



Troubleshooting MySQL Performance

Sveta Smirnova
Principal Support Engineer
November, 21, 2015

About Percona

- Performance Consulting Experts
- Free software: extended versions of MySQL and MongoDB (Percona Server), technical support, consulting, remote DBA (managed services)
- Since 2006, headquarters are in USA
- Distributed company with more than 100 employees in 30 countries
- 3000+ customers, including Cisco and Alcatel Lucent

Percona customers

- Data retrieved from free sources



1С-БИТРИКС

Selectel

Acronis



интернет
ХОСТИНГ
центр



ALAWAR
ENTERTAINMENT



SMARTRESPONDER

Сервис умного Email-маркетинга



MIRANTIS

CUSTIS®

заказные информсистемы

Table of Contents

- MySQL Server overview
- Single statement performance
- Internal concurrency issues
- Environment

MySQL Server overview

MySQL architecture

Connectors: C, JDBC, ODBC, Python, ...

Connection Pool: Authentication, Caches

SQL interface

Parser

Optimizer

Caches and Buffers:
Global
Engine-specific

Storage engines: InnoDB, TokuDB, ...

File system: Data, Index, logs, other files

- Base
 - Layout
 - Log files
- Connectors
- Optimizer
- Cache and buffers
- Storage engines
- Management

Important server parts

- Optimizer
- Storage engine
 - Study tools which offers your engine
- Variables

Troubleshooting instruments

- Log files
- Information Schema
 - SHOW commands
- Performance Schema
- Storage engine instruments

Troubleshooting instruments

- Log files
 - Error
 - Slow query
 - General query
 - Audit plugins
 - Your application
 - Operating system

Troubleshooting instruments

- Information Schema
 - Contains metadata information
 - Tables
 - Indexes
 - Other
 - Allows to create plugins
 - InnoDB plugins

Troubleshooting instruments

- Performance Schema
 - Monitors interval operations
 - Statements
 - Stages
 - Locks
 - Memory
 - Variables
 - Replication
 - IO
 - Mutexes and waits

Troubleshooting instruments

- Storage engine instruments
 - SHOW ENGINE ... STATUS
 - engine_* variables
 - Engine-specific tools

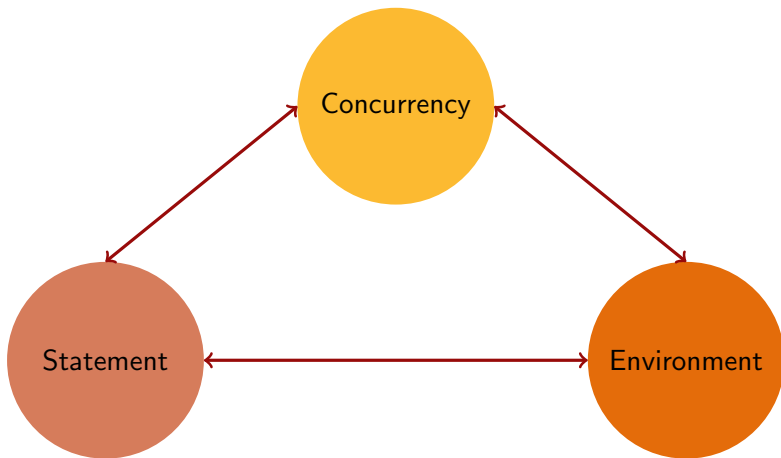
Monitoring solutions

- Command-line
 - Percona Toolkit
 - MySQL Utilities
- With GUI
 - MySQL Enterprise Monitor (MEM)
 - MEM plugin for Oracle Enterprise Manager
 - VividCortex
 - **Many others**

What can affect performance?

- Query processing
 - Optimization
 - Execution
- Concurrency
- Environment
 - Parallel processes
 - Hardware

Troubleshooting workflow



Single statement performance


Where to find slow statement?

- While designing application
- In slow query log
- In Performance Schema

```
mysql> select sql_text, (timer_end - timer_start)/1000000000 as time \  
-> from events_statements_history_long \  
-> where (timer_end - timer_start) > 4*10000000000 \  
-> and event_name='statement/sql/select' and current_schema='employees'\G  
***** 1. row *****  
sql_text: select * from titles  
time: 1929.8997  
***** 2. row *****  
sql_text: select count(*) from titles join employees using (emp_no)  
where title='Senior Engineer'  
time: 4493.3863  
2 rows in set (0.01 sec)
```

MySQL Indexes

- B-Tree (Mostly)
- Fractal Tree
- R-Tree (Spatial)
- Hash (Memory SE)
- Engine-dependent



