

Computer Science E-75

Building Dynamic, Scalable Websites

Harvard Extension School
<http://www.cs75.net/>

Lecture 12: Scalability

David J. Malan
malan@post.harvard.edu

0

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
- ...

1

Vertical v. Horizontal Scaling



Image from Sainfield

2

Vertical Scaling

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

3

Horizontal Scaling



Image from wikimedia.org

4

PHP Acceleration

- Code Optimization
- Opcode Caching
- ...

5

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/>
- ...

6

Load Balancing at Layer 4




Image from Building Scalable Websites

7

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
```

8

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



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




MySQL

	MyISAM	InnoDB	MEMORY	NDB
Multi-statement transactions, ROLLBACK	-	X	-	X
Foreign key constraints	-	X	-	-
Locking level	table	row	table	row
B+TREE indexes	X	X	-	X
FULLTEXT indexes	X	-	-	-
HASH lookups	-	X	X	X
Other in-memory tree-based index	-	-	4.1.0	-
GIS, B+TREE indexes	-	-	-	-
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.

18


Replication: Master-Master



Excerpted from High Performance MySQL.

19


Load Balancing + Replication



Excerpted from High Performance MySQL.

20


... + Partitioning



Excerpted from High Performance MySQL.

21


High Availability



Excerpted from High Performance MySQL.

22

MySQL Cluster



Excerpted from <http://www.mysql.com/news-and-events/on-demand-webinars/cluster-20061116.pdf>.

23

Computer Science E-75

Building Dynamic, Scalable Websites

Harvard Extension School
<http://www.cs75.net/>

Lecture 12: Scalability

David J. Malan
malan@post.harvard.edu

24