



PERCONA
Performance Consulting Experts

Percona Server: secret features

Date, time, place:

MySQL Conference & Expo 2011
12-Apr-2011

Reporter:

Vadim Tkachenko
Co-founder, CTO,
Percona Inc

- In this talk: some not high profile, but addictive features. Once you start to use them you can't give up.

Percona Server

- Percona Server is an enhanced drop-in replacement for MySQL
 - Performance: better and stable
 - Diagnostics
 - Operations

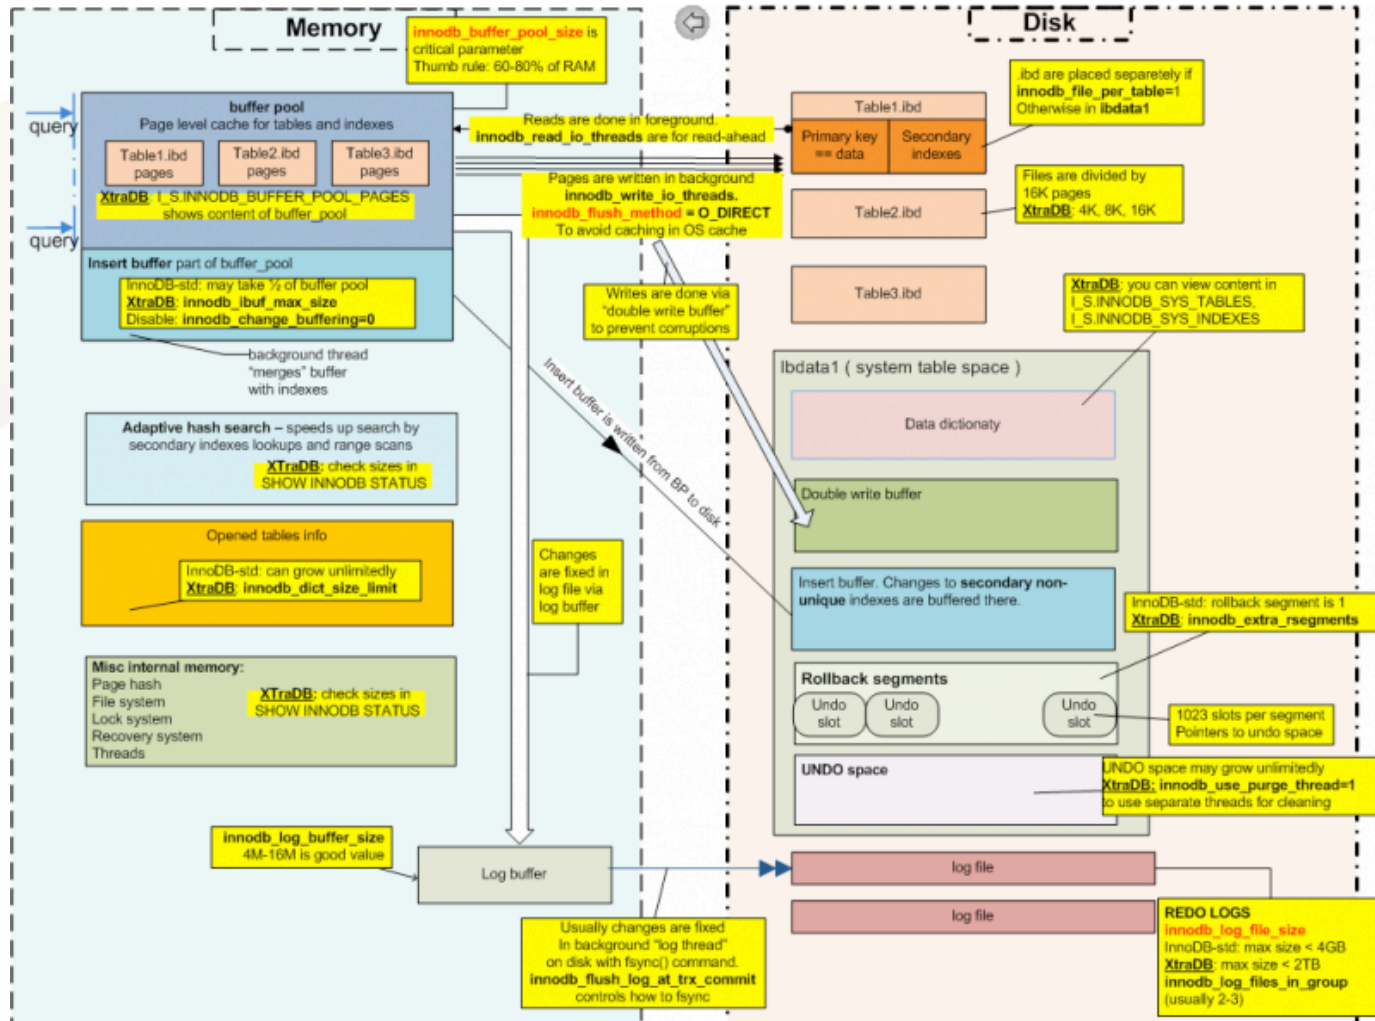


PERCONA
SERVER

What's inside

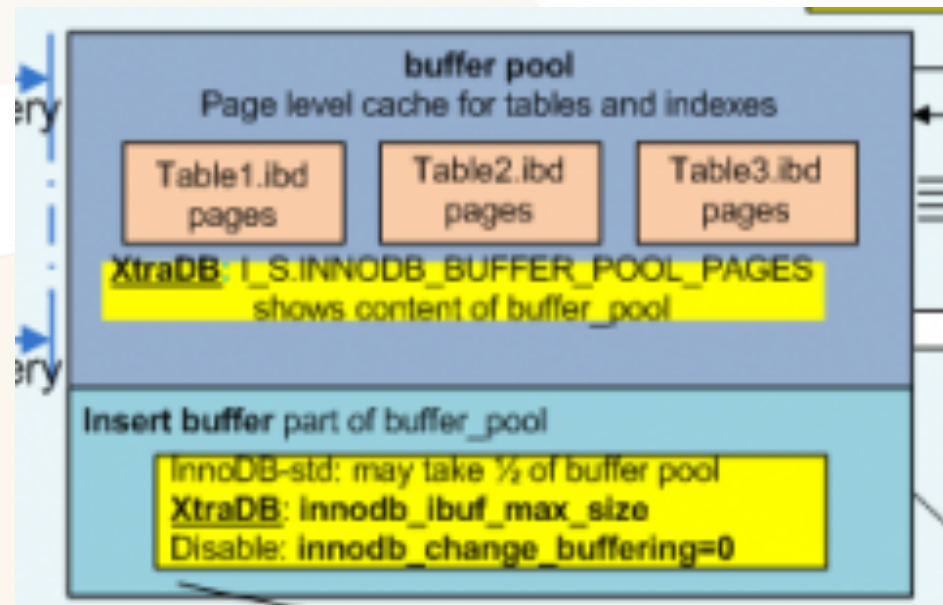
