
    }
  }
```

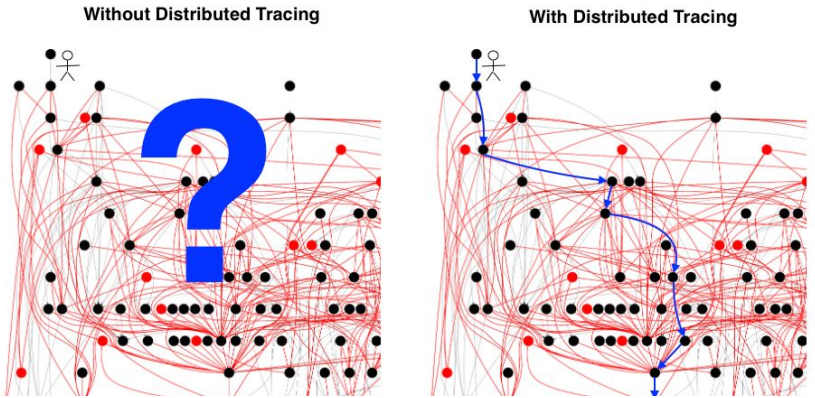

Strategies for reducing the amount of data

Sean McCarthy
Lyle Scott

Paramount

G&L

Sampling Data

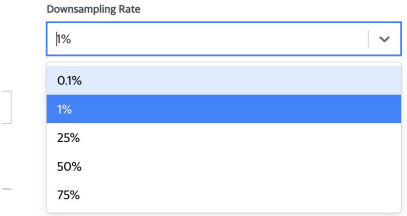

- Make the CDN do it!
  - Simple approach: sample “successes”, not “errors”
  - Sampling options vary greatly
    - If possible, don’t sample cache misses, logs with high transfer times, or buffer flags
  - Multiple destinations
  - Probability-based sampling breaks up sessions
    - We want CDN vendors to sample based on session ID (this is tough)
  - Changing sample rate in a live event can be risky for some Vendors
- Build your own middleware between CDN and backend?
  - Drop data at ingest before indexing
  - Possible, but adds complexity, cost, and risk.

Make the Data Whole Again? - Compensating for sampling

- If you've sampled logs at 50%, just multiply a metric by 2, right?
  - Works:
    - Concurrent viewers
      - Unique count of client_ip + user agent + referrer
    - Error %
  - Does not work:
    - Server_TTLB
    - Error Count -> assuming conditional sampling
  - Kind of works:
    - Bandwidth
    - RPS

Strategies for reducing cost and complexity

Sean McCarthy
Lyle Scott

Paramount

G&L

Big data, big bucks.

