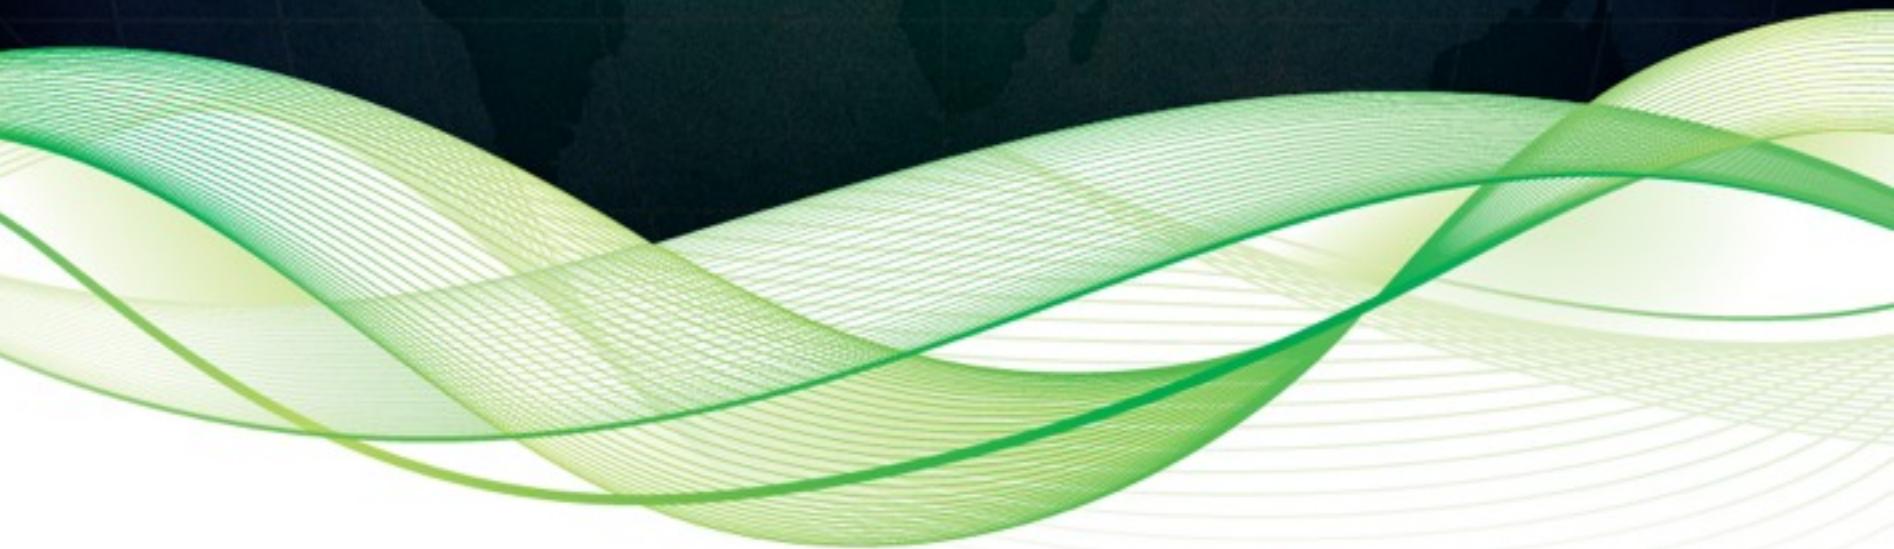


# Operational DBA In A Nutshell



Frédéric Descamps (@lefred)

Liz van Dijk (@lizztheblizz)

Kenny Gryp (@gryp)



# Introduction

2

- Welcome!
- Who are we?
  - Fans of devopsreactions.tumblr.com
- What will we talk about?



# Basic Housekeeping

3

- Full Day Tutorial
- HandsOn!, Laptop Required
- Breaks
- Wireless SSID = ??
- How to ask for help
- Using the documentation

# Beware!

4

- I know what `replicate_same_server_id` is!
- I know all GRANTS by heart!
- `pt-query-digest` is my favorite tool!
- I've been dealing with replication inconsistencies for a while already!
- My backup environment is the most awesome there is!

# Beware!

- I know what `replicate_same_server_id` is!
- I know all GRANTS by heart!
- `pt-query-digest` is my favorite tool!
- I've been dealing with replication inconsistencies for a while already!
- My backup environment is the most awesome there is!

**TAKE THE DAY OFF  
AND GET SOME BEER**

# Operational DBA In A Nutshell

- Setting Up Today's Environment
- Installation
- Logging In
- MySQL Privileges
- Diagnostics
- Troubleshooting
- Monitoring
- Backups
- Replication
- Schema Changes
- Configuration Optimization

How doing operations feels



# Operational DBA In A Nutshell

- Setting Up Today's Environment
- Installation
- Logging In
- MySQL Privileges
- Diagnostics
- Troubleshooting
- Monitoring
- Backups
- Replication
- Schema Changes
- Configuration Optimization



How doing operations feels



# Today's Application - Basic IMDB Interface

7

Login Signup

## MyMovies

There are 50005 users, 1543720 movies and 2812743 actors in the system.

**Random Movie**  
[Entre Calais et Douvres \(1897\)](#)

**Latest comments**

**Featured User**  
[Larry.B.Fargo@dodgit.com](#)

**Online right now**

**Being viewed right now**

Paul Calderon  
[Pulp Fiction \(1994\)](#)  
[Tight Security \(2001\)](#)  
[Assrar fi rammel \(1954\)](#)  
John Travolta

# Setting Up Our VM

- Copy Files From USB Stick
- Install VirtualBox, Import Virtual Machine
- Change ~/Library/VirtualBox/VirtualBox.xml
  - ~/.virtualbox/VirtualBox.xml

```
<MachineRegistry>
  <MachineEntry uuid="{d4a72cfa-d3f1-4d8e-8e32-f72315c48faf}" src="/Users/gryp/
VirtualBox VMs/PLMCE Slave (32bit)/PLMCE Slave (32bit).vbox"/>
  <MachineEntry uuid="{50109a84-d2f1-4992-b359-b1c1b40fbdbda}" src="/Users/gryp/
VirtualBox VMs/PLMCE master (32bit)/PLMCE master (32bit).vbox"/>
</MachineRegistry>
```

**{50109a84-d2f1-4992-b359-b1c1b40fbdbda}**

- Test Connectivity
  - ssh -p 2221 root@localhost (password: vagrant)
  - <http://localhost:8080/> (should see an apache test page)

# Operational DBA In A Nutshell

- Setting Up Today's Environment
- Installation
- Logging In
- MySQL Privileges
- Diagnostics
- Troubleshooting
- Monitoring
- Backups
- Replication
- Schema Changes
- Configuration Optimization



How doing operations feels



# Installation Of MySQL - HandsOn!

10

- Install Percona Server

```
# yum install Percona-Server-server-55  
Percona-Server-client-55
```

- Verify if you can connect

```
# mysql
```

- Verify the application:

- <http://localhost:8080/my-movies/>



gifbin.com

```
[root@node1 ~]# mysql  
Welcome to the MySQL monitor. C  
Your MySQL connection id is 270  
Server version: 5.5.30-30.2-log  
500
```

Copyright (c) 2000, 2013, Oracle

Oracle is a registered trademark  
affiliates. Other names may be t  
owners.

```
node1 mysql>
```

# Operational DBA In A Nutshell

- Setting Up Today's Environment
- Installation
- **Logging In**
- MySQL Privileges
- Diagnostics
- Troubleshooting
- Monitoring
- Backups
- Replication
- Schema Changes
- Configuration Optimization



How doing operations feels



# Logging In - HandsOn!

12

```
$ mysql  
mysql> SHOW DATABASES;  
mysql> USE imdb;  
mysql> SHOW TABLES;  
mysql> SELECT * FROM users LIMIT 1;  
mysql> SHOW PROCESSLIST;
```

# Operational DBA In A Nutshell

- Setting Up Today's Environment
- Installation
- Logging In
- MySQL Privileges
- Diagnostics
- Troubleshooting
- Monitoring
- Backups
- Replication
- Schema Changes
- Configuration Optimization



How doing operations feels



# MySQL Privileges

14

- Users
- Grants
- mysql database
- pt-show-grants
- Default Permissions
- Creating Application user

Asking the security team for a firewall exception



IWORM

PERCONA  
LIVE

# Users

15

- Identify users based on: user@host
  - user: username
  - host: hostname/ip of the client that connects
    - different host, different user, different ‘grants’
- Examples:

```
'fred'@'localhost',      'root'@'localhost'  
'lefred'@'app0001',     'dijkske'@'192.168.%'  
'ledijkske'@'192.168.1__', 'fredjen'@'app.fq.dn'
```
- Creating A User:

```
>CREATE USER 'lefred'@'app0001';
```
- Drop user: change CREATE into DROP

# Grants

16

- Grant the user some kind of privilege
- Grant ... to:

server,	column,	view,
database,	trigger,	index
table,	stored procedure,	
- Example: **INSERT**, **SELECT**, **UPDATE**, **DELETE**
- SQL Command:
  - >GRANT SELECT ON db.\* TO 'lefred'@'app0001';
  - >GRANT INSERT ON \*.\* TO 'lefred'@'app0001';
- Revoking privileges: change GRANT into REVOKE

# Table/Column Level Grants

17

- Possible:
  - > GRANT SELECT ON db.table TO 'lefred'@'app';
  - > GRANT SELECT (col) ON db.table to 'fr'@'app';
- Too much columns might make authentication slower
- Not commonly used

# Password

18

```
> SET PASSWORD FOR 'lefred'@'app0001' =  
PASSWORD('pass');  
  
> SELECT PASSWORD('pass')\G  
PASSWORD('pass'): *196BDEDE2AE4F84CA44C47D54D78478C7E2BD7B7  
  
> SET PASSWORD FOR 'lefred'@'app0001' =  
'*196BDEDE2...';  
  
> CREATE USER 'lefred'@'app0001' IDENTIFIED  
BY 'pass';  
> CREATE USER 'fred'@'app' IDENTIFIED BY  
PASSWORD '*196BDEDE2...';  
  
> GRANT SELECT ON db.* TO 'fred'@'app'  
IDENTIFIED BY 'pass';
```

# Grants

19

- Complete list of grants:

CREATE  
DROP  
GRANT OPTION  
LOCK TABLES  
EVENT  
ALTER  
DELETE  
INDEX  
INSERT  
SELECT  
UPDATE  
CREATE TEMPORARY TABLES  
TRIGGER  
CREATE VIEW  
SHOW VIEW

ALTER ROUTINE  
CREATE ROUTINE  
EXECUTE  
FILE  
CREATE USER  
PROCESS  
PROXY  
RELOAD  
REPLICATION CLIENT  
REPLICATION SLAVE  
SHOW DATABASES  
SHUTDOWN  
SUPER  
ALL [PRIVILEGES]  
USAGE

# SHOW GRANTS

20

```
> SHOW GRANTS;
+-----+
| Grants for root@localhost
+-----+
| GRANT ALL PRIVILEGES ON *.* TO 'root'@'localhost' IDENTIFIED BY
| PASSWORD '*81F5E21E35407D884A6CD4A731AEBFB6AF209E1B' WITH GRANT
| OPTION
+-----+

> SHOW GRANTS FOR 'lefred'@'app0001';
+-----+
| Grants for lefred@app0001
+-----+
| GRANT INSERT ON *.* TO 'lefred'@'app0001'
| GRANT SELECT ON `db`.* TO 'lefred'@'app0001'
+-----+
```

