Bootstrapping Solr search clusters and maintain them using Puppet

All you ever wanted to know about maintaining Solr for a search-critical Drupal application

Who Are We

- **Nick Veenhof**
  Nick_vh, http://drupal.org/user/122682
- **Peter Wolanin**
  pwolanin, http://drupal.org/user/49851

- Acquia Search - hosted Apache Solr in the cloud since 2009 has given us experience.

Acquia

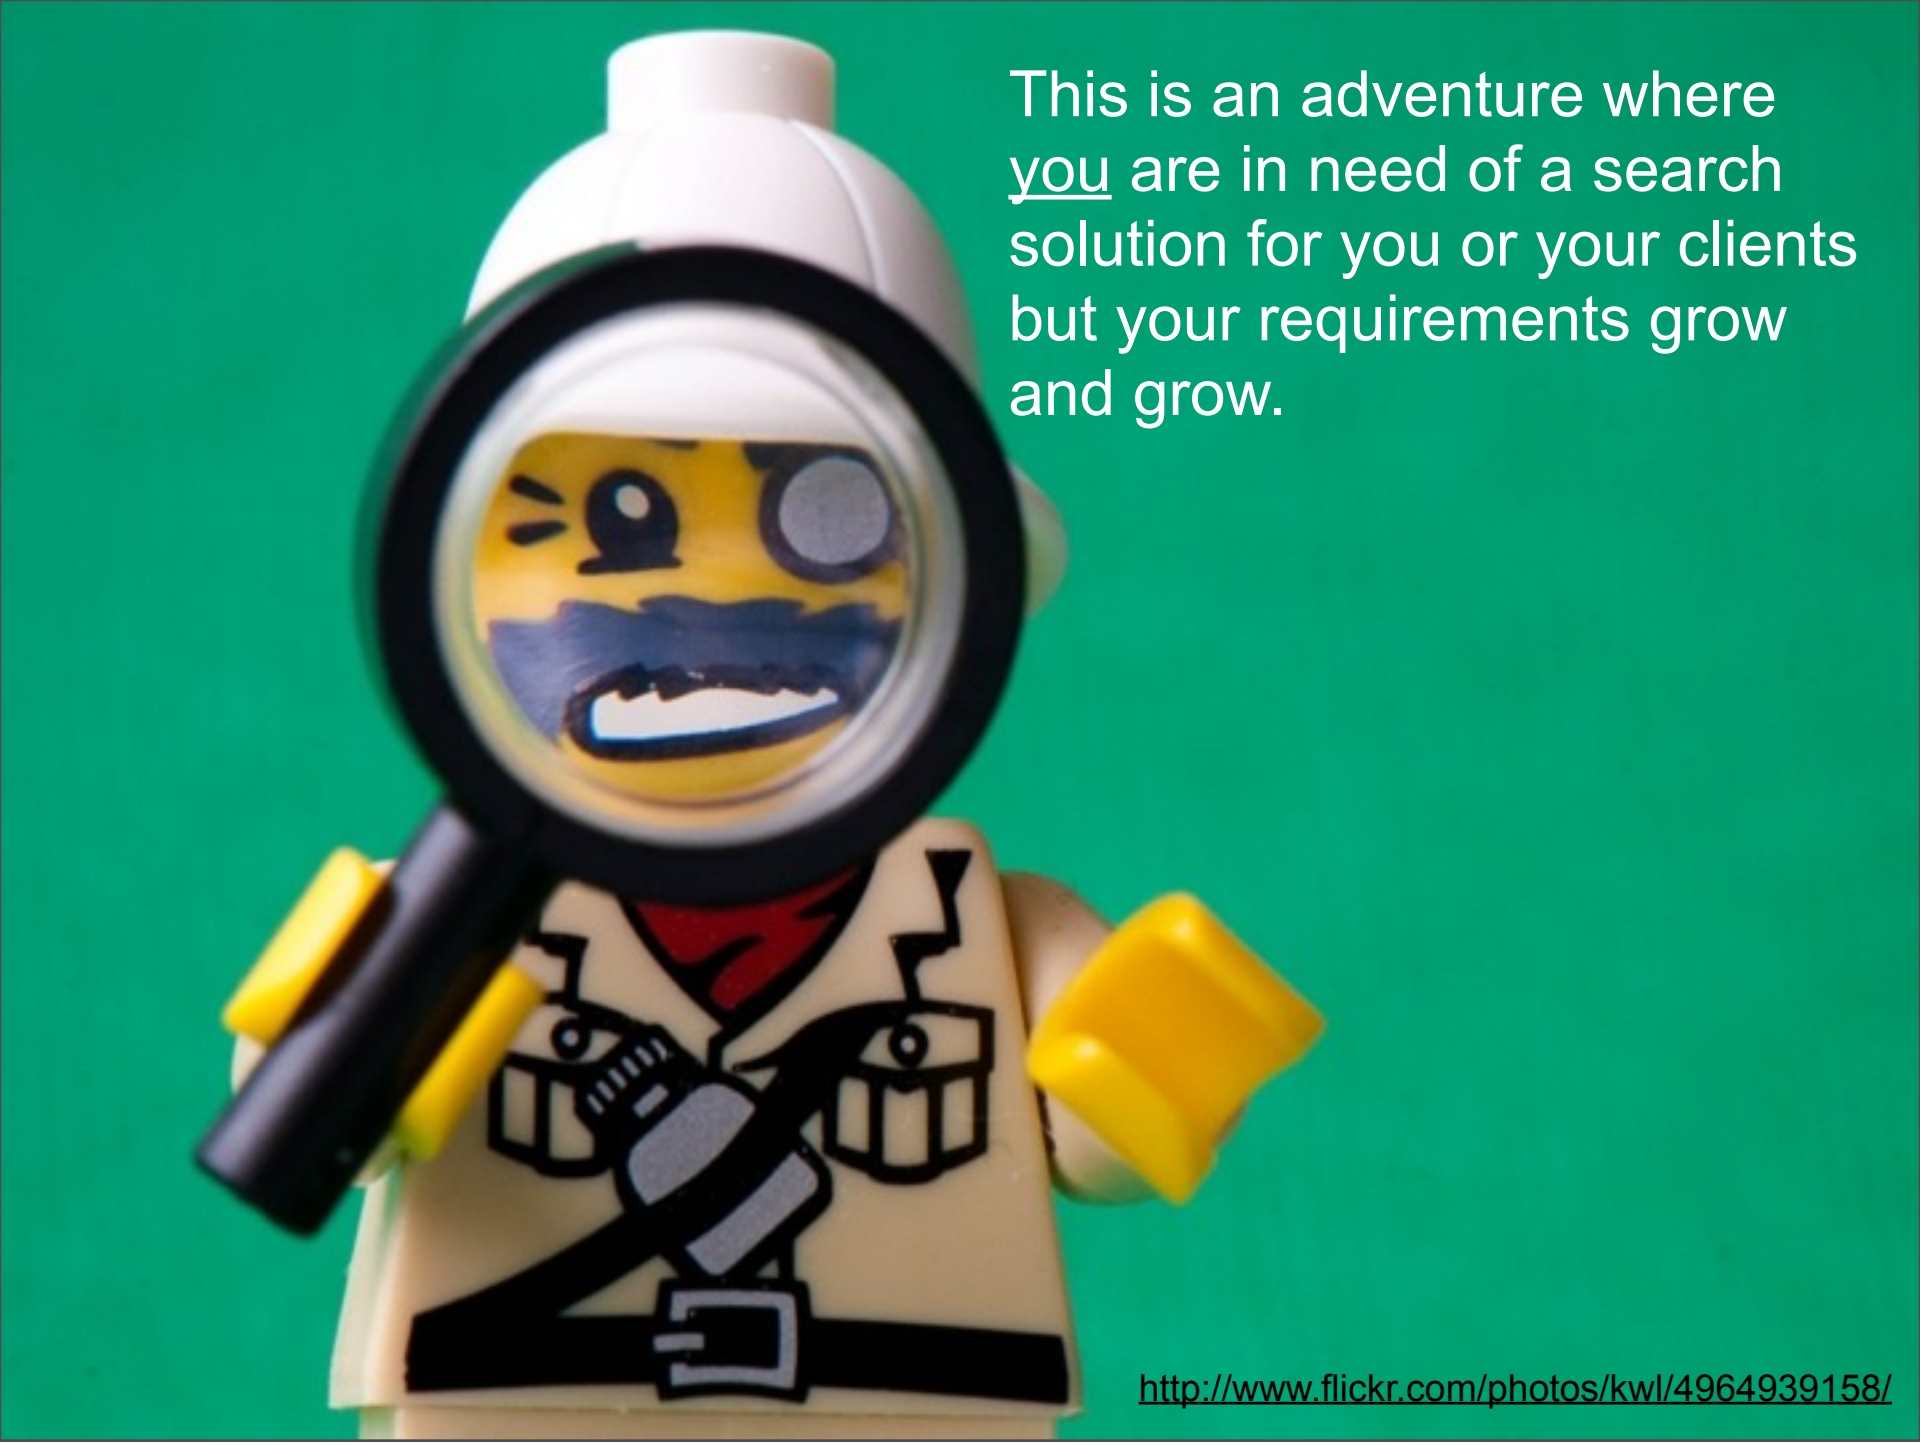
This is an adventure where you are in need of a search solution for you or your clients but your requirements grow and grow.
http://www.flickr.com/photos/kwl/4964939158/

