Centralized Logging Patterns

Philipp Krenn @xeraa

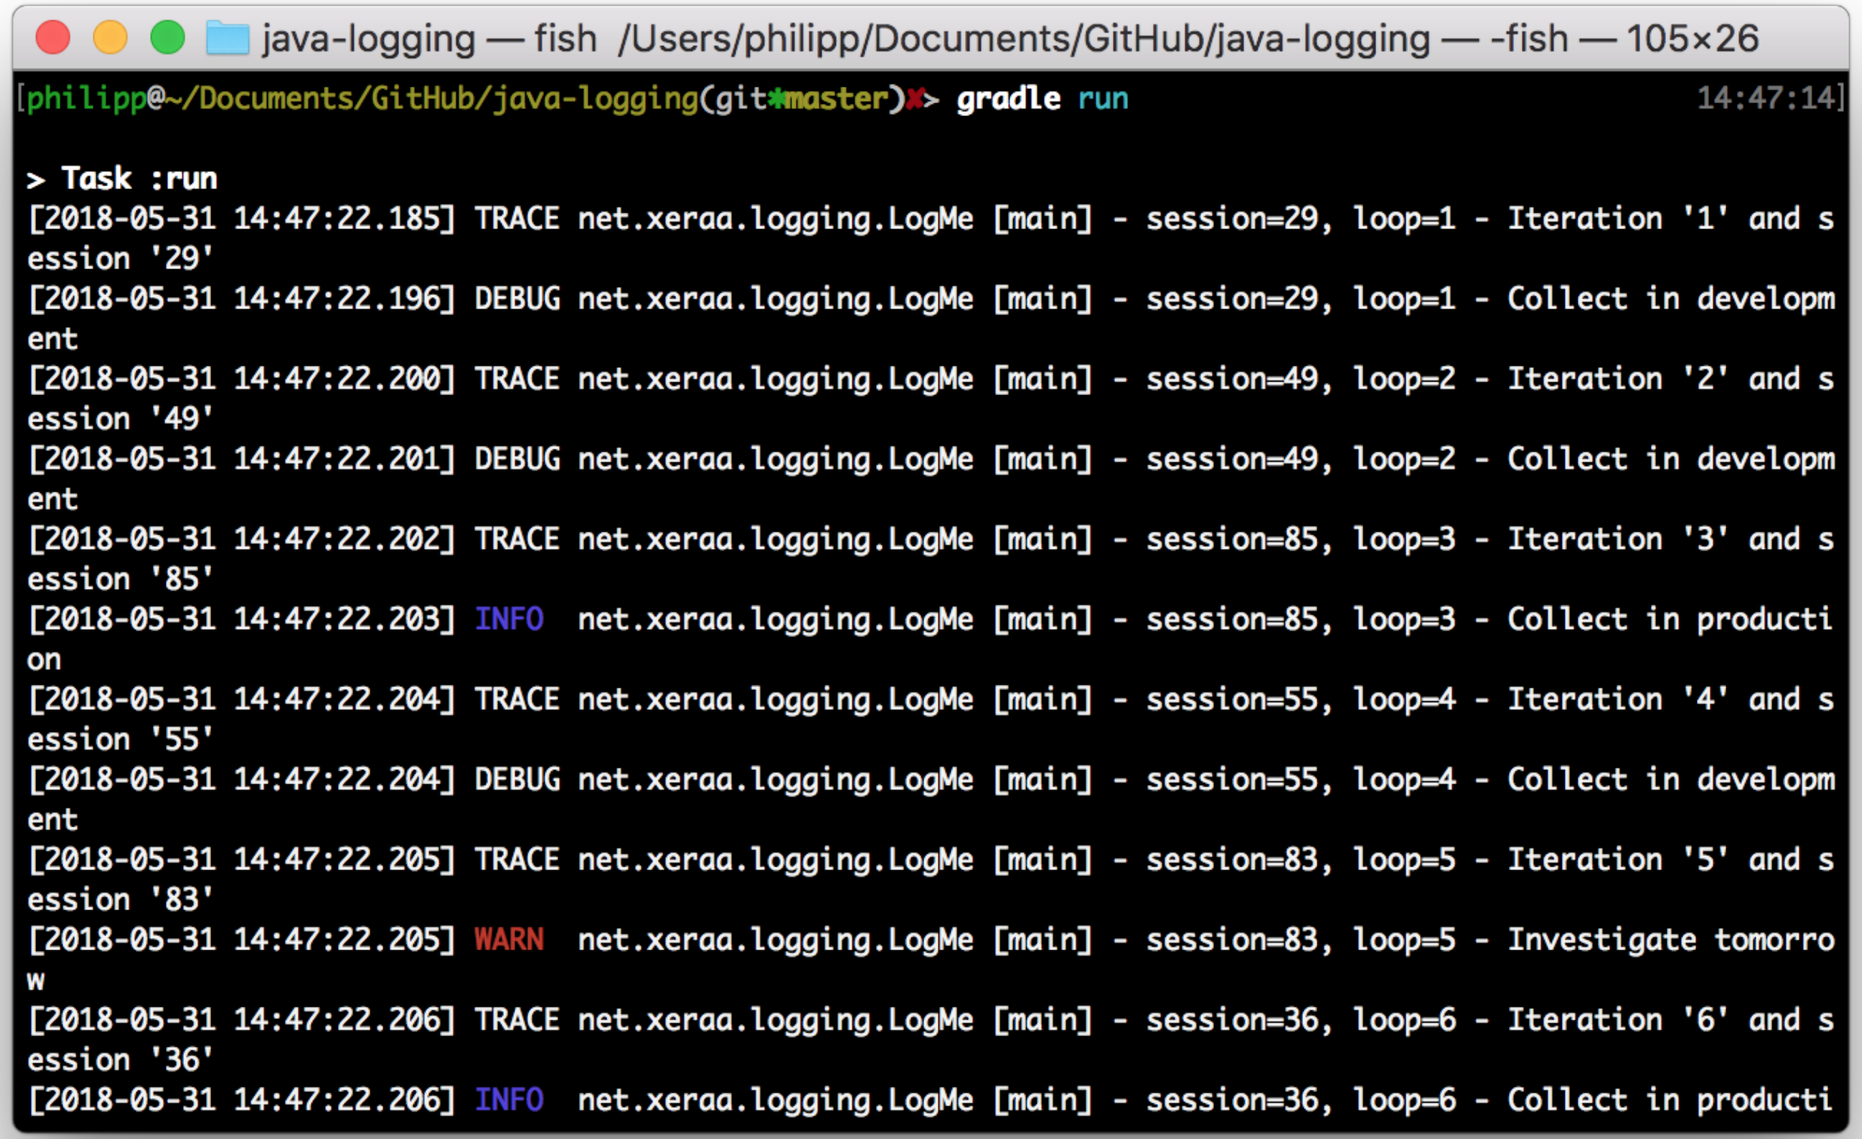
java-logging — fish /Users/philipp/Documents/GitHub/java-logging — -fish — 105×26
[philipp@~/Documents/GitHub/java-logging(git*master)✘> gradle run 14:47:14]
> Task :run
[2018-05-31 14:47:22.185] TRACE net.xeraa.logging.LogMe [main] - session=29, loop=1 - Iteration '1' and session '29'
[2018-05-31 14:47:22.196] DEBUG net.xeraa.logging.LogMe [main] - session=29, loop=1 - Collect in development
[2018-05-31 14:47:22.200] TRACE net.xeraa.logging.LogMe [main] - session=49, loop=2 - Iteration '2' and session '49'
[2018-05-31 14:47:22.201] DEBUG net.xeraa.logging.LogMe [main] - session=49, loop=2 - Collect in development
[2018-05-31 14:47:22.202] TRACE net.xeraa.logging.LogMe [main] - session=85, loop=3 - Iteration '3' and session '85'
[2018-05-31 14:47:22.203] INFO net.xeraa.logging.LogMe [main] - session=85, loop=3 - Collect in production
[2018-05-31 14:47:22.204] TRACE net.xeraa.logging.LogMe [main] - session=55, loop=4 - Iteration '4' and session '55'
[2018-05-31 14:47:22.204] DEBUG net.xeraa.logging.LogMe [main] - session=55, loop=4 - Collect in development
[2018-05-31 14:47:22.205] TRACE net.xeraa.logging.LogMe [main] - session=83, loop=5 - Iteration '5' and session '83'
[2018-05-31 14:47:22.205] WARN net.xeraa.logging.LogMe [main] - session=83, loop=5 - Investigate tomorrow
[2018-05-31 14:47:22.206] TRACE net.xeraa.logging.LogMe [main] - session=36, loop=6 - Iteration '6' and session '36'
[2018-05-31 14:47:22.206] INFO net.xeraa.logging.LogMe [main] - session=36, loop=6 - Collect in producti

