fluentbit

Log Forwarding at Scale

Scale 15X, Pasadena, CA

Eduardo Silva (@edsiper)
eduardo@treasure-data.com

TREASURE DATA

/about

- Open Source Engineer at Treasure Data
- Repositories / Projects
  - github.com/edsiper
  - fluentbit.io
  - duda.io
  - monkey-project.com

“Logging is Simple”

TREASURE DATA

“Logging is Simple”

“Logging is Simple”

Logging exists because of Analysis needs

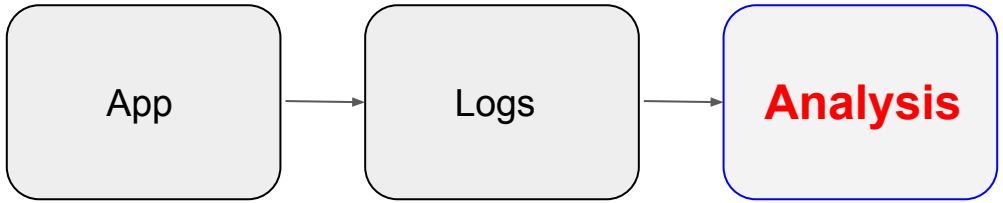

Before Analysis

Someone have to do some work

In a galaxy not so far away...

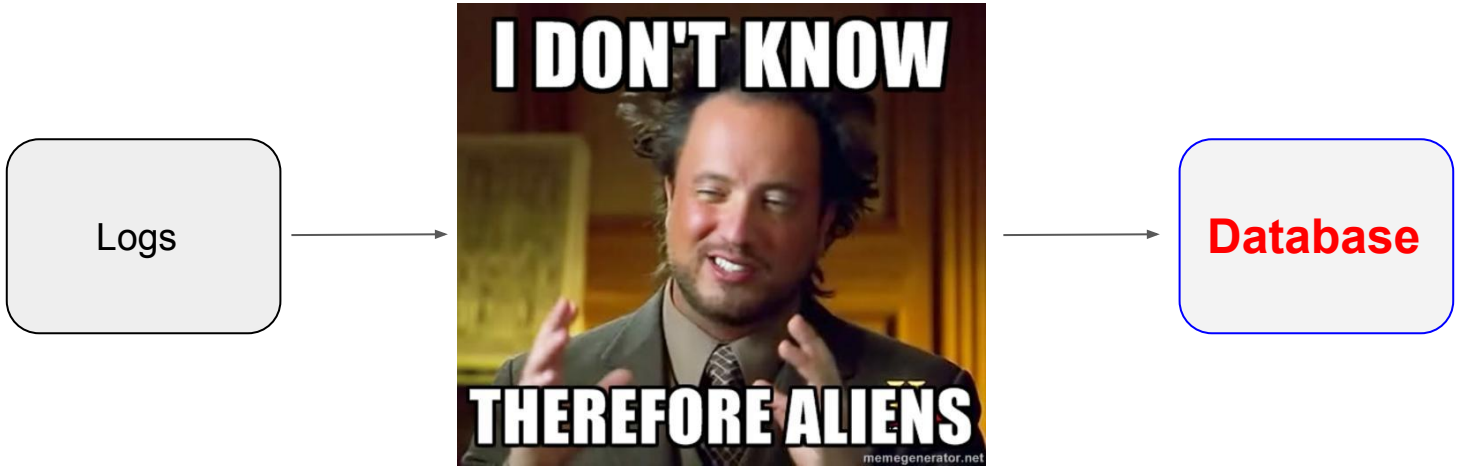

Analysis

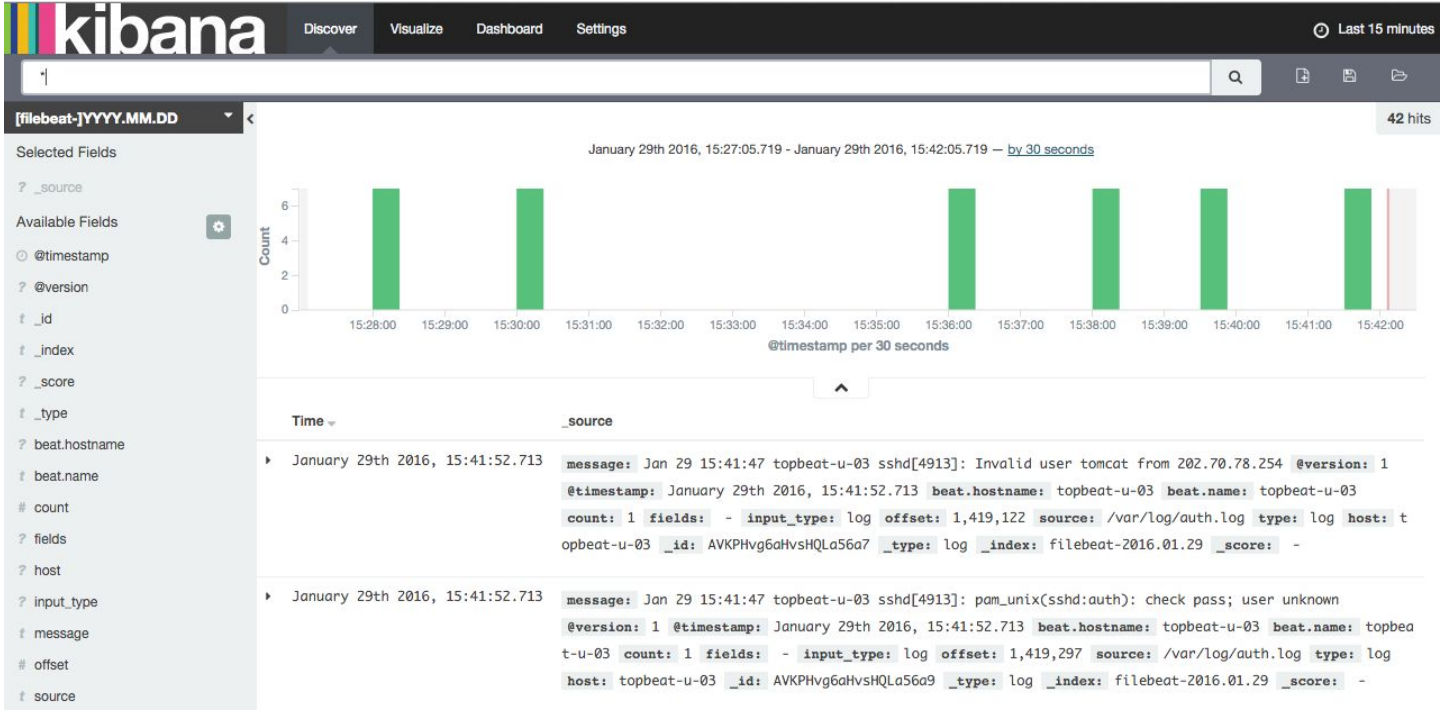
kibana
Discover
Visualize
Dashboard
Settings
Last 15 minutes
*
[filebeat-]YYYY.MM.DD
Selected Fields
? _source
Available Fields
@timestamp
? @version
t _id
t _index
? _score
t _type
? beat.hostname
t beat.name
count
? fields
? host
? input_type
t message
offset
t source
42 hits
January 29th 2016, 15:27:05.719 - January 29th 2016, 15:42:05.719 — by 30 seconds
Count
6
4
2
0
15:28:00
15:29:00
15:30:00
15:31:00
15:32:00
15:33:00
15:34:00
15:35:00
15:36:00
15:37:00
15:38:00
15:39:00
15:40:00
15:41:00
15:42:00
@timestamp per 30 seconds
Time
_source
January 29th 2016, 15:41:52.713
message: Jan 29 15:41:47 topbeat-u-03 sshd[4913]: Invalid user tomcat from 202.70.78.254 @version: 1 @timestamp: January 29th 2016, 15:41:52.713 beat.hostname: topbeat-u-03 beat.name: topbeat-u-03 count: 1 fields: - input_type: log offset: 1,419,122 source: /var/log/auth.log type: log host: topbeat-u-03 _id: AVKPHvg6aHvsHQLa56a7 _type: log _index: filebeat-2016.01.29 _score: -
January 29th 2016, 15:41:52.713
message: Jan 29 15:41:47 topbeat-u-03 sshd[4913]: pam_unix(sshd:auth): check pass; user unknown @version: 1 @timestamp: January 29th 2016, 15:41:52.713 beat.hostname: topbeat-u-03 beat.name: topbeat-u-03 count: 1 fields: - input_type: log offset: 1,419,297 source: /var/log/auth.log type: log host: topbeat-u-03 _id: AVKPHvg6aHvsHQLa56a9 _type: log _index: filebeat-2016.01.29 _score: -