Overview

- Basic Understanding
- Monitoring
- Optimizing your server
- Load balancing
- Template it in puppet
- Scaling up to +1000 cores
- Provision new cores automagically
- Keeping it secure

Getting started

```
java -jar start.jar
java -Dsolr.solr.home=multicore -jar start.jar
```

Tip: Easy local install guide at

http://nickveenhof.be/blog/simple-guide-install-apache-solr-3x-drupal-7

Caveats:

- No HA
- No restart on reboot
- No security

Problem 1

PID	Process Name	% CPU ▼	Threads	Real Mem	K
38375	httpd	62.8	1	81.9 MB	I
38082	java	34.2	40	138.9 MB	I

"My Server CPU and memory skyrockets whenever I start my indexing process. It looks like Solr eats up everything."

Likely cause: Solr, HTTPD and MySQL on the same server

Spread the load

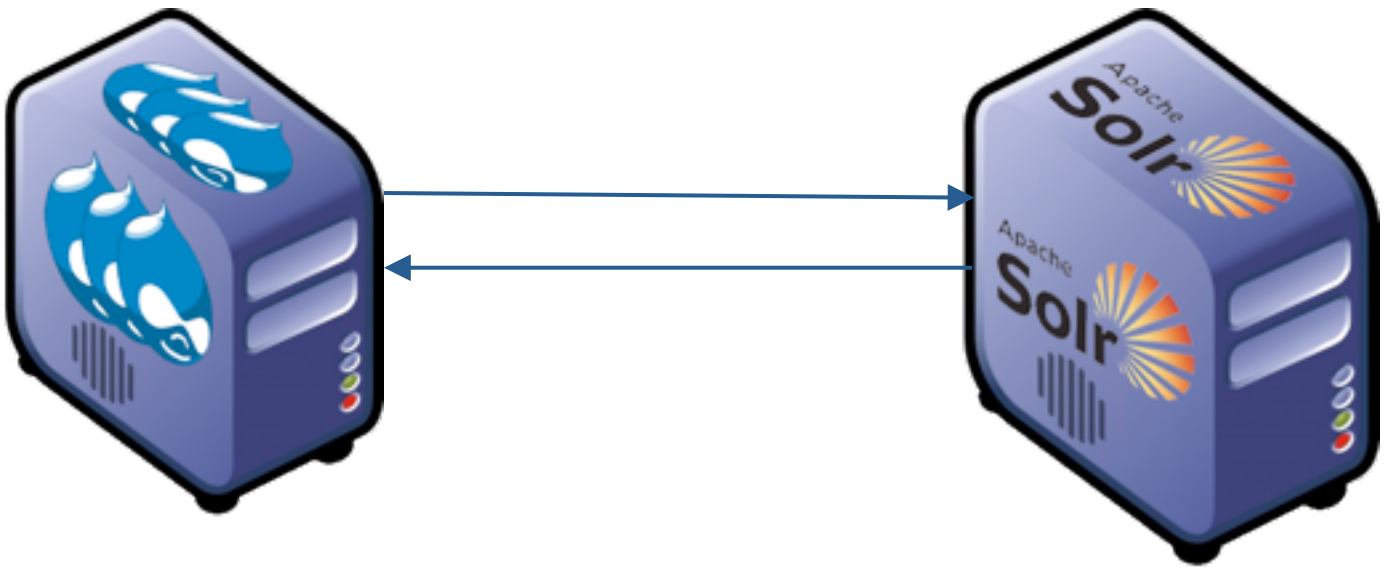

Problem 2

"My Server CPU and memory usage still skyrockets for some queries. What is going on?"

