Computer Science E-75

Building Dynamic Websites

Harvard Extension School

http://www.cs75.net/

Lecture 12: Scalability

David J. Malan
malan@post.harvard.edu

Recommended Reading

- Building Scalable Websites
  by Henderson
- High Performance MySQL
  by Zawodny and Balling
- MySQL Clustering
  by Davis and Fisk
- Scalable Internet Architectures
  by Schlossnagle
- . . .

Vertical Scaling

- CPU
  - cores, L2 Cache, ...
- Disk
  - PATA, SATA, SAS, ...
  - RAID
- RAM
- . . .

Horizontal Scaling

Image from wikimedia.org.

PHP Acceleration

- Code Optimization
- Opcode Caching
- . . .

PHP Accelerators

- Alternative PHP Cache (APC)
  http://pecl.php.net/package/APC
- eAccelerator
  http://eaccelerator.net/
- XCache
  http://xcache.lighttpd.net/
- Zend Platform
  http://www.zend.com/en/products/platform/
- . . .

Load Balancing at Layer 4

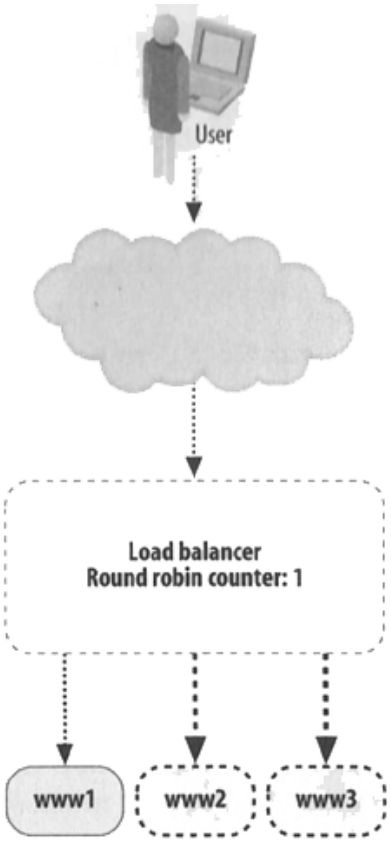

Image from *Building Scalable Websites*.

Load Balancing with BIND

```
www   IN   A   64.131.79.131
www   IN   A   64.131.79.132
www   IN   A   64.131.79.133
www   IN   A   64.131.79.134
```

Load Balancing at Layer 7

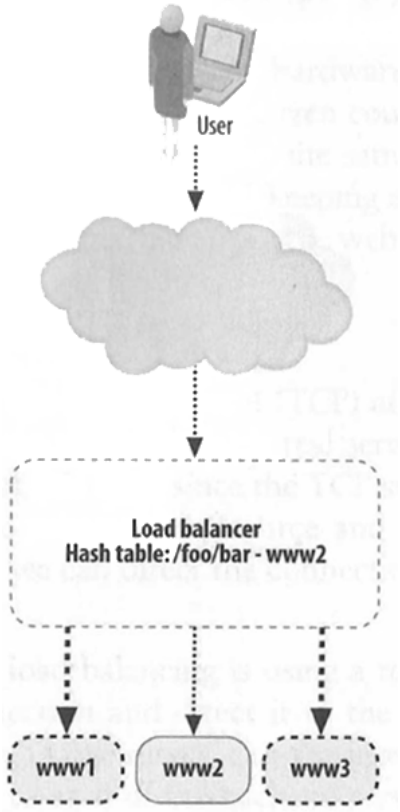

Image from *Building Scalable Websites*.

Sticky Sessions

- Layer-7 Load Balancing?
- Shared Storage?

  FC, iSCSI, NFS, *etc*.
- Cookies?

Load Balancers

- Software
  - LVS
  - Perlbal
  - Pirhana
  - Pound
  - Ultra Monkey
  - ...
- Hardware
  - Cisco
  - Citrix
  - F5
  - ...

Caching

- .html
- MySQL Query Cache
- memcached
- . . .

