



**loggly**

Using MongoDB for Logs  
05/24/2011:15:30:45PST

Kord Campbell  
@kordless



Loggly is a highly scalable hosted search engine for logs.



Loggly is a highly scalable hosted search engine for logs.



Home

Games

Jobs

Labs

Contact



226 people like this.

## News

- [Dragon Age Legends](#) has officially launched into open beta! Be legendary!
- Kotaku has a new [preview of Fancy Pants Adventures!](#)
- EA2D featured in Joystiq's "[What's In a Name?](#)"
- EA2D is [Going Browser and Beyond!](#)
- We've released [AWOL](#), an [EA2D Labs](#) project. See how far you can run and compete against your friends!

## Mission Statement

EA2D is a studio within Electronic Arts that develops high quality social games for gamers, focused on delivering connected experiences in the browser, on mobile and beyond. EA2D is driven by innovation, integrity and collaboration.

Browser based Dragon Age Journeys sees Electronic Arts returning to its roots. It could change the way you feel about browser games.

GamePro



Our team and partners define talent, passion and creativity. We are building long term value and doing it with integrity.

Mark Spenner VP & GM, EA2D



# about.me™

## A CUSTOM PROFILE & PERSONAL ANALYTICS DASHBOARD

Sign up to quickly build your personal profile page that points users to your content from around the web.

ENTER AN EMAIL ADDRESS

CREATE A PASSWORD

PICK AN ABOUT.ME URL

[Create Your Profile](#)

about.me  
It's all about you.

01:13



Loggly | Shell

https://geekceo.loggly.com/shell/#/\_graph 404

loggly

From: NOW-1HOUR    Until: NOW    Inputs: All

```

3.0.30729; .NET4.0C)" RG4AxX8AAAEAAAxAUaUEAAAAT 404
2011-04-26 14:48:51.076 loggly_web 174.129.233.147
207.46.204.180 - - [26/Apr/2011:14:48:50 -0700] "GET /robots.txt HT
" "Mozilla/5.0 (compatible; bingbot/2.0; +http://www.bing.com/bingb
Results from range 2011-04-26T21:20:04.463Z - 2011-04-26T21:48:51.076Z for term
kordless@geekceo> graph 404

```

Events from NOW-1HOUR until NOW

kordless@geekceo>

Loggly | Main Dashboard

https://geekceo.loggly.com

profile | help | logout

kordless@geekceo> search for broken things

dashboard inputs users account monitor

### Log History

73,608 events

### Inputs

Name	Events In Last Day
default	5,837
gecko_http	0
geekceo_http	0
kedge_http	6
logg_ly	51,038
loggly_web	15,212
port514	0
titus	1,571

