

Manage MySQL with Open source

周振兴@Taobao DBA
<http://orczhou.com>

2010-06





What MySQL do for us

System A



分享给好友

收藏该宝贝

- 950G data
- 650G index
- Total **1.6T**
- IOPS 18947
- Delta 3.2G
 - $8000 * 1024 = 800W$



What MySQL do for us

- System B
 - **15,0000,0000** records
 - Total: **4T**
 - Delta
 - **1000,0000**

16 *




Uptime:217 days 18 hours 34 min 2 sec




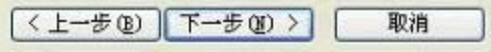


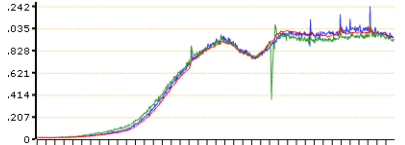
What about my topic

- How to manage 200+ MySQL Instance





What we do

-  Test
-  routine Install
-  Find Problem/debug
-  Backup
-  Monitor



Why

- Why we think about this a lot?
 - We have **200+ MySQL** Instance

<i>Test</i>	is about	M oney
<i>Install</i>	is about	T ime
<i>Debug</i>	is about	S tability
<i>Backup</i>	is about	S afety
<i>Monitor</i>	is about	S tatus

Which one you can drop?



Why Test

- Simple math
 - Performance increase 10%
 - $200 * 2W * 10\% = 40W$
 - $200 * 2W * 20\% = 80W$



What can 80W do ?



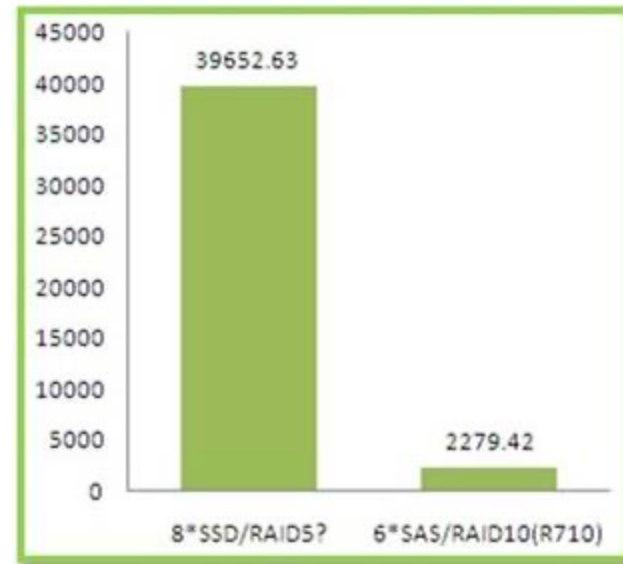
Why Test

- Performance
- Stable
 - Last year >85% accident can be avoided
 - This year < 5%



