DEVELOPERWEEK

Jae Kim, LogDNA - Sales Engineer

Logging - start to finish!

Loglines:

Create
Collect
Send
Parse
Store
View
More!

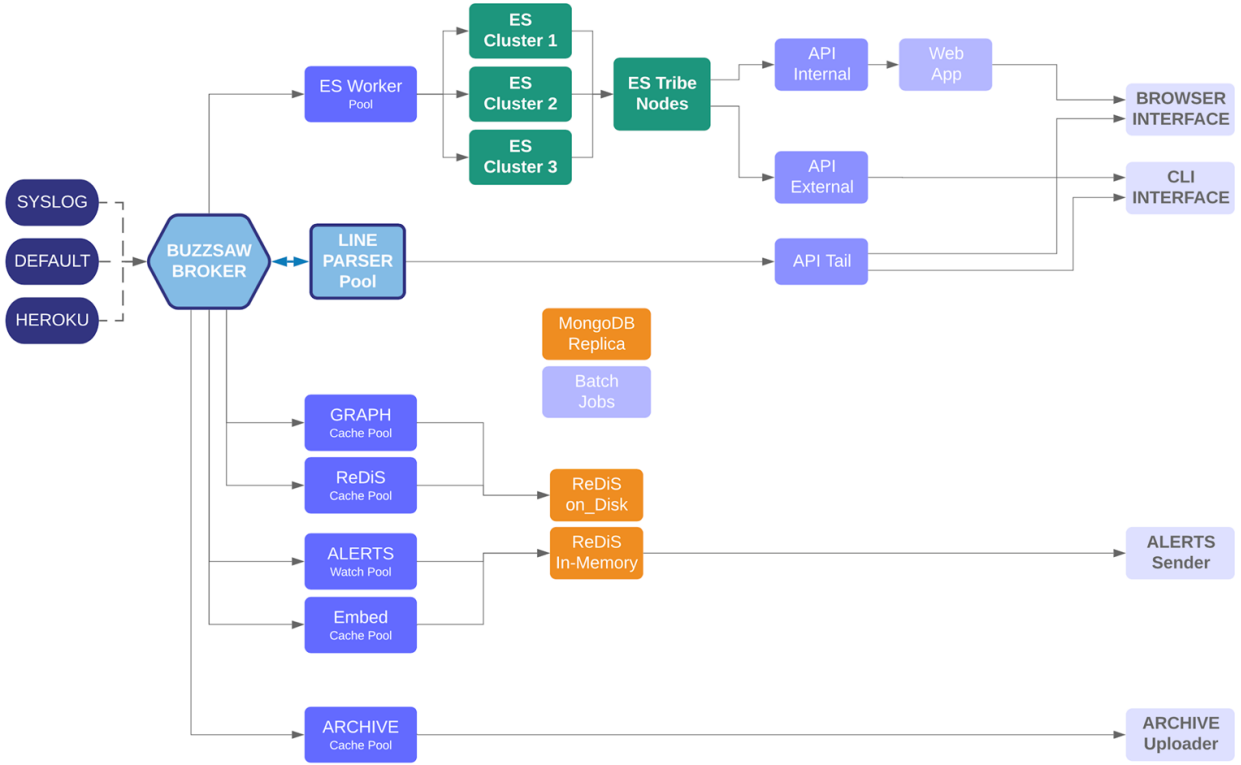

Create your logline

- Application type
  - Monolithic
    - Logs written to disk, inline
    - Limited primary insight into application state
  - Distributed / microservice
    - Event-based
- Plain language vs Structured logs
- Levels (debug > info > error > FATAL)
- Types (system, application, audit, security, …)

Collect your logline

- syslog
  - Compare to SNMP
- OS level (Linux, Win)
  - /var/log
  - %windir%/system32/...
- Platform level (Docker, K8S, etc)
  - Logspout
  - STDOUT and STDERR
- Environment level (AWS Cloudwatch, Azure Event Hub, etc)
  - Cloudwatch capture from: EC2, DynamoDB, S3, ECS, EKS…

Aggregate / stream your logline

- Local logs
  - printf, et al
- Stream to a repository!
  - Local or remote
  - Depends on data type / source
  - Efficiencies on sending side

