

# Effective MySQL Architecture and Design Practices

**Ronald Bradford**  
**May 2015**

# INFORMATION

These slides can be found at

<http://effectiveMySQL.com>

# **What is architecture?**

# **What is design?**

# EFFECTIVE TIPS

1. The MySQL warning
2. Do not start with default configuration
3. Start with MySQL 5.5 or newer
4. Design with replication in mind
5. Know the right data types

# **25+ YEARS EXPERIENCE**

**MySQL has many great features that are abused by “lack of good defaults,” “poor open source reference examples” and “lack of experienced software developers.”**

**Ronald Bradford  
August 2014**

# MySQL Warnings

```
CREATE SCHEMA IF NOT EXISTS test;  
USE test;  
  
DROP TABLE IF EXISTS stats;  
CREATE TABLE stats(  
    counter TINYINT(4) UNSIGNED NOT NULL,  
    val     CHAR(1) NULL  
);
```

## \* Simple table

- \* one mandatory integer column
- \* one optional character column



```
CREATE SCHEMA IF NOT EXISTS test;
USE test;

DROP TABLE IF EXISTS stats;
CREATE TABLE stats(
  counter TINYINT(4) UNSIGNED NOT NULL,
  val     CHAR(1) NULL
);
```

```
START TRANSACTION;
INSERT INTO stats(counter) VALUES (-1);
INSERT INTO stats(counter) VALUES (100);
INSERT INTO stats(counter) VALUES (999);
COMMIT;

SELECT * FROM stats;
```

# Effective MySQL Architecture and Design Practices

```
START TRANSACTION;  
INSERT INTO stats(counter) VALUES (-1);  
INSERT INTO stats(counter) VALUES (100);  
INSERT INTO stats(counter) VALUES (999);  
COMMIT;  
  
SELECT * FROM stats;
```

counter	val
-1	NULL
100	NULL
999	NULL

counter	val
0	NULL
100	NULL
255	NULL

counter	val
100	NULL

```
CREATE SCHEMA IF NOT EXISTS test;
USE test;


DROP TABLE IF EXISTS stats;
CREATE TABLE stats(
  counter TINYINT(4) UNSIGNED NOT NULL,
  val     CHAR(1) NULL
);
```

- \* TINYINT 1 byte = 8 bits
- \* (4) meaningless - for display only
- \* UNSIGNED (non negative)
  - \*  $2^8$  (256 values) Range 0 - 255



Discuss in  
Point 5

- \* TINYINT | byte = 8 bits
- \* UNSIGNED (non negative)
- \* Range 0 - 255


```
START TRANSACTION;  
... VALUES (-1);  
... VALUES (100);  
... VALUES (999);  
COMMIT;  
  
SELECT * FROM stats;
```



counter	val
-1	NULL
100	NULL
999	NULL



counter	val
0	NULL
100	NULL
255	NULL



counter	val
100	NULL

# THREE STATES

- \* Default MySQL Behavior for SQL

  - \* ERROR

  - \* Success

  - \* Success, but with warnings



Why????

- \* 90+% developers (and products) ignore warnings

# SHOW WARNINGS

```
mysql> INSERT INTO stats(counter) VALUES(-1);
Query OK, 1 row affected, 1 warning (0.03 sec)
```

-1 != 0

```
mysql> SHOW WARNINGS\G
```

```
***** 1. row *****
Level: Warning
Code: 1264
Message: Out of range value for column 'counter' at row 1
1 row in set (0.00 sec)
```

```
mysql> SELECT * FROM stats;
```

counter	val
0	NULL

**Complete loss of Data Integrity. Unrecoverable.**

# DEFAULT OPERATION

“MySQL by default performs silent truncations, making a best guess of what the value of data should be based on the schema design, rather than enforcing, data in equals data out.”