More in Performance Schema

- `events_statements_*` and `prepared_statements_instances` tables
 - Important field names
 - `CREATED_TMP_DISK_TABLES`
 - `CREATED_TMP_TABLES`
 - `SELECT_FULL_JOIN`
 - `SELECT_RANGE_CHECK`
 - `SELECT_SCAN`
 - `SORT_MERGE_PASSES`
 - `SORT_SCAN`
- Views in `sys` schema

More in Performance Schema

- `events_statements_*` and `prepared_statements_instances` tables
- Views in `sys` schema
 - Important view names
 - `statement_analysis`
 - `statements_with_full_table_scans`
 - `statements_with_runtimes_in_95th_percentile`
 - `statements_with_sorting`
 - `statements_with_temp_tables`
 - `statements_with_errors_or_warnings`

Stages shortcuts

- Everything, related to temporary tables
 - `EVENT_NAME LIKE 'stage/sql/%tmp%'`
- Everything, related to locks
 - `EVENT_NAME LIKE 'stage/sql/%lock%'`
- Everything in state "Waiting for"
 - `EVENT_NAME LIKE 'stage/%/Waiting for%'`
- Frequently met issues (in Support)

Stages shortcuts

- Everything, related to temporary tables
- Everything, related to locks
- Everything in state "Waiting for"
- Frequently met issues (in Support)
 - EVENT_NAME='stage/sql/end'
 - EVENT_NAME='stage/sql/freeing items'
 - EVENT_NAME='stage/sql/Sending data'
 - EVENT_NAME='stage/sql/cleaning up'
 - EVENT_NAME='stage/sql/closing tables'

How to fix? Optimize!

- EXPLAIN EXTENDED
 - Should be followed by SHOW WARNINGS
- EXPLAIN PARTITIONS
- EXPLAIN FORMAT=JSON
- Status variables 'Handler_%'
- INFORMATION_SCHEMA.TRACE

EXPLAIN in MySQL

```
mysql> EXPLAIN EXTENDED SELECT user, host FROM mysql.user\G
***** 1. row *****
      id: 1
  select_type: SIMPLE
        table: user
         type: index
possible_keys: NULL
          key: PRIMARY
    key_len: 228
         ref: NULL
         rows: 4
  filtered: 100.00
    Extra: Using index
1 row in set, 1 warning (0.00 sec)
mysql> SHOW WARNINGS\G
***** 1. row *****
  Level: Note
   Code: 1003
Message: /* select#1 */ select 'mysql'.'user'.'User' AS 'user',
'mysql'.'user'.'Host' AS 'host' from 'mysql'.'user'
```


EXPLAIN in details

```
mysql> explain extended select * from t1 join t2 where...
```

id	select_type	table	type	***
1	SIMPLE	t1	ref	***
1	SIMPLE	t2	index	***

```
2 rows in set, 1 warning (0.00 sec)
```

- system
- const
- eq_ref
- ref
- fulltext
- ref_or_null
- index_merge
- unique_subquery
- index_subquery
- range
- index
- ALL

SIMPLE; PRIMARY; UNION; DEPENDENT UNION; UNION RESULT;
SUBQUERY; DEPENDENT SUBQUERY; DERIVED; MATERIALIZED

EXPLAIN in details: keys

Keys, which can be used for resolving the query

Actual length of the key (Important for multiple-column keys)

```
mysql> explain extended select * from t1 join t2 where t1.int_key=1;
```

```
****+-----+-----+-----+-----+****
***| possible_keys | key      | key_len | ref  | ***
***+-----+-----+-----+-----+***
***| int_key, ik   | int_key | 5       | const ***
***| NULL         | pk      | 9       | NULL ***
***+-----+-----+-----+-----+***
```

```
2 rows in set, 1 warning (0.00 sec)
```

Constant
Numeric in our case

Index used
to resolve rows

Which columns were compared with the index

Only one key was actually used

EXPLAIN in details: rows

% of rows filtered

Number of rows accessed

Additional information:
how query is resolved

Using filesort
Using temporary
etc.

```
mysql> explain extended select * from t1 join t2 where t1.int_key=1;
```

```
****+-----+-----+-----+-----+
***| rows  | filtered | Extra |
***+-----+-----+-----+-----+
***| 4      | 100.00 | NULL |
***| 6      | 100.00 | Using index; Using join buffer (Block Nested Loop) |
***+-----+-----+-----+-----+
2 rows in set, 1 warning (0.00 sec)
```

4X6
=
24

All rows used

EXPLAIN FORMAT=JSON

- Gives more information than regular EXPLAIN
 - Real execution path of the query
 - Pushed conditions
 - Query cost
 - Temporary table and index creation are more precise
 - Reflects execution order of "group by" and "order by" operations
 - Displays table materializations