.html

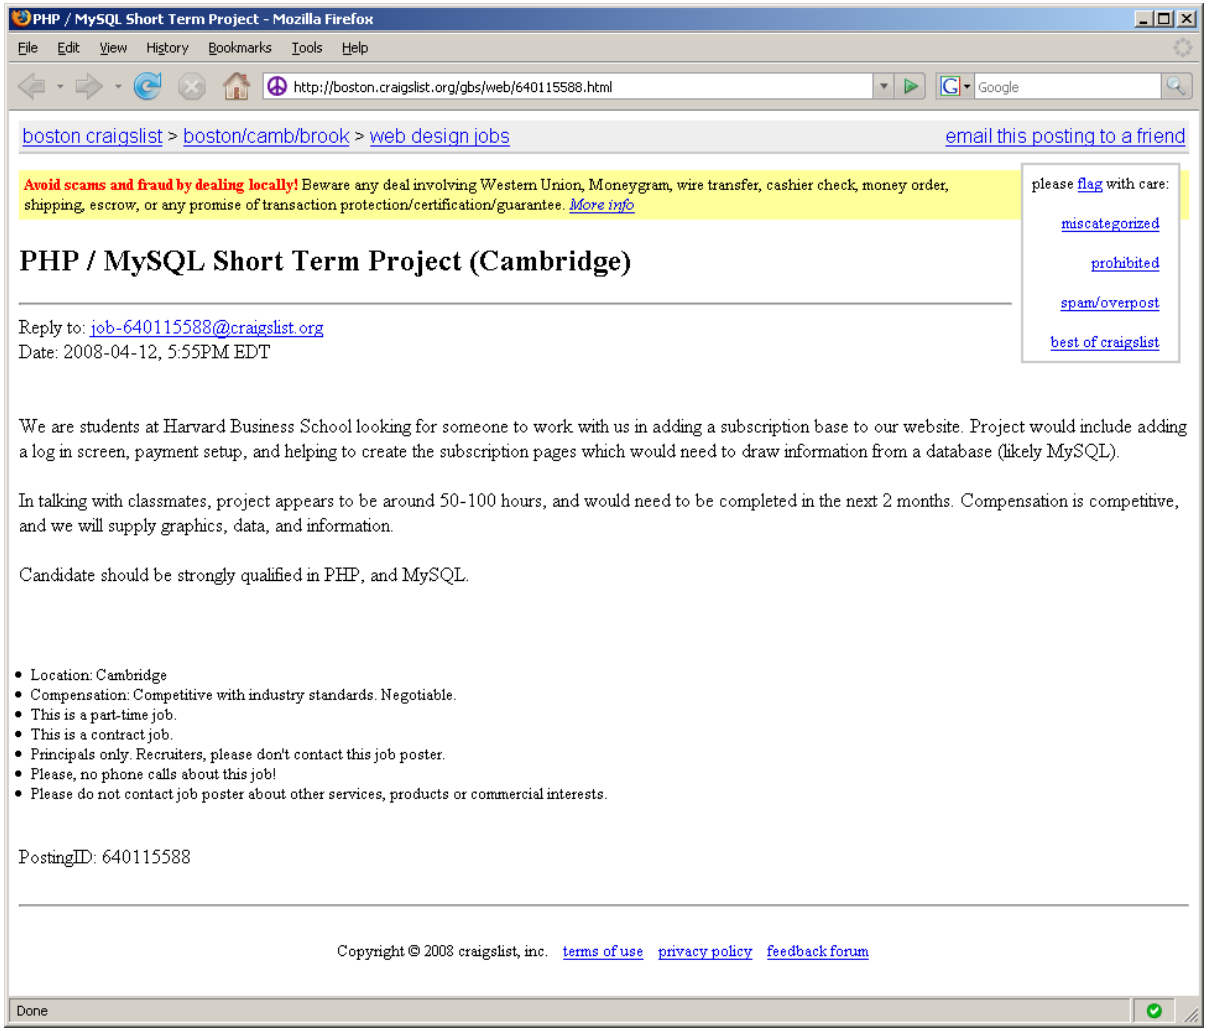
PHP / MySQL Short Term Project - Mozilla Firefox
File Edit View History Bookmarks Tools Help
http://boston.craigslist.org/gbs/web/640115588.html
Google
boston craigslist > boston/camb/brook > web design jobs
email this posting to a friend
Avoid scams and fraud by dealing locally! Beware any deal involving Western Union, Moneygram, wire transfer, cashier check, money order, shipping, escrow, or any promise of transaction protection/certification/guarantee. More info
please flag with care:
miscategorized
prohibited
spam/overpost
best of craigslist
PHP / MySQL Short Term Project (Cambridge)
Reply to: job-640115588@craigslist.org
Date: 2008-04-12, 5:55PM EDT
We are students at Harvard Business School looking for someone to work with us in adding a subscription base to our website. Project would include adding a log in screen, payment setup, and helping to create the subscription pages which would need to draw information from a database (likely MySQL).
In talking with classmates, project appears to be around 50-100 hours, and would need to be completed in the next 2 months. Compensation is competitive, and we will supply graphics, data, and information.
Candidate should be strongly qualified in PHP, and MySQL.
Location: Cambridge
Compensation: Competitive with industry standards. Negotiable.
This is a part-time job.
This is a contract job.
Principals only. Recruiters, please don't contact this job poster.
Please, no phone calls about this job!
Please do not contact job poster about other services, products or commercial interests.
PostingID: 640115588
Copyright © 2008 craigslist, inc. terms of use privacy policy feedback forum
Done