**Ronald Bradford**  
**August 2014**

# TRUNCATION SCOPE

- \* Not limited to Integers
  - \* Strings, Dates, Float, Decimal



# STRING EXAMPLE

```
CREATE TABLE tour_locations (place VARCHAR(10));
INSERT INTO tour_locations VALUES
('Helsinki, Finland'),
('Stockholm, Sweden'),
('Riga, Latvia'),
('Tallinn, Estonia');
```

```
SELECT place FROM tour_locations;
```

```
+-----+
| place |
+-----+
| Helsinki, |
| Stockholm, |
| Riga, Latv |
| Tallinn, E |
+-----+
```

## DATE EXAMPLE

```
CREATE TABLE tour_dates(event_date DATE);

INSERT INTO tour_dates VALUES ('2014-02-28'),
('2014-02-29'), ('2014-03-00'), ('2014-31-07');

SELECT event_date FROM tour_dates;
+-----+
| event_date |
+-----+
| 2014-02-28 |
| 0000-00-00 |
| 2014-03-00 | * ZERO DAY STORED
| 0000-00-00 |
+-----+
```

## OTHER WARNINGS

```
+-----+-----+-----+
| Level   | Code  | Message                                     |
+-----+-----+-----+
| Warning | 1364  | Field 'password' doesn't have a default value |
| Warning | 1364  | Field 'salt' doesn't have a default value     |
| Warning | 1364  | Field 'token' doesn't have a default value    |
+-----+-----+-----+
```

```
SELECT cnt AS methods, COUNT(*) AS cnt
...
GROUP BY cnt
```

```
+-----+-----+-----+
| Level   | Code  | Message                                     |
+-----+-----+-----+
| Warning | 1052  | Column 'cnt' in group statement is ambiguous |
+-----+-----+-----+
```

# IDEAL OPERATION

- \* MySQL defaults can be changed
- \* MySQL can and does support good data and referential integrity
- \* MySQL defaults change and improve with newer versions

# **MySQL Configuration**

# OBSERVATION

“Using MySQL configuration defaults leads to loss of data integrity that can be unrecoverable. Code that will fail when good defaults are enabled.”

**Huge cost to business.**

**Ronald Bradford  
August 2014**

# SERVER SQL MODE

Better practices exist

"Modes define what SQL syntax MySQL should support and what kind of data validation checks it should perform."

<http://dev.mysql.com/doc/refman/5.5/en/sql-mode.html>

# CONFIGURATION

## \* Via Configuration

```
[mysqld]
sql_mode=
'STRICT_ALL_TABLES,NO_ENGINE_SUBSTITUTION,NO_ZERO_DATE,
NO_ZERO_IN_DATE,ERROR_FOR_DIVISION_BY_ZERO,ONLY_FULL_GROU
P_BY'
```

## \* Dynamic

```
mysql > SET GLOBAL
sql_mode='STRICT_ALL_TABLES,NO_ENGINE_SUBSTITUTION,NO_Z
ERO_DATE,NO_ZERO_IN_DATE,ERROR_FOR_DIVISION_BY_ZERO,ONL
Y_FULL_GROUP_BY';
```



# SQL\_MODE OPTIONS

## \* sql\_mode

\* NO\_ENGINE\_SUBSTITUTION

5.6

\* STRICT\_ALL\_TABLES

5.7

\* NO\_ZERO\_DATE

deprecated in 5.7

\* NO\_ZERO\_IN\_DATE

\* ERROR\_FOR\_DIVISION\_BY\_ZERO

\* ONLY\_FULL\_GROUP\_BY

## STRICT\_ALL\_TABLES

```
SET SESSION sql_mode='STRICT_ALL_TABLES';
START TRANSACTION;
INSERT INTO stats(counter) VALUES(-1);
ERROR 1264 (22003): Out of range value for column
'counter' at row 1
INSERT INTO stats(counter) VALUES(100);
INSERT INTO stats(counter) VALUES(999);
ERROR 1264 (22003): Out of range value for column
'counter' at row 1
COMMIT;
SELECT * FROM stats;
+-----+-----+
| counter | val  |
+-----+-----+
|      100 | NULL |
+-----+-----+
```

## APPLICATION LOGIC

mysql interactive mode

```
SELECT * FROM stats;  
+-----+-----+  
| counter | val  |  
+-----+-----+  
|      100 | NULL |  
+-----+-----+
```

\* An application would ROLLBACK on ERROR

```
$ mysql < stats.sql  
$ mysql -vvv -e "SELECT * FROM stats"  
  
SELECT * FROM stats  
-----  
Empty set (0.00 sec)
```

