Beginner's Guide to Performance!

Jonathan Rowny

JonathanRowny

- Software Architect @ AboutWeb
- Speaking at RIACon next week about NodeJS

- jrowny.com
- @jrowny
- github.com/jrowny
- jrowny@aboutweb.com

We have a lot to talk about...

- Hosting Options
- Content Delivery Networks
- Drupal's Cache Mechanisms
- Opcode Caches: APC
- Memcached
- Varnish (And other Reverse Proxies)
- Static Caching
- Web Server Tweaks (ETags, gzip-compression, etc.)
- Front End (sprites, compressing images, minifying code)
- Benchmarking (are we going faster yet?)

Hosting Options

Shared	VPS/Cloud	Dedicated/CoLo
Managed Cheap!	Dedicated Resources Affordable Scales Excellent Performance	Dedicated Resources Excellent Performance
Unreliable Does not scale Limited Tuning Options Slow	Read/Write Speeds	Expensive More difficult to scale More difficult to manage

Managed/Unmanaged

Unmanaged

- You might have to wake up at 3am to fix a problem
- Security patches, updates, etc. are your responsibility
- All configuration (security and otherwise) is your responsibility
- Cheaper

Managed

- Someone else has to be awake at 3am to fix a problem
- Security patches done regularly
- You can still make it insecure, you can still break it
- Way more expensive

Managed/Unmanaged

Unmanaged

- Amazon Web Services (AWS)
- Linode
- MediaTemple (ve)
- Rackspace CloudServers

Managed

- Acquia Managed/Acquia DevCloud
- GetPantheon
- GetCadre
- MediaTemple (dv)
- Rackspace **Managed** CloudServers

Content Delivery Network (CDN)

- Host Static Content
  - Video/Music/Images
  - Static HTML/JS/CSS
  - Streaming*
- Global Network Reduces Latency
- Inexpensive (or free?)
- Requires Pressflow 6 or Drupal 7
- CNAME record
  - **media**.yourdomain.com
  - **cdn**.yourdomain.com
- CDN Module http://www.drupal.org/project/cdn
- Origin Pull/File Convey

I don't need a CDN!, my host provides **UNLIMITED BANDWIDTH!!**

Origin Pull

1. **User:** request (from USA)
   a. http://**mycdn**.mydomain.com/HugeFile.jpg
2. **CDN:** Uhh, I don't have this file, let me go ***pull*** it from the ***origin***
   a. *http://www.mydomain.com/HugeFile.jpg*
   b. Now I have the file, I'll keep it
   c. And I'll distribute it all over the world!

3. **User 2:** request (from Japan)
   a. http://**mycdn**.mydomain.com/HugeFile.jpg
4. **CDN:** I have that file! And I have a copy on a server in Japan already!

Popular CDNS

- Google
  - jquery
  - fonts
- Amazon CloudFront
- Rackspace CloudFiles
- CloudFlare (free plan)
- Akamai
- Limelight
- Azure CDN
- Internap (Voxel)
- MaxCDN

Origin Pull CDN Steup

1. Download/Install CDN Module
2. Choose CDN provider
3. Configure provider to point at your web server
4. Configure CDN Module with provider URL (or CNAME)
5. Your site is faster!

