Which Database is Better for Zabbix ? PostgreSQL vs MySQL

Yoshiharu Mori
SRA OSS Inc. Japan

About Myself

Yoshiharu Mori

Belongs To: SRA OSS,Inc.Japan

Division: OSS technical Group

Providing support and consulting services for 30 or more kind of OSS

Zabbix experience

- We started Zabbix support service from 2011
- Last year, we joined Zabbix patnership program.

- Our company systems is monitored by Zabbix.
- We design the monitoring system of our customer's.
  - Over 100 devices
  - Including virtulization system (vCenter /ESXi)

Agenda

- Introduction
- Simple test

  large number of items
- Partitioning test

  large number of items with Partitioning

Which database engine are you using for Zabbix?

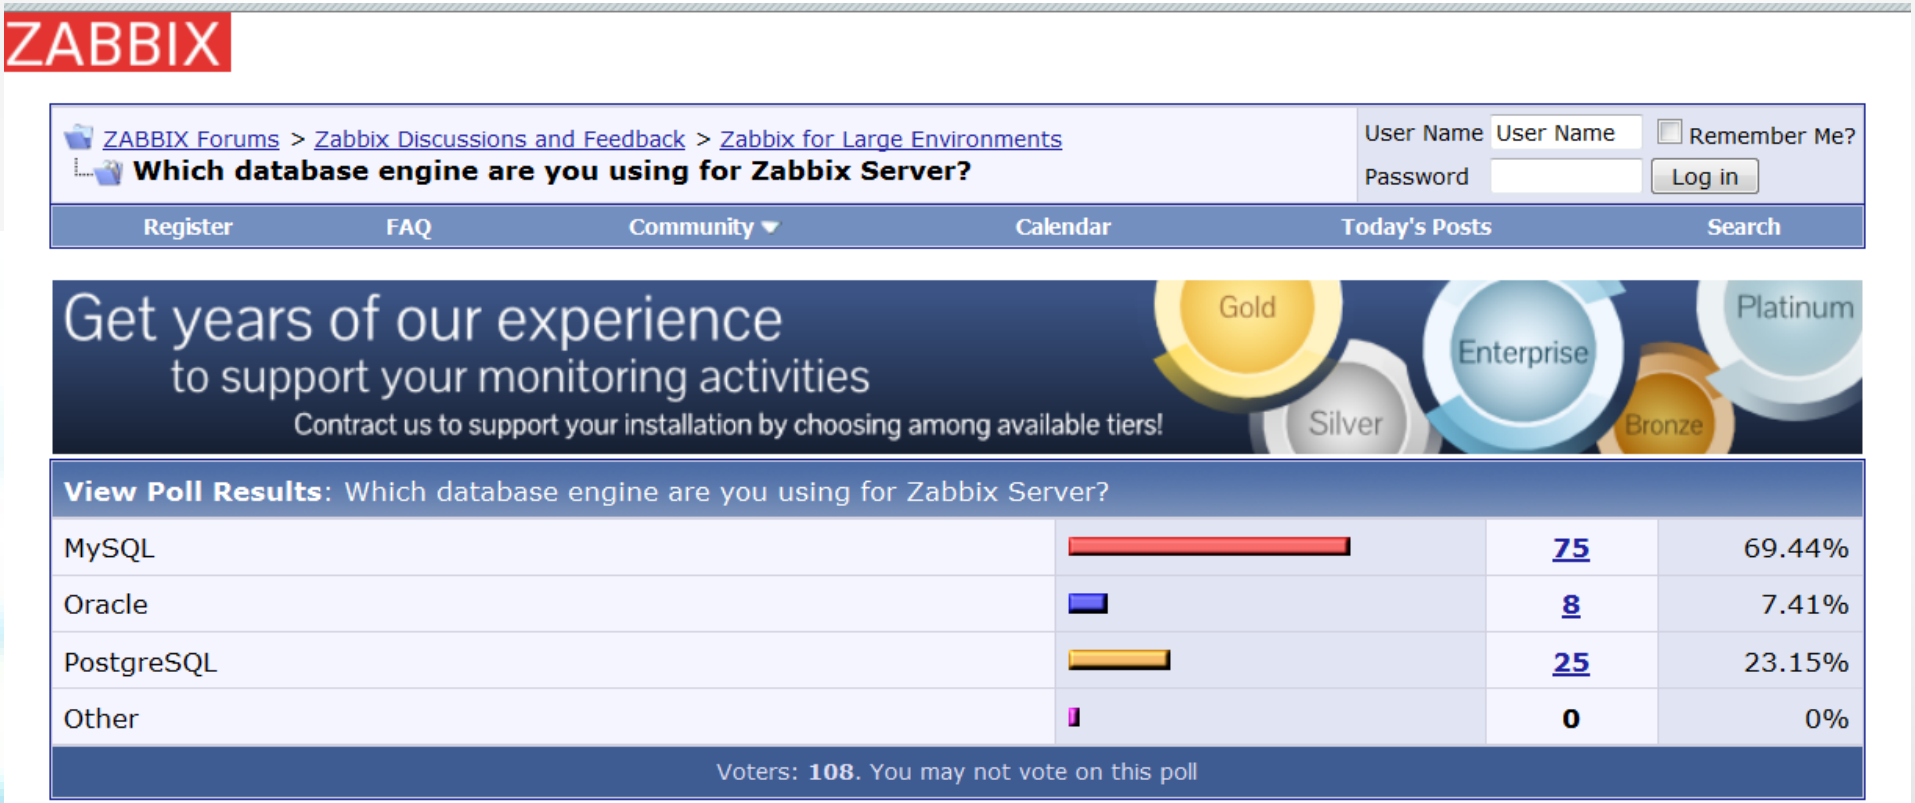

PostgreSQL is Slow?

- The Performance of PG was improved from 9.2!

  Compared with PG9.1
  - Read Query 4x Faster
  - Write Query 5x Faster

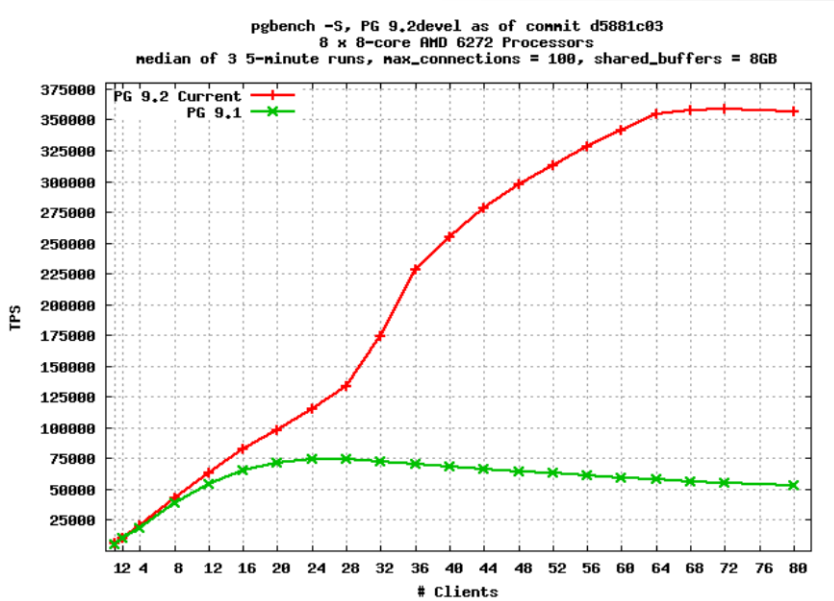

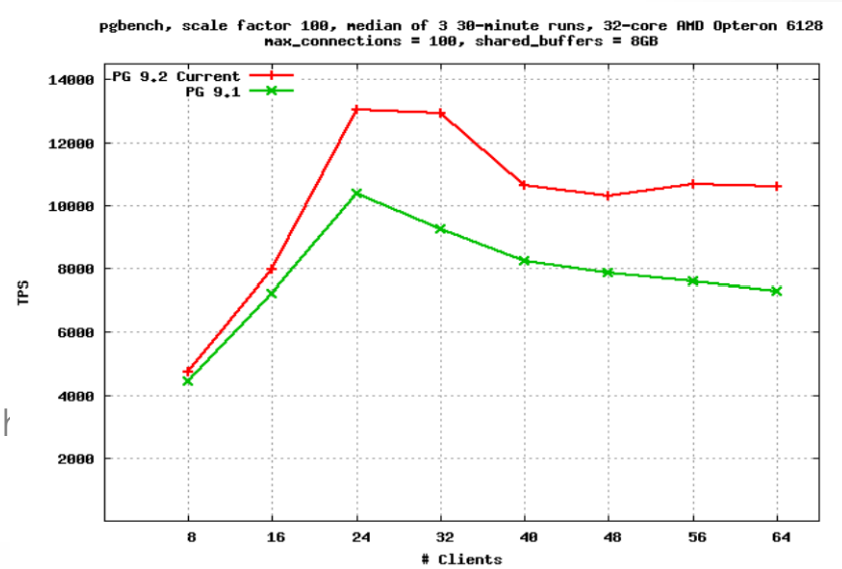

Reference:
PostgreSQL 9.2 release news
http://www.postgresql.org/about/news/1415/

Robert Haas@PGcon2012
Performance Improvement in PostgreSQL 9.2
http://www.pgcon.org/2012/schedule/events/416.en.h

MySQL vs PostgreSQL

- Zabbix benchmark with MySQL and PostgreSQL
- Environment
  - Amazon EC2 M1 medium instance
  - 3.75 Gib RAM
  - 2ECU (1 core)
  - 410 GB instance storage
  - OS: Amazon Linux AMI 2013.03 (64bit)
  - 3 servers (1 for Zabbix server and DB ,2 for Zabbix Agent)

