TURNING EVERY ENTRY INTO GOLD

Sam Kottler

THE TOPICS

- Why do we care about logs?
- Our toolkit
- Collection patterns
- Storage
- Security (briefly)

WHAT IS A LOG ENTRY?

Details

Type	system
Date	Thursday, July 5, 2012 - 22:09
User	admin
Location	http://localhost:8080/?q=admin/modules/list/confirm
Referrer	http://localhost:8080/?q=admin/modules
Message	*dblog* module enabled.
Severity	info
Hostname	10.0.2.2
Operations	

A LOG ENTRY IS AN EVENT

DRUPAL ISN'T THE ONLY THING THAT LOGS...

- Apache2 (or php-fpm)
- Nginx
- Varnish
- MySQL
- MongoDB
- Redis

LOGS ARE PRODUCED BY MOST (ALL) SOFTWARE

```
alternatives.log  lastlog
apache2           mail.err
apt               mail.log
auth.log          mysql
boot              mysql.err
boot.log          mysql.log
btmp              news
dist-upgrade      syslog
dmesg             udev
dmesg.0           ufw.log
dmesg.1.gz        unattended-upgrades
dpkg.log          vboxadd-install.log
faillog           vboxadd-install-x11.log
fsck              VBoxGuestAdditions.log
installer         wtmp
kern.log
```

LOG MESSAGES ARE A LENS INTO WHAT IS HAPPENING ACROSS ALL SYSTEMS

1. LOGS ARE UBIQUITOUS

2. LOGS ARE ACCESSIBLE

3. THEY ARE EASY TO READ

3.5 (HOPEFULLY)