# mysql Database

21

```
node1 mysql> show tables;
+-----+-----+
| Tables_in_mysql | plugin
| columns_priv    | proc
| db              | procs_priv
| event           | proxies_priv
| func            | servers
| general_log     | slow_log
| help_category   | tables_priv
| help_keyword    | time_zone
| help_relation   | time_zone_leap_second
| help_topic      | time_zone_name
| host            | time_zone_transition
| ndb_binlog_index| time_zone_transition_type
| user            |
+-----+
24 rows in set (0.00 sec)
```

# mysql Database

22

- Do not give rights for app or general users to the MySQL database
- DML statements are possible on mysql database
  - use FLUSH PRIVILEGES to apply

# MySQL Privileges - HandsOn!

23

- Install Percona Toolkit

```
# yum install percona-toolkit  
# pt-show-grants
```

- show all grants on a system (in order)
  - good for version control

```
[root@node1 lib]# pt-show-grants  
-- Grants dumped by pt-show-grants  
-- Dumped from server Localhost via UNIX socket, MySQL  
5.5.30-30.2-log at 2013-04-19 00:18:28  
-- Grants for ''@'localhost'  
GRANT USAGE ON *.* TO ''@'localhost';  
-- Grants for ''@'node1'  
GRANT USAGE ON *.* TO ''@'node1';  
-- Grants for 'cactiuser'@'localhost'  
GRANT USAGE ON *.* TO 'cactiuser'@'localhost' IDENTIFIED BY  
PASSWORD '*43DD7940383044FBDE5B177730FAD3405BC6DAD7';
```

# MySQL Privileges - HandsOn!

```
[root@node1 lib]# pt-show-grants
-- Grants dumped by pt-show-grants
-- Dumped from server Localhost via UNIX socket, MySQL
5.5.30-30.2-log at 2013-04-19 00:18:28
-- Grants for ''@'localhost'
GRANT USAGE ON *.* TO ''@'localhost';
-- Grants for ''@'node1'
GRANT USAGE ON *.* TO ''@'node1';
-- Grants for 'cactiuser'@'localhost'
GRANT USAGE ON *.* TO 'cactiuser'@'localhost' IDENTIFIED BY
PASSWORD '*43DD7940383044FBDE5B177730FAD3405BC6DAD7';
GRANT ALL PRIVILEGES ON `cacti`.* TO 'cactiuser'@'localhost';
-- Grants for 'root'@'127.0.0.1'
GRANT ALL PRIVILEGES ON *.* TO 'root'@'127.0.0.1' WITH GRANT
OPTION;
-- Grants for 'root'@'::1'
GRANT ALL PRIVILEGES ON *.* TO 'root'@'::1' WITH GRANT
OPTION;
-- Grants for 'root'@'localhost'
```

# Default Permissions Are Bad - HandsOn!

25

```
# mysql_secure_installation
```

```
Enter current password for root (enter for none):
```

```
OK, successfully used password, moving on...
```

```
Set root password? [Y/n] y
```

```
New password: root
```

```
Re-enter new password: root
```

```
Remove anonymous users? [Y/n] y
```

```
Disallow root login remotely? [Y/n] y
```

```
Remove test database and access to it? [Y/n] y
```

```
Reload privilege tables now? [Y/n] y
```

```
All done! If you've completed all of the above steps, your  
MySQL installation should now be secure.
```

```
Thanks for using MySQL!
```

# MySQL Privileges - HandsOn!

26

- Now we have 'secure' privileges

```
$ mysql -u root -p  
create a ~/.my.cnf
```

- Create a user for 'my-movies' application

```
mysql> GRANT ALL PRIVILEGES ON imbd.* TO  
'mymovies'@'localhost' IDENTIFIED BY 'password';
```

Change user&password:

```
/var/www/html/my-movies/lib/config.inc.php
```

- Verify Application:

```
http://localhost:8080/my-movies
```

- Doublecheck with pt-show-grants

```
# pt-show-grants -u root -proot
```

# Operational DBA In A Nutshell

- Setting Up Today's Environment
  - Installation
  - Logging In
  - MySQL Privileges
  - **Diagnostics**
  - **Troubleshooting**
  - **Monitoring**
  - Backups
- Replication
  - Schema Changes
  - Configuration Optimization



How doing operations feels



# Troubleshooting

28

- Diagnostics
  - Operating System
  - MySQL
- Queries
- Percona Toolkit
- Trending
- Alerting

Troubleshooting a problem that I don't understand



# Add Some Load First - HandsOn!

29

- Open up New Terminal window  
\$ /home/percona/add\_load.py

# Diagnostics - HandsOn!

30

- MySQL:
  - Errorlog: /var/lib/mysql/node1.err
  - SHOW GLOBAL STATUS
  - mysqladmin extended-status
  - SHOW ENGINE INNODB STATUS
  - SHOW PROCESSLIST
- OS metrics (memory, cpu...)
  - vmstat, iostat, mpstat, meminfo, top, ps
  - Trending (Cacti and newer (-: )
  - Upper Layer Instrumentation (Application, Front-end, Cacheing)

# Percona Toolkit

31

- pt-mext
- pt-diskstats
- pt-query-digest

# pt-mext - HandsOn!

32

\$ pt-mext -r -- mysqladmin ext -i 10 -c 3			
Binlog_cache_disk_use	0	0	0
Binlog_cache_use	0	0	0
Bytes_received	2875788973602	1738235	346057
Bytes_sent	863929033790	588078	536398
Com_begin	6298644573	3516	5102
Com_delete	23721852	26	51
Com_insert	4454794705	1518	3287
Com_replace	527848577	197	121
Com_select	6993291133	8114	7594
Com_set_option	5112076	250	262
Connections	7331059	250	262
Created_tmp_disk_tables	113568	0	0
Created_tmp_files	7803	0	0
Created_tmp_tables	729281259	1816	479

# pt-mext

33

Handler_commit	4002481284	5295	4911
Handler_delete	7256841	10	25
Handler_discover	0	0	0
Handler_prepare	0	0	0
Handler_read_first	47274	0	0
Handler_read_key	42993091324	34920	27522
Handler_read_next	19633194815	16911	10142
Handler_read_prev	2440127	0	0
Handler_read_rnd	488760449	40	12
Handler_read_rnd_next	2731205271	268	231
Handler_rollback	5781	0	0
Handler_savepoint	0	0	0
Handler_savepoint_rollback	0	0	0
Handler_update	7022320034	10047	3329
Handler_write	7334430104	1945	3638

# pt-mext

34

<u>Qcache_free_blocks</u>	2899	100	-15
<u>Qcache_free_memory</u>	519642808	164104	-8080
<u>Qcache_hits</u>	325634530	0	0
<u>Qcache_inserts</u>	978847229	194	104
<u>Qcache_lowmem_prunes</u>	19158357	0	0
<u>Qcache_not_cached</u>	211301010	806	798
<u>Qcache_queries_in_cache</u>	3677	-112	9
<u>Qcache_total_blocks</u>	10277	-131	6
<u>Threads_cached</u>	9	1	0
<u>Threads_connected</u>	11	-1	0
<u>Threads_created</u>	294	0	0
<u>Threads_running</u>	5	-3	0
<u>Uptime</u>	21912350	10	10

- Look at current global behavior of database
- Query Optimization necessary? (sorting\_%, handler\_%, range\_%, tmp\_table\_%)
- Innodb misbehaving?
- ...

# Disk Subsystem Statistics

36

- *iostat* commonly used:

```
$ iostat -x -m 1
```

Device:	rrqm/s	wrqm/s	r/s	w/s	rsec/s	wsec/s	avgrq-sz	avgqu-sz	await	svctm	%util
sda	0.00	73.27	0.00	54.46	0.00	1061.39	19.49	4.84	88.80	18.36	100.00
sda1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
sda2	0.00	73.27	0.00	54.46	0.00	1061.39	19.49	4.84	88.80	18.36	100.00
sdb	0.00	451.49	0.99	<b>338.61</b>	7.92	6368.32	18.78	144.23	<b>420.93</b>	<b>2.94</b>	100.00
sdb1	0.00	451.49	0.99	338.61	7.92	6368.32	18.78	144.23	420.93	2.94	100.00

- %util: how many % of time at least one request was busy
- await+svctm:
  - response time
  - writes and reads combined

# pt-diskstats - HandsOn!

37

- reads */proc/diskstats*, shows wr/rd response time

```
$ pt-diskstats --devices-regex sdb1
```

<i>device</i>	<i>rd_s</i>	<i>rd_avkb</i>	<i>rd_mb_s</i>	<i>rd_mrg</i>	<i>rd_cnc</i>	<i>rd_rt</i>
<i>sda</i>	0.1	4.0	0.0	0%	0.0	5.0
<i>sda2</i>	0.1	4.0	0.0	0%	0.0	5.0
<i>sdb</i>	119.8	5.3	0.6	0%	0.5	4.1
<i>sdb1</i>	119.8	5.3	0.6	0%	0.5	4.1

# pt-diskstats - HandsOn!

38

<i>device</i>	<i>wr_s</i>	<i>wr_avkb</i>	<i>wr_mb_s</i>	<i>wr_mrg</i>	<i>wr_cnc</i>	<i>wr_rt</i>
<i>sda</i>	<b>23.5</b>	<b>35.6</b>	<b>0.8</b>	<b>89%</b>	<b>1.2</b>	<b>5.9</b>
<i>sda2</i>	<b>23.5</b>	<b>35.6</b>	<b>0.8</b>	<b>89%</b>	<b>1.2</b>	<b>5.9</b>
<i>sdb</i>	<b>160.3</b>	<b>7.5</b>	<b>1.2</b>	<b>47%</b>	<b>18.3</b>	<b>61.0</b>
<i>sdb1</i>	<b>160.3</b>	<b>7.5</b>	<b>1.2</b>	<b>47%</b>	<b>18.3</b>	<b>61.0</b>

# pt-diskstats - HandsOn!

39

<i>device</i>	<i>busy</i>	<i>in_prg</i>	<i>io_s</i>	<i>qtime</i>	<i>stim</i>
<b>sda</b>	<b>7%</b>	<b>0</b>	<b>23.6</b>	<b>5.6</b>	<b>0.3</b>
<b>sda2</b>	<b>7%</b>	<b>0</b>	<b>23.6</b>	<b>5.6</b>	<b>0.3</b>
<b>sdb</b>	<b>47%</b>	<b>0</b>	<b>280.0</b>	<b>43.6</b>	<b>1.1</b>
<b>sdb1</b>	<b>47%</b>	<b>0</b>	<b>280.0</b>	<b>43.6</b>	<b>1.1</b>

# Enhanced Slow Log Statistics - HandsOn!

40

```
> SET GLOBAL slow_query_log=on;
> SET GLOBAL slow_query_log_file='/var/
lib/mysql/slow.log';
> SET GLOBAL long_query_time=0;;
> SET GLOBAL log_slow_verbosity=full;
```