Software

- Zabbix 2.1.1 (alpha release)
- Use latest release version for the components
  - MySQL 5.6.12
  - PostgreSQL 9.2.4
  - Apache httpd 2.4.6
  - PHP 5.4.17

Test Target

- Simulate 600 hosts (2 agent servers)
- 26400 items
  (44 items/1 host,5 sec interval, keep history 1day)
- 10200 triggers (17 triggers / 1 host)
- Starting with empty history
- Monitoring for:
  4 hours (short run)
  36 hours (long run)
- Zabbix:
  modified CacheSize=64M in zabbix_server.conf

Simple Test

- Test 1
  - Default DB setting
- Test 2
  - Same buffer size and Transaction log size
- Each Test is
  - Monitoring 4 hours
  - comparing CPU utilization
    and Zabbix performance

Test 1: default DB Config

- MySQL
  - character-set-server=utf8
  - skip-character-set-client-handshake
  - innodb_file_per_table

- PostgreSQL
  - ALL default

    PostgreSQL uses files per table

Test 1: CPU utilization

MySQL

PostgreSQL

User System iowait

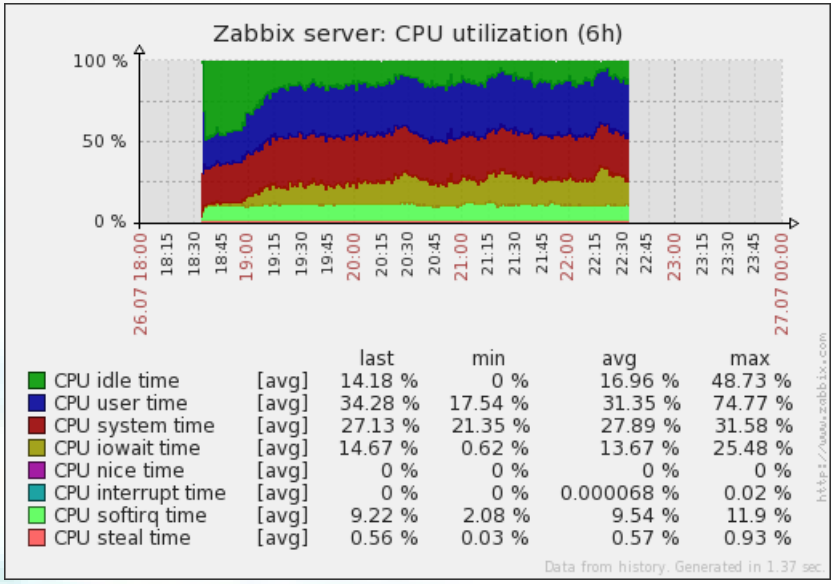

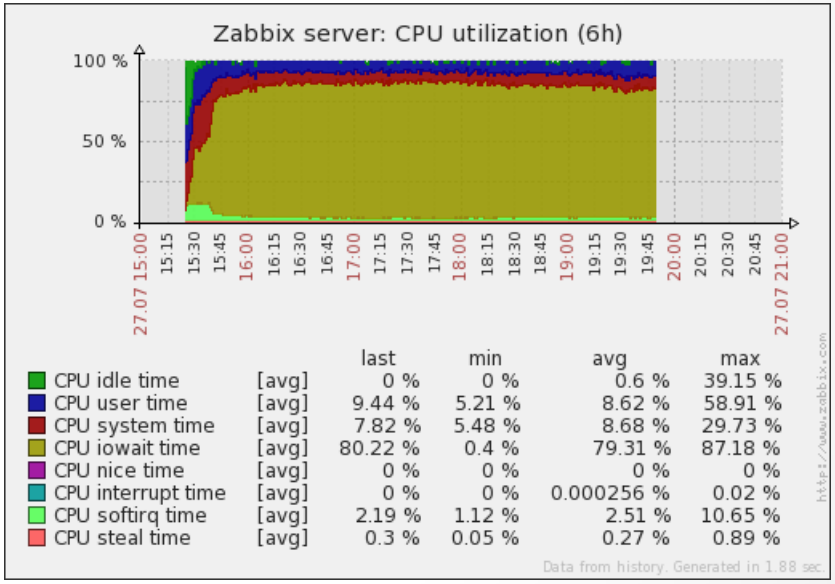

PG: Heavy io wait

Test 1: Zabbix Performance

MySQL

PostgreSQL

zabbix[wcache,values] :Values processed by Zabbix server per second

zabbix[queue]:Zabbix queue

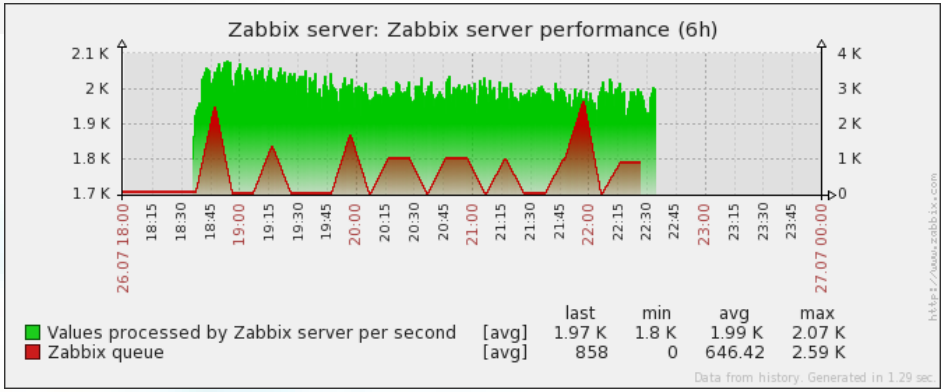

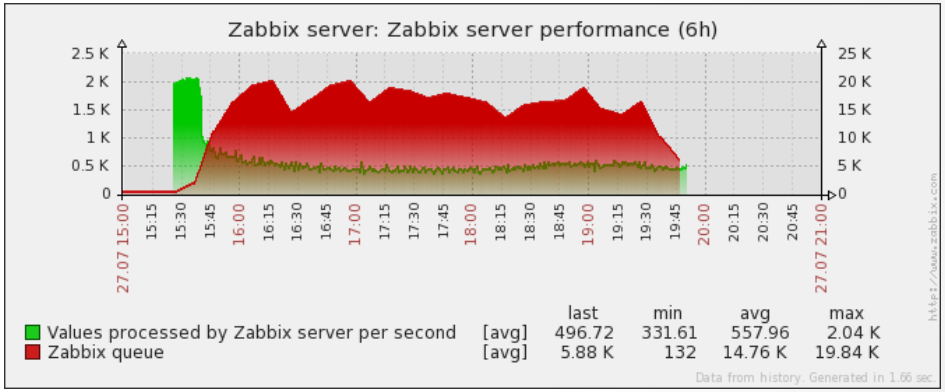

PG cannot keep up with processing

Test 2: Tuning

- Same Buffer and Transaction log size
- MySQL
  - Innodb_buffer_pool_size = 512MB
  - Innodb_log_file_size = 256MB
    x innodb_log_files_in_group=2(default)
    = log size 512MB
- PostgreSQL
  - shared_buffers = 512MB
  - checkpoint_segments = 32
    16MB/each segmnet →log size 512MB

Test 2: CPU utilization

MySQL

PostgreSQL

User System iowait

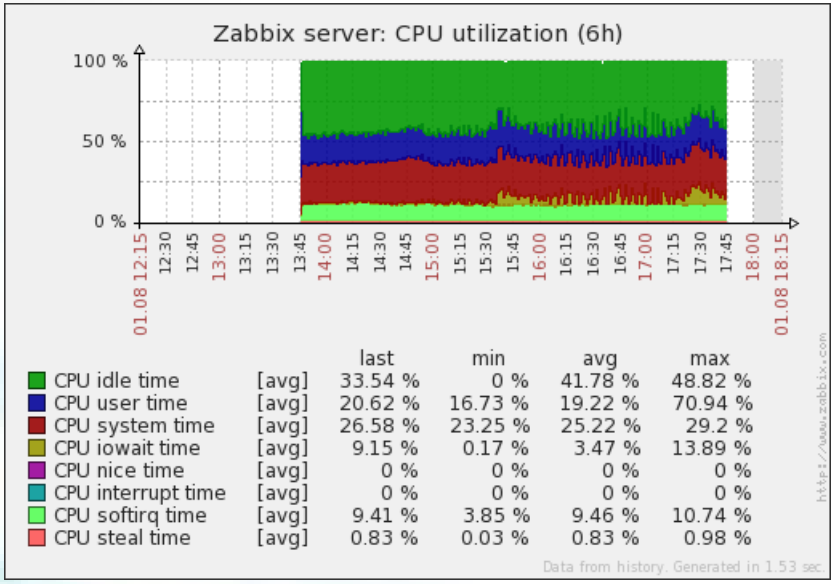

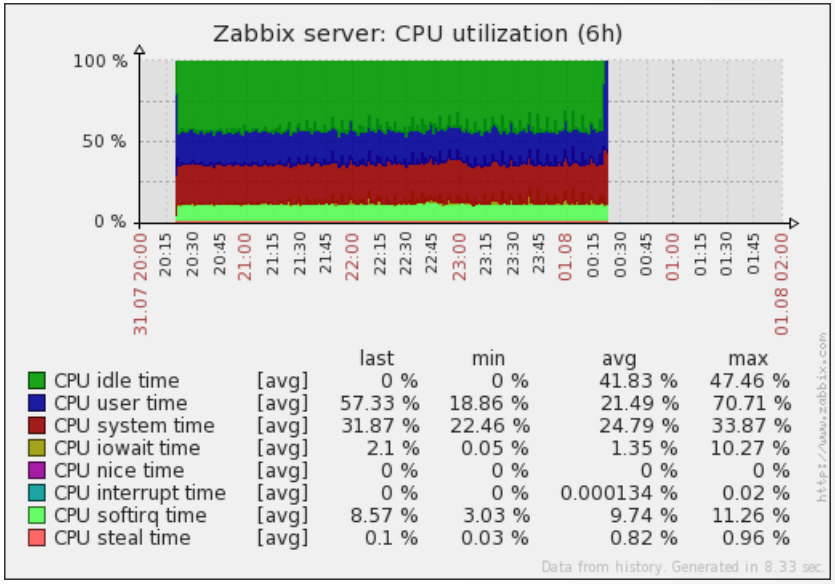

Almost same load (PG is slightly stable and low io)