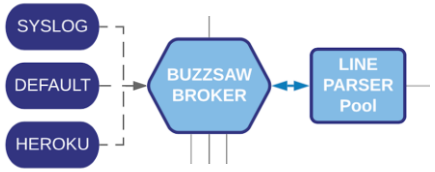

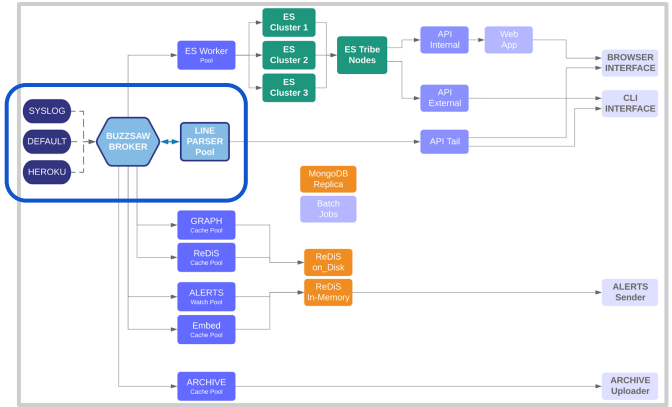

Parse / route / process your logline

- Messaging
  - Message-queue vs Pub-Sub
  - Synchronous vs async
    - Use both!
- Parsing / searching
  - grok / REGEX
  - Parse to create field:value objects
- Alerting
- Archiving
- ...