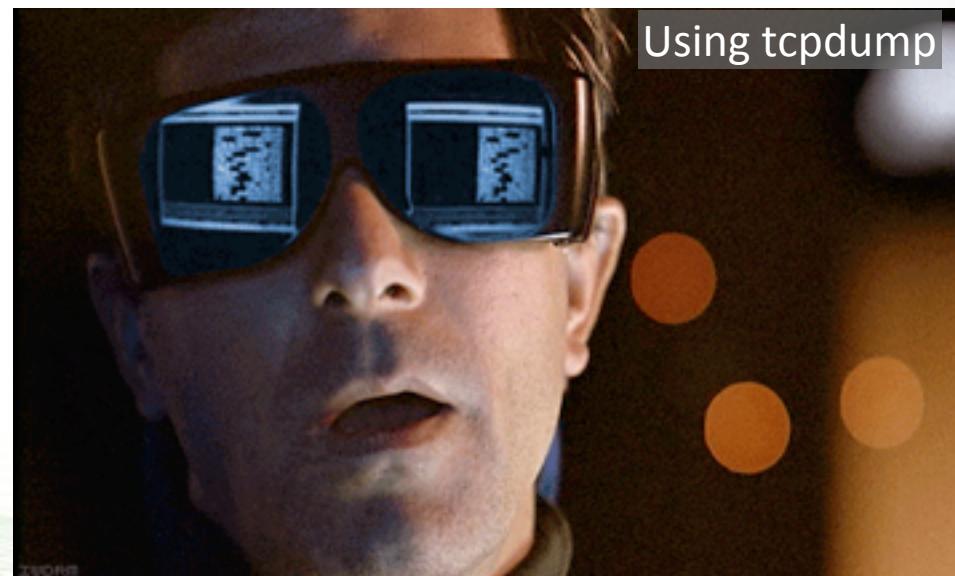
```
tail -n 200 /var/lib/mysql/slow.log
```

```
# Time: 100924 13:58:47
# User@Host: root[root] @ localhost []
# Thread_id: 10 Schema: imbd Last_errno: 0 Killed: 0
# Query_time: 399.563977 Lock_time: 0.000110 Rows_sent: 1 Rows_examined: 46313608
Rows_affected: 0 Rows_read: 1
# Bytes_sent: 131 Tmp_tables: 1 Tmp_disk_tables: 1 Tmp_table_sizes: 25194923
# InnoDB_trx_id: 1403
# QC_Hit: No Full_scan: Yes Full_join: No Tmp_table: Yes Tmp_table_on_disk: Yes
# Filesort: Yes Filesort_on_disk: Yes Merge_passes: 5
# InnoDB_IO_r_ops: 1064749 InnoDB_IO_r_bytes: 17444847616
# InnoDB_IO_r_wait: 26.935662
# InnoDB_rec_lock_wait: 0.000000 InnoDB_queue_wait: 0.000000
# InnoDB_pages_distinct: 65329
SET timestamp=1285336727;
select STRAIGHT_JOIN count(*) as c, person_id FROM cast_info FORCE INDEX(person_id) INNER JOIN
```

# pt-query-digest - HandsOn!

41

- generate reports from
  - slow query log  
**pt-query-digest /var/lib/mysql/slow.log | less**
  - general log (not so useful)
  - tcpdump
- group-by& order-by: db/ip/host/query\_time:sum/max/min/count



# pt-query-digest - HandsOn!

42

```
# Profile
# Rank QID      Response time Calls  R/Call  Apdx  V/M Item
# === =========
#   1 ... 1349.6240 62.4%  11976 0.1127 1.00  0.03 SELECT table1 table9 table2
table3 table1
#   2 ... 114.9014 5.3%    437  0.2629 1.00  0.50 SELECT table5 table6 table8
table7 table6 table8 table10
#   3 ... 92.9441 4.3%    791  0.1175 1.00  0.06 SELECT table13
#   4 ... 77.5712 3.6%    43   1.8040 0.65  0.73 SELECT table11 table12
table9 table2 table14 table15 table16 table14 table17
#   5 ... 67.1673 3.1%    296  0.2269 1.00  0.17 SELECT table8 table4 table14
table8 table18
#   6 ... 49.0330 2.3%   15630 0.0031 1.00  0.00 ADMIN CONNECT
#   7 ... 43.4990 2.0%    274  0.1588 1.00  0.12 SELECT table19
#   8 ... 30.0898 1.4%    416  0.0723 1.00  0.07 SELECT table13
#   9 ... 19.6506 0.9%   13424 0.0015 1.00  0.01 UPDATE table20
```

# pt-query-digest - HandsOn!

43

```
# Query 1: 17.06 QPS, 1.92x concurrency, ID 0x3928FBFF36663F33 at byte 1417466467
# Attribute      pct      total      min      max      avg     95%    stddev   median
# ====== ====== ====== ====== ====== ====== ====== ====== ======
# Count          1    11976
# Exec time     62   1350s    25ms   395ms   113ms   219ms   54ms    91ms
# Rows affected    0      20      0      25      0.00      0      0.32      0
# Query size     23   28.75M   2.46k   2.46k   2.46k   2.38k      0   2.38k
# warning count   11   51.51k      0   12.80k   4.40      0   255.99      0
```

# pt-query-digest - HandsOn!

44

```
# Query_time distribution
#    1us
#    10us #####
#   100us #####
#    1ms ##
#   10ms #
# 100ms #####
#    1s
#   10s+
# Tables

#      SHOW TABLE STATUS LIKE 'table19' \G
#      SHOW CREATE TABLE `table19` \G
# EXPLAIN /*!50100 PARTITIONS*/
SELECT user_agent_id, search_engine FROM table19 WHERE user_agent='Mozilla/4.0
(compatible; MSIE 7.0; Windows NT 5.1; .NET CLR 1.0.3705) '\G
```

# pt-query-digest - HandsOn!

45

#	Attribute	pct	total	min	max	avg	95%	stddev	median
#	# Count	0	519						
#	# Exec time	2	148s	11us	33s	285ms	53ms	2s	26us
#	# Lock time	0	5ms	0	334us	9us	66us	32us	0
#	# Rows sent	0	41	0	1	0.08	0.99	0.27	0
#	# Rows examine	1	4.97M	0	445.49k	9.80k	5.73k	49.33k	0
#	# Rows affected	0	2	0	1	0.00	0	0.06	0
#	# Rows read	1	2.01M	0	250.47k	3.96k	1.96	27.94k	0.99
#	# Bytes sent	0	241.20k	11	8.01k	475.89	918.49	689.98	258.32
#	# Merge passes	0	0	0	0	0	0	0	0
#	# Tmp tables	0	15	0	1	0.03	0	0.17	0
#	# Tmp disk tbl	0	3	0	1	0.01	0	0.08	0

# pt-query-digest - HandsOn!

46

# Tmp tbl size	0	4.78k	0	4.78k	9.43	0	211.60	0
# Query size	0	100.95k	19	2.71k	199.17	363.48	206.60	151.03
# InnoDB:								
# IO r bytes	0	0	0	0	0	0	0	0
# IO r ops	0	0	0	0	0	0	0	0
# IO r wait	0	0	0	0	0	0	0	0
# pages distin	1	67.99k	0	10.64k	1.26k	3.88k	2.47k	31.70
# queue wait	0	0	0	0	0	0	0	0
# rec lock wai	0	0	0	0	0	0	0	0
# Boolean:								
# Filesort	0%	yes,	99%	no				
# Full scan	7%	yes,	92%	no				
# QC Hit	78%	yes,	21%	no				
# Tmp table	2%	yes,	97%	no				
# Tmp table on	0%	yes,	99%	no				

# pt-query-digest - HandsOn!

47

- store reports in db: **--review, --review-table8**
- enhanced filtering capabilities
  - '\$event->{fingerprint} =~ m/^select/'
  - '\$event->{Warning\_count} > 1'
  - '\$event->{InnoDB\_IO\_r\_ops} > 50'
  - '\$event->{QC\_hit} = "Yes"'
  - '\$event->{Bytes} >= m/^1\_048\_576/'

# Trending

48

- Go to the cacti installation:

<http://localhost:8080/cacti/>

user: admin  
pass: plmce

- Graphs ->
  - Preset: Last Half Hour
  - Graphs Per Page: 100
  - GO!

# Optimizing This Application

- The problem is that the database is not optimized, indexing is not done properly
- We have Some Suggestions:

```
$ less /home/percona/.data/billkarwin-opt.sql
```

- Apply the changes:  

```
$ cat ~percona/.data/billkarwin-opt.sql | mysql
```
- Have another look with pt-query-digest

```
$ cd /var/lib/mysql/; mv slow.log slow.log_old  
$ killall -HUP mysqld  
$ /home/percona/add_load.py  
$ pt-query-digest /var/lib/mysql/slow.log | less
```

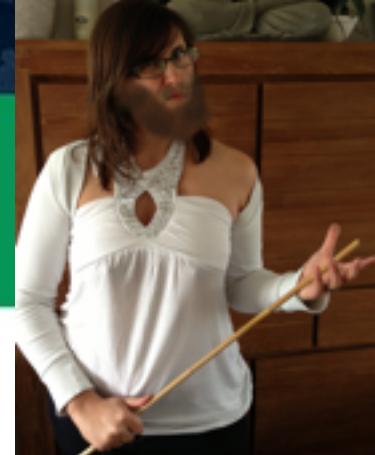
# Alerting

50

- Nagios, Zabbix, Icinga, ZenOss....
- Things To Alert on:
  - Can Connect to MySQL.
  - Replication IO Thread is OK
  - Replication SQL Thread is OK
  - Replication Seconds Behind Master.
  - Connections
  - InnoDB is enabled.
  - No tables are marked as corrupt.
- Not the best: tmp tables, ratios, qps
- Focus on what matters: usually in upper layers

# Operational DBA In A Nutshell

- Setting Up Today's Environment
- Installation
- Logging In
- MySQL Privileges
- Diagnostics
- Troubleshooting
- Monitoring
- Backups
- Replication
- Schema Changes
- Configuration Optimization



How doing operations feels



# Backing up your Data - Challenges

52

- MySQL is multi-engine
  - Different storage engines call for different strategies
- MySQL can be quite busy
  - Do we need to shut it down?
  - How much load will we generate?
  - How quickly can we restore?

# Backing up your Data - Concepts

53

- Impact of Backup
  - **Hot backup** - does not prevent readers or writers from performing operations.
  - **Warm backup** - readers may continue, but write activity must be queued until the backup is completed.
  - **Cold backup** - system is unavailable during backup window.

# Backing up your Data - Concepts

54

- Type of Backup
  - **Physical** - stores data in the “native” format.
  - **Logical** - stores data in a generic representation, such as a SQL dump.

# A few available backup solutions

55

	Impact	Warmth	Backup Time	Restore Time
Cold Backup	very high	cold	very fast	fast
<b>mysqldump</b>	high	warm	medium	slow
Snapshotting	high/medium	<b>hot/warm</b>	fast	<b>fast</b>
MySQL Enterprise Backup	<b>low/medium</b>	warm	fast	<b>fast</b>
XtraBackup	<b>low/medium</b>	warm	fast	<b>fast</b>

# mysqldump

56

- Included by default with MySQL
- Logical backup solution - creates SQL file to reproduce data
- Pro's
  - Easy to use, reliable
- Cons
  - High performance impact, long restore time

# mysqldump - HandsOn!

57

- Take backup:

```
# mysqldump --all-databases --single-transaction  
> /var/backup/dump.sql
```

- --all-databases for every database (including mysql).
- --single-transaction is only safe when all InnoDB.
- This ensures a hot backup.
- --master-data=1 (optional) records the binary log coordinates.

