DATA DRIVEN EVENT MONITORING

- Fabien Wernli
- wernli@in2p3.fr
- Sysadmin

- Monitoring, out-of-band management, …

FJPPL MEETING LYON 2015-03-11

DISCLAIMER

This is not a technical talk, let's leave this up for discussion later

EVENT MANAGEMENT

- **LOGS**
  - systems
  - applications
  - appliances
  - ...

  push

  syslog-ng

- **PERFORMANCE DATA**
  - system metrics
  - application counters
  - error counters
  - ...

  pull

  collectd

SOURCES

- **WHEREABOUTS**

- ~400'000 metrics
- ~2000 nodes
- ~200 PDUs
- ~50 ACUs
- ~10 UPS'
- 1 Generator

- **THROUGHPUT**

- logs: ~2k/s (30k/s max)
- metrics: ~10k/s

~ 1 BILLION (NOISE) EVENTS PER DAY

LHC GENERATES 10^3 MORE (REAL DATA) ∴

EXAMPLES

LOGS

```
kernel: Killed process 29959, UID 42046, (hadd) total-vm:202363492kB, a
ata2.00: exception Emask 0x0 SAct 0xffff SErr 0x0 action 0x0
puppet-agent[16528]: Finished catalog run in 44.06 seconds
```

METRICS

ccsvli64.in2p3.fr	cpu/cpu-idle	ok	99.53
ccosvms0034.in2p3.fr	interface-eth0/if_octets/rx	ok	172.1
ccwntest14.in2p3.fr	df-var/percent_bytes-free	critical	0.004

USE CASES

1. Control Room / On-call
   - *Be informed in real-time*
2. Post-incident analysis
   - *Understand & Learn*

LEGACY SYSTEM

1. REAL-TIME: CCZE, SWATCH, ...

LEGACY SYSTEM

2. ANALYSIS:

SYSLOG-NG

```
# /etc/syslog-ng/syslog-ng.conf:

destination d_remote_by_host {
  file("/var/syslog-ng/remote/$YEAR/$MONTH/$DAY/by-host/${HOST}");
};
```

GREP

```
$ grep -rP 'kernel:.* ata' ./2014/09/29/by-host

./2014/09/29/by-host/ccwsge0622:2014-09-22T14:54:24+02:00 ccwsge0622 <k
./2014/09/29/by-host/ccwsge0622:2014-09-22T14:54:24+02:00 ccwsge0622 <k
./2014/09/29/by-host/ccwsge0622:2014-09-22T14:54:24+02:00 ccwsge0622 <k
```

PROBLEMS...

- grep is slow
- fs is slow
- requires shell access

- does have low footprint though, when compressed

MANY NEW(ISH) TOOLS...

- logstash
- elasticsearch
- graylog2
- syslog-ng
- rsyslog
- kafka
- redis
- rabbitmq
- ...