- **Diagnostic / monitoring**
- Operations / SaaS / Shards
- Still performance

Monitoring / Diagnostic



Buffer pool

- INFORMATION_SCHEMA.
innodb_buffer_pool_pages_index



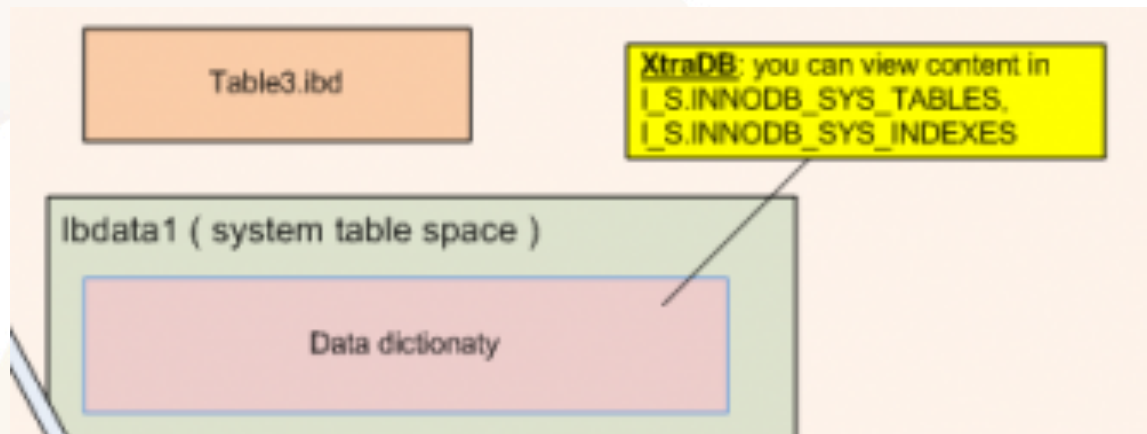
Buffer pool

```
mysql> SELECT index_id, COUNT(*) cnt, SUM(dirty = 1) dirty FROM  
INFORMATION_SCHEMA.innodb_buffer_pool_pages_index GROUP BY index_id;
```

index_id	cnt	dirty
33	1144941	0
36	24314	0
37	72327	0
38	50281	0
39	129414	0
40	247860	0
43	4494461	0
44	665262	0
46	2083240	0
47	100130	0
48	110772	0
49	56997	0