- Restore as follows:

```
# mysql < dump.sql
```

# mysqldump - HandsOn!

```
[root@node1 ~]# vmstat 5
procs -----memory----- ---swap-- -----io---- --system-- -----cpu-----
 r b swpd free buff cache si so bi bo in cs us sy id wa st
0 0 4316 3844 936 74820 2 9 8438 7121 501 1810 17 12 64 7 0
0 3 4316 4020 1092 74504 0 0 39 54 37 65 0 0 94 5 0
0 0 4316 3596 1112 75016 0 0 107 10 38 71 0 0 95 5 0
0 0 4316 3692 1112 75032 0 0 0 0 25 47 0 0 100 0 0
0 0 4316 3992 1156 74512 0 0 50 50 42 78 0 0 99 0 0
1 5 4316 10464 1768 67396 0 0 4816 0 273 1371 9 6 83 3 0
2 3 4316 4176 492 74600 0 0 20922 19623 1186 4066 57 35 0 8 0
2 3 4316 3080 508 75664 0 0 19482 15578 1089 5046 54 38 0 8 0
1 3 4540 3076 716 74828 0 45 17785 17338 1293 4411 54 36 0 10 0
2 4 6068 3856 736 71172 0 306 15406 10082 1256 4365 62 35 0 2 0
3 3 8280 3048 184 66460 0 0 26261 19267 1135 4789 51 35 0 14 0
1 2 10036 4144 452 66584 0 351 35832 31802 1472 1981 53 34 0 12 0
0 2 10180 4088 200 67304 0 29 38532 30722 1376 2997 56 35 0 9 0
2 2 12468 4140 148 69028 6 77 29004 26610 1387 2785 60 34 0 6 0
1 0 14628 4328 2552 59096 0 432 21390 15741 1397 4629 51 35 0 14 0
0 0 14628 20048 2868 61040 0 0 396 71 635 2187 26 19 54 1 0
0 0 14628 19932 2884 61220 0 0 34 16 36 58 0 0 99 1 0
0 0 14628 19628 2916 61588 0 0 72 54 39 74 0 0 90 10 0
0 0 14624 19564 2940 61636 6 0 16 11 34 62 0 0 100 0 0
0 0 14624 19620 2964 61636 0 0 0 41 34 62 0 0 90 10 0
0 0 14624 18900 2988 62308 0 0 137 34 43 80 0 0 99 1 0
```

# mysqldump - HandsOn!

```
[root@node1 ~]# vmstat 5
procs -----memory-----  ---swap--  -----io----  --system--  -----cpu-----
 r b swpd free buff cache si so bi bo in cs us sy id wa st
 0 0 4316 3844 936 74820 2 9 8438 7121 501 1810 17 12 64 7 0
 0 3 4316 4020 1092 74504 0 0 39 54 37 65 0 0 94 5 0
 0 0 4316 3596 1112 75016 0 0 107 10 38 71 0 0 95 5 0
 0 0 4316 3692 1112 75032 0 0 0 0 25 47 0 0 100 0 0
 0 0 4316 3992 1156 74512 0 0 50 50 42 78 0 0 99 0 0
 1 5 4316 10464 1768 67396 0 0 4816 0 273 1371 9 6 83 3 0
 2 3 4316 4176 492 74600 0 0
 2 3 4316 3080 508 75664 0 0
 1 3 4540 3076 716 74828 0 45
 2 4 6068 3856 736 71172 0 306
 3 3 8280 3048 184 66460 0 0
 1 2 10036 4144 452 66584 0 351
 0 2 10180 4088 200 67304 0 29
 2 2 12468 4140 148 69028 6 77
 1 0 14628 4328 2552 59096 0 432
 0 0 14628 20048 2868 61040 0 0
 0 0 14628 19932 2884 61220 0 0
 0 0 14628 19628 2916 61588 0 0
 0 0 14624 19564 2940 61636 6 0
 0 0 14624 19620 2964 61636 0 0
 0 0 14624 18900 2988 62308 0 0
```

Waiting on MySQLdump to Finish



# XtraBackup

59

- Completely free and open source
- Physical backup solution - Stores data in similar format as the original.
- Pro's
  - Powerful, scriptable, feature-packed
- Cons
  - Multiple steps required to restore, complexity

# XtraBackup - Features

60

- Non-blocking
- Support for MyISAM
- Compression
- Partial Backups
- Throttling
- Incremental Backups
- Exporting/Importing individual tables
- Streaming
- Parallel copying
- ...

# XtraBackup - Operation

61

- 3 separate "manual" steps are required:
  - (1) - Backing up
  - (2) - Preparing InnoDB files for recovery
  - (3) - Restoring the files

# XtraBackup - Operation (2)

62

- Performing and restoring a mixed full backup:
  - Backing up (1):
    - Copy datafiles while recording transaction log changes
  - FLUSH TABLES WITH READ LOCK;
  - Get binlog position
  - Copy all .MYD, .MYI, .TRG, .TRN, ... files
  - Stop recording transaction log changes
  - UNLOCK TABLES;
- Preparing/Restoring:
  - Prepare backup by applying recorded tlog changes (2)
- Restoring files to original location (3)

# XtraBackup - HandsOn!

63

- Simple full backup and restore:

(1)

```
# innobackupex /var/backup
```

(2)

```
# innobackupex --apply-log /var/backup/
[backupfolder]/
```

(3)

```
# innobackupex --copy-back /var/backup/
[backupfolder]/
```

```
# chown -R mysql:mysql /var/lib/mysql
```

# XtraBackup - HandsOn!

64

## (1) - Full Backup

```
[root@node1 ~]# innobackupex /var/backup/
```

```
InnoDB Backup Utility v1.5.1-xtrabackup; Copyright 2003, 2009 Innobase Oy  
and Percona Ireland Ltd 2009-2012. All Rights Reserved.
```

```
This software is published under  
the GNU GENERAL PUBLIC LICENSE Version 2, June 1991.
```

```
130421 18:26:20 innobackupex: Starting mysql with options: --unbuffered --  
130421 18:26:20 innobackupex: Connected to database with mysql child process (pid=63)  
130421 18:26:26 innobackupex: Connection to database server closed  
IMPORTANT: Please check that the backup run completes successfully.  
At the end of a successful backup run innobackupex  
prints "completed OK!".
```

```
innobackupex: Using mysql Ver 14.14 Distrib 5.5.30, for Linux (i686) using readline 5  
innobackupex: Using mysql server version Copyright (c) 2000, 2013, Oracle and/or its  
All rights reserved.
```

```
innobackupex: Created backup directory /var/backup/2013-04-21_18-26-26  
130421 18:26:26 innobackupex: Starting mysql with options: --unbuffered --  
130421 18:26:26 innobackupex: Connected to database with mysql child process (pid=63)
```

# XtraBackup - HandsOn!

65

## Backup Folder Structure

```
[root@node1 2013-04-21_18-26-26]# ls  
backup  backup-my.cnf  cacti  ibdata1  imbd  mysql  performance_schema  
xtrabackup_binary  xtrabackup_binlog_info  xtrabackup_checkpoints  xtrabackup_logfile
```

## (2) - Applying the transaction logs

```
[root@node1 2013-04-21_18-26-26]# innobackupex --apply-log /var/backup/2013-04-21_18-26-26/
```

```
InnoDB Backup Utility v1.5.1-xtrabackup; Copyright 2003, 2009 Innobase Oy  
and Percona Ireland Ltd 2009-2012. All Rights Reserved.
```

```
This software is published under  
the GNU GENERAL PUBLIC LICENSE Version 2, June 1991.
```

```
IMPORTANT: Please check that the apply-log run completes successfully.  
At the end of a successful apply-log run innobackupex  
prints "completed OK!".
```

```
130421 18:34:38 innobackupex: Starting ibbackup with command: xtrabackup_55 --defaults-file="/var/backup/2013-04-21_18-26-26/backup-my.cnf" --defaults-group="mysqld" --prepare --target-dir=/var/lib/mysql tmpdir=/tmp
```

```
xtrabackup_55 version 2.0.6 for Percona Server 5.5.16 Linux (i686) (revision id: 521)  
xtrabackup: cd to /var/lib/mysql/backup/2013-04-21_18-26-26  
xtrabackup: This target seems to be not prepared yet.  
xtrabackup: xtrabackup_logfile detected: size=2097152, start_lsn=(8915647136)  
xtrabackup: Temporary instance for recovery is set as followings.
```

# Understanding Backup Requirements

66

- Recovery Time Objective - How quickly must I restore?
- Recovery Point Objective - How much data can I lose?
- Risks - What failures should this data be protected against?

# Advanced Backup Topics

67

- Incremental Backups
  - HandsOn!
- Delayed Slaves
  - Using pt-slave-delay to keep a slave behind on its master by a certain amount of time.
  - Good help against potential disastrous user mistakes
  - Example:

```
# pt-slave-delay --delay 1m --interval 15s
```

# Advanced Backup Topics - HandsOn!

68

- Incremental Backup:

1. Create a full backup

```
2. # innobackupex --incremental /var/backup/inc/ --incremental-
   basedir=/var/backup/[backupfolder]
```

- Incremental Restore:

```
3. # innobackupex --apply-log --redo-only /var/backup/[backupfolder]
```

- Preparing the base backup FIRST

```
4. # innobackupex --apply-log --redo-only /var/backup/[backupfolder] --
   incremental-basedir=/var/backup/inc/[incbackupfolder]
```

- Applying each incremental backup to the base

```
5. # innobackupex --apply-log /var/backup/[backupfolder]
```

- Finish preparing the now completed base backup

6. Restore as usual

# Advanced Backup Topics

69

- Point-In-Time recovery
  - Making use of the binary log
  - mysqlbinlog can be used to pipe commands into mysql
  - Use in combination with incremental backups to restore right up to the point where the disaster occurred.