Likely cause: Query doesn't make sense or stopwords are not defined

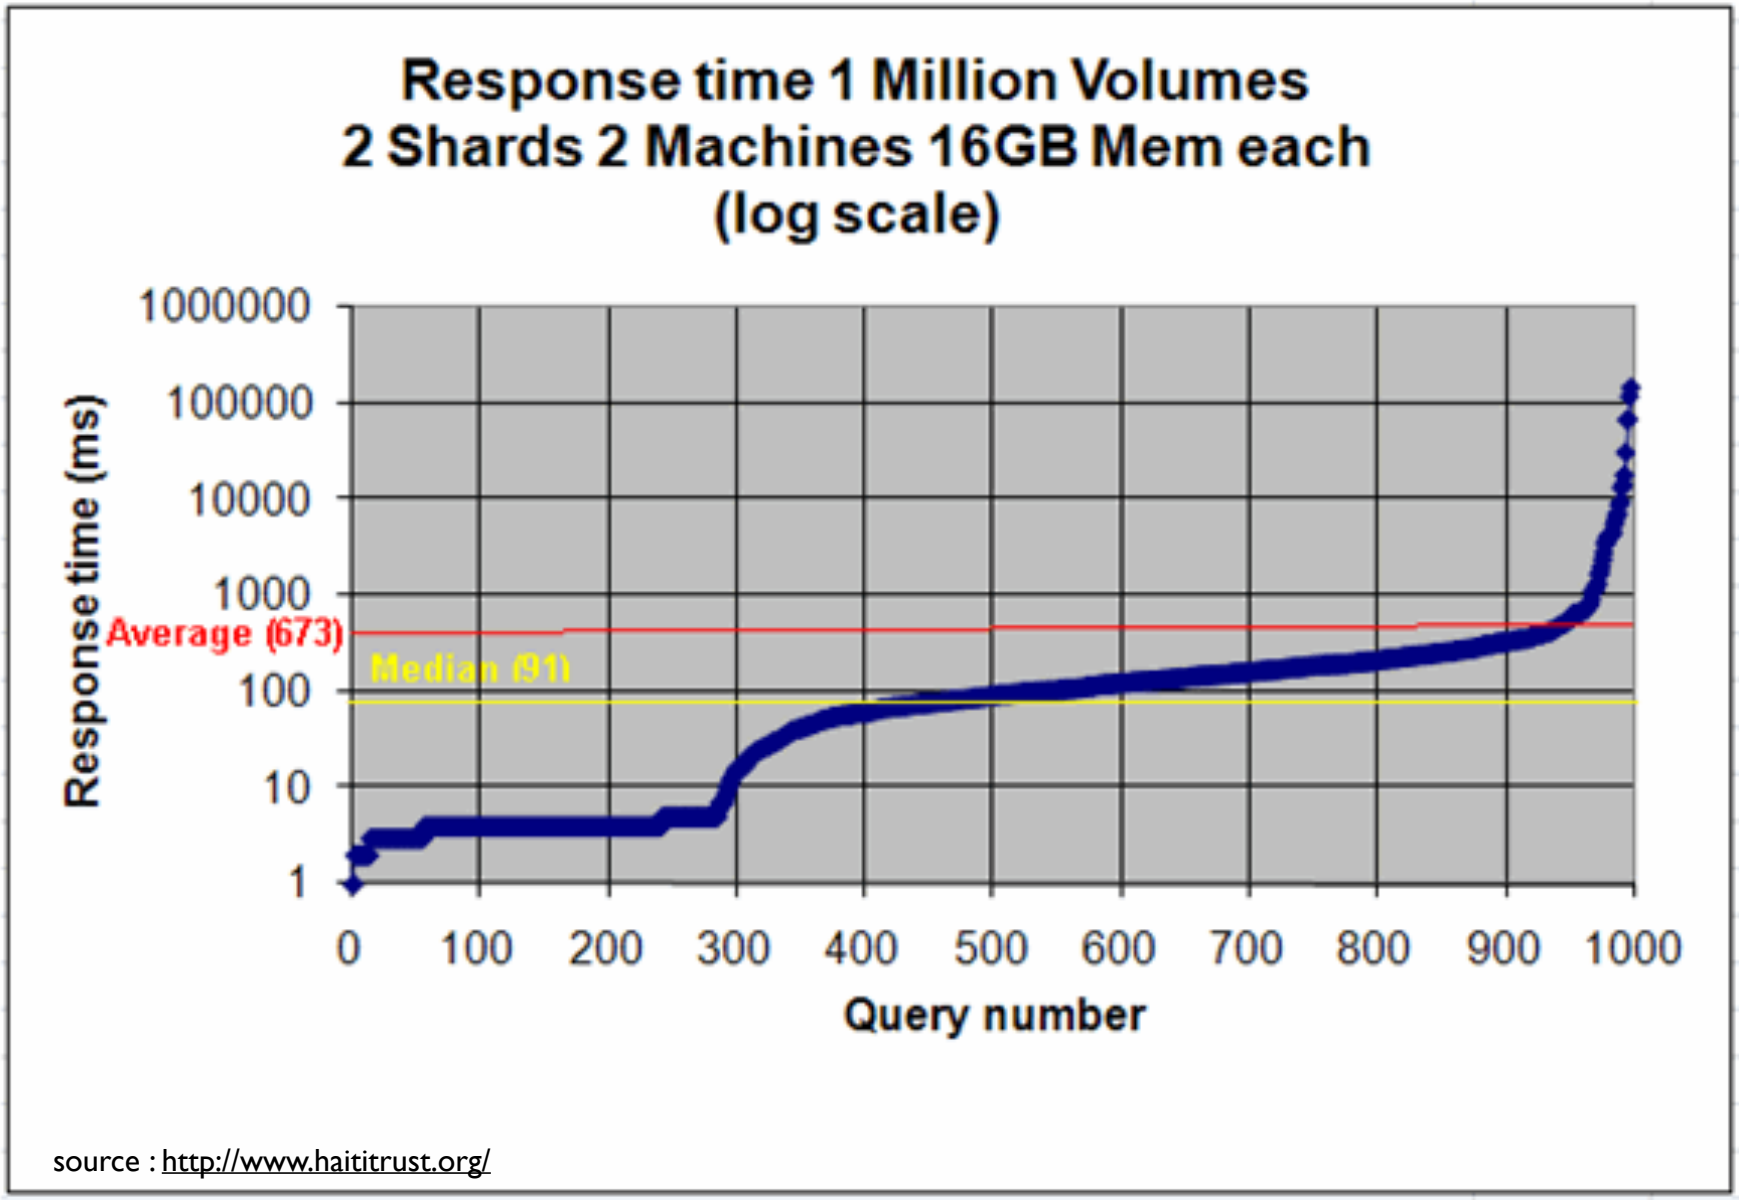

source : http://www.haititrust.org/

How can I debug Solr?

- Enable extra debugging info

```
select/
?q=Robin+Hood&debugQuery=on&debug=on
```

- Indentation!

```
select/?q=Robin+Hood&indent=true
```

- admin/analysis.jsp?highlight=on
- Tomcat logs, jetty logs!

What do these params mean?

- Query (q)

```
select/?q=superhero
```

- sort, start, rows

```
select/?q=superhero&start=0&rows=10&
sort=sort_name+asc
```

What do these params mean?

- Filter Query (fq)

```
select/?q=superhero&fq=bundle:person&
fq=attribute:cape
```

- Fields (to return) (fl)

```
select/?q=superhero&fl=id,entity_id,
name,attribute,score
```

What about dismax/edismax?

- Highlighting (hl, hl.q, hl.fl)

```
select/?q=superhero&hl=true&hl.q=super&
hl.fl=name,content,comments
```

- defType

```
select/?q=superhero+AND+evil&
defType=edismax
```

What about dismax/edismax?

- Alternative Query (q.alt)

```
select/?q.alt=bundle:person
```

- Query fields (qf)

```
select/?q=Superhero&qf=teaser^2.0
```

- Phrase Fields (pf)

```
select/?q=Robin Hood&pf=name^10
```

Problem 3

"Now I want to take advantage of this separate server and host search indexes for several sites"

" How can I be certain that everything is actually loaded and is working fine?"

