

# How to Stop Hating MySQL: Fixing Common Mistakes and Myths

Sheeri K. Cabral  
Database Team Lead  
The Pythian Group, [www.pythian.com](http://www.pythian.com)

[cabral@pythian.com](mailto:cabral@pythian.com)  
LISA 2008



*your database maestros*



# Who I Am

- MySQL DBA
- MySQL User Group
- MySQL Podcast (OurSQL, on hiatus), videos (technocation.org)
- Lots of community stuff



# Myths about MySQL

- Uses too much memory
- Slow
- I need more features!



# More Myths about MySQL

- Don't use ENUM
- Schema changes take forever
- You have to restart to log
- No partitioning



[www.pythian.com/blogs/1168/why-you-want-to-switch-to-mysql-51](http://www.pythian.com/blogs/1168/why-you-want-to-switch-to-mysql-51)



# MySQL uses too much memory!

- INFORMATION\_SCHEMA.GLOBAL\_VARIABLES (5.1)
- SHOW GLOBAL VARIABLES (5.0)
- LIKE
  - '%cache%'
  - '%buffer%'

# Common Mistake: Wasting Memory

- `bdb_cache_size`
- `bdb_log_buffer_size`
- 32-bit operating system



# Myth: Query Cache Is Faster!

- Use memcached for caching common queries
- query\_cache\_type
  - 1 or ON
  - 2 or DEMAND – SQL\_CACHE
  - SQL\_NO\_CACHE
- Default query\_cache\_size=0





# MyISAM Index Cache

- `key_buffer_size` – globally allocated on startup
- How much is being used?
  - STATUS variables
  - `Key_blocks_unused` vs. `Key_blocks_used`
  - $(\text{Key\_blocks\_}\% * \text{key\_cache\_block\_size}) / \text{key\_buffer\_size}$

# InnoDB Buffer Pool

- `innodb_buffer_pool_size` – globally allocated on startup
- How much is being used?
  - STATUS variables
  - `Innodb_buffer_pool_pages_free`
  - `Innodb_buffer_pool_pages_total`



# Allocated Per-thread

- `binlog_cache_size` = 32 Kb
- `max_binlog_cache_size` = 4 Gb
- `net_buffer_length` = 16 Kb
  - 1 Mb max
  - `max_allowed_packet`



# Per-Thread, Allocated As Needed

- `join_buffer_size` – memory buffer for joins not using indexes
- `read_buffer_size` – memory buffer for sequential table scans
- `read_rnd_buffer_size` – memory buffer for random table seeks



# Per-Thread, Allocated As Needed

- `preload_buffer_size` – when pre-loading indexes
- `sort_buffer_size` – for sorting
- `mysam_sort_buffer_size`
  - Index sorting only – REPAIR, OPTIMIZE, creating indexes



# Temporary Tables

- tmp\_table\_size
- max\_heap\_table\_size
- Large rows, BLOBs written to disk



# Performance Caches

- Global Variables
  - table\_cache
    - STATUS variables Opened\_tables
  - thread\_cache\_size
    - STATUS variables Connections and Threads\_created



# Size Matters

- Larger/fragmented data/indexes use more memory
- Larger/fragmented data/indexes take more time to search
- Clustered indexes
- `innodb_file_per_table`





# Size Matters – Data Diet

- OPTIMIZE – how often?
- Purge/archive regularly
- Follow large deletes with OPTIMIZE



# Size Matters – Schema Diet

- IP addresses
- TIMESTAMP vs DATETIME
- Strings vs. Numbers vs. ENUM/SET



# Myth: MySQL is Slow

- Memory/disk tradeoffs
- RAID
- Temporary tables
  - Created\_tmp\_tables
  - Created\_tmp\_disk\_tables



# Disk I/O

- Many data changes
  - Binary logs
  - Data
  - Index
- INSERT DELAYED
- Batch update/deletes



# Network

- Large queries vs. small
  - CPU compute time
  - Network traffic
  - Large data sets
- INSERT....ON DUPLICATE KEY UPDATE



# Size Matters, So Does Performance

- Choosing Correct Table Types:
  - MyISAM
  - InnoDB
  - BLACKHOLE
  - MERGE
  - MEMORY
  - ARCHIVE



# Query Optimization

- Use EXPLAIN

<http://dev.mysql.com/doc/refman/5.1/en/using-explain.html>

[www.pythian.com/blogs/wp-content/uploads/explain-diagram.pdf](http://www.pythian.com/blogs/wp-content/uploads/explain-diagram.pdf)



# Know What You're Doing

- Subqueries
- VIEW
- TRIGGER





# Know What You're Doing

- LIMIT
- Stored Procedures
  - compiled per thread
- Indexes
  - Selectivity
  - Functions
  - Overhead vs. utility



# More Features: Do Not Want!

- Look into Drizzle, a MySQL fork
- <http://www.drizzleproject.org>



# Feedback

- Questions?
- Comments?
- Why else do you hate MySQL?