- Store data at rest in a compressed / convenient format
- Query resource isolation → ensures best-fit resourcing for data consumers
- Understand data consumers to develop cost-effective strategies
- Lifecycle management (ie, retention on storage, data)
    - Archive “important” data you don’t want to expire to cold storage
- Dashboard query caching
- Scale dynamically
    - Predict audience sizes and viewing patterns

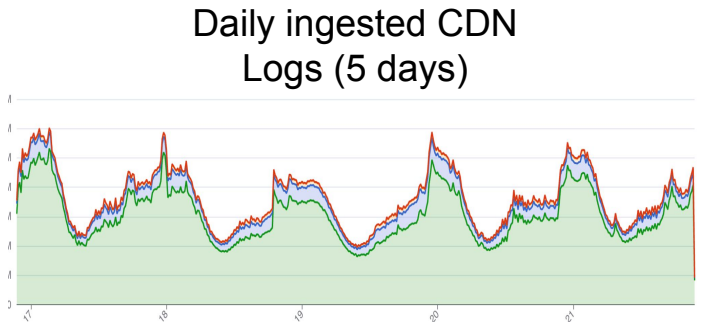

Make Data Approachable

Enrich data to make it more approachable and relevant to your Users' contexts to increase ROI

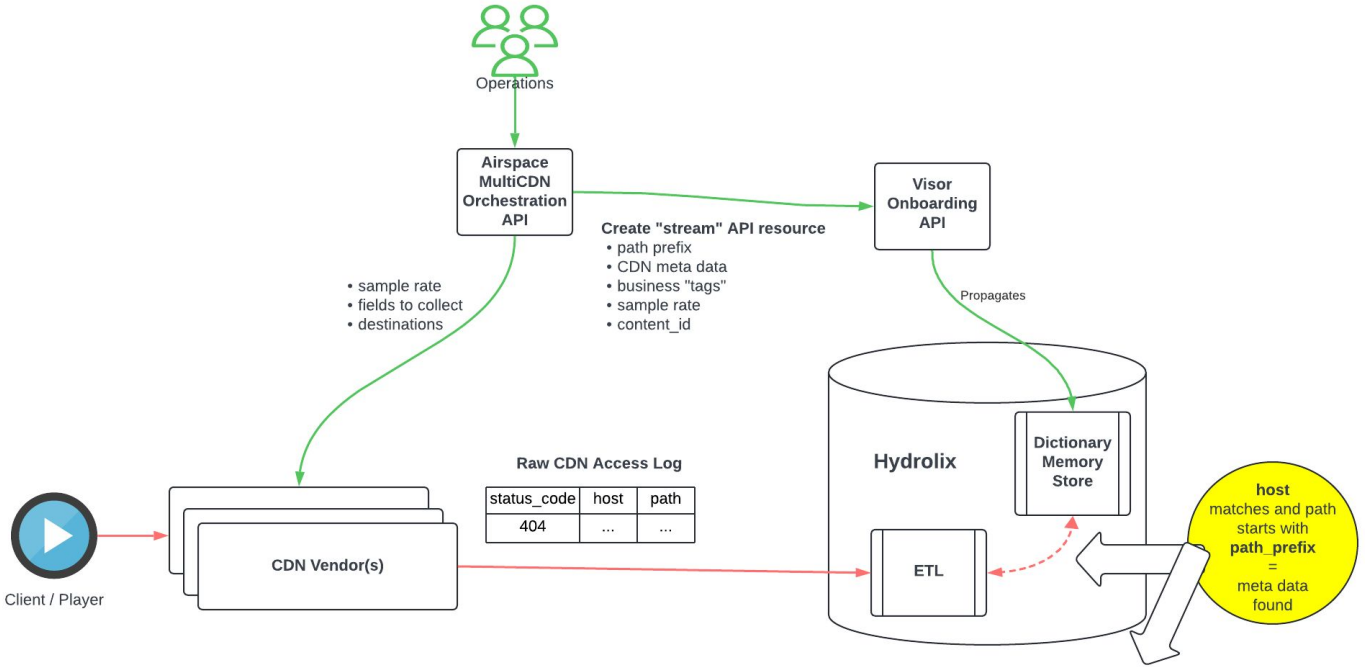

status_code	host	path	business_unit	workflow	stream	sample_rate	content_id
404	...	...	ParamountPlus	live-event	uefa-foo-vs-bar	100	foo
403	...	...	CBS News	live-us	boston-sdr	50	bar
403	...	...	CBS News	live-us	boston-hdr	50	baz

Transformed CDN Access Log

G&L | Paramount

Reduce Complexity

- Summarizing high-cardinality data is useful, but a balancing act
- To queue or not to queue
- Unified and normalized schemas (ie, across CDNs)

Going beyond the capabilities of standard log analytics tools

Brenton Ough

touchstream

G&L

Custom Data Transformation and Dashboards

- Going beyond capabilities of standard log analytics tools
  - Data Interchange