THE AGGREGATE OF OUR LOG MESSAGES TELL THE STORY OF WHAT'S HAPPENING

THE TOOLS

- syslog.module
- rsyslogd
- Logstash
- Kibana
- ElasticSearch

DBLOG VS. SYSLOG

/ETC/RSYSLOG.CONF

```
local0.* /var/log/drupal.log
```

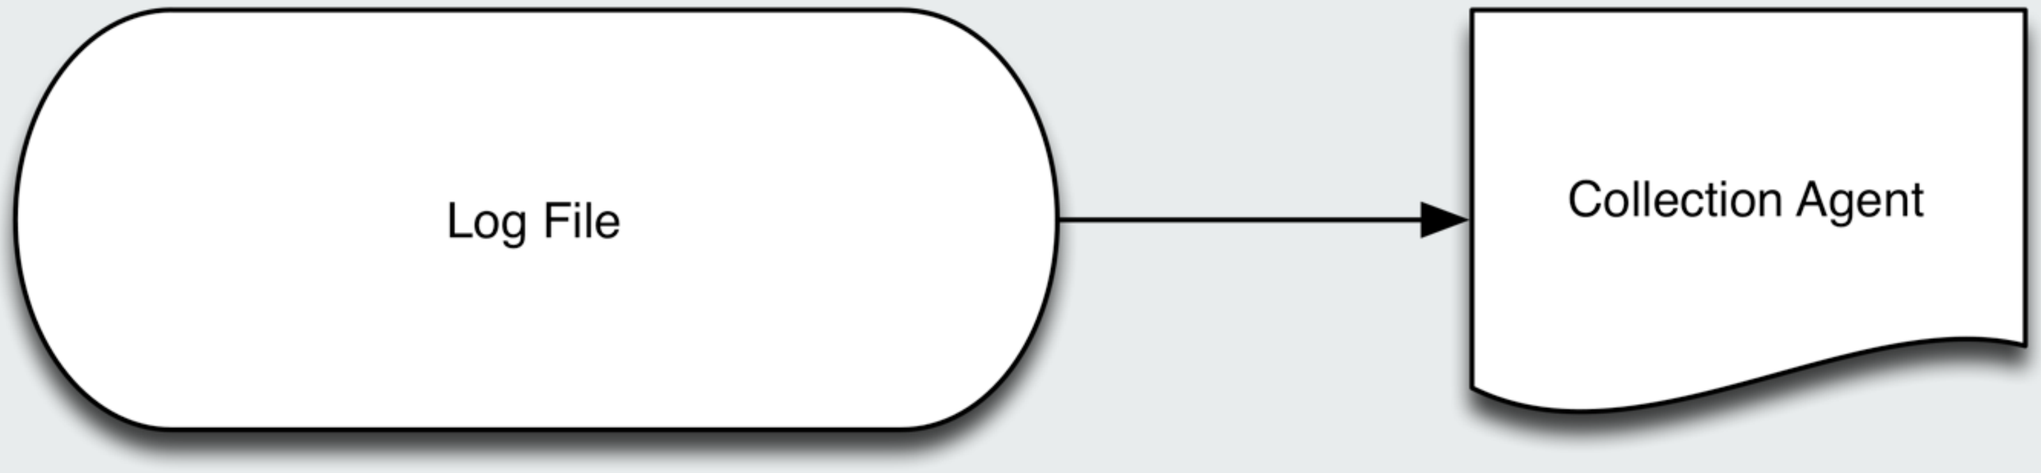

/var/log/drupal.log

rsyslogd

/etc/rsyslog.d/60-apache2.conf

```
$ModLoad imfile
$ModLoad imklog
$ModLoad imuxsock

$InputFileName /var/log/drupal.log
$InputFileTag drupal