### Devices

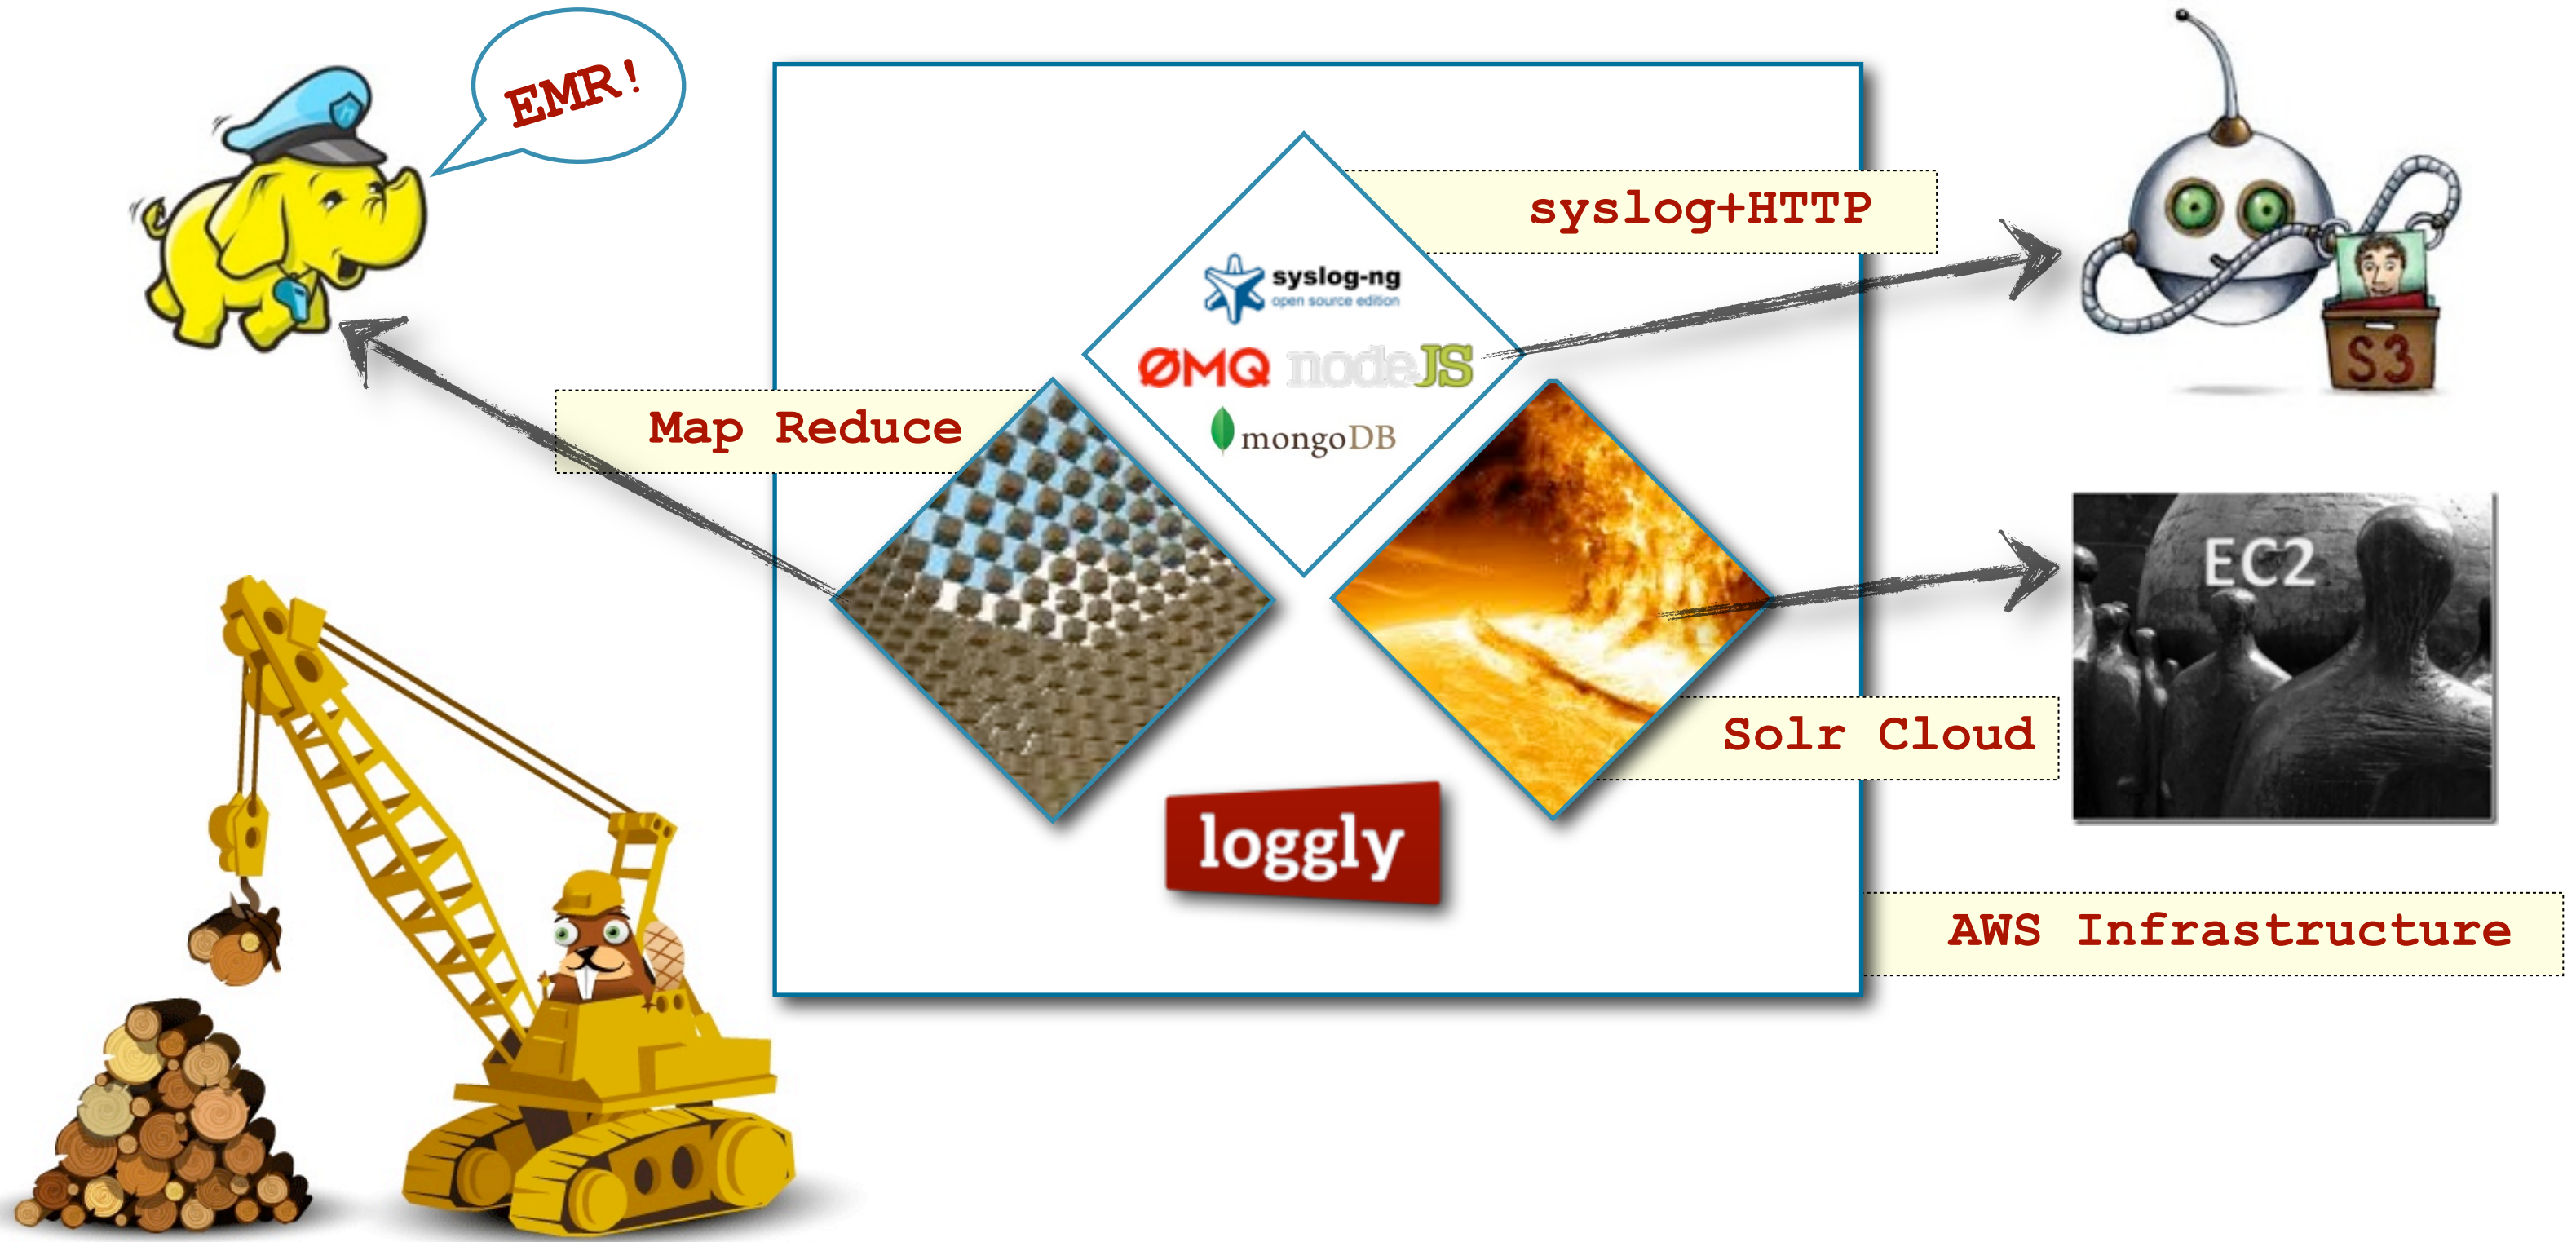
Name/IP	Events In Last Day
72.14.194.33	34,175
72.14.194.17	16,863
Loggly's web server	15,212
Loggly Office	5,837
98.210.51.152	1,577