How Test

- Sysbench
 - How to use sysbench
 - How many row
 - Read write rate





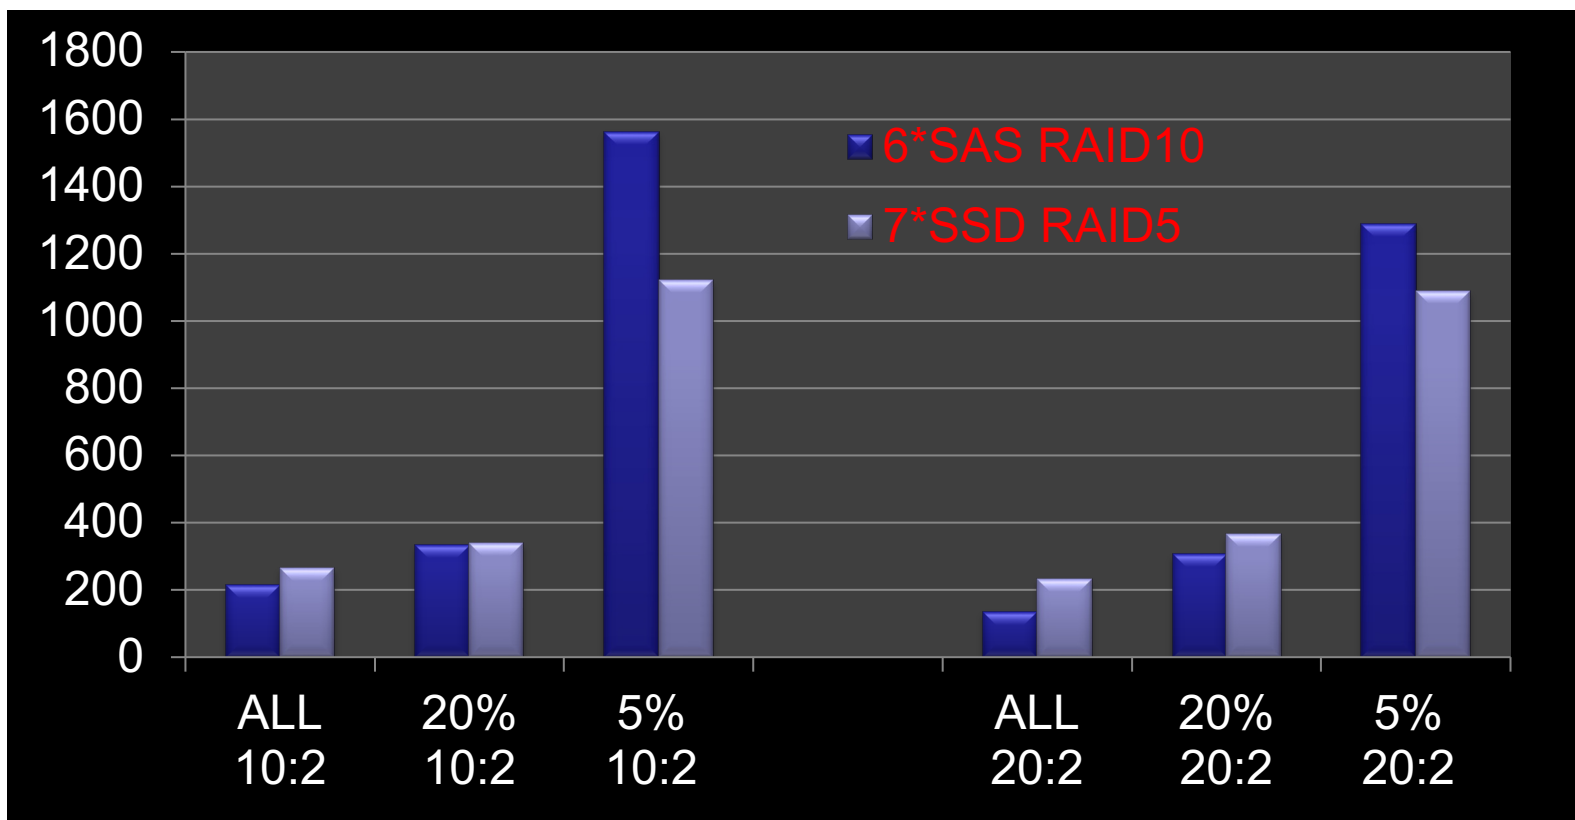
How Test

- Why sysbench?
- Why not sysbench?
- Why supersmack?
 - Data set > 120G
 - 95% query for 5% data
 - 80% query for 20% data