EXPLAIN FORMAT=JSON

```
mysql> EXPLAIN FORMAT=JSON SELECT user, host FROM mysql.user\G
***** 1. row *****
EXPLAIN: {
  "query_block": {
    "select_id": 1,
    "table": {
      "table_name": "user",
      "access_type": "index",
      "key": "PRIMARY",
      "used_key_parts": [
        "Host",
        "User"
      ],
      "key_length": "228",
      "rows": 8,
      "filtered": 100,
      "using_index": true
    }
  }
}
```

EXPLAIN PARTITIONS

```
mysql> explain partitions select count(*) \  
-> from employees_part where hire_date > '1991-01-01'\G  
***** 1. row *****  
      id: 1  
select_type: SIMPLE  
      table: employees_part  
partitions: p1,p2  
      type: index  
possible_keys: NULL  
      key: PRIMARY  
     key_len: 7  
      ref: NULL  
     rows: 135214  
Extra: Using where; Using index
```

When EXPLAIN lies: Handler_%

```
mysql> flush status;
Query OK, 0 rows affected (0.00 sec)
mysql> select titles.* from titles join employees using (emp_no) \
-> where title='Senior Engineer' and to_date='9999-01-01' \
-> order by from_date limit 10;
```

```
+-----+-----+-----+-----+
| emp_no | title           | from_date | to_date |
+-----+-----+-----+-----+
| 235233 | Senior Engineer | 1985-02-01 | 9999-01-01 |
...
```

```
mysql> SHOW STATUS LIKE 'Handler_read_%';
```

```
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Handler_read_first | 1 |
| Handler_read_key | 11 |
...
| Handler_read_rnd_next | 443309 |
+-----+-----+
```

```
7 rows in set (0.00 sec)
```

I_S.OPTIMIZER_TRACE

- Debug trace of optimizer
- Prints
 - join_preparation
 - join_optimization
 - considered_execution_plans
 - refine_plan
 - More
 - join_execution

I_S.OPTIMIZER_TRACE example

```
mysql> SET optimizer_trace="enabled=on";  
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> select titles.* from titles join employees using (emp_no) \  
-> where title='Senior Engineer' and to_date='9999-01-01' \  
-> order by from_date limit 10\G
```

```
***** 1. row *****  
emp_no: 235233  
title: Senior Engineer  
from_date: 1985-02-01  
to_date: 9999-01-01  
...
```

I_S.OPTIMIZER_TRACE example cont.

```
mysql> SELECT * FROM INFORMATION_SCHEMA.OPTIMIZER_TRACE\G
***** 1. row *****
QUERY: select titles.* from titles join employees using (emp_no) where ...
TRACE: {
  "steps": [
    {
      "join_preparation": {
        "select#": 1,
        "steps": [
          {
            "expanded_query":
...
        {
          "join_optimization": {
            "select#": 1,
            "steps": [
              {
                "condition_processing": {
...

```

I_S.OPTIMIZER_TRACE example cont.

```
"considered_execution_plans": [
{
  "plan_prefix": [
  ],
  "table": "'employees'",
  "best_access_path": {
    "considered_access_paths": [
      {
        "access_type": "ref",
        "index": "PRIMARY",
        "usable": false,
        "chosen": false
      },
      {
        "rows_to_scan": 269763,
        "access_type": "scan",
        "resulting_rows": 269763,
        "cost": 54754,
        "chosen": true
      }
    ]
  }
}
```

I_S.OPTIMIZER_TRACE example cont.

```
"condition_filtering_pct": 100,
"rows_for_plan": 269763,
"cost_for_plan": 54754,
...
"rest_of_plan": [
{
  "plan_prefix": [
    "'titles'"
  ],
  "table": "'employees'",
  "best_access_path": {
    "considered_access_paths": [
      {
        "access_type": "eq_ref",
        "index": "PRIMARY",
        "rows": 1,
        "cost": 5222.9,
        "chosen": true,
        "cause": "clustered_pk_chosen_by_heuristics"
      },

```

I_S.OPTIMIZER_TRACE example cont.

```
    {
      "access_type": "scan",
      "chosen": false,
      "cause": "covering_index_better_than_full_scan"
    }
  ],
  "condition_filtering_pct": 100,
  "rows_for_plan": 4352.4,
  "cost_for_plan": 93459,
  "chosen": true
...

