

Fluentd

Open Source Data Collector



Jan 23, 2016

Scale14x, Pasadena!



Eduardo Silva

eduardo@treasuredata.com

[@edsiper](https://twitter.com/edsiper)



spread the word!

#scale14x #fluentd

@edsiper

About Me



Eduardo Silva

- Github & Twitter
- Personal Blog

@edsiper

<http://edsiper.linuxchile.cl>

Treasure Data

- Open Source Engineer
- Fluentd / Fluent Bit

<http://github.com/fluent>

Projects

- Monkey HTTP Server
- Duda I/O

<http://monkey-project.com>


<http://duda.io>

Logging

Logging Matters

Pros

- Application status
- Debugging
- General information about anomalies: errors
- Troubleshooting / Support
- Local or Remote (network)



Logging Matters

From a business point of view

- Input data → Analytics
- User interaction / behaviors
- Improvements



Assumptions

Logging Matters

Assumptions



- I have enough disk **space**
- I/O operations will **not** block
- Log messages are **human readable**
- My logging mechanism **scale**

Logging Matters

Assumptions

| *Basically, yeah.. it **should** work.*

Concerns

Logging Matters

Concerns

- Logs **increase** = data **increase**
- Message format get more **complex**
- Did the Kernel flush the buffers ? (sync(2))
- Multi-thread application ?, **locking** ?
- **Multiple Applications** = **Multiple Logs**

Logging Matters

Concerns

*If **Multiple Applications** = **Multiple logs***

Multiple Hosts** **x** **Multiple Applications** = **???

OK, so:

1. Logging matters
2. It's really beneficial
3. but...

 **It needs to
be done right.**

Logging

Common sources & inputs

- **Application** Logs
 - Apache
 - NginX
 - Syslog (-ng)
- **Custom** applications / Languages
 - C, Ruby, Python, PHP, Perl, NodeJS, Java, etc.



In a galaxy
not so far away...

Access logs

Apache



Metrics


Blueflood

App logs

Frontend



Backend



Analysis


MongoDB

MySQL

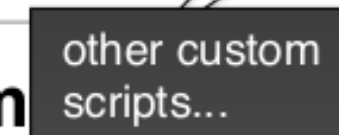
Hadoop

System logs

syslogd



Your system



Archiving

Amazon S3

How to parse/store multiple data sources ?

| note: performance matters!



