

EffectiveMySQL.com

Its all about Performance and Scalability

Common MySQL Scalability Mistakes

Ronald Bradford
<http://ronaldbradford.com>

2011.04

EffectiveMySQL.com - Its all about Performance and Scalability

AUTHOR

- 2010 - Oracle ACE Director (first in MySQL)
- 2009 - MySQL community member of the year
- Co Author of Expert PHP & MySQL
- Top individual blog contributor to Planet MySQL
- 22 years of RDBMS experience, 12 years with MySQL
 - MySQL Inc (2006-2008), Oracle Corp (96-99)
- Provide independent consulting/ Available NOW



EffectiveMySQL.com - Its all about Performance and Scalability

EffectiveMySQL.com - Its all about Performance and Scalability

PROBLEM

My website seems to freeze or responds randomly?

CAUSE

The default MyISAM storage engine uses exclusive table locks for DML.

SOLUTION

- Optimize blocking query performance
- Use a transactional engine with MVCC and row based locking to address the LOCKING issue

EXAMPLE

End user report that selects all customer, order, order lines and order history data and performs poor joins. This takes shared read locks blocking future write locks then future reads.

WHY

MySQL is unique in that it offers different mechanisms for storing and retrieving data, each with strengths and weaknesses.

The DEFAULT is not always the best.

HOW

MySQL PROCESSLIST

- Blocked have State = Locked
- Blocker - Same table, larger Time

```
mysql> SHOW PROCESSLIST;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Id | User | Host | db | Command | Time | State | Info |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 13 | apl | localhost | odtug | Query | 144 | User sleep | UPDATE emp ... |
| 14 | apl | localhost | odtug | Query | 116 | Locked | select c from emp |
| 15 | apl | localhost | odtug | Query | 89 | Locked | select c from emp |
```

HOW

Optimize Blocker

- Indexes
- Limit query
- Summary table
- Change storage engine



PROBLEM

Why is my database so large?

SOLUTION

Don't store large static objects in the database

EXAMPLE

- 80% of data is email content/ attachments
- 60% of data is PDF documents
- 30% of data is uncompressed large XML objects

WHY

Maximize memory usage
for important data
Reduce database recovery
time

HOW

- Compress large text data
 - 90% saving on XML data
- Store static data in files
 - Avoids DB handling overhead

HOW

- Table Size per schema

```
# Schema Table Usage
SELECT table_schema, table_name, engine, row_format, table_rows, avg_row_length,
       (data_length+index_length)/1024/1024 as total_mb,
       (data_length)/1024/1024 as data_mb,
       (index_length)/1024/1024 as index_mb,
       CURDATE() AS today
FROM   information_schema.tables
WHERE  table_schema = DATABASE()
ORDER BY 7 DESC;
```

<http://ronaldbradford.com/mysql-dba/>



PROBLEM

I can't access my website?

TRUE STORY

Question:

How do you know when your server is down or not accessible?

Answer:

The users will let us know.

SOLUTION

Integrated monitoring including graphical interface, real time analysis and notification

HOW

- Monitoring/Alerting
 - Graphical
 - Historical
 - Necessary
 - Generally missing/incomplete
 - Useless for real-time analysis

HOW

- Dashboard
 - The state of NOW
 - Sampling at 1s/3s/5s
 - e.g. 0.1% of throughput



HOW

• Instrumentation

- Important to business viability
 - e.g. orders per minute
 - page load time
- Seamless implementation
 - i.e. no code changes to view real-time
 - extensible



PROBLEM

My replication slave can't keep up?

SOLUTION

Know the weakest link(s) of MySQL replication and don't exceed that, or cheat.

WHY

If replication can't catchup,
slaves are useless.

Backup & recovery may
also suffer.

HOW

Master

- DML Statement
- Write Data/Redo Log
- Write Binary Log
- Return OK to client

HOW

Slave

- Detect master log change
- Retrieve binary log entry
- Write relay log (IO_THREAD)
- Read relay log
- **Apply DML (SQL_THREAD)**
- Write Data/redo log

HOW

Replication workarounds

- Restrict queries executed
 - --ignore
- Different storage engines
- Different index structures

EXPERT TIP

Advanced workarounds

- RAID 0 (large number of slaves)
- Pre fetch thread



PROBLEM

My server has crashed with
a hard drive failure

TRUE STORY

Question:

Have you ever performed a database recovery?

Answer:

No, why?

TRUE STORY

Consultant:

Do you know that your daily backups only recover the data up to that time, e.g. 1am. You know you have lost all your sales and changes since then.

Customer:

No, I didn't know that.

SOLUTION

- Have a DR Plan
 - Documented
 - Tested
 - Timed
 - Verified - End to End

HOW

- Do you pass the MySQL backup/recovery quiz?

<http://rb42.com/mysql-backup-quiz>

QUIZ

1. Do you have MySQL backups in place?
2. Do you backup ALL your MySQL data?
3. Do you have consistent MySQL backups?
4. Do you have backups that include both static snapshot and point in time transactions?
5. Do you review your backup logs EVERY SINGLE day or have tested backup monitoring in place?
6. Do you perform a test recovery of your static backup?
7. Do you perform a test recovery to point in time?
8. Do you time your backup and recovery process and review over time?
9. Do you have off-site copies of your backups?
10. Do you backup your primary binary logs?

<http://rb42.com/mysql-backup-quiz>

4

PROBLEM

Why is my database executing 1,200 qps for 50 users?

SOLUTION

Determine what queries are running and why they are running?