How Test

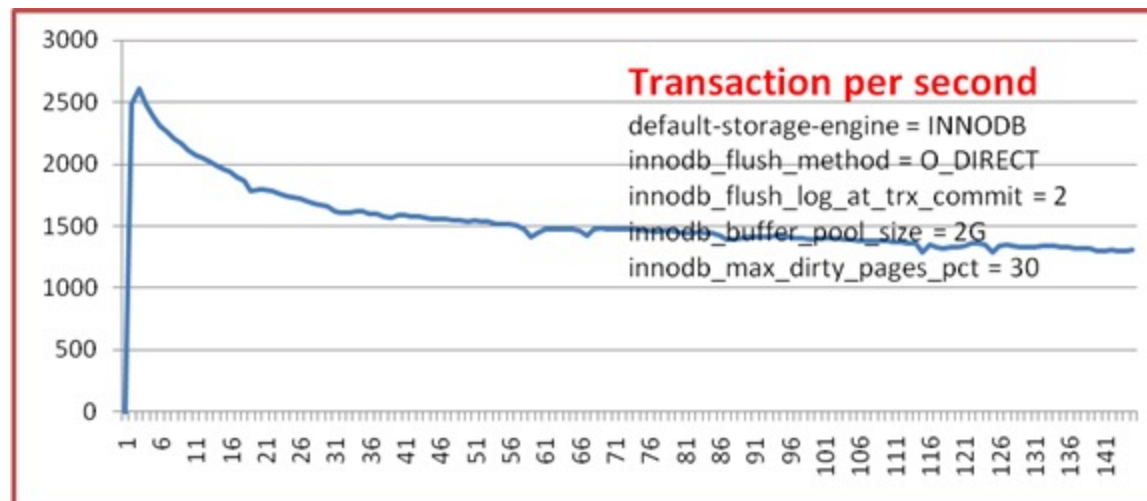
- Supersmack





How Test

- Supersmack
 - Dell PowerEdge 2950
 - Mem:24G 6*300 SAS RAID10
 - We test different PC Server/Hardware






What Test

- Decide which value we choose

InnoDB_IO_CAPACITY



One hour average in blue
Two hour average in red



Install

<i>Test</i>	Money	Sysbench/supersmack
<i>Install</i>	Time	
<i>Debug</i>	Stability	
<i>Backup</i>	Safety	

```
>mysql_install_auto -h 172.23.119.257 -v 5.1.45
```

File descriptor
Swappiness

.....



Debug

- It's a real complicated work
- One point: `tcpdump + mk-query-digest`
 - It's a talent work



Debug

- tcpdump + mk-query-digest

```
# Query 1: 355.85 QPS, 0.80x concurrency, ID 0x4C4E1DAC3F814B78 at byte 3539233984
# This item is included in the report because it matches --limit.
#          pct  total      min      max      avg      95%  stddev  median
# Count          47  440277
# Exec time      44   990s    71us    40s     2ms     8ms    63ms   185us
# Hosts          6 172.16.0.28.184 (91448), 172.16.0.23.62 (91139)... 4 more
# Time range 2010-06-08 09:11:34.646121 to 2010-06-08 09:32:11.898820
# bytes        43 126.34M    291    322  300.89  313.99    9.12  284.79
# Errors        2 none (440259), #1264 (18)
# Rows affe    71 429.35k     0      1    1.00    0.99    0.04    0.99
# Warning c     0      0      0      0      0      0      0      0
# Query_time distribution
# 1us
# 10us #####
# 100us #####
# 1ms #####
# 10ms ###
# 100ms #
# 1s #
# 10s+ #
# Tables
# SHOW TABLE STATUS FROM `beidou` LIKE 'king_vip_account_record_02_0123'\G
# SHOW CREATE TABLE `beidou`.`king_vip_account_record_02_0123`\G
insert into king_vip_account_record_02_0123 ( user_num, cnt_rule_id
score_date, state, gmt_create ) values ( 185544315, 21,
06-08 17:24:31.0', 1, now() ) \G
```




Backup

- mysqldump is cool
 - 39G .sql takes 80 min
 - Restore takes 206 min(3.4)
- Still we need Xtrabackup
 - On line
 - No lock
 - Some environment
 - it's quick && restore is quick



Xtrabackup

- Using InnoDB *Redo* (Transaction log)
- No effect to the *Online* system
- Also help u to setup a new slave
- System can be more *Available*




More backup strategy

- Last 4 Day Last 3 month Last 3 Year
- 1 3 5 mysqldump 2 4 6 xtrabackup



MySQL Monitor

- Quick way: Cacit + RRDTool
 - Screenshot





What we monitor

- Com_select
 - Com_insert
 - Com_update
 - Com_delete
- InnoDB
 - Logical read/write
 - Physical read/write
 - Row operation
- Principle:

The bigger the system is,
the more simple the basic layer is



Finally

<i>Test</i>	Money	Sysbench/supersmack
<i>Install</i>	Time	Linux/bash
<i>Debug</i>	Stability	mk-query-digest / iostat / ...
<i>Backup</i>	Safety	Xtrabackup
<i>Monitor</i>	Status	Cacti / RRDTool

- Enjoy....