...WITH MANY POSSIBLE COMBINATIONS

OUR SOLUTION

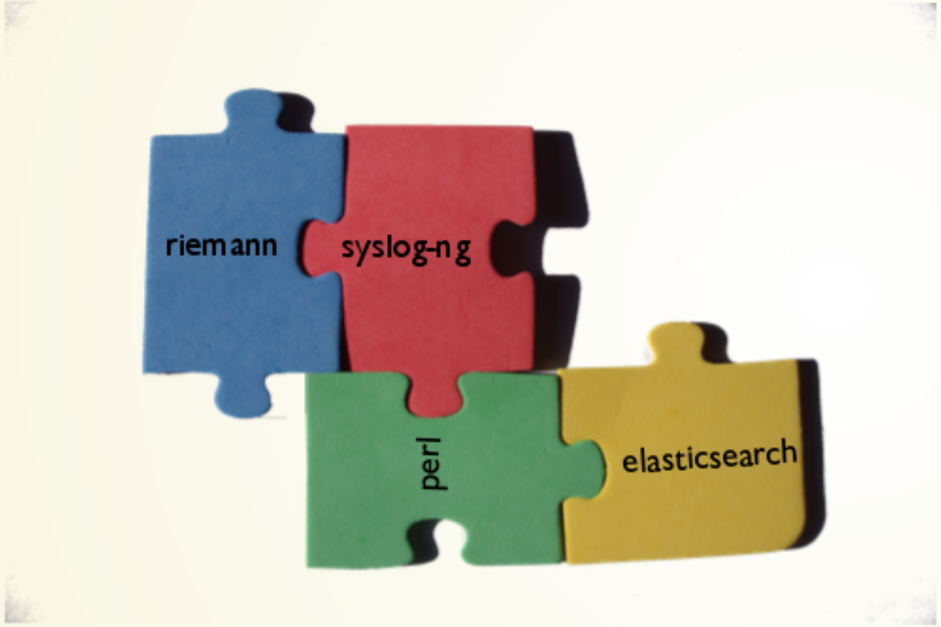

- *syslog-ng*: Collector Analyzer Router
- *Riemann*: Realtime stream processor
- *Elasticsearch*: Storage Indexing Query