Internally, Logging is **not** Simple

Scale Logging Requires **Understanding**

Logging Pipeline

Logging Pipeline

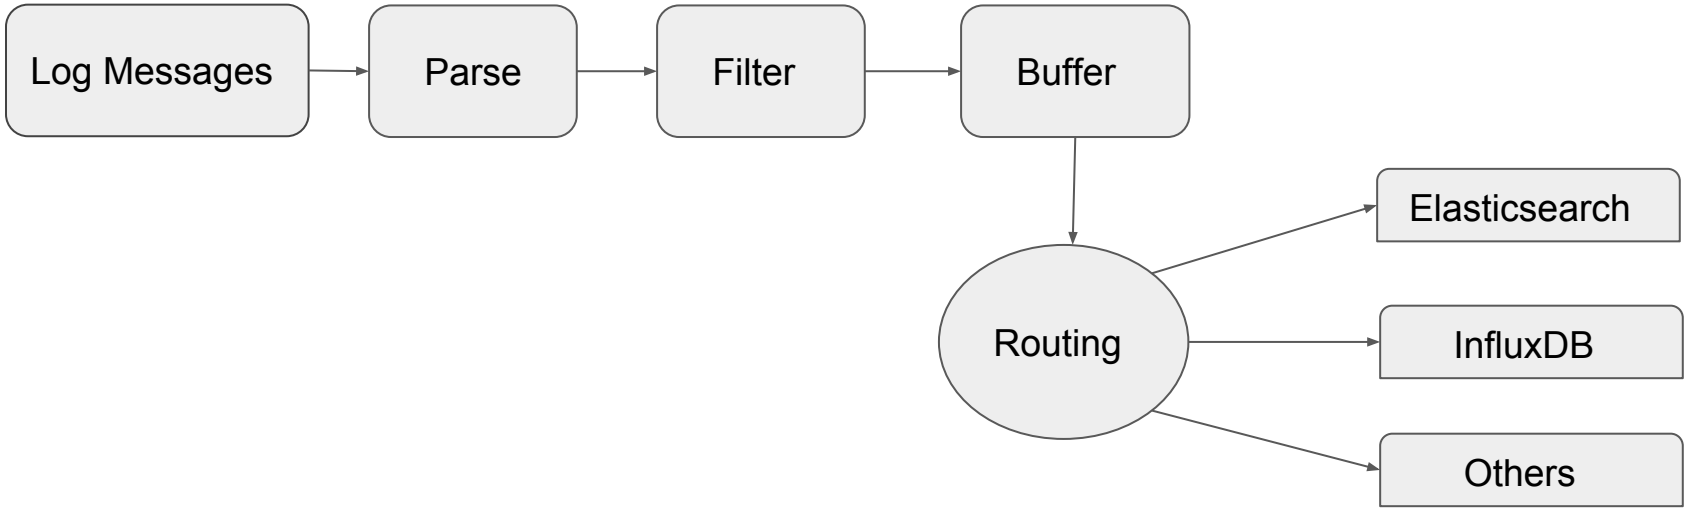

Logging Pipeline

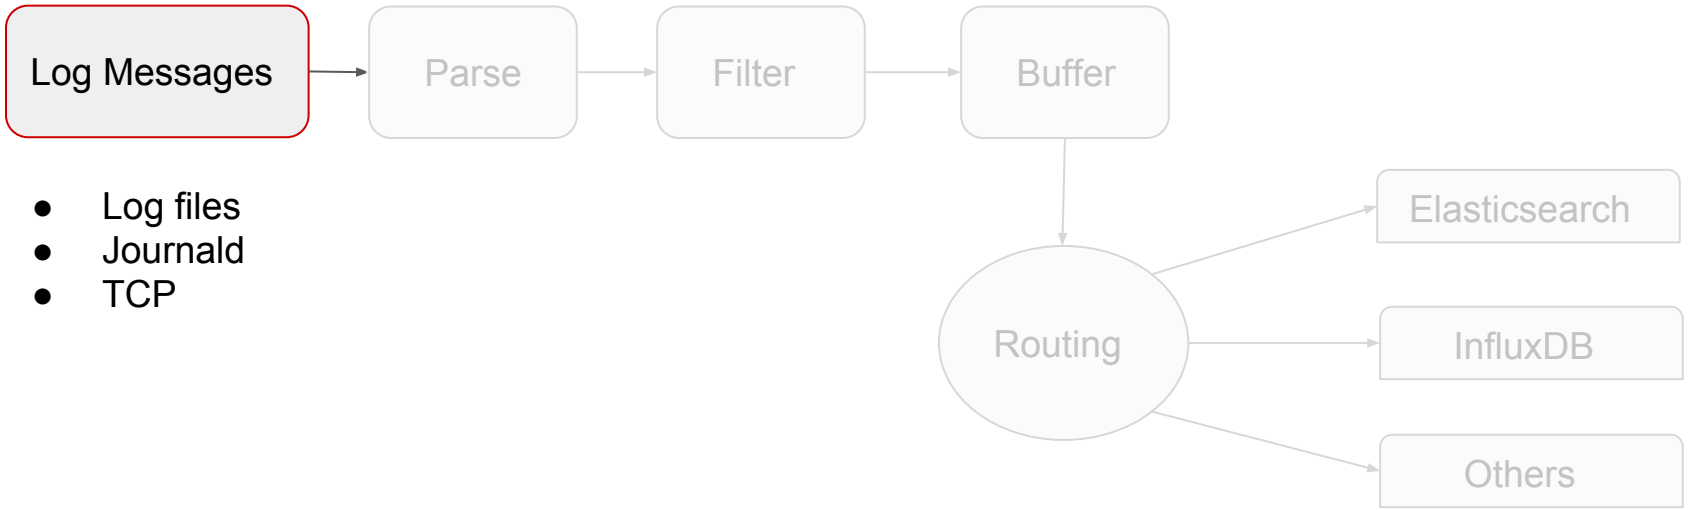

- Log files
- Journald
- TCP

Logging Pipeline

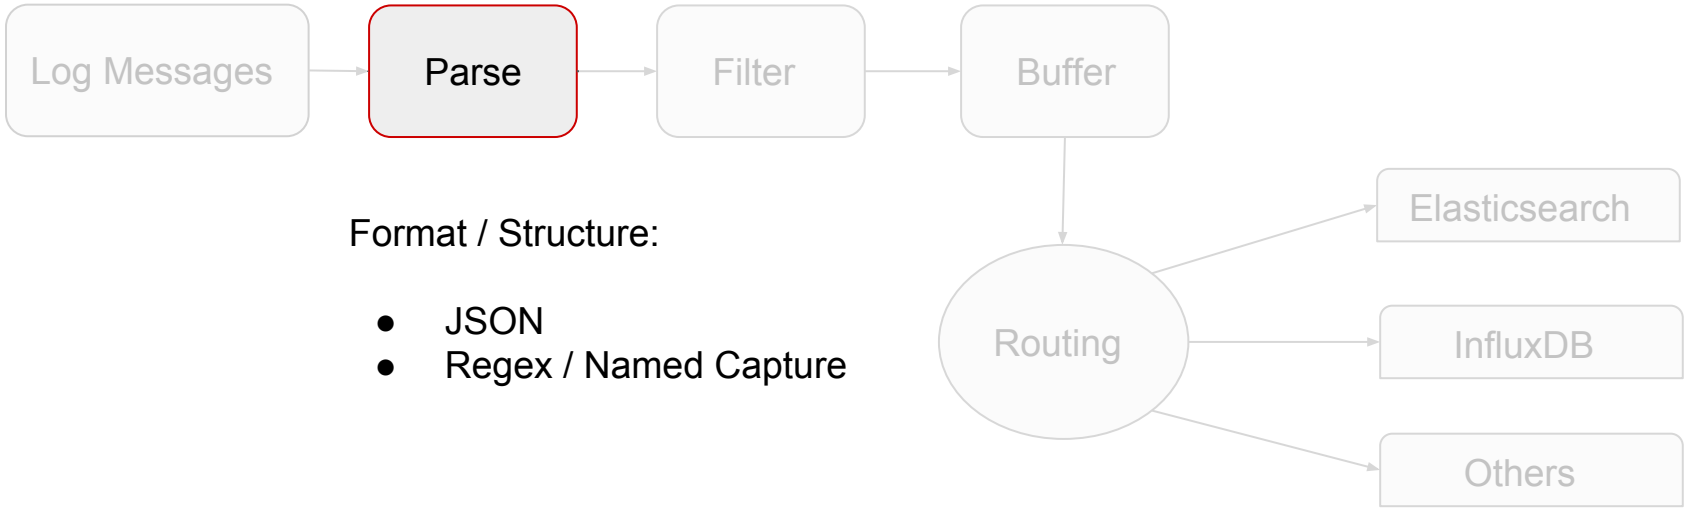

Format / Structure:

- JSON
- Regex / Named Capture

Logging Pipeline

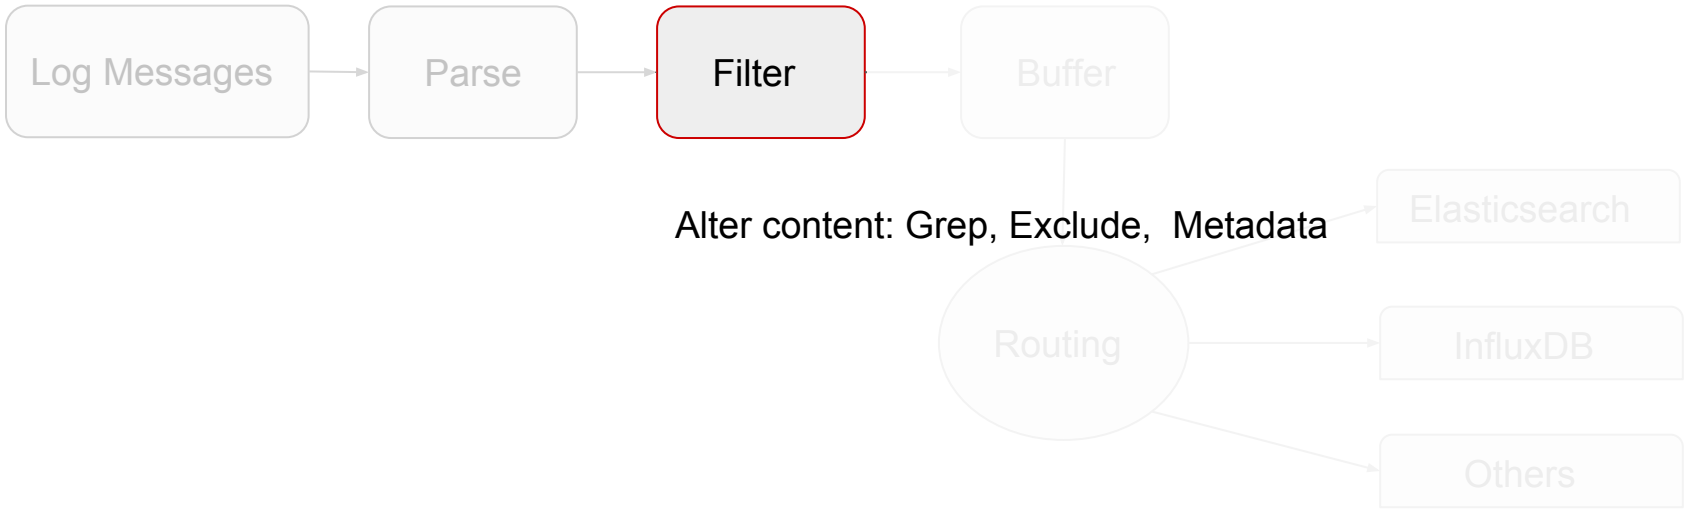

Logging Pipeline

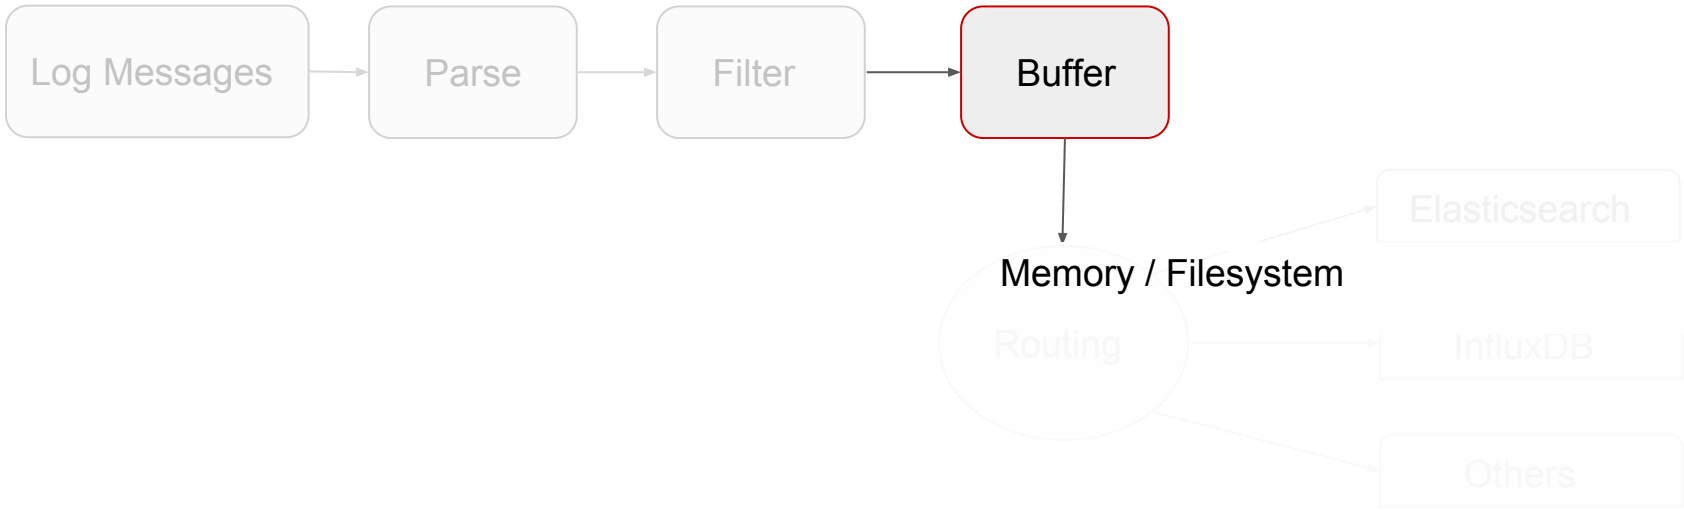

Logging Pipeline

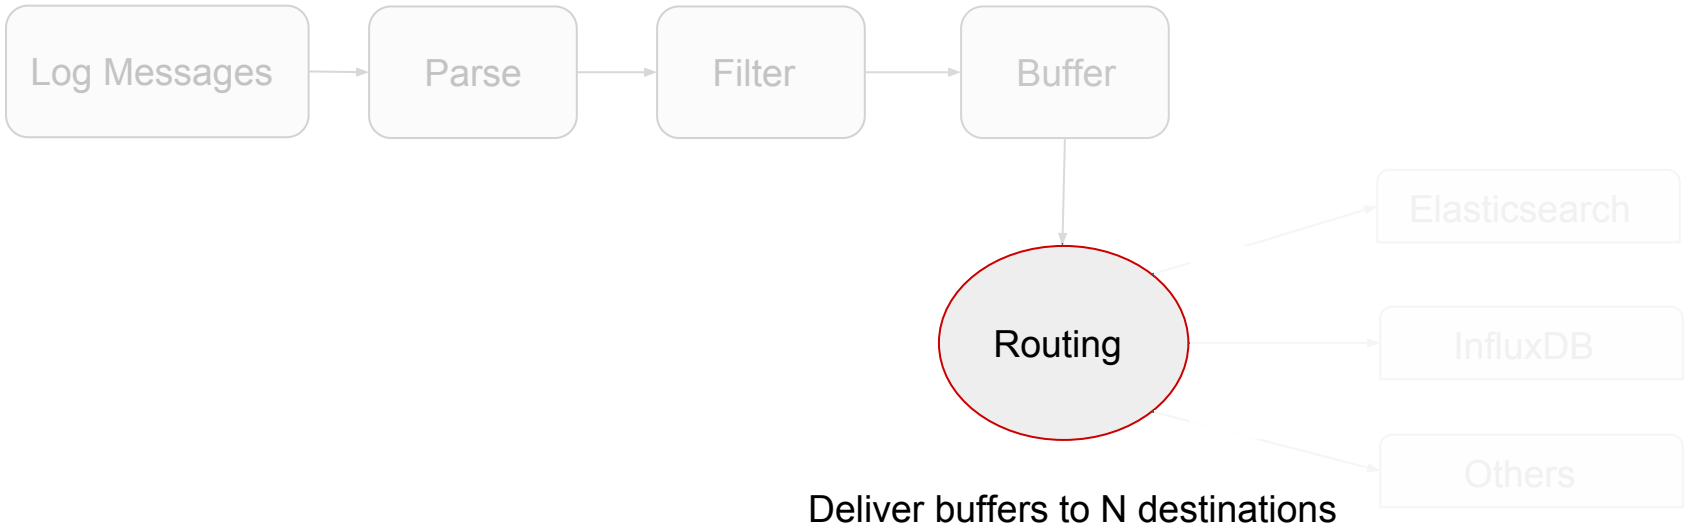

Deliver buffers to N destinations

Logging Pipeline

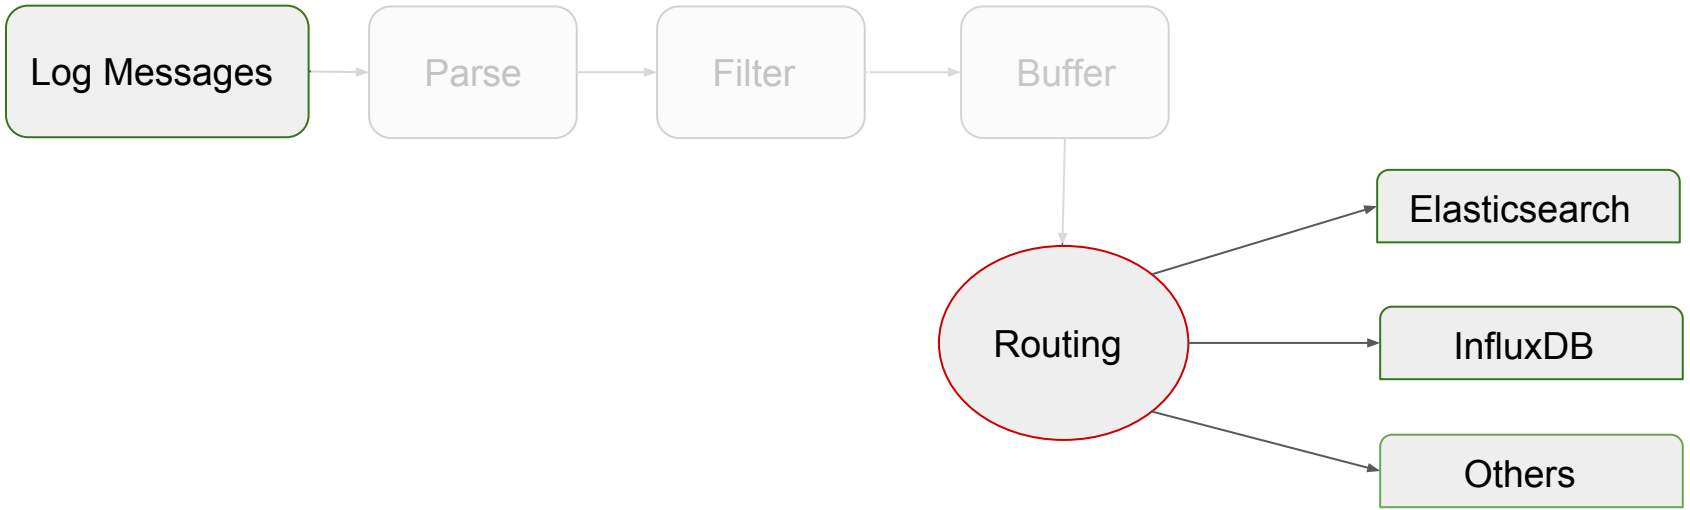

Logging Pipeline

- How to deal with the Logging Pipeline ?
- Is there any solution around ?


fluentd

About Fluentd

- Created by

TREASURE DATA

- Now hosted at

CLOUD NATIVE
COMPUTING FOUNDATION

About Fluentd

- More than 600 plugins available
- Pluggable Architecture
- Built-in Reliability
- Full integration with Docker and Kubernetes
- Written in Ruby + C

Fluentd Modes

- Log **F**orwarder
- Log **A**ggregator

Log Aggregator =
(Forwarder + Buffering Capabilities)