  - Data Transformation Pipeline
  - Data Storage
  - Dashboarding

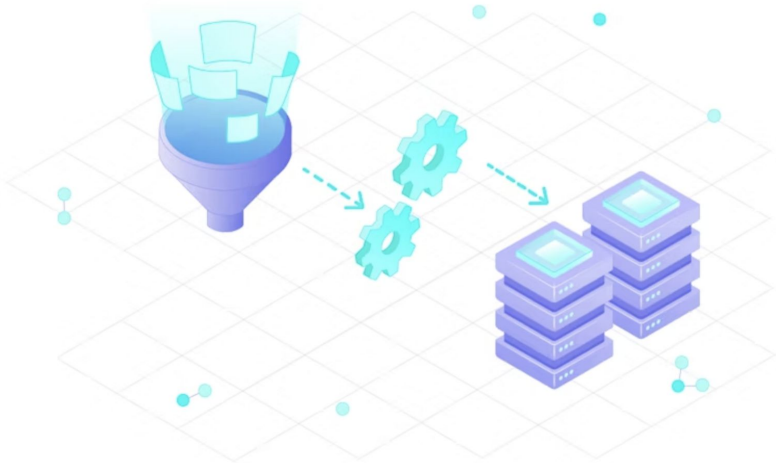

Strategy and planning

- Understand the use cases you are addressing
- Know your data
- Know what you are looking for to support your use cases
- Choose what's important - metrics
- Data sampling - understand your strategy and its implications

Data Interchange

- Have a solid robust plan
- Include metadata + timings
- Have a strategy for data recovery / data gaps
- Use the most appropriate tooling, eg S3 or pub/sub
- Establish clear requirements on both sides of the interchange
- Focus on efficiency and speed
- Don’t query raw data tables

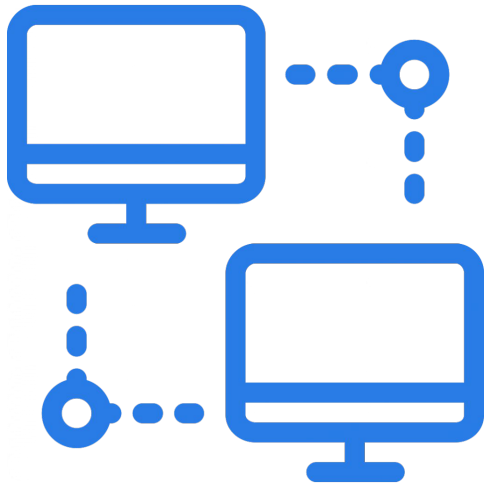

Custom data transformation pipeline

- What are the requirements: inputs → outputs
- Break pipeline up to discrete components
- Keep services simple and separate where possible
- Evaluate cost and speed
- High level of optimisation
  - Type data in optimal ways to speed the processing
  - Use appropriate data structures, eg Python hashmaps & _slots_
  - Use mem-cache whenever possible, eg Redis
  - Use async rest API's when practical, eg FastAPI with Redis caching
  - When using S3 (or similar) at scale look at async Get/Put
  - Build timing in from the start, time everything, experiment and optimize