# GROUP BY EXAMPLE

```
SELECT customer_id, sum(amount)
FROM order
WHERE payment_method='CreditCard';
```

...

```
SET SESSION sql_mode='ONLY_FULL_GROUP_BY';
```

```
SELECT customer_id, sum(amount)
FROM order
WHERE payment_method='CreditCard';
```

```
ERROR 1140 (42000): Mixing of GROUP columns  
(MIN(),MAX(),COUNT(),...) with no GROUP columns is  
illegal if there is no GROUP BY clause
```

# OBSERVATION

“Poor defaults are not limited to SQL syntax, but data recoverability, data consistency, backup and recovery abilities, data access and locking”

# CONFIGURATION

## \* InnoDB

```
[mysqld]
default_storage_engine=InnoDB      # now default in 5.5
innodb_file_per_table              # now default in 5.6
innodb_log_file_size= N            # default of 5M
innodb_flush_method = O_DIRECT
```

\* innodb\_log\_file\_size ~ 256M - 1G

\* depending on version, buffer size

# ANTI DEFAULTS

Discussed in Web Operations presentation

\* mysqldump

**DEFAULT**

\* --lock-tables

**NOT DEFAULT**

\* --single-transaction

\* --master-data

\* --dump-slave (5.5)

\* --routines



# **MySQL Versions**



# ORACLE SUPPLIED

\* Version 5.5 GA (12/2010)

\* Version 5.6 GA (3/2013)

\* Version 5.7 DMR (3/2014)

\* Release cycle around every 2 years

# REPOS

## ORACLE PROVIDED

- \* RedHat/CentOS/OL Yum
- \* Debian/Ubuntu Apt

Server/Connectors/  
Workbench/Utilities/...