Test 2: Zabbix Performance

MySQL

PostgreSQL

zabbix[wcache,values] :Values processed by Zabbix server per second

zabbix[queue]:Zabbix queue

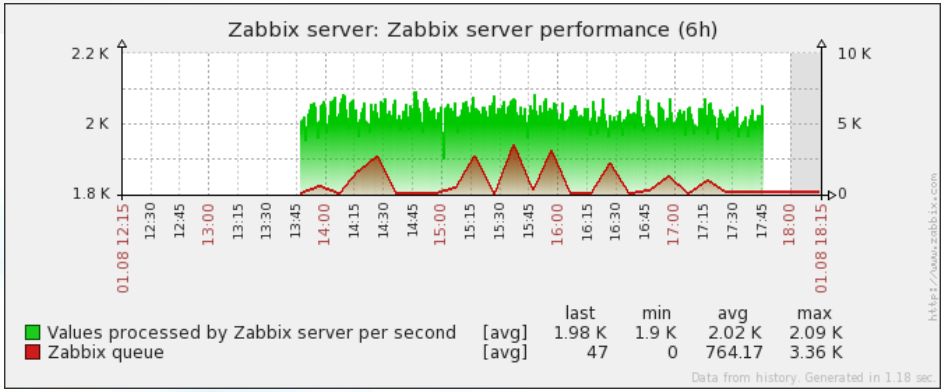

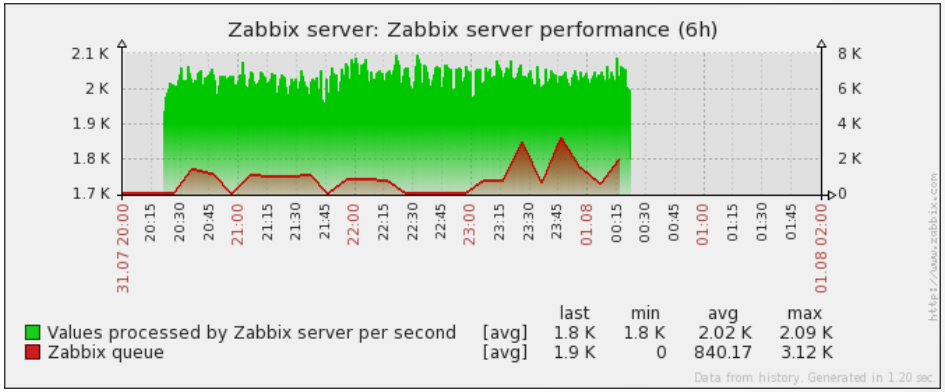

Compare DB Size

- Test 2 Case
- Exclude Transaction log

- MySQL : 3.3 GB
- PostgreSQL : 3.2 GB

- Almost Same Size

Study of Simple Test

- MySQL default setting is better than PostgreSQL.
- By Increasing Buffer and Transaction log size, MySQL and PostgreSQL are almost same performance.
- The size of DB is same in this Simple Test.