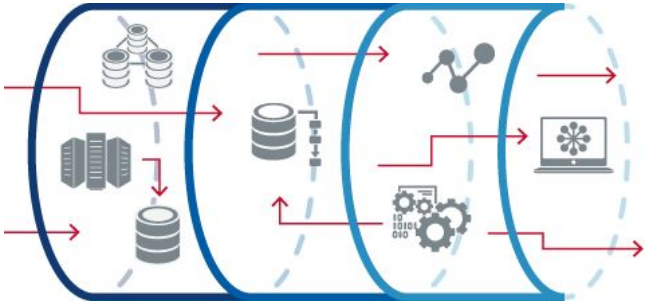

Data Storage

- Multi level storage
  - High speed access to summary level data (Mongo)
  - Easy fast access to summary history + pointers to details (Postgres)
  - Low cost storage for detail data infrequently accessed (S3)
- Use smart self managing data management - partitions & TTLs

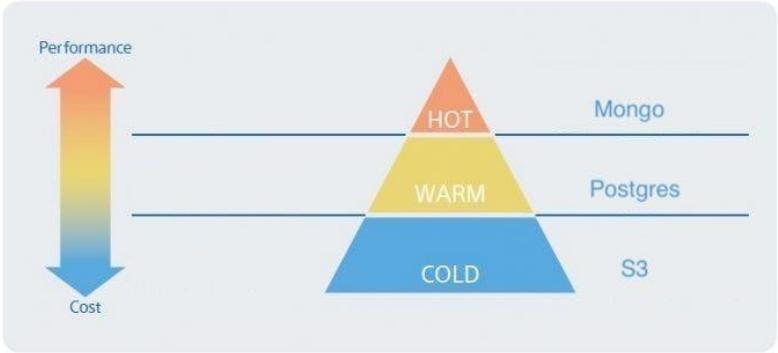

Custom dashboarding

- Design Intuitive visualisations
- Integrate context-sensitive links to related data and tools
- Store data in a visualisation friendly way

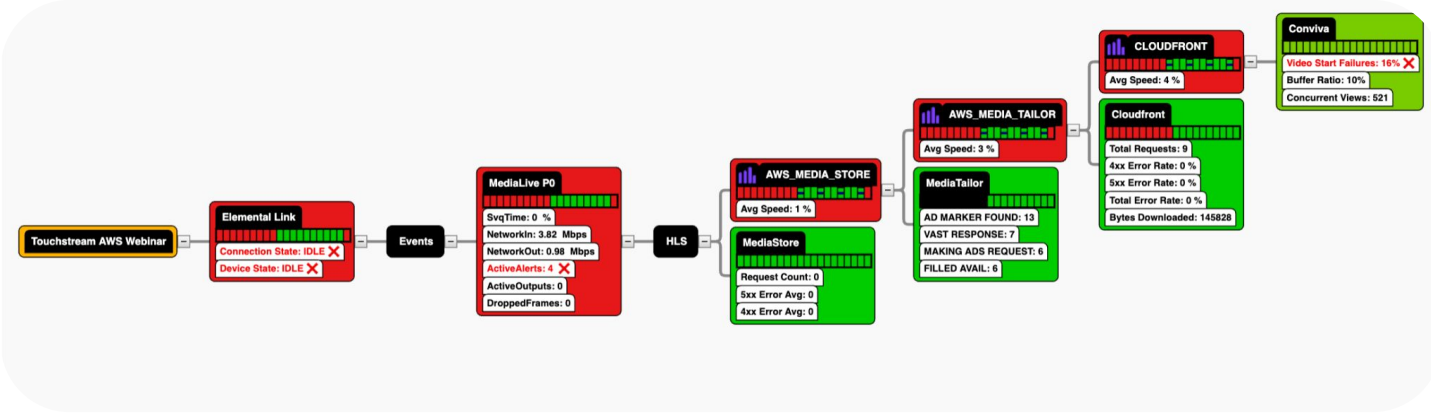

Simple Example

Needed to add real time metrics

- add second Akamai datastream
- use opensource tools to summarise and create metrics

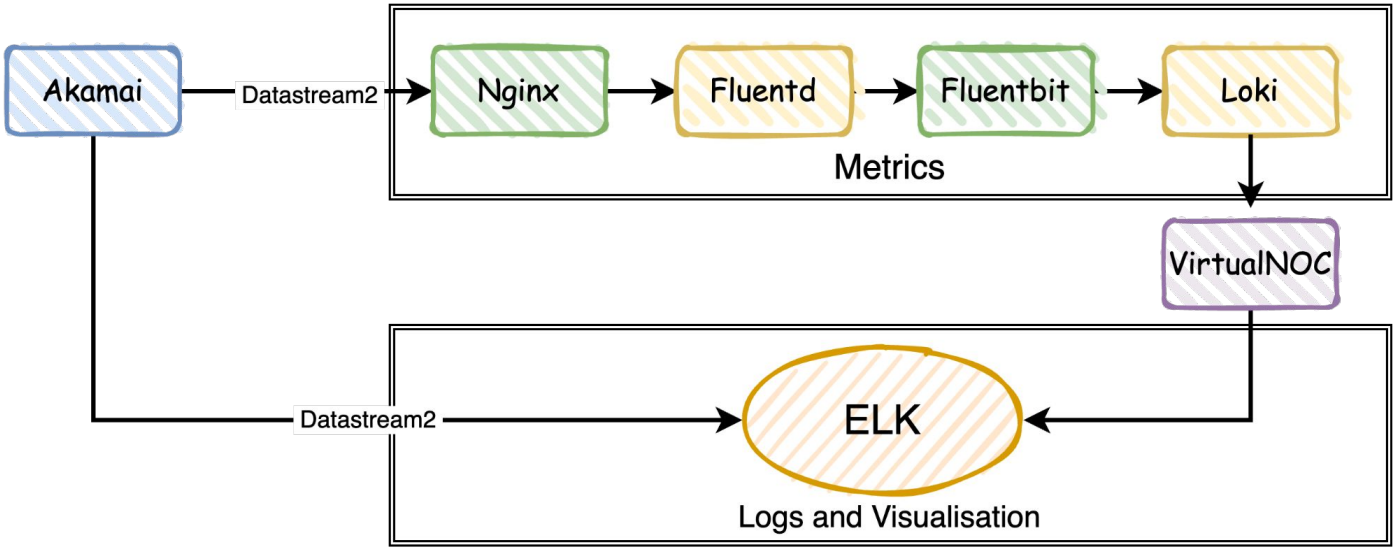

Lessons Learnt

- It's way harder than you think!
- Your prototype may not be the best design for production
  - Use prototyping phase to prove that you can extract valuable insights from data
  - Start again from blank page to design robust cost effective production model
- Understand the use cases better
  - Prove you can find valuable insights and that it solves problems
- Tried to solve too complex problem at the beginning
  - Better to break into smaller "bits"