MySQL Query Cache

```
query_cache_type = 1
```

http://dev.mysql.com/doc/refman/5.0/en/query-cache.html

memcached

```
$memcache = memcache_connect(HOST, PORT);
$user = memcache_get($memcache, $id);
if (is_null($user))
{
    mysql_connect(HOST, USER, PASS);
    mysql_select_db(DB);
    $result = mysql_query("SELECT * FROM users WHERE id=$id");
    $user = mysql_fetch_object($result, User);
    memcache_set($memcache, $user->id, $user);
}
```

http://www.danga.com/memcached/

http://us2.php.net/memcache

MySQL

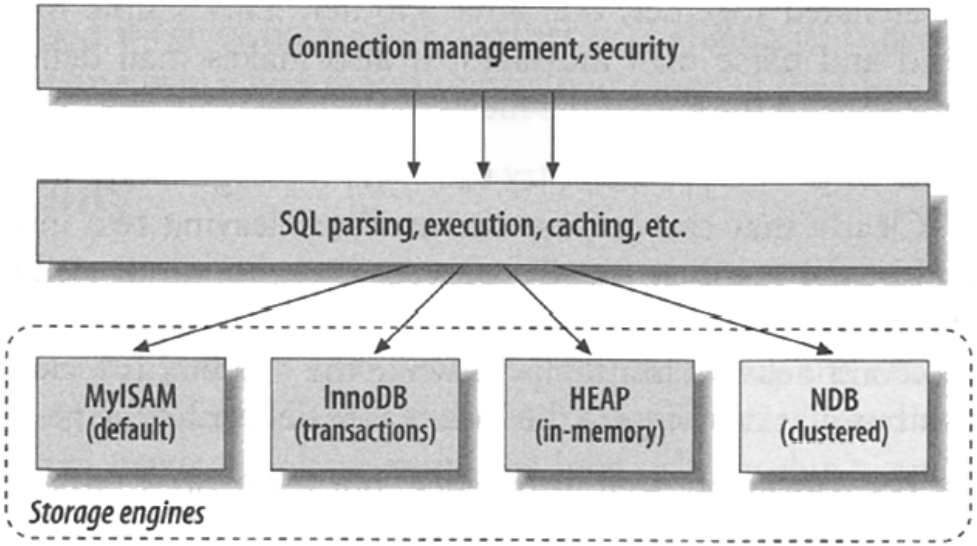

Image from *High Performance MySQL*.

MySQL

	MyISAM	InnoDB	MEMORY	NDB
Multi-statement transactions, ROLLBACK	-	X	-	X
Foreign key constraints	-	X	-	-
Locking level	table	row	table	row
BTREE indexes	X	X	-	X
FULLTEXT indexes	X	-	-	-
HASH lookups	-	X	X	X
Other in-memory tree-based index	-	-	4.1.0	-
GIS, RTREE indexes	4.1.0	-	-	-
Unicode	4.1.0	4.1.2	-	-
Merge (union views)	X	-	-	-
Compress read-only storage	X	-	-	-
Relative disk use	low	high	-	low
Relative memory use	low	high	low	high

Excerpted from http://dev.mysql.com/tech-resources/articles/storage-engine/part_3.html.