InnoDB data-dictionary

- innodb_sys_tables, innodb_sys_indexes



InnoDB data-dictionary

```
select * from innodb_sys_tables;
```

TABLE_ID	SCHEMA	NAME	FLAG	N_COLS	SPACE
11		SYS_FOREIGN	0	7	0
12		SYS_FOREIGN_COLS	0	7	0
31	tpcc1000	customer	1	24	28
30	tpcc1000	district	1	14	27
33	tpcc1000	history	1	11	30
28	tpcc1000	item	1	8	25
34	tpcc1000	new_orders	1	6	31
37	tpcc1000	order_line	1	13	34
35	tpcc1000	orders	1	11	32
39	tpcc1000	stock	1	20	36
22	tpcc1000	warehouse	1	12	19

InnoDB data-dictionary

```
select * from innodb_sys_indexes;
```

INDEX_ID	NAME	TABLE_ID	TYPE	N_FIELDS	PAGE_NO	SPACE
24	PRIMARY	22	3	1	3	19
30	PRIMARY	28	3	1	3	25
32	PRIMARY	30	3	2	3	27
33	PRIMARY	31	3	3	3	28
48	idx_customer	31	0	4	4	28
36	GEN_CLUST_INDEX	33	1	0	3	30
37	fkey_history_1	33	0	3	4	30
38	fkey_history_2	33	0	2	5	30
39	PRIMARY	34	3	3	3	31
40	PRIMARY	35	3	3	3	32
49	idx_orders	35	0	4	5	32
43	PRIMARY	37	3	4	3	34
44	fkey_order_line_2	37	0	2	4	34
46	PRIMARY	39	3	2	3	36
47	fkey_stock_2	39	0	1	4	36

Buffer pool

```
SELECT st.SCHEMA,st.NAME,si.NAME, si.index id, COUNT(*) cnt, SUM(dirty = 1) dirty FROM
innodb_buffer_pool_pages_index bp JOIN innodb_sys_indexes si ON (bp.index_id = si.index_id) JOIN
innodb_sys_tables st ON (st.TABLE_ID = si.TABLE_ID) GROUP BY index_id;
```

SCHEMA	NAME	NAME	index_id	cnt	dirty
tpcc1000	warehouse	PRIMARY	24	8	0
tpcc1000	item	PRIMARY	30	592	0
tpcc1000	district	PRIMARY	32	70	0
tpcc1000	customer	PRIMARY	33	1144941	0
tpcc1000	history	GEN_CLUST_INDEX	36	24314	0
tpcc1000	history	fkey_history_1	37	72327	0
tpcc1000	history	fkey_history_2	38	50281	0
tpcc1000	new_orders	PRIMARY	39	129414	0
tpcc1000	orders	PRIMARY	40	247860	0
tpcc1000	order_line	PRIMARY	43	4494461	0
tpcc1000	order_line	fkey_order_line_2	44	665262	0
tpcc1000	stock	PRIMARY	46	2083240	0
tpcc1000	stock	fkey_stock_2	47	100130	0
tpcc1000	customer	idx_customer	48	110772	0
tpcc1000	orders	idx_orders	49	56997	0

Index stats / optimizer

```
mysql> select * from INNODB_INDEX_STATS;
```

table_schema	table_name	index_name	fields	rows_per_key	index_total_pages	index_leaf_pages
tpcc1000	district	PRIMARY	2	10, 1	97	69
tpcc1000	orders	PRIMARY	3	3872368, 2661, 1	119872	104684
tpcc1000	orders	idx_orders	4	3872368, 5076, 1, 0	62904	54848
tpcc1000	new_orders	PRIMARY	3	3835, 1920, 1	23538	20411
tpcc1000	stock	PRIMARY	2	11121643, 1	2087680	2085308
tpcc1000	stock	fkey_stock_2	2	1000, 0	114496	100000
tpcc1000	order_line	PRIMARY	4	34461209, 34461209, 10, 1	1969408	1720147
tpcc1000	order_line	fkey_order_line_2	6	34461209, 3, 3, 1, 0, 0	732672	639402
tpcc1000	warehouse	PRIMARY	1	1	8	7
tpcc1000	history	GEN_CLUST_INDEX	1	1	156992	149167