- Using one tool for everything may not be ideal
  - Experiment with tools for specific purposes

Selected Tools

Alexander Leschinsky

G & L

Commercial challenges of general observability platforms

- Only use in reasonable shape and size
- Not for high-volume data - see following back-of-an-envelope calculations on list prices for a hypothetical daily 10 TB of raw log data
- Limited query options and visualizations
- Limited/expensive retention
- Limited ETL capabilities
- Add latency
- **Can serve a purpose to augment internal observability systems**
    - i.e. alert workflows and integrations

Commercial challenges: Datadog

Ingest

STARTING AT

$ 0.10

Per ingested or scanned GB, per month*

Ingest, process, live tail, and archive all logs

- **Enrich and structure log data**
- **Parse on ingestion**
- **Generate log-based metrics**
- **Self-hosted archives, with the option to rehydrate**
- **Dynamic index routing**

*Per GB of uncompressed data ingested for processing, or compressed data scanned for rehydrating.

Retain or Rehydrate

30-DAY RETENTION

$ 2.50

Per million log events, per month*

Retain logs based on their value and rehydrate from archives on-demand

- **Define log retention based on tags or facets**
- **Simplified pricing based on retention for better cost control**
- **Log patterns and analytics**
- **Log Rehydration™ for audits and historical analysis**

*Billed annually or $3.75 on-demand

Retain or Rehydrate

15-DAY RETENTION

- 3-day retention
- 7-day retention
- 15-day retention
- 30-day retention
- Greater than 30 days

based

-dema

d re

10 TB daily raw logs
= 300 TB monthly
= $30,000

1 mio log events ~ 0.25 GB raw logs
1 GB raw logs ~ 4 mio log events
300 TB monthly raw logs ~ **$3 mio**

Commercial challenges: New Relic

10 TB daily raw logs = 300 TB monthly
= **$89,970** / month for ingest

= **$149,950** / month for ingest

\+ User costs

	Pro	Enterprise
DATA COSTS (MONTHLY)		
Option 1: Original data ingest ⓘ 30 days retention	$0.30/GB beyond free 100 GB limit	$0.30/GB beyond free 100 GB limit
Option 2: Data Plus data ingest ⓘ 120 days retention	$0.50/GB beyond free 100 GB limit	$0.50/GB beyond free 100 GB limit (also includes FedRAMP Moderate ⓘ and HIPAA eligibility)
USER COSTS (MONTHLY) ⓘ		
Basic users	$0	$0
Core users	$49/user	$49/user
Full platform users	$349/user (for annual commitments) $418.80/user (for monthly pay as you go) ⓘ	$549/user (for annual commitments) $658.80/user (for monthly pay as you go) ⓘ

G&L

Commercial challenges: dynatrace

Ingest

