MidCamp

High Performance Drupal

CHRISTOPHER J. WELLS
REDFIN SOLUTIONS, LLC

tweet @sceo • http://redfinsolutions.com

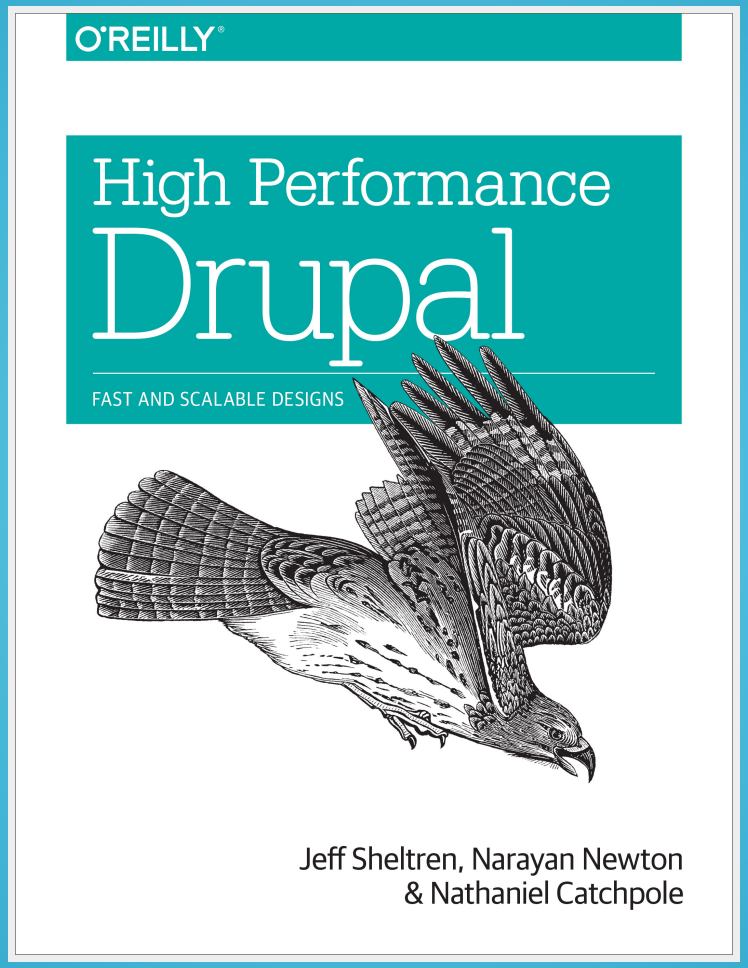
O'REILLY
High Performance
Drupal
FAST AND SCALABLE DESIGNS
Jeff Sheltren, Narayan Newton
& Nathaniel Catchpole

THE BOOK

A QUICK REFRESHER

How does all this work again?