# Advanced Backup Topics

70

- Checking backup binlog position

```
[root@node1 2013-04-21_18-58-22]# cat xtrabackup_binlog_info
node1-bin.000006      243118
```

- Checking binlog contents

```
[root@node1 mysql]# mysqlbinlog --start-position=243118 --stop-position=243300 node1-bin.000006
/*!50530 SET @@SESSION.PSEUDO_SLAVE_MODE=1*/;
/*!40019 SET @@session.max_insert_delayed_threads=0*/;
/*!50003 SET @OLD_COMPLETION_TYPE=@@COMPLETION_TYPE,COMPLETION_TYPE=0*/;
[...]
```

- Applying contents to mysql

```
[root@node1 mysql]# mysqlbinlog --start-position=243118 --stop-position=243300
node1-bin.000006 | mysql
```

# Advanced Backup Topics

71

- Managing & Retaining your backups
  - Storing offsite in case of disaster
  - Recovery testing
  - Retention policies
  - Monitoring Backup & Restore success rate

# XtraBackup - HandsOn!

72

- Please run the following commands to prep your environment for the next parts:

```
# cd /var/lib/mysql/backup  
# ln -s `ls -t1 | grep 2013* | head -n1` latest_full
```

# Operational DBA In A Nutshell

- Setting Up Today's Environment
  - Installation
  - Logging In
  - MySQL Privileges
  - Diagnostics
  - Troubleshooting
  - Monitoring
  - Backups
- Replication
  - Schema Changes
  - Configuration Optimization



How doing operations feels



# MySQL Replication

74

- Replication?
- Binary Logs
- Setting Up Replication
- Commands
- Breaking & Fixing Replication
- Inconsistencies

# What is Replication?

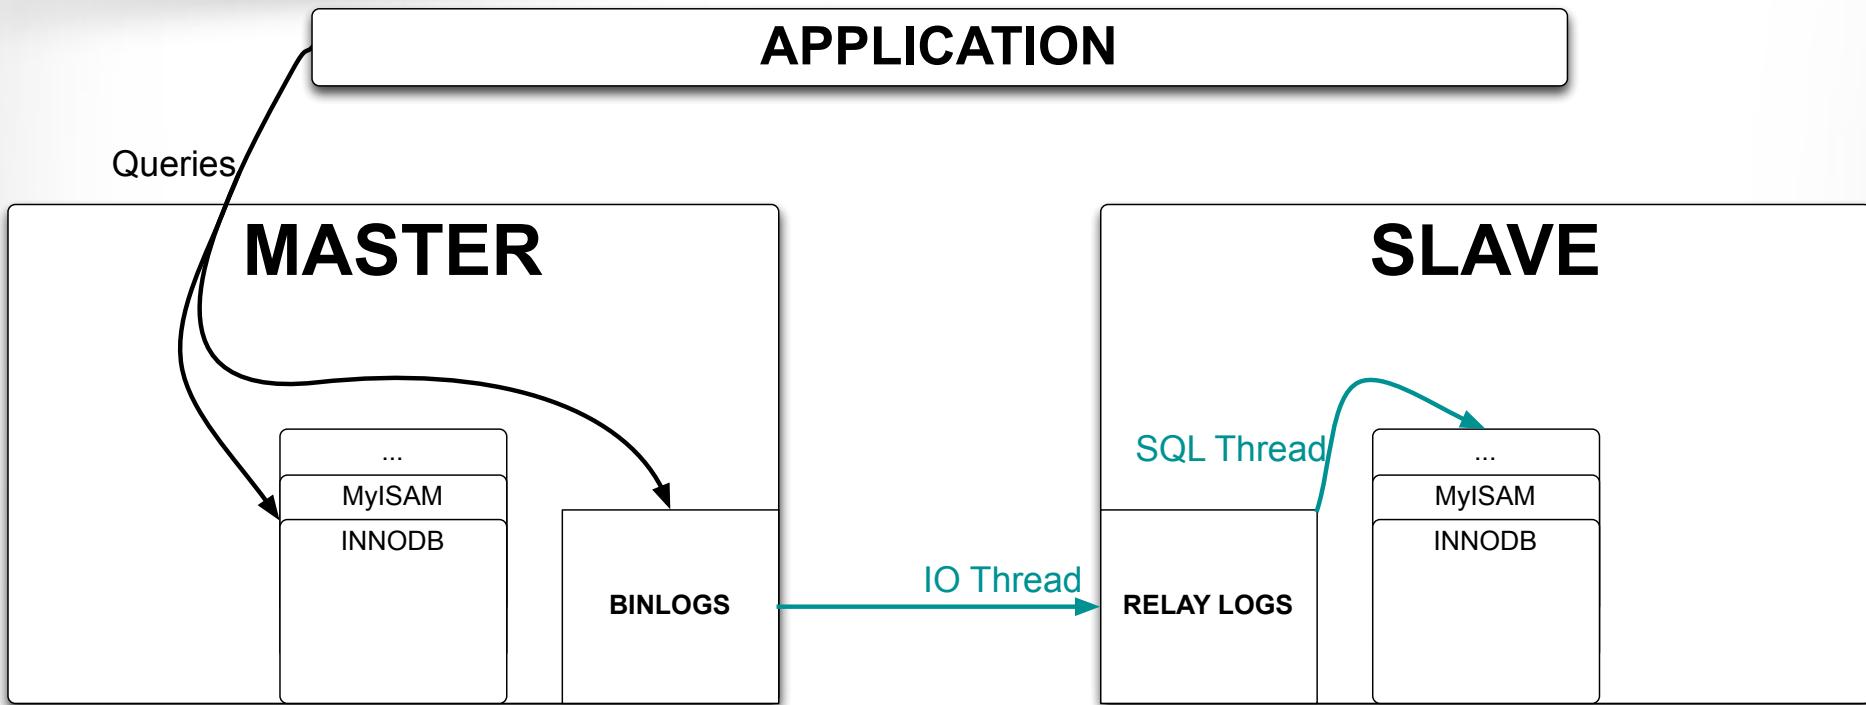
75

Replication enables data from one MySQL database server (the master) to be replicated to one or more MySQL database servers (the slaves)

(Source: <http://dev.mysql.com/doc/refman/5.5/en/replication.html>)

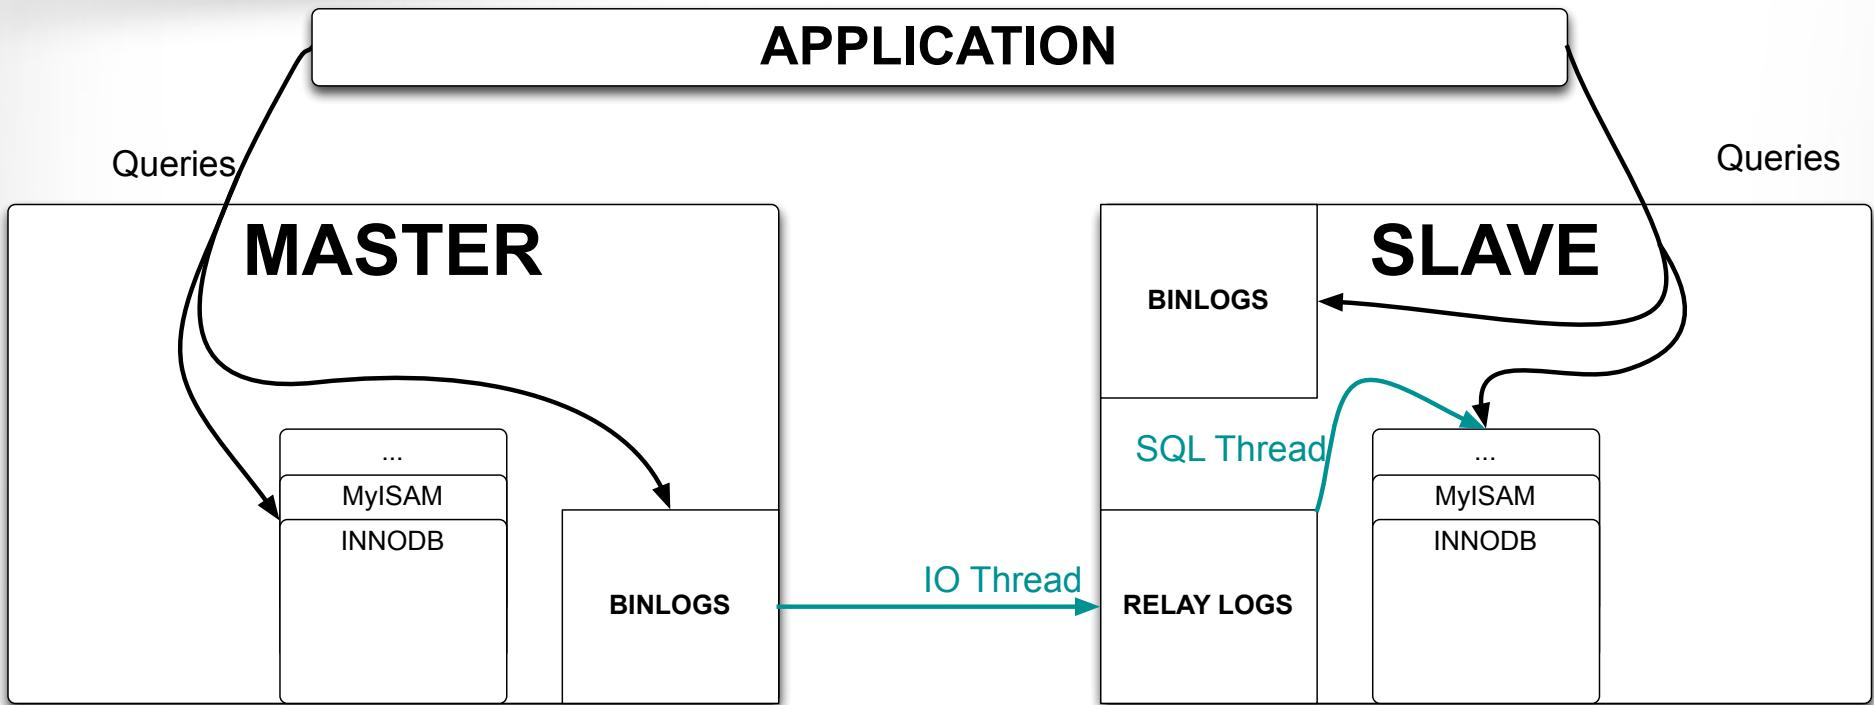
# Replication Diagram

76



# Replication Diagram

77



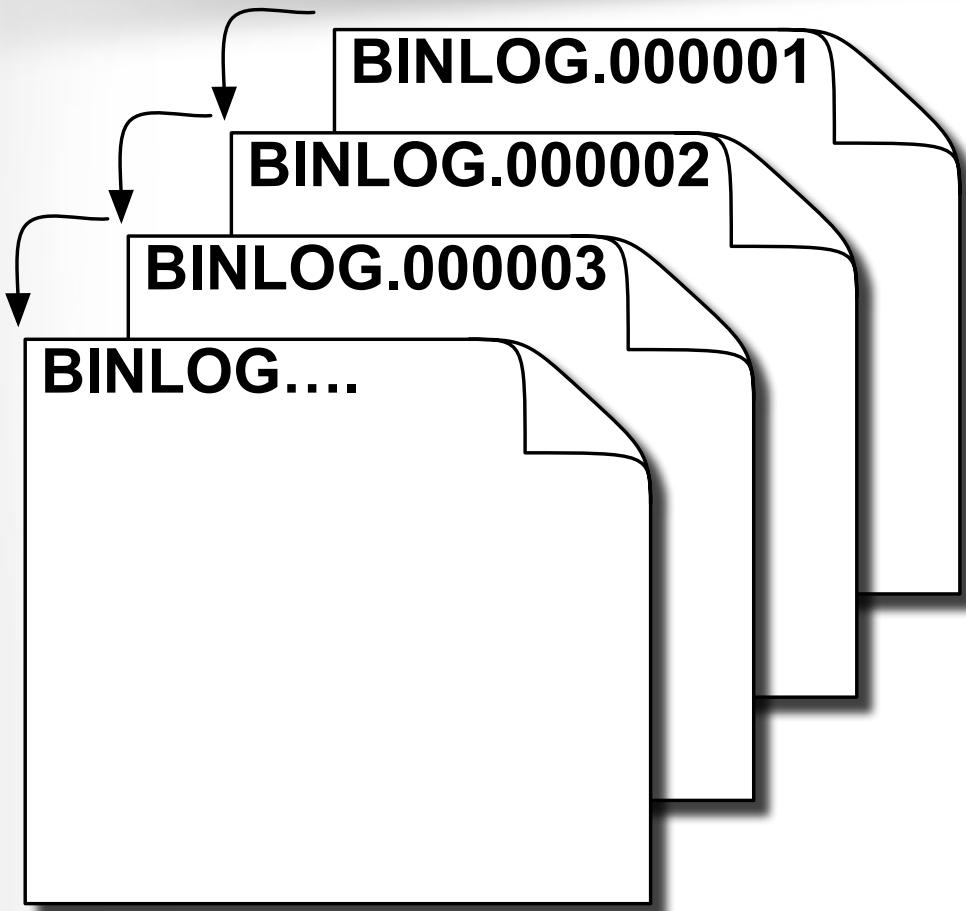
# Replication

78

- Happens at MySQL level, not Storage Engine Level
- Asynchronous! (Semi-sync available in 5.5)
- A server can have only 1 master
- IO Thread: Fetches from master
- SQL Thread: Executes on slave
- Single Threaded Execution (Until 5.6)
- Different schema's are possible between master and slave (Watch out!!):
  - different indexes
  - storage engines
  - data types, columns