10 TB daily raw logs = 300 TB monthly
= 279 396 GiB
= **$55,879.20**

\+ Retention

30 days retention
279,396 GiB x 30 x $0.0007 = **$5,867**

120 days retention
279,396 GiB x 120 x $0.0007 = **$23,469**

\+ Query

“Databases” that can work with cheap object storage

- Logs
  - **hydrolix.io**
    - Highest compression rate
    - Scale components independently
    - De-couple compute from object storage
  - Loki by Grafana: https://github.com/grafana/loki
  - BigQuery
- Metrics
  - Prometheus
    - **Cortex**: https://github.com/cortexproject/cortex (in use at G&L for metrics, mainly for multi-tenancy support)
    - Mimir: https://github.com/grafana/mimir
    - Thanos: https://github.com/thanos-io/thanos

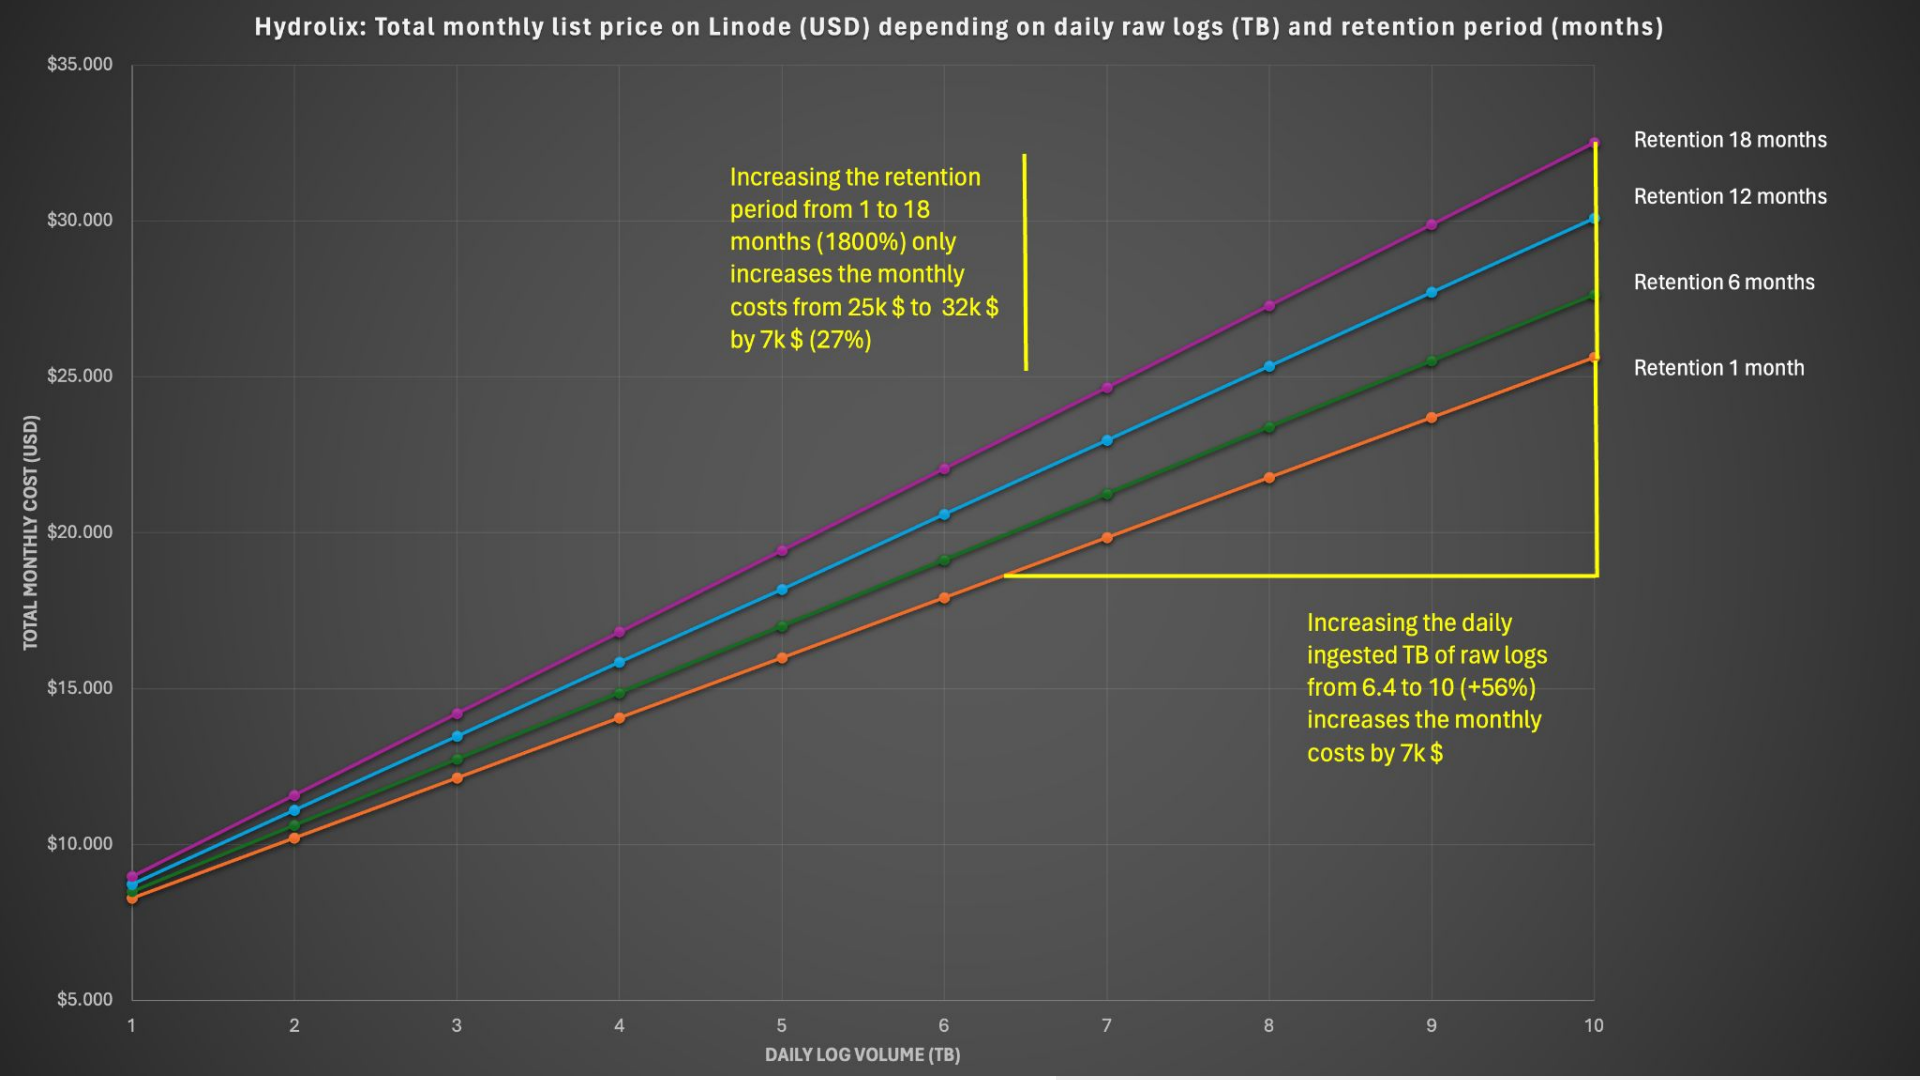
Hydrolix: Total monthly list price on Linode (USD) depending on daily raw logs (TB) and retention period (months)
$35.000
$30.000
$25.000
$20.000
$15.000
$10.000
$5.000
TOTAL MONTHLY COST (USD)
1
2
3
4
5
6
7
8
9
10
DAILY LOG VOLUME (TB)
Retention 18 months
Retention 12 months
Retention 6 months
Retention 1 month
Increasing the retention period from 1 to 18 months (1800%) only increases the monthly costs from 25k $ to 32k $ by 7k $ (27%)
Increasing the daily ingested TB of raw logs from 6.4 to 10 (+56%) increases the monthly costs by 7k $