```

Internal concurrency issues

Common concurrency issues

- Query or transaction waits a lock, held by another one
- Fight for system resources
- Symptoms
 - Many processes are waiting
 - You see query in the slow log, but it runs fast if single-thread environment
 - Query execution time varies

Lock types and transactions

- Lock types
- Levels
 - MDL
 - Table-level
 - Row-level
- What do they lock
 - Read locks
 - Block writes
 - Write locks
 - Block reads and writes

- Transactions
 - Server-level
 - MDL locks
 - Table locks
 - Engine-level
 - Table locks
 - Row locks
 - AUTOCOMMIT
 - supported

Locks diagnostic

- SHOW [FULL] PROCESSLIST
- SHOW ENGINE INNODB STATUS
- INFORMATION SCHEMA
 - PROCESSLIST
 - InnoDB tables
- Performance Schema
 - METADATA_LOCKS
 - TABLE_HANDLES
 - EVENTS_TRANSACTIONS_*
 - Both server-level and engine-level
 - Contain GTID information

Locks diagnostic

- SHOW [FULL] PROCESSLIST
- For any kind of locks

```
mysql> select id, state, info from information_schema.processlist\G
***** 1. row *****
  id: 5
state:
  info: NULL
***** 2. row *****
  id: 4
state: Waiting for table metadata lock
  info: alter table titles add column f1 int
***** 3. row *****
  id: 2
state: executing
  info: select id, state, info from information_schema.processlist
```

Locks diagnostic

- Performance Schema
- Metadata locks

```
mysql> select processlist_id, object_type, lock_type, lock_status, source
  -> from metadata_locks join threads on (owner_thread_id=thread_id)
  -> where object_schema='employees' and object_name='titles'\G
***** 1. row *****
processlist_id: 4
  object_type: TABLE
    lock_type: EXCLUSIVE
  lock_status: PENDING -- waits
      source: mdl.cc:3263
***** 2. row *****
processlist_id: 5
  object_type: TABLE
    lock_type: SHARED_READ
  lock_status: GRANTED -- holds
      source: sql_parse.cc:5707
```

Locks diagnostic

- Performance Schema
- Table locks

```
mysql2> select * from table_handles\G
***** 1. row *****
      OBJECT_TYPE: TABLE
      OBJECT_SCHEMA: employees
      OBJECT_NAME: employees
OBJECT_INSTANCE_BEGIN: 140544885988272
      OWNER_THREAD_ID: 23
      OWNER_EVENT_ID: 818320
      INTERNAL_LOCK: NULL
      EXTERNAL_LOCK: READ EXTERNAL -- Table lock!
1 row in set (0.00 sec)
```

Locks diagnostic

- Performance Schema
- Transactions

```
mysql> select thread_id, processlist_id, access_mode, gtid \  
-> from events_transactions_current join threads using(thread_id) \  
-> where state = 'active'\G  
***** 1. row *****  
  thread_id: 25  
processlist_id: 2  
  access_mode: READ WRITE  
      gtid: AUTOMATIC  
***** 2. row *****  
  thread_id: 28  
processlist_id: 5  
  access_mode: READ WRITE  
      gtid: 07aff4c6-87e7-11e5-9a51-30b5c2208a0f:202  
2 rows in set (0.00 sec)
```

InnoDB Monitors

- SHOW ENGINE INNODB STATUS
- Permanent output
 - innodb_status_output
 - innodb_status_output_locks
 - innodb-status-file - Deleted on normal shutdown!

InnoDB Monitors

- SHOW ENGINE INNODB STATUS

```
-----  
TRANSACTIONS  
-----
```

```
Trx id counter 0 26243837
```

```
Purge done for trx's n:o < 0 26243834 undo n:o < 0 0
```

```
History list length 2
```

```
LIST OF TRANSACTIONS FOR EACH SESSION:
```

```
---TRANSACTION 0 26243836, ACTIVE 4 sec, OS thread id 101514240 starting index
```

```
mysql tables in use 1, locked 1
```

```
LOCK WAIT 2 lock struct(s), heap size 320, 1 row lock(s)
```

```
MySQL thread id 4485, query id 25022137 localhost root Updating
```

```
update t set a=36 where a=6
```

Other InnoDB diagnostic

- INFORMATION_SCHEMA
 - INNODB_TRX
 - INNODB_LOCKS
 - INNODB_LOCK_WAITS
 - INNODB_METRICS
 - Options innodb_monitor_*
- Option innodb_print_all_deadlocks

Other InnoDB diagnostic

• INFORMATION_SCHEMA

```
mysql> SELECT TRX_ID, TRX_REQUESTED_LOCK_ID, BLOCKING_TRX_ID, \
-> BLOCKING_LOCK_ID, TRX_QUERY \
-> FROM INNODB_TRX JOIN INNODB_LOCK_WAITS ON \
-> (INNODB_TRX.TRX_ID=INNODB_LOCK_WAITS.REQUESTING_TRX_ID) \
-> WHERE TRX_STATE = 'LOCK WAIT'\G
***** 1. row *****
          TRX_ID: 2318
TRX_REQUESTED_LOCK_ID: 2318:28:5:2
          BLOCKING_TRX_ID: 2316
          BLOCKING_LOCK_ID: 2316:28:5:2
          TRX_QUERY: update titles set title='Principal Engineer'
                    where title='Senior Engineer'
1 row in set (0.00 sec)
```

```
mysql> select TRX_ID, trx_mysql_thread_id from INNODB_TRX WHERE TRX_ID=2316\G
***** 1. row *****
          TRX_ID: 2316
trx_mysql_thread_id: 2
1 row in set (0.00 sec)
```