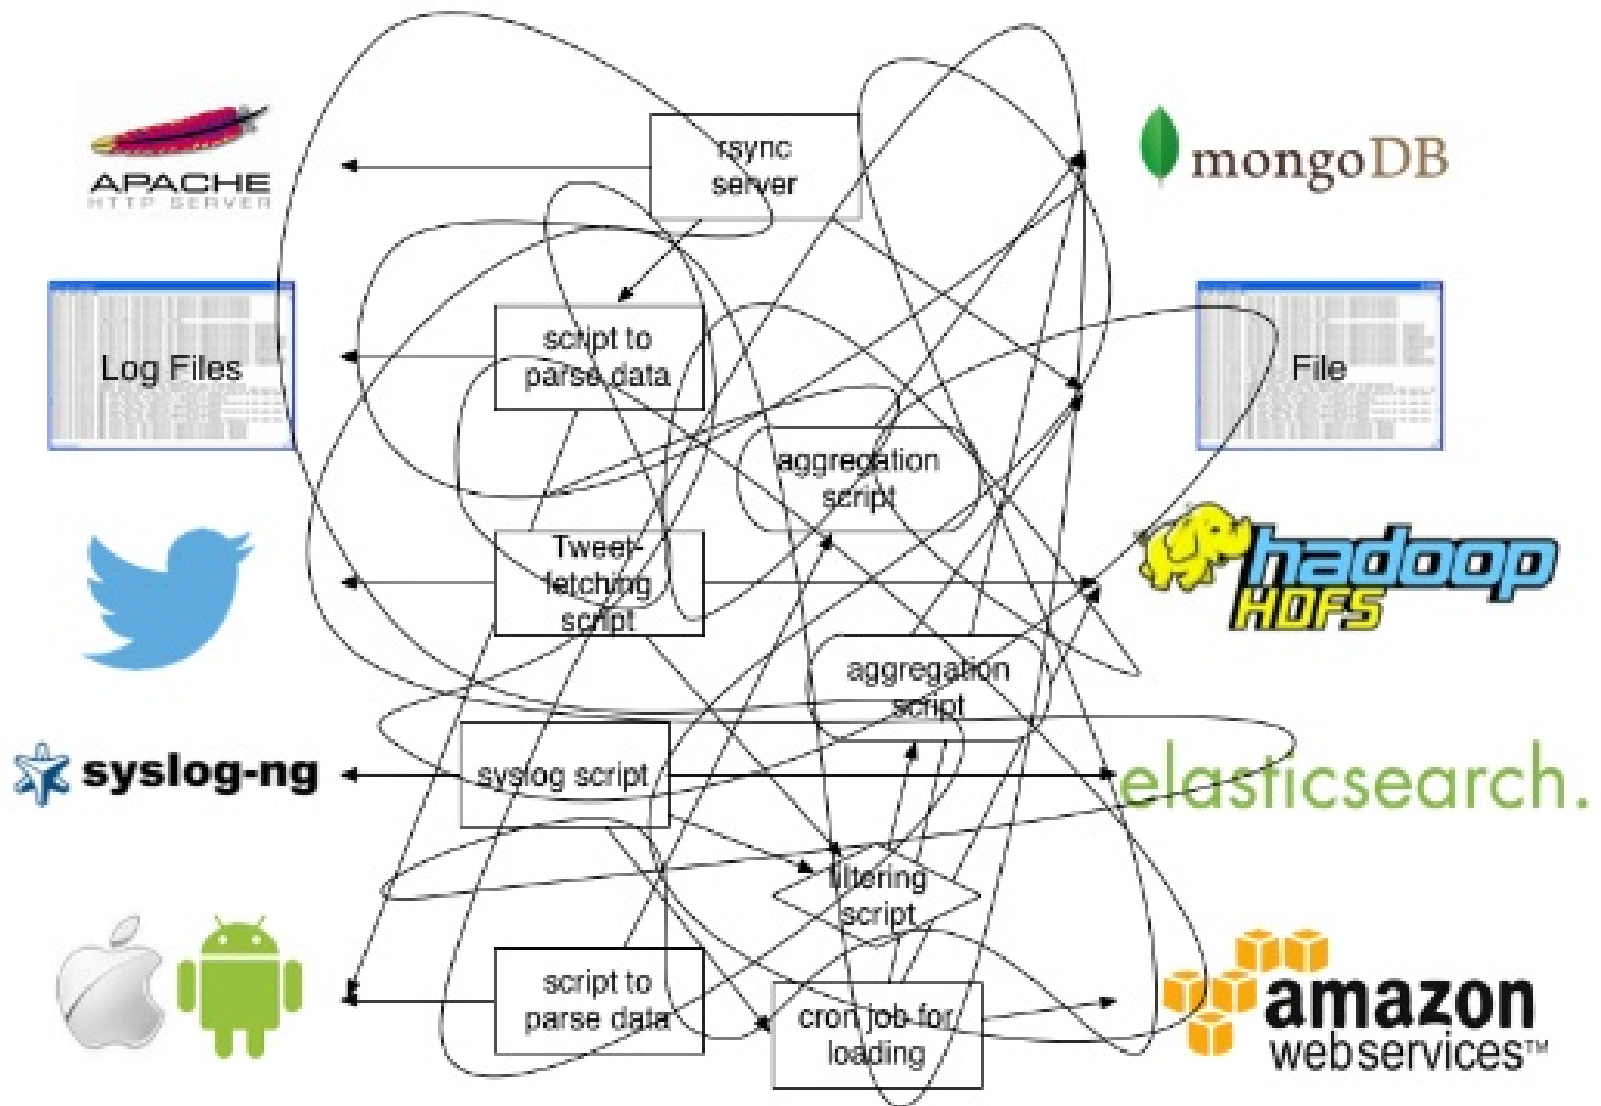
fluentd



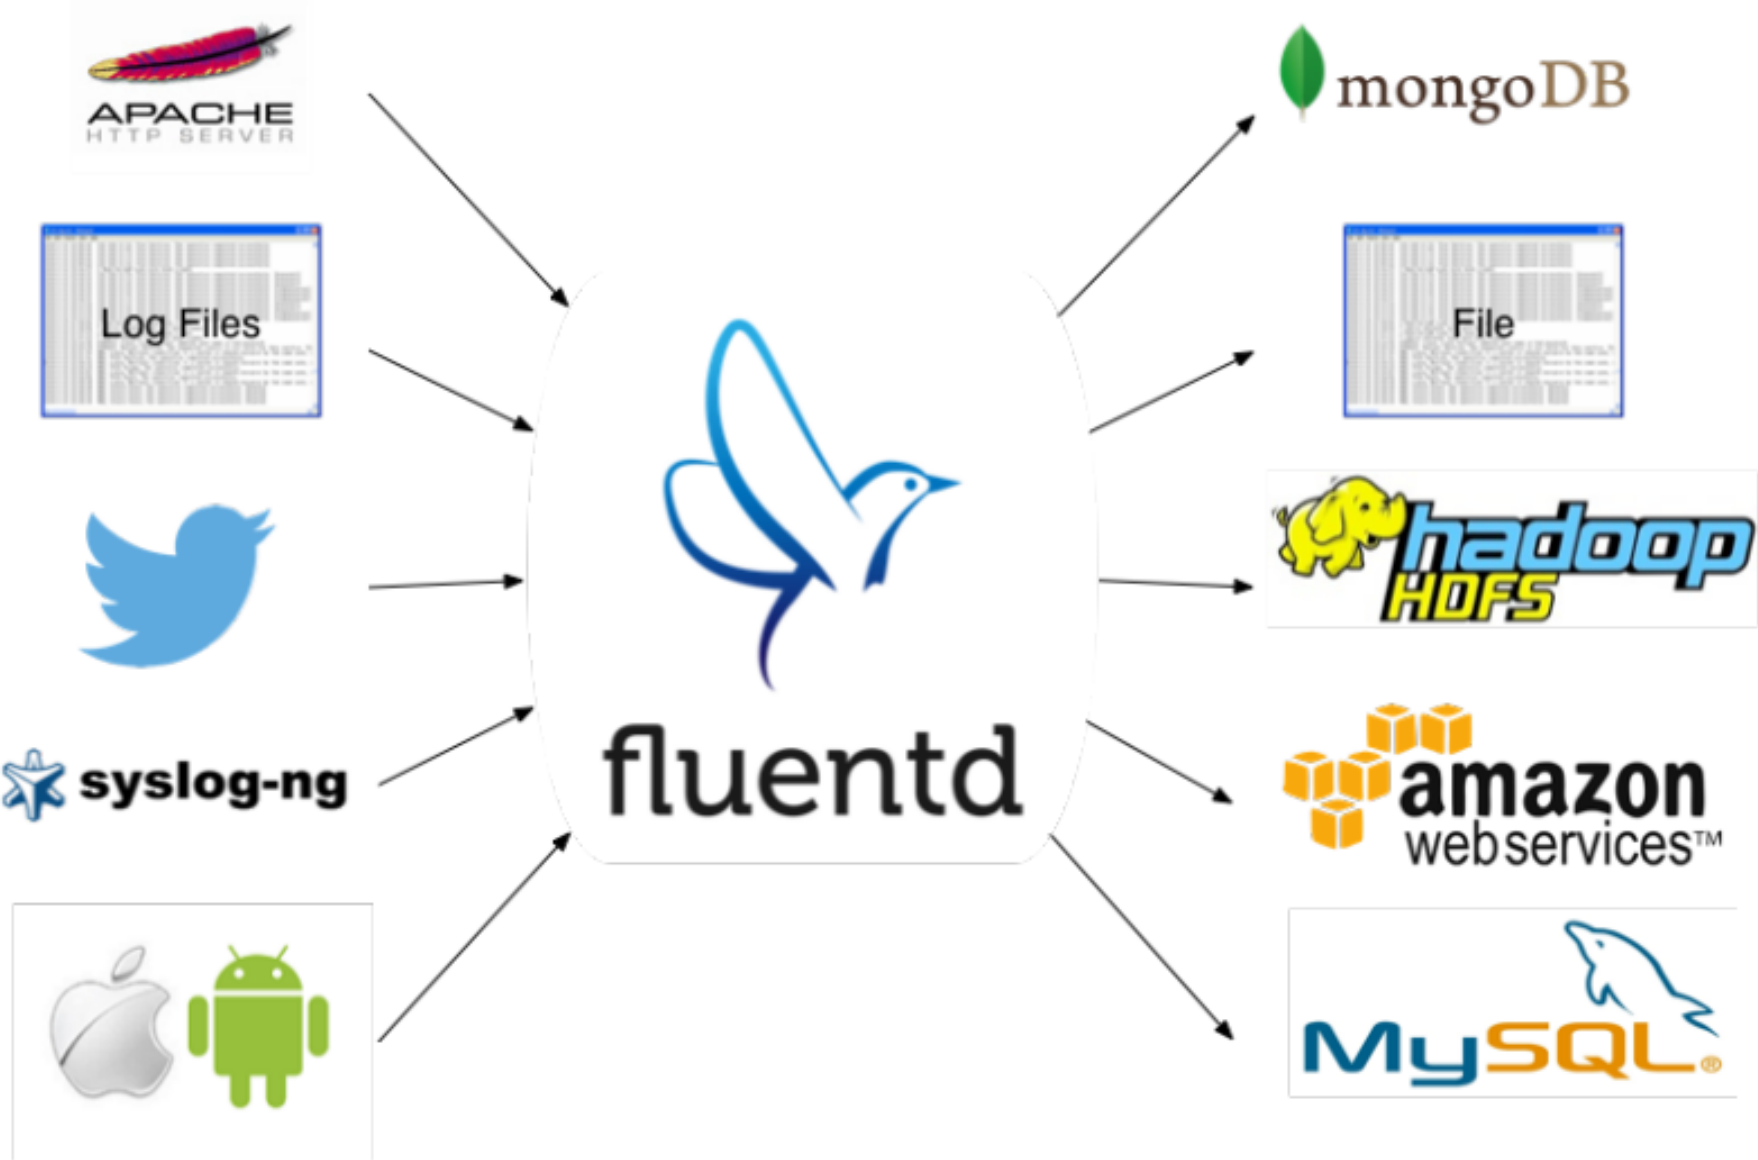
Fluentd is an open source data collector

It let's you unify the data collection for a better use and understanding of data.