Monitoring

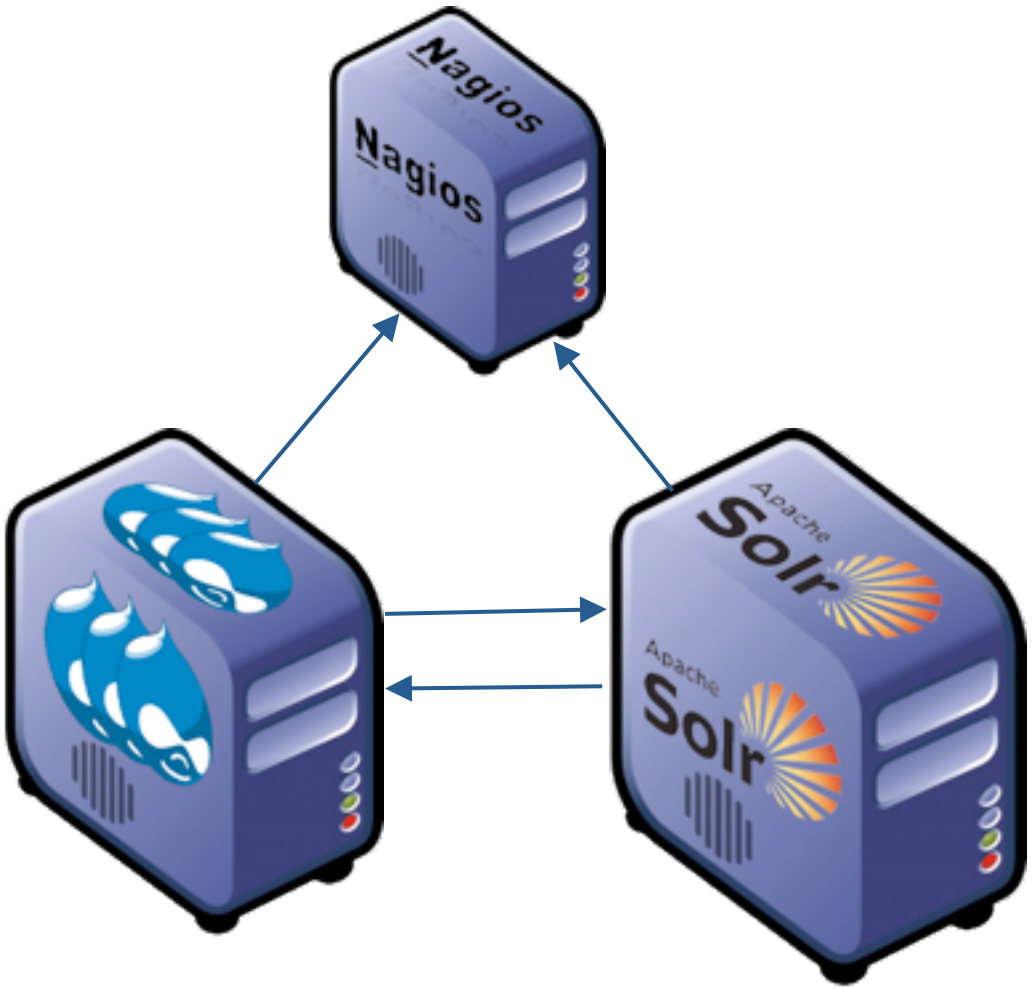

Monitoring

- Average Time Per Request & Requests per second

```
solr/core_name/admin/mbeans?
wt=json&stats=true&
key=org.apache.solr.handler.component.SearchHa
ndler& stats=true&cat=QUERYHANDLER
```

Monitoring


```
{
"responseHeader":{
  "status":0,
  "QTime":1},
"solr-mbeans":[
  "QUERYHANDLER",{
    "org.apache.solr.handler.component.SearchHandler":{
      ...
      "docs":null,
      "stats":{
        "handlerStart":1345463690388,
        "requests":2,
        "errors":0,
        "timeouts":0,
        "totalTime":75,
        "avgTimePerRequest":37.5,
        "avgRequestsPerSecond":0.0013287809}}}]}
```

Monitoring

- Number and identity of the cores

/admin/cores?wt=json&action=STATUS


```
{
  "responseHeader":{
    "status":0,
    "QTime":6},
  "status":{
    "core0":{
      "name":"core0",
      "instanceDir":"multicore/core0/",
      "dataDir":"multicore/core0/data/",
      "startTime":"2012-08-20T11:54:50.275Z",
      "uptime":2015408,
      "index":{
        "numDocs":887,
        "maxDoc":1279,
        "version":1323430446081,
        "segmentCount":5,
        "current":true,
        "hasDeletions":true,
        "lastModified":"2012-08-02T15:43:12Z"}},
```

Monitoring

- Size of each core on the server
  Check ['solr-mbeans'][1]['/replication']['stats']['indexSize']

```
solr/core_name/admin/mbeans&wt=json&
key=/replication&
stats=true&cat=QUERYHANDLER
```

- Document Size
  Check ['solr-mbeans'][1]['searcher']['stats']['numDocs']

```
solr/core_name/admin/mbeans&wt=json&
key=searcher&stats=true
```

Monitoring - New Relic

- New Relic is a useful tool to have a deeper insight. Comes with full Solr support.
- New relic does not allow a per-core granularity. So not appropriate to hand over the credentials to customers.
- Performance impact has not been proven nor tested yet. Be careful when using this tool