Locks diagnostic summary

- Table-level
 - PROCESSLIST: "Waiting for table lock"
 - P_S.TABLE_HANDLES
- Row-level
 - InnoDB monitors
 - SHOW ENGINE INNODB STATUS
 - Tables in INFORMATION_SCHEMA
 - P_S.EVENTS_TRANSACTIONS
 - Option `-innodb_print_all_deadlocks`
- MDL
 - PROCESSLIST
 - "Waiting for metadata lock"
 - P_S.METADATA_LOCKS

Environment

Memory usage

- What uses memory in MySQL
 - Buffers
 - Temporary tables
 - Internal structures - **out of user control!**
 - OS did not show memory as freed yet?

Memory diagnostic before 5.7

- free
- top
- vmstat
- Investigation
- There was no way to know how exactly memory was allocated

Memory diagnostic before 5.7

- free

```
$free
total      used      free      shared    buffers    cached
Mem:      16149184 6223916 9925268 317536    1048    3655160
-/+ buffers/cache: 2567708 13581476
Swap:      2110460      0    2110460
```

- top

- vmstat

- Investigation

Memory diagnostic before 5.7

- free
- top

\$top

```
Tasks: 295 total,   3 running, 292 sleeping,   0 stopped,   0 zombie
%Cpu(s):  3.0 us,  0.8 sy,  0.1 ni, 95.4 id,  0.8 wa,  0.0 hi,  0.0 si,  0.0 st
KiB Mem: 16149184 total, 6231688 used, 9917496 free,   1048 buffers
KiB Swap: 2110460 total,   0 used, 2110460 free. 3670752 cached Mem
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
1914	mysql	20	0	670m	95m	1296	S	0.7	1.2	2:42.14	mysqld

- vmstat
- Investigation

Memory diagnostic before 5.7

- free
- top
- vmstat

```
$vmstat -t 5 3
```

```
procs -----memory-----swap-----
r  b   swpd   free   buff  cache   si   so   bi   bo   in   cs  us  sy  id  wa...
2  0     0 9923160 1048 3662724  0   0  168   86  167  674  3  1 87...
0  0     0 9923252 1048 3662904  0   0   30  122 1168 5264  3  1 96...
0  0     0 9922864 1048 3663120  0   0   25  128 1191 5342  2  1 96...
```

- Investigation

Memory diagnostic before 5.7

- free
- top
- vmstat
- Investigation
 - Total size of buffers
 - Number of temporary tables
 - Number of parallel connections

Memory diagnostic in 5.7

```
mysql> select thread_id tid, user, current_allocated ca, total_allocated  
-> from sys.memory_by_thread_by_current_bytes;
```

tid	user	ca	total_allocated
1	sql/main	2.53 GiB	2.69 GiB
150	root@127.0.0.1	4.06 MiB	32.17 MiB
146	sql/slave_sql	1.31 MiB	1.44 MiB
145	sql/slave_io	1.08 MiB	2.79 MiB
...			
60	innodb/io_read_thread	0 bytes	384 bytes
139	innodb/srv_purge_thread	-328 bytes	754.21 KiB
69	innodb/io_write_thread	-1008 bytes	34.28 KiB
68	innodb/io_write_thread	-1440 bytes	298.05 KiB
74	innodb/io_write_thread	-1656 bytes	103.55 KiB
4	innodb/io_log_thread	-2880 bytes	132.38 KiB
72	innodb/io_write_thread	-7632 bytes	1.10 MiB