before



after



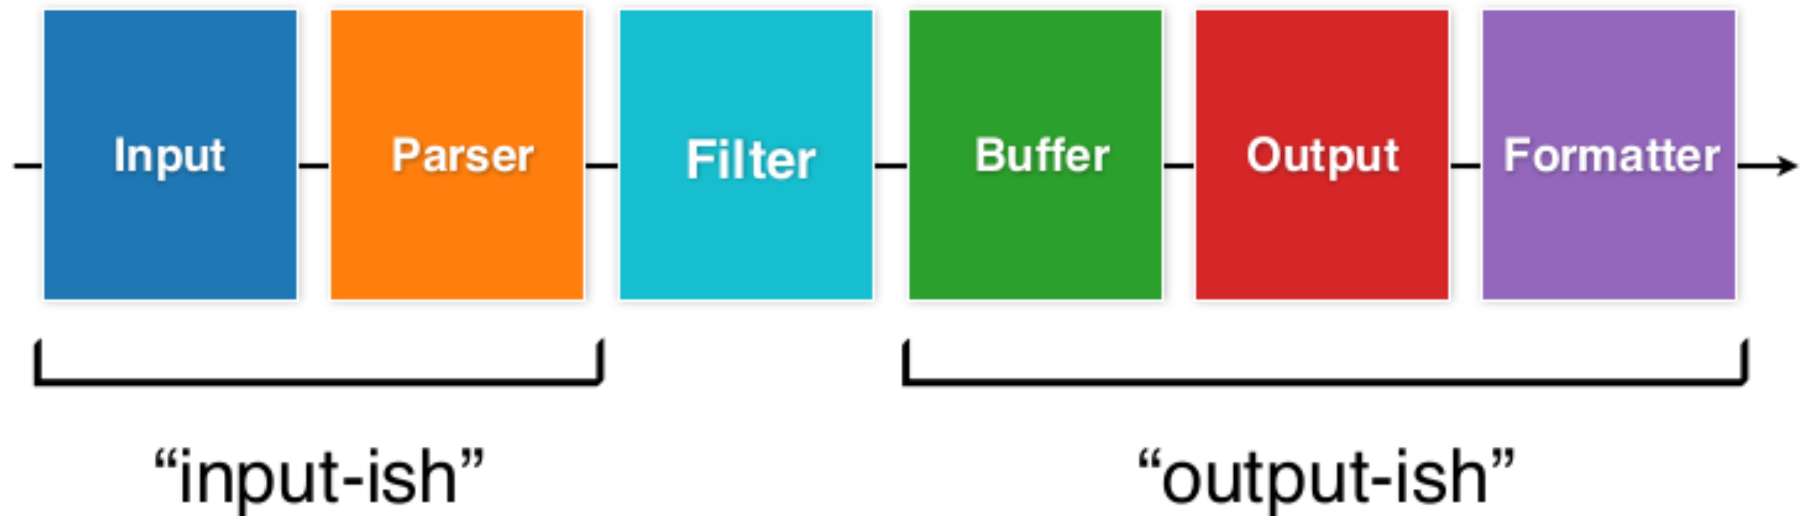
Fluentd

Highlights

- High **Performance**
- Built-in **Reliability**
- **Structured** Logs
- **Pluggable** Architecture
- More than 300 plugins! (input/filtering/output)

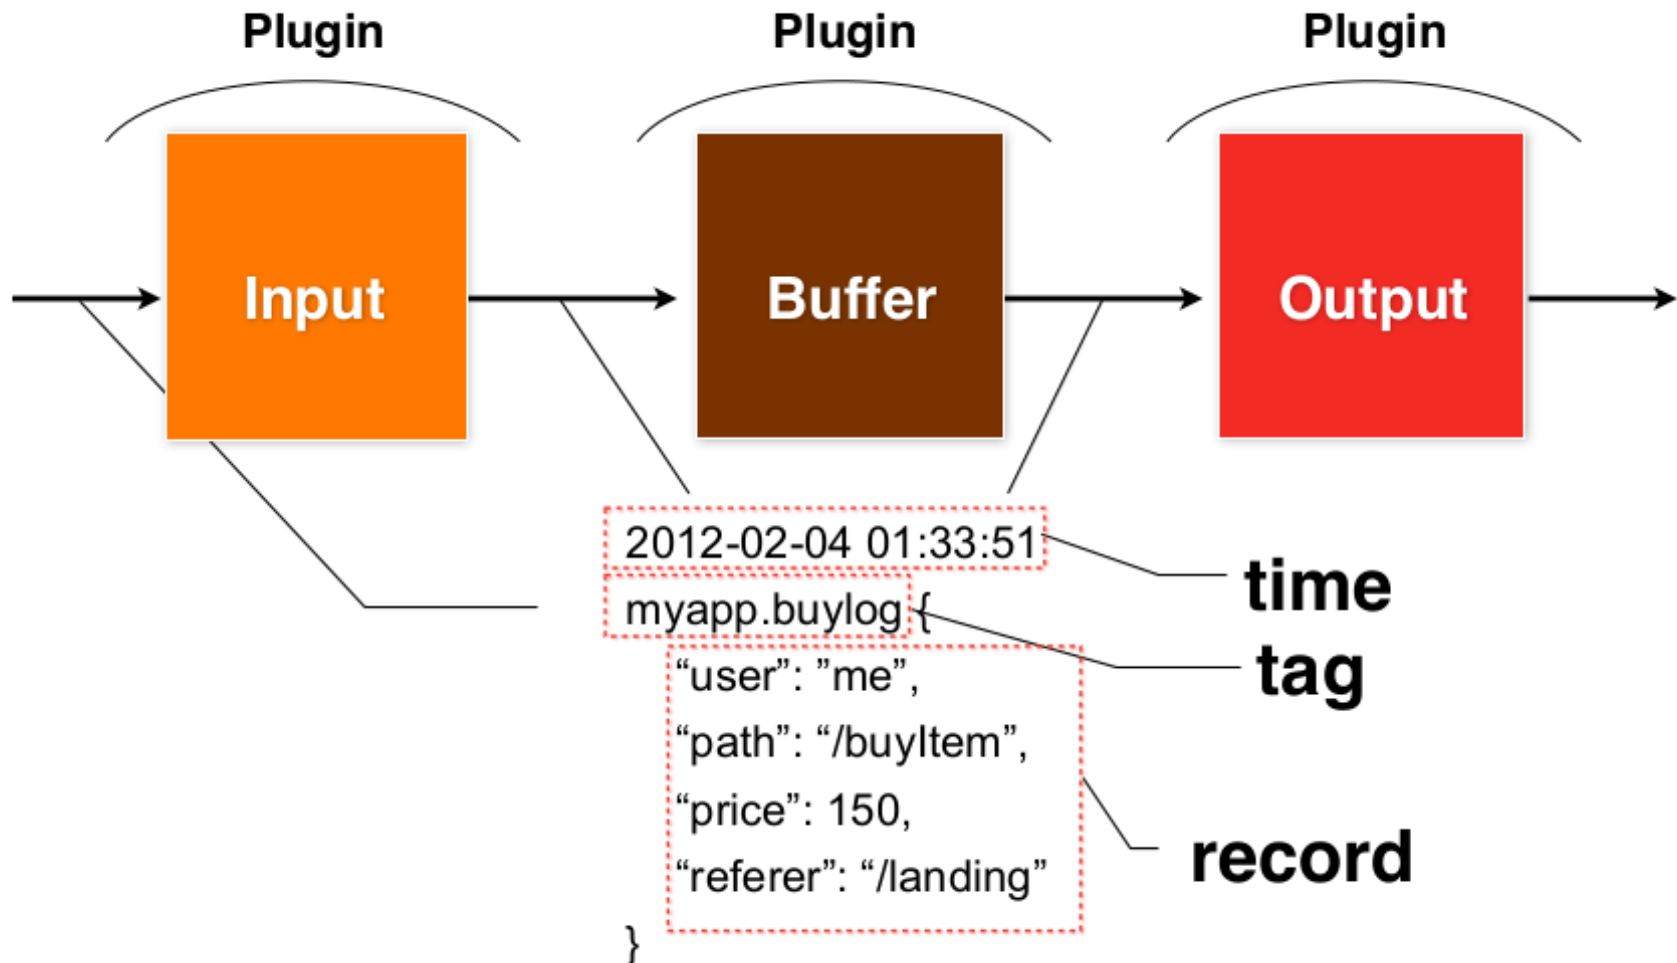
Fluentd

Architecture



Fluentd

Internals simplified



Fluentd

Input plugins



✓ Receive logs

✓ Or pull logs from data


sources

✓ in non-blocking manner

Fluentd

Output plugins

Plugin



✓ Write or send event logs

File (out_file)

