

EffectiveMySQL.com

Performance, Scalability & Business Continuity



MySQL Disasters, and how to avoid yours.

Ronald Bradford
<http://ronaldbradford.com>

2012.05



EffectiveMySQL.com - Performance, Scalability & Business Continuity

**"No one cares if you can
backup, only that you
can restore."**

Adapted from W. Curtis Preston - *Backup & Recovery* (O'Reilly 2009)



EffectiveMySQL.com - Performance, Scalability & Business Continuity

AGENDA

- Lots of Quotes
- The Five Steps to Recovery
- Memorable Examples
- Point in Time
- MySQL Replication



EffectiveMySQL.com - Performance, Scalability & Business Continuity

ABOUT THE AUTHOR

RONALD BRADFORD

- **Provide independent consulting - Available NOW**
- Top invited MySQL speaker (2011- 13 countries, 30+ talks)
- All time top MySQL blogger
- Published Author (3++ books)
- Oracle ACE Director
- MySQL community member of the year
- 23 years of RDBMS experience, 13 years with MySQL
 - MySQL Inc (2006-2008)
 - Oracle Corporation (1996-1999)



Five Step B&R Overview

BACKUP BASICS



Static Consistent Backup
+
Master Binary Logs

RECOVERY BASICS



Static Restore
+
Point in Time Recovery

VERIFICATION



- Necessary at EVERY step
- Commands complete without error
- No errors in logs
- Results match expectations

TESTING



Testing is about trying to break your software, not checking that it works!

REDUNDANCY



Difficult to restore your backup when the server with the backup is unavailable.

**"Disaster is inevitable.
Total failure is
avoidable"**

Ronald Bradford & Paul Carlstroem - 2011

Memorable Examples

EXAMPLE 1 ENVIRONMENT

- Master/Slave Configuration
- Dedicated backup on slave
- Regular backups of binary logs
- Tested recovery processes
- System/MySQL monitoring

EXAMPLE 1 PROBLEM

```
[ERROR] mysqld: Disk is full writing '/mysql/binlog/log-bin.000020'
(Errcode: 28). Waiting for someone to free space... Retry in 60 secs
[ERROR] Could not use /mysql/binlog/log-bin for logging (error 28).
Turning logging off for the whole duration of the MySQL server process.
To turn it on again: fix the cause, shutdown the MySQL server and restart
it.
```

<http://ronaldbradford.com/blog/never-let-your-binlog-directory-fill-up-2009-07-15/>

EXAMPLE 1 BACKGROUND

Insert long and interesting story here.

EXAMPLE 1

TRIAGE

- Bureaucracy
- Adequate staff training
- Open access to information

EXAMPLE 2

From: Existing client
Subject: Emergency

Are you around? My production system is crashed, I'm traveling, and have an emergency.

EXAMPLE 2

ENVIRONMENT

- Single server
- Binary logging *disabled*


EXAMPLE 2

PROBLEM

```
[ERROR] Got error 127 when reading table './cust/tbl1'
[ERROR] Got error 127 when reading table './cust/tbl1'
[ERROR] Got error 127 when reading table './cust/tbl1'
[ERROR] Got error 127 when reading table './cust/tbl1'
[Note] Retrying repair of: './cust/tbl1' with keycache
[Note] Retrying repair of: './cust/tbl1' failed. Please
try REPAIR EXTENDED or myisamchk
[ERROR] /var/lib/mysql5/bin/mysqld: Table './cust/tbl1'
is marked as crashed and last (automatic?) repair failed
[ERROR] /var/lib/mysql5/bin/mysqld: Table './cust/tbl2'
is marked as crashed and should be repaired
[Warning] Checking table: './cust/tbl2'
```


EXAMPLE 2


PROBLEM



```
$ myisamchk cust/*.MYI
Checking MyISAM file: /var/lib/mysql/cust/tblX.MYI
Data records:      0   Deleted blocks:      0
- check file-size
- check record delete-chain
- check key delete-chain
- check index reference
- check data record references index: 1
- check data record references index: 2
- check record links
```