$InputFileStateFile state-drupal
$InputRunFileMonitor
```

```
$ModLoad imfile
```
← File handler module

```
$ModLoad imklog
```
← Log handler module

```
$ModLoad imuxsock
```
← Socket module

```
$InputFileName /var/log/drupal.log

$InputFileTag drupal

$InputFileStateFile state-drupal

$InputRunFileMonitor
```

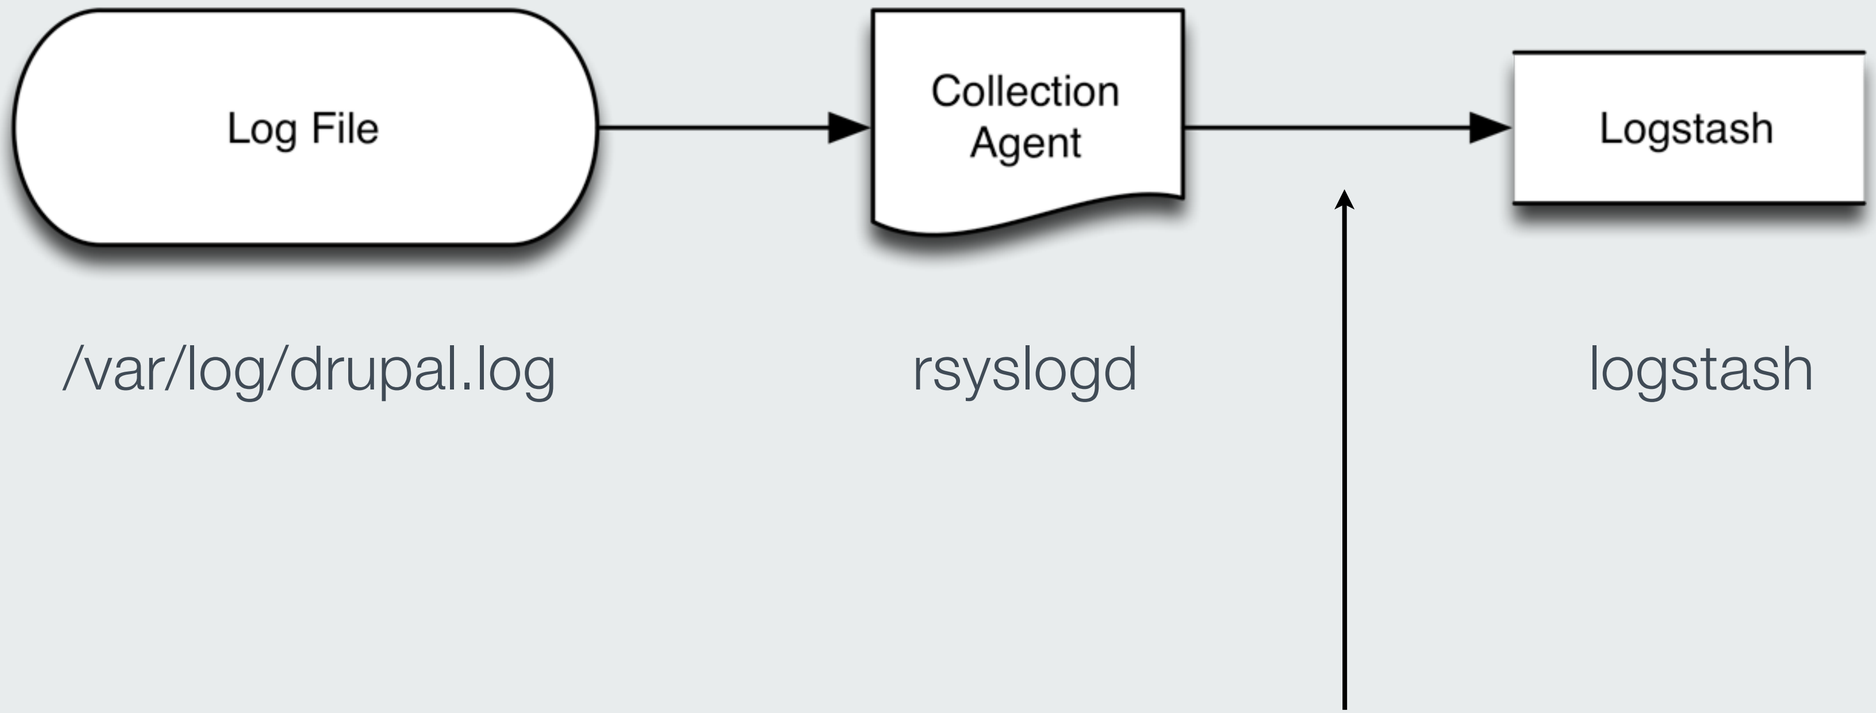
Log File
Collection Agent
Logstash
/var/log/drupal.log
rsyslogd
logstash
UDP/TCP connection

