



Computer Science E-75

Building Dynamic Websites

Harvard Extension School

<http://www.cs75.net/>

Lecture 9: Scalability

David J. Malan
dmalan@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 techui.com.




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



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
```


Sticky Sessions

- Shared Storage?
FC, iSCSI, MySQL, NFS, *etc.*
- Cookies?

Load Balancers

- Software

- HAProxy
- LVS
- Perlbal
- Pirhana
- Pound
- Ultra Monkey
- ...

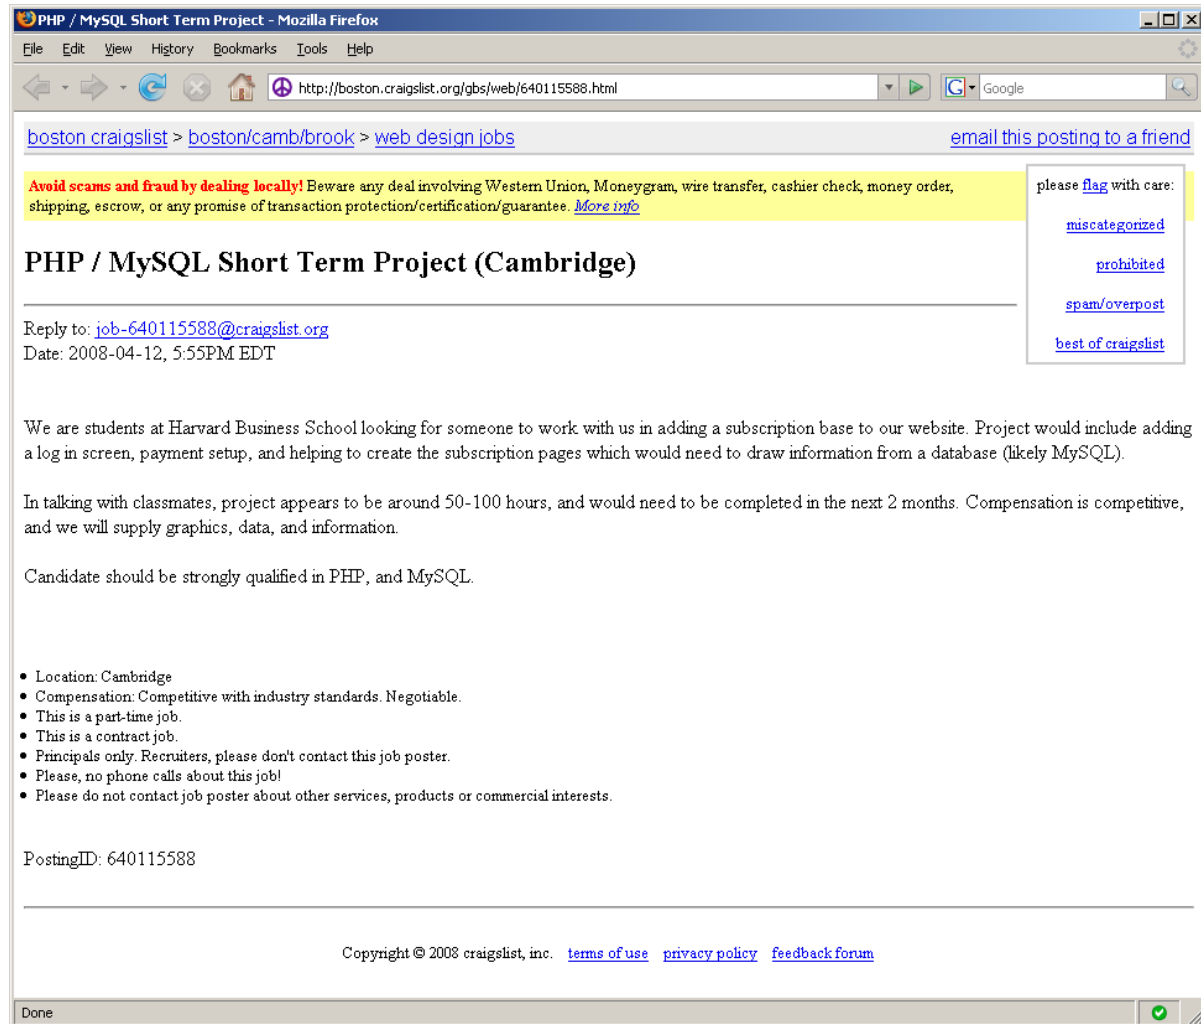
- Hardware

- Barracuda
- Cisco
- Citrix
- F5
- ...

Caching

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

.html



MySQL Query Cache

```
query_cache_type = 1
```

<http://dev.mysql.com/doc/refman/5.1/en/query-cache.html>

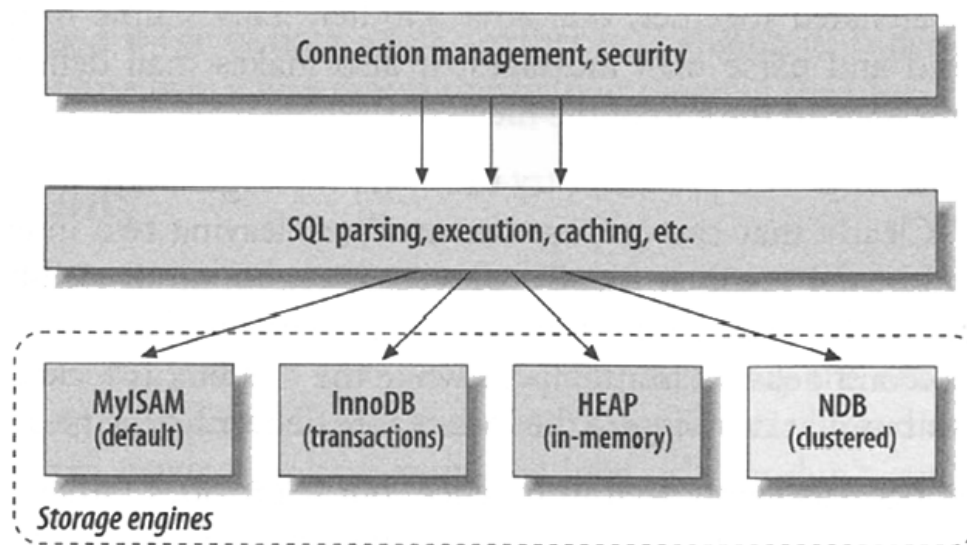
memcached

```
$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




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


Replication: Master-Slave




Replication: Master-Master




Load Balancing + Replication



... + Partitioning



High Availability





Computer Science E-75

Building Dynamic Websites

Harvard Extension School

<http://www.cs75.net/>

Lecture 9: Scalability

David J. Malan
dmalan@harvard.edu