Index stats / usage

```
mysql> select * from INDEX_STATISTICS;
```

TABLE_SCHEMA	TABLE_NAME	INDEX_NAME	ROWS_READ
tpcc1000	order_line	PRIMARY	6682478
tpcc1000	customer	idx_customer	805912
tpcc1000	district	PRIMARY	666834
tpcc1000	orders	idx_orders	16265
tpcc1000	item	PRIMARY	1625832
tpcc1000	customer	PRIMARY	738535
tpcc1000	stock	PRIMARY	6481682
tpcc1000	warehouse	PRIMARY	487929
tpcc1000	orders	PRIMARY	341484
tpcc1000	new_orders	PRIMARY	325120

Usage statistics

- **USER_STATISTICS/CLIENT_STATISTICS/
TABLE_STATISTIC**

Table stats / usage

```
mysql> select * from TABLE_STATISTICS;
```

TABLE_SCHEMA	TABLE_NAME	ROWS_READ	ROWS_CHANGED	ROWS_CHANGED_X_INDEXES
tpcc1000	order_line	15170566	7380949	14761898
tpcc1000	customer	3503133	738082	1476164
tpcc1000	history	0	369059	738118
tpcc1000	warehouse	1107185	369063	369063
tpcc1000	item	3689250	0	0
tpcc1000	district	1513149	738122	738122
tpcc1000	stock	14712182	3689244	7378488
tpcc1000	new_orders	737830	737862	737862
tpcc1000	orders	811868	738088	1476176

Processlist

```
SELECT * FROM PROCESSLIST\G
***** 1. row *****
      ID: 49176
     USER: root
     HOST:
      DB: tpcc1000
  COMMAND: Query
      TIME: 25
   STATE: Sending data
     INFO: select count(*) from order_line
  TIME_MS: 25226
  ROWS_SENT: 0
ROWS_EXAMINED: 0
  ROWS_READ: 2776173
```

Temporary tables

```
select * from information_schema.global_temporary_tables;
```

SESSION_ID	DATABASE	TABLE_NAME	ENGINE	NAME
6	test	t3	MyISAM	#sql29da_6_0
5	test	t2	MEMORY	#sql29da_5_3
5	test	t1	MyISAM	#sql29da_5_2