Edge Nodes / Forward to Aggregators

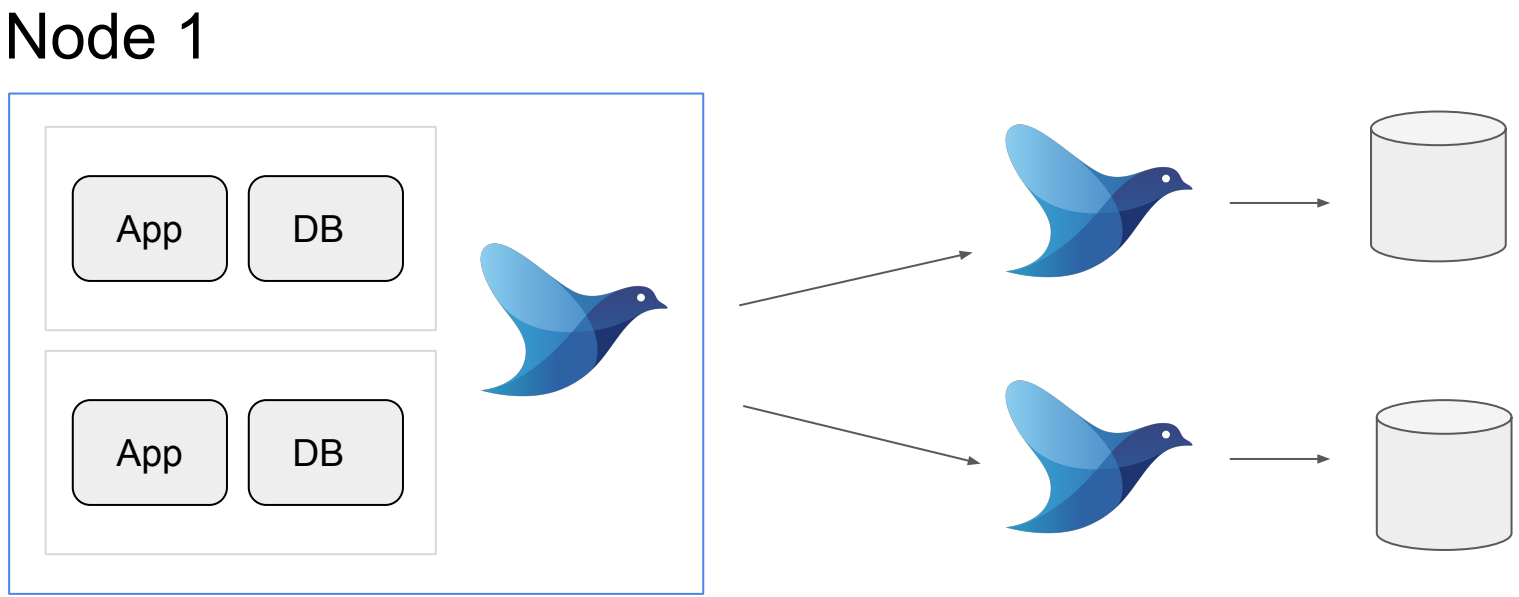

Edge Nodes & Costs

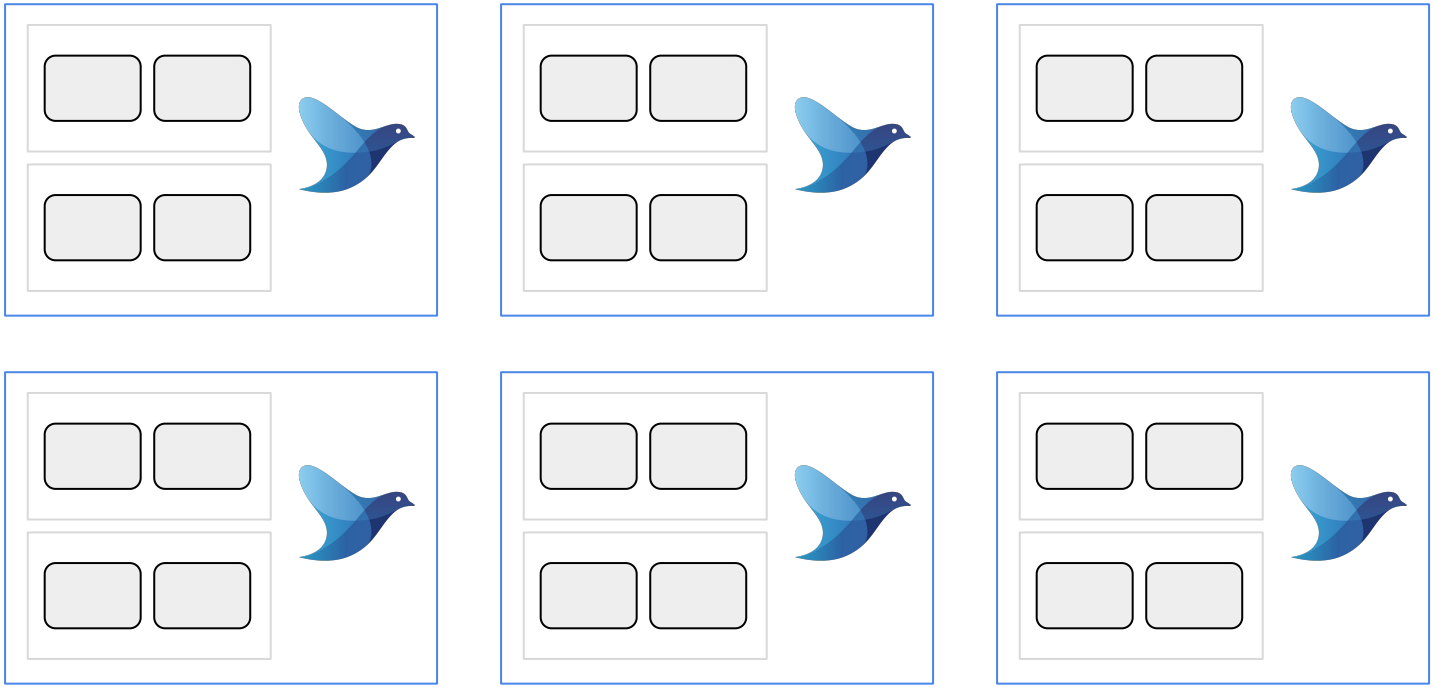

Edge Nodes & Costs

- Fluentd requires ~40MB as minimum
- Deploying a few hundred could be expensive
- Can we make **Forward cheaper** ?