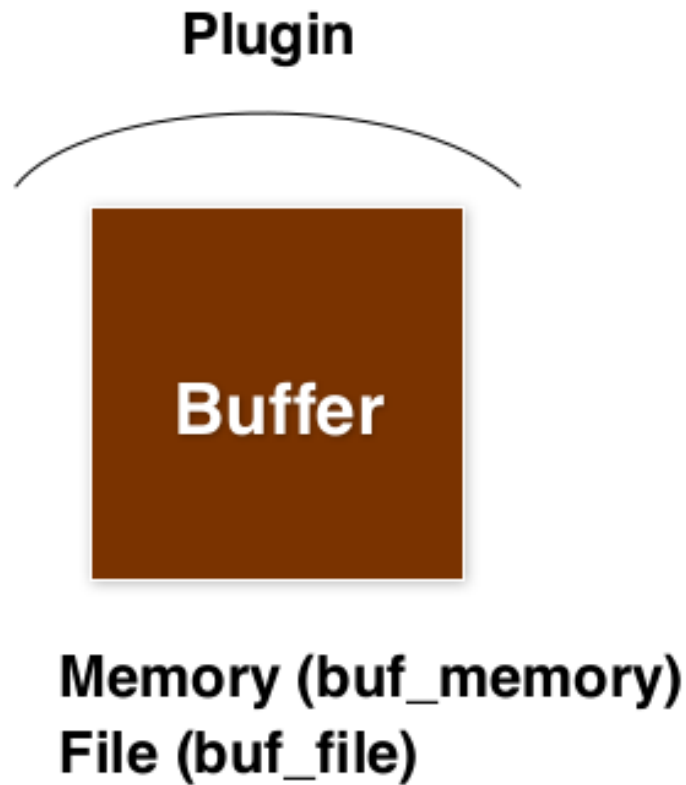
Amazon S3 (out_s3)

MongoDB (out_mongo)

...

Fluentd

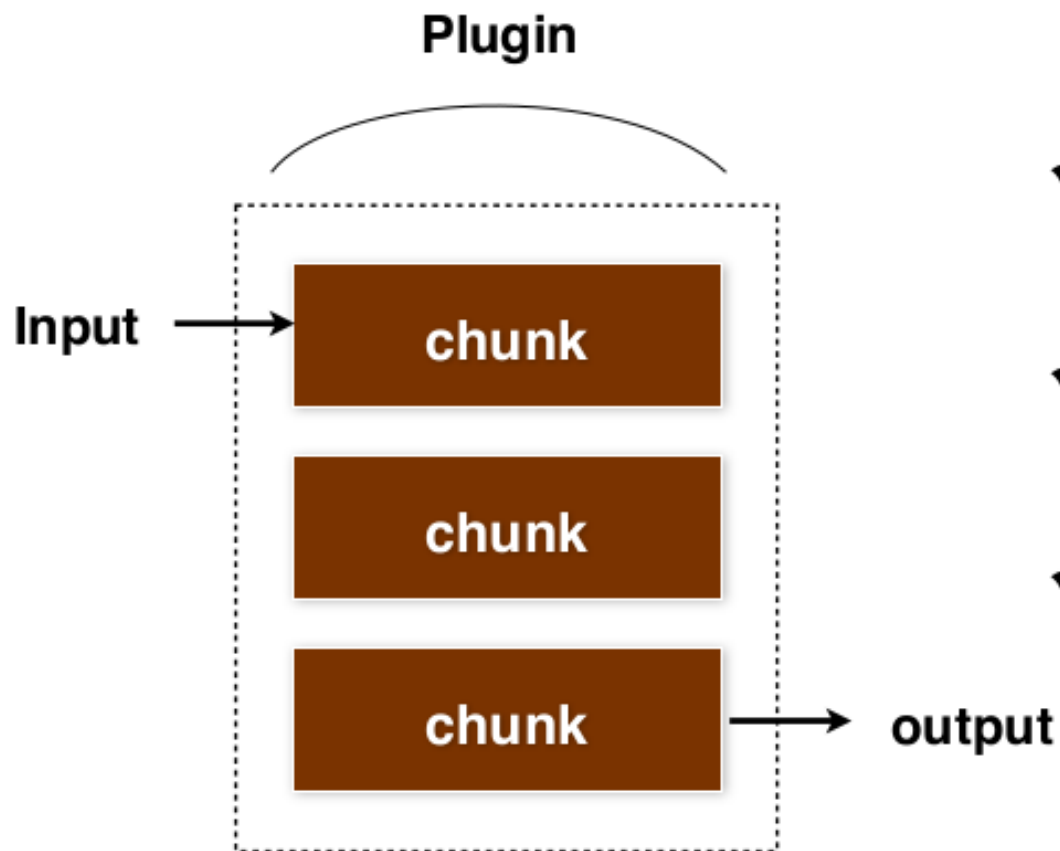
Buffer plugins



- ✓ Improve performance
- ✓ Provide reliability
- ✓ Provide thread-safety

Fluentd

Buffer plugins



✓ Improve performance

✓ Provide reliability

✓ Provide thread-safety

Input / Filter

Output



emit

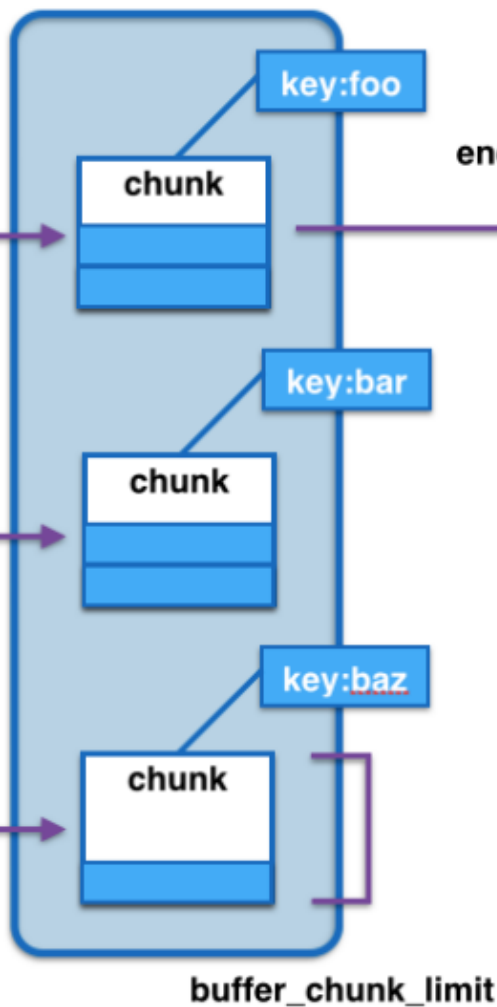
Router

emit

Key pattern:

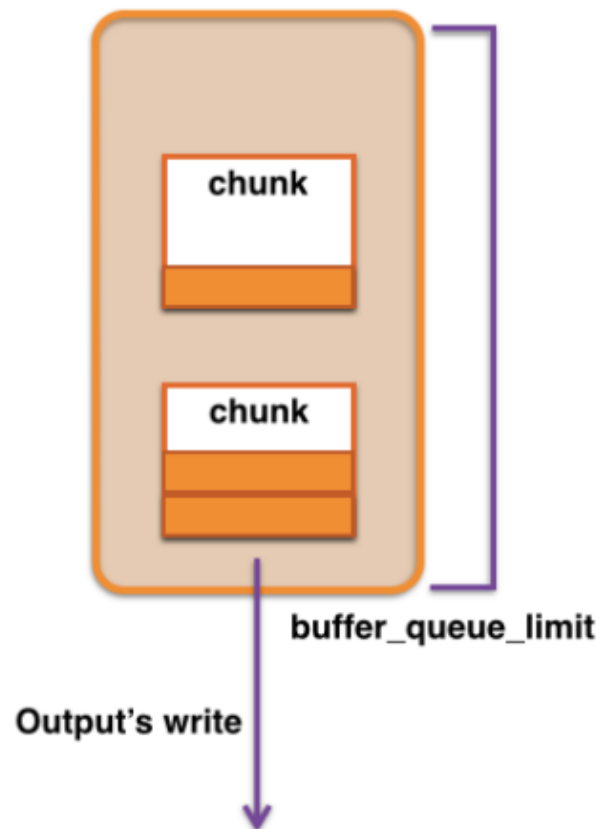
- BufferedOutput
empty string or specified key
- ObjectBufferedOutput
tag
- TimeSlicedOutput
time slice

Buffer



enqueue: exceed flush_interval
or buffer_chunk_limit

Queue



$$M \times N \rightarrow M + N$$

Access logs

Apache

App logs

Frontend

Backend

System logs

syslogd

Databases

Alerting

Nagios

Analysis

MongoDB

MySQL

Hadoop

Archiving

Amazon S3



fluentd

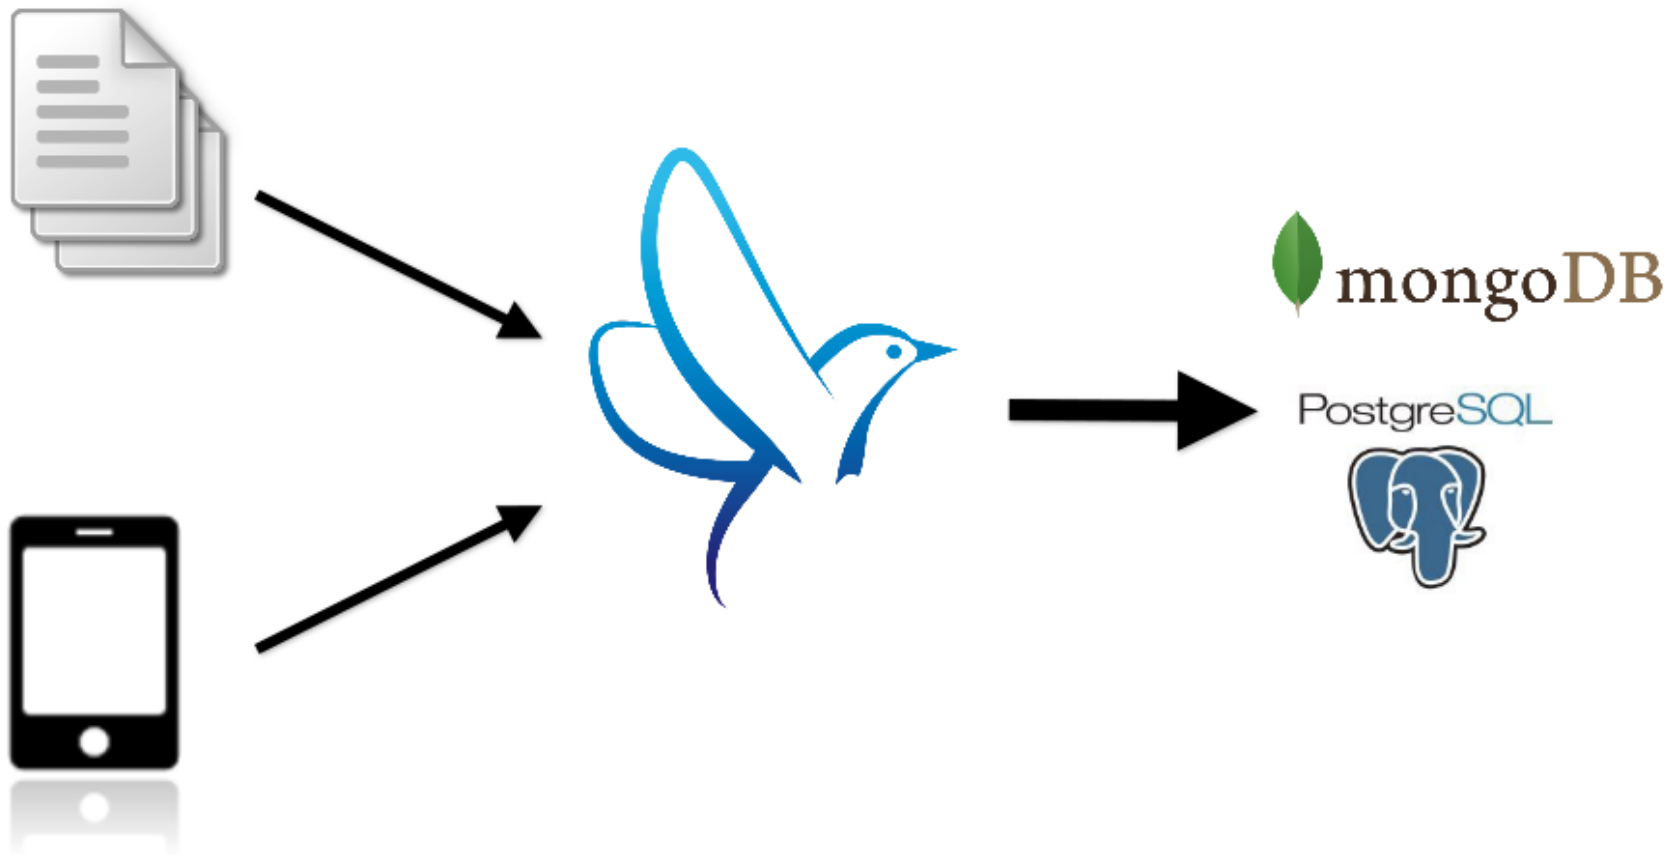


buffer/filter/route



Fluentd

Simple Forwarding



Fluentd

Simple Forwarding: configuration

logs from a file

<source>

```
type    tail
path    /var/log/httpd.log
format  apache2
tag     backend.apache
```

</source>

logs from client libraries

<source>

```
type    forward
port    24224
```

</source>

store logs to MongoDB


<match backend.>*

```
type      mongo
database  fluent
collection test
```

</match>


Fluentd

Less Simple Forwarding



Fluentd

Lambda Architecture



Fluentd

logs from a file

```
<source>
  type    tail
  path    /var/log/httpd.log
  format  apache2
  tag     backend.apache
</source>
```

logs from client libraries

```
<source>
  type    forward
  port    24224
</source>
```

store logs to MongoDB

```
<match *.*>
  type    copy
  <store>
    type    elasticsearch
    logstash_format true
  </store>
  <store>
    type    webhdfs
    host    192.x.y.z
    port    50070
    path    /path/to/hdfs
  </store>
</match>
```

Who uses Fluentd
in production ?