CAUSE

Excessive SQL statements

- Duplicate
- Redundant
- Cachable
- Row at a time (RAT)

WHY

Reducing SQL load both improves performance now and provides greater capacity as you scale

EXAMPLE

RAT

```
SELECT * FROM activities_theme WHERE theme_parent_id=0
SELECT * FROM activities_theme WHERE theme_parent_id=1
SELECT * FROM activities_theme WHERE theme_parent_id=2
SELECT * FROM activities_theme WHERE theme_parent_id=11
SELECT * FROM activities_theme WHERE theme_parent_id=16
```

CAT

```
SELECT *
FROM activities_theme
WHERE theme_parent_id in (0,1,2,11,16)
```

<http://ronaldbradford.com/blog/optimizing-sql-performance-the-art-of-elimination-2010-07-08/>

EXAMPLE

Unnecessary

Duplicate

Only 2 queries necessary or 1 CAT

```
5 Query SELECT * FROM `artist`
5 Query SELECT * FROM `artist`
5 Query SELECT * FROM `artist`
5 Query SELECT * FROM `artist`
5 Query SELECT * FROM `artist`
5 Query SELECT * FROM `artist` WHERE (ArtistID = 196 )
5 Query SELECT * FROM `artist` WHERE (ArtistID = 2188 )
5 Query SELECT * FROM `artist`
5 Query SELECT * FROM `artist`
```

<http://ronaldbradford.com/blog/optimizing-sql-performance-the-art-of-elimination-2010-07-08/>

EXAMPLE

Unnecessary

```
SELECT pages_id, pages_livestats_code, pages_title,
       pages_parent, pages_exhibid, pages_theme,
       pages_accession_num
FROM pages WHERE pages_id = 0
```

- 5 minutes
- 6000 executions
- 0 is out of bounds

<http://ronaldbradford.com/blog/optimizing-sql-performance-the-art-of-elimination-2010-07-08/>

HOW

- Capture & Analyze
- DML is easy
- SELECT is harder

<http://www.slideshare.net/ronaldbradford/capturing-analyzing-and-optimizing-mysql>

HOW

- MySQL Binary Log (Archive Redo)
 - mysqlbinlog
 - mk-query-digest
 - One Liner

<http://ronaldbradford.com/blog/mysql-dml-stats-per-table-2009-09-09/>

HOW

```
55463 update sessions
25574 insert into sessions
12820 update items
11636 insert into item_categories
7532 update users
5168 delete from item_categories
4076 update extended_item_infos
3701 insert into sphinx_new_items
3701 insert into mini_items
2190 update sweet_bars
1922 update chat_users
1662 update item_shipping_infos
1265 update search_terms
1260 insert into images
931 delete from item_shipping_infos
825 update booths
713 update booth_stats
574 update topics
540 update offers
```

Top 2 queries = 57%

81k of 141k

HOW

- Process List
- General Log
- tcpdump
- Application

HOW

Capturing, Analyzing and Optimizing your SQL

<http://www.slideshare.net/ronaldbradford/capturing-analyzing-and-optimizing-mysql>

EXPERT TIP

`/* Comment your queries */`

The more products you have, the more developers you have, the more time you spend in code identification before you can even determine a resolution

3

PROBLEM

The database is slow.
My webpage takes five seconds to load.

SOLUTION

Evaluate the time taken in the database and all stages in between

EXAMPLE

Client example showed a webpage taking 5 seconds to load. The html component was taking only 400 ms. Any MySQL performance improvement will only tune 8% of the total time.

WHY

Performance is important to end user
Performance is perception

HOW

- Firebug - <http://getfirebug.com/>
- Httpwatch - <http://httpwatch.com/>
- Page speed - <http://code.google.com/speed/page-speed/>
- YSlow - <http://developer.yahoo.com/yslow/>
- wget/curl
- Application code instrumentation

EXPERT TIP

- <http://www.stevesouders.com/>
- <http://developer.yahoo.com/performance/rules.html>



PROBLEM

I want to add new H/W.
How do I change my
application to support this?

SOLUTION

Develop a seamless
integration that requires no
code changes, no downtime
and very little additional
physical resources.

HOW

- Integrated monitoring and instrumentation
 - Deployed from Day 1

HOW

- Seamless automated server deployment
 - Version Control
 - Build & Release
 - Runtime config management
 - Automated discovery

HOW

- API
 - One code path for business functionality
 - Implied business documentation
 - Enforced data exchange standard
 - **Testability**

HOW

- Different levels of data availability
 - Read & Write
 - Read Only
 - No Access
 - Cached

HOW

- Different principles for scalability
 - Read Scalability
 - Write Scalability
 - Caching

REFERENCE

Successful MySQL Scalability

1. Integrated monitoring & instrumentation
2. Seamless automated server deployment
3. Disaster is inevitable
4. Application Programming Interface
5. Support different levels of data availability
6. Support different scalability principles

<http://ronaldbradford.com/blog/successful-mysql-scalability-presentation-2010-09-17/>

1

PROBLEM

My website is slow?

SOLUTION

Seek professional advice.

Hire for example Ronald Bradford.

- 22+ years of system architecture, database design and performance tuning.
- Employment as Consultant for Oracle Corporation (96-99)
- Employment as Senior Consultant for MySQL Inc (06-08)

PR

RECAP

- Monitoring. Before, during and after NOW.
- You may not be able to predict the future but you can preempt the future.
- Choose the best product and features for you needs.
- **The best SQL statement is the one you never have to execute.**

RECAP

- 3 levels of real time data access.
 - Read/Write, Read and no access
- 3 aspects of scalability.
 - Read, Write and Caching
- Operate below 90% capacity. That 10% is your insurance.

CONTACT

- <http://ronaldbradford.com>
- me@ronaldbradford.com

$$E_M = p s^n$$