Forwarder & Aggregator

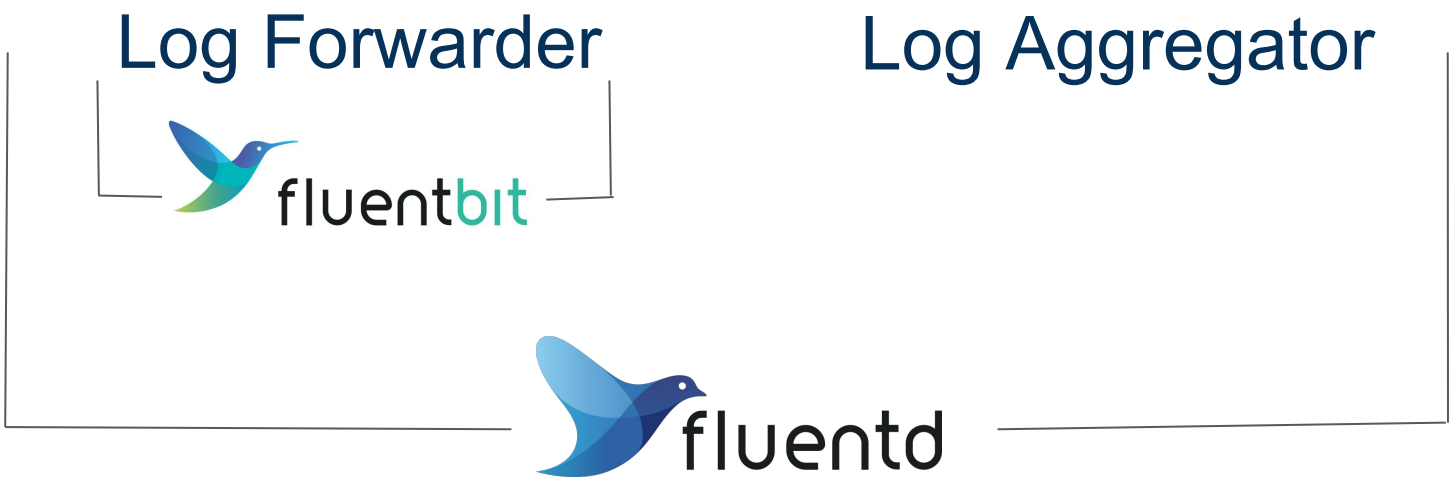

fluentbit

TREASURE DATA

About Fluentd

- Written in **C**
- Pluggable Architecture
- Built-in Reliability
- Event Driven - Async I/O

Why Fluent Bit as a Forwarder

- Features
  - Input, Filter and Output Plugins
  - Built-in parsing support
  - Minimum memory required **450KB**