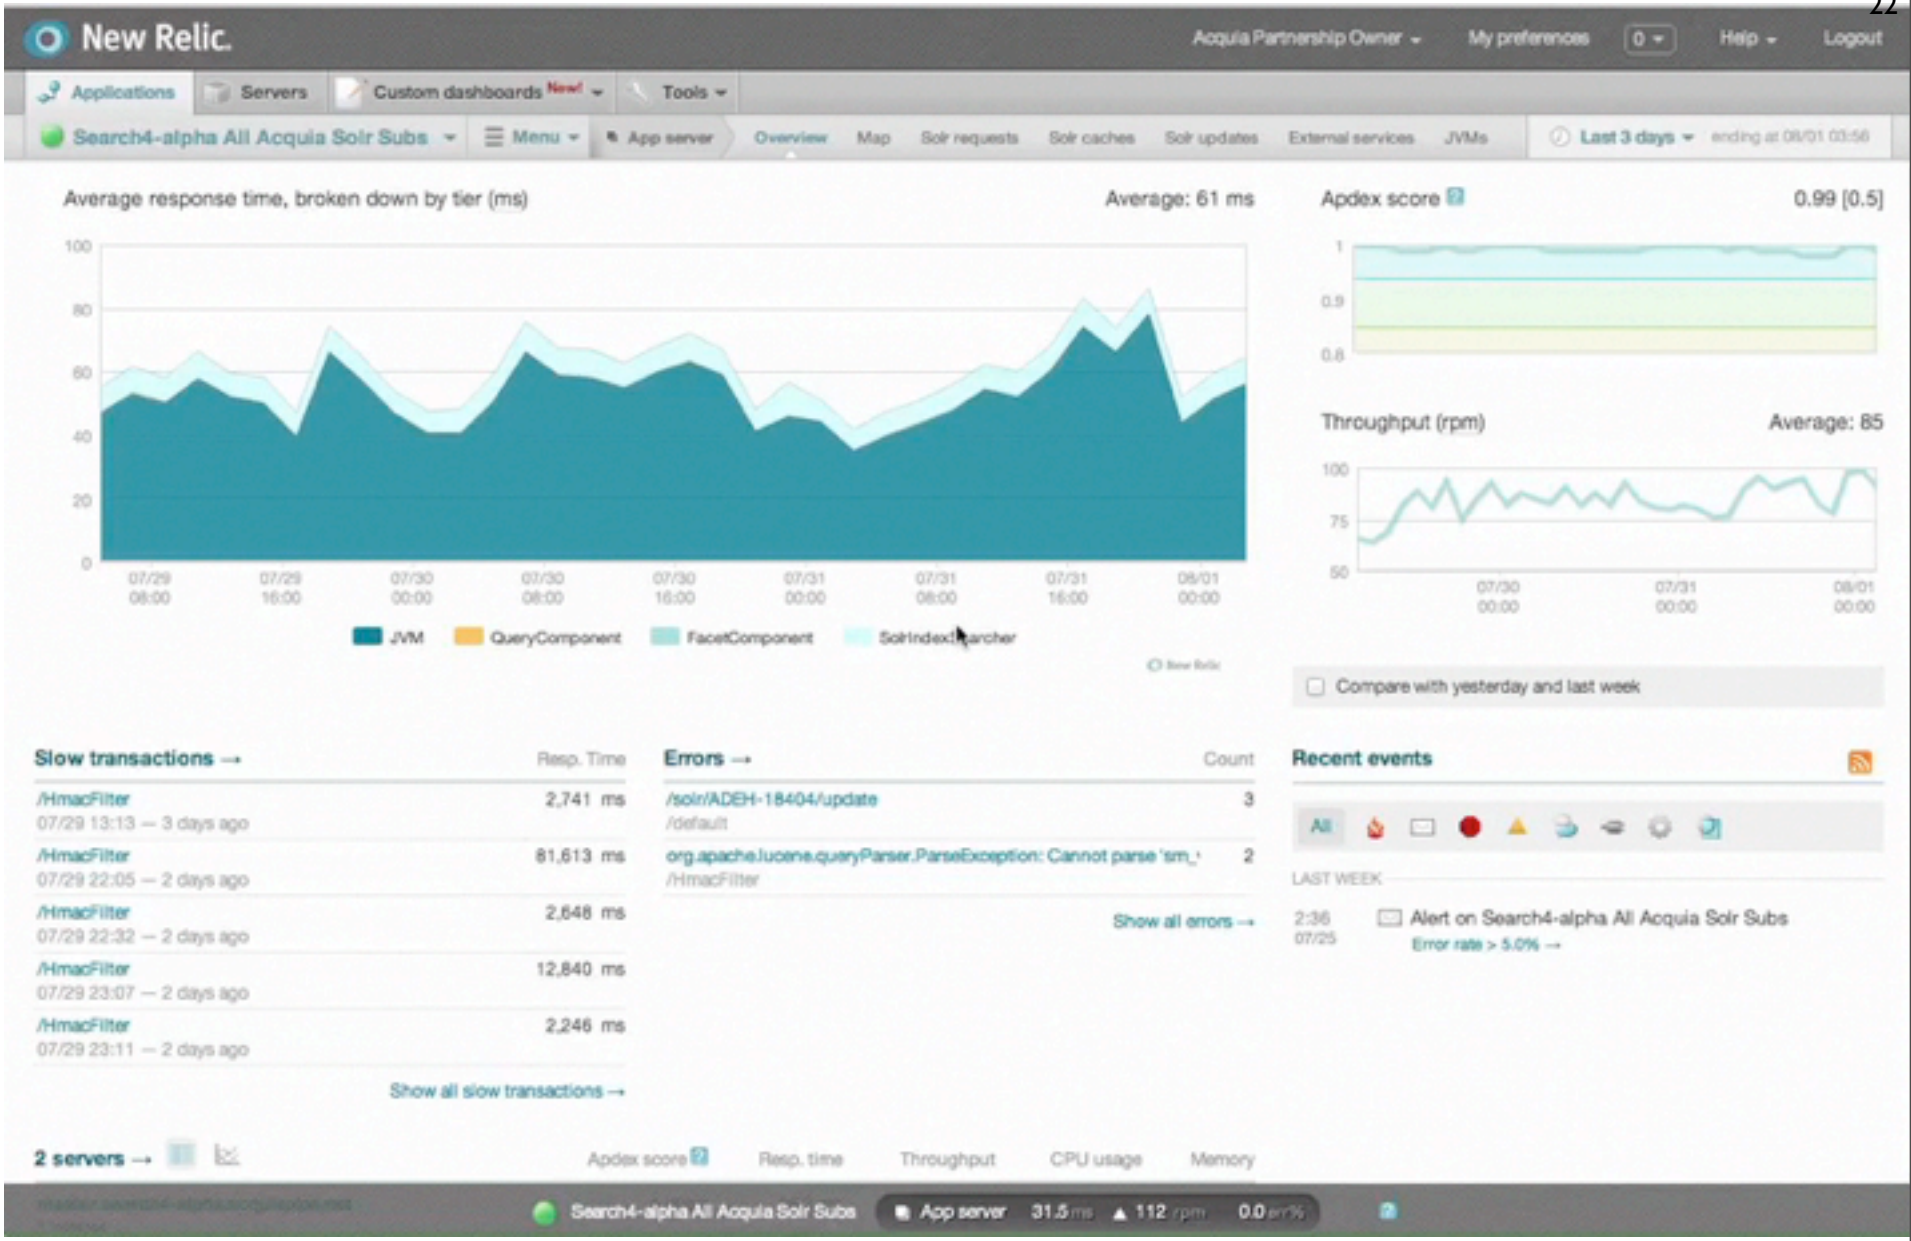
New Relic.
Acquia Partnership Owner
My preferences
Help
Logout
Applications
Servers
Custom dashboards
Tools
Search4-alpha All Acquia Solr Subs
Menu
App server
Overview
Map
Solr requests
Solr caches
Solr updates
External services
JVMs
Last 3 days
Average response time, broken down by tier (ms)
Average: 61 ms
JVM
QueryComponent
FacetComponent
Apdex score
0.99 [0.5]
Throughput (rpm)
Average: 85
Compare with yesterday and last week
Slow transactions
Resp. Time
/HmacFilter
2,741 ms
81,613 ms
2,648 ms
12,840 ms
2,246 ms
Show all slow transactions
Errors
Count
/solr/ADEH-18404/update
/default
org.apache.lucene.queryParser.ParseException: Cannot parse 'sm_'
/HmacFilter
Show all errors
Recent events
LAST WEEK
Alert on Search4-alpha All Acquia Solr Subs
Error rate > 5.0%
2 servers
Apdex score
Resp. time
Throughput
CPU usage
Memory
Search4-alpha All Acquia Solr Subs
App server
31.5 ms
112 rpm
0.0 err%

Problem 4

"So I monitored my servers now, but am I utilizing my server at it highest capacity?"

Optimizing your server

- Pick a reasonable fraction of your machine's memory (30-70%) depending on how it's used

```
JAVA_OPTS="-server -
Djava.awt.headless=true -Xms1000m -
Xmx1000m"
```

- Depending on your amount of CPU'S

```
JAVA_OPTS="$JAVA_OPTS -XX:
+CMSIncrementalMode"
```

Optimizing your solrconfig

```
<luceneMatchVersion>LUCENE_35</
luceneMatchVersion>


<mergePolicy
class="org.apache.lucene.index.LogByteS
izeMergePolicy" />
```

- Tip: More performance info can be found @ http://nickveenhof.be/blog/upgrading-apache-solr-14-35-and-its-implications

Problem 5

"Help, how do I spread the load of my solr cores. My hardware has been maxed out?!?"