### Daily Usage

Today's usage: 11 MB

Legend: Dropped (red), Indexed (blue)

# The Loggly Stack



# Loggly's Infrastructure

- Amazon's [AWS](#) for infrastructure
- [Syslog-NG](#) for syslog/TLS input services
- [Node.js](#) for HTTP/HTTPS input services
- [OMQ](#) for event queuing & work distribution
- [S3 bucket](#) writer for local store and archives
- [MongoDB](#) for statistics and API /stat methods
- [Solr Cloud](#) for scalable search and facets
- [Django/Python](#) for middleware/app





Loggly | Main Dashboard <https://geekceo.loggly.com> profile | help | logout

kordless@geekceo> search

dashboard inputs users account monitor

Powered by  
**mongoDB**

### Log History

33,870 events

Time	Events
12:00	500
14:00	500
16:00	500
18:00	500
20:00	500
22:00	500
00:00	1,000
02:00	500
04:00	500
06:00	500
08:00	1,500
09:00	5,800
10:00	2,800
12:00	1,000
14:00	1,000

### Inputs

Name	Events In Last Day
castor_loggly	32
default	1,729
gecko_http	0
geekceo_http	0
kedge_http	0
logg_ly	11,072
loggly_web	18,432
port514	0

### Devices

Name/IP	Events In Last Day
Loggly's web server	18,432
74.125.75.20	4,100
64.233.172.17	3,190
98.210.51.152	2,088
Loggly Office	1,756
64.233.172.6	1,614
74.125.64.81	75
74.125.64.83	73

Loggly | Main Dashboard <https://geekceo.loggly.com> profile | help | logout

kordless@geekceo> search

dashboard inputs users account monitor

Powered by  
 Apache  
**Solr**

33,870 events

### Log History

Time	Events
12:00	500
14:00	500
16:00	500
18:00	500
20:00	500
22:00	500
00:00	1,000
02:00	500
04:00	500
06:00	500
08:00	1,500
09:00	5,800
10:00	2,800
12:00	1,000
14:00	1,000

### Inputs

Name	Events In Last Day
castor_loggly	32
default	1,729
gecko_http	0
geekceo_http	0
kedge_http	0
logg_ly	11,072
loggly_web	18,432
port514	0

### Devices

Name/IP	Events In Last Day
Loggly's web server	18,432
74.125.75.20	4,100
64.233.172.17	3,190
98.210.51.152	2,088
Loggly Office	1,756
64.233.172.6	1,614
74.125.64.81	75
74.125.64.83	73