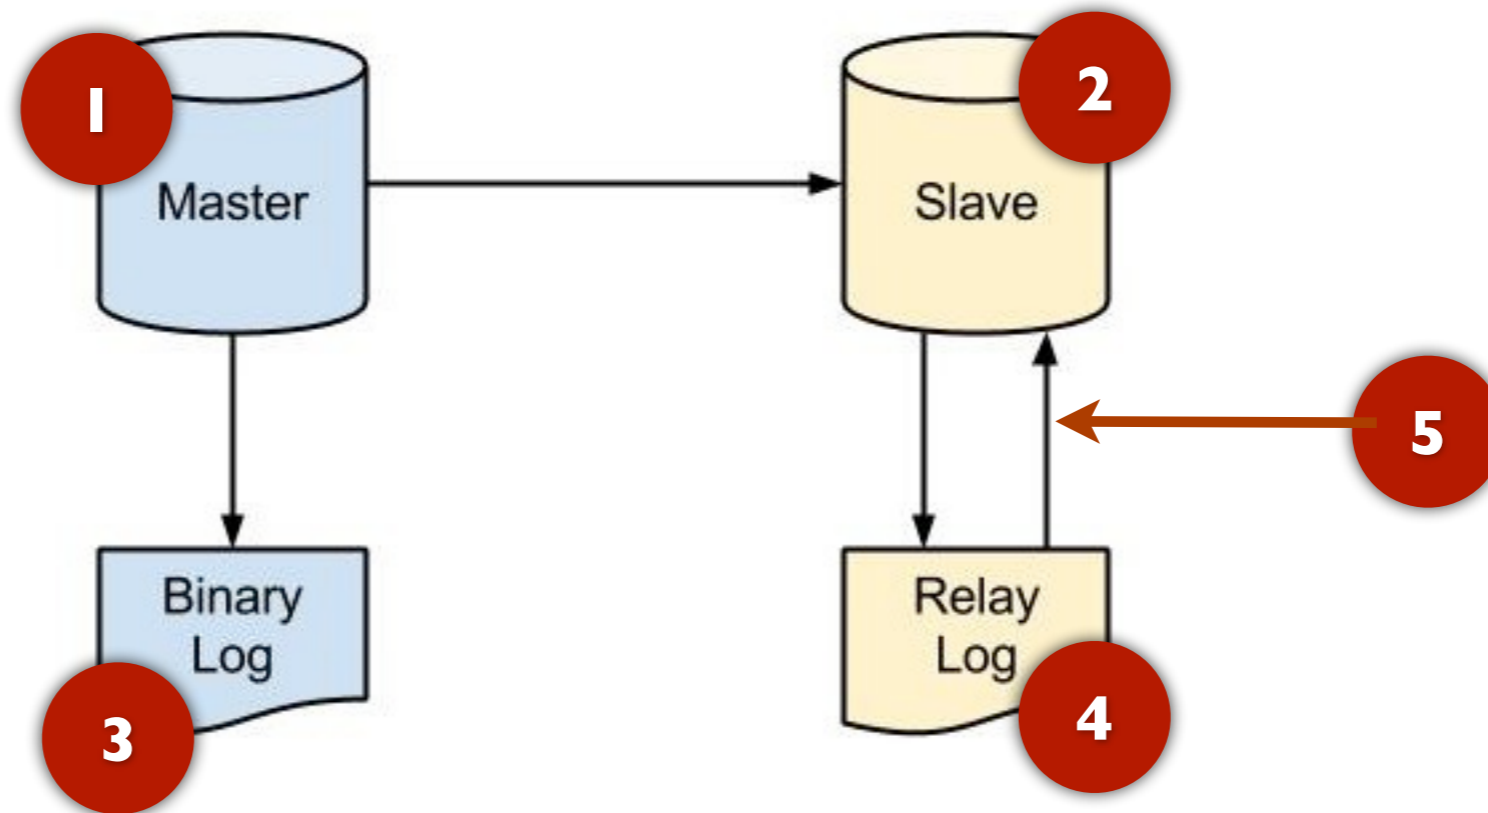
## 3rd PARTY

- \* Legacy mysql.com rpm's
- \* Others

Migration  
complexities

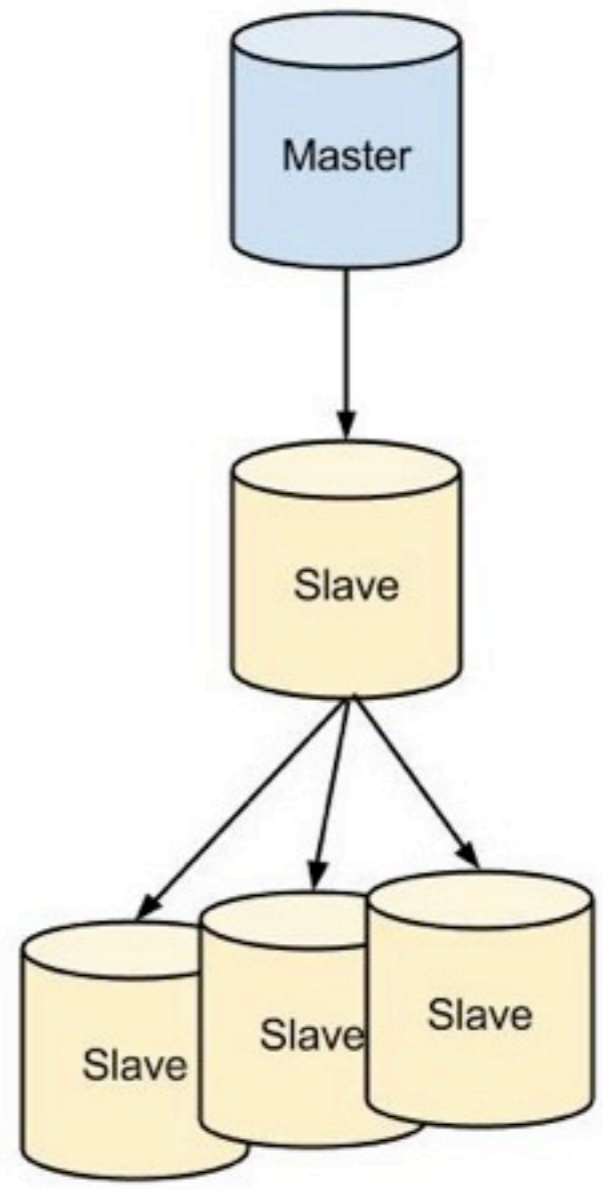
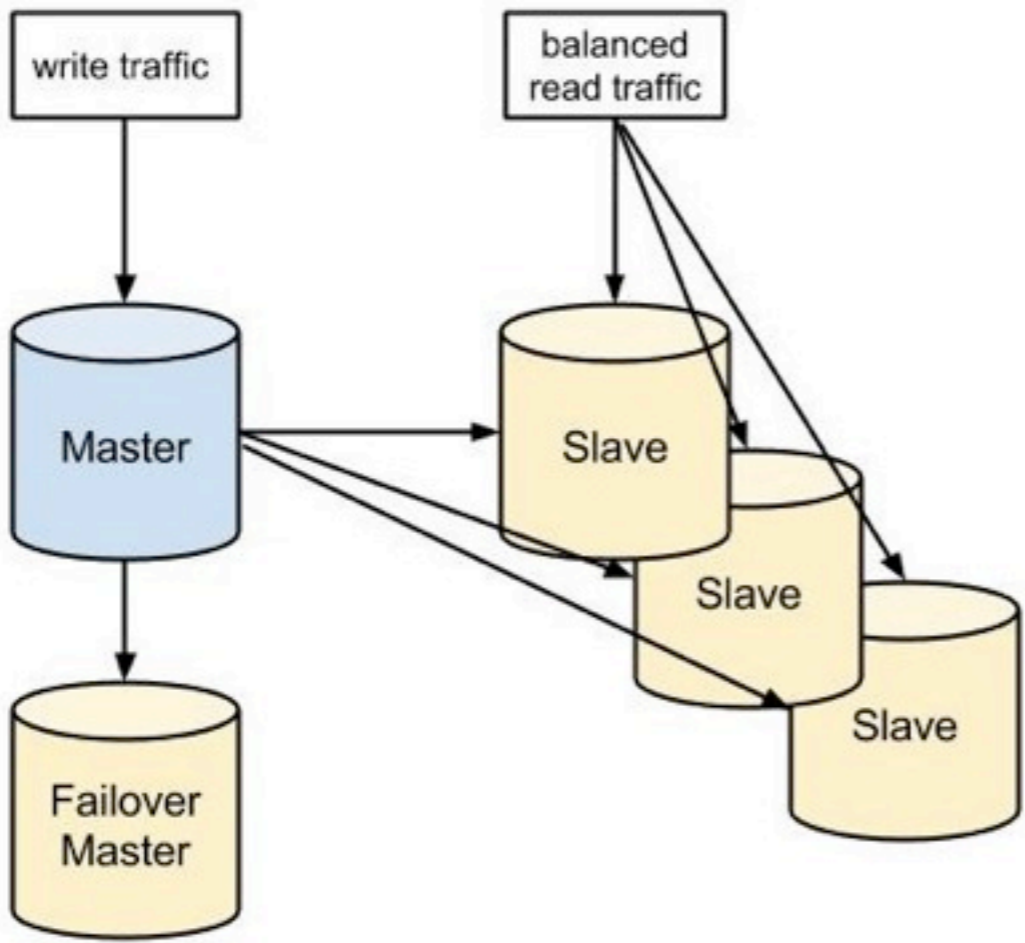


# **MySQL Replication**



<http://dev.mysql.com/doc/refman/5.6/en/replication.html>

# REALITY



# MYSQL REPLICATION

- \* Not configured ready by default
- \* Essential in all production systems
- \* Asynchronous (default)
  - \* Delay
  - \* Drift
- \* Semi-Synchronous (5.6+)

# CONFIGURATION

## \* Master Config (Required)

```
[mysqld]  
server_id=N  
log-bin
```

## \* Master user privileges (Required)

```
CREATE USER repl@99.99.99.99 IDENTIFIED BY 'somepassword';  
GRANT REPLICATION SLAVE ON *.* TO repl@99.99.99.99;
```

## \* Master Config (Recommended)

```
binlog_format=MIXED  
expire_logs_days=14
```

# 5.6 CONFIGURATION

\* Server UUID

Only for 5.6+

\* Auto created on startup

```
$ cat [datadir]/auto.cnf
[auto]
server-uuid=6fc5b8e6-1b43-11e4-b9f8-4f3c408e0a83
```

Be wary when cloning a slave. You must remove this file



# CONFIGURATION

## \* Slave (required)

```
[mysqld]  
server_id=M
```

## \* Slave (Recommended)

```
read_only=TRUE  
master-info-file=/mysql/binlog/master.info  
relay-log=/mysql/binlog/mysql-relay-bin  
relay-log-index=/mysql/binlog/mysql-relay-bin.index  
relay-log-info-file=/mysql/binlog/relay-log.info
```