EXAMPLE 2

PROBLEM



```
$ myisamchk cust/*.MYI
Checking MyISAM file: /var/lib/mysql/cust/tblY.MYI
Data records: 68384   Deleted blocks:      0
- check file-size
- check record delete-chain
- check key delete-chain
- check index reference
- check data record references index: 1
- check data record references index: 2
- check data record references index: 3
- check data record references index: 4
- check data record references index: 5
- check record links
myisamchk: error: Found wrong record at 1644072
MyISAM-table '/var/lib/mysql/cust/tblY.MYI' is corrupted
Fix it using switch "-r" or "-o"
```

EXAMPLE 2


PROBLEM



```
$ myisamchk -r /var/lib/mysql/cust/tblY.MYI
- recovering (with sort) MyISAM-table '/var/lib/mysql/cust/tblY.MYI'
Data records: 68384
- Fixing index 1
myisamchk: Duplicate key for record at 315344 against record at 315296
myisamchk: Duplicate key for record at 312272 against record at 312224
- Fixing index 2
- Fixing index 3
- Fixing index 4
- Fixing index 5
Data records: 68382
myisamchk: warning: 2 records have been removed
```

EXAMPLE 2

PROBLEM



```
$ myisamchk -r /var/lib/mysql/cust/tblZ.MYI
- recovering (with sort) MyISAM-table '/var/lib/mysql/cust/tblZ.MYI'
Data records: 22528
- Fixing index 1
- Fixing index 2
Wrong bytesec: 2- 0- 38 at 386620; Skipped
MyISAM-table '/var/lib/mysql/cust/tblZ.MYI' is not fixed because of errors
Try fixing it by using the --safe-recover (-o), the --force (-f)
option or by not using the --quick (-q) flag
```

EXAMPLE 2

PROBLEM

```
$ myisamchk -r /var/lib/mysql/cust/tbl1.MYI
- recovering (with sort) MyISAM-table '/var/lib/mysql/cust/tbl1.MYI '
Data records: 584
- Fixing index 1
Key 1 - Found wrong stored record at 14541448
Found wrong packed record at 14542372
Wrong aligned block at 109106305
Delete link points outside datafile at 109106305
...
Several hundred more lines
...
Found block with too small length at 109330004; Skipped
Wrong bytesec: 14- 0- 25 at 65272324; Skipped
Key 1 - Found wrong stored record at 114760168
Key 1 - Found wrong stored record at 114761084
Segmentation fault
```

EXAMPLE 2

BACKGROUND

Insert another long and interesting story here.

EXAMPLE 2

TRIAGE

- Anywhere, anytime
- Single server
- Contingency plan

MORE EXAMPLES

1. No Binary Logging
2. Single Production Server
3. Lack of MySQL security
4. Deleting MySQL data
5. Deleting the InnoDB data file
6. Deleting Binary Logs

MORE EXAMPLES

7. Failure after software upgrade
8. Handling MySQL corruption
9. Missing schema data in backup
10. Data dictionary inconsistency
11. Restoring static on a running server
12. Auto crash recovery failure

B&R Options

B&R OPTIONS

- | | |
|-----------------------|----------------------|
| ● Included | ● Open Source |
| ● mysqldump | ● XtraBackup |
| ● OS filecopy | ● mydumper |
| ● OS Specific | ● Commercial |
| ● filesystem snapshot | ● MEB |

B&R OPTIONS

- “MySQL Backup and Recovery Essentials”

Slides at <http://j.mp/EM-BandR>

ADVANCED FEATURES

- Compression
- Incremental
- Remote
- Parallel
- Partial

COMPRESSION

Utility	Comp (s)	Dec (s)	Saving
lzo (-3)	21	34	48%
pigz (-1)	43	33	64%
pigz [-6]	105	25	69%
gzip [-6]	232	78	69%
bzip2	540	175	74%
lzo (-9)	20m	82	58%
lzma	58m	180	78%
xz	59m	160	78%