```
145 rows in set (2.65 sec)
```

Threads Statistics

```
mysql> select * from sys.memory_by_thread_by_current_bytes
      -> order by current_allocated desc\G
***** 1. row *****
      thread_id: 152
         user: lj@127.0.0.1
current_count_used: 325
  current_allocated: 36.00 GiB
  current_avg_alloc: 113.43 MiB
  current_max_alloc: 36.00 GiB
    total_allocated: 37.95 GiB
...

```

RAW Performance Schema tables

- `memory_summary_by_account_by_event_name`
- `memory_summary_by_host_by_event_name`
- `memory_summary_by_thread_by_event_name`
- `memory_summary_by_user_by_event_name`
- `memory_summary_global_by_event_name`
- `sys` schema also includes information about user name

Users in sys.memory_* tables

- NAME@HOST - regular user
- System users
 - sql/main
 - innodb/*
 - ...
- Data comes from table THREADS

Memory usage: how to fix?

- mysqld uses too low memory
 - Buffers are not tuned for performance
 - Operating system limits resources
- mysqld uses too much memory
 - Buffers are too large
 - General use buffers: allocated once for whole server
 - Connection buffers: allocated for every connection
 - Task buffers: allocated for the particular task
 - Prioritize!
 - Allocations, which cannot be controlled
 - Watch
 - Find workaround
 - Implement

Disk diagnostics

- df
- iostat
- ls -l /proc/{PID_OF_MYSQLD}/fd
- InnoDB status
- performance_schema.table_io_waits_%
- sys.io_%

Disk diagnostics

- df

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
/dev/sda6	41944064	28429816	11433384	72%	/
devtmpfs	8067632	0	8067632	0%	/dev
tmpfs	8074592	41300	8033292	1%	/dev/shm
tmpfs	8074592	2376	8072216	1%	/run
tmpfs	8074592	0	8074592	0%	/sys/fs/cgroup
/dev/sda6	41944064	28429816	11433384	72%	/.snapshots
/dev/sda6	41944064	28429816	11433384	72%	/var/tmp
/dev/sda6	41944064	28429816	11433384	72%	/var/spool
/dev/sda6	41944064	28429816	11433384	72%	/var/opt
/dev/sda6	41944064	28429816	11433384	72%	/var/log
/dev/sda6	41944064	28429816	11433384	72%	/var/lib/pgsql
/dev/sda6	41944064	28429816	11433384	72%	/var/lib/named
/dev/sda6	41944064	28429816	11433384	72%	/var/lib/mailman
/dev/sda6	41944064	28429816	11433384	72%	/var/crash
/dev/sda6	41944064	28429816	11433384	72%	/usr/local
/dev/sda6	41944064	28429816	11433384	72%	/tmp
...					

Disk diagnostics

- iostat

```
sveta@thinkie:~/src/test-db/employees_db> iostat -x
Linux 3.16.7-92-desktop (thinkie) 11/11/2015 _x86_64_ (8 CPU)

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           14.47    1.77    2.54    0.94    0.00   80.28

Device: rrqm/s  wrqm/s     r/s     w/s    rkB/s    kB/s avgrq-sz avgqu-sz ...
sda      0.28    0.70    4.66   21.59   93.69   346.10   33.51    4.14 ...
sdb      0.00    0.00    0.00    0.00    0.02    0.00   15.69    0.00 ...
```

Disk diagnostics

- `ls -l /proc/{PID_OF_MYSQLD}/fd`

```
sveta@thinkie:~/src/test-db/employees_db> ls -l /proc/26676/fd
total 0
lrwx----- 1 sveta users 64 Nov 11 15:11 0 -> /dev/pts/1
l-wx----- 1 sveta users 64 Nov 11 15:11 1 -> /log/mysqlld.1.err
lrwx----- 1 sveta users 64 Nov 11 15:11 10 -> /data/ib_logfile1
lrwx----- 1 sveta users 64 Nov 11 15:11 11 -> /data/ibtmp1
lrwx----- 1 sveta users 64 Nov 11 15:11 12 -> /tmp/mysqlld.1/ibx2zTxG (deleted)
...
lrwx----- 1 sveta users 64 Nov 11 15:11 27 -> /data/sys/sys_config.ibd
l-wx----- 1 sveta users 64 Nov 11 15:11 28 -> /data/master-bin.000002
lrwx----- 1 sveta users 64 Nov 11 15:11 29 -> socket:[1245731]
lrwx----- 1 sveta users 64 Nov 11 15:11 3 -> /data/master-bin.index
...
lrwx----- 1 sveta users 64 Nov 11 15:11 5 -> /tmp/mysqlld.1/ibT4odfr (deleted)
lrwx----- 1 sveta users 64 Nov 11 15:11 6 -> /tmp/mysqlld.1/ibJoMdZu (deleted)
lrwx----- 1 sveta users 64 Nov 11 15:11 7 -> /tmp/mysqlld.1/ibvUBeJy (deleted)
lrwx----- 1 sveta users 64 Nov 11 15:11 8 -> /tmp/mysqlld.1/ibP6JxvC (deleted)
lrwx----- 1 sveta users 64 Nov 11 15:11 9 -> /data/ib_logfile0
```

Disk diagnostics

- InnoDB status