Characteristic of Zabbix DB

Large write and small read

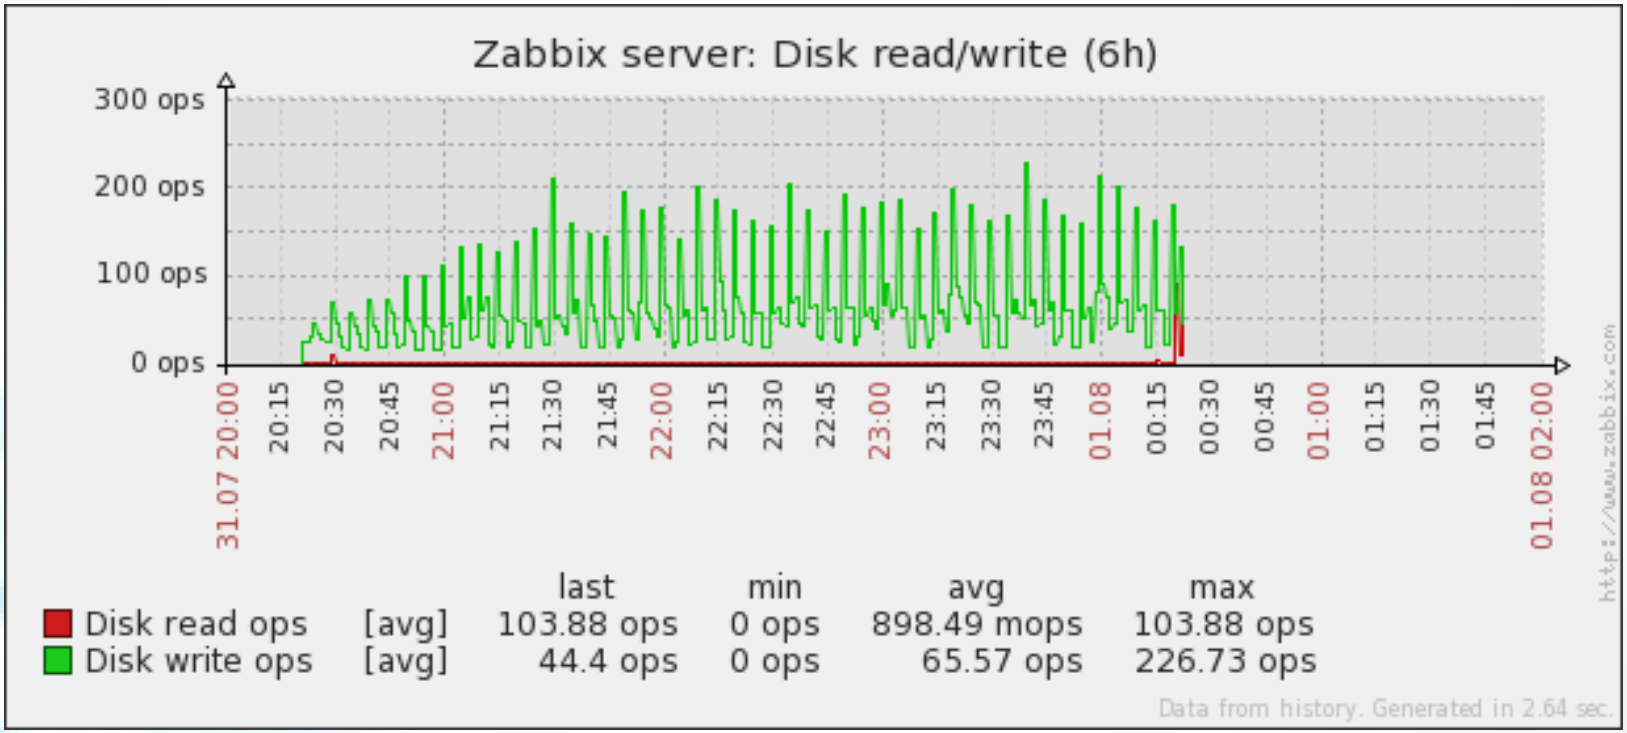

Tuning Point

- Checkpoints tuning is very important in write case.
  - Incorrect Checkpoints increase write load
  - The opportunity of checkpoints are
    - full of Buffers
    - full of Transaction log
  - To Improve Checkpoint,
    it's important of Buffer and Transaction log tuning

Zabbix DB Tuning

- Zabbix DB continue write operation by same rate.
- Estimate the growth rate of DB size

  In this test case 3.2GB/4h ≒ 70MB/5min

- Set a plenty of buffer and transaction log size from a growth rate of DB size.

  In this test case

  512MB buffers and 512MB transaction log size

Partitioning Test

- Same DB Tuning
  - Buffer and Total Transaction log size 512MB
- Test 3
  - Non-Partitioning and housekeeper is enabled.
  - Monitoring 36 hours
- Test 4
  - Partitioning daily and housekeeper is disabled.
  - Monitoring 36 hours
- Comparing CPU utilization and Zabbix performance