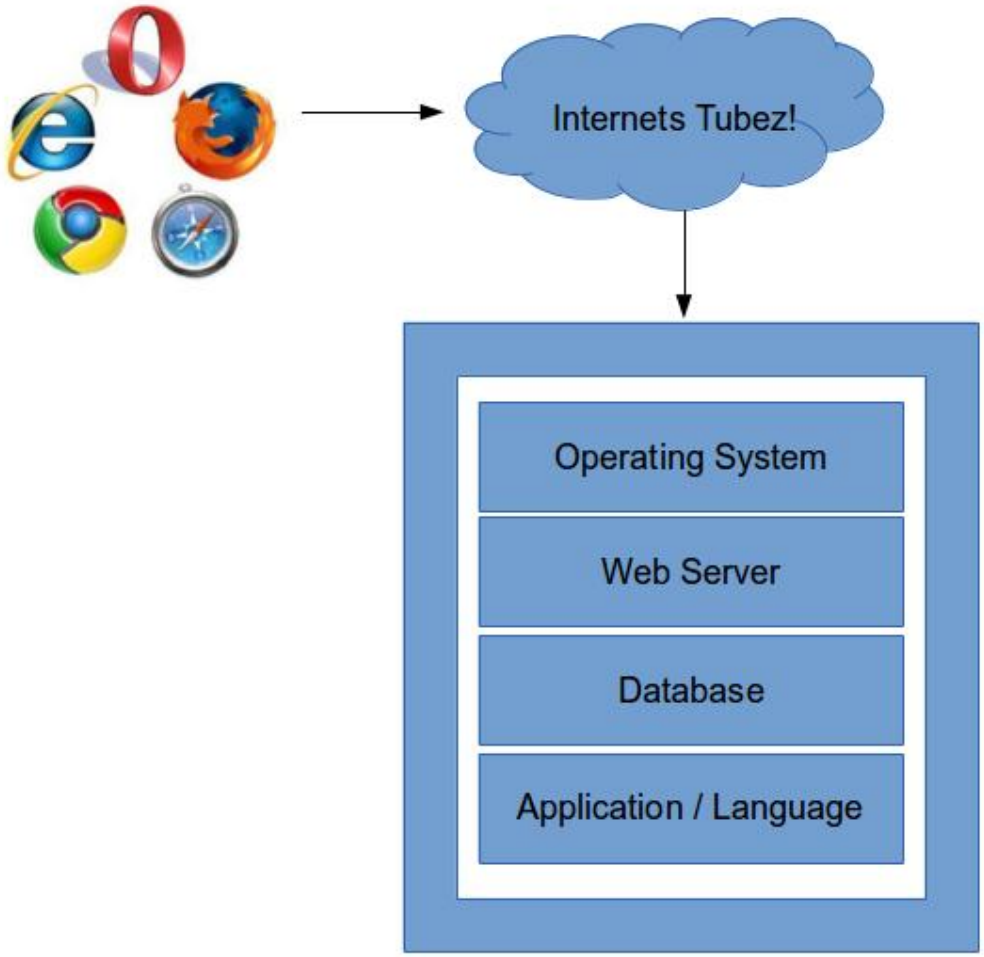
Internets Tubez!
Operating System
Web Server
Database
Application / Language

And we call this, typically, the LAMP stack.

WHAT COULD POSSIBLY GO WRONG?

ESTABLISH A BASELINE

- get basic readings
- find averages
- identify what improvement means
- only then do we begin the diagnosis

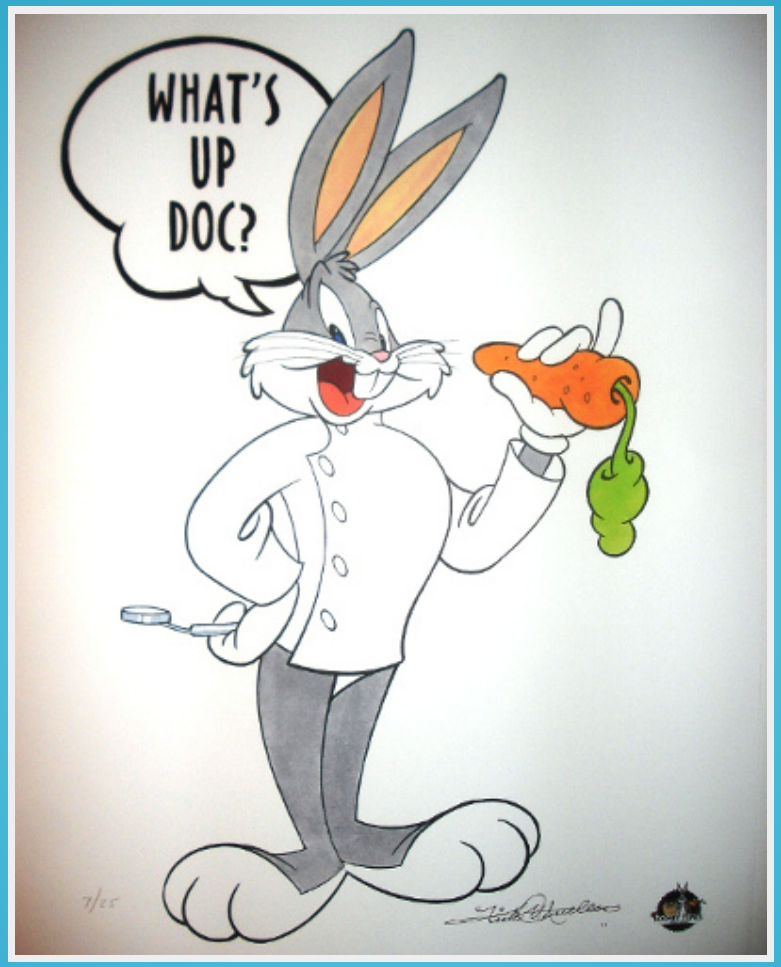
WHAT'S
UP
DOC?