[cabral@pythian.com](mailto:cabral@pythian.com)



*your database maestros*



# How to Stop Hating MySQL: Fixing Common Mistakes and Myths

Sheeri K. Cabral  
Database Team Lead  
The Pythian Group, [www.pythian.com](http://www.pythian.com)

[cabral@pythian.com](mailto:cabral@pythian.com)  
LISA 2008



*your database maestros*

# Who I Am

- MySQL DBA
- MySQL User Group
- MySQL Podcast (OurSQL, on hiatus), videos (technocation.org)
- Lots of community stuff



# Myths about MySQL

- Uses too much memory
- Slow
- I need more features!



# More Myths about MySQL

- Don't use ENUM
- Schema changes take forever
- You have to restart to log
- No partitioning



[www.pythian.com/blogs/1168/why-you-want-to-switch-to-mysql-51](http://www.pythian.com/blogs/1168/why-you-want-to-switch-to-mysql-51)





# MySQL uses too much memory!

- INFORMATION\_SCHEMA.GLOBAL\_VARIABLES (5.1)
- SHOW GLOBAL VARIABLES (5.0)
- LIKE
  - '%cache%'
  - '%buffer%'



# Common Mistake: Wasting Memory

- `bdb_cache_size`
- `bdb_log_buffer_size`
- 32-bit operating system



# Myth: Query Cache Is Faster!

- Use memcached for caching common queries
- query\_cache\_type
  - 1 or ON
  - 2 or DEMAND – SQL\_CACHE
  - SQL\_NO\_CACHE
- Default query\_cache\_size=0



# MyISAM Index Cache

- `key_buffer_size` – globally allocated on startup
- How much is being used?
  - STATUS variables
  - `Key_blocks_unused` vs. `Key_blocks_used`
  - $(\text{Key\_blocks\_}\% * \text{key\_cache\_block\_size}) / \text{key\_buffer\_size}$



*your database maestros*

# InnoDB Buffer Pool

- `innodb_buffer_pool_size` – globally allocated on startup
- How much is being used?
  - STATUS variables
  - `InnoDB_buffer_pool_pages_free`
  - `InnoDB_buffer_pool_pages_total`



# Allocated Per-thread

- `binlog_cache_size` = 32 Kb
- `max_binlog_cache_size` = 4 Gb
- `net_buffer_length` = 16 Kb
  - 1 Mb max
  - `max_allowed_packet`



## Per-Thread, Allocated As Needed

- `join_buffer_size` – memory buffer for joins not using indexes
- `read_buffer_size` – memory buffer for sequential table scans
- `read_rnd_buffer_size` – memory buffer for random table seeks



# Per-Thread, Allocated As Needed

- `preload_buffer_size` – when pre-loading indexes
- `sort_buffer_size` – for sorting
- `mysam_sort_buffer_size`
  - Index sorting only – REPAIR, OPTIMIZE, creating indexes





# Temporary Tables

- `tmp_table_size`
- `max_heap_table_size`
- Large rows, BLOBs written to disk



# Performance Caches

- Global Variables
  - table\_cache
    - STATUS variables Opened\_tables
  - thread\_cache\_size
    - STATUS variables Connections and Threads\_created



# Size Matters

- Larger/fragmented data/indexes use more memory
- Larger/fragmented data/indexes take more time to search
- Clustered indexes
- innodb\_file\_per\_table



# Size Matters – Data Diet

- OPTIMIZE – how often?
- Purge/archive regularly
- Follow large deletes with OPTIMIZE



# Size Matters – Schema Diet

- IP addresses
- TIMESTAMP vs DATETIME
- Strings vs. Numbers vs. ENUM/SET



# Myth: MySQL is Slow

- Memory/disk tradeoffs
- RAID
- Temporary tables
  - Created\_tmp\_tables
  - Created\_tmp\_disk\_tables



# Disk I/O

- Many data changes
  - Binary logs
  - Data
  - Index
- INSERT DELAYED
- Batch update/deletes



# Network

- Large queries vs. small
  - CPU compute time
  - Network traffic
  - Large data sets
- INSERT....ON DUPLICATE KEY UPDATE





# Size Matters, So Does Performance

- Choosing Correct Table Types:
  - MyISAM
  - InnoDB
  - BLACKHOLE
  - MERGE
  - MEMORY
  - ARCHIVE



# Query Optimization

- Use EXPLAIN

<http://dev.mysql.com/doc/refman/5.1/en/using-explain.html>

[www.pythian.com/blogs/wp-content/uploads/explain-diagram.pdf](http://www.pythian.com/blogs/wp-content/uploads/explain-diagram.pdf)



# Know What You're Doing

- Subqueries
- VIEW
- TRIGGER



# Know What You're Doing

- LIMIT
- Stored Procedures
  - compiled per thread
- Indexes
  - Selectivity
  - Functions
  - Overhead vs. utility



# More Features: Do Not Want!

- Look into Drizzle, a MySQL fork
- <http://www.drizzleproject.org>



# Feedback

- Questions?
- Comments?
- Why else do you hate MySQL?

[cabral@pythian.com](mailto:cabral@pythian.com)