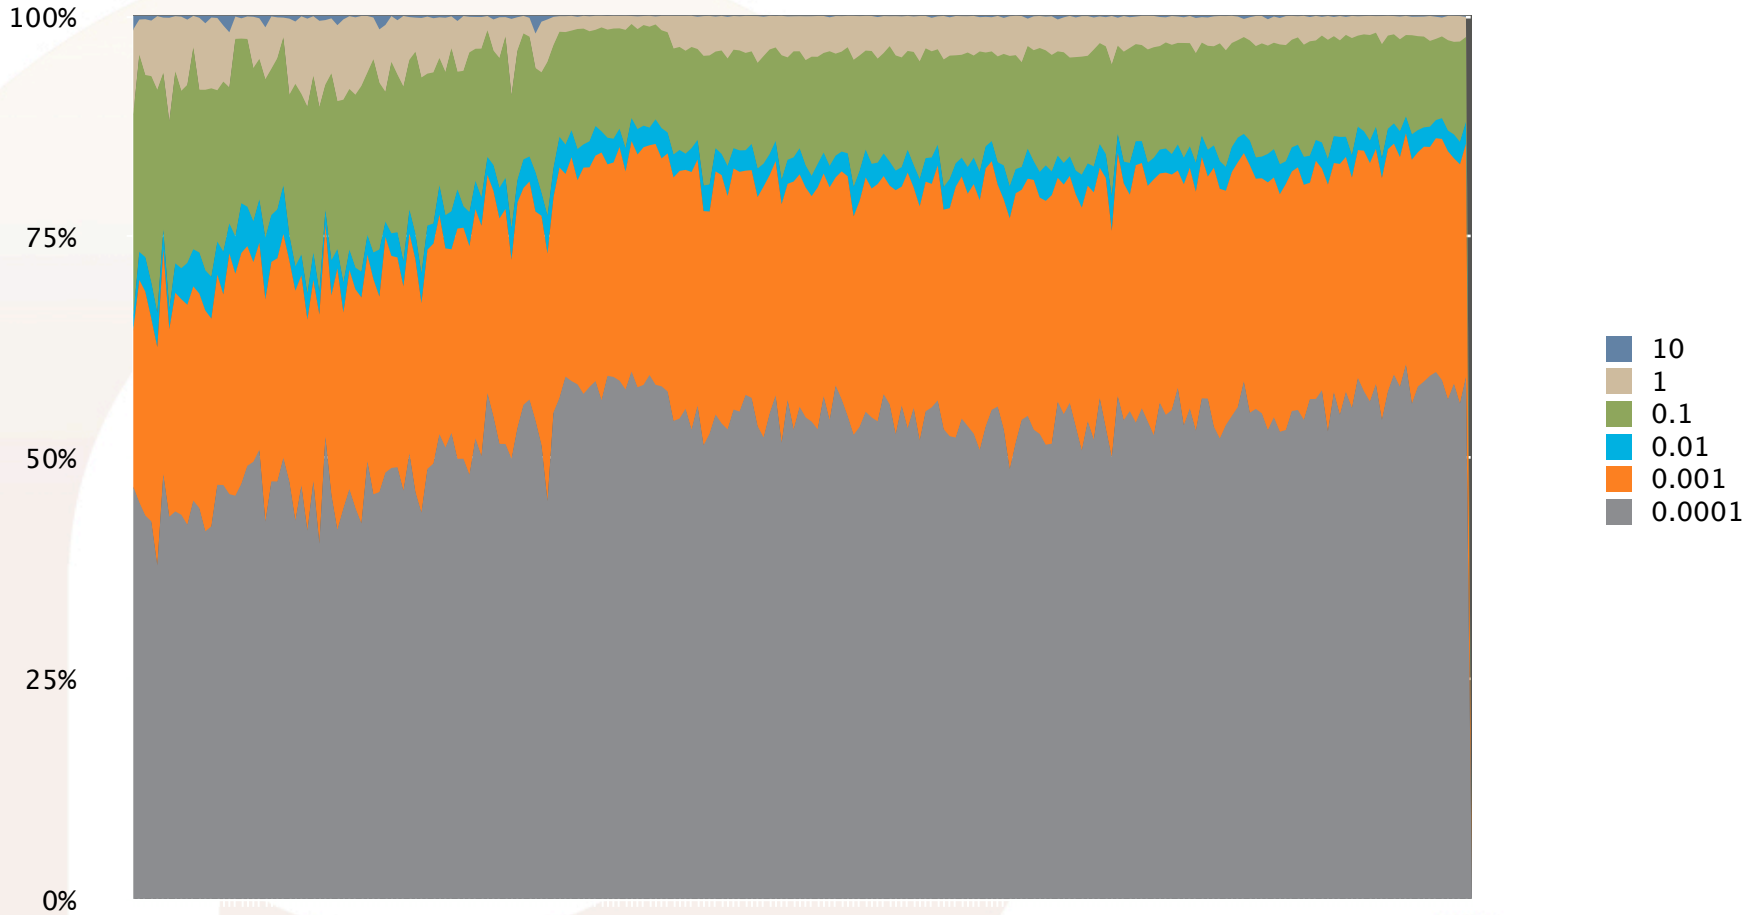
Credits to Venu Anuganti for contribution

Response time

```
mysql> select * from QUERY_RESPONSE_TIME;
```

time	count	total
0.000001	8	0.000000
0.000010	54	0.000060
0.000100	166264	10.416637
0.001000	13256	2.398679
0.010000	2479	12.299993
0.100000	2853	108.834296
1.000000	202	28.999465
10.000000	0	0.000000
100.000000	0	0.000000
1000.000000	0	0.000000
10000.000000	0	0.000000
100000.000000	0	0.000000
1000000.000000	0	0.000000
TOO LONG	0	TOO LONG

Response time



12-Apr-2011

Percona Server: secret features

What's inside

- Diagnostic / monitoring
- **Operations / SaaS / Shards**
- Still performance

Crash tolerant

- `innodb_pass_corrupt_table (5.1) / innodb_corrupt_table_action=assert|warn (5.5)`
- `SELECT data FROM corrupted_table (work in progress)`