Replication: Master-Slave

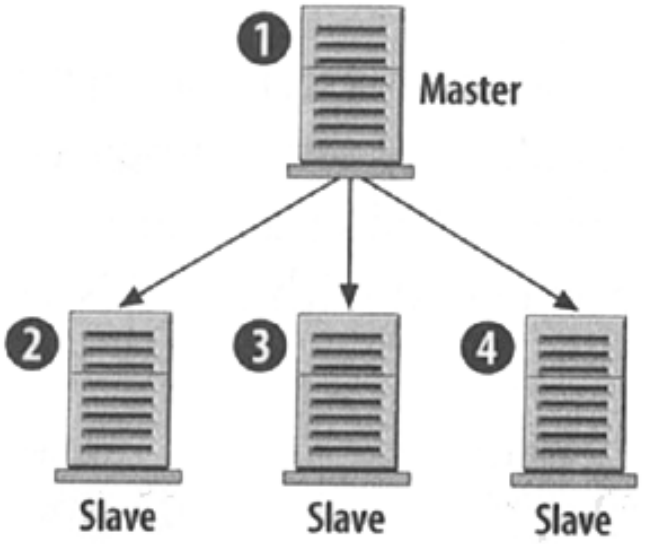

Excerpted from *High Performance MySQL*.

Replication: Master-Master

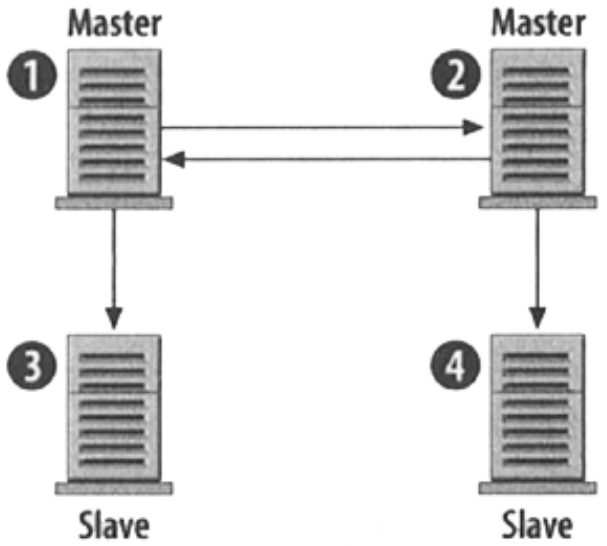

Excerpted from *High Performance MySQL*.

Load Balancing + Replication

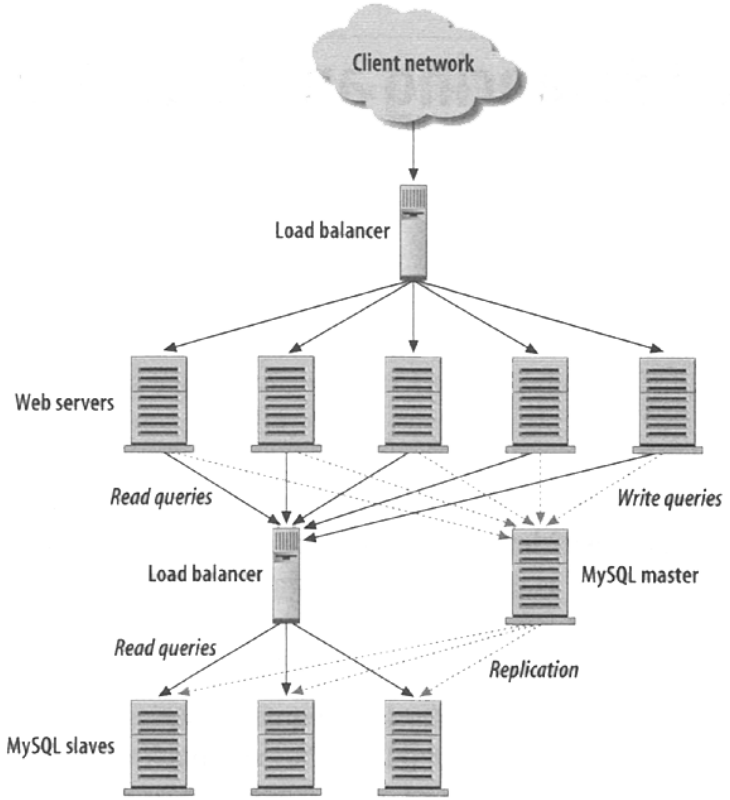

Excerpted from *High Performance MySQL*.