1. REAL-TIME: RIEMANN

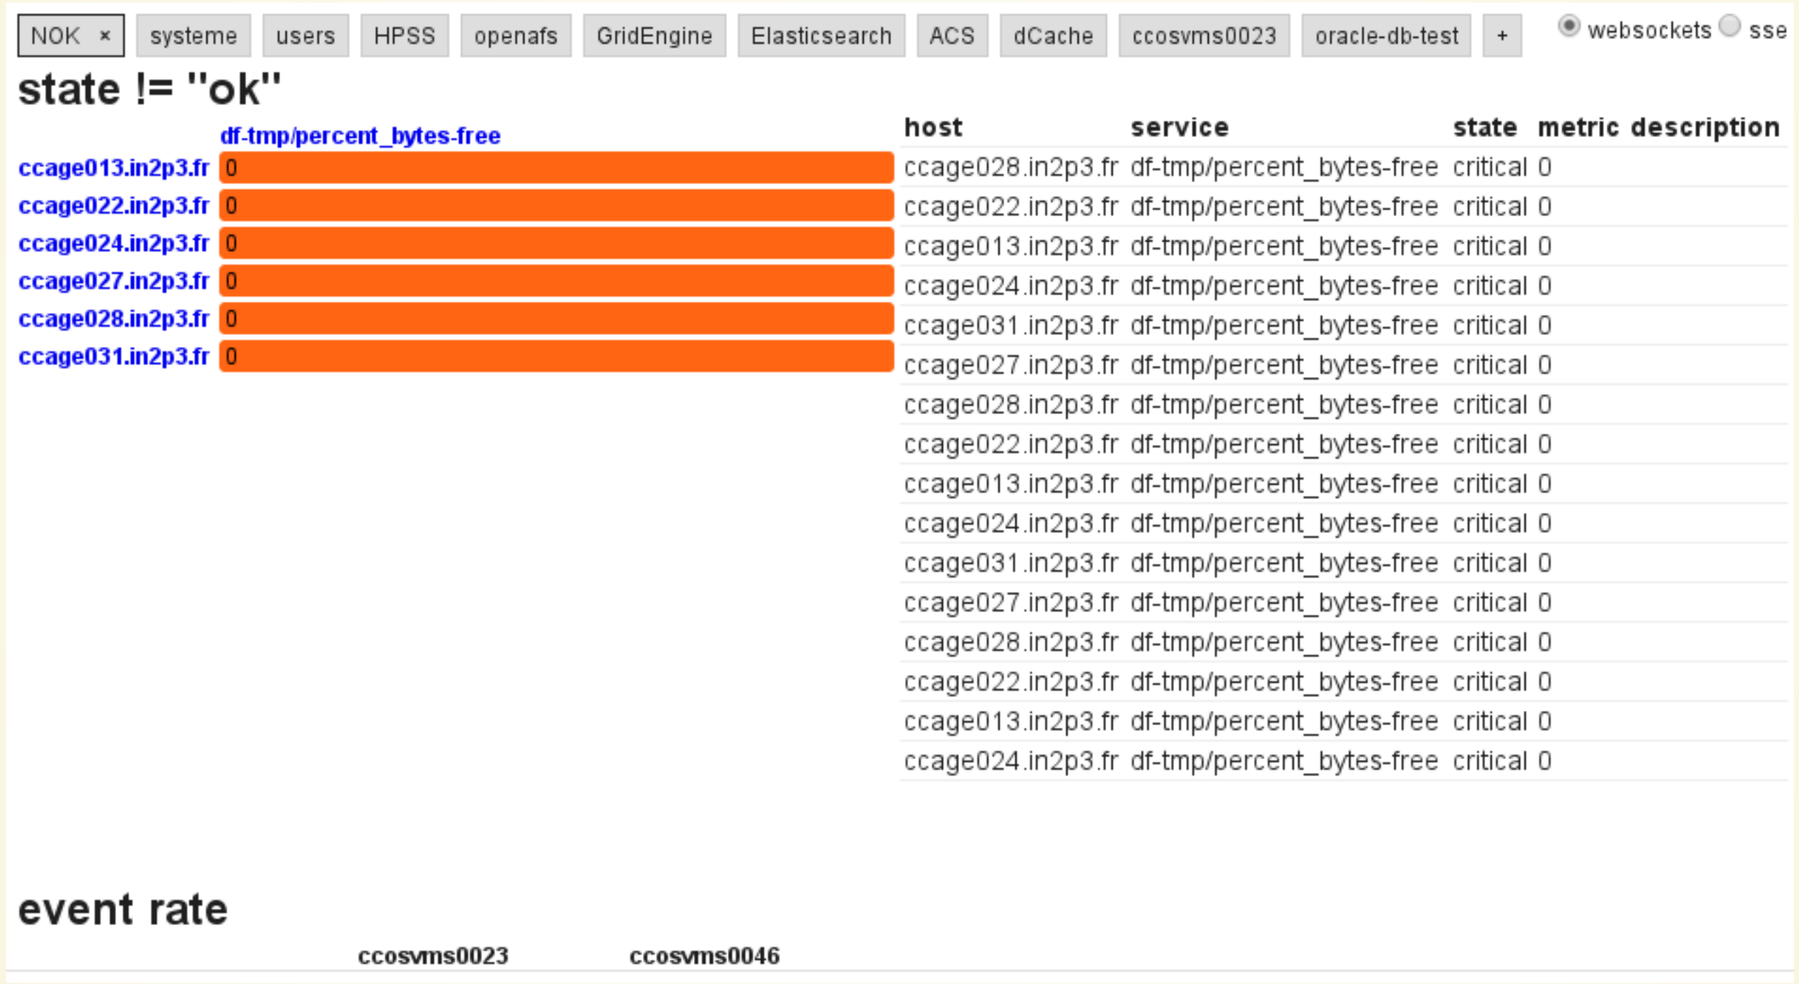

2. ANALYSIS: ELASTICSEARCH/KIBANA