LINE




BACKPLANE





TREASURE
DATA

We collect
1M events **per** second !



Internet of Things

Internet of Things

Facts

- IoT will grow to many **billions** of devices over the next decade.
- Now it's about **device** to **device** connectivity.
- Different **frameworks** and **protocols** are emerging.
- It needs **Logging**.

Internet of Things

Alliances

Vendors formed alliances to join forces and develop generic software layers for their products:



OPEN
INTERCONNECT
CONSORTIUMSM



SIEMENS



**ALLSEEN
ALLIANCE**

Canon



Panasonic

SONY

SHARP.



Internet of Things

Solutions provided

Alliance



Framework



IoT & Big Data

Analytics

IoT requires a **generic solution** to collect events and data from different sources for further analysis.

Data can come from a specific framework, radio device, sensor or other. How do we collect and **unify** data **properly** ?



@fluentbit




Fluent Bit is an open source data collector

It let's you collect data from IoT/Embedded devices and transport It to third party services.

Fluent Bit

Targets

- Services
- Sensors / Signals / Radios
- Operating System information
- Automotive / Telematics



Fluent Bit

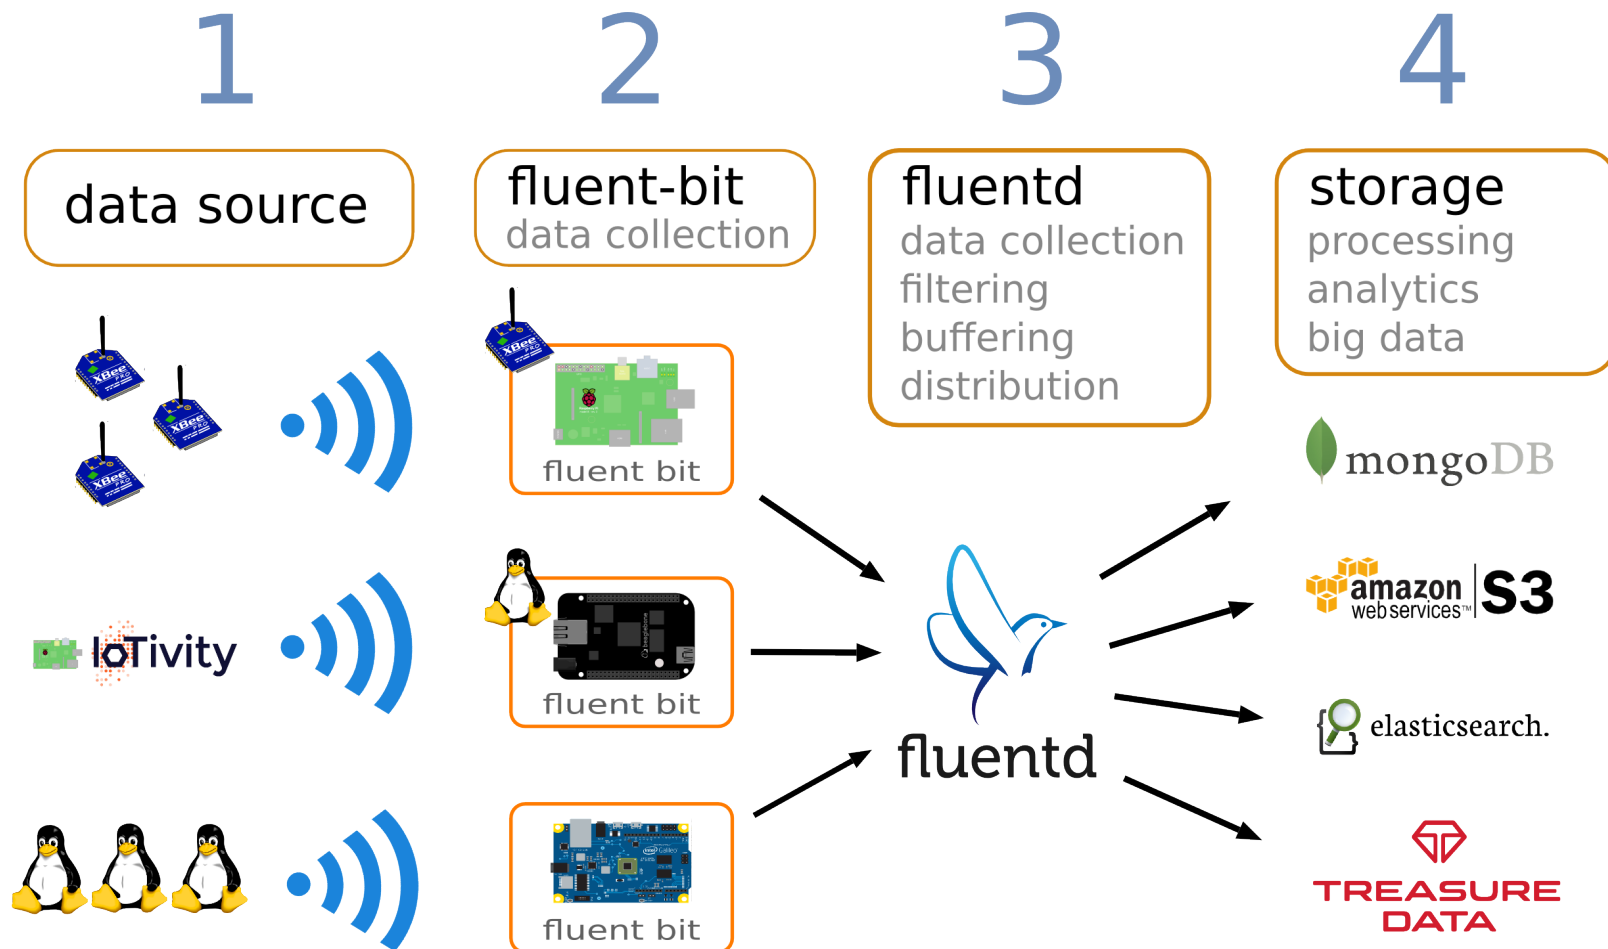
Requirements

IoT and Embedded environment requires special handling, specifically on performance and resource utilization:

- Lightweight
- Written in **C** Language
- Customizable, **pluggable** architecture
- Full integration with **Fluentd**