Caching in General

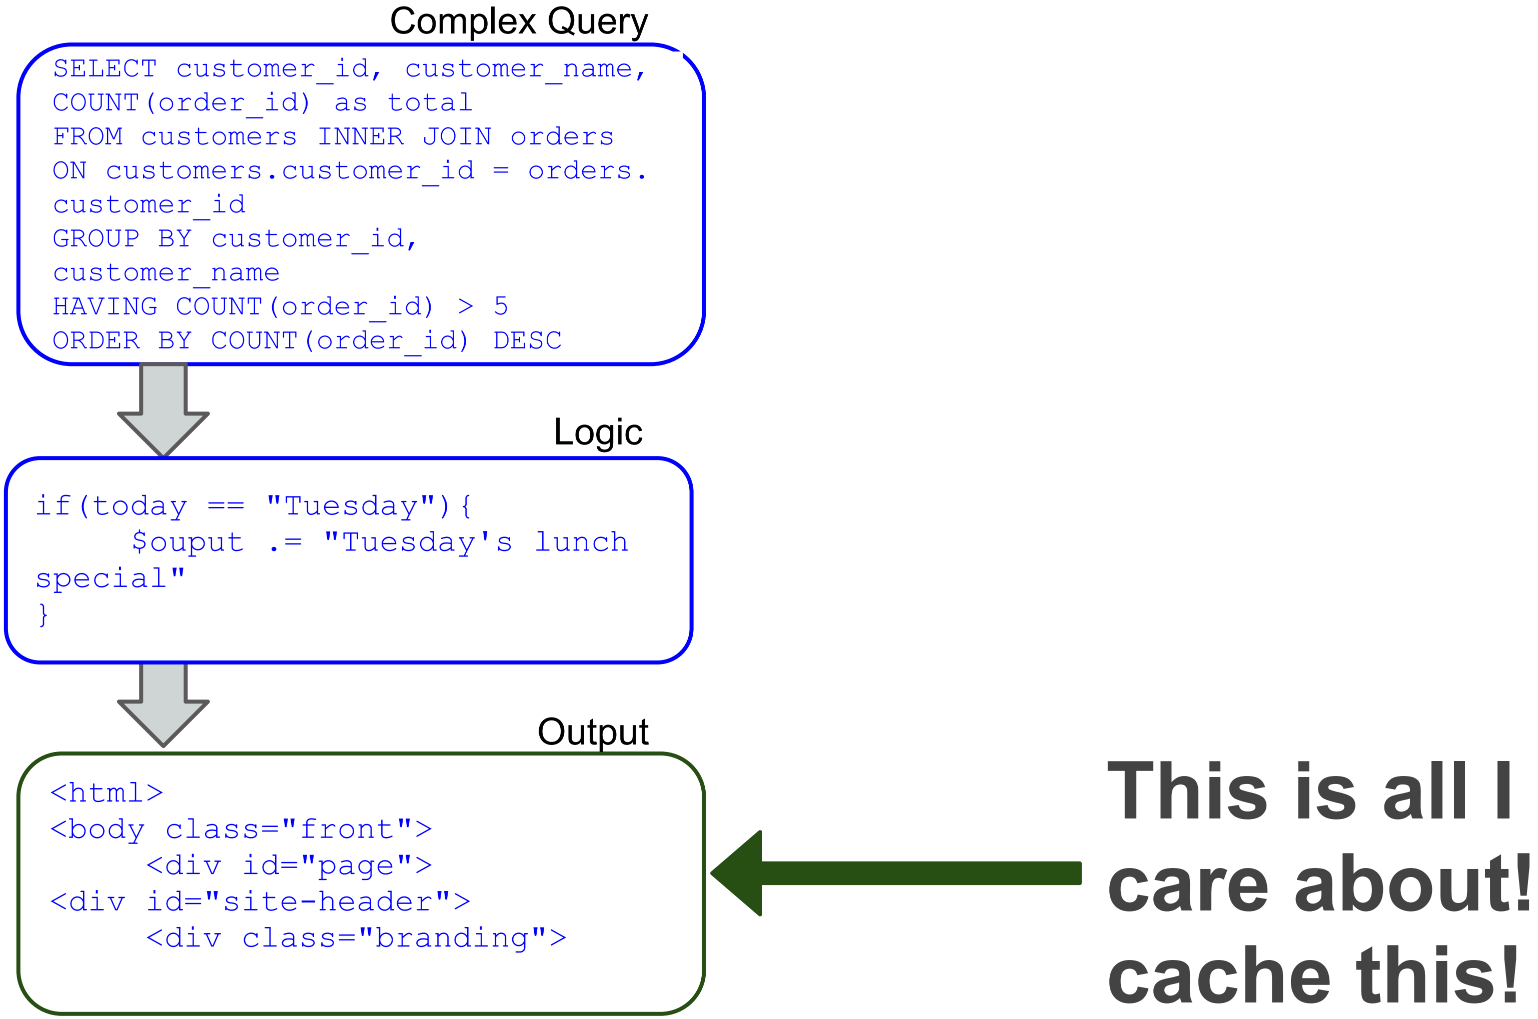

Complex Query

```
SELECT customer_id, customer_name,
COUNT(order_id) as total
FROM customers INNER JOIN orders
ON customers.customer_id = orders.
customer_id
GROUP BY customer_id,
customer_name
HAVING COUNT(order_id) > 5
ORDER BY COUNT(order_id) DESC
```

Logic

```
if(today == "Tuesday"){
     $ouput .= "Tuesday's lunch
special"
}
```

Output

```
<html>
<body class="front">
     <div id="page">
<div id="site-header">
     <div class="branding">
```

This is all I care about! cache this!

Cahing in Drupal

- Rules for caching
  - Page
  - Role
  - User
  - Global
  - Custom
- Lots and lots of caching modules!*

-
  ```
  cache_set($id, $data, $bin, $expiration);
              who   what     where  when
  ```

*http://groups.drupal.org/node/21897

PHP Accelerators

APC
eAccellerator
xCache

- Store compiled code
- Put compiled code in memory
- Can be used to cache other stuff in memory
    - Yes, there's a module for that

APC (Alternative PHP Cache)

Instantly improve drupal performance:

```
pecl install apc
```

in php.ini:

```
extension = apc.so
apc.shm_size = 48
```

source: http://2bits.com/articles/benchmarking-drupal-with-php-op-code-caches-apc-eaccelerator-and-xcache-compared.html

Memcached

- Store values in memory
- key/value
- Service (download from http://memcached.org/)
- PHP extension (memcache)
  - PECL like APC
- Drupal Module
  - Memcache
  - CacheRouter (no D7)
- Many alternatives (like APC, Redis)

Varnish!

- Caches entire pages (Static content)
- Reverse Proxy
  - Runs before a webserver
  - Apache's "body guard"
  - Because apache is kind of slow
- Drupal module for invalidating cache
- Caches to memory

BOOST! (Static Cache)

- Make an HTML copy of your website
- **Anonymous** users served plain-old-html
- Super easy to configure!

1. Download/Enable BOOST module
2. Configure
3. Copy rules into .htaccess file
4. Test

Alternative Web Servers

- Nginx
- Lighttpd
- Custom apache configs

ETags

- Tell your browser if they have the latest
- Don't work well in clustered environments
  - Set to "none"
- Instead, use Expires

gZip

- Unziping Text = faster than downloading
- Many modules do this for you!
  - Boost
  - AdvAgg

Reduce HTTP Requests

- Combine JS
  - AdvAgg/D7
- Combine CSS
  - AdvAgg/D7
- Combine Images (Sprites)

Sprites

- One big image
  - buttons
  - icons
  - logo
  - navigation
  - design elements
- NOT for content images
- Set background/position

Minify Code

- Remove whitespace
- Every little bit counts

Optimize Images

- "Save for Web" from photoshop sucks
- http://kraken.io/
- http://smush.it/
- ImageOptim
- Shrink O'Matic

Misc Tweaks

- Move js to bottom, css to top
- no Inline css/js
- don't scale images with CSS

Benchmarks

- Yahoo! ySlow
- Google Webmaster Tools
- https://developers.google.com/speed/pagespeed/insights