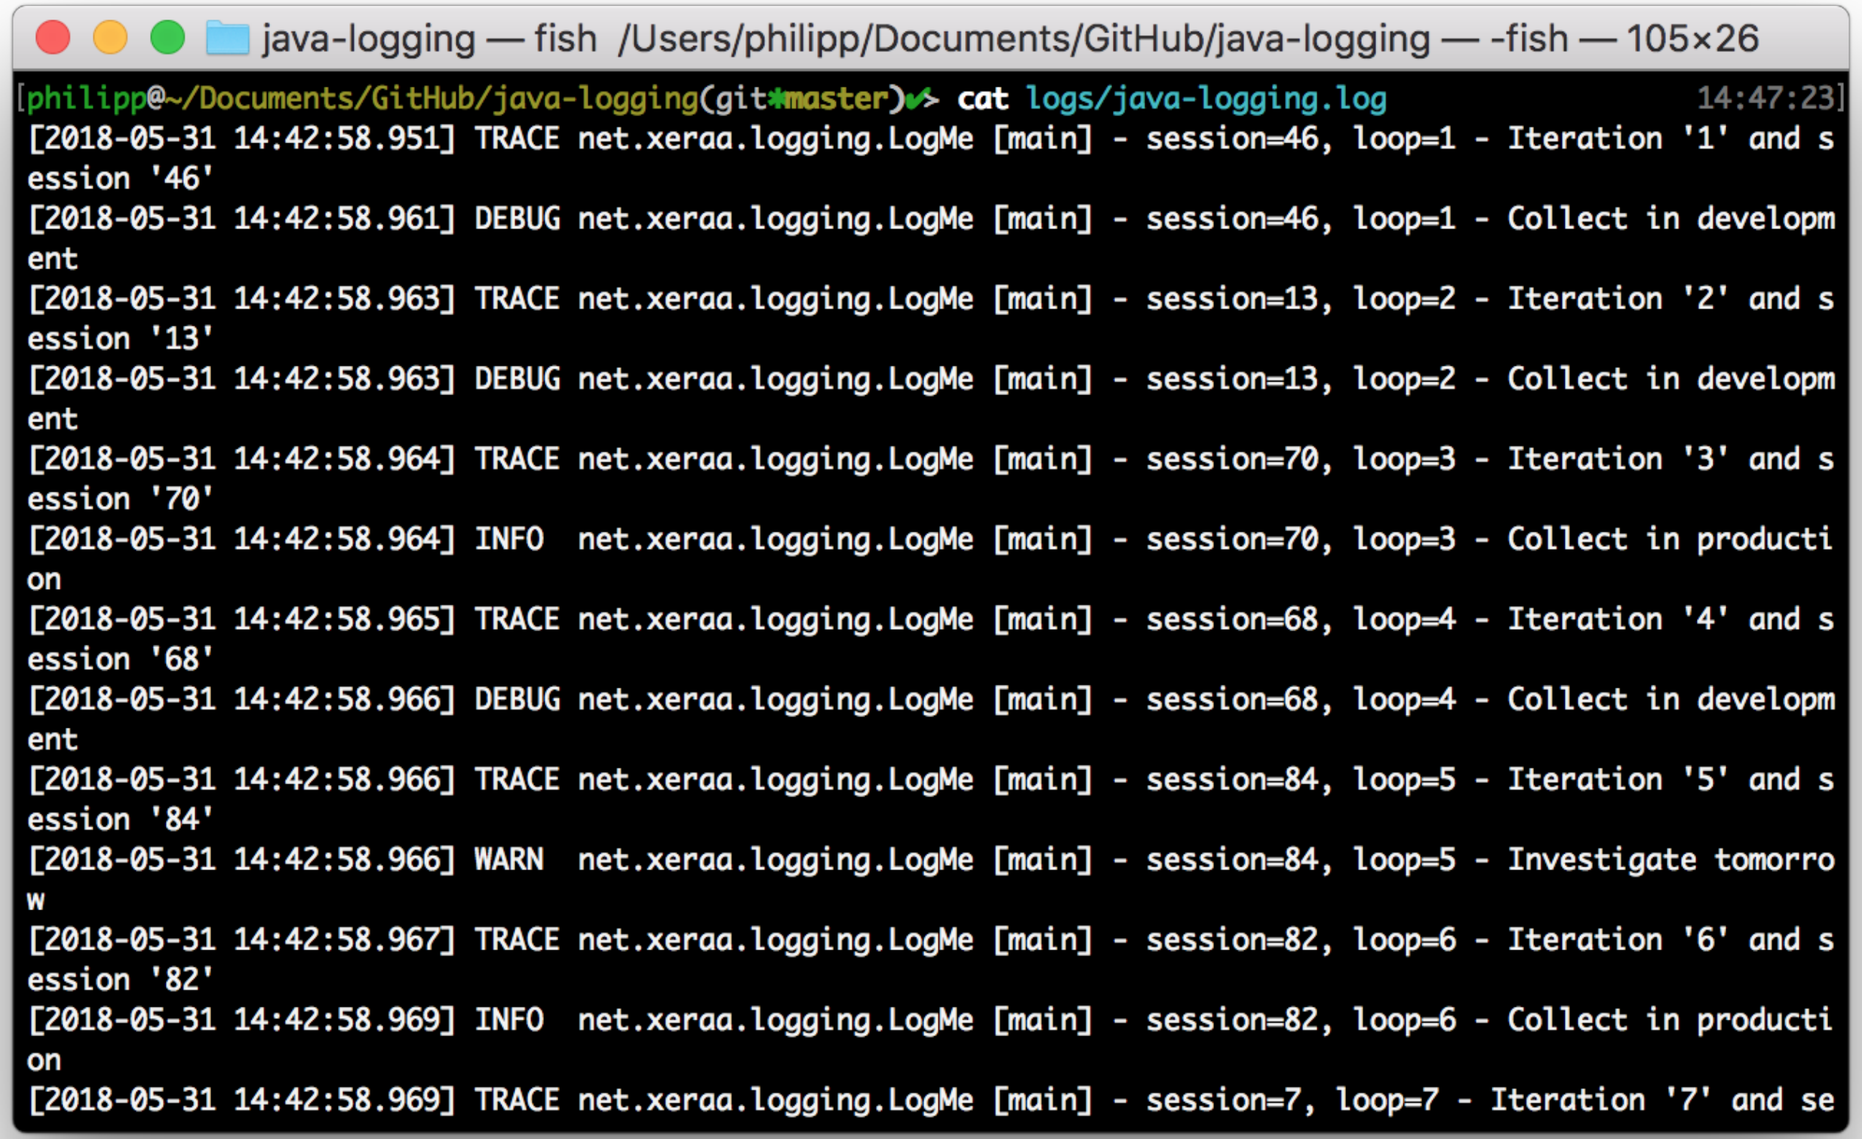
java-logging — fish /Users/philipp/Documents/GitHub/java-logging — -fish — 105×26
[philipp@~/Documents/GitHub/java-logging(git*master)✔> cat logs/java-logging.log 14:47:23]
[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and session '46'
[2018-05-31 14:42:58.961] DEBUG net.xeraa.logging.LogMe [main] - session=46, loop=1 - Collect in development
[2018-05-31 14:42:58.963] TRACE net.xeraa.logging.LogMe [main] - session=13, loop=2 - Iteration '2' and session '13'
[2018-05-31 14:42:58.963] DEBUG net.xeraa.logging.LogMe [main] - session=13, loop=2 - Collect in development
[2018-05-31 14:42:58.964] TRACE net.xeraa.logging.LogMe [main] - session=70, loop=3 - Iteration '3' and session '70'
[2018-05-31 14:42:58.964] INFO net.xeraa.logging.LogMe [main] - session=70, loop=3 - Collect in production
[2018-05-31 14:42:58.965] TRACE net.xeraa.logging.LogMe [main] - session=68, loop=4 - Iteration '4' and session '68'
[2018-05-31 14:42:58.966] DEBUG net.xeraa.logging.LogMe [main] - session=68, loop=4 - Collect in development
[2018-05-31 14:42:58.966] TRACE net.xeraa.logging.LogMe [main] - session=84, loop=5 - Iteration '5' and session '84'
[2018-05-31 14:42:58.966] WARN net.xeraa.logging.LogMe [main] - session=84, loop=5 - Investigate tomorrow
[2018-05-31 14:42:58.967] TRACE net.xeraa.logging.LogMe [main] - session=82, loop=6 - Iteration '6' and session '82'
[2018-05-31 14:42:58.969] INFO net.xeraa.logging.LogMe [main] - session=82, loop=6 - Collect in production
[2018-05-31 14:42:58.969] TRACE net.xeraa.logging.LogMe [main] - session=7, loop=7 - Iteration '7' and se

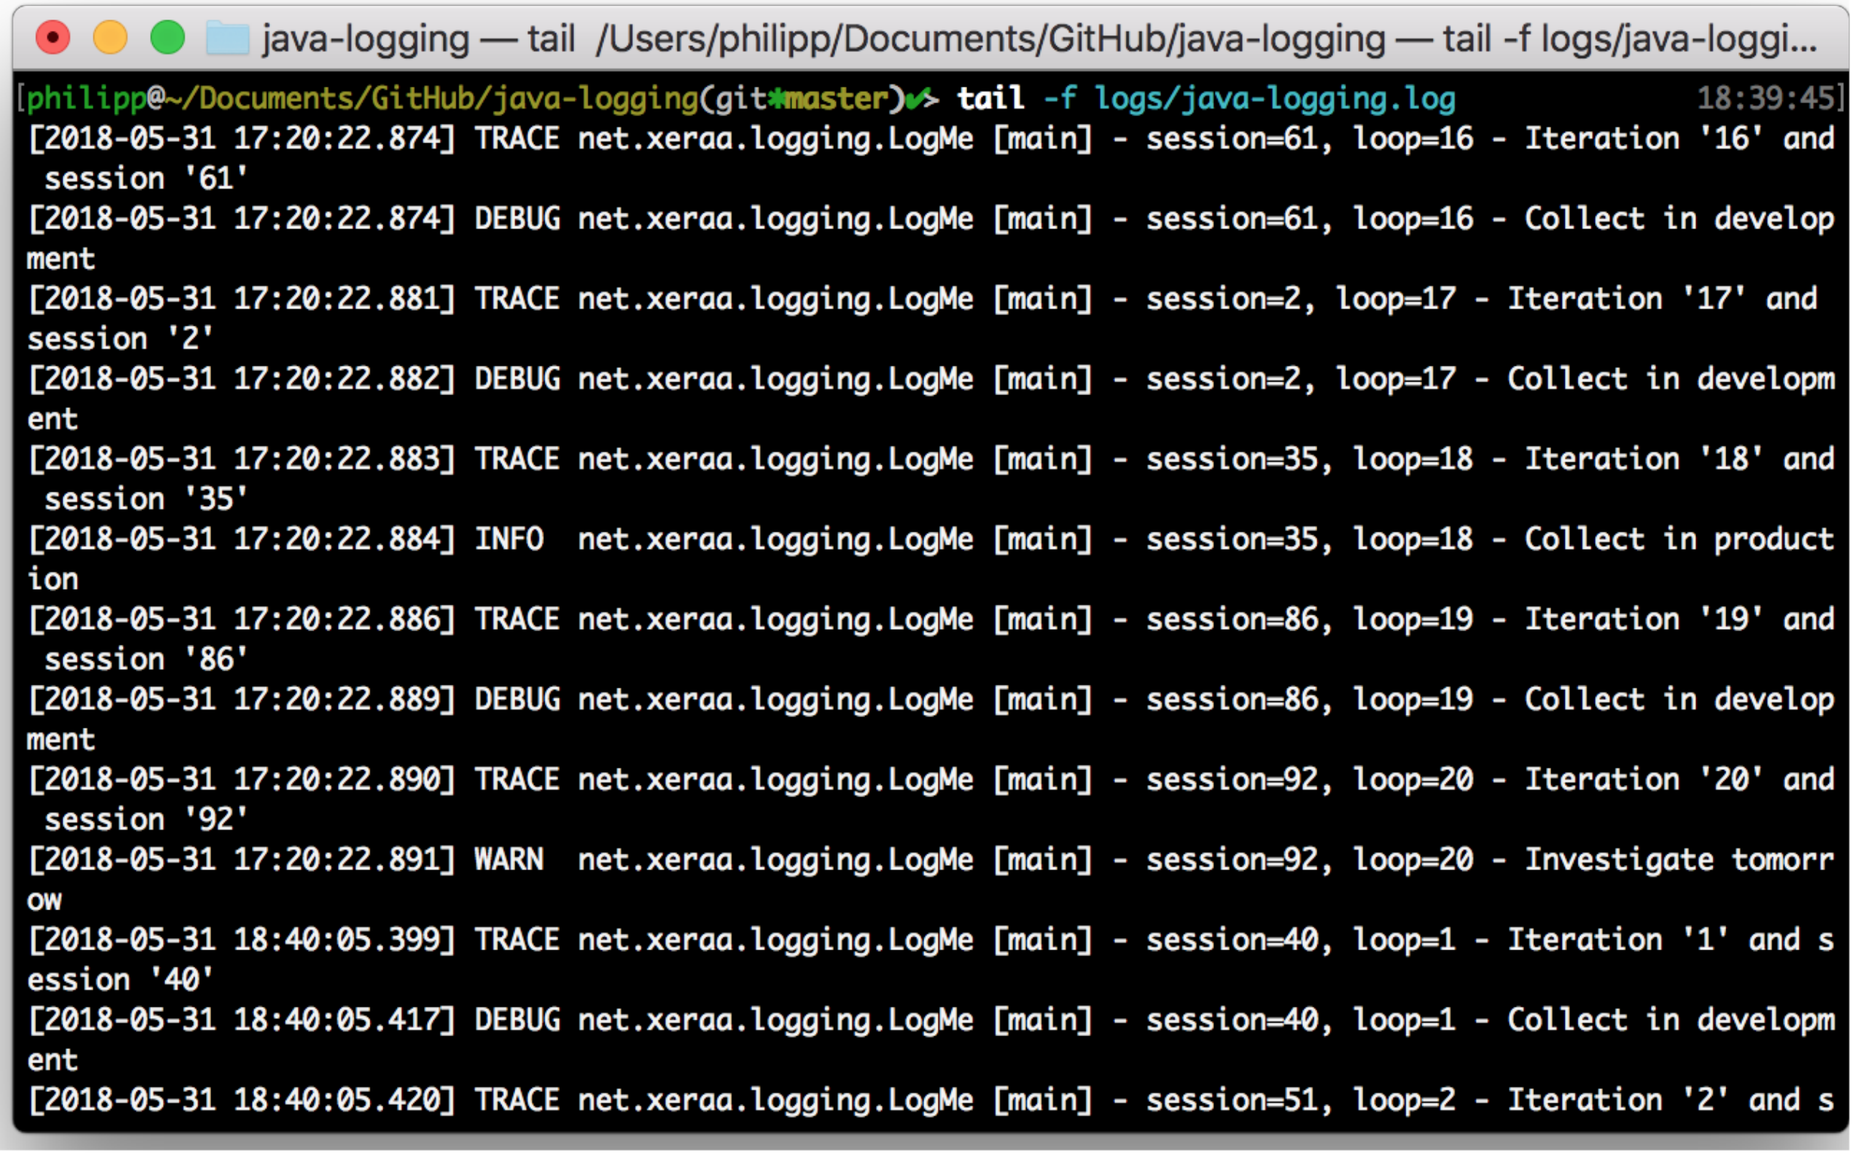