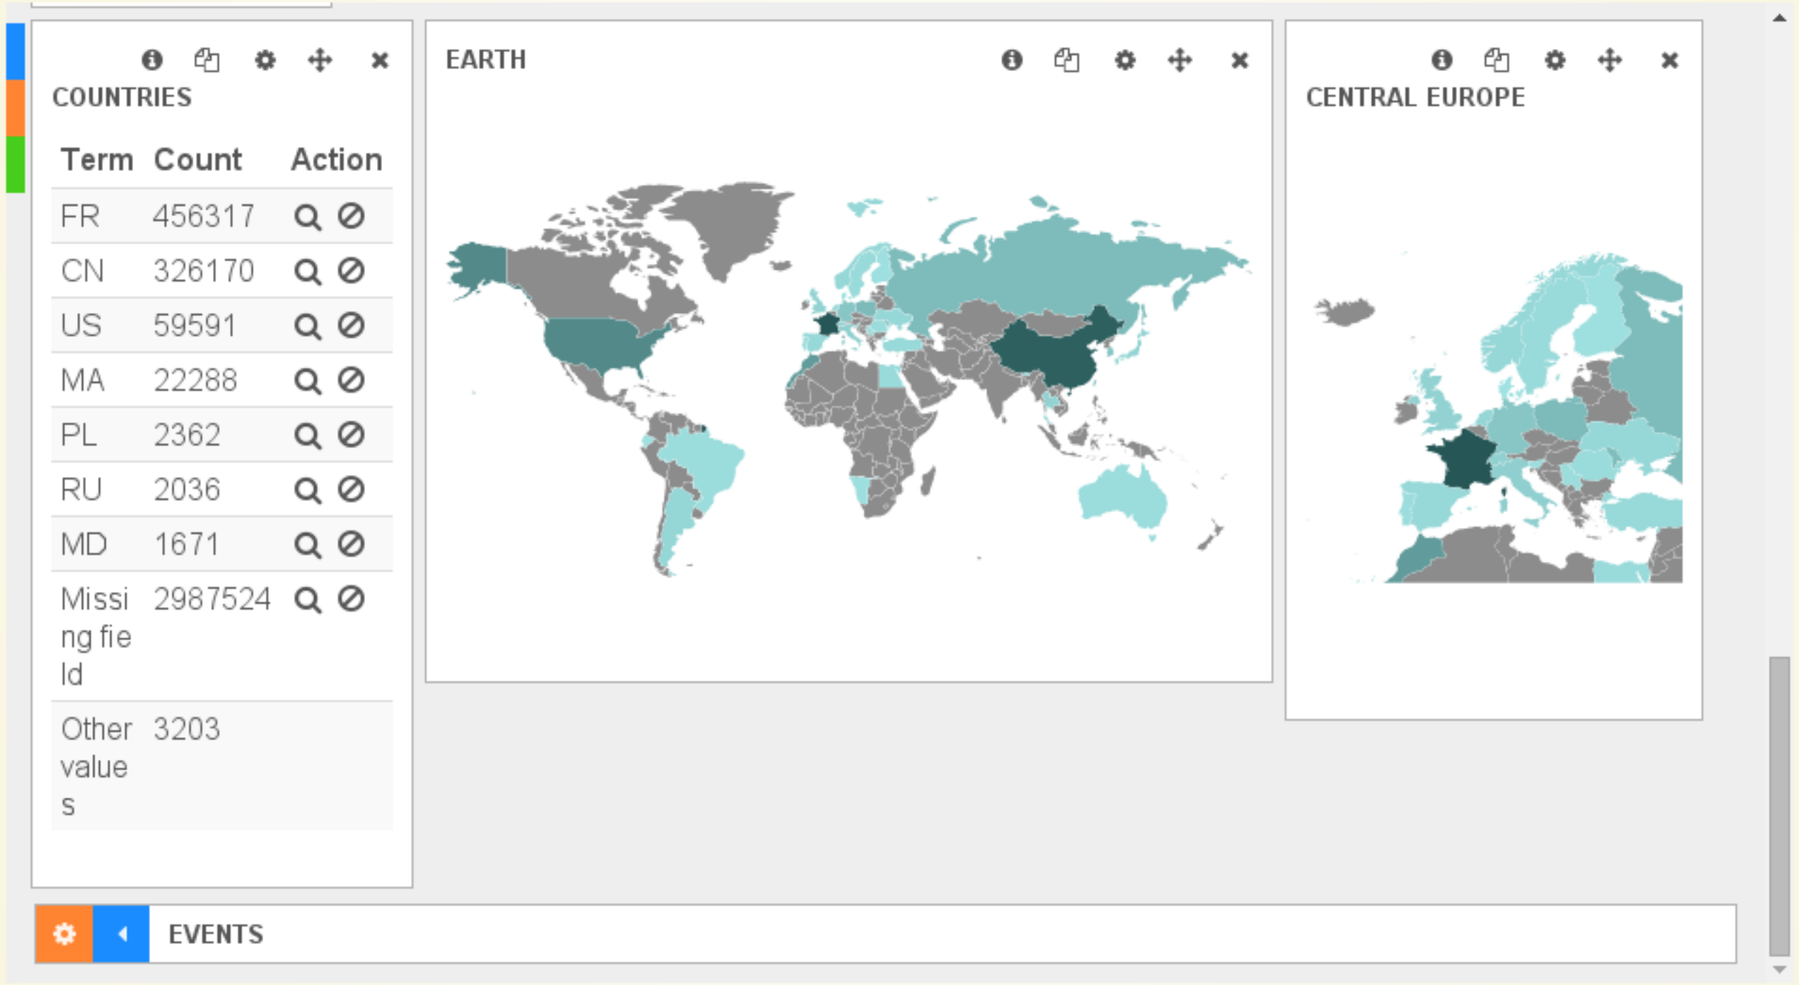

ARCHITECTURE

EVENT FLOW