# Binary Logs

79



**BINLOG.index**

BINLOG.000001  
BINLOG.000002  
BINLOG.000003

# Binary Logs

80

- Set of files
- Contains all writes and schema changes
- != REDO/Transaction log
- Rotated when full (Set `max_binlog_size`)
- Incrementing numbers  
(000001,000002,000003,...)
- Relay Logs are also binary logs
- 2 Formats:
  - Statement Based (SBR)
  - Row Based (RBR, since MySQL 5.1)
- Binary file

# Binary Logs - Info Files

81

- `master.info`
  - Contains IO Thread information & connection information
- `relay.info`
  - Contains SQL Thread information

# Binary Log - Formats (`binlog_format`)

82

- Statement Based Replication (SBR):
  - Writes statements to binary logs, slave executes the stmts  
(An update stmt that changes 1 row but reads many will also do the same on the slave)
  - Some functions are non-deterministic and cause inconsistencies:
    - `UUID()`, `SYSDATE()`, `FOUND_ROWS()`, `UPDATE ... LIMIT` without `ORDER BY...`
    - Works for `NOW()`: timestamp included in binary log
  - More complete list of issues:  
<http://dev.mysql.com/doc/refman/5.5/en/replication-features.html>

# Binary Log - Formats (`binlog_format`)

83

- Row Based Replication (RBR, since 5.1):
  - Write row changes (larger binlogs)
  - Check `Binlog_cache_disk_use`, possibly increase `binlog_cache_size`
  - Does not need to parse/execute queries, just make the changes necessary
  - Much less/different issues compared to SBR:  
<http://dev.mysql.com/doc/refman/5.5/en/replication-rbr-usage.html>
- Mixed: Combination of both: defaults to SBR, use RBR when necessary

# Looking at Binary Log Contents

84

- mysqlbinlog
- SHOW BINLOG EVENTS

# Example SBR

```
> SHOW GLOBAL VARIABLES LIKE 'binlog_format';
+-----+-----+
| Variable_name | Value      |
+-----+-----+
| binlog_format | STATEMENT |
+-----+-----+
1 row in set (0.00 sec)
```

```
> CREATE DATABASE replication;
Query OK, 1 row affected (0.14 sec)
```

```
> use replication
Database changed
```

```
> CREATE TABLE repl (a int) ENGINE=innodb;
Query OK, 0 rows affected (0.25 sec)
```

```
> INSERT INTO repl VALUES (1);
Query OK, 1 row affected (0.14 sec)
```

# Example SBR - mysqlbinlog

```
# mysqlbinlog mysql-bin.000193
...
# at 106
#120106 15:19:13 server id 9999    end_log_pos 203        Query
      thread_id=11
CREATE DATABASE replication
/*!*/;
# at 203
#120106 15:19:32 server id 9999    end_log_pos 312        Query
      thread_id=11 exec_time=1    error_code=0
use replication/*!*/;
SET TIMESTAMP=1325859572/*!*/;
CREATE TABLE repl (a INT) ENGINE=innodb
/*!*/;
# at 312
#120106 15:19:55 server id 9999    end_log_pos 387        Query
      thread_id=11 exec_time=0    error_code=0
SET TIMESTAMP=1325859595/*!*/;
BEGIN
/*!*/.
```

> SHOW BINLOG EVENTS FROM 106\G

\*\*\*\*\* 1. row \*\*\*\*\*

Log\_name: mysql-bin.000193

Pos: 106

Event\_type: Query

Server\_id: 1

End\_log\_pos: 203

Info: **CREATE DATABASE replication**

\*\*\*\*\* 2. row \*\*\*\*\*

Log\_name: mysql-bin.000193

Pos: 203

Event\_type: Query

Server\_id: 1

End\_log\_pos: 312

Info: use `replication`; CREATE TABLE repl (a INT)

**ENGINE=innodb**

\*\*\*\*\* 3. row \*\*\*\*\*

Log\_name: mysql-bin.000193

Pos: 312

Event\_type: Query

Server\_id: 1

End\_log\_pos: 387

Info: **BEGIN**

\*\*\*\*\* 4. row \*\*\*\*\*

# Example RBR

```
> SHOW GLOBAL VARIABLES LIKE 'binlog_format';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| binlog_format | ROW   |
+-----+-----+
1 row in set (0.00 sec)
```

```
> CREATE DATABASE replication;
Query OK, 1 row affected (0.14 sec)
```

```
> use replication
Database changed
```

```
> CREATE TABLE repl (a int) ENGINE=innodb;
Query OK, 0 rows affected (0.25 sec)
```

```
> INSERT INTO repl VALUES (1);
Query OK, 1 row affected (0.14 sec)
```

# Example RBR - mysqlbinlog

89

```
# mysqlbinlog mysql-bin.000193 --start-position=606
...
# at 606
#120106 15:54:54 server id 1    end_log_pos 703    Query
    thread_id=11  exec_time=0   error_code=0
CREATE DATABASE replication
/*!*/;
# at 703
#120106 15:55:02 server id 1    end_log_pos 812    Query
    thread_id=11  exec_time=0   error_code=0
use replication/*!*/;
SET TIMESTAMP=1325861702/*!*/;
CREATE TABLE repl (a int) ENGINE=innodb
/*!*/;
# at 812
...
# at 937
#120106 15:55:06 server id 1    end_log_pos 937    Table_map:
`replication`.`repl` mapped to number 17
#120106 15:55:06 server id 1    end_log_pos 971    Write_rows::table
```

# Example RBR - mysqlbinlog --verbose

90

```
# mysqlbinlog mysql-bin.000193 --verbose --verbose
...
# at 937
#120106 15:55:06 server id 1    end_log_pos 937    Table_map:
`replication`.`repl` mapped to number 17
#120106 15:55:06 server id 1    end_log_pos 971    Write_rows: table
id 17 flags: STMT_END_F

BINLOG '
SgsHTxMBAAAAMgAAAKkDAAAAABEAAAAAAEAC3JlcGxpY2F0aW9uAARyZXBsAAEDAA
E=
SgsHTxcBAAAIfgAAAMsDAAAAABEAAAAAAEAAf/+AQAAAA=='
/* */;
### INSERT INTO replication.repl
### SET
###   @1=1 /* INT meta=0 nullable=1 is_null=0 */
# at 971
```

> SHOW BINLOG EVENTS FROM 606\G

\*\*\*\*\* 1. row \*\*\*\*\*

Log\_name: mysql-bin.000193

Pos: 606

Event\_type: Query

Server\_id: 1

End\_log\_pos: 703

Info: **CREATE DATABASE replication**

\*\*\*\*\* 2. row \*\*\*\*\*

Log\_name: mysql-bin.000193

Pos: 703

Event\_type: Query

Server\_id: 1

End\_log\_pos: 812

Info: use `replication`; CREATE TABLE repl (a int)

**ENGINE=innodb**

...

\*\*\*\*\* 4. row \*\*\*\*\*

Log\_name: mysql-bin.000193

Pos: 887

Event\_type: Table\_map

Server\_id: 1

End\_log\_pos: 937

Info: table\_id: 17 (replication.repl)

# Setting up Replication

92

- Prerequisites
- Change master/slave MySQL configuration
- Configure Replication
- Start Replication/Check Status

# Prerequisites - HandsOn!

93

- No way to easily create slave with 1 command
- It's required to Create/Restore **consistent** backup
  - with mysqldump, use --master-data or xtrabackup
- Binlog file and position from backup should be recorded
  - Example
    - File: mysql-bin.000008, Pos: 106
- ```
slave# rm -rf /var/lib/mysql/*
slave# scp -r root@192.168.70.1:/var/backup/
latest_full /var/lib/mysql/
slave# chown -R mysql:mysql /var/lib/mysql/
```

# Change master/slave Configuration - HandsOn!

94

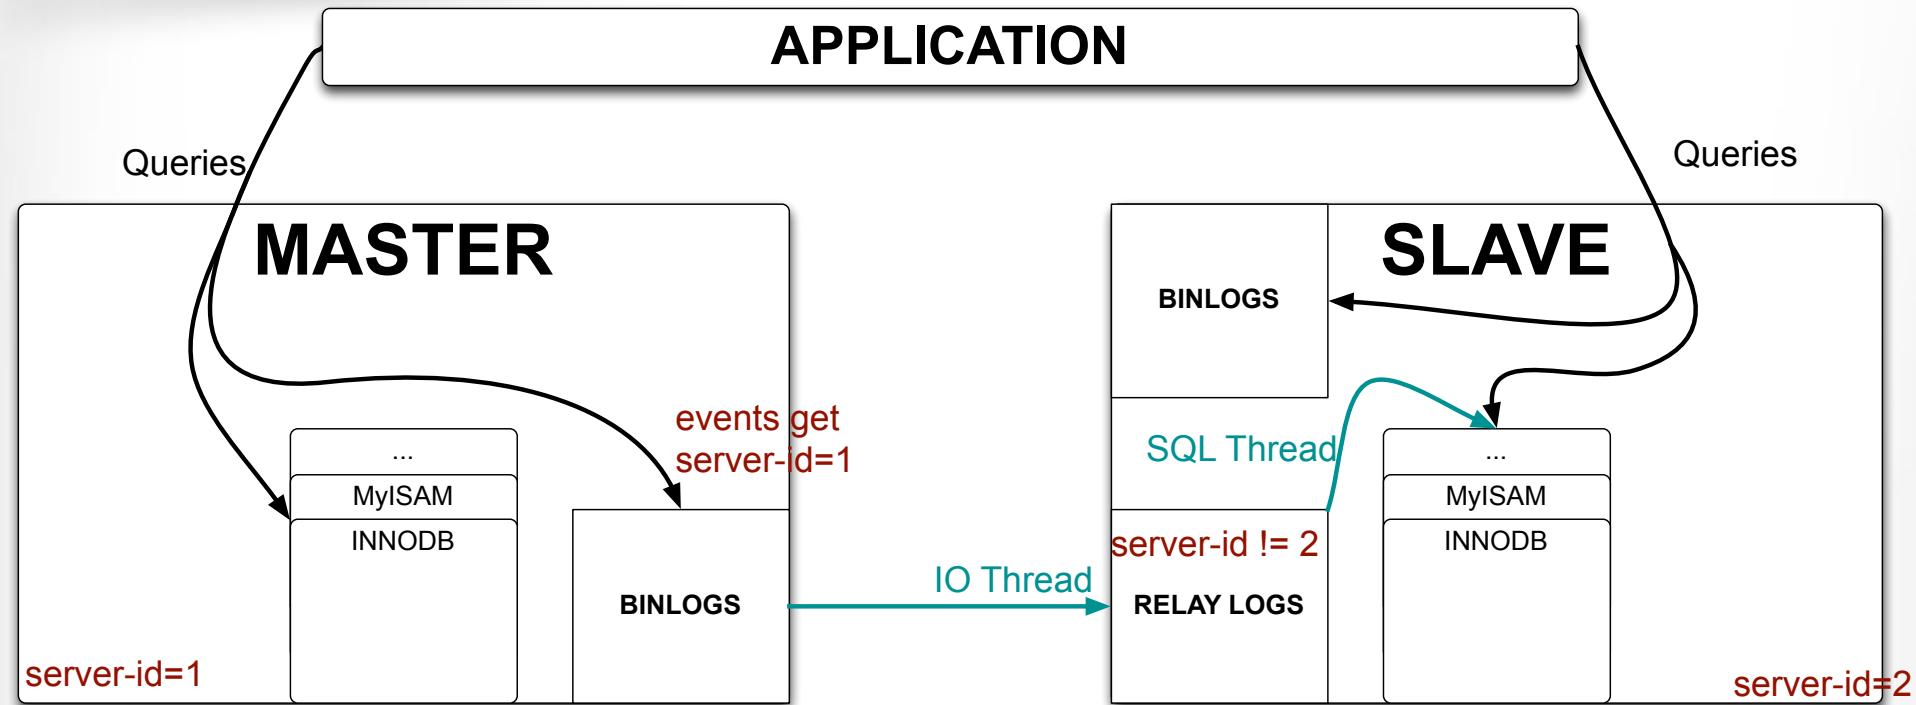
- On the master:
  - Enable Binary Logging: `log-bin=log-bin`
  - Set Server ID: `server-id=1` (already done)
- On the slave:
  - Set Server ID:

```
$ vi /etc/my.cnf
```

```
server-id=2
```

