Analyzing logs using-Elasticsearch, Fluentd, and Kibana(EFK)

- Ravdeep S Pasricha
ravdeep003

Agenda

- Why EFK
- Elasticsearch, Fluentd, Kibana Overview
- Demo
- High Availability

Why EFK

Before EFK

- Log files across the server
- Grep/shell scripting mess
- Production system access

EFK:

- Detecting log patterns
- Easily Searchable
- Trouble-shooting issues

Elasticsearch

- Elasticsearch is a search server based on Lucene.

- It provides a distributed, multitenant-capable full-text search engine with a RESTful web interface and schema-free JSON documents.

Kibana

- Kibana is an open source (Apache Licensed), browser based analytics and search dashboard for Elasticsearch.
- Explore and visualize your data

Fluentd

- Fluentd is an open source data collector for unified logging layer.
- Fluentd allows you to unify data collection and consumption for a better use and understanding of data.

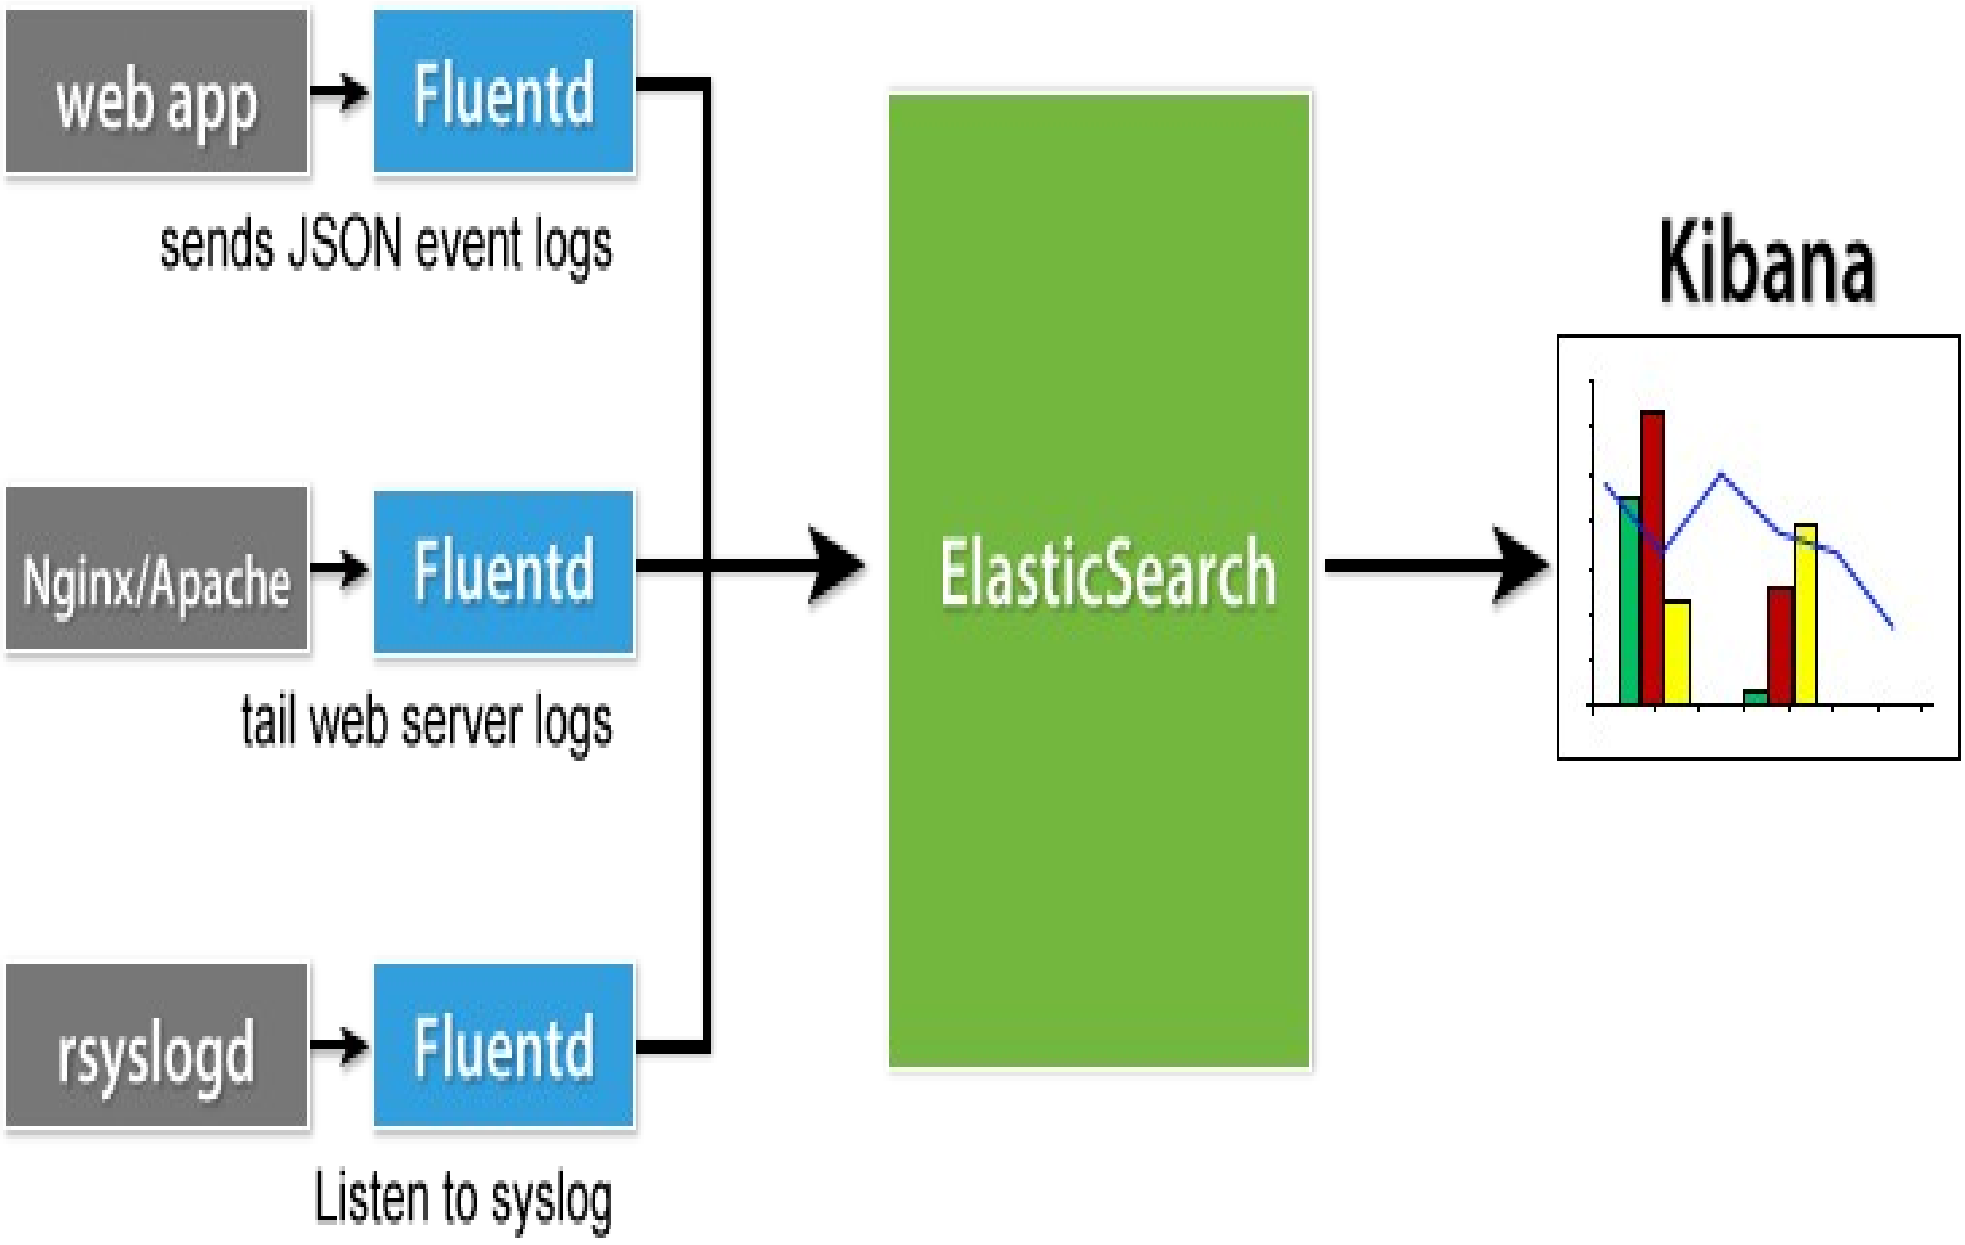

web app

Fluentd

sends JSON event logs

Nginx/Apache

Fluentd

tail web server logs

rsyslogd

Fluentd

Listen to syslog

ElasticSearch

Kibana

DEMO

Fluentd

- Fluentd has 5 types of plugins: Input, Parser, Output, Formatter and Buffer

  Input Plugins:
  - in_forward
  - in_http
  - in_tail
  - in_exec
  - in_syslog

Output Plugins:

- out_forward
- out_mongo
- out_file
- out_s3

Buffer Plugins:

Buffer plugins are used by buffered output plugins, such as out_file, out_forward, etc

- buf_memory
- buf_file

Fluentd Setup

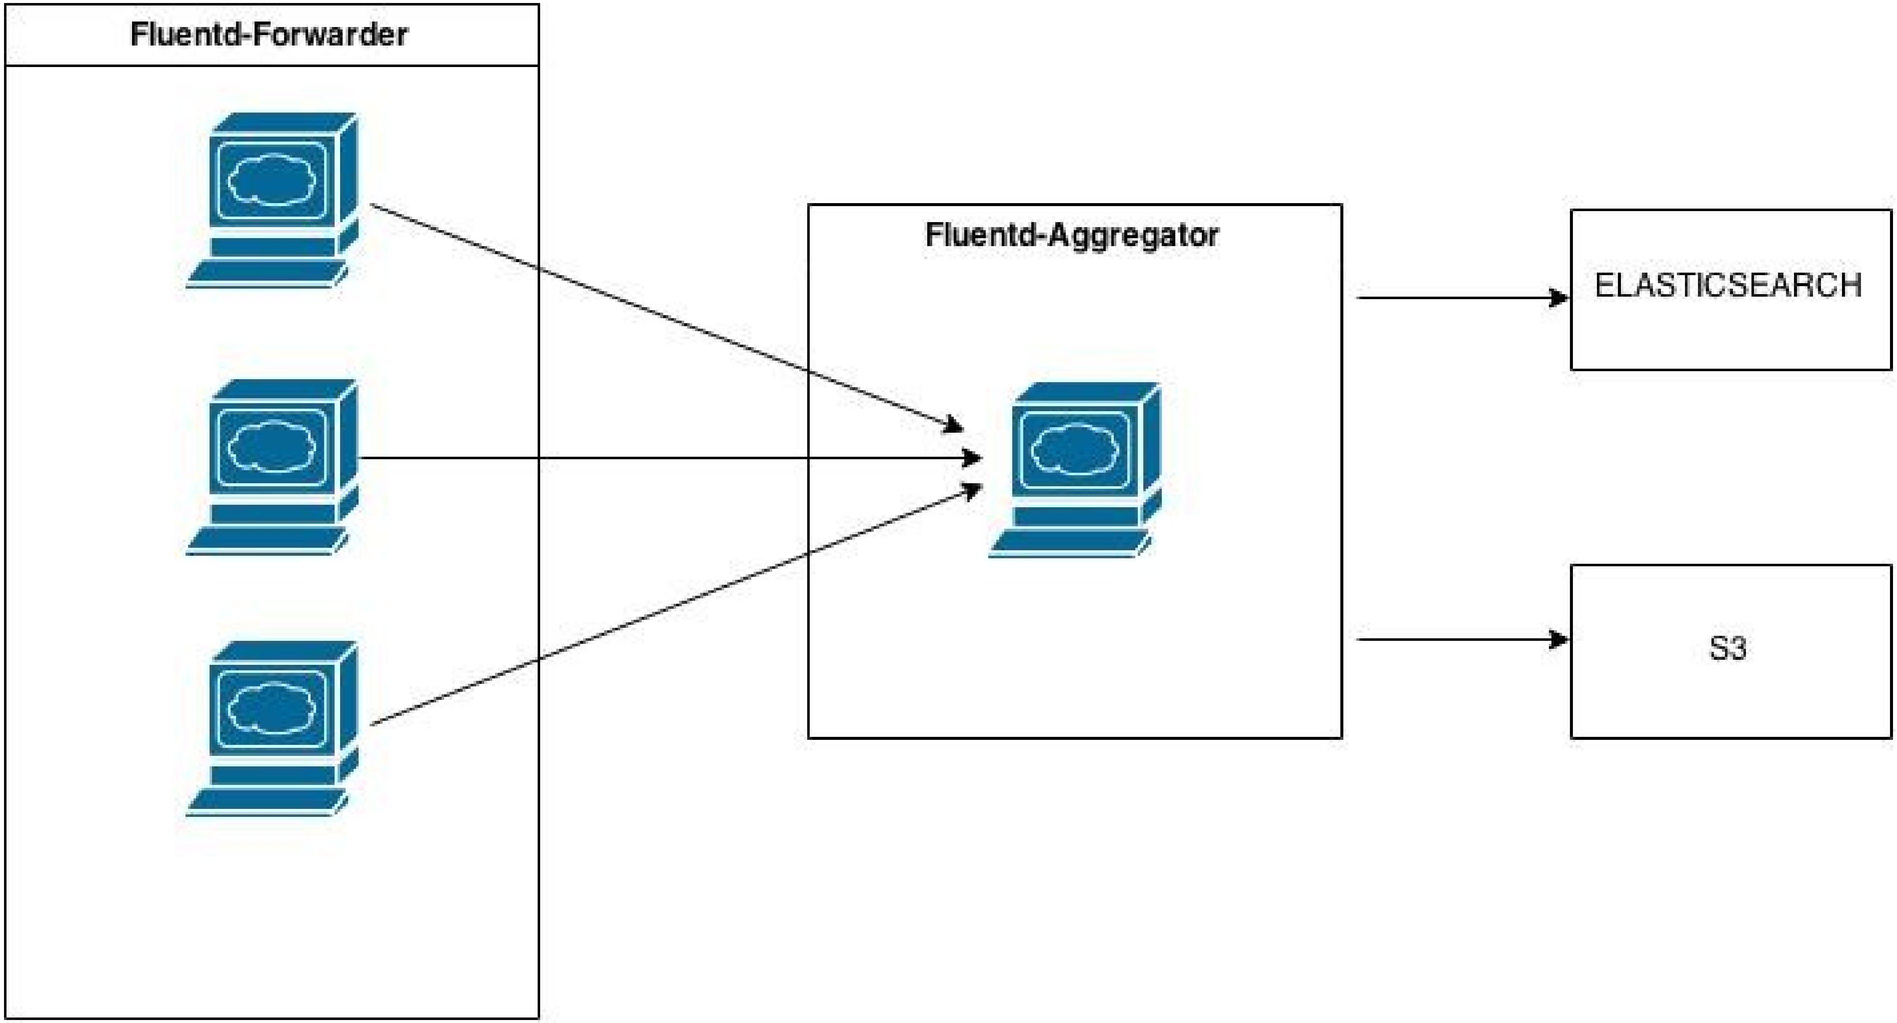

HIGH AVAILABILITY

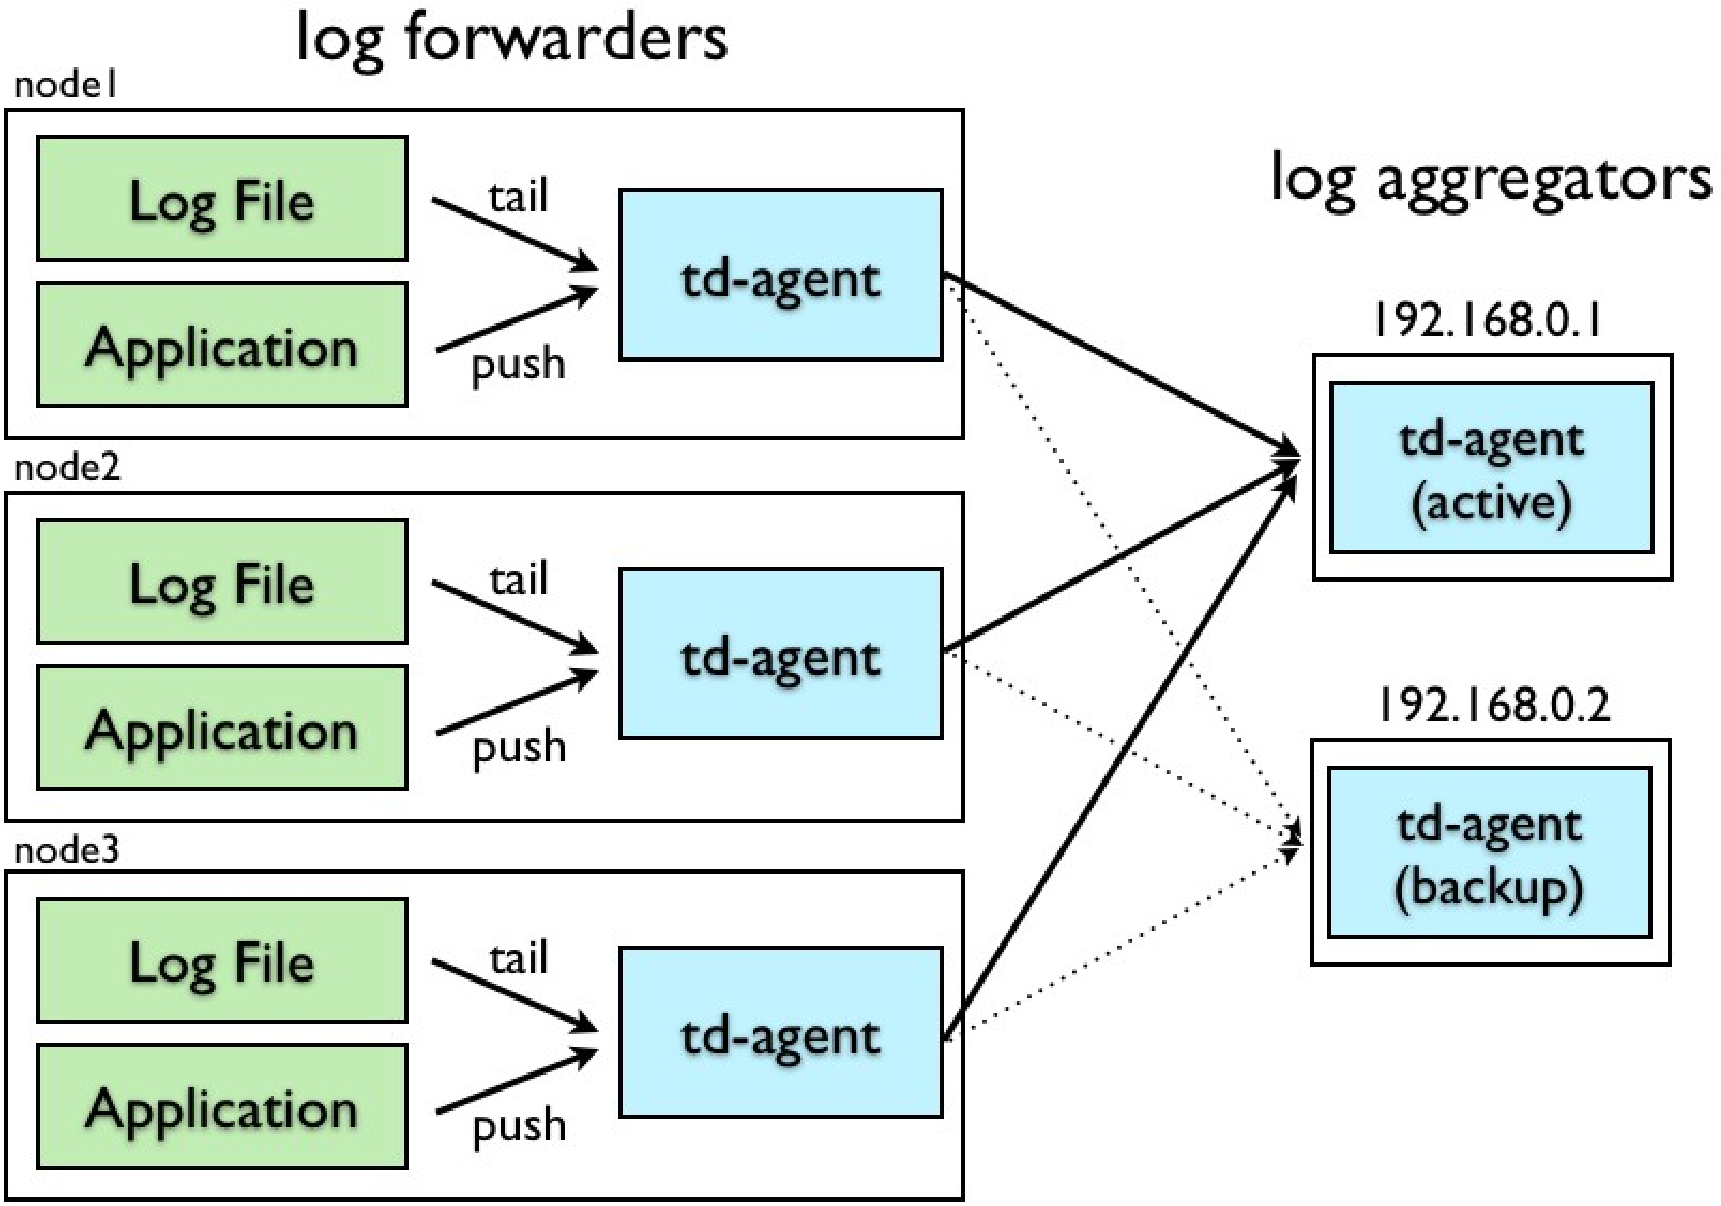

THANK YOU