/etc/rsyslog.d/60-drupal.conf

```
$ModLoad imfile
$ModLoad imklog
$ModLoad imuxsock

$InputFileName /var/log/drupal.log
$InputFileTag drupal
$InputFileStateFile state-drupal
$InputRunFileMonitor

*.* @<%= @rsyslog_host %>:<%= @rsyslog_port %>
```

```
*.* @<%= @rsyslog_host %>:<%= @rsyslog_port %>
```

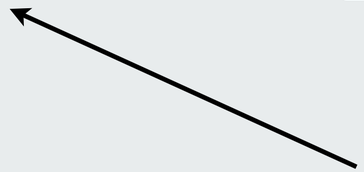

Send logs via UDP to <host>:<port>

```
*.* @@<%= @rsyslog_host %>:<%= @rsyslog_port %>
```

Send logs via TCP to <host>:<port>

WORD OF WARNING

- Setup logrotate (please!)
- Rsyslogd has no sense of where it started
- All entries get resent when you restart rsyslogd

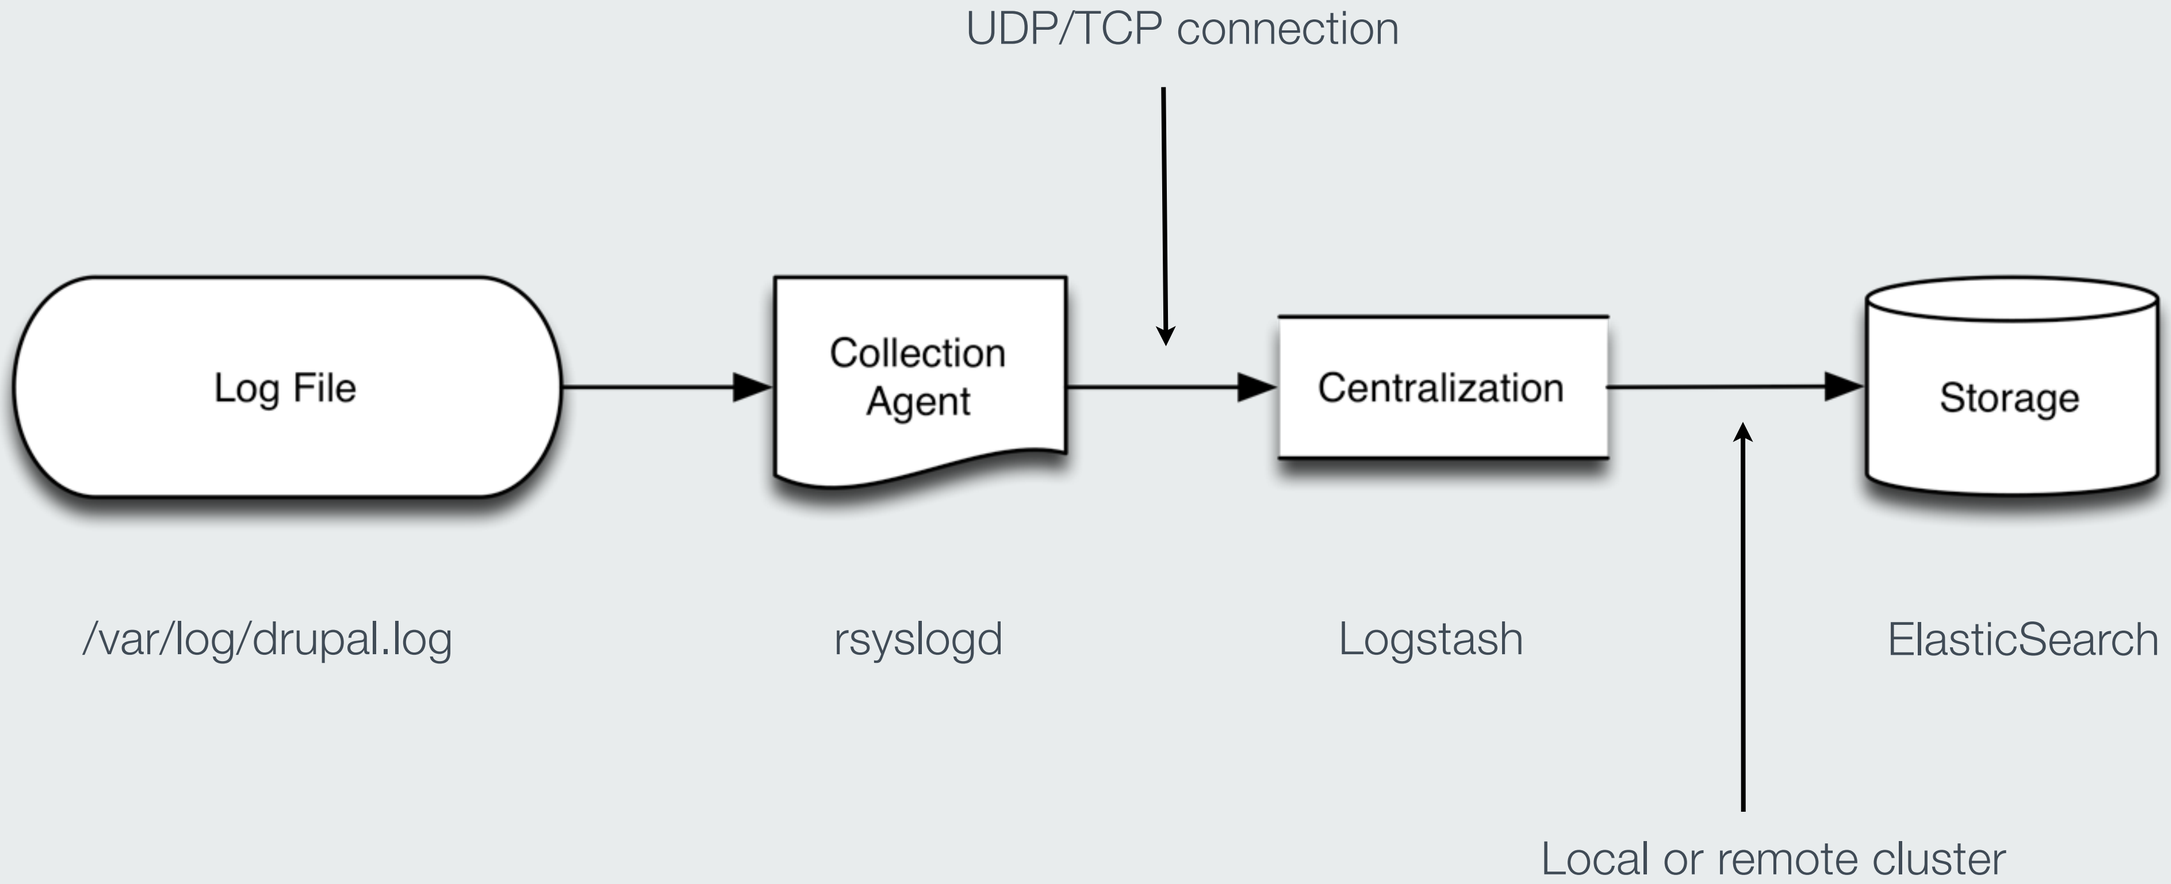
UDP/TCP connection
Log File
Collection Agent
Centralization
Storage
/var/log/drupal.log
rsyslogd
Logstash
ElasticSearch
Local or remote cluster