java-logging — tail /Users/philipp/Documents/GitHub/java-logging — tail -f logs/java-loggi...
[philipp@~/Documents/GitHub/java-logging(git*master)✔> tail -f logs/java-logging.log 18:39:45]
[2018-05-31 17:20:22.874] TRACE net.xeraa.logging.LogMe [main] - session=61, loop=16 - Iteration '16' and session '61'
[2018-05-31 17:20:22.874] DEBUG net.xeraa.logging.LogMe [main] - session=61, loop=16 - Collect in development
[2018-05-31 17:20:22.881] TRACE net.xeraa.logging.LogMe [main] - session=2, loop=17 - Iteration '17' and session '2'
[2018-05-31 17:20:22.882] DEBUG net.xeraa.logging.LogMe [main] - session=2, loop=17 - Collect in development
[2018-05-31 17:20:22.883] TRACE net.xeraa.logging.LogMe [main] - session=35, loop=18 - Iteration '18' and session '35'
[2018-05-31 17:20:22.884] INFO net.xeraa.logging.LogMe [main] - session=35, loop=18 - Collect in production
[2018-05-31 17:20:22.886] TRACE net.xeraa.logging.LogMe [main] - session=86, loop=19 - Iteration '19' and session '86'
[2018-05-31 17:20:22.889] DEBUG net.xeraa.logging.LogMe [main] - session=86, loop=19 - Collect in development
[2018-05-31 17:20:22.890] TRACE net.xeraa.logging.LogMe [main] - session=92, loop=20 - Iteration '20' and session '92'
[2018-05-31 17:20:22.891] WARN net.xeraa.logging.LogMe [main] - session=92, loop=20 - Investigate tomorrow
[2018-05-31 18:40:05.399] TRACE net.xeraa.logging.LogMe [main] - session=40, loop=1 - Iteration '1' and session '40'
[2018-05-31 18:40:05.417] DEBUG net.xeraa.logging.LogMe [main] - session=40, loop=1 - Collect in development
[2018-05-31 18:40:05.420] TRACE net.xeraa.logging.LogMe [main] - session=51, loop=2 - Iteration '2' and s