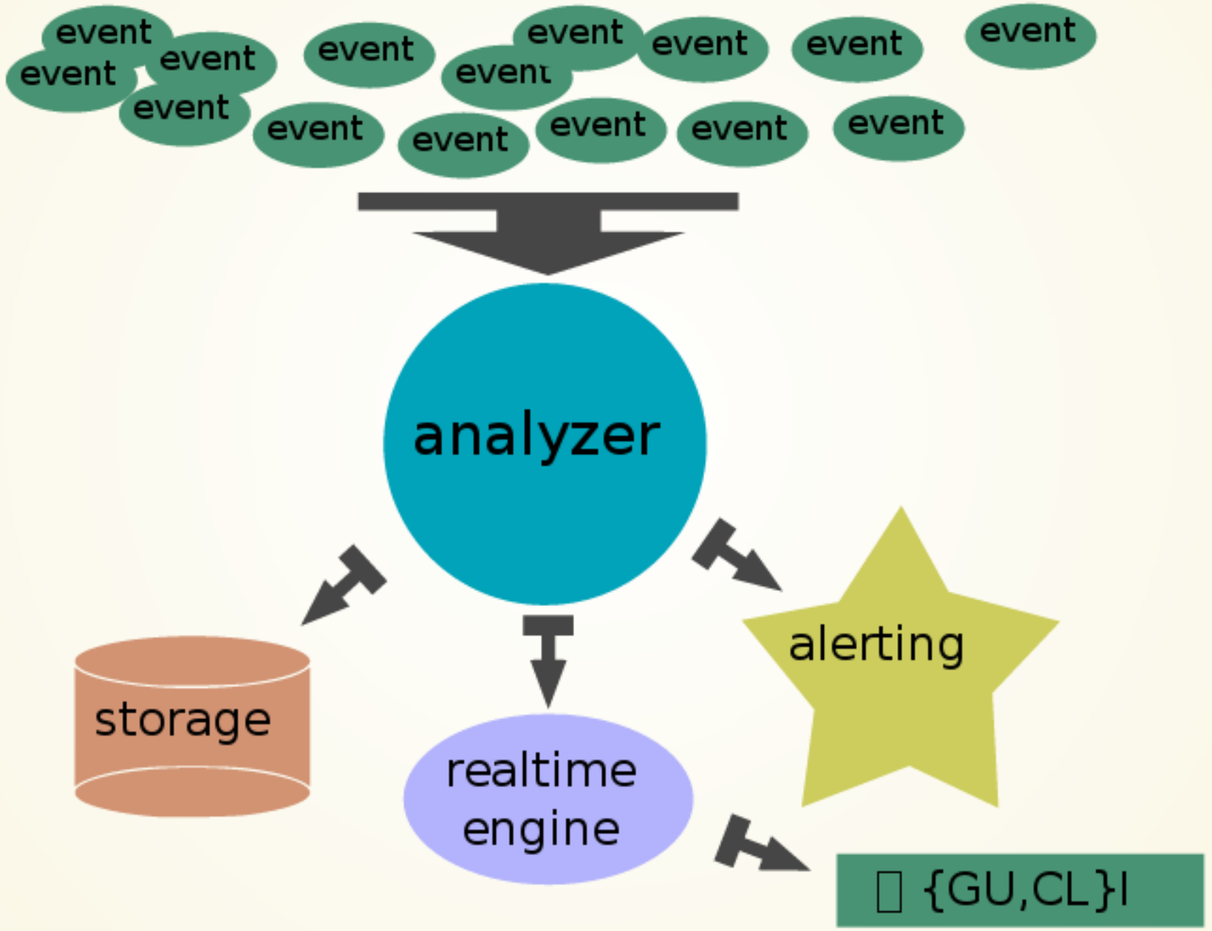
event
event
event
event
event
event
event
event
event
event
event
event
event
event
event
event
analyzer
storage
realtime
engine
alerting
{GU,CL}I