ANY GOOD DOCTOR HAS TOOLS.

SO HOW TO START THE DIAGNOSIS?

- Browser's built in debugger
- Google Analytics
- Google PageSpeed / Yahoo's YSlow
- Real User Monitoring (New Relic, ex.g.)

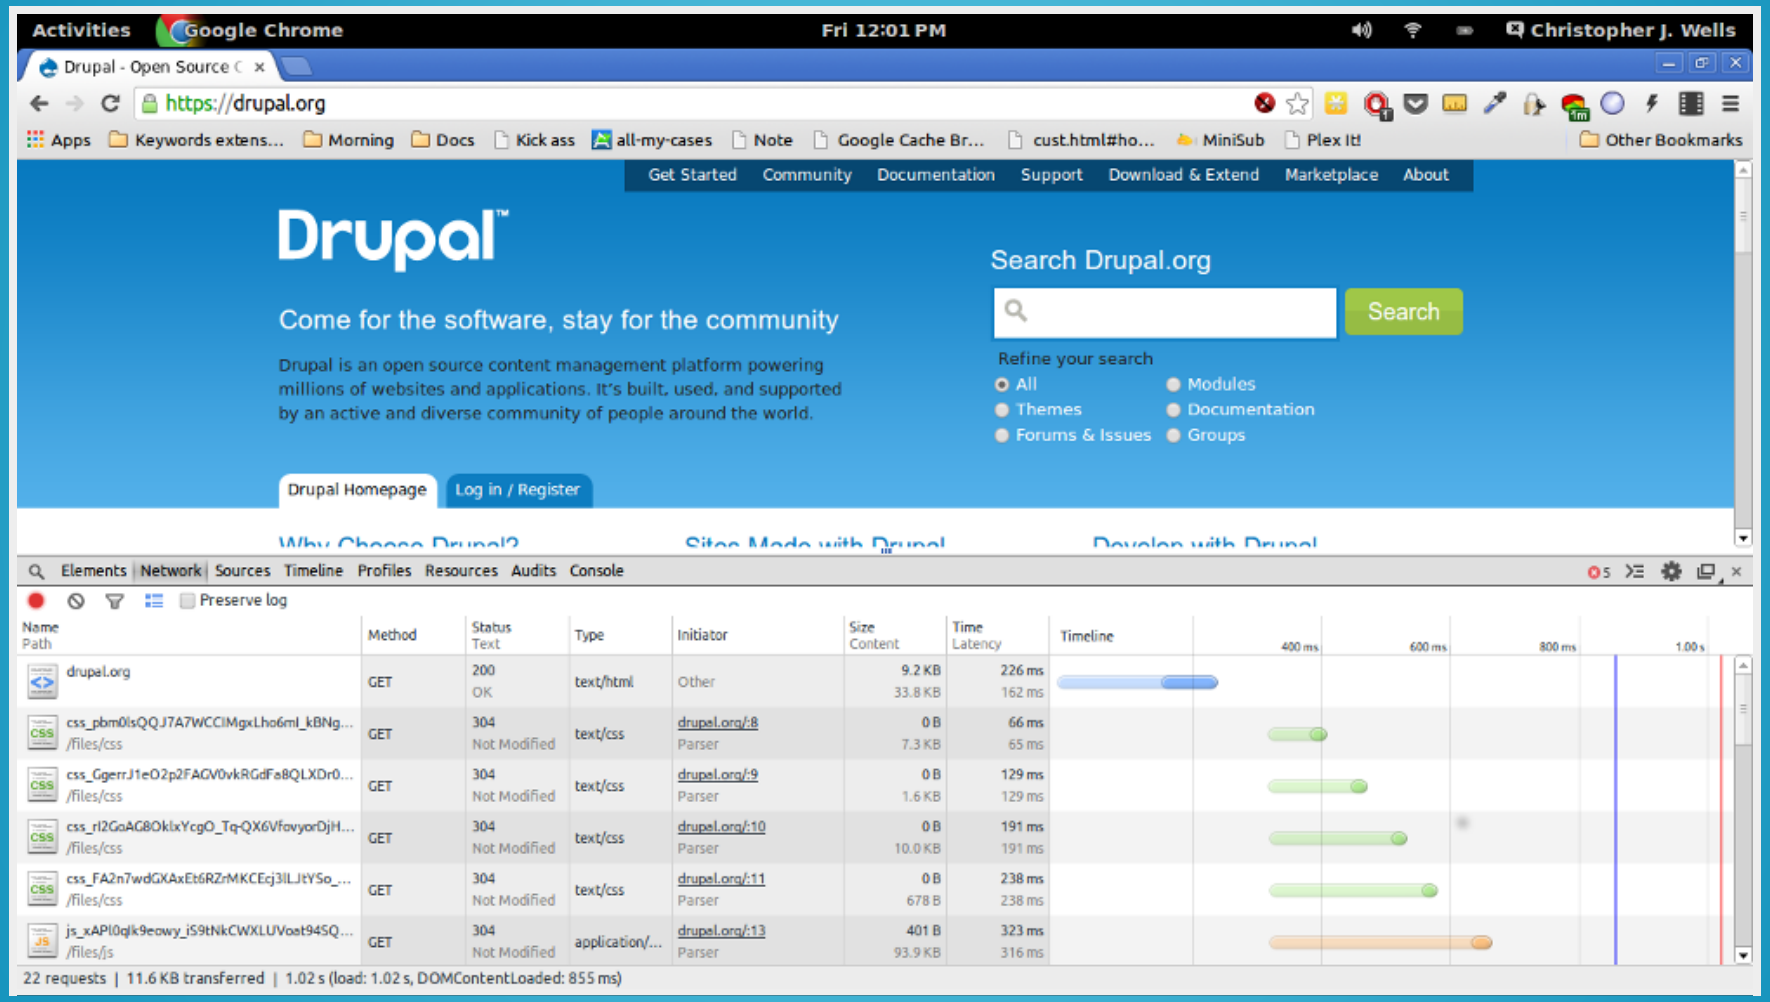
Activities
Google Chrome
Fri 12:01 PM
Christopher J. Wells
Drupal - Open Source
https://drupal.org
Apps
Keywords extens...
Morning
Docs
Kick ass
all-my-cases
Note
Google Cache Br...
cust.html#ho...
MiniSub
Plex It!
Other Bookmarks
Get Started
Community
Documentation
Support
Download & Extend
Marketplace
About
Drupal
Come for the software, stay for the community
Drupal is an open source content management platform powering millions of websites and applications. It's built, used, and supported by an active and diverse community of people around the world.
Search Drupal.org
Search
Refine your search
All
Modules
Themes
Documentation
Forums & Issues
Groups
Drupal Homepage
Log in / Register
Elements
Network
Sources
Timeline
Profiles
Resources
Audits
Console
Preserve log
Name
Path
Method
Status
Text
Type
Initiator
Size
Content
Time
Latency
Timeline
400 ms
600 ms
800 ms
1.00 s
drupal.org
GET
200
OK
text/html
Other
9.2 KB
33.8 KB
226 ms
162 ms
/files/css
304
Not Modified
text/css
drupal.org/:8
Parser
0 B
7.3 KB
66 ms
65 ms
drupal.org/:9
1.6 KB
129 ms
drupal.org/:10
10.0 KB
191 ms
drupal.org/:11
678 B
238 ms
/files/js
application/...
drupal.org/:13
401 B
93.9 KB
323 ms
316 ms
22 requests | 11.6 KB transferred | 1.02 s (load: 1.02 s, DOMContentLoaded: 855 ms)