Depends greatly on data types

Point in Time

POINT IN TIME

● Possible with Binary Logs

WARNING: If you care about your data, enable binary logging

```
#my.cnf

[mysqld]
log-bin=mysql-bin
expire-logs-days=5
```

```
mysql> SHOW BINARY LOGS;
+-----+-----+
| Log_name          | File_size |
+-----+-----+
| mysql-bin.019662  | 104857736 |
| mysql-bin.019663  | 104857699 |
| mysql-bin.019664  | 104857850 |
```

```
$ ls -ltr /var/log/mysql | tail
-rw-rw---- 1 mysql adm 104857848 2011-09-04 22:00 mysql-bin.020607
-rw-rw---- 1 mysql adm 104857953 2011-09-04 22:08 mysql-bin.020608
```

BACKUP OPTIONS

- cp
- rsync
- Slave --log-slave-updates
- DRBD
- mysqlbinlog --read-from-remote-server
(New in 5.6)

BACKUP NEEDS

- Static Backup
- Useless without Binary Log position

```
mysql> SHOW MASTER STATUS\G
      File: mysql-bin.020616
      Position: 63395562
      Binlog_Do_DB:
      Binlog_Ignore_DB:
```

WARNING: Can work on slave and provide the wrong information

BACKUP NEEDS

```
$ mysqldump --master-data (or --dump-slave)
CHANGE MASTER TO MASTER_HOST='10.0.0.1', MASTER_USER=
'repl',MASTER_PASSWORD='*****', MASTER_LOG_FILE=
'mysql-bin.000146', MASTER_LOG_POS=810715371;
```

Xtrabackup

```
$ cat xtrabackup_binlog_info
mysql-bin.000001 37522
```

MEB

```
$ grep binlog meta/backup_variables.txt
binlog_position=mysql-bin.000017:5555
```

mydumper

```
$ cat export-20120407-230027/metadata
Log: mysql-bin.000017 Pos: 8328
```

RESTORE OPTIONS

- mysqlbinlog

```
$ mysqlbinlog /path/to/mysql-bin.000146 \
--start-position=810715371 | mysql -uroot -p
```

```
$ mysqlbinlog /path/to/mysql-bin.000147 \
/path/to/mysql-bin.00148 ... etc | mysql -uroot -p
```

- MySQL replication

MySQL Replication

REPLICATION

- Replication is not a backup solution
- Important part of a total strategy
- Important for HA
 - Only when failover capable

CONSISTENCY

- Asynchronous Nature
- SQL_SLAVE_SKIP_COUNTER
- --slave-skip-errors
- SUPER privilege
- binlog-[do/ignore]-db
- replicate-[do/ignore]-db

MYSQL 5.6

- Server UUID
- Global Transaction Id
- Binlog API
- Multi-threaded slaves
- Crash safe slaves
- Replication Checksums
- Remote binlog backup
- Failover utilities

Conclusion

“A SAN is not a backup solution.”

<http://it.toolbox.com/blogs/database-soup/a-san-is-not-a-highavailability-solution-47644>

CONCLUSION

Human error due to lack of knowledge and procedures causes many cascading failures.

CONCLUSION

"In most cases, emergencies analyzed could have been prevented best by a systematic, organization-wide effort to find and remove latent problems [before they occur]. Many of the activities involved in this effort could appear to be unproductive, and might be unrewarding for people to do."

<http://www.percona.com/about-us/mysql-white-paper/causes-of-downtime-in-production-mysql-servers>


CONCLUSION

- Advanced features are important
- Best option depends
- Replication is important in your strategy
- Test, Test, Test. - “Chaos Monkey”

PRESENTATIONS

More presentations at

<http://ronaldbradford.com/mysql-presentations/>
<http://effectivemysql.com/presentation/>



220 pages dedicated to B&R

<http://j.mp/EM-book2>

$EM = ps^n$

Ronald Bradford
<http://effectivemysql.com>