... + Partitioning

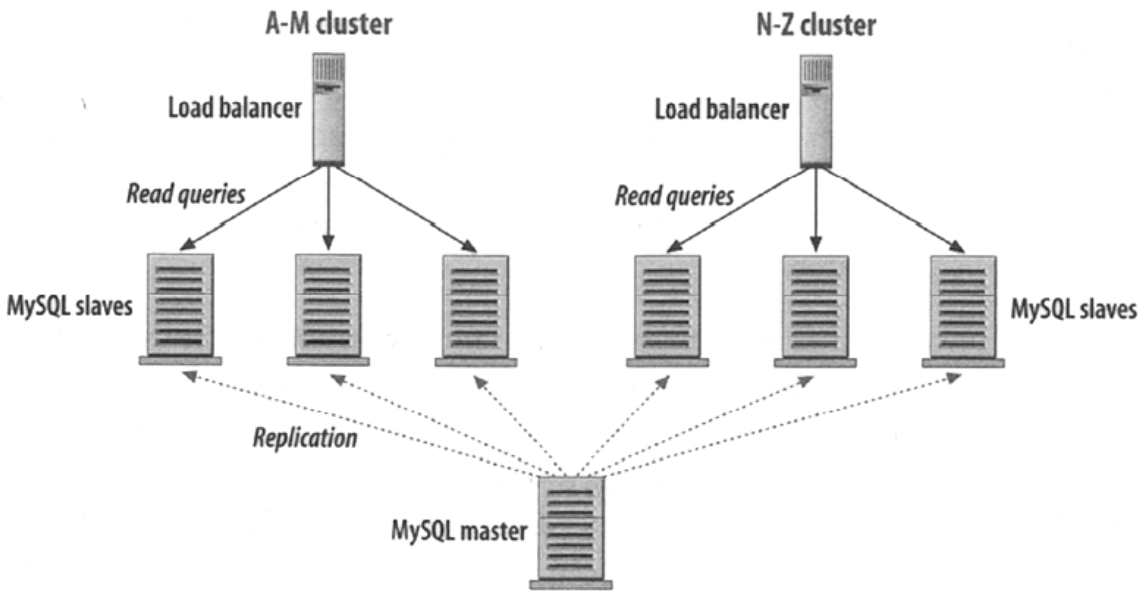

Excerpted from *High Performance MySQL*.

High Availability

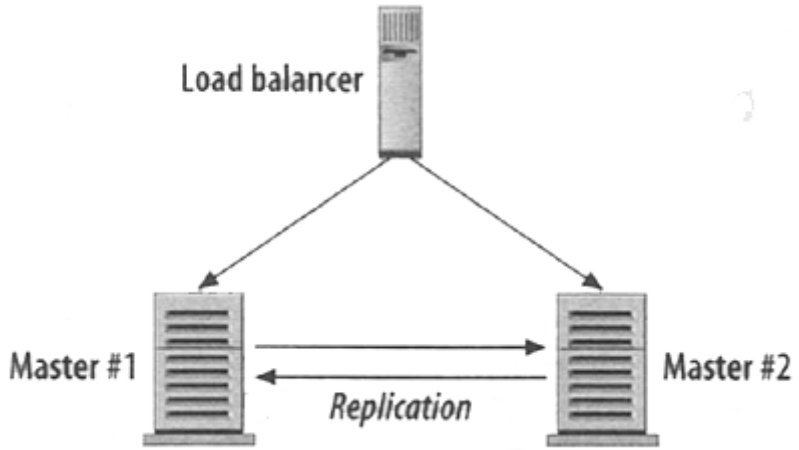

Computer Science E-75

Building Dynamic Websites

Harvard Extension School

http://www.cs75.net/

Lecture 12: Scalability

David J. Malan
malan@post.harvard.edu