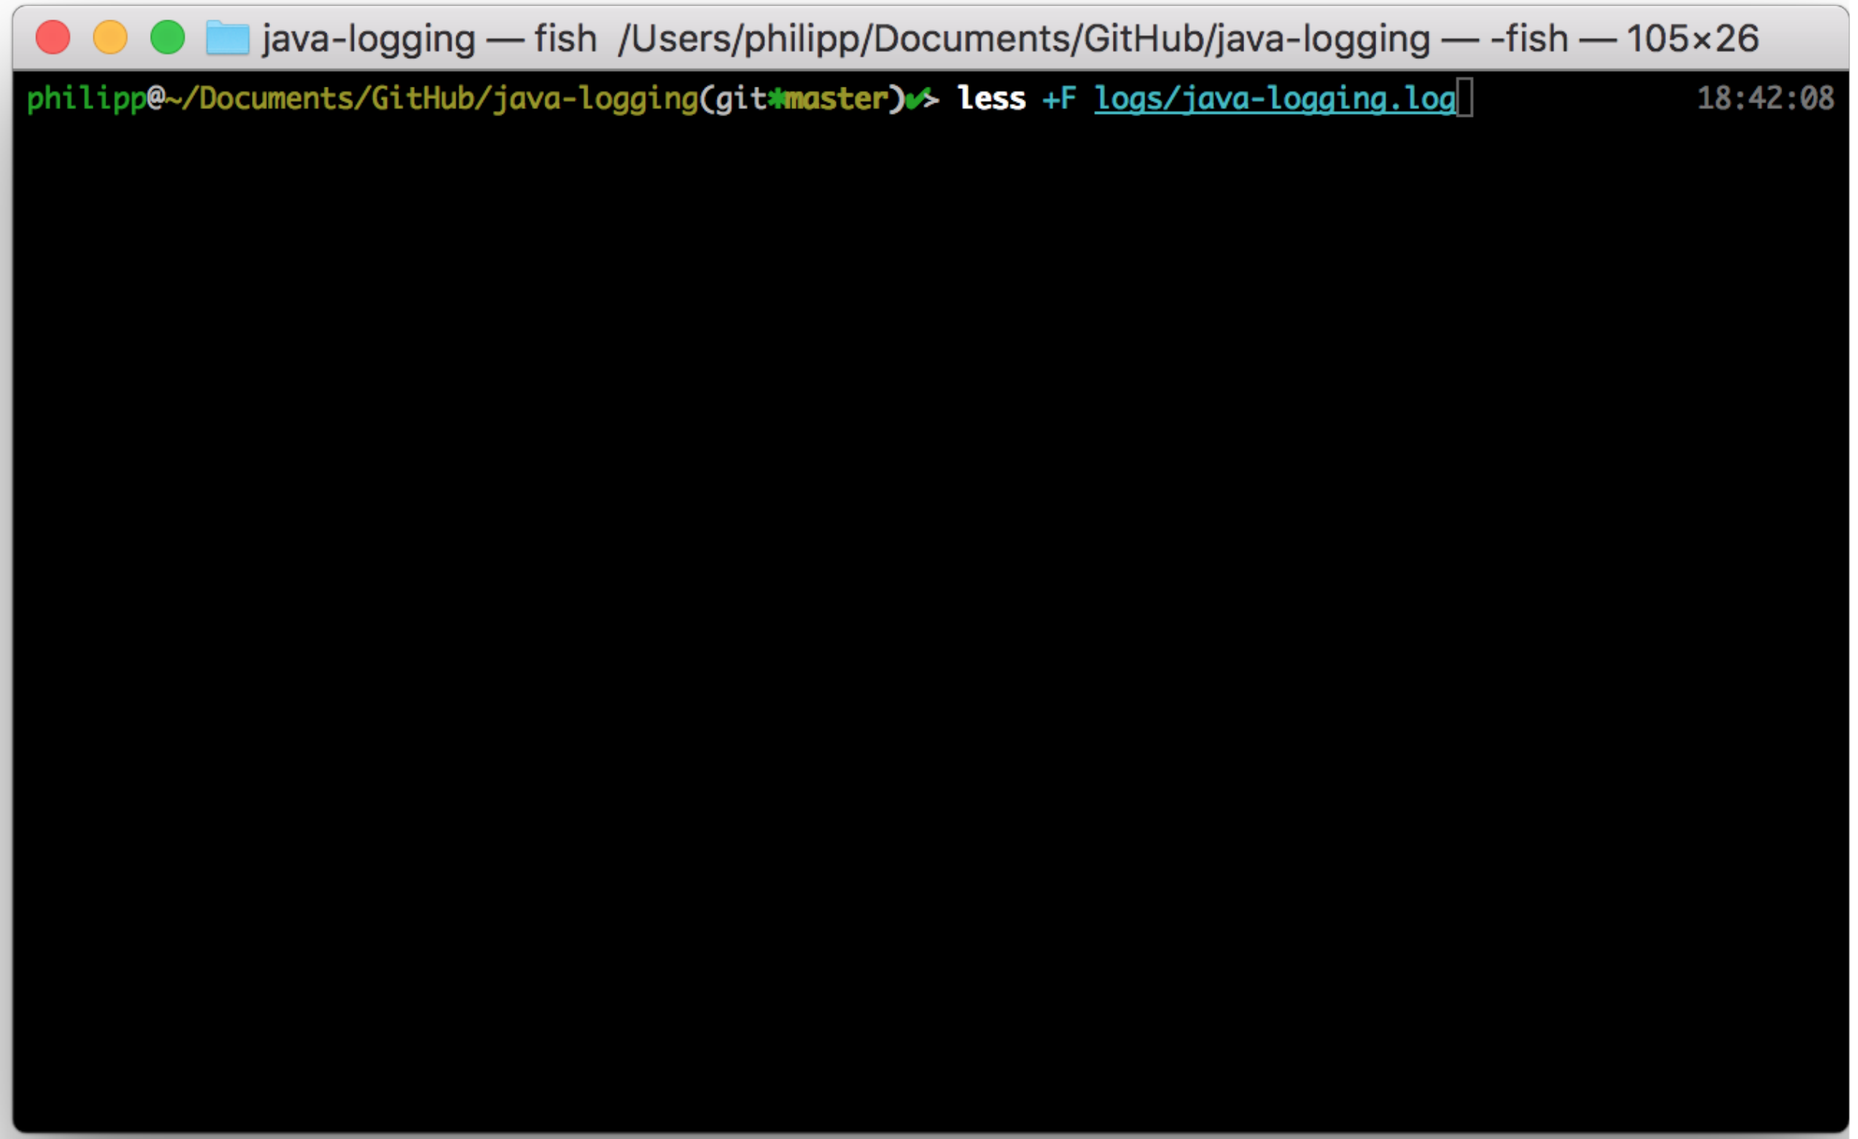
java-logging — fish /Users/philipp/Documents/GitHub/java-logging — -fish — 105×26
philipp@~/Documents/GitHub/java-logging(git:master)> less +F logs/java-logging.log
18:42:08


THIS IS FINE.

