

Best Practices for Database Administrators

Sheeri K. Cabral
Database Administrator
The Pythian Group, www.pythian.com

cabral@pythian.com
2008 MySQL User Conference & Expo



MIRE

- Make It Really Easy
- Automate
- Document
- As for your brain.....

Use your brain for CPU, not storage!

(use a request tracking system!)

Monitoring Basics

- Graph (Cacti)
- Alert (Nagios)
- Check your checks

Tradeoffs

- Tradeoffs always exist
- Think about them

Most Commonly Given Advice

- Backup
- Restore
- What do you use backups for?

In Case of Failure.....

- Restore from backup
- Master/Slave
- Master/Master
- Cluster

DR? HA!

- Test your DR/HA plans
- What scenarios do they cover?
- What scenarios do they NOT cover?

Ounces of Prevention

- Configuration vs. reality
- Possible memory usage
- Disk space / tablespace size

Pounds of Cure

- Error Logs
- Slow Query Logs
- Query Review

mysqldumpslow

Count: 1 Time=12772.00s (12772s) Lock=0.00s (0s) Rows=59493.0 (59493),
prod[prod]@[192.168.217.53]

```
select * from properties_new prop left join nodes node on prop.node_id  
= node.id where node.class_type = 'S' and qname like 'S' and node_id  
not in ( select ancestor from history_links ) and node.id >= N order  
by node_id
```

Count: 7 Time=160.57s (1124s) Lock=0.00s (0s) Rows=5567304.1
(38971129), test[test]@localhost SELECT /*!N SQL_NO_CACHE */ * FROM
`properties_new`

Count: 2 Time=51.00s (102s) Lock=0.00s (0s) Rows=0.0 (0),
prod[prod]@[192.168.217.53]

```
update nodes set vers=N, version_id=N, guid='S', creator='S',  
owner='S', lastModifier='S', createDate=N, modDate=N, accessDate=N,  
is_root=N, store_new_id=null, acl_id=null where id=N and vers=N
```

Count: 7 Time=32.00s (224s) Lock=0.00s (0s) Rows=698414.6 (4888902),
test[test]@localhost

```
SELECT /*!N SQL_NO_CACHE */ * FROM `nodes` |
```

Reading slow log '/opt/mysql/mysql/data/mysql-slow.log'.
25851 total queries, 291 unique.
Sorting by 'at'.

001

Count : 2 (0%)
Time : 19677 s total, 9838 s avg, 6905 s to 12772 s max
95% of Time : 6905 s total, 6905 s avg, 6905 s to 6905 s max
Lock Time : 0 s total, 0 s avg, 0 s to 0 s max
Rows sent : 29746 avg, 0 to 59493 max
Rows examined : 75011 avg, 0 to 150023 max
User : user1@/192.168.217.53 (96%)
Database : Unknown

```
SELECT * FROM properties_new prop LEFT JOIN nodes node ON prop.node_id =  
node.id WHERE node.class_type = 'S' AND qname LIKE 'S' AND node_id NOT IN  
(S0) AND node.id >= N ORDER BY node_id;
```

Count : 13 (0%)
Time : 24042 s total, 1849 s avg, 5 s to 19839 s max
95% of Time : 4203 s total, 350 s avg, 5 s to 2055 s max
Lock Time : 0 s total, 0 s avg, 0 s to 0 s max
Rows sent : 7692 avg, 0 to 100000 max
Rows examined : 32573 avg, 0 to 423454 max
User : prod@/196.168.217.53 (96%)
Database : prod

```
SELECT * FROM properties_new prop LEFT JOIN nodes node ON prop.node_id =  
node.id WHERE node.class_type = 'S' AND qname LIKE 'S' AND node_id NOT IN  
(S0) AND node.id >= N ORDER BY node_id LIMIT N;
```

Query Profile

- mk-query-profiler

Press <ENTER> when the external program is finished

```
+-----+
|                1 (476.2167 sec)                |
+-----+
```

Overall stats	Value
Total elapsed time	476.217
Questions	5431
COMMIT	1723
DELETE	8
DELETE MULTI	0
INSERT	8
INSERT SELECT	0
REPLACE	0
REPLACE SELECT	0
SELECT	2833
UPDATE	0
UPDATE MULTI	0
Data into server	884177
Data out of server	1013231

Table and index accesses	Value
Table locks acquired	3959
Table scans	26
Join	0
Index range scans	0
Join without check	0
Join with check	0
Rows sorted	17489
Range sorts	203
Merge passes	0
Table scans	21
Potential filesorts	11

Query Profile

- mk-query-profiler
- SHOW STATUS before and after
- mysqlreport

MySQL 5.0.45-log

uptime 52 21:53:13

Tue Apr 15 14:37:23 2008

Key

Buffer used 2.16M of 16.00M %Used: 13.51
Current 3.32M %Usage: 20.73
Write hit 0.00%
Read hit 100.00%