Test 3: CPU utilization

MySQL

PostgreSQL

User System iowait

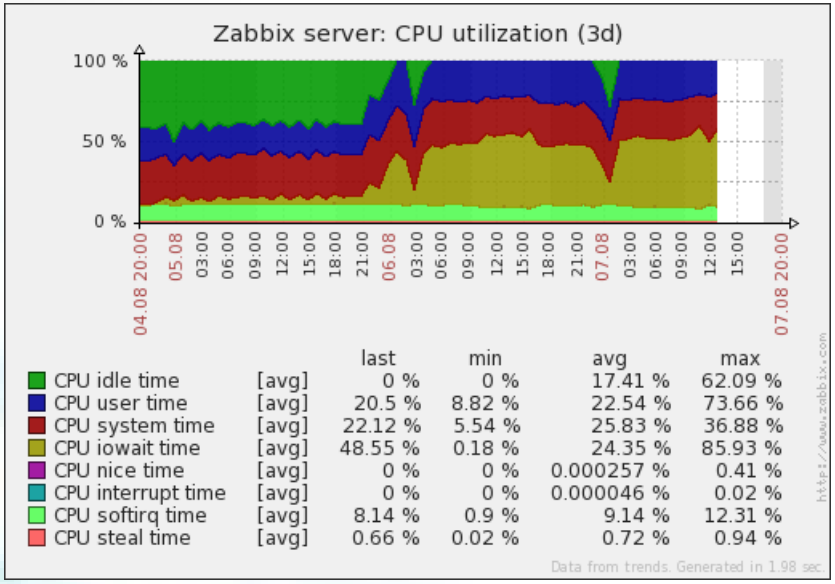

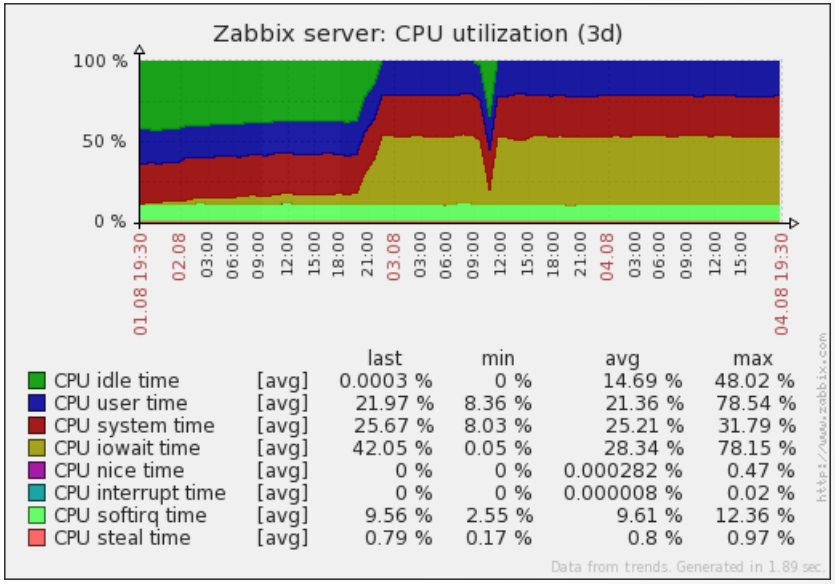

Housekeeper's Large io wait

Test 3: Zabbix Performance

MySQL

PostgreSQL

zabbix[wcache,values] :Values processed by Zabbix server per second

zabbix[queue]:Zabbix queue

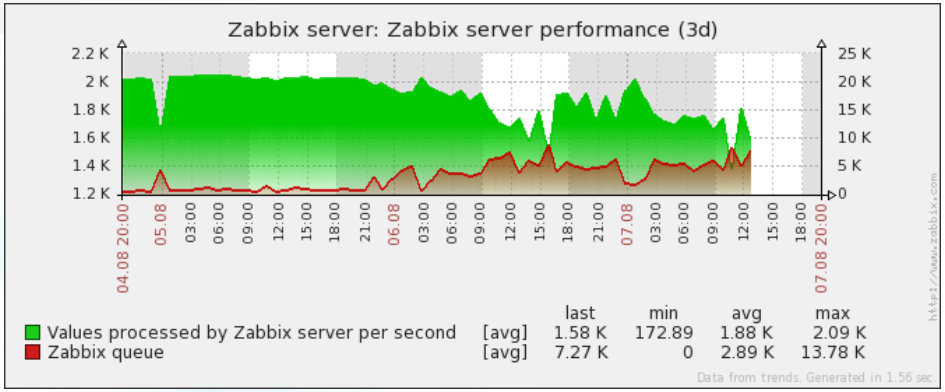

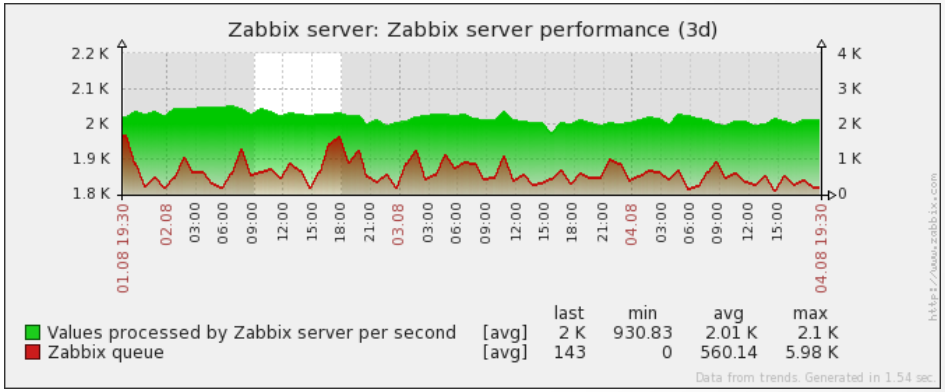

PostgreSQL is more stable in heavy io situation.