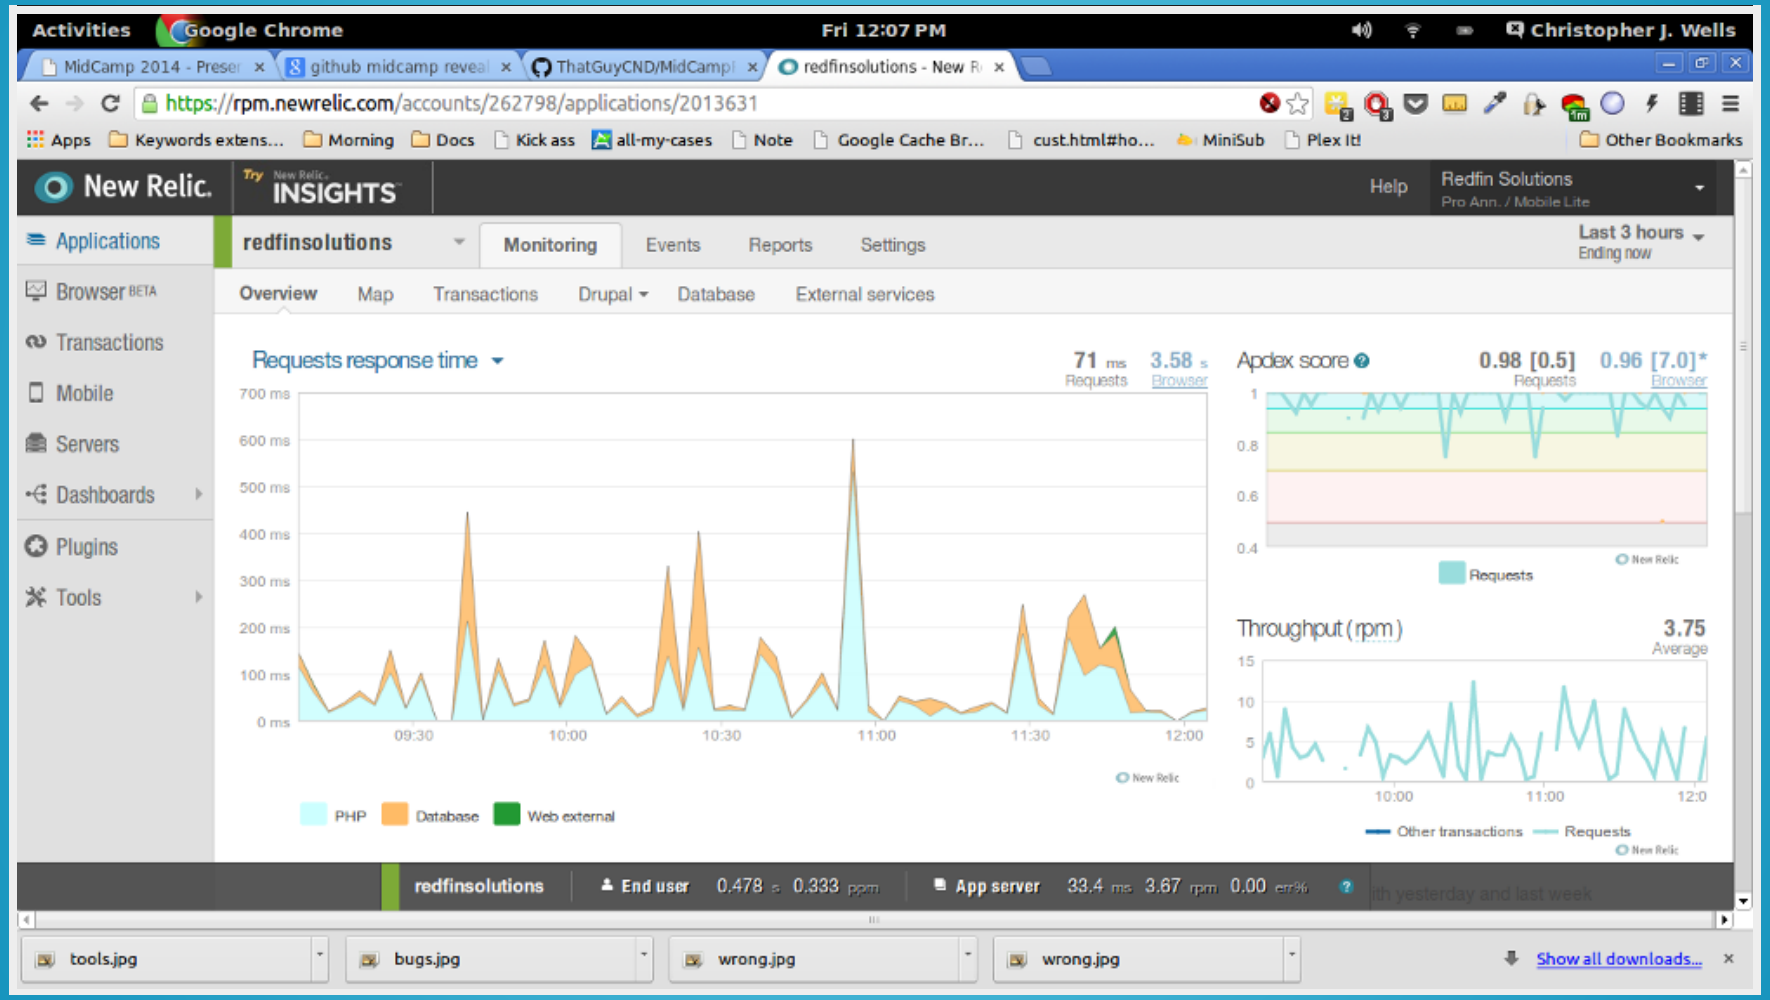
Activities
Google Chrome
Fri 12:07 PM
Christopher J. Wells
https://rpm.newrelic.com/accounts/262798/applications/2013631
New Relic.
INSIGHTS
Help
Redfin Solutions
Pro Ann. / Mobile Lite
Applications
Browser BETA
Transactions
Mobile
Servers
Dashboards
Plugins
Tools
redfinsolutions
Monitoring
Events
Reports
Settings
Last 3 hours
Ending now
Overview
Map
Transactions
Drupal
Database
External services
Requests response time
71 ms
Requests
3.58 s
Browser
700 ms
600 ms
500 ms
400 ms
300 ms
200 ms
100 ms
0 ms
09:30
10:00
10:30
11:00
11:30
12:00
PHP
Database
Web external
Apdex score
0.98 [0.5]
Requests
0.96 [7.0]*
Browser
1
0.8
0.6
0.4
Requests
Throughput (rpm)
3.75
Average
15
10
5
0
10:00
11:00
Other transactions
Requests
redfinsolutions
End user
0.478 s
0.333 rpm
App server
33.4 ms
3.67 rpm
0.00 err%
tools.jpg
bugs.jpg
wrong.jpg
wrong.jpg
Show all downloads...