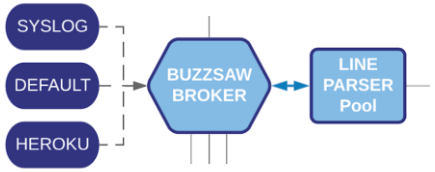

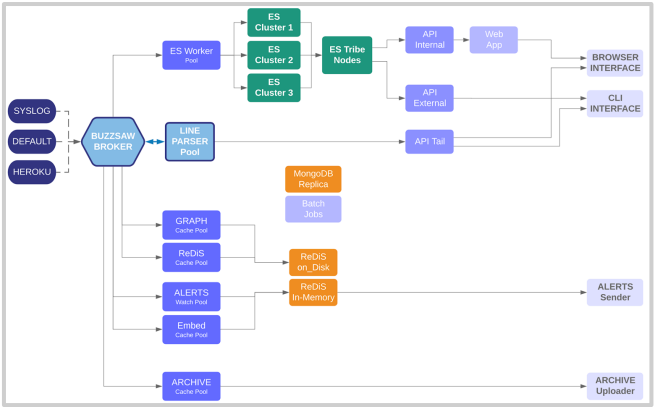

Parse / route / process

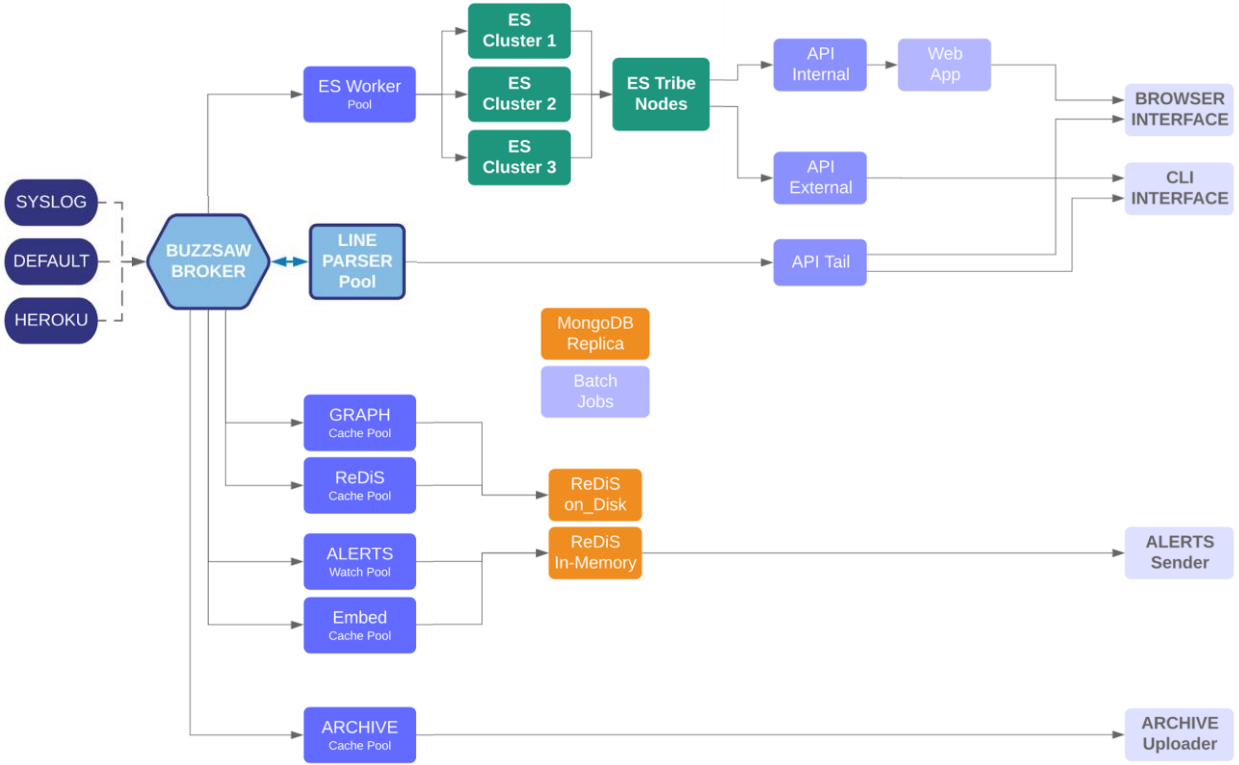

Store your logline

- Log volume?
  - MB of data > write to file(s)
  - GB of data > write to DB
  - TB of data > write to NoSQL
    - <-----> scalability
    - very fast search
- Other considerations
  - Distribution > sharding / replicas

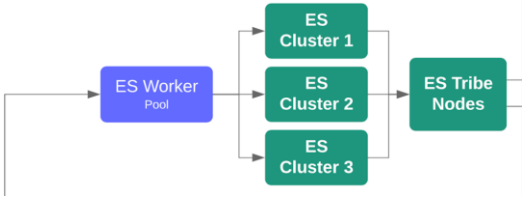

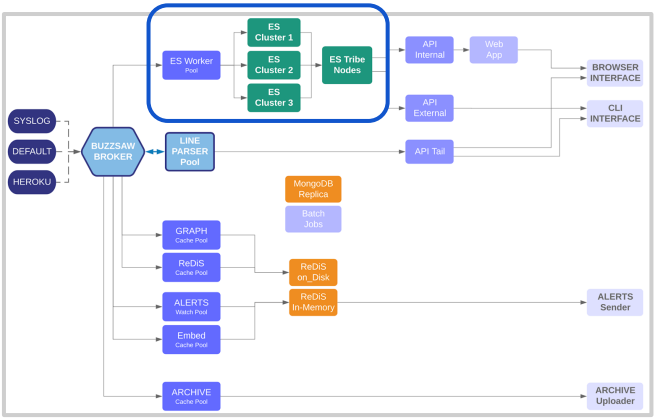

View your logline

- Search
  - grok / REGEX
  - Full-text search
- Visualize
  - Graphing
  - Dashboards

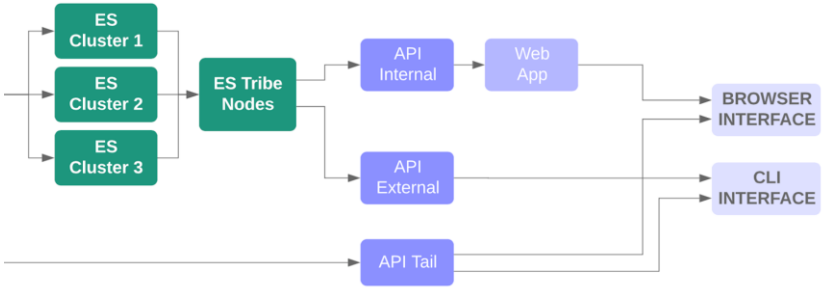

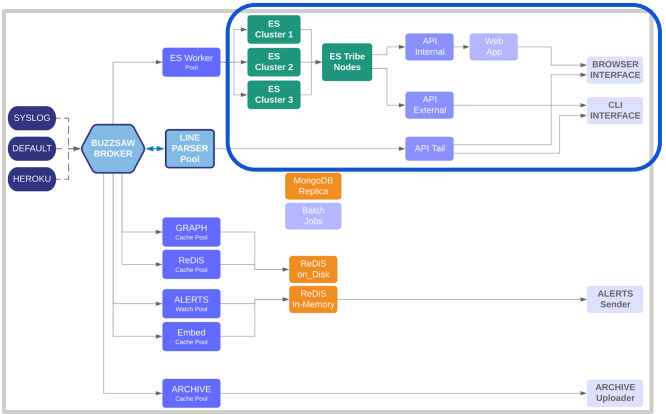

More!

- Data analytics
- Advanced visualization
  - Tableau, et al
- Next chapter of DB technology

Resources

- https://charity.wtf/2019/02/05/logs-vs-structured-events/
- https://logdna.com/bring-structure-to-your-logs-with-custom-parsing-on-logdna/
- https://www.elastic.co/guide/en/logstash/current/plugins-filters-grok.html

Q&A

Find us at Booth 208!