IMPLEMENTATION

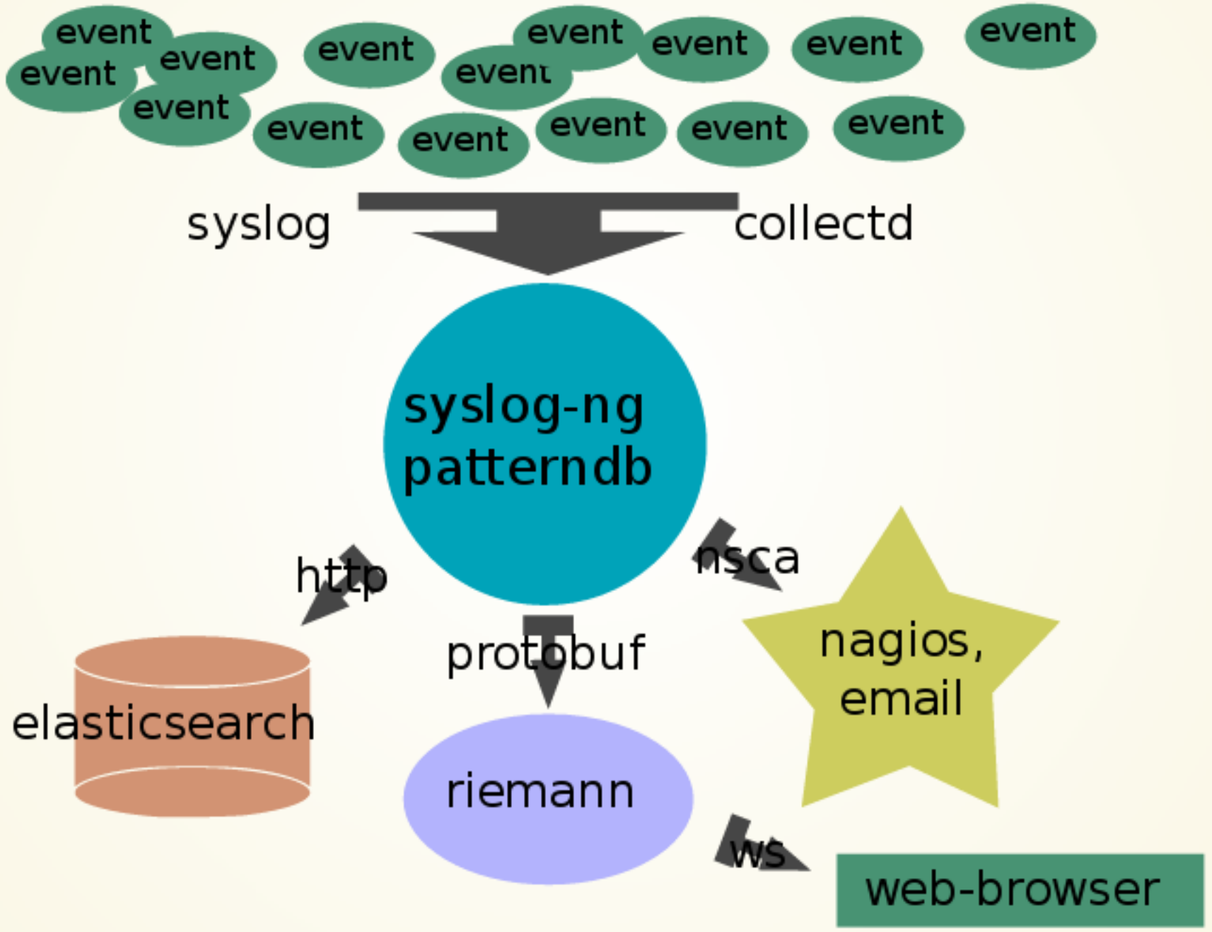

HARDWARE

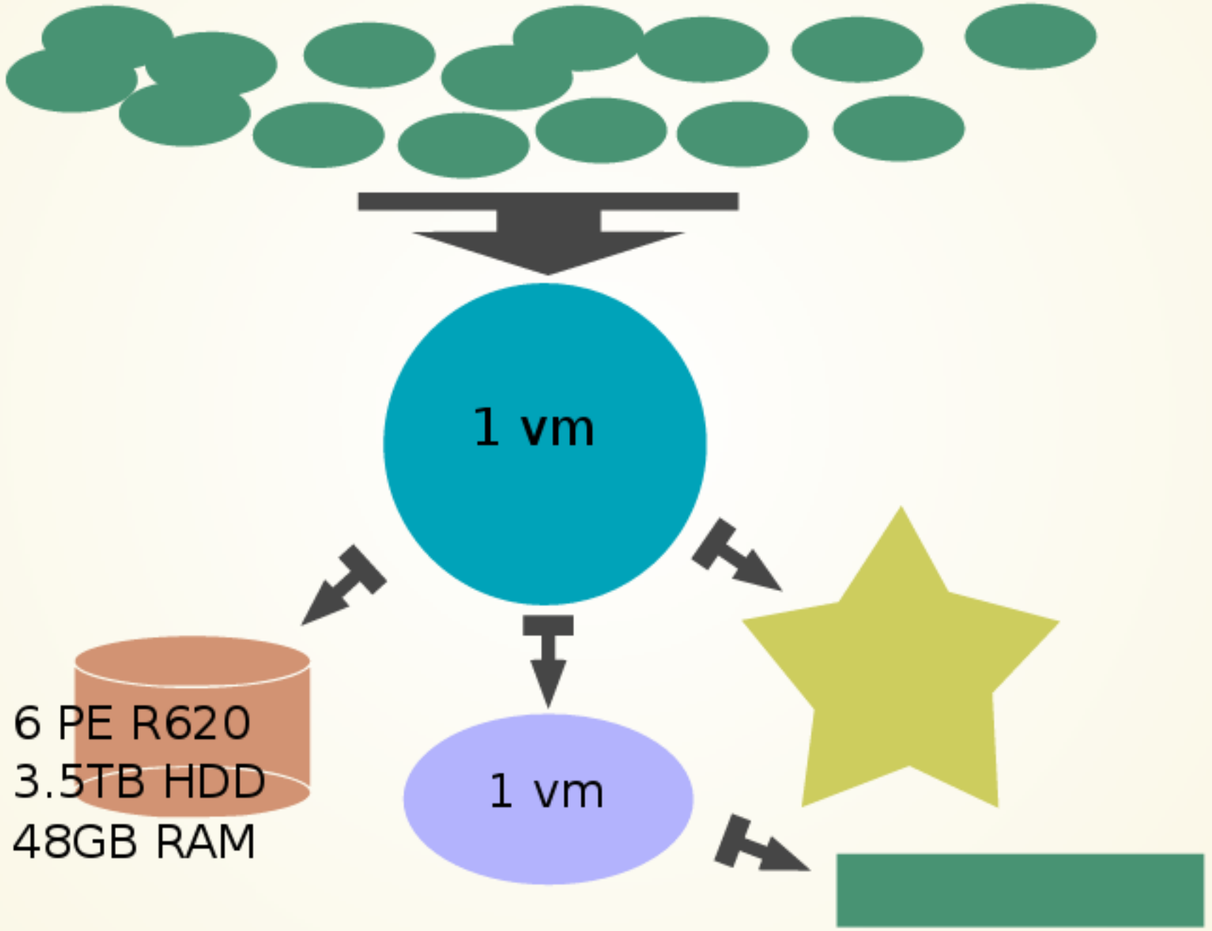

ANALYZER: *SYSLOG-NG*

- Opensource edition
- Events:
  - collect
  - tag
  - rewrite
  - route, alert, …and a lot more
- *patterndb*
  - Pattern matching engine
  - Radix tree
  - Correlation engine
- Scales extremely well

REALTIME ENGINE: *RIEMANN*

- Opensource
- Stream processor
- Synchronous
- Events:
  - collect
  - munge
  - route
  - aggregate
  - alert
  - subscribe: websockets
  - …anything really: config is a programming language

STORAGE BACKEND: *ELASTICSEARCH*

- Opensource
- Distributed
- Search Engine
- Events:
  - store
  - index
  - query
  - …"You Know, for Search"
- Interfaces:
  - Kibana
  - RESTful API

PERFORMANCE DATA: *COLLECTD*

- Opensource
- Events
  - collect
  - store
  - munge
  - forward
  - store
- Threshold alerting
- Flexible: plugin system

EVENT NORMALIZATION

DATA MODEL

- Timestamp (UTC)
  timestamp
- 2 primary keys (mandatory)
  host,service
- State (like Nagios)
  state
- Freshness (expiration)
  ttl
- KV-pairs (other stuff)
  tag: "foo", attributes: { "foo" : "bar"}

STATE

- Default

```
"ok"
```

- If considered unusual

```
("warning"|"critical")
```

Basis of alerting

SERVICE

- Default

```
<progname>
```

- Logs: after matching and correlation

```
<app>-<instance>/<type>-<instance>
```

- Metrics: *collectd-threshold*

```
<plugin>-<p_instance>/<type>-<t_instance>