WHAT ARE WE LOOKING FOR?

- start by identifying where on the path we're off
- "time to first byte" is the dividing line between client and server
- Continue the diagnosis server-side
  - Devel Module
  - Xdebug
  - XHProf [1]

[1] *this is amazing*

CLIENT-SIDE OPTIMIZATIONS

- optimize the requests (make fewer, or make them faster)
- optimize your markup
- beware of SPOFs
- Keep your libraries updated

NETWORK

- do not discount that "it's just slow"
- network hardware fails
- traceroute is your friend

APACHE

- How to invoke PHP - mod_php vs. fcgi
- Configuration settings
  - StartServers, MinSpareServers, MaxSpareServers
  - MaxClients, ServerLimit, MaxRequestsPerChild
  - KeepAlive, KeepAliveTimeout

MYSQL

- MySQL, MariaDB, Percona Server
- Configuration Settings
  - Global settings - max_connections, key_buffer, query_cache_size
  - Per-Thread settings - tmp_table_size, max_heap_table_size
  - Storage Engine settings - innodb_buffer_size, innodb_flush_log_at_trx_commit, innodb_log_file_size, innodb_log_files

MYSQL: MANAGING & MONITORING

- Managing & Monitoring
  - Percona Toolkit
  - Openark Kit
  - mysqlreport

MYSQL OPTIMIZING

- Indexes
- Joins
- Views Pitfalls
  - Order by a non-base table
  - using DISTINCT "just in case"
  - not requiring relationships
  - not using views caching (views_cache_bully)
  - using node_access (p.s. no way out)

APPLICATION (DRUPAL)

- Performance page settings (caching, aggregation)
- Logging
- Cron

WRITING AWESOME DRUPAL CODE

- Do the easy thing first (ex.g. call user_access before user_load)
- Use the Entity API (especially entity_load_multiple)
- Cache (static using drupal_static(), persistent using Cache API)

WRITING TERRIBLE DRUPAL CODE

- Abuse variable_set()
- Rely on external services
- Use $_SESSION (you cannot be served cached pages)
- Write finely-cached MySQL queries
- Ignore PHP Warnings
- Commit debug code to production
- Use development settings in production

OK, I DID ALL THAT. NOW WHAT?

- APC
- Memcache (or APC again, or Redis)
- Varnish
- CDNs
- Nginx
- Solr
- Load Testing (JMeter)
- WRITE A GROWTH PLAN

SCALING & INFRASTRUCTURE

- Horiztonal vs. Vertical Scaling
- High Availability/Failover (IP-based vs. DNS based) - Heartbeat
- Challenges: Syncing...
  - Sessions
  - Content files (simple v. complex - GlusterFS)
  - Other caches (i.e Memcache/Radis)
- "A Note About the 'Cloud'"

QUESTIONS

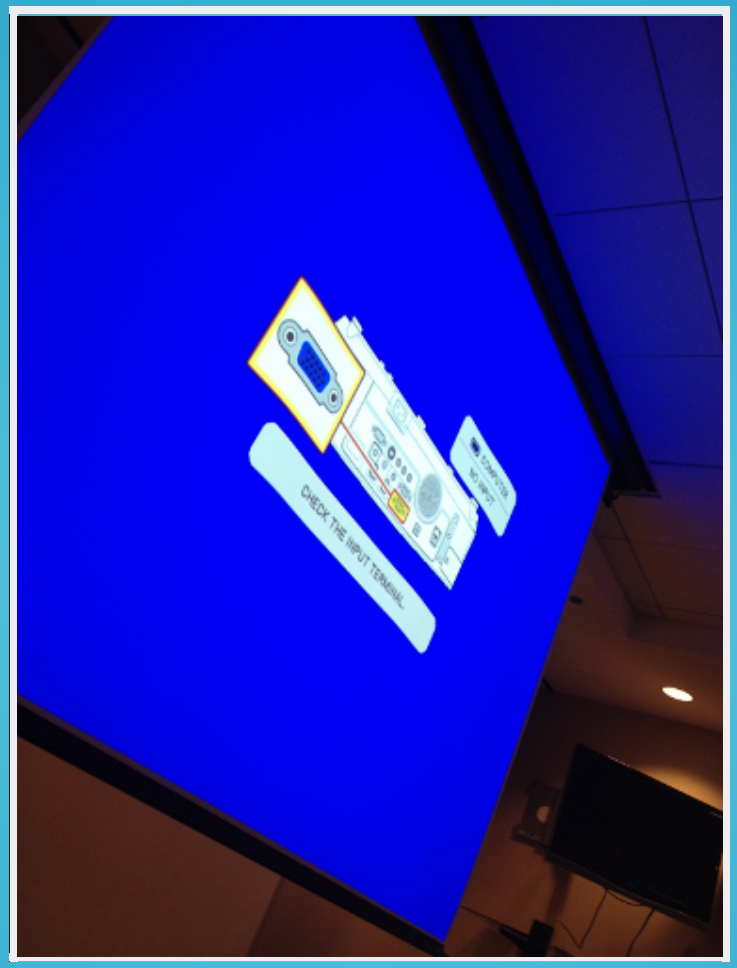

SORRY!