# MongoDB for Statistics

- switched because it scales well & easy to do upserts
- runs all of Loggly's primary stats storage
- use master-slave replication
- java bindings + jetty to serve aggregated data
  - use the same parsing methods for dates as solr
  - presents data the same as solr does to the middleware



mongoDB

# MongoDB for Statistics

- stats job watches OMQ for events
- as things fly by, stuff gets counted
- by ip, destination(port), time
- stuff that into mongoDB
- rollups occur every hour via upserts
- can rebuild the entire stats store in about an hour



mongoDB

# MongoDB for Statistics

- MongoDB is GREAT for logging statistics

```
{  
  u'endtime': 1299948299,  
  u'name': u's3countbyip',  
  u'value': 1,  
  u'width': 60,  
  u'cust_id': u'1000',  
  u'split': u'229.94.176.132',  
  u'starttime': 1299948240,  
  u'_id': ObjectId('4d7ba317635d28d2cdb4f337')  
}
```



mongoDB

# More Logging w/ MongoDB

- MongoDB is PERFECT for logging stuff if it's semi-structured


```
{  
  "version": "1.0",  
  "host": "webhead2",  
  "short_message": "Short message",  
  "full_message": "Backtrace here\n\nmore stuff",  
  "timestamp": 1291899928,  
  "level": 1,  
  "facility": "payment-backend",  
  "file": "/var/www/somefile.rb",  
  "line": 356,  
  "_user_id": 42,  
  "_something_else": "foo"  
}
```



mongoDB

# More Logging w/ MongoDB

- Voxify's Robert Stewart @ <http://logg.ly/tpM>
- Wrote log4mongo-java parked on Github



mongoDB

# More Logging w/ MongoDB


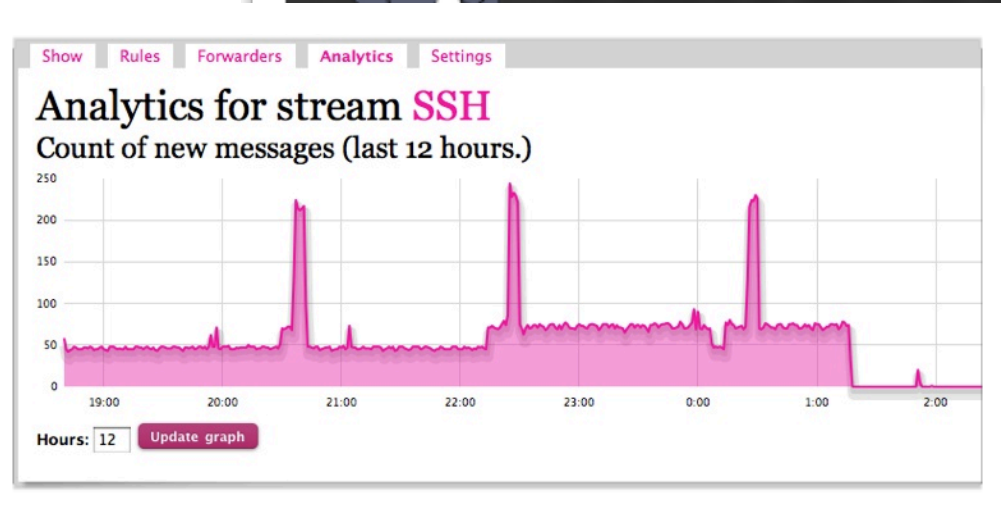
- logstash by Jordan Sissel (Loggly's devops extraordinaire)
- <http://logstash.net/>
- uses MongoDB for output channel





# More Logging w/ MongoDB

- graylog2 by Lennart Koopmann (@\_lennart)
- <http://graylog2.org/>
- uses MongoDB for storage and search (with Elasticsearch)



# Visualization Stuff

- Highcharts @ <http://highcharts.com/>
- Protoviz @ <http://vis.stanford.edu/protovis/>
- Smoothie Charts @ <http://smoothiecharts.org/>
- Google Earth plugin @ <http://code.google.com/apis/earth/>
- Google Charts @ <http://code.google.com/apis/chart/>



# Free Stuff

- Free accounts at <http://loggly.com/signup>
- Free stickers
- Free shirts
- Free lunch
- Free advice



# loggly

Follow me  
[@loggly](#) on  
Twitter!



If you like beavers, you should  
work here! <http://loggly.ly/jobs>