

Achieving Optimal MySQL Performance for Drupal



Overview of MySQL



- Founded in 1995
- Operations in 25 countries and growing
- World's Most Popular Open Source Database
- Up to 60,000 downloads/day
- Reduces TCO of database management

Reliable, High Performance, Easy to Use

- Embedded by over 100 companies



- Hundred's of Partners and Service Integrators



January 16th 2008



Sun to Acquire MySQL



[Read more »](#)

Powering the modern online world



Web and Web 2.0



SaaS &
Managed Hosting Providers



Telecom & Communications



Enterprise 2.0



OEM & Embedded



- www.observer.com www.theonion.com drupal.org
www.britney.com www.jenniferlopez.com dead.net
www.nowpublic.com officepranks.forbes.com
spreadfirefox.com www.amnesty.org kerneltrap.org
www.relayforlife.org www.mtv.co.uk www.eff.org
airamerica.com www.ubuntu.com greenparty.ca
warnerbrosrecords.com musicbox.sonybmg.com
standagainstpoverity.org bigideas.berkeley.edu
letsgreenthiscity.com www.conservationfund.org
research.yahoo.com dev.aol.com osuosl.org
teamsugar.com www.echoinggreen.org typophile.com





MySQL

- MySQL websites, estimated monthly page views:
 - drupal.org: 18 million
 - slashdot.org: 50 million
 - youtube.com: 75 million
 - digg.com: 200 million
 - livejournal.com: 475 million
 - wikipedia.org: 7 billion
 - www.facebook.com: 65 billion

Achieving Optimal MySQL Performance for Drupal

- Monitoring MySQL, Finding Performance Bottlenecks
- Choosing a MySQL Storage Engine
- Optimizing MySQL for Drupal

About The Presenter

- Name: Jeremy Andrews
- Started KernelTrap.org in 2001. 
- Began working with Drupal in January of 2002. 
- Wrote core statistics module, throttle module, pager.inc, various patches.
- Wrote dozens of contrib modules.
- Founded Tag1 Consulting in October of 2007. 
- Founded the Ad Bard Network in January of 2008. 

Monitoring MySQL

- SHOW FULL PROCESSLIST
- SHOW GLOBAL STATUS
- mytop, innotop
- mysqlreport
- top M, vmstat, iostat, sar, ...
- slow query log, no index log

SHOW FULL PROCESSLIST

- Need PROCESS privilege to see info on all threads
- FULL shows you more than the first 100 characters
- KILL to end processes or query
- Equivalent: mysqladmin processlist

```
mysql> SHOW FULL PROCESSLIST;
```

```
+-----+-----+-----+-----+-----+-----+-----+-----+
| Id    | User | Host      | db    | Command | Time | State | Info                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 4340  | root | localhost | NULL  | Query   | 0    | NULL  | SHOW FULL PROCESSLIST            |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

SHOW GLOBAL STATUS

- GLOBAL aggregates all connections
- Can filter with LIKE '%name%'
- Equivalent: mysqladmin extended-status

```
mysql> SHOW GLOBAL STATUS LIKE 'Qcache%';
+-----+-----+
| Variable_name          | Value          |
+-----+-----+
| Qcache_free_blocks    | 1              |
| Qcache_free_memory    | 0              |
| Qcache_hits            | 2404827826    |
| Qcache_inserts        | 621290216     |
| Qcache_lowmem_prunes  | 171912228     |
| Qcache_not_cached     | 58710443      |
| Qcache_queries_in_cache | 0              |
| Qcache_total_blocks   | 0              |
+-----+-----+
8 rows in set (0.04 sec)
```

mytop

- console perl MySQL tool similar to *top*
- t: thread view
- c: command view