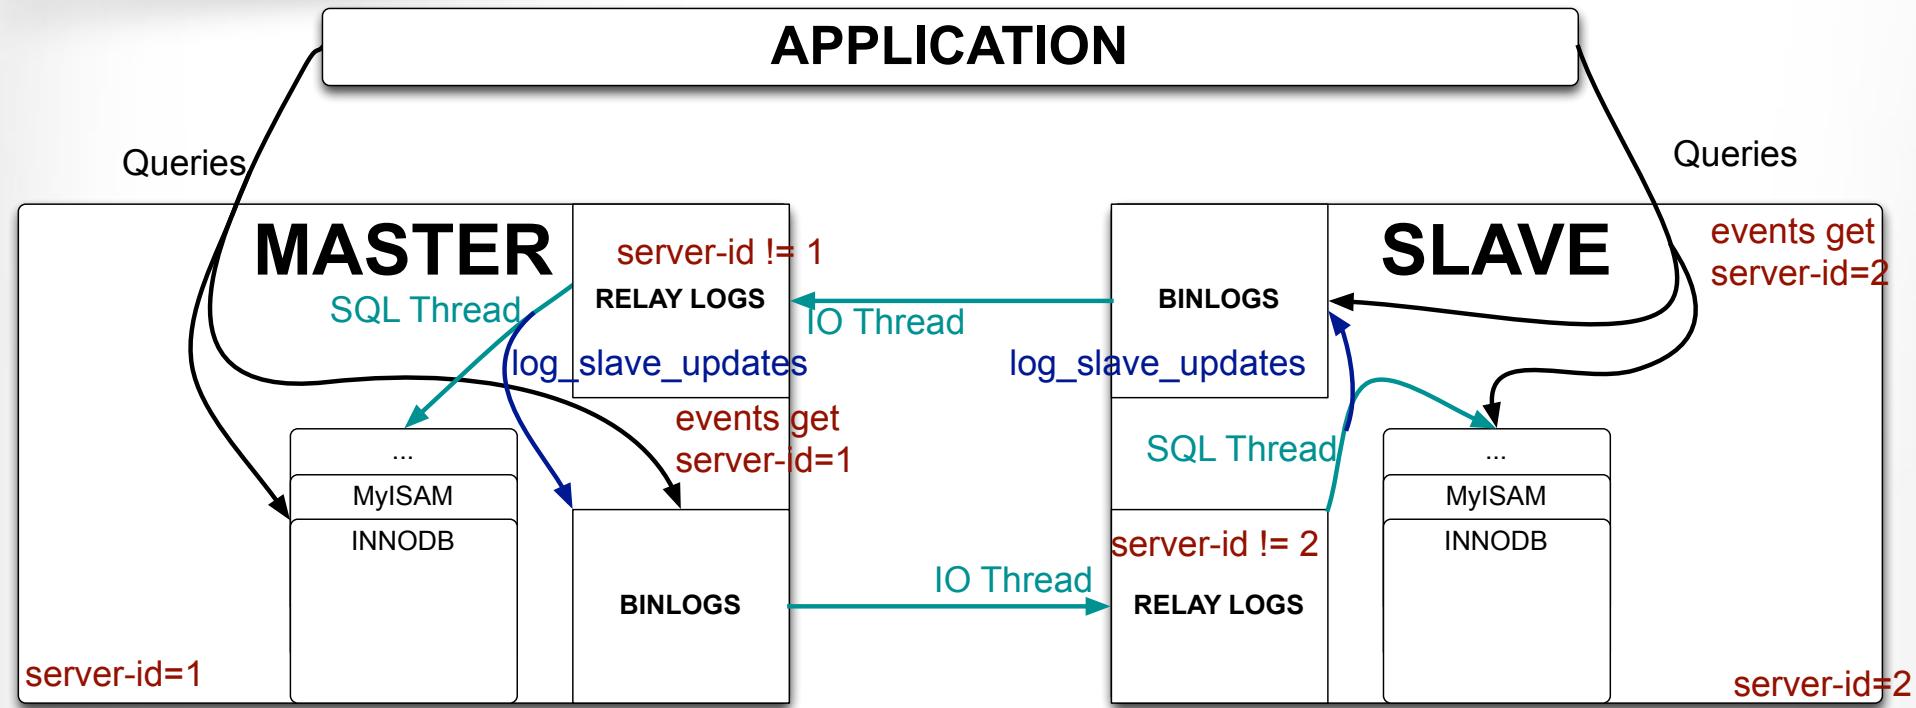
# Why server-id ?

95



# Why server-id ?

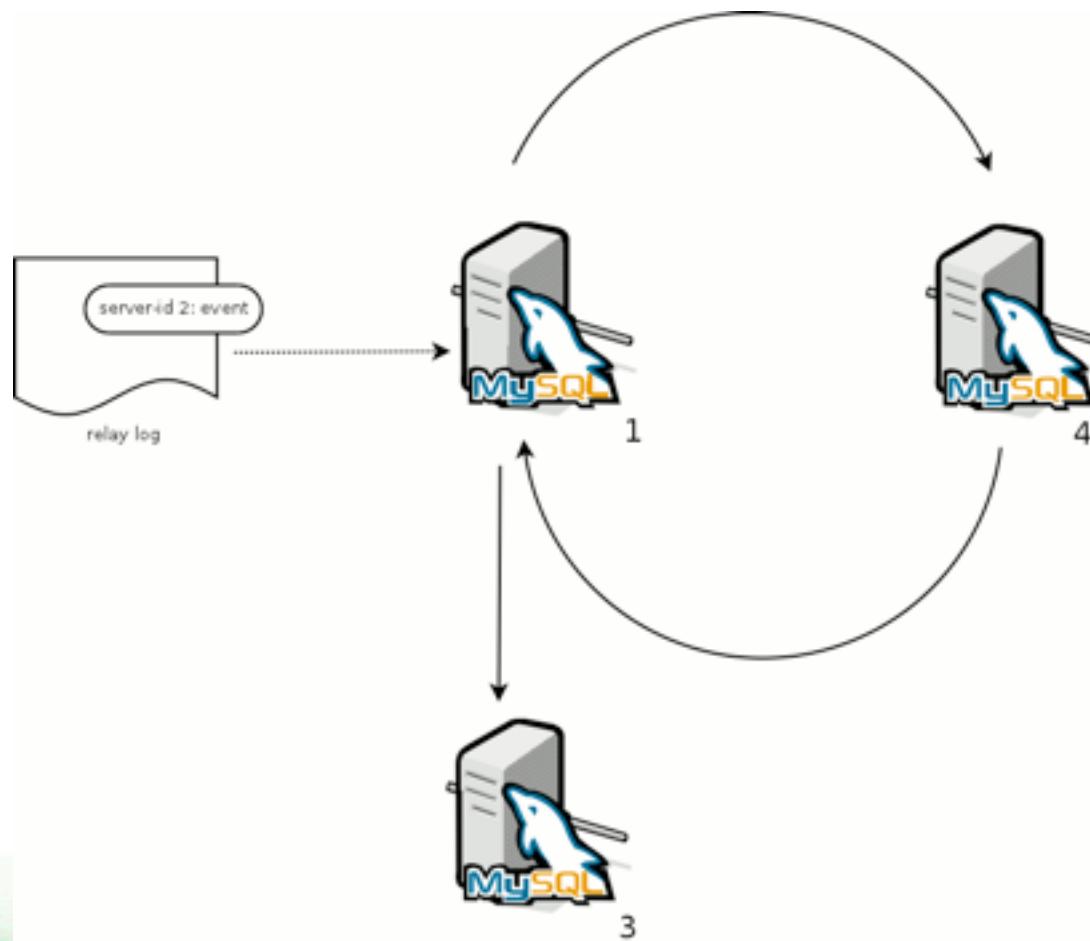
96



# Why server-id ?

97

- lefred didn't like my diagram, here is his version



# Round 1

# Round 1

98



# Round 1

98



VS

# Round 1

98



VS



# Round 1

98



VS



**Who Wins?**

# Why server-id ?

99

- Avoid events to be written more than once
- `replicate_same_server_id` does what it says

# Configure Replication - HandsOn!

100

- On Master, add permissions:

```
node1> GRANT REPLICATION SLAVE ON *.* TO  
'repl'@'%' IDENTIFIED BY 'slaveme';
```

- Get Binlog file/Position from backup:

```
# cat /var/lib/mysql/  
xtrabackup_binlog_info
```

- On Slave, configure replication:

```
node2> CHANGE MASTER TO  
MASTER_HOST='192.168.70.1',  
MASTER_USER='repl',  
MASTER_PASSWORD='slaveme',  
MASTER_LOG_FILE='__yourfile__',  
MASTER_LOG_POS='__yourpos__';
```

```
slave> START SLAVE;
slave> SHOW SLAVE STATUS\G
    Slave_IO_State: Waiting for master to send event
    Master_Host: master
    Master_User: repl
    Master_Port: 3306
    Connect_Retry: 60
    Master_Log_File: mysql-bin.000008
    Read_Master_Log_Pos: 254
    Relay_Log_File: relay-bin.000002
    Relay_Log_Pos: 399
    Relay_Master_Log_File: mysql-bin.000008
    Slave_IO_Running: Yes
    Slave_SQL_Running: Yes
    Replicate_Do_DB:
    Replicate_Ignore_DB:
    Replicate_Do_Table:
    Replicate_Ignore_Table:
    Replicate_Wild_Do_Table:
    Replicate_Wild_Ignore_Table:
        Last_Error:
        Skip_Counter: 0
    Exec_Master_Log_Pos: 254
    Relay_Log_Space: 566
    Until_Condition: None
    Until_Log_File:
    Until_Log_Pos: 0
    Master_SSL_Allowed: No
    Master_SSL_CA_File:
    Master_SSL_CA_Path:
        Master_SSL_Cert:
        Master_SSL_Cipher:
        Master_SSL_Key:
    Seconds_Behind_Master: 0
Master_SSL_Verify_Server_Cert: No
    Last_IO_Errno: 0
    Last_IO_Error:
    Last_SQL_Errno: 0
    Last_SQL_Error:
```

IO Thread State

IO Thread Read Up To 101

SQL Thread Read Up To

Threads Running?