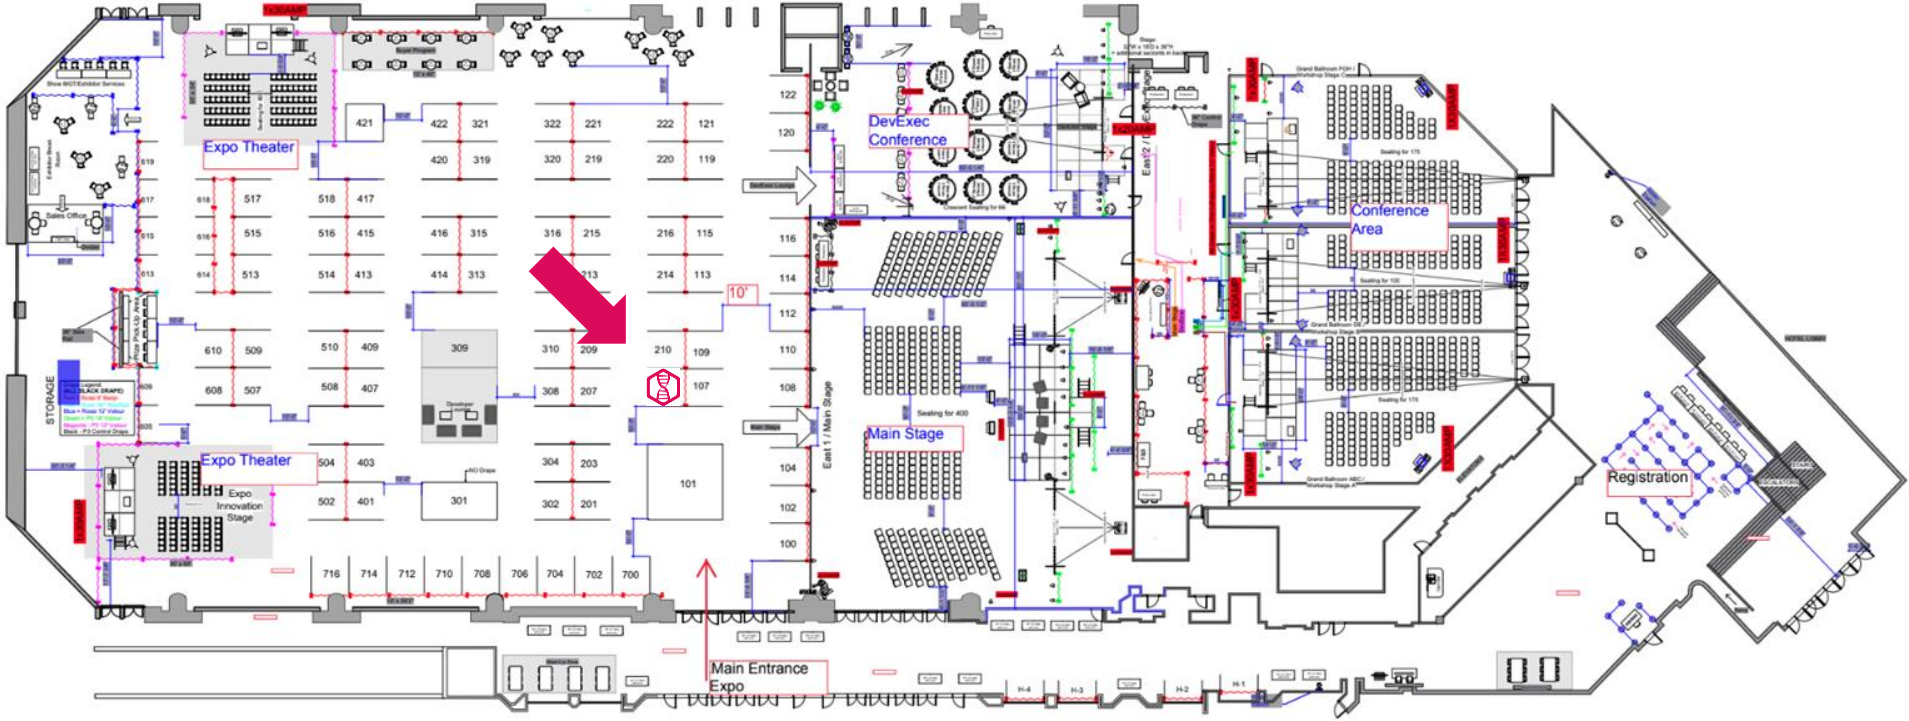

thank you!

Sponsored Session

Jae Kim

Systems Engineer, LogDNA

Jae wears several hats at LogDNA, including Sales/Solutions Engineer and Product Evangelist. He comes from Zerto and DellEMC where he consulted on virtualization data storage/protection solutions.

On the Nature of Logging

Log aggregation is no longer cutting-edge, rather a needed part of the infrastructure toolkit, as well as a prudent business decision. How were logs leading to actionable conclusions in the past and how do we do it today? From log collector to aggregator to ingester, from parsing to indexing and storage, what are the top pain points and how can we do it [better]?

And what's next...?

Booth MockUp

Booth #208