```
MySQL on kerneltrap.db.osuosl.org (5.0.46-log) up 51+07:11:59 [02:17:34]
Queries: 420.1M qps: 99 Slow: 0.0 Se/In/Up/De(%): 34/00/00/00
          qps now: 312 Slow qps: 0.0 Threads: 31 ( 2/ 89) 66/00/00/00
Key Efficiency: 97.7% Bps in/out: 0.0/ 0.0 Now in/out: 8.3/ 1.3k
```

Id	User	Host/IP	DB	Time	Cmd Query or State
11795234	kerneltra	kt1	test	0	Query show full processlist
11795646	kerneltra	kt1	kerneltrap	0	Query SELECT MIN(mid) AS mi
11795648	kerneltra	kt1	kerneltrap	0	Sleep

innotop

- inspired by mytop, much more powerful
- Originally for InnoDB, now does much more
- Server groups: monitor multiple MySQL servers
- 13 different modes, derived from SHOW VARIABLES, SHOW GLOBAL STATUS, SHOW FULL PROCESSLIST, SHOW ENGINE INNODB STATUS, SHOW MASTER STATUS, SHOW SLAVE STATUS, SHOW OPEN TABLES, EXPLAIN, and more.
- <http://innotop.sourceforge.net/>

mysqlreport

- perl script
- Report derived from SHOW STATUS values
- earlier versions required --all for complete report
- --email and cron for automated reports
- <http://hackmysql.com/mysqlreport>

SysAdmin 101

- `top M` – sort by memory usage
- `vmstat 1 10` – virtual memory statistics
- `sar` – collecting historical data

Query Logs

- Slow Query Logs
 - my.cnf: `log-slow-queries=/var/log/mysql-slow.log`
 - my.cnf: `long_query_time=1`
- No Index Logs
 - my.cnf: `log_queries_not_using_indexes=TRUE`
- `mysqlsa`
 - Trim log as necessary
 - Summarizes logs
 - <http://hackmysql.com/mysqlsa>

MySQL Storage Engines

- Underlying component for creating, retrieving, updating and deleting data.
- Each has own tunings, may not be optimal to mix.
- 5.1 introduces dynamically loaded storage engines.
- MyISAM, InnoDB, Falcon, Maria

MyISAM Storage Engine

- Default storage engine since MySQL 3.23.
- Based on original ISAM storage engine.
- Tables stored in three simple files:
 - table.frm (format)
 - table.MYD (data)
 - table.MYI (index)

MyISAM Limitations

- Not ACID compliant, no transactions.
- No foreign key support.
- Table level locking, no row level locking support.
- File corruption.
- Relies on OS for caching reads and writes.

InnoDB Storage Engine

- Introduced in MySQL 3.23.
- ACID-compliant (except DDL), supports transactions
- Foreign key constraints.
- Row level locking.
- Automatic crash recovery.
- Owned by Oracle since October 2005

InnoDB Limitations

- ACID compliance requires high-end IO hardware
- Tables take 2-3x more space than MyISAM tables
- No full text search support
- No compression support
- Drupal not optimized for InnoDB

Falcon Storage Engine

- MySQL 6.0
- Based on the Netrastructure database engine.
- Optimized for multi-core and 64-bit architectures.
- Early benchmarks show good and bad.
- http://forge.mysql.com/wiki/Falcon_Feature_Preview.html
- <http://mysql.com/mysql60/>

Maria Storage Engine

- Announced January 27'th, 2008
- Michael "Monty" Wideneus, MySQL founder
"Maria 1.0 is basically a crash-safe non-transactional version of MyISAM."
http://forge.mysql.com/wiki/Maria_Docs
- MySQL 5.1, (goal: stable, not high performing)
- Possible replacement for MyISAM in MySQL 6.0
- Eventually to include:
 - ACID, MVCC
 - Commit/Rollback, group commits
 - Row locking, Concurrent Selects

Optimizing MySQL

- Tuning with mysqlreport
- Fixing slow queries / EXPLAIN
- devel module
- Replication

Tuning with mysqlreport (1)

- Header:

```
MySQL 5.0.46-log      uptime 52 7:26:48      Mon Jan 28 02:32:23 2008
```

- version, uptime

- Key (MyISAM)

```
__ Key _____  
Buffer used      82.27M of 100.00M  %Used:  82.27  
  Current        54.70M           %Usage:  54.70  
Write hit        84.27%  
Read hit         97.81%
```

- Buffer used, high water mark
- Current, actual usage
- Write / Read hits rates are ratio of hard drive : RAM

Tuning with mysqlreport (2)

- Questions: SQL queries & MySQL protocol

Questions			
Total	471.35M	104.3/s	
DMS	236.42M	52.3/s	%Total: 50.16
QC Hits	169.87M	37.6/s	36.04
Com_	53.43M	11.8/s	11.34

- DMS, Data Manipulation Statements (includes: SELECT, INSERT, UPDATE, DELETE)
- QC Hits, query cache
- Com, MySQL commands & protocol
- Below, break down of each

Tuning with mysqlreport (3)

- SELECT and sort

SELECT and Sort				
Scan	5.09M	1.1/s	%SELECT:	2.68
Range	15.64M	3.5/s		8.22
Full join	40.96k	0.0/s		0.02
Range check	6.33k	0.0/s		0.00
Full rng join	32.48k	0.0/s		0.02
Sort scan	9.94M	2.2/s		
Sort range	13.19M	2.9/s		
Sort mrg pass	109.82k	0.0/s		

- Scan: entire table
- Full join: multiple full tables (tunable: join_buffer_size *)
- Sorts: monitor "SHOW STATUS LIKE Sort_merge_passes"
(tunable: sort_buffer_size *)
- * per-connection memory allocations

Tuning with mysqlreport (4)

- Query cache