Replication

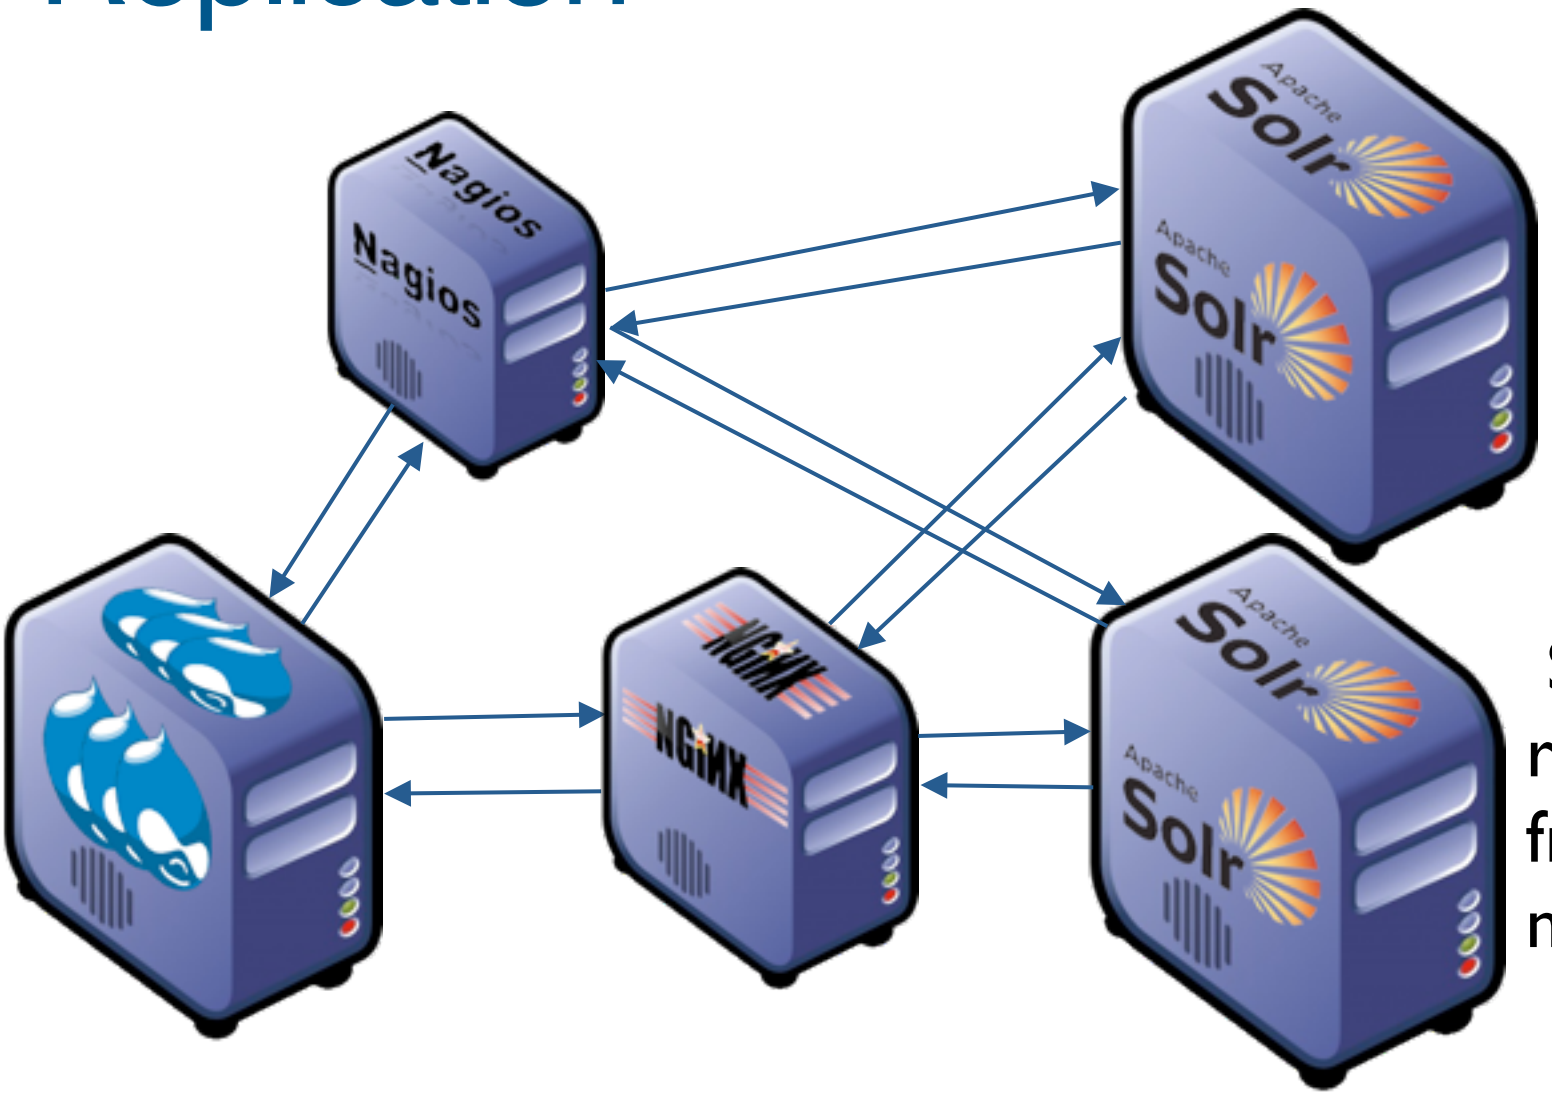

Replication

```
#solrcore.properties file
enable.master=false
enable.slave=true
poll_time=00:02:00
master_core_url=http://localhost:8983/solr/MYMASTERCORE
```

This file or support is not yet committed to both projects, but the common solrconfig/schema initiative is making sure it will.
http://drupal.org/sandbox/cpliakas/1600962

Problem 6

"What about availability? My hosting provider says they never have downtime, I know. But what if?"