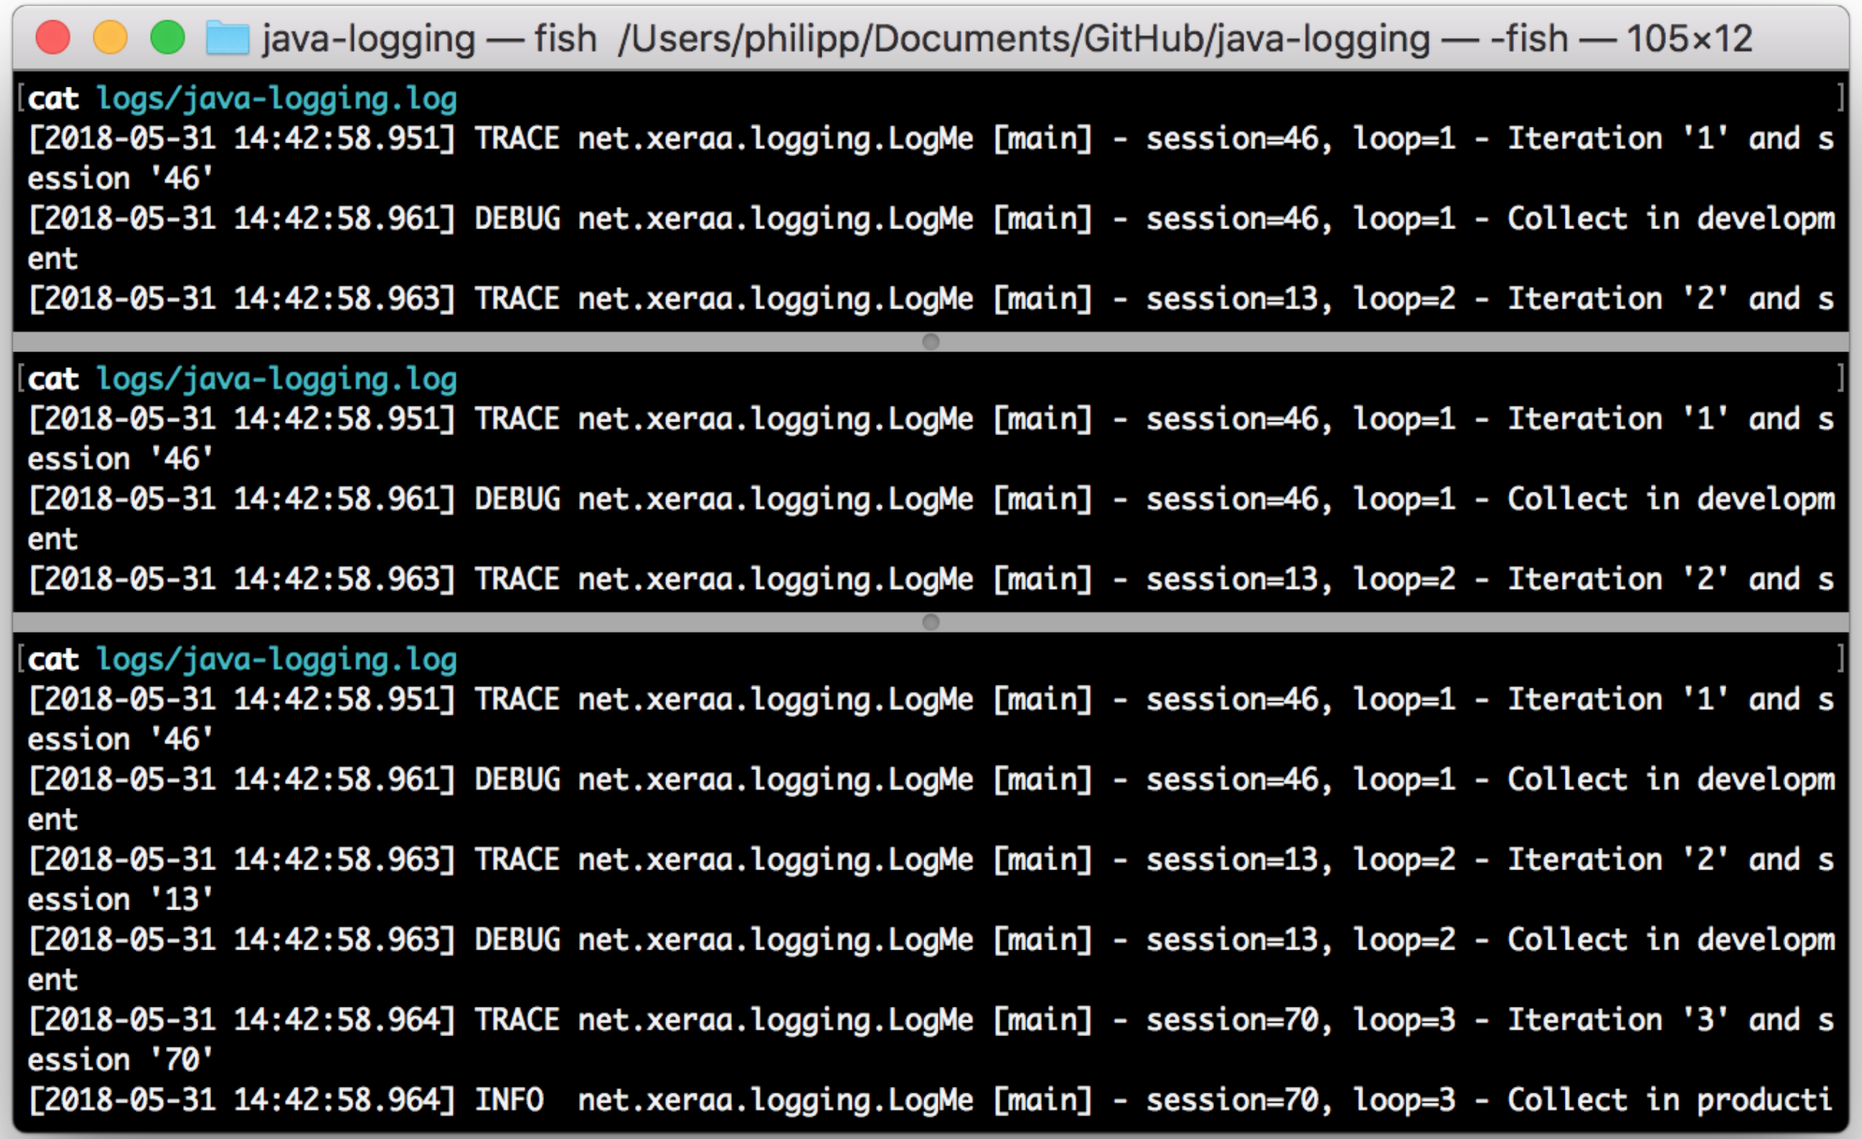
java-logging — fish /Users/philipp/Documents/GitHub/java-logging — -fish — 105×12
cat logs/java-logging.log
[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and session '46'
[2018-05-31 14:42:58.961] DEBUG net.xeraa.logging.LogMe [main] - session=46, loop=1 - Collect in development
[2018-05-31 14:42:58.963] TRACE net.xeraa.logging.LogMe [main] - session=13, loop=2 - Iteration '2' and s
cat logs/java-logging.log
[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and session '46'
[2018-05-31 14:42:58.961] DEBUG net.xeraa.logging.LogMe [main] - session=46, loop=1 - Collect in development
[2018-05-31 14:42:58.963] TRACE net.xeraa.logging.LogMe [main] - session=13, loop=2 - Iteration '2' and s
cat logs/java-logging.log
[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and session '46'
[2018-05-31 14:42:58.961] DEBUG net.xeraa.logging.LogMe [main] - session=46, loop=1 - Collect in development
[2018-05-31 14:42:58.963] TRACE net.xeraa.logging.LogMe [main] - session=13, loop=2 - Iteration '2' and session '13'
[2018-05-31 14:42:58.963] DEBUG net.xeraa.logging.LogMe [main] - session=13, loop=2 - Collect in development
[2018-05-31 14:42:58.964] TRACE net.xeraa.logging.LogMe [main] - session=70, loop=3 - Iteration '3' and session '70'
[2018-05-31 14:42:58.964] INFO net.xeraa.logging.LogMe [main] - session=70, loop=3 - Collect in producti

java-logging — fish /Users/philipp/Documents/GitHub/java-logging — -fish — 105×1

```
[cat logs/java-logging.log                                                                              ]
[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and s

[cat logs/java-logging.log                                                                              ]
[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and s

[cat logs/java-logging.log                                                                              ]
[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and s

[cat logs/java-logging.log                                                                              ]
[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and s

[cat logs/java-logging.log                                                                              ]
[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and s

[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and s
ession '46'

[cat logs/java-logging.log                                                                              ]
[2018-05-31 14:42:58.951] TRACE net.xeraa.logging.LogMe [main] - session=46, loop=1 - Iteration '1' and s

[cat logs/java-logging.log                                                                              ]
```


me looking
for the bug
7.2 GB
of log file


TH-
THIS IS NOT FINE!!
OH MY GOD EVERY THINGS ON FIRE
WHAT THE HELL IS MY PROBLEM


ALL THE THINGS!

elastic

Developer 🥑

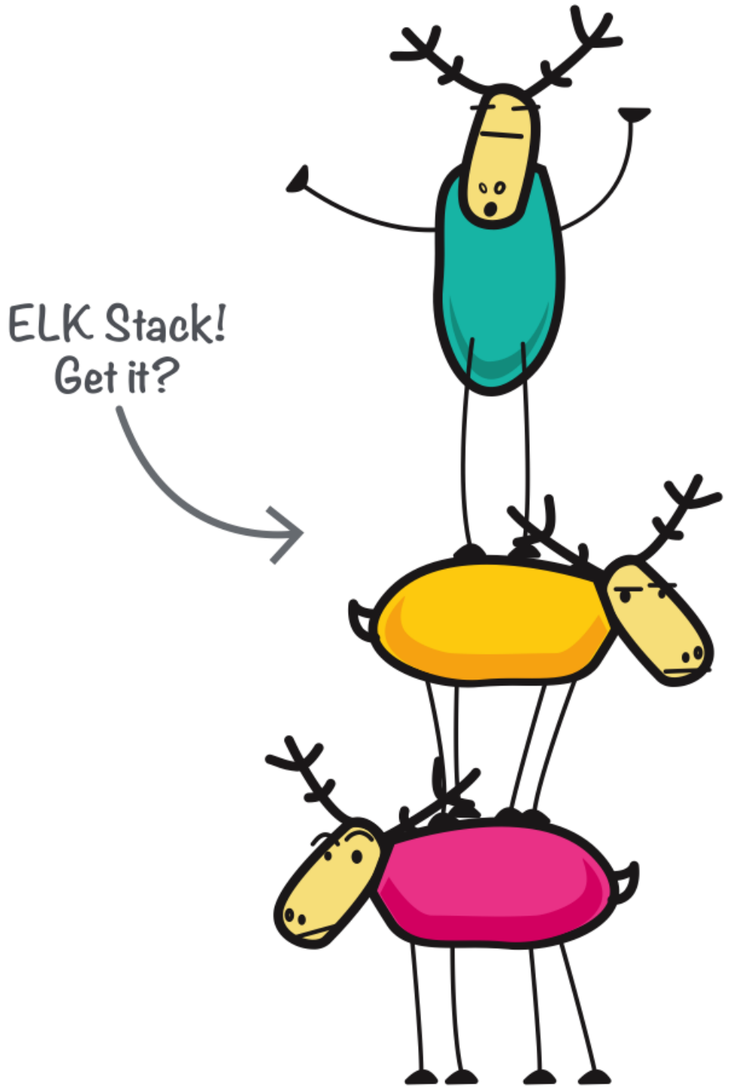

E Elasticsearch

L Logstash

K Kibana


lyft


slack


fitbit

Apparently, I'm an ELKB personality.