```

same model for logs and metrics

FRESHNESS

- Logs: default

```
ttl = 300
```

- Logs: after matching and correlation

```
ttl = <persistence in realtime dash>
```

- Metrics:

```
ttl = <collectd-timeout> * 4
```

a fresh event sticks
an expired event vanishes

EXAMPLE

raw

```
2014-05-30T14:34:53 node01 ata2.00: excepti
Emask 0x0 SAct 0xffff SErr 0x0 action 0x0
```

normalized

```
{
  "timestamp": "2014-05-30T14:34:53",
  "host":      "node01",
  "service":   "kernel-drivers/ata-2.(
  "state":     "warning",
  "ttl":       300
  "kernel": {
    "type":   "exception",
    "emask":  "0x0",
    "sact":   "0xffff",
    "serr":   "0x0",
    "action": "0x0"
  }
}
```

PATTERNDB

EVENT CORRELATIONS

- *Syslog-ng / PatternDB*
- Triggers: timeout or conditional
- Combine multiple messages *e.g.* with same identifier
- Alert

TODO

- Use java driver for syslog-ng
- Elasticsearch & Riemann auth{entication,orization}
- Automate pattern generation: collaboration with Blue Waters?
- Collaboration with Balabit to redesign patterndb (Tibor Benke)
- Integrate with message broker
- …and some documentation

REFERENCES

http://	www.syslog-ng.org
	github.com/balabit/syslog-ng-incubator
	riemann.io
	riemann.io/dashboard.html
	en.wikipedia.org/wiki/WebSocket
	collectd.org
	elasticsearch.org
	github.com/elasticsearch/kibana
	forge.puppetlabs.com/ccin2p3/patterndb
	devops.com/features/guide-modern-monitoring-

ありがとう