Multi Data-center Replication For the Truly Demanding

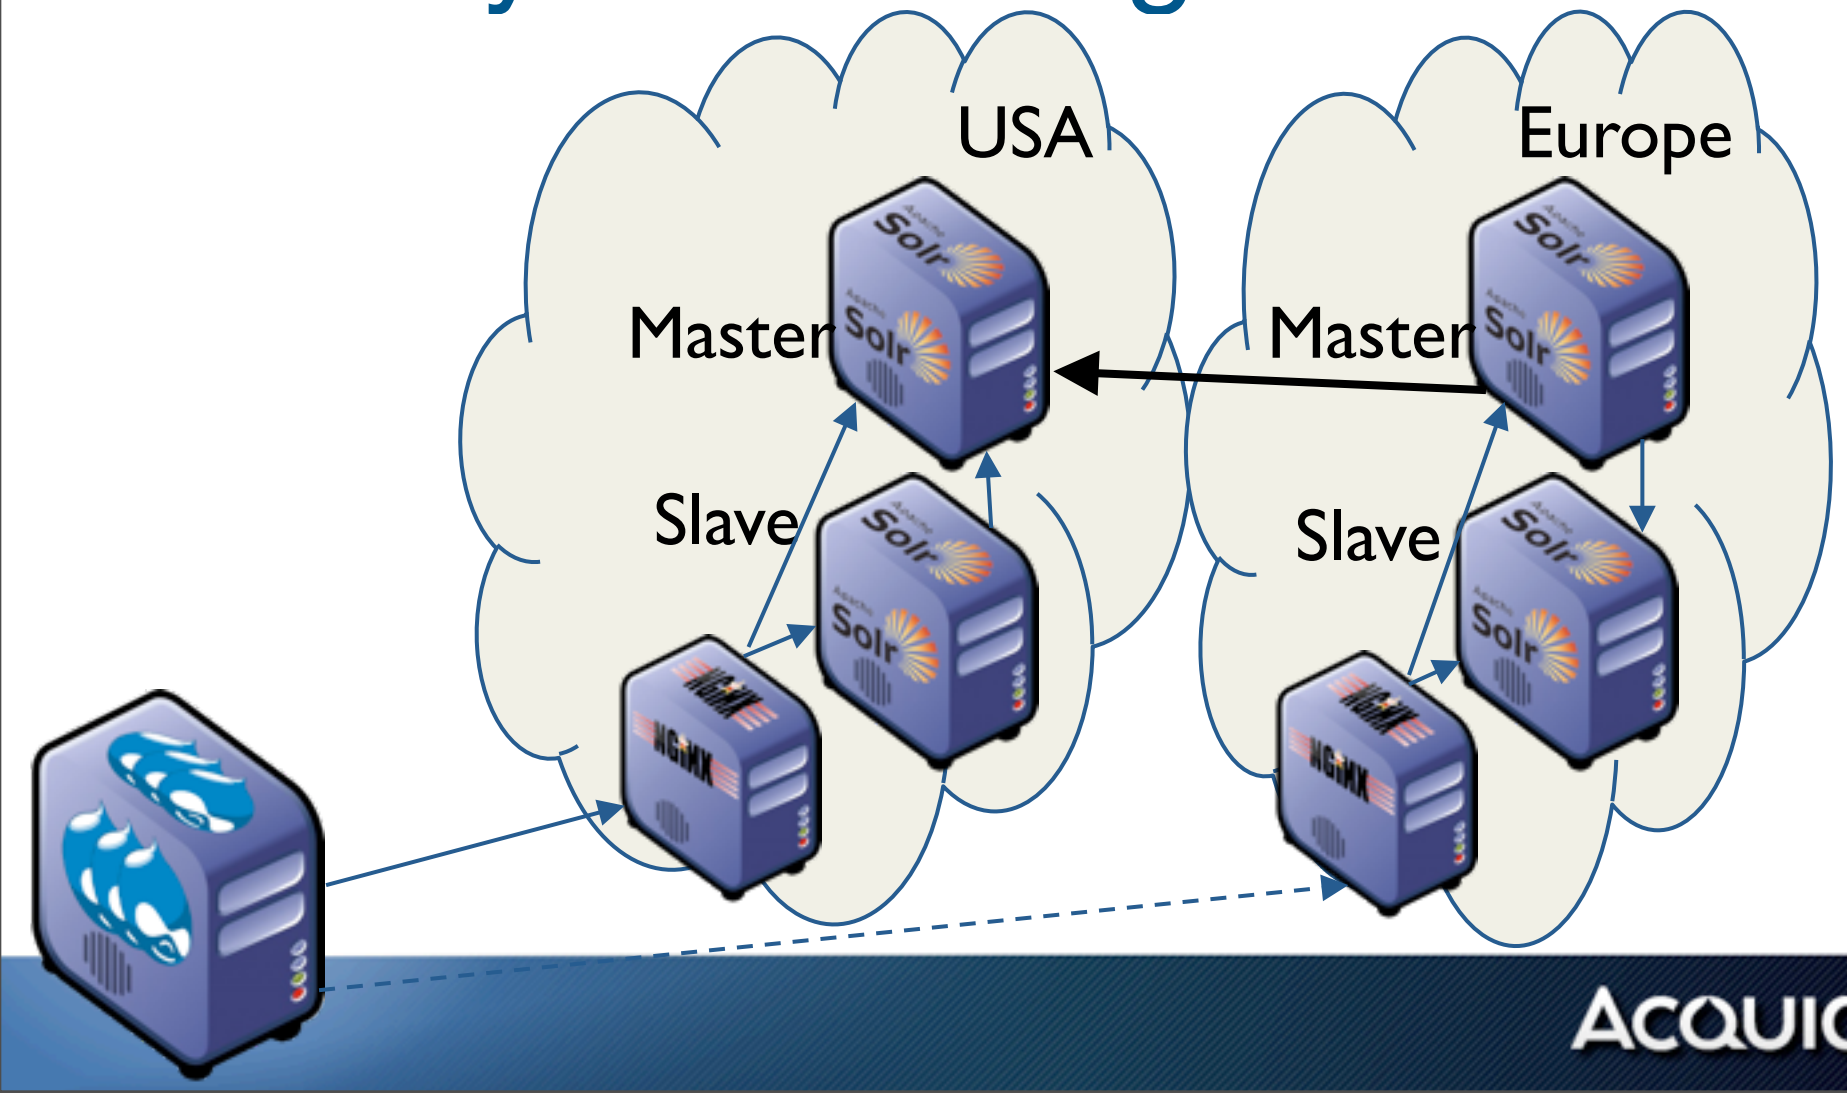

Acquia

Automate it all

"Nice, I have the perfect setup. However, it is tiresome to always set up new server and repeat what I've done."

Template it in puppet

```
class tomcat {
  package {"openjdk-6-jdk":
    ensure => installed,
  }
  package { "tomcat6":
    require => [ Package["openjdk-6-jdk"] ],
    ensure => installed,
  }
  package { "libtcnative-1":
    require => [ Package["tomcat6"] ],
    ensure => installed,
  }
  service { "tomcat6":
    require => [ Package["tomcat6"], Package["libtcnative-1"] ],
    ensure => running,
  }
}
```

Template it in puppet

```
class solr {

  package { "solr":
     require => [ Package["tomcat"] ],
     ensure => installed,
     name => "solr-common",
  };
  file { "solr initscript":
    ensure => present,
    path => "/etc/init.d/solr",
    owner => root,
    group => root,
    mode  => "0755",
    content => template("solr/solr.init.erb");

  file { "solr conf":
    ensure => present,
    path => "/usr/share/solr/solr.xml",
    owner => root,
    group => root,
    mode => "0755",
    content => template("solr/solr.xml.erb");
  }
  service { "solr":
    ensure => "running",
    enable => "true",
    require => File["solr initscript", "solr conf"];
  }

}
```

Example was adjusted, but more extensive example at: https://github.com/KrisBuytaert/puppet-solr/

Template it in puppet

Local scripts to run new machines:

```
./provision --server-allocate searchsrv --size m1.large
```

Server types are a combination of different puppet classes:

```
'searchsrv'    => [ 'tomcat', 'solr', 'package::subclass' ],
```

Problem 7

"So, I hate the fact that I have to create a new folder and restart solr. There must be a way to automate the creation of new Solr Cores?"