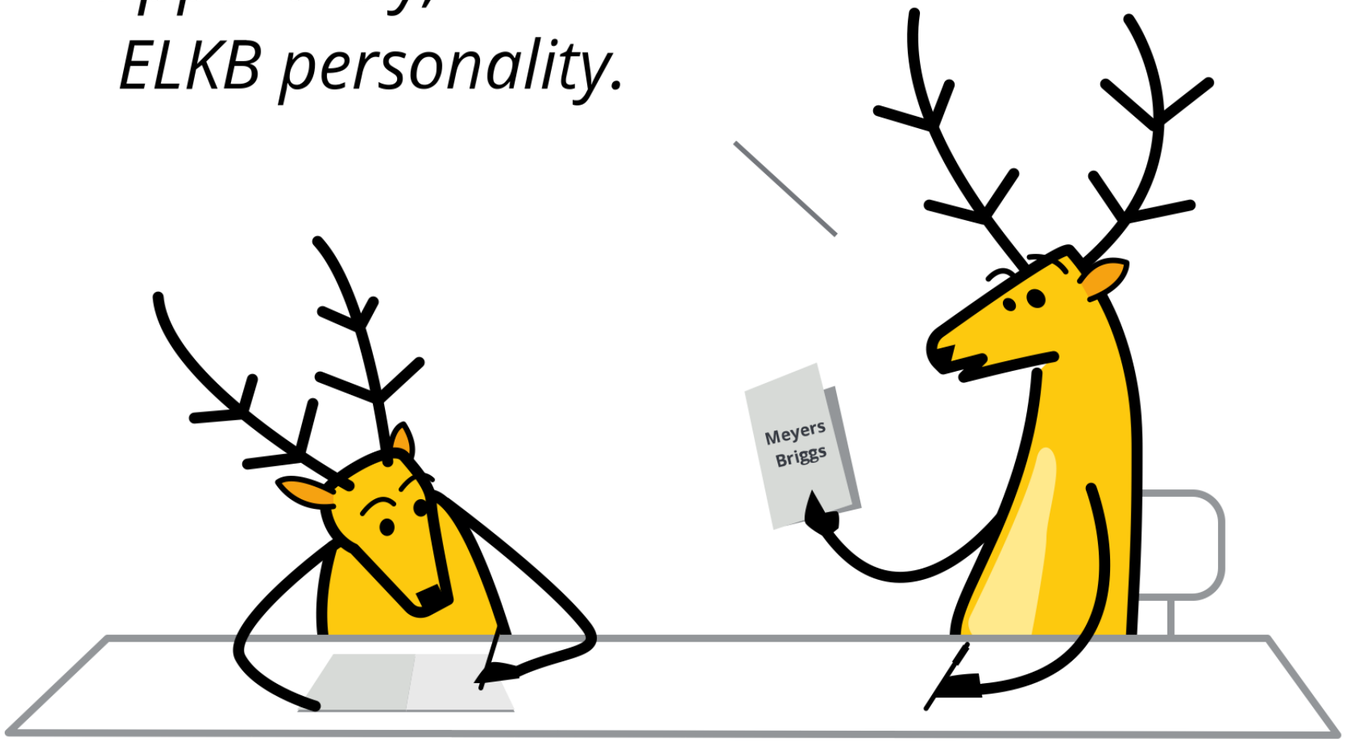

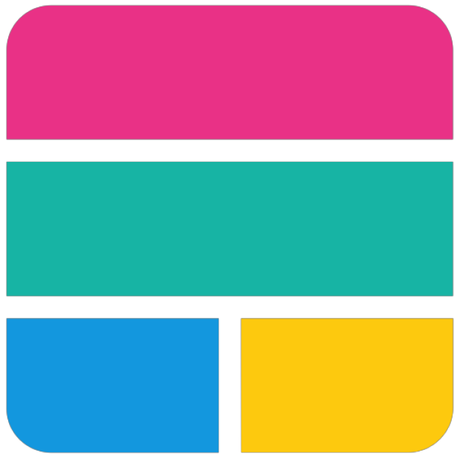

elastic stack

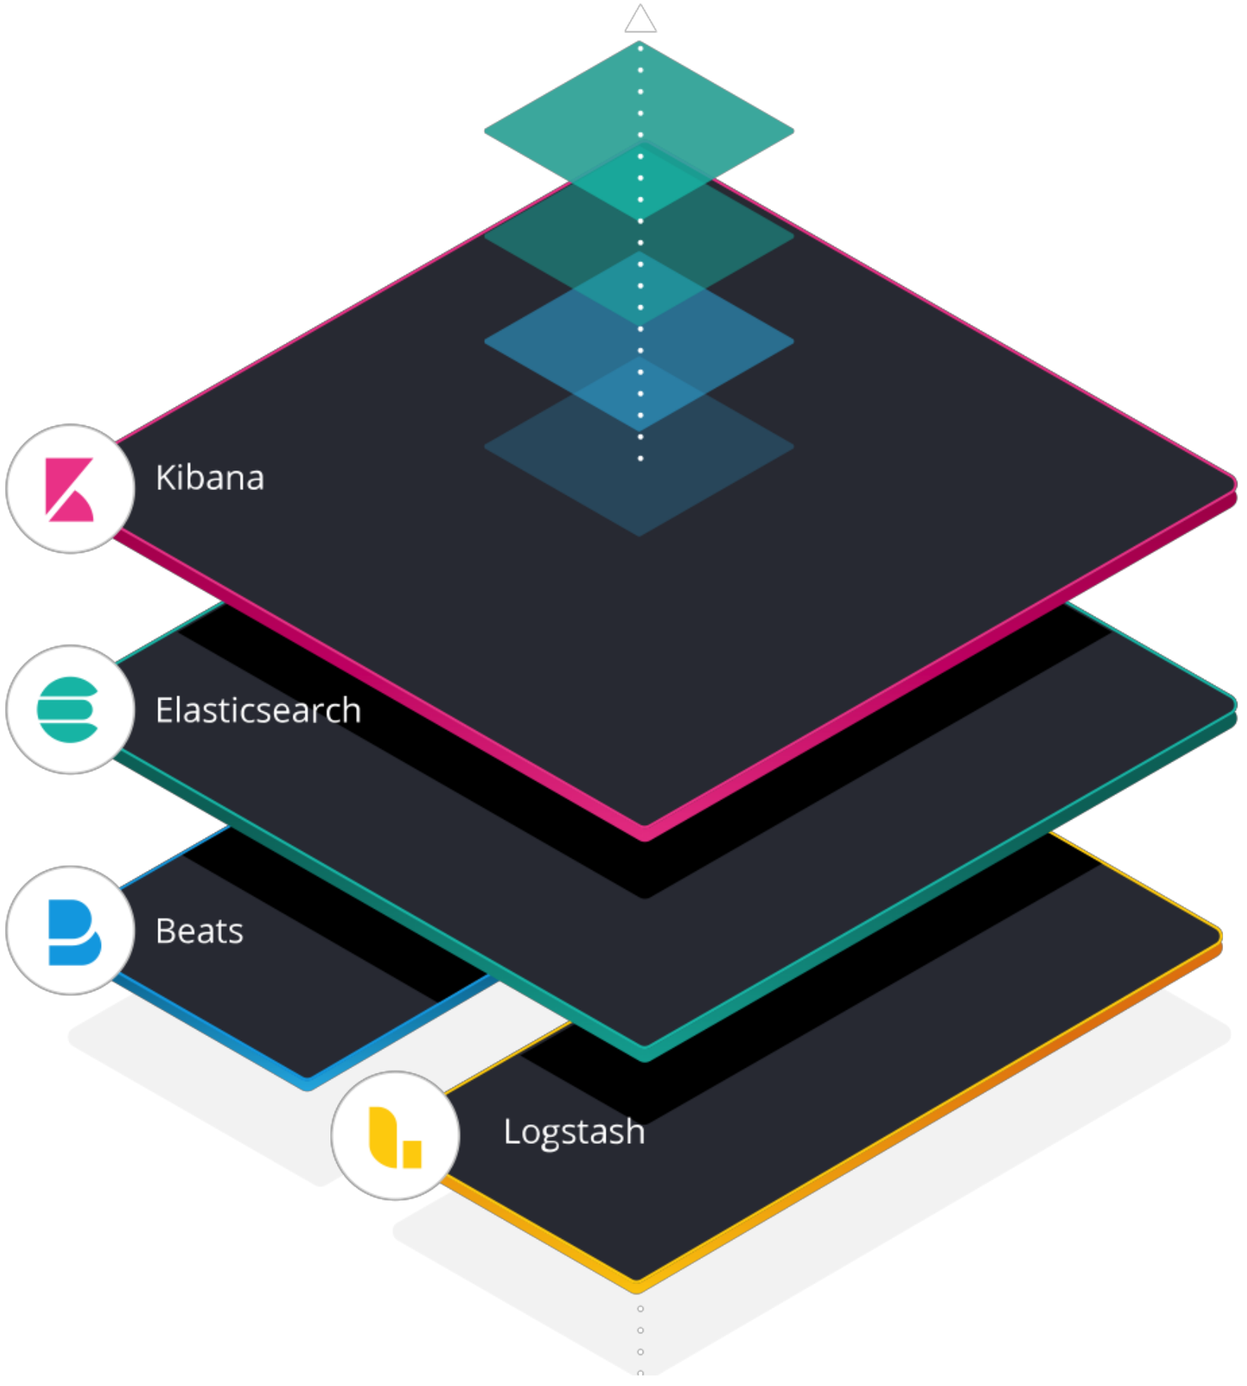
Kibana
Elasticsearch
Beats
Logstash

Disclaimer

I build **highly** monitored Hello World apps

Example: Java

SLF4J, Logback, MDC
with logstash-logback-encoder

Alternative https://github.com/vy/log4j2-logstash-layout

And Everywhere Else

.NET: NLog

JavaScript: Winston

Python: structlog

PHP: Monolog

Anti-Pattern: `print`

```
System.out.println("Oops");
```

Anti-Pattern: Coupling

Parse

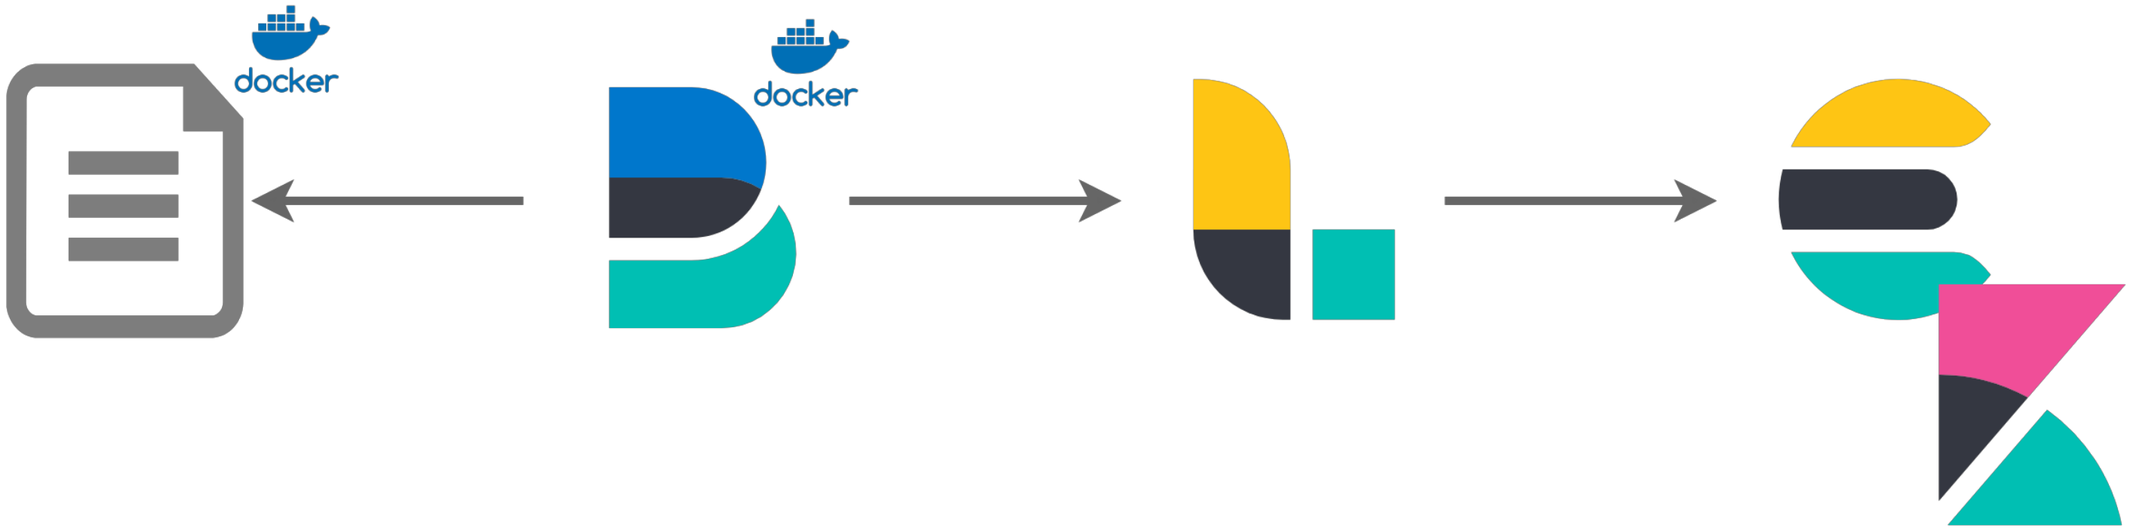
docker
docker

Collect Log Lines

```
filebeat.inputs:
- type: log
  paths:
    - /mnt/logs/*.log
```

Setting for `negate`	Setting for `match`	Result	Example `pattern: ^b`
`false`	`after`	Consecutive lines that match the pattern are appended to the previous line that doesn't match.	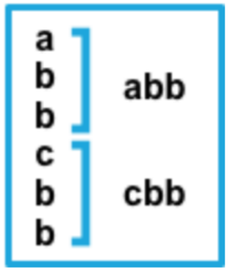
`false`	`before`	Consecutive lines that match the pattern are prepended to the next line that doesn't match.	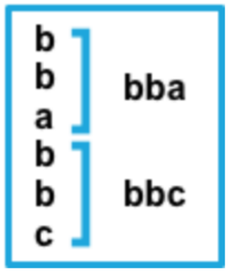
`true`	`after`	Consecutive lines that don't match the pattern are appended to the previous line that does match.	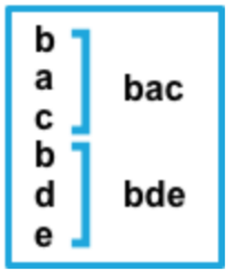
`true`	`before`	Consecutive lines that don't match the pattern are prepended to the next line that does match.	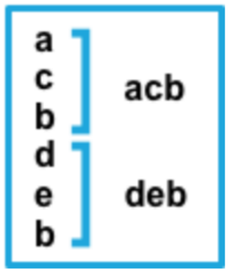

Grok

https://github.com/logstash-plugins/logstash-patterns-core/blob/master/patterns/grok-patterns

Dev Tools

Grok Debugger

elastic @xeraa

```
[2018-09-28 10:30:38.516] ERROR net.xeraa.logging.LogMe [main] -
                           user_experience=🤬, session=46, loop=15 -
                           Wake me up at night
java.lang.RuntimeException: Bad runtime...
  at net.xeraa.logging.LogMe.main(LogMe.java:30)
```

```
^\[%{TIMESTAMP_ISO8601:timestamp}\]%{SPACE}%{LOGLEVEL:level}
  %{SPACE}%{USERNAME:logger}%{SPACE}\[%{WORD:thread}\]
  %{SPACE}-%{SPACE}%{GREEDYDATA:mdc}%{SPACE}-%{SPACE}
  %{GREEDYDATA:themessage}(?:\n+(?<stacktrace>(?:.|\r|\n)+))?
```

Logstash Key Value Filter for MDC

```
kv {
  source => "labels"
  field_split => ","
  trim_key => " "
}
```