Hydrolix (logs)

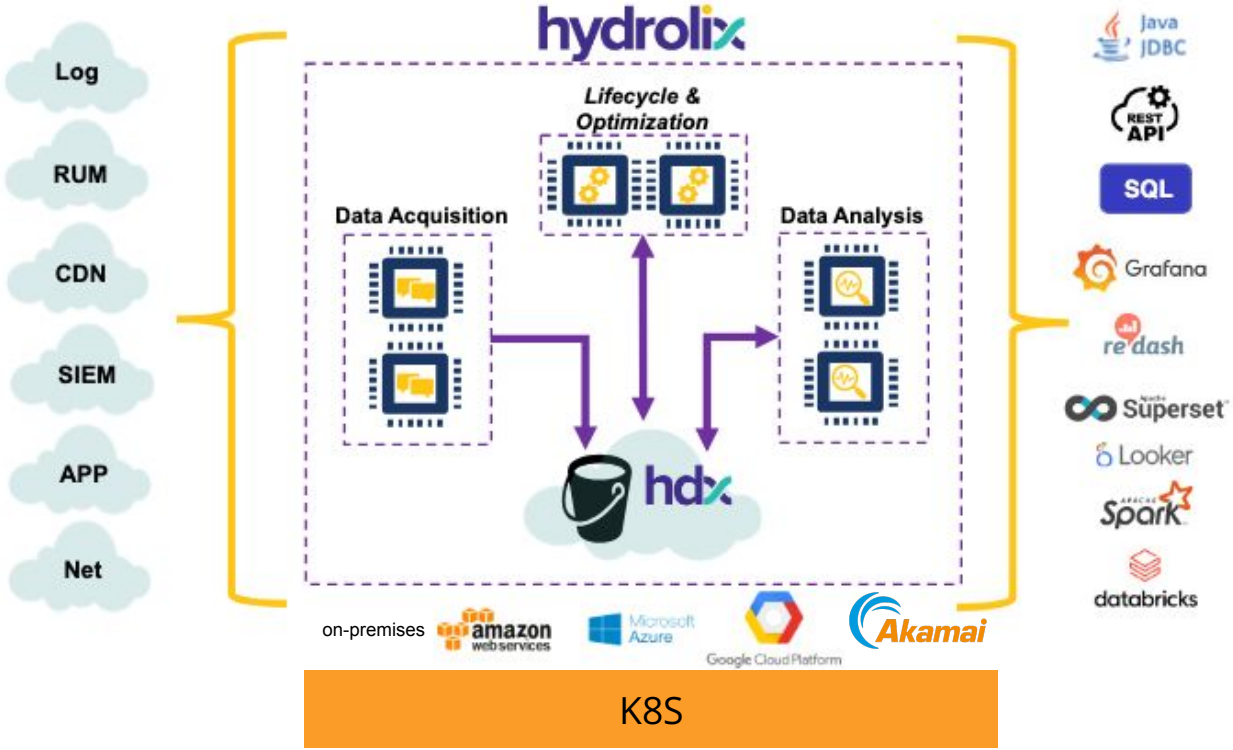

Hydrolix

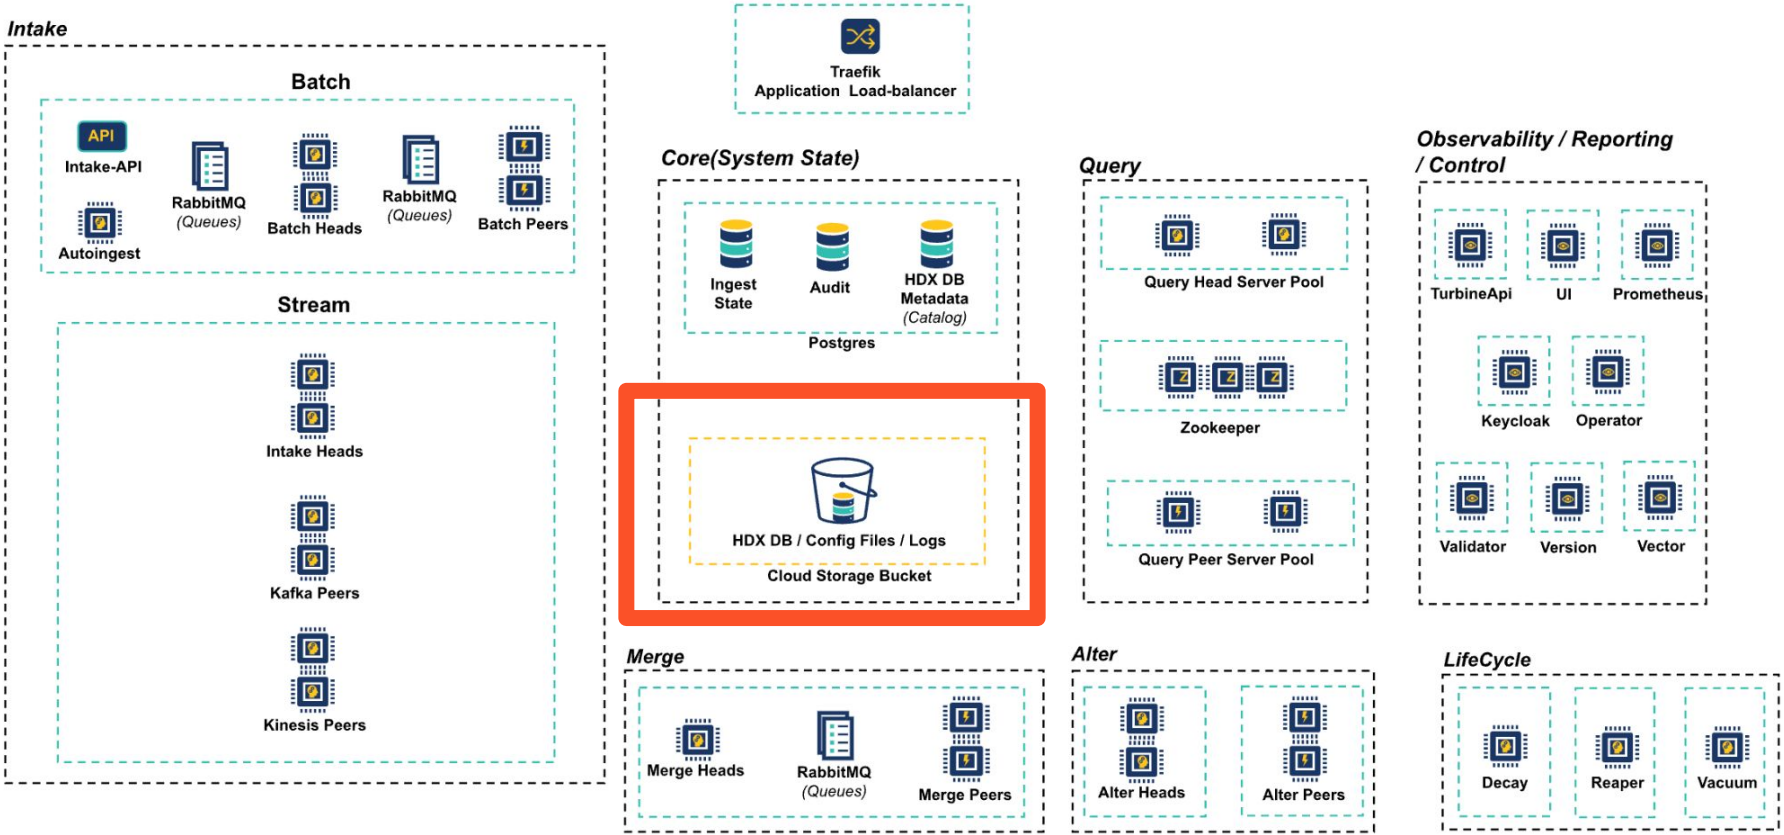
Intake
Batch
API
Intake-API
Autoingest
RabbitMQ
(Queues)
Batch Heads
RabbitMQ
(Queues)
Batch Peers
Stream
Intake Heads
Kafka Peers
Kinesis Peers
Traefik
Application Load-balancer
Core(System State)
Ingest
State
Audit
HDX DB
Metadata
(Catalog)
Postgres
HDX DB / Config Files / Logs
Cloud Storage Bucket
Query
Query Head Server Pool
Zookeeper
Query Peer Server Pool
Observability / Reporting
/ Control
TurbineApi
UI
Prometheus
Keycloak
Operator
Validator
Version
Vector
Merge
Merge Heads
RabbitMQ
(Queues)
Merge Peers
Alter
Alter Heads
Alter Peers
LifeCycle
Decay
Reaper
Vacuum

Cortex (metrics)

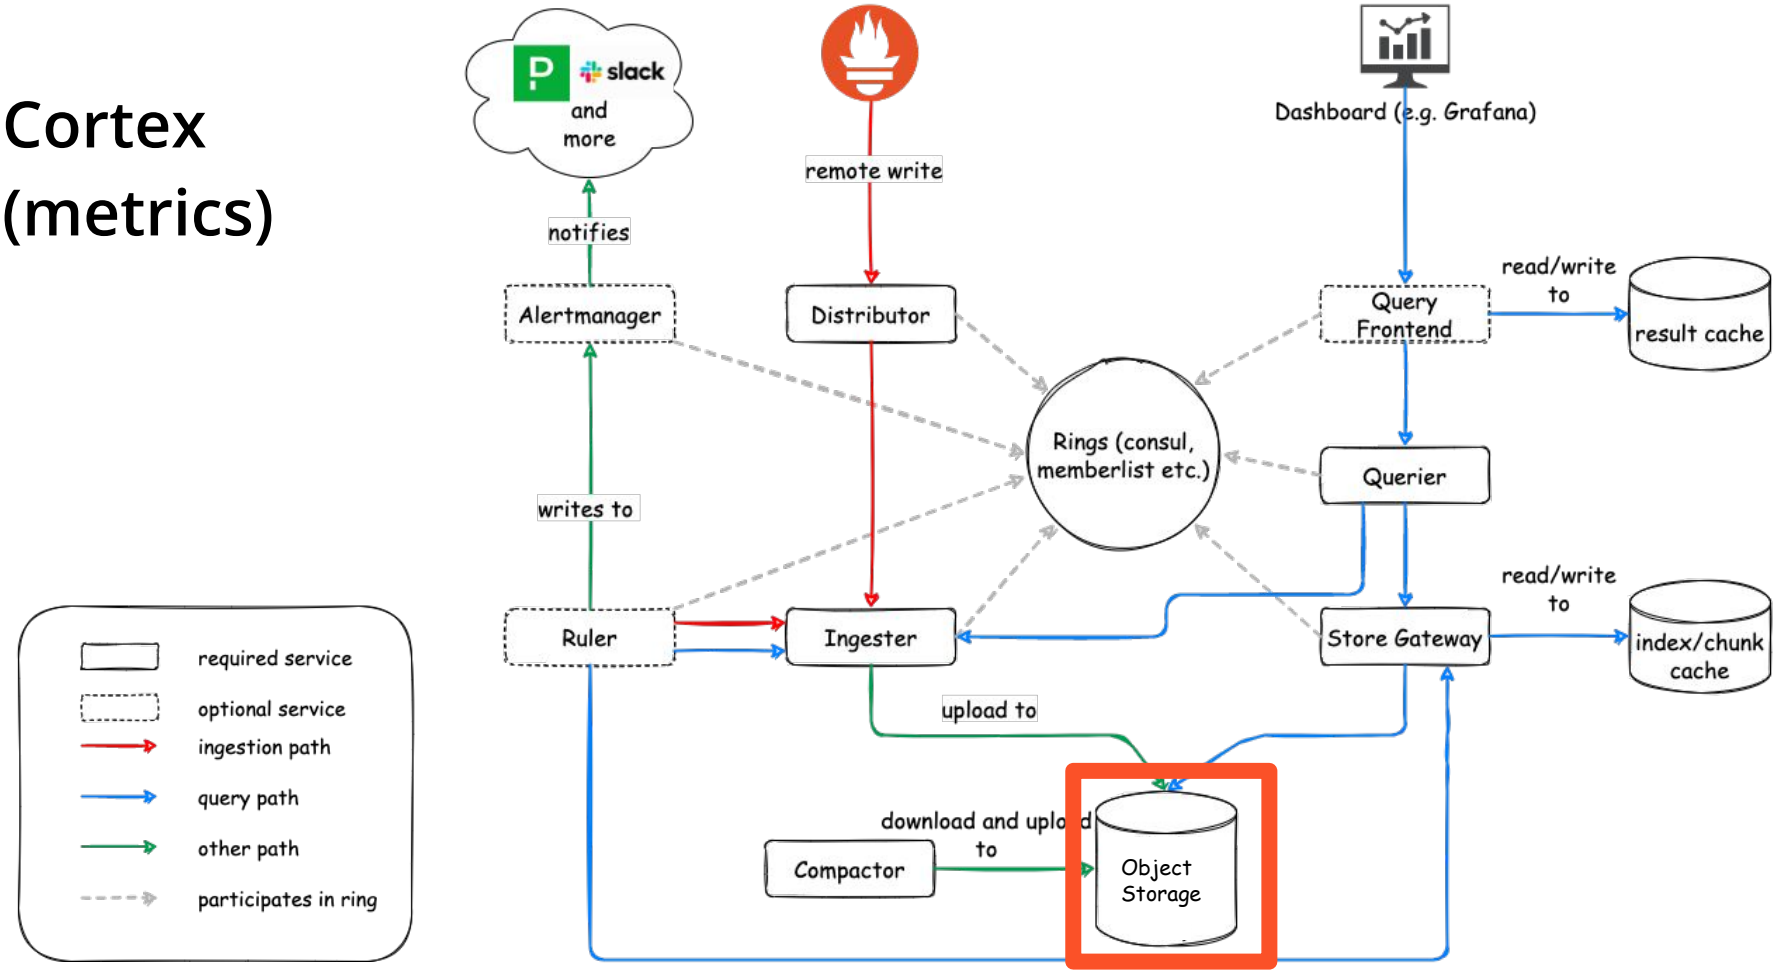

Sending Hydrolix Metrics to Cortex

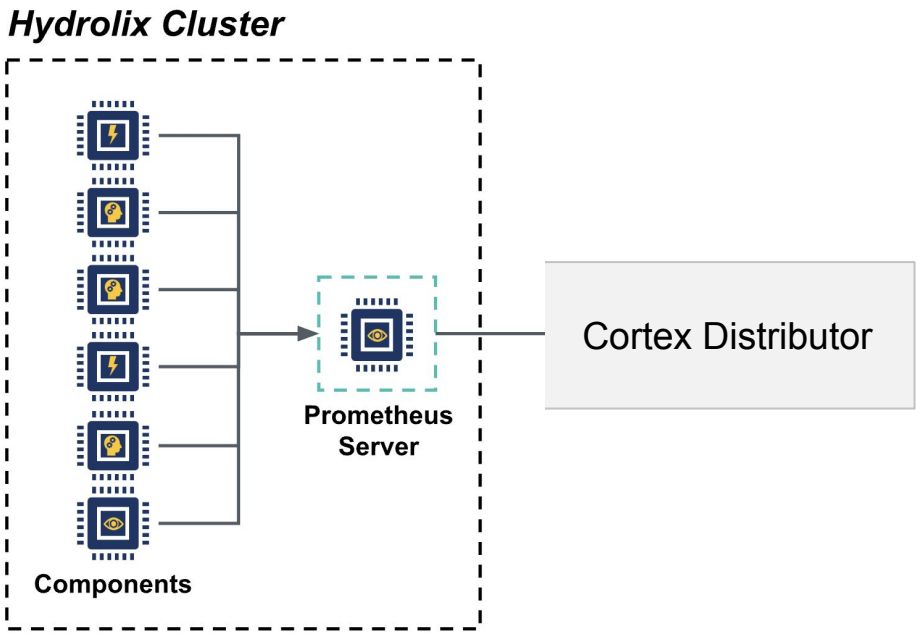

quesma.com

- Bridges ES tooling with Hydrolix and others
- Use ES queries, Kibana, Beats/Logstash with Hydrolix/Clickhouse
- Works in G&L PoC with Hydrolix

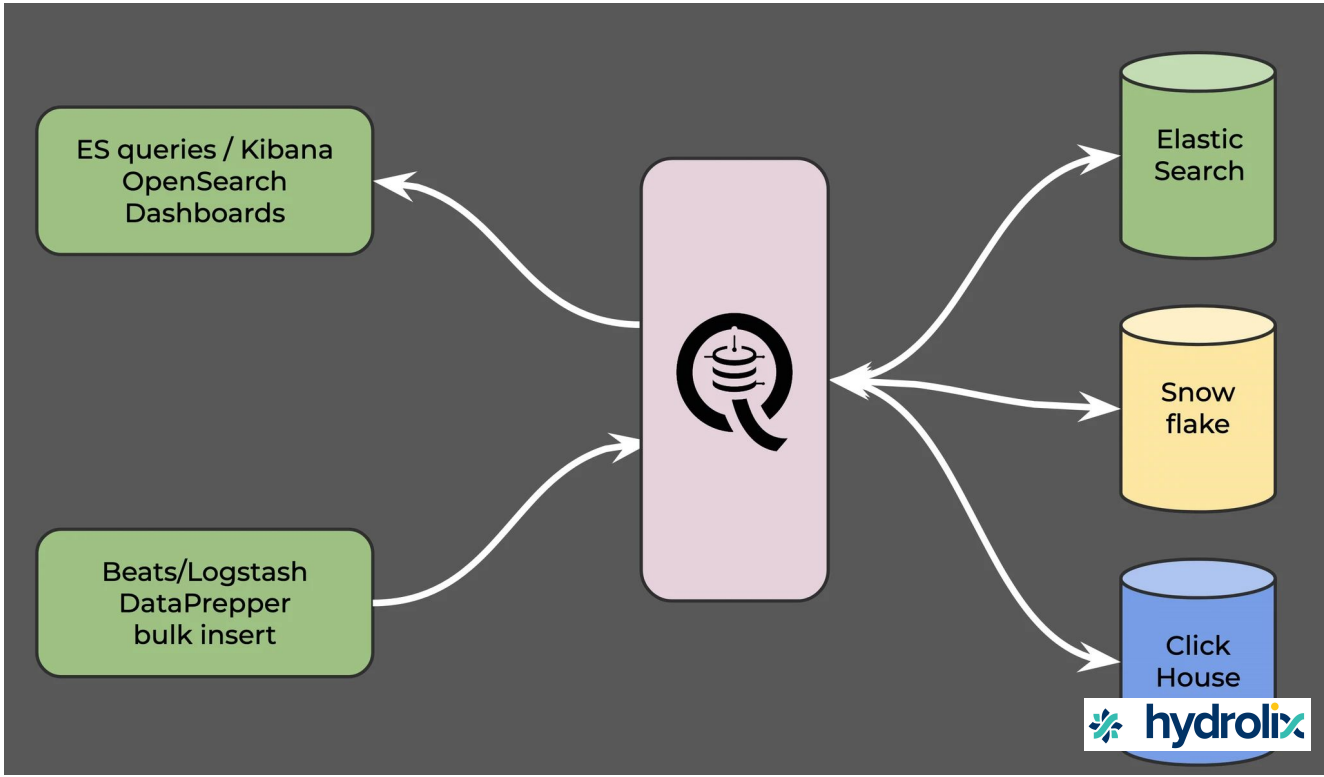

Specific G&L setup for logs

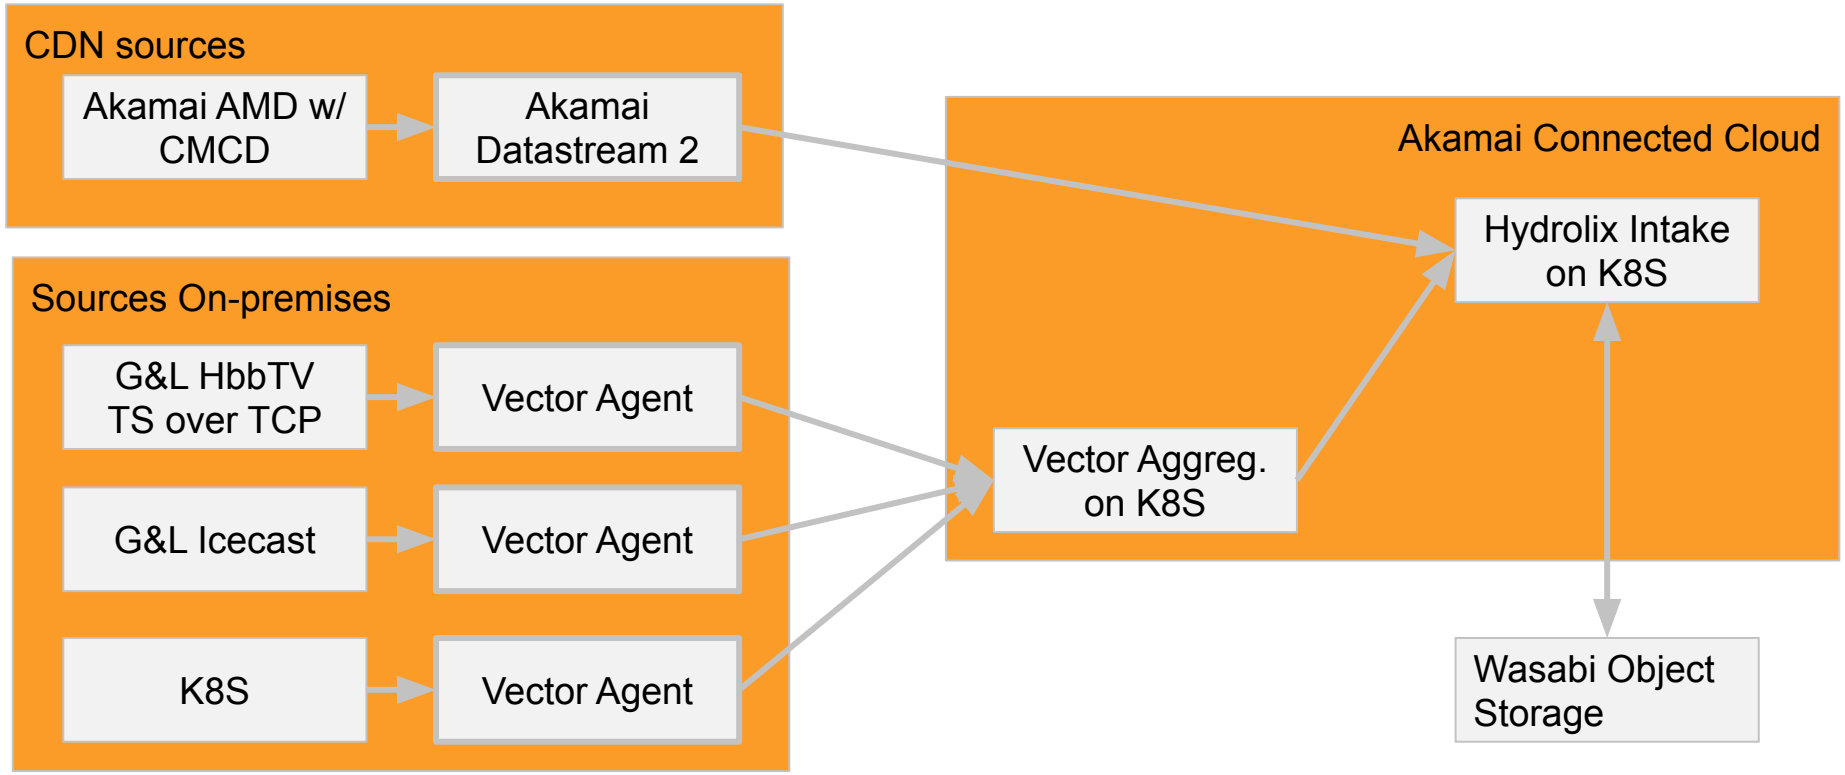

Specific G&L setup for logs

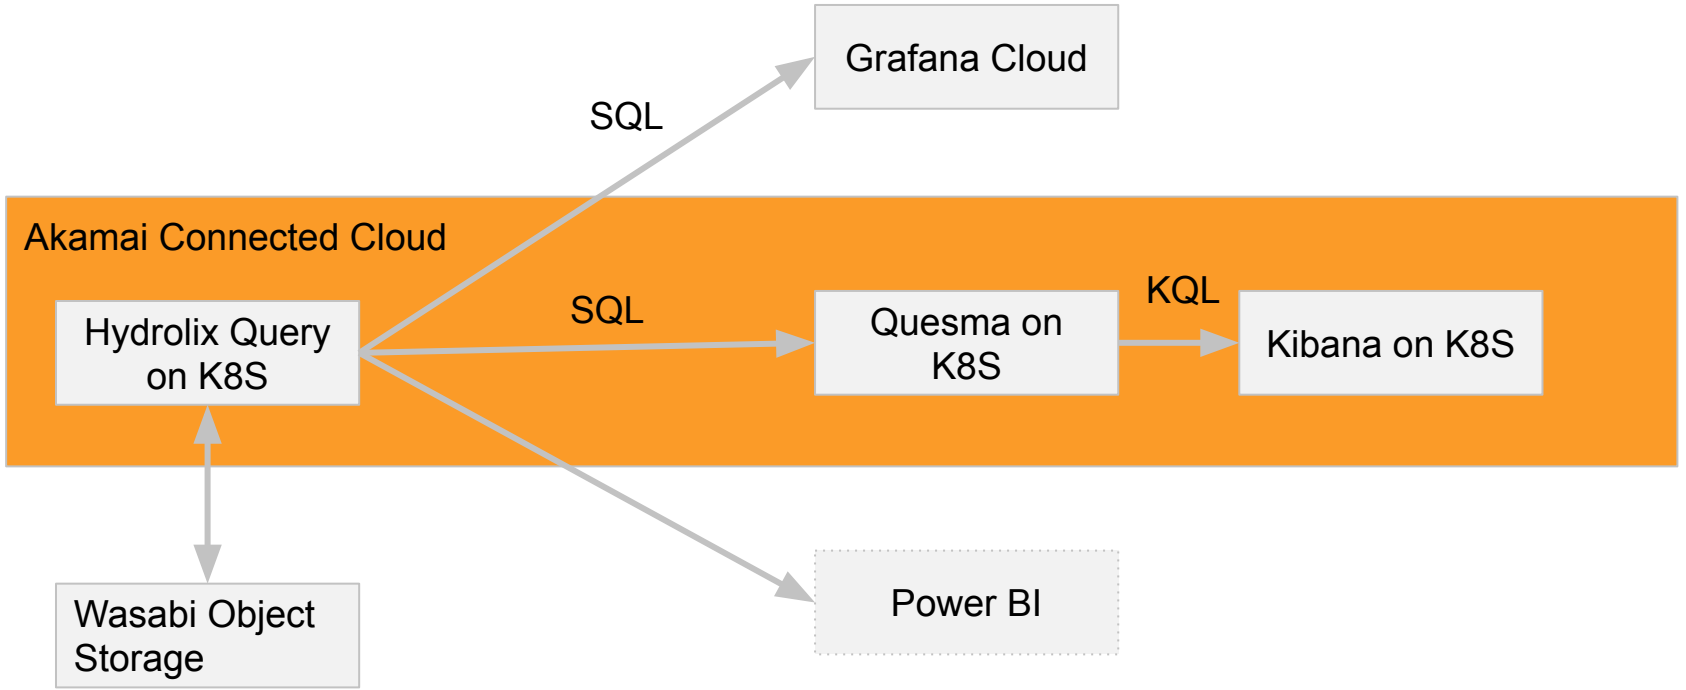

Q&A - Sli.do

All speakers

G&L

Thank you

More questions?

contact@gl-systemhaus.de