```
mysql> show engine innodb status\G
...
-----
FILE I/O
-----
I/O thread 0 state: waiting for i/o request (insert buffer thread)
I/O thread 1 state: waiting for i/o request (log thread)
I/O thread 2 state: waiting for i/o request (read thread)
I/O thread 3 state: waiting for i/o request (read thread)
I/O thread 4 state: waiting for i/o request (write thread)
I/O thread 5 state: waiting for i/o request (write thread)
Pending normal aio reads: [0, 0] , aio writes: [0, 0] ,
ibuf aio reads:, log i/o's:, sync i/o's:
Pending flushes (fsync) log: 0; buffer pool: 0
11468 OS file reads, 102 OS file writes, 17 OS fsyncs
11.80 reads/s, 16384 avg bytes/read, 0.00 writes/s, 0.00 fsyncs/s
...
```

Disk diagnostics

- InnoDB status

```
-----  
ROW OPERATIONS  
-----  
0 queries inside InnoDB, 0 queries in queue  
0 read views open inside InnoDB  
Process ID=26676, Main thread ID=140424906065664, state: sleeping  
Number of rows inserted 19, updated 0, deleted 0, read 5843307  
0.00 inserts/s, 0.00 updates/s, 0.00 deletes/s, 1584.23 reads/s
```

Disk diagnostics

```
mysql> select OBJECT_SCHEMA, OBJECT_NAME, INDEX_NAME \
      -> from table_io_waits_summary_by_index_usage where COUNT_STAR=0;
```

OBJECT_SCHEMA	OBJECT_NAME	INDEX_NAME
employees	departments	PRIMARY
employees	departments	dept_name
employees	dept_emp	PRIMARY
employees	dept_emp	emp_no
employees	dept_emp	dept_no
employees	dept_manager	PRIMARY
employees	dept_manager	emp_no
employees	dept_manager	dept_no
employees	salaries	PRIMARY
employees	salaries	emp_no
sys	sys_config	PRIMARY

```
11 rows in set (0.00 sec)
```

Disk diagnostics

- `sys.io_%`

```
mysql> select event_name, count_read, total_read, total_requested \  
-> from io_global_by_wait_by_bytes;
```

event_name	count_read	total_read	total_requested
innodb/innodb_data_file	17349	274.36 MiB	289.22 MiB
sql/binlog	16696	130.42 MiB	130.42 MiB
sql/io_cache	1556	37.61 MiB	100.49 MiB
sql/FRM	1137	626.03 KiB	626.03 KiB
innodb/innodb_log_file	8	132.50 KiB	141.50 KiB
sql/ERRMSG	3	72.23 KiB	72.23 KiB
sql/query_log	0	0 bytes	44.02 KiB
myisam/dfile	20	21.99 KiB	21.99 KiB
mysys/charset	1	18.27 KiB	18.27 KiB
myisam/kfile	24	6.62 KiB	6.62 KiB
sql/file_parser	3	2.11 KiB	2.11 KiB
sql/slow_log	0	0 bytes	1.64 KiB
mysys/cnf	3	56 bytes	56 bytes
...			

Disk diagnostics

- sys.io_%
- Other useful views
 - host_summary_by_file_io
 - host_summary_by_file_io_type
 - io_by_thread_by_latency
 - io_global_by_file_by_bytes - **Per-file**
 - io_global_by_file_by_latency
 - io_global_by_wait_by_bytes
 - io_global_by_wait_by_latency
 - latest_file_io - **Latest IO operations**
 - user_summary_by_file_io
 - user_summary_by_file_io_type

How to fix IO issues

- Optimize queries
- InnoDB options, affecting writes
- Binary log

How to fix IO issues

- Optimize queries
- InnoDB options, affecting writes
 - innodb_adaptive_flushing
 - innodb_doublewrite
 - innodb_flush_log_at_trx_commit
 - innodb_flush_method - Test!
 - innodb_io_capacity - Follow docs
 - innodb_max_dirty_pages_pct

How to fix IO issues

- Optimize queries
- Binary log
 - Formats
 - ROW
 - STATEMENT
 - MIXED
 - binlog_row_image
 - sync_binlog

CPU diagnostics

- top
- iostat
- ps

CPU diagnostics

- top

```
Tasks: 298 total,  2 running, 296 sleeping,  0 stopped,  0 zombie
%Cpu(s): 15.2 us,  2.7 sy,  1.8 ni, 79.4 id,  0.9 wa,  0.0 hi,  0.0 si,  0.0 st
KiB Mem: 16149184 total, 15075560 used,  1073624 free,    1596 buffers
KiB Swap: 2110460 total,    8252 used,  2102208 free. 9994956 cached Mem
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
26676	sveta	20	0	1963204	155108	23300	S	100.0	0.960	0:43.73	mysqld
10365	sveta	20	0	3601524	1.555g	108488	S	25.00	10.10	224:21.91	firefox
29915	sveta	20	0	15476	2612	2120	R	12.50	0.016	0:00.02	top

...