Machine Learning

Data Visualizer

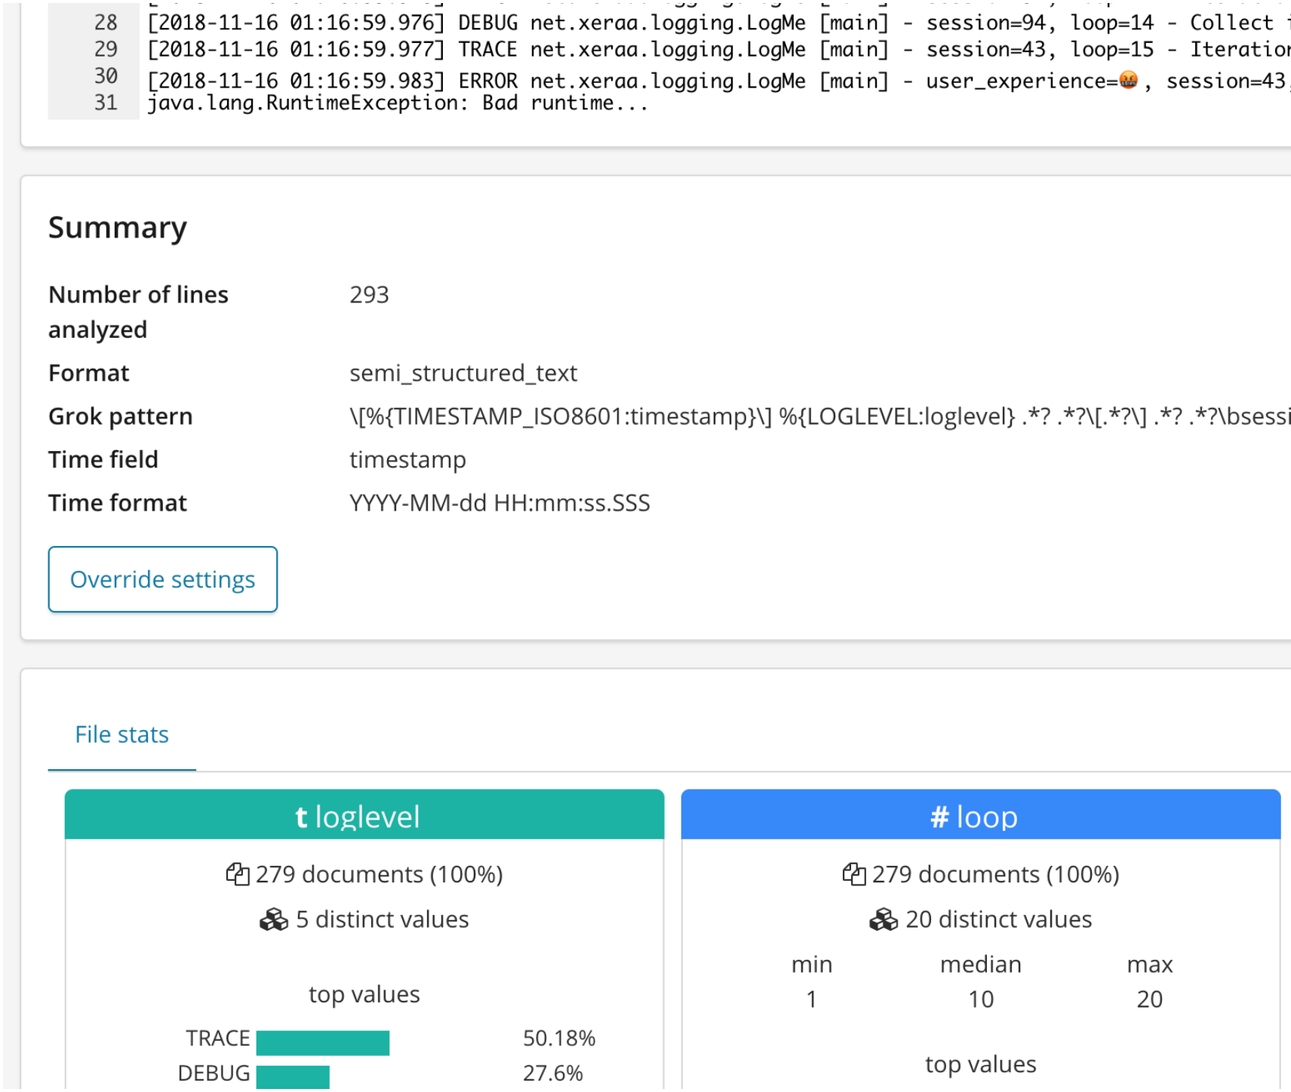

elastic @xeraa

Pro: No change

Con: Regular expression, multiline, format changes

Send 

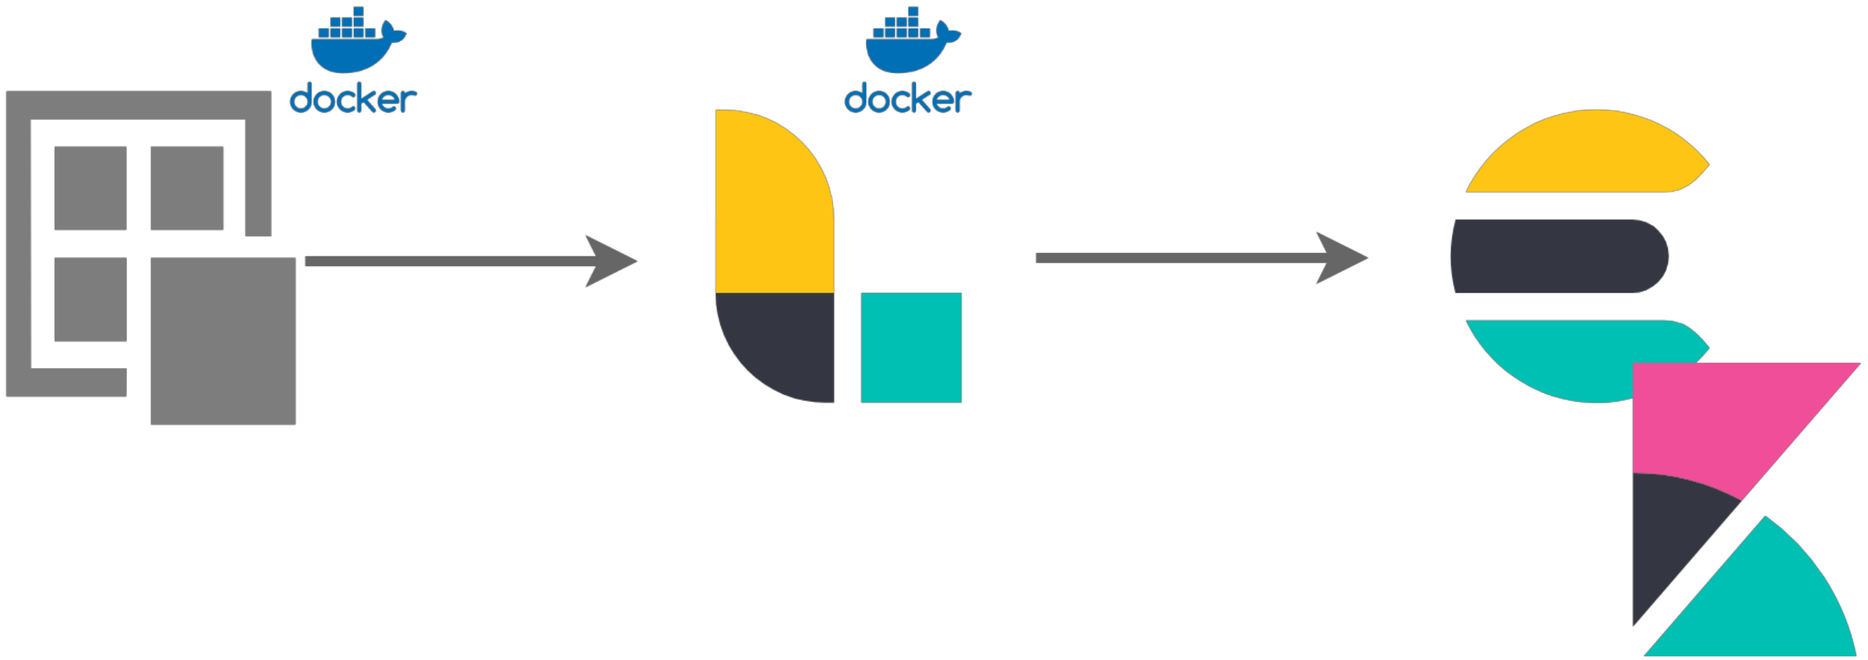
docker
docker

logback.xml

```
<appender name="logstash" class="net.logstash.logback.appender.LogstashAccessTcpSocketAppender">
    <destination>logstash:4560</destination>
    <encoder class="net.logstash.logback.encoder.LogstashEncoder"/>
</appender>
```

Pro: No files

Con: Outages & coupling

Structure

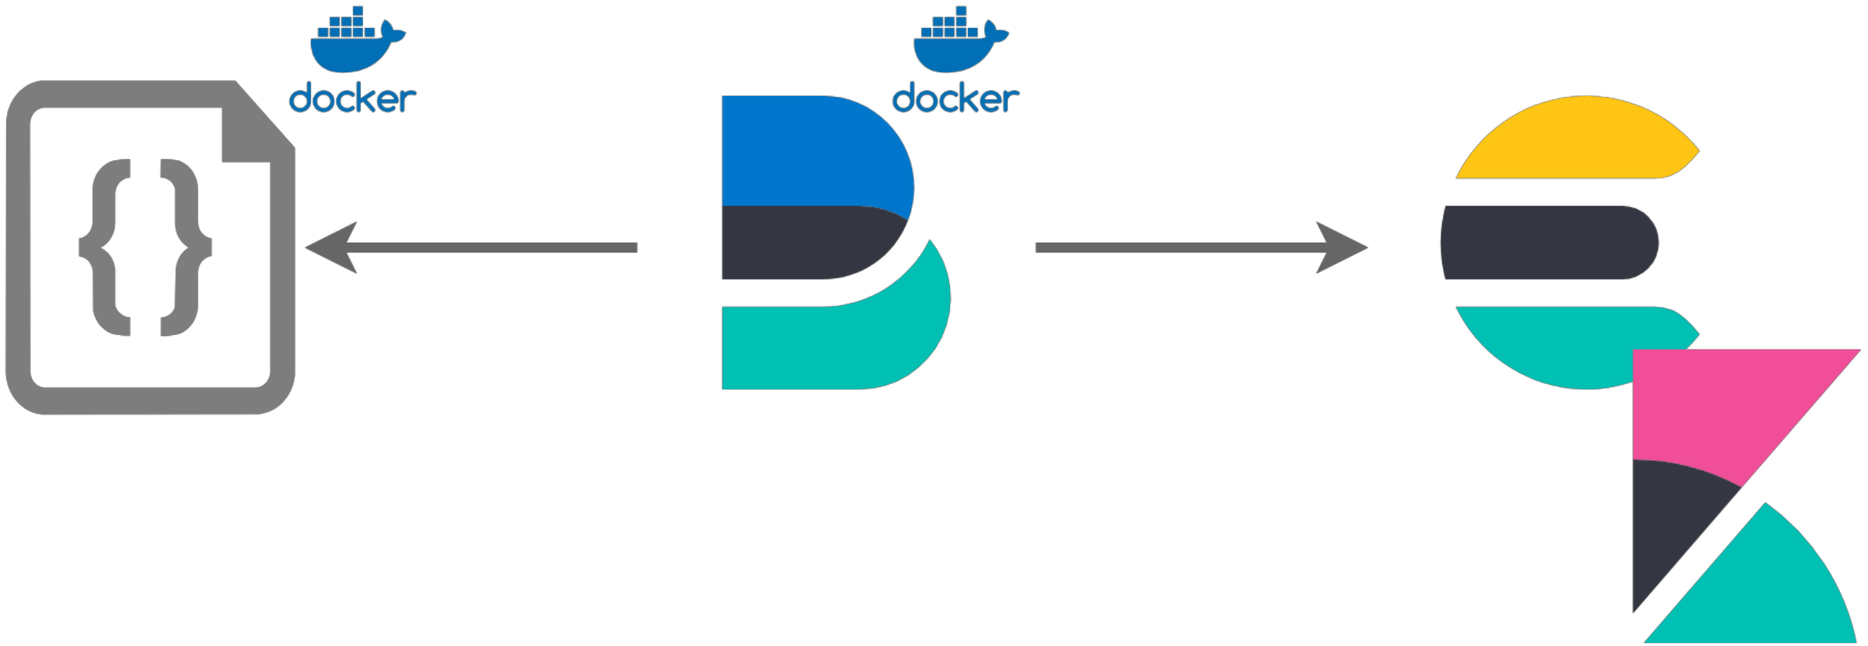
docker
docker

Collect JSON

```
filebeat.input:
- type: log
  paths:
    - /mnt/logs/*.json
  json:
    message_key: message
    keys_under_root: true
```

Stack(trace) Hash

Pro: Right format

Con: JSON serialization overhead

Containerize 

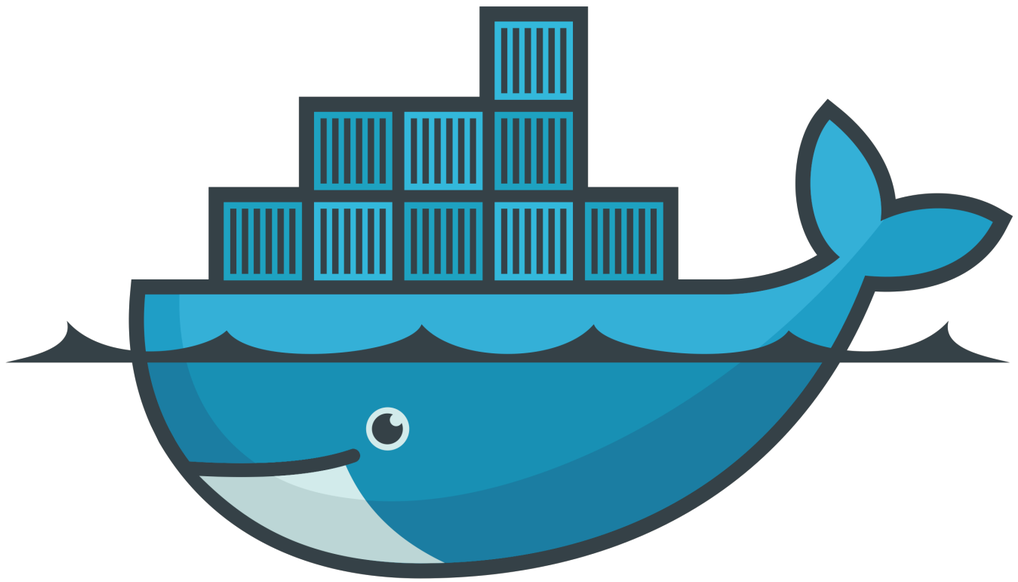

docker

elastic

@xeraa

Where to put Filebeat?