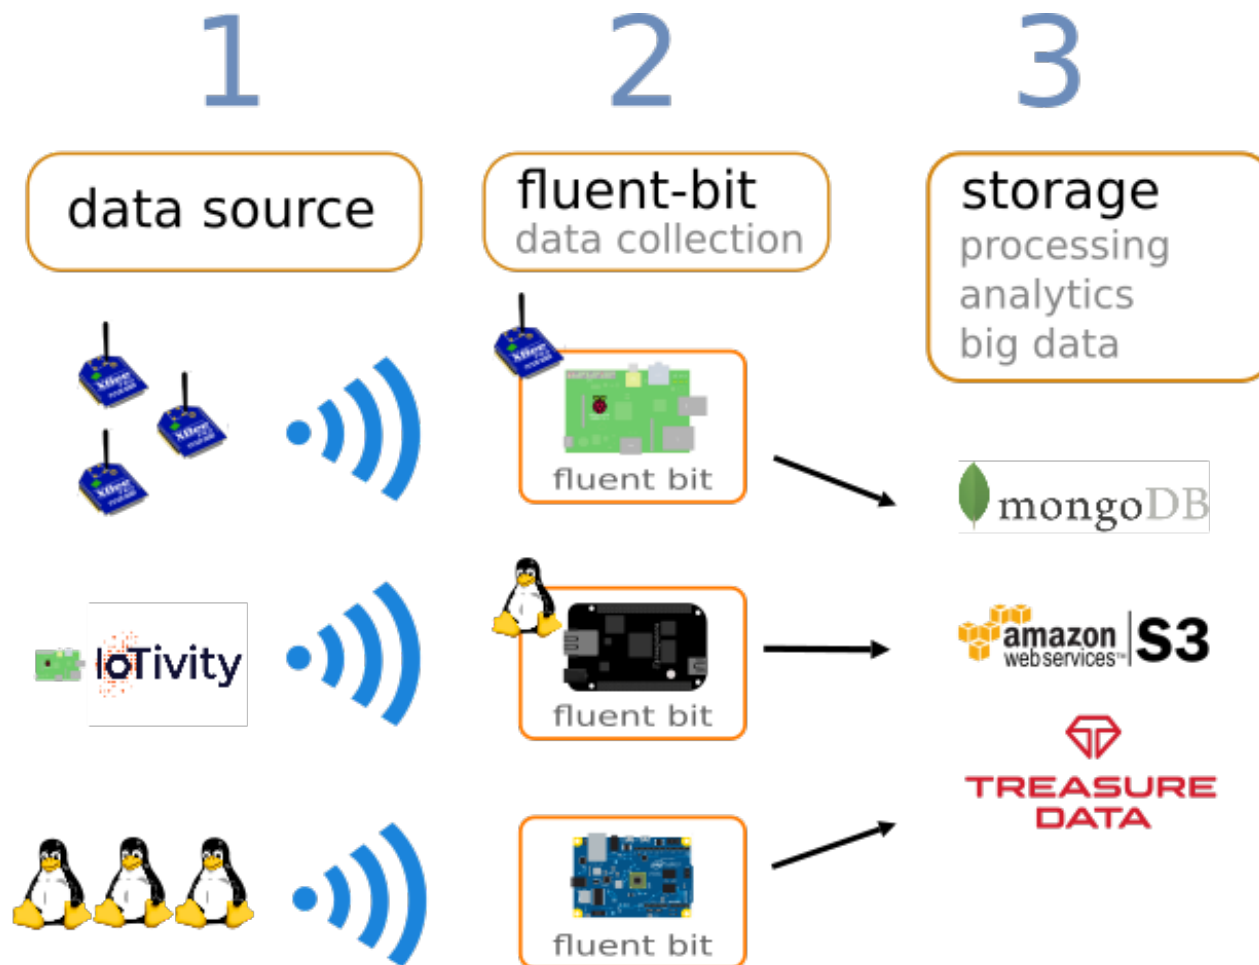
Fluent Bit

Integration




Fluent Bit

Direct Output



Fluent Bit

Elastic Search support



Configure an index pattern

In order to use Kibana you must configure at least one index pattern. Index patterns are used to identify the Elasticsearch index to run search and analytics against. They are also used to configure fields.

Index contains time-based events

Use event times to create index names [DEPRECATED]