```
__ Query Cache
Memory usage 13.37M of 20.00M %Used: 66.84
Block Fragmnt 13.30%
Hits 169.87M 37.6/s
Inserts 183.65M 40.6/s
Insrt:Prune 1.59:1 15.0/s
Hit:Insert 0.92:1
```

- If too big (64M+), can cause MySQL server freezes
- Fragmented? Increase memory (`query_cache_size`), or decrease tunable: `query_cache_min_res_unit` (2k is a good place to start with Drupal)
- Review Insert:Prune and Hit:Insert
- Can control QC with `SQL_NO_CACHE` and `SQL_CACHE`

Tuning with mysqlreport (5)

- Tables

```
__ Tables _____
Open          4048 of 4048   %Cache: 100.00
Opened        573.88k     0.1/s
```

- Same table can be opened by many threads
- More than 1/s? Increase tunable: `table_cache`
- Monitor memory usage closely when increasing cache

- Connections

```
__ Connections _____
Max used      64 of 100   %Max: 64
Total         12.22M  2.7/s
```

- Default of 100 is generally sane, watch #/s
- Needing more is often sane of poor tuning

Tuning with mysqlreport (6)

- Created Temp

Created Temp			
Disk table	2.10M	0.5/s	
Table	12.27M	2.7/s	Size: 100.0M
File	146.59k	0.0/s	

- Want most as table (RAM)
- Tunables: tmp_table_size, max_heap_table_size (*per-connection memory allocations*)
- Queries with blobs and large text fields won't fit in RAM

- Threads

- Tunable: thread_cache_size

Threads			
Running	1 of	29	
Cached	90 of	100	%Hit: 99.95
Created	6.20k	0.0/s	

Tuning with mysqlreport (7)

- InnoDB Buffer Pool

```
___ InnoDB Buffer Pool _____
Usage          3.00G of  3.00G  %Used: 100.00
Read hit       99.99%
Pages
  Free         1          %Total:  0.00
  Data        185.02k     94.11 %Drty:  0.28
  Misc        11588      5.89
  Latched     0          0.00
Reads          163.33G   36.1k/s
  From file   14.27M        3.2/s          0.01
  Ahead Rnd   868753       0.2/s
  Ahead Sql   634949       0.1/s
Writes        670.00M   148.2/s
Flushes       16.82M    3.7/s
```

- Rule of thumb: 70% of available memory
- Increase: over 80% used, lower than .1 read ratio, significant reads from file
- Tunable: innodb_buffer_pool_size

Tuning with mysqlreport (8)

- InnoDB Data, Pages, Rows

```
___ InnoDB Data, Pages, Rows ___  
Data  
Reads      21.33M    4.7/s  
Writes     12.74M    2.8/s  
fsync      3.60M     0.8/s
```

- Ratio of Writes to fsync?
- InnoDB defaults to ACID
- Can you afford data loss?
- Tunable: `innodb_flush_log_at_trx_commit` (1 = flush every write, 0 = flush every second)

Fixing Slow Queries

- Pick lowest hanging fruit first
- Review mysqlsla top 10
- EXPLAIN [EXTENDED] (SHOW WARNINGS)
- SHOW INDEX, ANALYZE TABLE
- Simplify...

Drupal devel module

Executed 68 queries in 12.4 milliseconds. Queries taking longer than 5 ms and queries executed more than once, are **highlighted**. Page execution time was 204.54 ms.

ms	#	where	query
3.74	1	cache_get	SELECT data, created, headers, expire FROM cache_filter WHERE ci
0.68	1	cache_get	SELECT data, created, headers, expire FROM cache_menu WHERE c
0.46	1	node_tag_new	UPDATE history SET timestamp = 1201572501 WHERE uid = 1 AND
0.42	1	node_load	SELECT n.nid, n.vid, n.type, n.status, n.created, n.changed, n.comm n.nid = 11
0.32	2	node_load	SELECT n.nid, n.vid, n.type, n.status, n.created, n.changed, n.comm n.nid = '11'

- Time in database versus total execution time.
- Slow queries, duplicate queries.
- <http://drupal.org/project/devel>

Replication

- Cons:
 - Slows down performance ~1% or more
 - Best effort, no built-in validation
- Pros:
 - Backups from slave
 - Master = read/write, Slave = read
 - Slaves can use different storage engine, tuning
 - High availability

Questions?

Fine Print

- This presentation was prepared by Jeremy Andrews of Tag1 Consulting (jeremy@tag1consulting.com) and is provided under the Creative Commons Attribution-ShareAlike2.0 license.

<http://tag1consulting.com/drupal/performance>

- These slide were presented on January 31st, 2008.
- Thanks to MySQL AB for hosting the webinar!