How Far Is Slave Behind?

- SQL Commands
  - Administrative Commands
  - Diagnostics Commands
- Shell Commands
  - mysqlbinlog

# Administrative Commands - HandsOn!

103

- Rotate binary log: FLUSH BINARY LOGS
- Rotate relay log: FLUSH RELAY LOGS
- START/STOP SLAVE IO\_THREAD/  
SQL\_THREAD
- Remove binary logs:  

```
PURGE MASTER LOGS TO 'mysql-bin.000005';
PURGE MASTER LOGS
        BEFORE '2012-01-01 00:00:00';
```
- Remove all binary logs: RESET MASTER
- Remove slave configuration and files: RESET SLAVE

# Diagnostics Commands

104

- On Master
  - SHOW MASTER STATUS
  - SHOW PROCESSLIST
  - SHOW SLAVE HOSTS
  - SHOW BINLOG EVENTS
- On Slave
  - SHOW SLAVE STATUS
  - SHOW PROCESSLIST

# On Master:SHOW MASTER STATUS - HandsOn!

105

- Current binary log file and position:

```
master> SHOW MASTER STATUS\G
***** 1. row *****
    File: mysql-bin.000008
    Position: 254
    Binlog_Do_DB:
Binlog_Ignore_DB: 1 row in set (0.00 sec)
```

# On Master: SHOW PROCESSLIST - HandsOn!

106

- Find Connected Slaves using SHOW PROCESSLIST:

```
master> SHOW PROCESSLIST\G
...
***** 2. row *****
    Id: 4
    User: repl
    Host: localhost:43537
    db: NULL
Command: Binlog Dump
    Time: 1264
    State: Has sent all binlog to slave; waiting for binlog to be
updated
    Info: NULL
...
...
```

# On Slave: SHOW PROCESSLIST - HandsOn!

107

- Slave Thread Status:

```
SLAVE> SHOW PROCESSLIST\G
***** 1. row *****
    Id: 5
  User: system user
   Host:
      db: NULL
Command: Connect
     Time: 88611
   State: Waiting for master to send event
     Info: NULL
***** 2. row *****
    Id: 8
  User: system user
   Host:
      db: NULL
Command: Connect
     Time: 83
   State: Has read all relay log; waiting for the slave I/O thread to update
          it
     Info: NULL
```

# Other Common Configuration Options

108

- Filtering: `binlog-%`, `replicate-%`
- Don't start replication at start:  
`skip_slave_start`
- Put relay log events in it's own binary log:  
`log_slave_updates`
- Disallow writing on slaves: `read_only`
- Automatically remove binlogs:  
`expire_logs_days`
- Change binlog format: `binlog_format`

# Filtering on Master

109

- `binlog-do-db=...`, `binlog-ignore-db=...`
- Warning! Different behavior between SBR and RBR:
  - SBR: Log all statements executed if the ‘default database’ (`use database`) should be binlogged
  - RBR: Log all row changes of the databases that should be binlogged

# Filtering on Slave

110

- `replicate-do-db=` & `replicate-ignore-db=`
- Warning! Different behavior between SBR and RBR  
Similar to `binlog-do/ignore-db`
- `replicate-do/ignore-table`: Filter on table level
- `replicate-wild-do/ignore-table`: Use wildcards (%,\_ ) to filter on table level
- `replicate-rewrite-db=db1->db2`

# Breaking & Fixing Replication - HandsOn!

111

- Create table on Slave First, then Master:

```
CREATE TABLE imdb.plmce
(
    id int auto_increment primary key,
    name varchar(20)
);
```

- Check Slave Status:

```
slave> SHOW SLAVE STATUS\G
```

- Replication broke :(

# Breaking & Fixing Replication - HandsOn!

112

- Replication is Broken.
- Here we understand what went wrong. Let's just skip this record:

```
slave> SET GLOBAL  
SQL_SLAVE_SKIP_COUNTER=1;  
slave> START SLAVE  
SQL_THREAD;
```



Listening to someone  
saying that he skipped a  
record and now  
replication works fine

# Inconsistencies - HandsOn!

113

- Let's delete some data from the slave:

```
node2> DELETE FROM imdb.users LIMIT 10;
```

- On the Master:

```
# pt-table-checksum --tables imdb.users  
--no-check-binlog-forma
```

- After Checksumming,  
let's verify:

```
# pt-table-checksum \  
--replicate-check-only
```



Almost ran an UPDATE without  
WHERE clause

# Fixing the Inconsistencies - HandsOn!

114

- Verify what data is wrong:

```
pt-table-sync --print --replicate percona.checksums \
localhost
```

- Approve and Execute:

```
pt-table-sync --execute --replicate percona.checksums \
localhost
```

- Warning! Split Brain situations may need different fixes

# Operational DBA In A Nutshell

- Setting Up Today's Environment
- Installation
- Logging In
- MySQL Privileges
- Diagnostics
- Troubleshooting
- Monitoring
- Backups
- Replication
- Schema Changes
- Configuration Optimization



# Schema Changes - Techniques

116

- Downtime
  - Con: Downtime
- Disable functionality in Application
  - Con: Partly Downtime
- Through replication
  - Con: A lot of Master/Slave roles being switches, time consuming, troublesome
- Online Schema Changes
  - Pro: Online!

# Online Schema Changes - HandsOn!

117

- Open second terminal

```
node1> insert into title  
(title, series_years) values  
('plmce !','2013');
```

- pt-online-schema-change

```
# time pt-online-schema-change \  
--execute --alter \  
"ADD INDEX new_idx (series_years)" \  
D=imdb,t=title
```

- Users triggers/views

# Operational DBA In A Nutshell

- Setting Up Today's Environment
- Installation
- Logging In
- MySQL Privileges
- Diagnostics
- Troubleshooting
- Monitoring
- Backups
- Replication
- Schema Changes
- **Configuration Optimization**



# Optimization - Things Not To Optimize

119

- `innodb_io_capacity`: don't oversize
- `sort_buffer_size`, `join_buffer_size`,  
`read_buffer_size`,  
`read_rnd_buffer_size`: leave default globally
- `innodb_thread_concurrency`: Default is fine  
for most environments
- `query_cache_type`, `query_cache_size`:  
disable or do not oversize (+64MB)
- `innodb_flush_method=O_DIRECT` and ext3

# Optimization - Things To Optimize

120

- `innodb_log_file_size`: Increase InnoDB Transaction Logs
  - mysql> pager grep sequence  
mysql> show engine innodb status\G  
select sleep(60); show engine innodb status\G
- `innodb_stats_on_metadata=0`
- `innodb_buffer_pool_size`
- `innodb_io_capacity`
- change the shared memory allocator (jemalloc)
  - [mysqld\_safe]  
`malloc-lib=/usr/lib64/libjemalloc.so.1`

# Configuration Wizard - HandsOn!

121

- <http://tools.percona.com>

## Optimize your MySQL Server

Leverage our expertise to create a good starting configuration for a MySQL server.

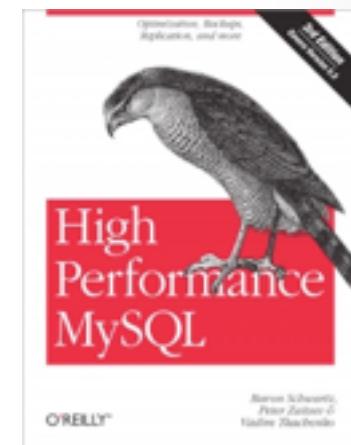
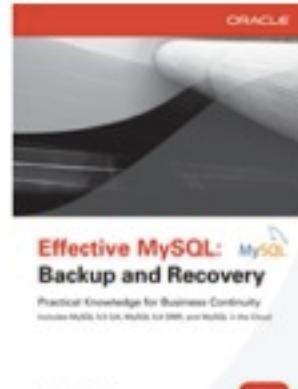
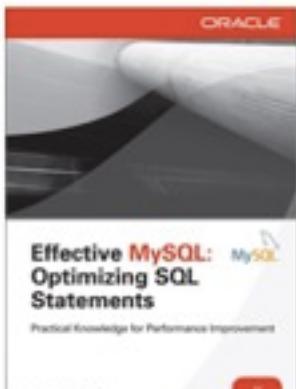
This tool will walk you through all of the steps necessary to create a ready to use MySQL configuration file in about 5 minutes.

**Generate A Configuration Now**

# Want to learn more?

122

- Some good resources to help you along:
  - Books
    - High Performance MySQL, 3rd Edition (Baron Schwartz, Peter Zaitsev, Vadim Tkachenko)
    - Effective MySQL series (Ronald Bradford)
    - Instant InnoDB (Matt Reid)



# Want to learn more?

123

- Some good resources to help you along:
  - Blogs
    - Planet MySQL: [planet.mysql.com](http://planet.mysql.com)
    - MySQL Performance Blog:  
[www.mysqlperformanceblog.com](http://www.mysqlperformanceblog.com)
  - Recorded Webinars
    - [www.percona.tv](http://www.percona.tv)

# Want to learn more?

124

- Some good entry-level talks at this conference
  - How to Design Indexes, Really - Bill Karwin
  - Backups in the Real World: Challenges and Solution - Ryan Huddleston
  - Choosing Hardware for MySQL - Kenny Gryp
  - Getting started with Tungsten Replicator - Neil Armitage
  - Load Balancing MySQL with HAProxy - Peter Boros
  - myq\_gadgets - Simple scripts to expose the inner workings of your MySQL server - Jay Janssen
  - You know SSD's make MySQL fast, but do you know why? - Jared Hulbert

# Operational DBA In A Nutshell - Q&A

125

- Setting Up Today's Environment
- Installation
- Logging In
- MySQL Privileges
- Diagnostics
- Troubleshooting
- Monitoring
- Backups
- Replication
- Schema Changes
- Configuration Optimization

How doing operations now feels

# Operational DBA In A Nutshell - Q&A

125

- Setting Up Today's Environment
- Installation
- Logging In
- MySQL Privileges
- Diagnostics
- Troubleshooting
- Monitoring
- Backups
- Replication
- Schema Changes
- Configuration Optimization

How doing operations now feels