CoreAdmin

```
solr/admin/cores?
action=CREATE&name=coreX&
instanceDir=path_to_instance_dir
```

- More : http://wiki.apache.org/solr/CoreAdmin
- No upload functionality.
- Look at the module test code for example calls.
- Core admin changes my be persistent or temporary depending on the solr.xml settings:
  `<solr persistent="true">`

Automate core provisioning

- What's your data source for the list of cores: DB, XML file, RPC result?
- Every core that's created needs configuration and data directories - need to copy the conf in place in advance.
- Script (bash, ruby, java, etc) any copying and core admin http calls to iterate through the cores specified.
- Look at Solr Cloud for Solr 4.0.

Scaling up to +1000 cores

- Calculate and do performance testing on multiple cores at once. jmeter is perfect to analyze the results
- Use customized load balancers to redirect the right core to the right machine

Nginx as load balancer

```
location ~ ^/solr/core_name/select/?$ {
  proxy_read_timeout 20;
  proxy_next_upstream error timeout http_404 http_500;
  proxy_pass   http://search_farm_1_select;
}

location ~ ^/solr/core_name/update(|/csv)/?$ {
  proxy_pass   http://search_farm_1_master;
}
...
```

Keeping it secure

- NO security by default, not present at all for per-core security. Google this: “[SCHEMA] [CONFIG] [ANALYSIS] [SCHEMABROWSER]"
- firewall rules
- SSL
- SSL + basic auth
- Acquia Search uses HMAC authentication & validation that is secure with or without SSL

Search master servers

Your web-server

Acquia
Network

index
replication

Search slave servers

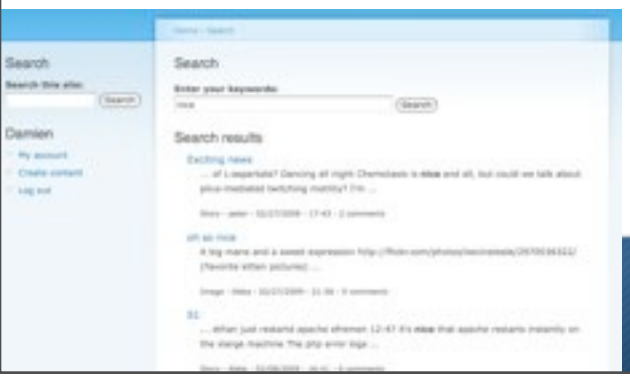

Acquia

Search master servers

Your web-server

authenticated request

content to index

SSL, HMAC

index replication

Acquia Network

Search slave servers

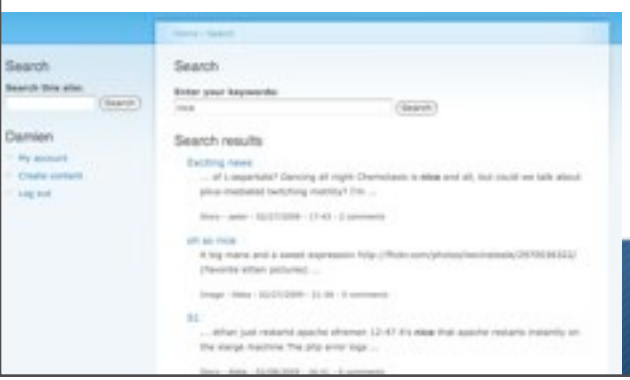

Acquia

Search master servers

Your web-server

Acquia
Network

index
replication

Search slave servers

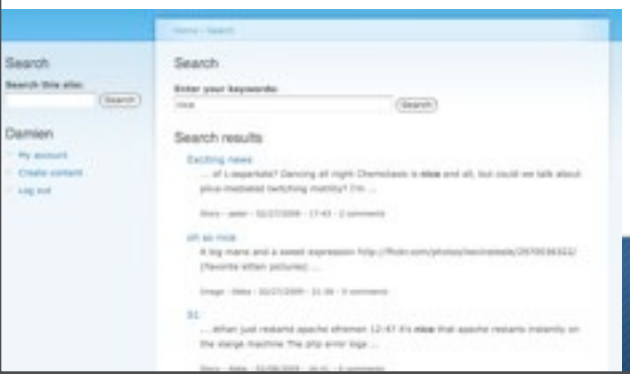

Acquia

Search master servers

Your web-server

SSL, HMAC

search request

Acquia Network

authenticated results

index replication

Search slave servers

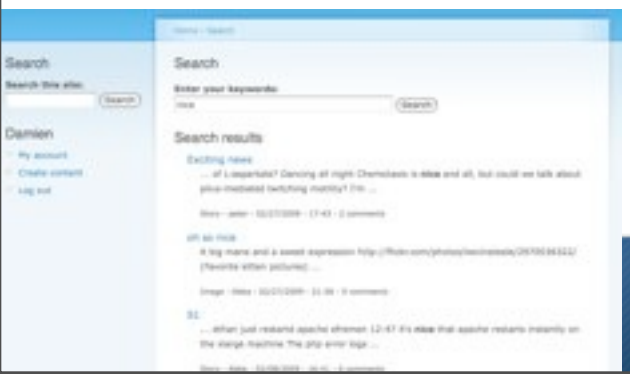

Acquia

Search master servers

Your web-server

Acquia
Network

index
replication

Search slave servers

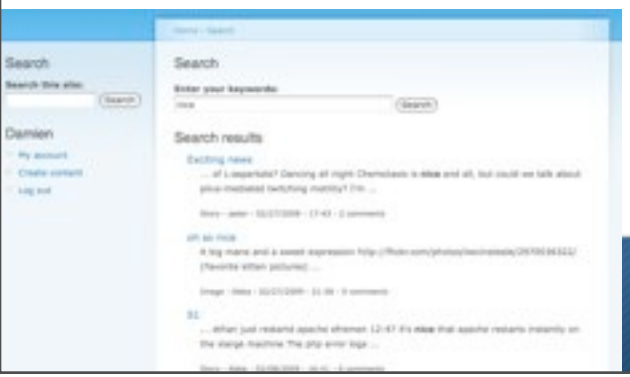

Acquia

Search master servers

Your web-server

SSL, HMAC

Acquia Network

index replication

Search slave servers

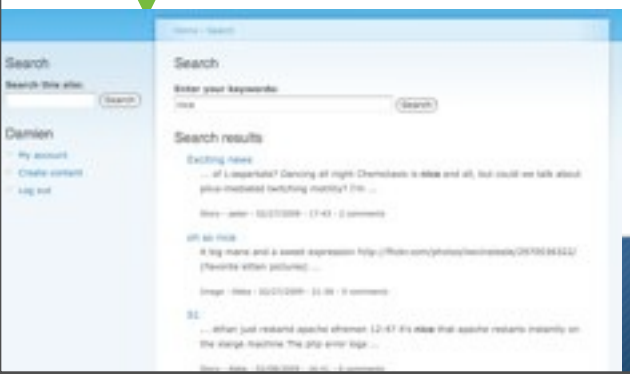

Acquia

Wrapping up

- If you want to try Apaches Solr, it only takes about 5 minutes to install locally.
- Putting Solr into production takes work, and we've covered some of the challenges.
- Or, Acquia Search already has these problems solved and it can be used locally or with any server.
- New free tier for Acquia Network services.

What did you think?

Please rate this session on the DrupalCon Munich website:

http://munich2012.drupal.org/node/678

Thank you!