CPU diagnostics

- iostat

```
sveta@thinkie:~/src/test-db/employees_db> iostat -x
Linux 3.16.7-92-desktop (thinkie) 11/11/2015 _x86_64_ (8 CPU)

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           14.47    1.77    2.54    0.94    0.00   80.28

Device: rrqm/s  wrqm/s     r/s     w/s    rkB/s    kB/s avgrq-sz avgqu-sz ...
sda      0.28    0.70    4.66   21.59   93.69   346.10   33.51    4.14 ...
sdb      0.00    0.00    0.00    0.00    0.02    0.00   15.69    0.00 ...
```

CPU diagnostics

- ps

```
sveta@thinkie> ps -eo pid,user,comm,pcpu,pmem,vsz | grep mysqld
26676 sveta      mysqld          1.4  0.9 1963204
26678 sveta      mysqld          0.0  0.7 1765772
```

CPU: problem solving

- Low load
 - Is your mysqld really active?
 - Check OS limits
 - ulimit
 - User options
 - Global options
 - Increase concurrency-related options
 - innodb_thread_concurrency
 - innodb_read_io_threads
 - innodb_write_io_threads
 - Max (2 X (number of CPU cores))!
- High load
 - Move job to memory: increase buffers
 - Tune queries

Network: what is worth attention?

- Stability
- Bandwidth
- Speed
 - RTT

```
sveta@thinkie> ping www.percona.com -c 2
PING www.percona.com (74.121.199.234) 56(84) bytes of data.
64 bytes from www.percona.com (74.121.199.234): icmp_seq=1 ttl=54 time=181.380ms
64 bytes from www.percona.com (74.121.199.234): icmp_seq=2 ttl=54 time=181.685ms

--- www.percona.com ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 181.380/181.685/181.991/0.524 ms
```


Network:diagnostics

- `-log-warnings=2`
- Send huge file (1 G or larger)
- `tcpdump`
- `P_S.accounts, hosts, users`
- `P_S.host_cache`

Network:diagnostics

- tcpdump

```
$tcpdump -i en1 host master.mysql.com and port 33051
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on en1, link-type EN10MB (Ethernet), capture size 96 bytes
22:28:12.195270 IP master.mysql.com.33051 > apple.60344: P
1752426772:1752426864(92) ack 1474226199 win 91 <nop,nop,timestamp 1939999898
649946687>
22:28:12.195317 IP apple.60344 > master.mysql.com.33051: . ack 92 win 65535
<nop,nop,timestamp 649946998 1939999898>
^C
2 packets captured
37 packets received by filter
0 packets dropped by kernel
```

Network:diagnostics

- P_S.accounts, hosts, users

```
mysql> select user, host, current_connections as cur, \  
        -> total_connections as total from accounts;
```

```
+-----+-----+-----+-----+  
| user | host      | cur | total |  
+-----+-----+-----+-----+  
| foo  | localhost | 0   | 3     |  
| root | localhost | 1   | 3     |  
| NULL | NULL      | 14  | 17    |  
+-----+-----+-----+-----+
```

```
3 rows in set (0.01 sec)
```

Network:diagnostics

- P_S.host_cache

```
mysql> select IP, HOST, HOST_VALIDATED, SUM_CONNECT_ERRORS, \  
  -> COUNT_NAMEINFO_PERMANENT_ERRORS, COUNT_PROXY_USER_ERRORS, \  
  -> COUNT_DEFAULT_DATABASE_ERRORS, FIRST_SEEN, LAST_SEEN from host_cache\G  
***** 1. row *****  
      IP: 10.159.169.27  
      HOST: foo.bar.com  
      HOST_VALIDATED: YES  
      SUM_CONNECT_ERRORS: 1  
COUNT_NAMEINFO_PERMANENT_ERRORS: 0  
      COUNT_PROXY_USER_ERRORS: 0  
COUNT_DEFAULT_DATABASE_ERRORS: 1  
      FIRST_SEEN: 2014-06-10 20:29:45  
      LAST_SEEN: 2015-11-11 15:32:44
```

More information

Further reading

- MySQL Troubleshooting book
- High Performance MySQL book
- Planet MySQL
- MySQL User Reference Manual
- Bug trackers
 - <http://bugs.mysql.com>
 - <https://bugs.launchpad.net/percona-server/>

Current Percona vacancies

www.percona.com/about-percona/careers

Place for your questions

???

Thank you!

<http://www.slideshare.net/SvetaSmirnova>

<https://twitter.com/svetsmirnova>