Questions

Total 84.87M 18.6/s
DMS 48.18M 10.5/s %Total: 56.77
Com_ 36.54M 8.0/s 43.06
COM_QUIT 143.59k 0.0/s 0.17
+Unknown 581 0.0/s 0.00
Slow (4) 799 0.0/s 0.00 %DMS: 0.00 Log: ON
DMS 48.18M 10.5/s 56.77
SELECT 47.88M 10.5/s 56.42 99.37
INSERT 165.96k 0.0/s 0.20 0.34
DELETE 79.91k 0.0/s 0.09 0.17
UPDATE 55.47k 0.0/s 0.07 0.12
REPLACE 23 0.0/s 0.00 0.00
Com_ 36.54M 8.0/s 43.06
commit 24.27M 5.3/s 28.59
set_option 11.57M 2.5/s 13.63
rollback 543.18k 0.1/s 0.64



__ SELECT and Sort

Scan	472.67k	0.1/s	%SELECT: 0.99
Range	9.49k	0.0/s	0.02
Full join	4.94k	0.0/s	0.01
Range check	0	0/s	0.00
Full rng join	146.47k	0.0/s	0.31
Sort scan	194.32k	0.0/s	
Sort range	2.80M	0.6/s	
Sort mrg pass	10.12k	0.0/s	

__ Table Locks

Waited	0	0/s	%Total: 0.00
Immediate	1.12k	0.0/s	

__ Tables

Open	422 of 1000	%Cache: 42.20
Opened	7.16k	0.0/s

__ Connections

Max used	42 of 100	%Max: 42.00
Total	143.59k	0.0/s

__ Created Temp

Disk table	9.32k	0.0/s	
Table	238.23k	0.1/s	Size: 32.0M
File	32	0.0/s	



Threads

Running	2 of 42		
Cached	0 of 0	%Hit:	0
Created	143.59k	0.0/s	
Slow	0	0/s	

Aborted

Clients	0	0/s
Connects	0	0/s

Bytes

Sent	2.83G	620.1/s
Received	2.81G	614.6/s

InnoDB Buffer Pool

Usage	1.00G of 1.00G	%Used:	100.00
Read hit	99.84%		
Pages			
Free	1	%Total:	0.00
Data	63.90k	97.50 %Drty:	0.00
Misc	1638	2.50	
Latched	0	0.00	
Reads	326.29M	71.4/s	
From file	515.08k	0.1/s	0.16
Ahead Rnd	16422	0.0/s	
Ahead Sql	30699	0.0/s	
Writes	21.52M	4.7/s	
Flushes	631.57k	0.1/s	
Wait Free	0	0/s	

__ InnoDB Lock

Waits	0	0/s
Current	0	
Time acquiring		
Total	0 ms	
Average	0 ms	
Max	0 ms	

__ InnoDB Data, Pages, Rows

Data

Reads	787.19k	0.2/s
Writes	1.35M	0.3/s
fsync	1.16M	0.3/s
Pending		
Reads	0	
Writes	0	
fsync	0	

Pages

Created	21.68k	0.0/s
Read	2.19M	0.5/s
Written	631.57k	0.1/s

Rows

Deleted	4.64M	1.0/s
Inserted	4.95M	1.1/s
Read	1.81G	396.5/s
Updated	32.92k	0.0/s

Query Profile

- mk-query-profiler
- SHOW STATUS before and after
- mysqlreport
- Is every index being used?

What Does This Query Do?

- `SELECT amount, created FROM payments where user='sheeri'`

What Does This Query Do?

- `SELECT amount, created FROM payments where user='sheeri'`

VS

- `SELECT /* find payments for customer */ amount, created FROM payments where user='sheeri'`

Data Profile

- PROCEDURE ANALYSE()
- What does each column, table do?
- What's in a name?

Does It Make Sense?

- Can you “read” the data?
- Question “best practices”
- Then, make some!

Schema Profile

- Start normal
- Denormalize if necessary
 - Descriptive foreign keys can prevent denormalization
- Stored procedures for developers

Replication

- Be careful of TRIGGERS
 - And any DML from SP or UDF
- Sync often with mk-table-sync
- Handle duplicates carefully

Maintenance

- Partition
- Archive
- Purge
- Static data

Manage User Expectations

- Constant report refreshes
- Aggregate data once every hour or 15 minutes
- Split off processing:
 - customer/non-customer
 - internal/external

Creative Suggestions – Forums

- COUNT(*) for paging
- LIMIT 0,n+1
- Cache each page of forum
 - First page can have up to n+5 entries

Creative Suggestions – Calculations

- Limited set of calculations?
 - Distance between zip codes in the US
- Calculate the same thing more than once?
- Calculating easy constants?

Foreign Keys

- Application still has to handle problems

Summary



When I Walk In the Door

- alerting
- request tracking
- graphing
- documentation

My Toolbox

- EXPLAIN
- Ma'atkit
- hackmysql.com (mysqsla, mysqlreport)
- MySQL Manual

Also Useful

- mytop
- innotop

Testing, Testing...is this thing on?

- Functional
- Load
- Start now!

Consider


- master/master vs. read_only slave
- Learn the foundations
- Semi-dynamic data

Redundancy

- Make everything replaceable
 - Kickstart
 - Automated config management
- Including yourself!
 - Don't you want a vacation?






Be a Good DBA

- Prove your work to yourself
- Be clear



FoxTrot

by Bill Amend



Sheeri Cabral

cabral@pythian.com

www.sheeri.com