SECURITY

- Rsyslog sends logs in plain text
- Rsyslogd versions 4+ have (working) TLS support
- bit.ly/SgEXaa

DO I REALLY NEED TLS?


Crying kitten
alert

DON'T MAKE KITTENS CRY!

LOGSTASH

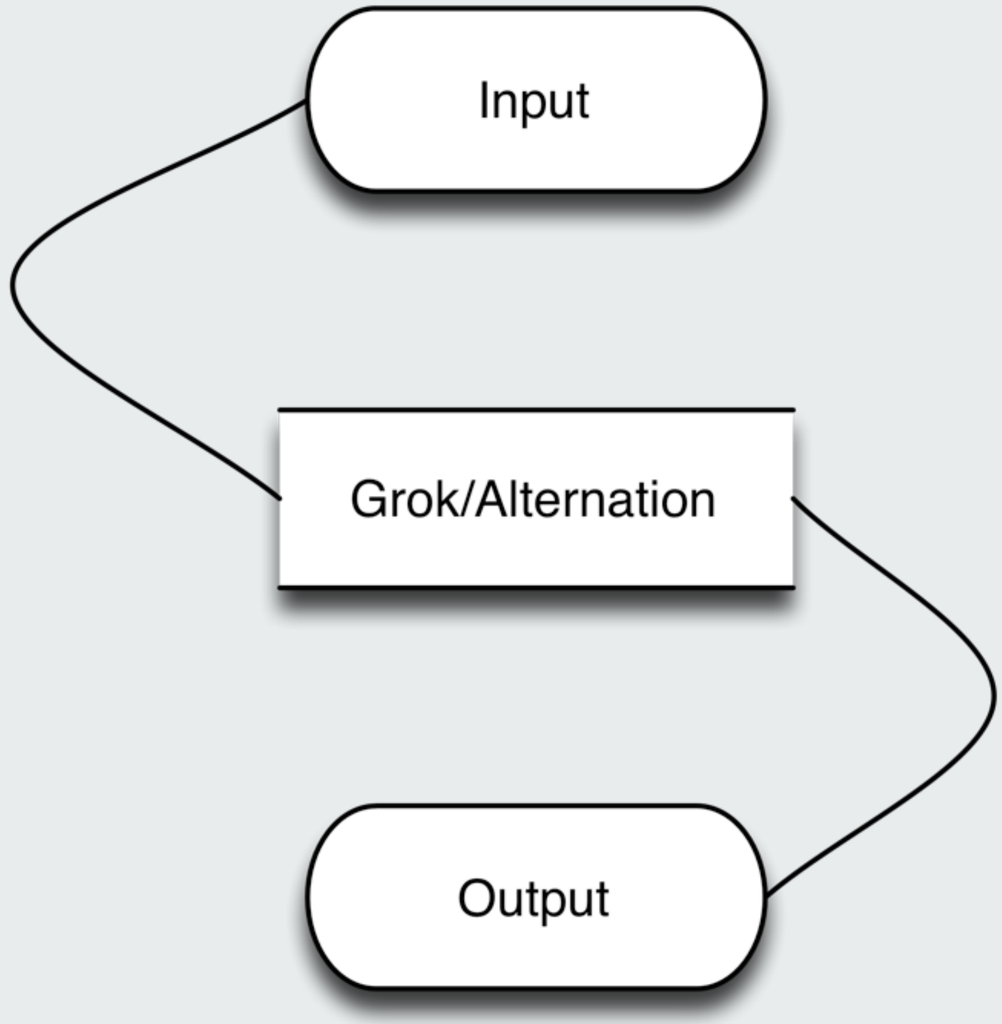

INPUT

- exec
- file
- gelf
- stdin
- syslog
- tcp
- zmq

```
input {
  syslog {
    port => 5544
    type => "remote-syslog"
  }
  stdin {}
}
```

2. root@(none):~ (ssh)

```
[root@(none) ~]#
```

FILTERS

- date
- dns
- gelf
- grep
- grok
- json
- split

```
grok {
  type => "apache"
  pattern => "%{COMBINEDAPACHELOG}"
}

date {
  type => "apache"
  timestamp => "dd/MMM/yyyy:HH:mm:ss Z"
}
```

OUTPUT

- elasticsearch
- file
- GELF
- graphite
- mongodb
- nagios
- tcp
- zmq

```
output {
  elasticsearch {
    host => "localhost"
    port => 9309
  }
  stdout {}
}
```

2. root@(none):~ (ssh)

```
[root@(none) ~]#
```

2. root@(none):/opt/elasticsearch (ssh)

```
[root@(none) elasticsearch]#
```

2. root@(none):~ (ssh)

```
[root@(none) ~]#
```

KIBANA

- Bootstrap-based web interface
- Connects directly to the ES cluster
- Formats logs for graphing and other purposes
- 689 lines of PHP


Kibana
localhost:8091

Kibana

localhost:8091/#eyJzZWFyY2giOiliLCJmaWVsZHMiOltdLCJvZmZzZXQiOjAsInRpbWVmcmFtZSI6IjE1IG1pbnV0ZXMiLCJncmFwaG1vZGUiOiJjb3VudCJ9

Last 15m | Search | Search | Reset | Hits 13,879 | Indexed 75,760,473

Show Fields

@message @source @source_host @source_path @tags @timestamp @type device message timestamp

Analyze Field

@message
@source
@source_host
@source_path
@tags
@timestamp
@type
device
message
timestamp