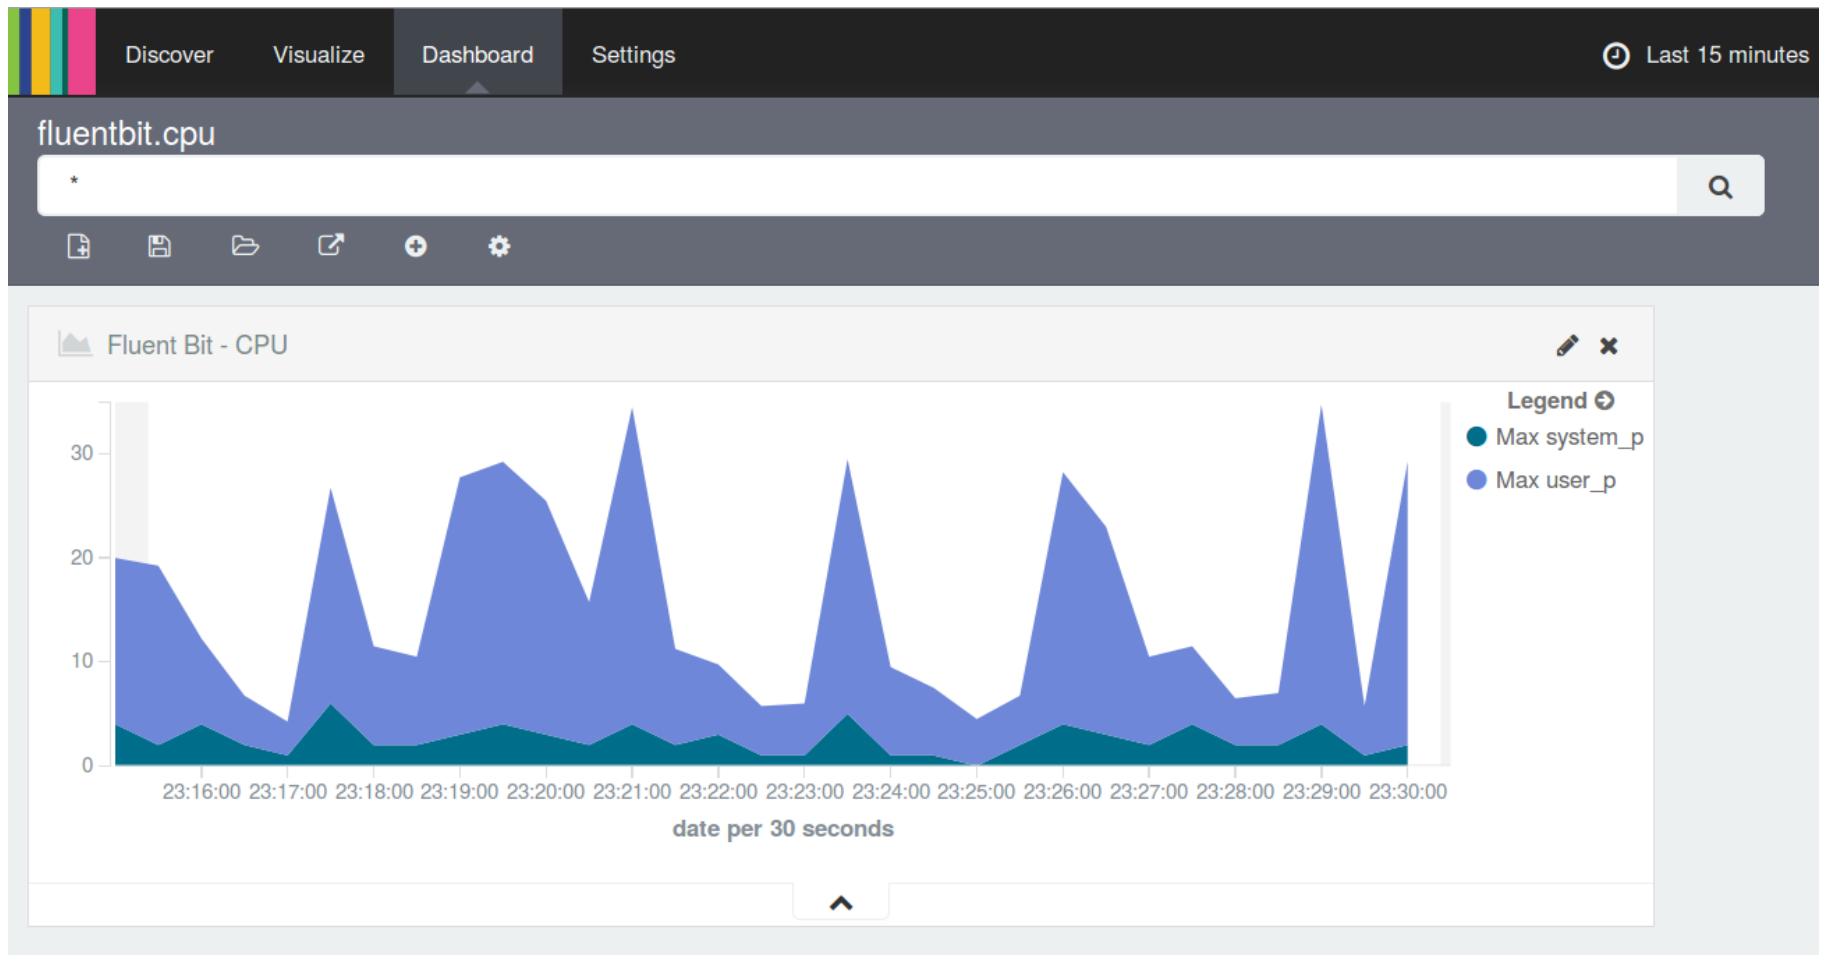
Index name or pattern

Patterns allow you to define dynamic index names using * as a wildcard. Example: logstash-*

Time-field name ⓘ refresh fields

Fluent Bit

Elastic Search: Dashboard



Containers





fluentd



docker

Docker


Logging driver


- Docker v1.6 released the concept of logging drivers
- Route container output
- Fluentd ?

Docker


 docker / docker


 Watch ▾ 2,081

 Star 23,470


 Fork 5,853

Add new Logging driver "fluentd" #12876

 **Merged** LK4D4 merged 2 commits into `docker:master` from `tagomoris:logger-driver-fluentd` 26 days ago

 Conversation 167

 Commits 2

 Files changed 33

+7,763 -3 



tagomoris commented on Apr 29

This patch provides `fluentd` logging plugin for docker, which is mentioned at #12540.

How to confirm behavior of this patch/new logging driver.

1. Setup dev docker container
2. Build repository with this patch by `make BINDDIR=. binary`
3. Copy binary and execute it in debug mode `docker -dD`
4. Attach from external environment by `docker exec -it kickass_heisenberg /bin/bash`
5. Build/execute another docker container from other console, built on Dockerfile and configuration file below:

```
# sh
```

Labels

status/4-merge


Milestone

No milestone

Assignee

No one assigned

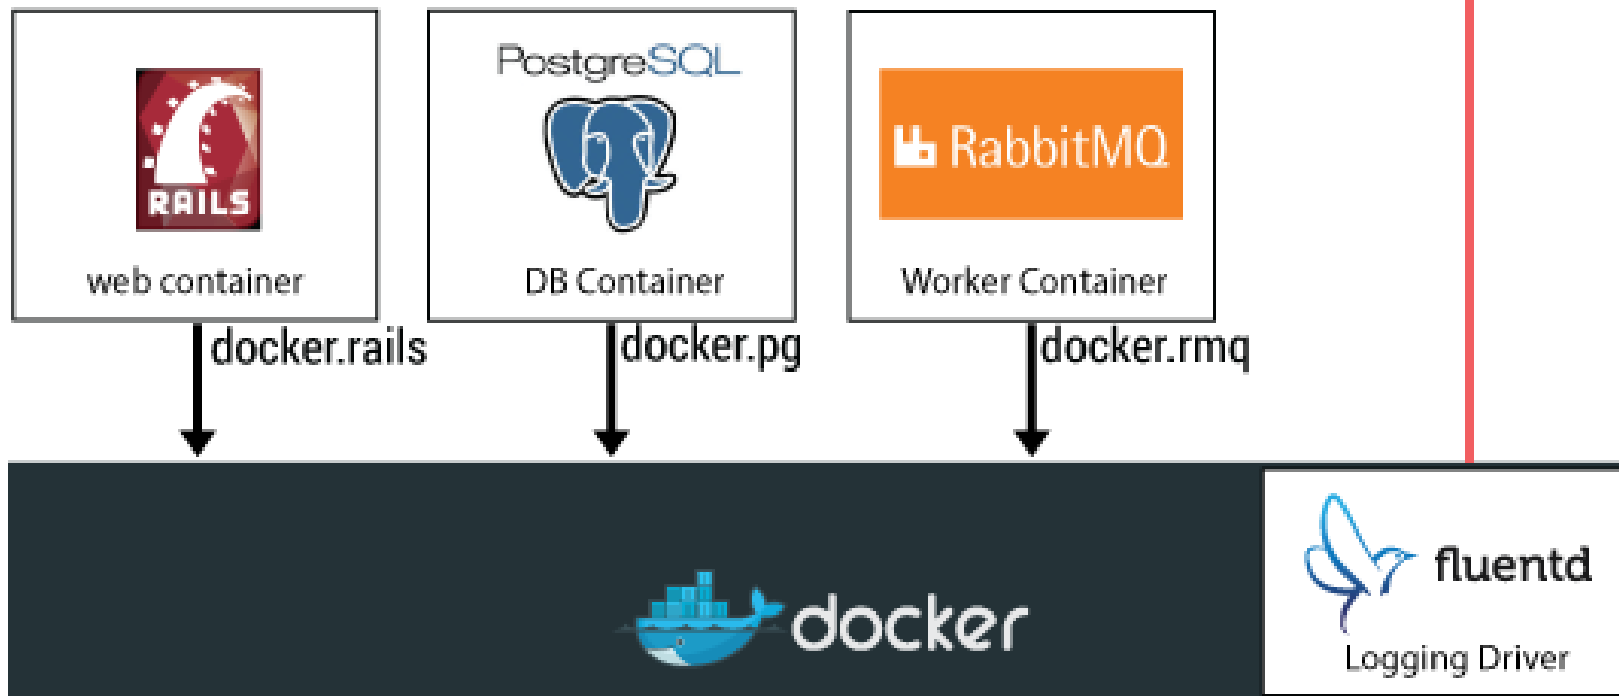
Notifications

 Subscribe

Docker v1.8

Fluentd Logging driver!

```
docker run --log-driver=fluentd \  
--log-opt fluentd-tag=docker.{{.Name}}
```



Docker

Data Stream

tag	time	source	container_id	container_name	log message
docker.3fd...	1441402468	stdout	3fd8678d487e...	/angry_kamam	Hello world!



docker.3fd...	1441402468	stdout
3fd8678d4...	/angry_kam	message




docker.3fd...	1441402469	stdout
3fd8678d4...	/angry_kam	message

docker.3fd...	1441402470	stdout
3fd8678d4...	/angry_kam	message

docker.3fd...	1441402471	stdout
3fd8678d4...	/angry_kam	message

Docker

Data Stream





NodeJS


Fluent-Logger (NPM)

```
var logger = require('fluent-logger')

logger.configure('tag', {
  | host: 'localhost',
  | port: 24224,
  | timeout: 3.0
});

logger.emit('label', {record: 'this is a log'});
```

We Love Data!



TREASURE
DATA

- <http://fluentd.org>
- <http://fluentbit.io>
- <https://docs.docker.com/reference/logging/fluentd/>
- <http://github.com/fluent/fluentd>



Thank you!