Resource control

- **innodb dict size limit**

Export/import tables

- Move one InnoDB tables from one server to another
- In combination with XtraBackup
 - More in my tomorrow's talk

InnoDB statistics

- Stop “divings”
- **innodb use sys stats table**
- **innodb stats auto update**

Crash-safe slaves

- `innodb_overwrite_relay_log_info`

- + InnoDB: In a MySQL replication slave the last master binlog file
 - + InnoDB: position 0 468, file name gauntlet3-bin.000015

What's inside

- Diagnostic / monitoring
- Operations / SaaS / Shards
- **Still performance**

Fast warmup

- `innodb_buffer_pool_restore_at_startup` (5.5) /
`innodb_auto_lru_dump` (5.1)
 - Now in Drizzle also

Transactional logs

- `innodb_use_global_flush_log_at_trx_commit (5.5)`
- `SET innodb_flush_log_at_trx_commit`
- `innodb_log_block_size = 512 | 4096`
- `Innodb_log_file_size > 4GB`

SLOW LOG

- Extended stats

```
# User@Host: mailboxer[mailboxer] @ [192.168.10.165]
# Thread_id: 11167745 Schema: board
# QC_Hit: No Full_scan: No Full_join: No Tmp_table: Yes Disk_tmp_table: No
# Filesort: Yes Disk_filesort: No Merge_passes: 0
# Query_time: 0.000659 Lock_time: 0.000070 Rows_sent: 0 Rows_examined: 30
Rows_affected: 0 Rows_read: 30
# InnoDB_IO_r_ops: 1 InnoDB_IO_r_bytes: 16384 InnoDB_IO_r_wait: 0.028487
# InnoDB_rec_lock_wait: 0.000000 InnoDB_queue_wait: 0.000000
# InnoDB_pages_distinct: 5
select count(distinct author_id) from art87.article87 force index (forum_id) where
forum_id = 240215 and thread_id = '710575'
```

SLOW LOG + SHOW PROFILE

```
# Query_time: 50.603038  Lock_time: 0.000142  Rows_sent: 19  Rows_examined: 9119006
Rows_affected: 0  Rows_read: 19
# Bytes_sent: 1163  Tmp_tables: 4  Tmp_disk_tables: 1  Tmp_table_sizes: 1046064
# Profile_starting: 0.000077  Profile_starting_cpu: 0.000076
Profile_checking_permissions: 0.000003  Profile_checking_permissions_cpu: 0.000003
Profile_checking_permissions: 0.000001  Profile_checking_permissions_cpu: 0.000002
Profile_checking_permissions: 0.000005  Profile_checking_permissions_cpu: 0.000005
Profile_Opening_tables: 0.000056  Profile_Opening_tables_cpu: 0.000056
Profile_System_lock: 0.000005  Profile_System_lock_cpu: 0.000004
Profile_init: 0.000024  Profile_init_cpu: 0.000024
Profile_optimizing: 0.000011  Profile_optimizing_cpu: 0.000011
Profile_statistics: 0.000011  Profile_statistics_cpu: 0.000011
Profile_preparing: 0.000030  Profile_preparing_cpu: 0.000030
Profile_Creating_tmp_table: 0.000025  Profile_Creating_tmp_table_cpu: 0.000025
Profile_executing: 0.315203  Profile_executing_cpu: 0.314331
Profile_converting_HEAP_to_MyISAM: 0.170516  Profile_converting_HEAP_to_MyISAM_cpu: 0.161116
Profile_executing: 10.745184  Profile_executing_cpu: 10.714887
Profile_Copying_to_tmp_table: 39.194304  Profile_Copying_to_tmp_table_cpu: 39.085046
# Profile_total: 50.603045  Profile_total_cpu: 50.452554
use information_schema;
SET timestamp=1302580533;
SELECT st.SCHEMA,st.NAME,si.NAME, si.index_id, COUNT(*) cnt, SUM(dirty = 1) dirty
FROM innodb_buffer_pool_pages_index bp JOIN innodb_sys_indexes si ON (bp.index_id = si.index_id)
JOIN innodb_sys_tables st ON (st.TABLE_ID = si.TABLE_ID) GROUP BY index_id;
```

The end

- Slides will be online. <http://www.percona.com/about-us/presentations/>
- vadim@percona.com
- Your questions ?
- We are hiring!