Edge Nodes / Forward to Aggregators

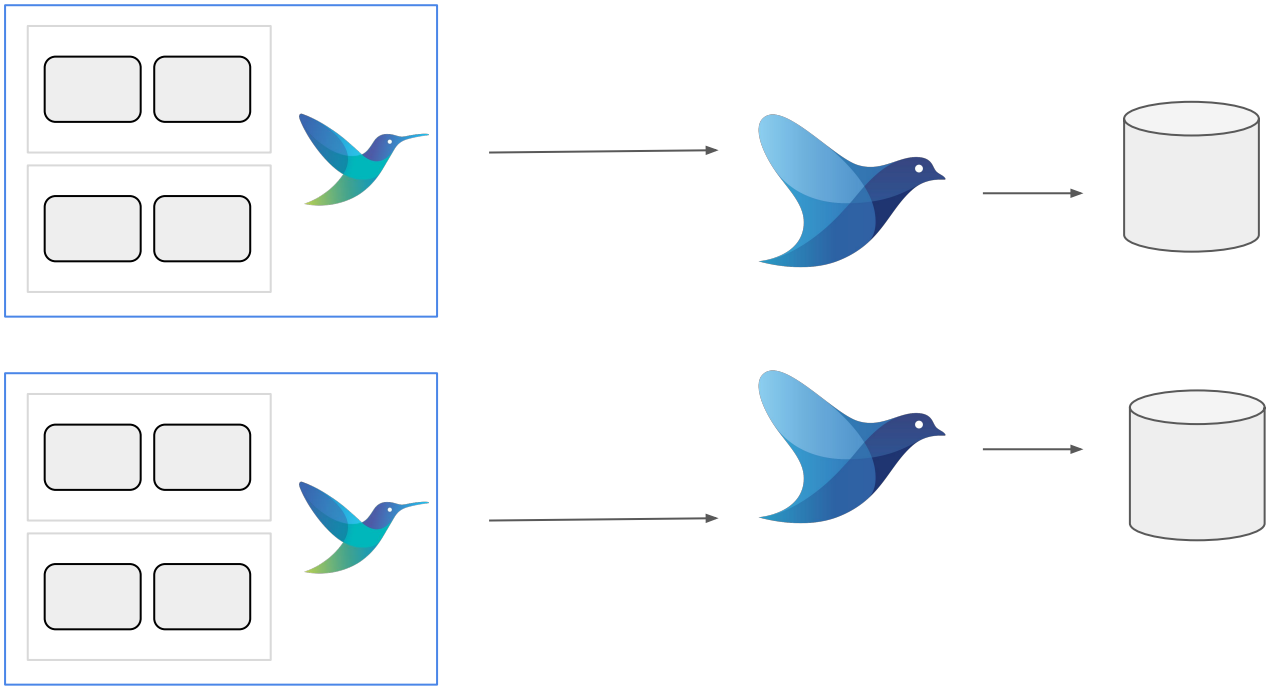

Cheap Forwarding

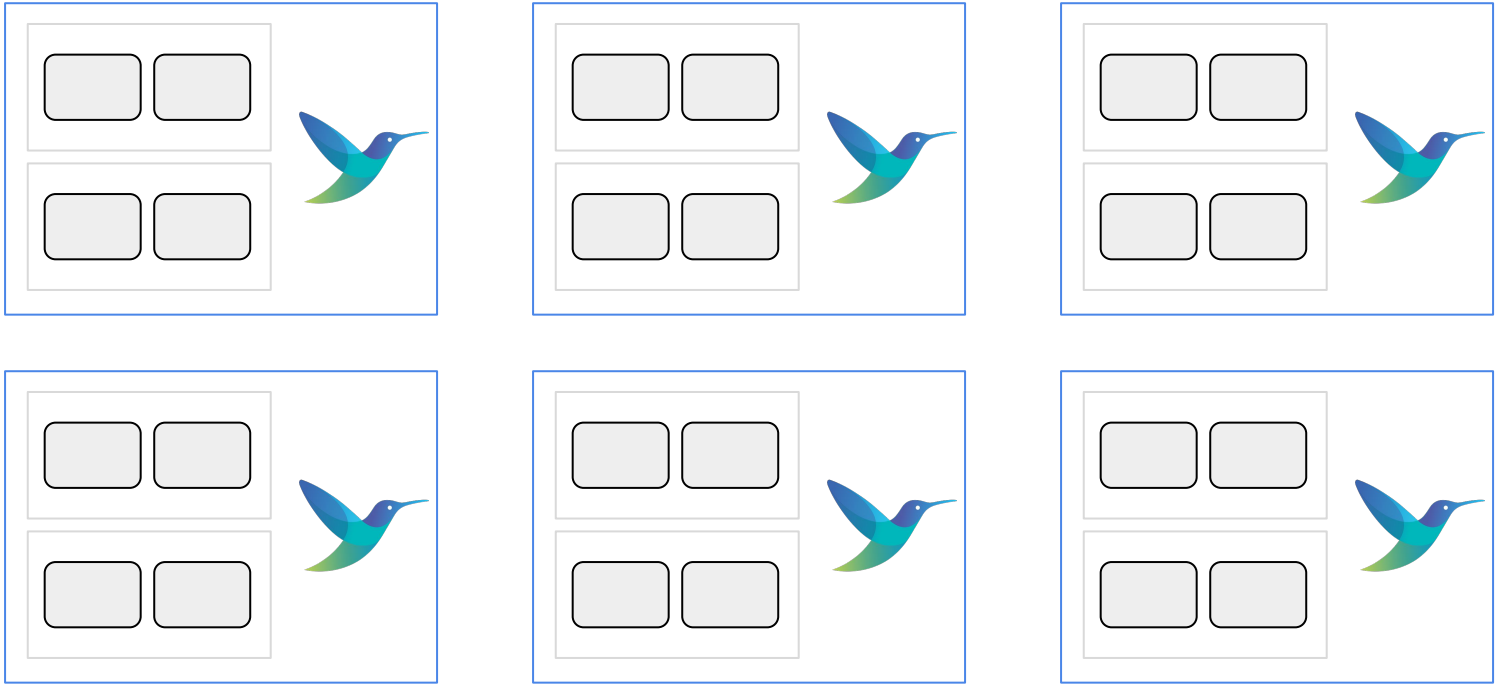

Cloud Native Features

- Docker & Kubernetes Support
- Buffering fully controled
  - pause() / resume() for input plugins
- Easy to containerize
  - Small memory footprint
  - No dependencies (all are built-in)

Hands on!

DEMO #1

Unstructured vs Structured data

Unstructured v/s Structured

- Why
  - Structured data have a schema
  - Easy to convert to different representations
  - It can be filtered

Hands on!

DEMO #2

Process Docker Logs

Kubernetes use case

- Applications runs in Containers
- Containers runs in a POD
- Multiple PODs can exists in a Node
- How to solve logging ?

Hands on!

DEMO #3

Kubernetes: parse logs and append Metadata

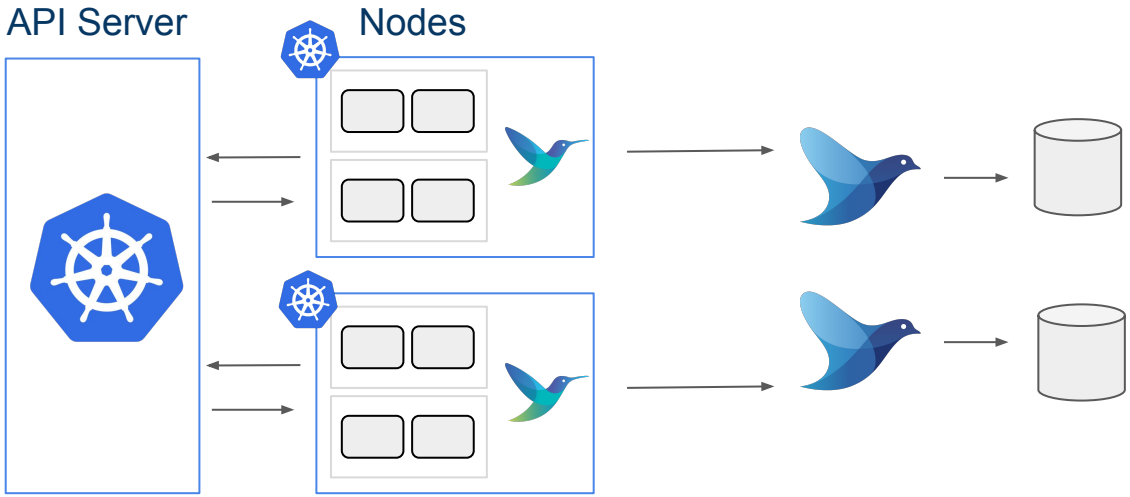

Kubernetes

Metadata Support Status

The new kubernetes filter takes care of the following metadata handling:

- Local data: POD Name, Namespace, Container Name and Container ID.
- Remote (API Server): Labels and Annotations

Fluent Bit, what else ?

Networking and Co-routines

Easier implementation of output plugins that interact with networking operations like socket(), connect(), read(), write(), etc.

Fluent Bit provides non-blocking networking API that uses the event-loop with co-routines to implement:

- Network I/O
- TLS/SSL usage
- HTTP Client

Kubernetes DaemonSet

Github Repository

- https://github.com/fluent/fluent-bit-kubernetes-daemonset

Docker Image (ubuntu-slim)

- quay.io/fluent/fluent-bit-kubernetes-daemonset

Roadmap

Next Release v0.11 (March 2017)

- Kubernetes support (filter_kubernetes)
- Parsers & Filters
- Memory optimizations

Release v0.12 (May 2017)

- in_tail + Multiline support
- Monitoring - re-enable HTTP service end-point: memory, records flow, others.

Thanks!

Project information

- Web site Fluentbit.io
- Documentation http://fluentbit.io/documentation/
- Github http://github.com/fluent/fluent-bit

Contact

- Slack http://slack.fluentd.org (fluent-bit channel)
- Twitter @fluentbit

TREASURE DATA