Sidecar

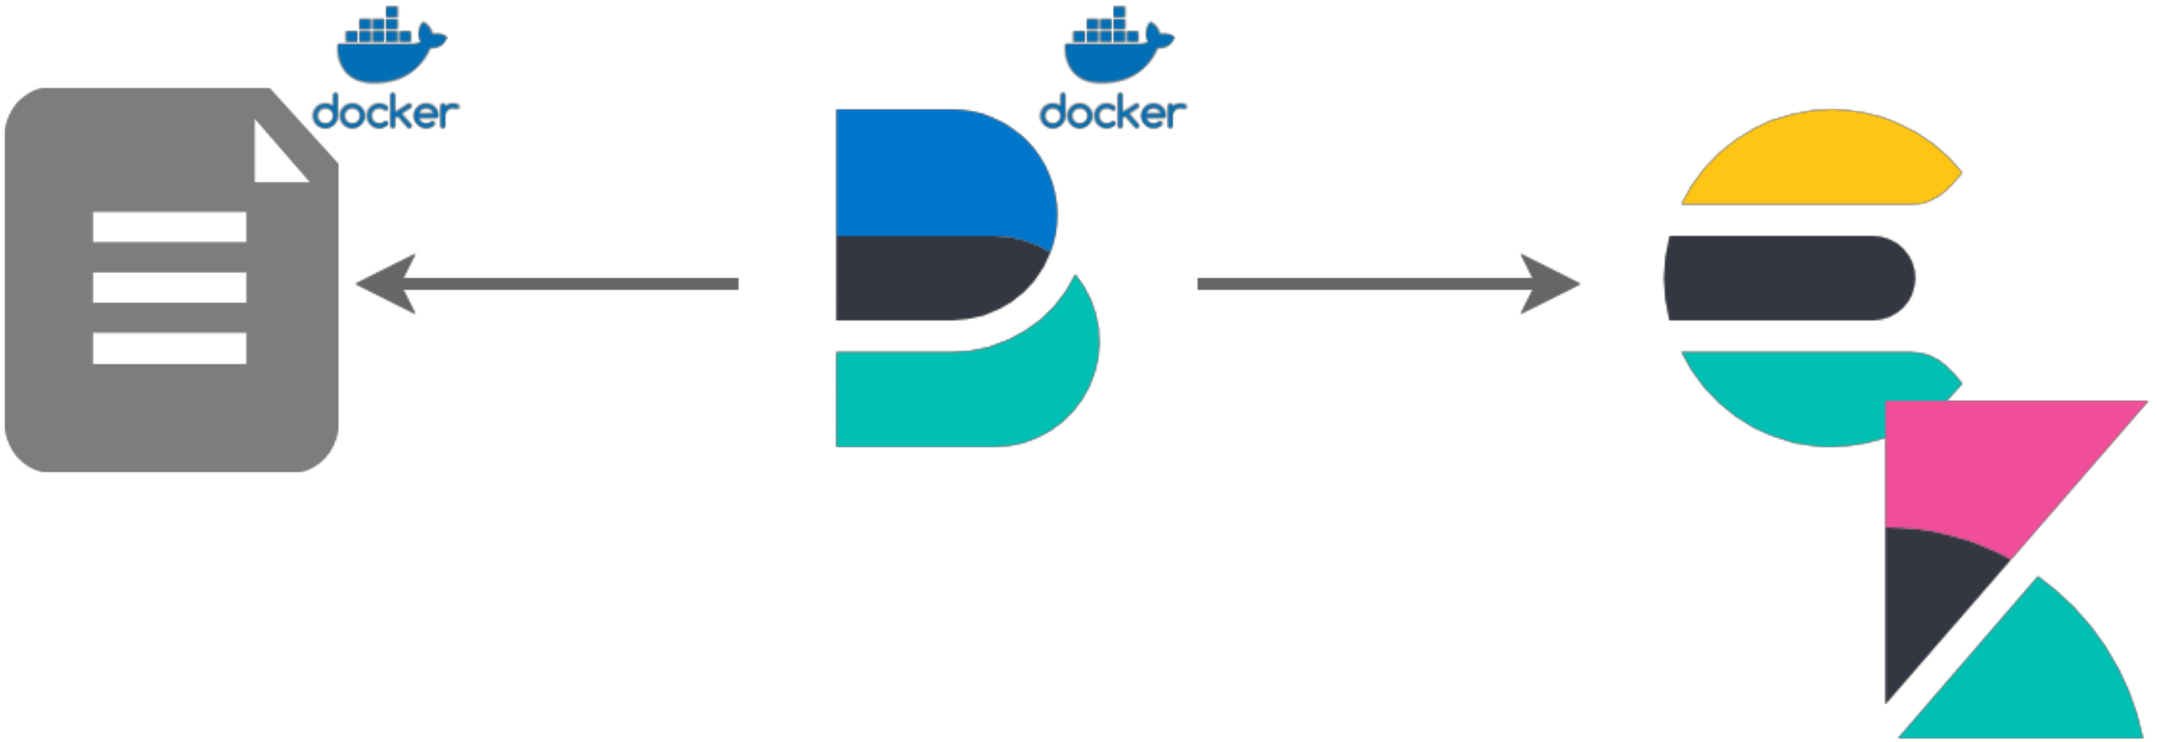
docker
docker

https://github.com/elastic/beats/tree/master/deploy/docker

Docker Logs

```
filebeat.autodiscover:
  providers:
    - type: docker
      hints.enabled: true

processors:
  - add_docker_metadata: ~
```

Metadata

No Docker metadata with the other methods


```
{
  "docker": {
    "container": {
      "image": "java-logging_java_app",
      "labels": {
        "com": {
          "docker": {
            "compose": {
              "container-number": "1",
              "project": "java-logging",
              "service": "java_app",
              "version": "1.23.2",
              "oneoff": "False",
              "config-hash": "2b38df3c73c6   1a68a37443c2006f3f3e4fc16c3c   2a1d7793f2a38841e274b607"
            }
          }
        },
        "app": "fizzbuzz"
      },
      "id": "9d6d5a7640a457a1e08c422cb0a08   f96ff3631fb5356f749b2ac7d8f3719687f"   ,
      "name": "java_app"
    }
  }
}
```

Missing the Last Line

Waiting for the newline

Hints

```
labels:
  - "app=fizzbuzz"
  - "co.elastic.logs/multiline.pattern=^\\["
  - "co.elastic.logs/multiline.negate=true"
  - "co.elastic.logs/multiline.match=after"
```

Registry File

```
filebeat.registry_file: /usr/share/filebeat/data/registry
```

Ingest Pipeline

```
output.elasticsearch:
  hosts: ["http://elasticsearch:9200"]
  index: "docker"
  pipelines:
    - pipeline: "parse_java"
      when.contains:
        docker.container.name: "java_app"
```

Ingest Pipeline

```
{
  "description" : "Parse Java log lines",
  "processors": [ {
      "grok": {
        "field": "message",
        "patterns": [ "^\\[%{TIMESTAMP_ISO8601:timestamp}\\]%{SPACE}%{LOGLEVEL:log.level}
            %{SPACE}%{USERNAME:log.package}%{SPACE}\\[%{WORD:log.method}\\]%{SPACE}-
            %{SPACE}%{GREEDYDATA:labels}%{SPACE}-%{SPACE}%{GREEDYDATA:message_rest}
            (?:\\n+(?<stacktrace>(?:.|\\r|\\n)+))?" ],
        "ignore_failure": true
      }
  } ]
}
```

Note: `\\`, `message` **vs** `message_rest`, `@timestamp` **vs** `timestamp`, `ignore_failure`

ASCII Art

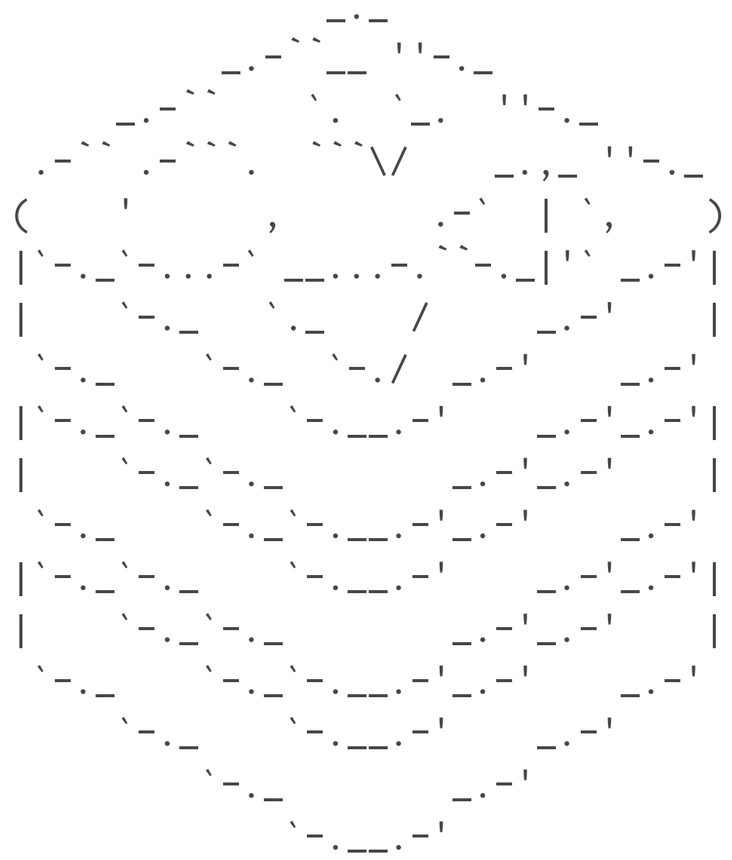

```
Redis 4.0.9 (00000000/0) 64 bit

Running in stand alone mode
Port: 6379
PID: 55757


      http://redis.io
```

Configuration Templates

```
filebeat.autodiscover:
  providers:
    - type: docker
      templates:
        - condition:
            equals:
              docker.container.image: redis
          config:
            - type: docker
              containers.ids:
                - "${data.docker.container.id}"
              exclude_lines: ["^\\s+[\\-`('.|_]"]
```

Who Logs the Logger

Avoid loops

Process without `-e`

filebeat.yml: `logging.to_files: true`

Pro: Hot 💩

Con: Complexity

Orchestrate 

kubernetes

elastic

@xeraa

Where to put Filebeat?

DaemonSet

https://github.com/elastic/beats/tree/master/deploy/kubernetes

Metadata

Either in cluster or not

```
processors:

- add_kubernetes_metadata:
    in_cluster: true

- add_kubernetes_metadata:
    in_cluster: false
    host: <hostname>
    kube_config: ${HOME}/.kube/config
```

Metadata


```
{
  "host": "172.17.0.21",
  "port": 9090,
  "kubernetes": {
    "container": {
      "id": "382184ecdb385cfd5d1f1a65f78911054c8511ae009635300ac28b4fc357ce51",
      "image": "my-java:1.0.0",
      "name": "my-java"
    },
    "labels": {
      "app": "java",
    },
    "namespace": "default",
    "node": {
      "name": "minikube"
    },
    "pod": {
      "name": "java-2657348378-k1pnh"
    }
  },
}
```

Configuration Templates

```
filebeat.autodiscover:
  providers:
    - type: kubernetes
      templates:
        - condition:
            equals:
              kubernetes.namespace: redis
          config:
            - type: docker
              containers.ids:
                - "${data.kubernetes.container.id}"
              exclude_lines: ["^\\s+[\\-`('.|_]"]
```

Customize Indices

```
output.elasticsearch:
  index: "%{[kubernetes.namespace]:filebeat}-%{[beat.version]}-%{+yyyy.MM.dd}"
```

Pro: Hot 💩💩💩

Con: Complexity++

Moar 

Index Patterns

Time based (default: daily)

Versioned

Sizing

Daily volume x retention x replication

Index Lifecycle Management

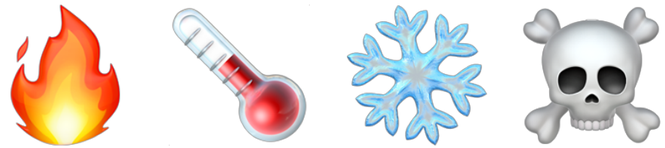

elastic @xeraa

Frozen Indices

https://www.elastic.co/guide/en/elasticsearch/reference/6.6/frozen-indices.html

Conclusion

Examples

https://github.com/xeraa/java-logging

Parse 🔪

Send ✉️

Structure 🕸

Containerize 📦

Orchestrate 🎻

Questions?

Philipp Krenn @xeraa