## \* Slave (MySQL MHA)

```
relay_log_purge=0
```

- \* Replication is Asynchronous
  - \* Slaves pull from the Master
  - \* Delay can occur
    - \* A single long running query
  - \* SQL apply is single threaded
    - \* 5.6+ Multi-threading options
- \* Not slave crash safe (5.6 improvements)


- \* **STATEMENT** format
  - \* **ROW/MIXED** format (from 5.1)
- \* **Not Read Only** (`--read_only`)
- \* **SUPER** privilege
- \* `--skip-slave-errors`
- \* **SQL\_SLAVE\_SKIP\_COUNTER**

# REPLICATION MODES

- \* Classic Replication
- \* Transaction Based Replication (5.6+)
  - \* GTID
  - \* Transaction based
  - \* Improved consistency
  - \* Simplified replication management

# MORE READING

- \* Common problems
- \* Improving features
- \* Multi-master replication
- \* Replication tools
- \* Extending replication
- \* Configuration
- \* Monitoring



# **MySQL Data Types**

# NUMBERS

\* ~~BIT~~

\* ~~TINYINT/BOOLEAN,~~

\* ~~INT,BIGINT,SMALLINT,MEDIUMINT~~

\* ~~FLOAT~~

\* ~~DOUBLE~~

\* ~~DECIMAL~~

# MONEY

## How do you store money?

\* FLOAT/DOUBLE

\* DECIMAL(13,2)

\* BIGINT

ISO Standard 13 digits for cents

\* SIGNED or UNSIGNED



# CHARACTER

\* CHAR

\* VARCHAR

\* ~~TEXT, TINYTEXT, MEDIUMTEXT, LONGTEXT~~

\* ENUM, SET



TEXT stores 64K

# BLOB

\* VARBINARY

\* ~~BLOB, TINYBLOB, MEDIUMBLOB, LONGBLOB~~

# DATE/TIME

\* DATE

\* TIME

\* YEAR

\* DATETIME

\* TIMESTAMP

Millisecond support (5.6+)

# TIMESTAMP

- \* Seconds since EPOCH
- \* Supports DEFAULT and ON UPDATE CURRENT\_TIMESTAMP syntax
- \* Complexities with sql\_mode and zero dates

# CONSTRAINTS

## \* UNSIGNED

\* Free negative number constraint

# CONSTRAINTS

## \* NOT NULL

- \* Define columns NOT NULL first
- \* Consider NULL when unknown

# CONSTRAINTS

## \* ENUM

- \* Set of known values

- \* 1 Byte length, yet offers long value results

# CONSTRAINTS

## \* ENUM

```
CREATE TABLE weather(  
  season ENUM('Spring', 'Summer', 'Autumn', 'Winter') NOT NULL  
);  
SET SESSION SQL_MODE='';  
INSERT INTO weather(season) VALUES ('Fall');  
Query OK, 1 row affected, 1 warning (0.00 sec)  
SELECT * FROM weather;  
+-----+  
| season |  
+-----+  
|       |  
+-----+  
  
SET SESSION SQL_MODE='STRICT_ALL_TABLES';  
ERROR 1265 (01000): Data truncated for column 'season' at row 1
```



# AUTOINCREMENT

- \* Use INT UNSIGNED not BIGINT
  - \* 4 bytes v 8 bytes
  - \* Doubles size of every secondary index



**What next?**

# MORE DISCUSSION

- \* Why you need a System Architect
- \* Using SSL (0.9.8 v 1.0.1)
- \* High Availability (HA) options
- \* Cloud Usage/Variants/Other solutions
- \* Integration with other products
- \* Development tools,resources,techniques
- \* ...

## TOP TIPS

- \* Is data integrity important?
- \* Scale out and HA can be complex
- \* Develop your developers
- \* Get to Version 5.5/5.6 ASAP

# **25+ YEARS EXPERIENCE**

**“To right certain wrongs we must look at the foundations. Simple and correct decisions have a momentous impact in the productivity of resources and success of business systems”.**

**Ronald Bradford  
August 2014**

# CONCLUSION

Copies of these slides can be found at

<http://effectiveMySQL.com>

[ronald@EffectiveMySQL.com](mailto:ronald@EffectiveMySQL.com)