rss export stream

2012-08-21T08:44:14 to 2012-08-21T08:59:14

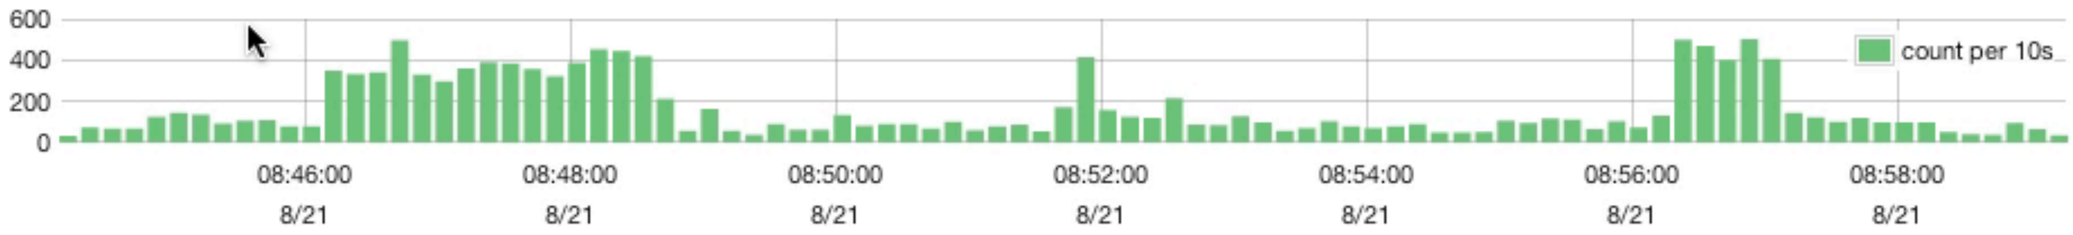

0 TO 50 Next

Time	@message
08/21 12:59:14	Aug 21 12:59:13 www2.drupal.org drupal: http://drupal.org|1345553953|search|127.0.0.1|http://drupal.org/search/apachesolr_multisitesearch/redhen|http://drupal.org/|0|re
08/21 12:59:14	Aug 21 12:59:13 www5.drupal.org drupal: http://drupal.org|1345553953|page not found|180.76.5.156|http://drupal.org/node/1732470|0||node/1732470
08/21 12:59:13	Aug 21 12:59:13 www1.drupal.org drupal: http://drupal.org|1345553953|page not found|14.140.152.18|http://drupal.org/files/js/js_46ce20a122f37f4e8507f080ae2aa5ca.js|http url=http%3A%2F%2Fdrupal.org%2Fproject%2Fmoodle|2261118||files/js/js_46ce20a122f37f4e8507f080ae2aa5ca.js
08/21 12:59:13	Aug 21 12:59:12 www1.drupal.org drupal: http://drupal.org|1345553952|page not found|188.24.124.10|http://drupal.org/aggregator/node/2008/10/24/node/node/www.soc page=2||0||aggregator/node/2008/10/24/node/node/www.socialmonkee.com/get-feed.php
08/21 12:59:13	Aug 21 12:59:13 www7.drupal.org postfix/qmgr[3747]: 30360110049: from=<bounces@drupal.org>, size=1150, nrcpt=1 (queue active)
08/21 12:59:13	Aug 21 12:59:13 www7.drupal.org postfix/smtp[16777]: 30360110049: to=<christina.hustedde@adu.edu>, relay=smtp4.osuosl.org[140.211.166.137]:25, delay=0.04, delay
08/21 12:59:13	Aug 21 12:59:13 www7.drupal.org postfix/qmgr[3747]: 30360110049: removed
08/21 12:59:13	Aug 21 12:59:13 www7.drupal.org postfix/cleanup[15643]: 30360110049: message-id=<20120821125913.30360110049@www7.drupal.org>
08/21 12:59:13	Aug 21 12:59:13 www6.drupal.org drupal: http://drupal.org|1345553953|page not found|14.140.152.18|http://drupal.org/files/css/css_74acbc51d541a9dfb04fc792720c383 url=http%3A%2F%2Fdrupal.org%2Fproject%2Fgradebook|2261118||files/css/css_74acbc51d541a9dfb04fc792720c3837.css
08/21 12:59:13	Aug 21 12:59:13 www7.drupal.org postfix/pickup[16137]: 30360110049: uid=48 from=<bounces@drupal.org>
08/21 12:59:13	Aug 21 12:59:12 www6.drupal.org drupal: http://drupal.org|1345553952|search|188.25.40.41|http://drupal.org/search/apachesolr_search/cck%20video?filters=ss_meta_t videofilters=ss_meta_type:module (Search).

DEMO TIME!

OTHER TOOLS

- Graylog2
- Flume
- Splunk (proprietary)

QUESTIONS?

QUESTIONS?

skottler@redhat.com

github: skottler

@samkottler