Test 4 Partitioning

- Disable Housekeeper
- Daily partitioning at history* tables.
- *PostgreSQL* supports *partitioning*
  - via table inheritance
  - define a trigger or rule.
    for redirecting data inserted into the parent table
    to the appropriate child partition.
  - The performance is also depend on the trigger or rule.
  - In this test case, we define the trigger
    - Using pl/pgsql
    - Using C language (heap_insert method)
      https://github.com/matheusoliveira/pg_partitioning_tests
      Developed by Matheus de Oliveira

Test 4: CPU utilization

MySQL

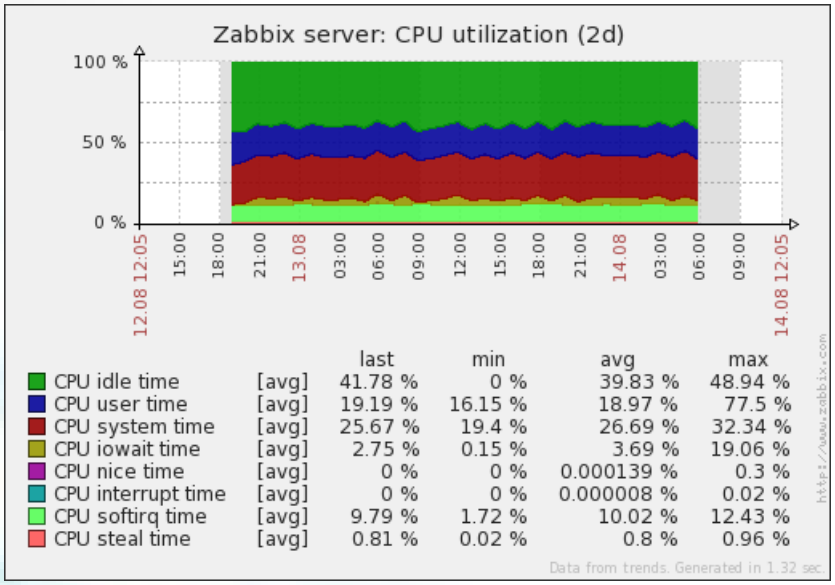

PG(pl/pgsql)

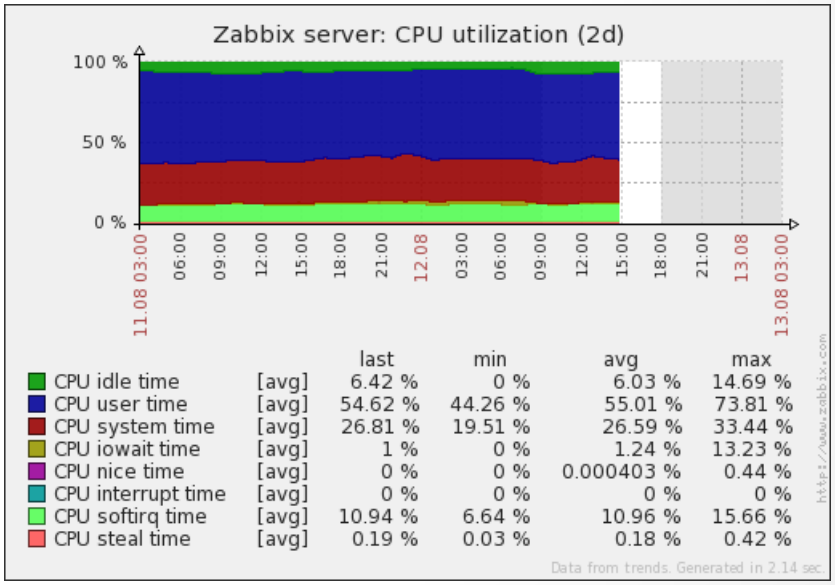

User System iowait

PG(pl/pgsql) need more 3x CPU(user) power

Test 4: Zabbix Performance

MySQL　　PG(pl/pgsql)

zabbix[wcache,values] :Values processed by Zabbix server per second

zabbix[queue]:Zabbix queue

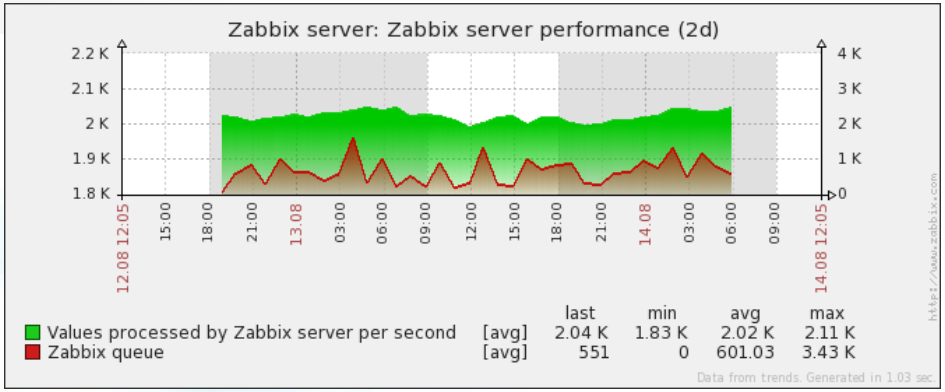

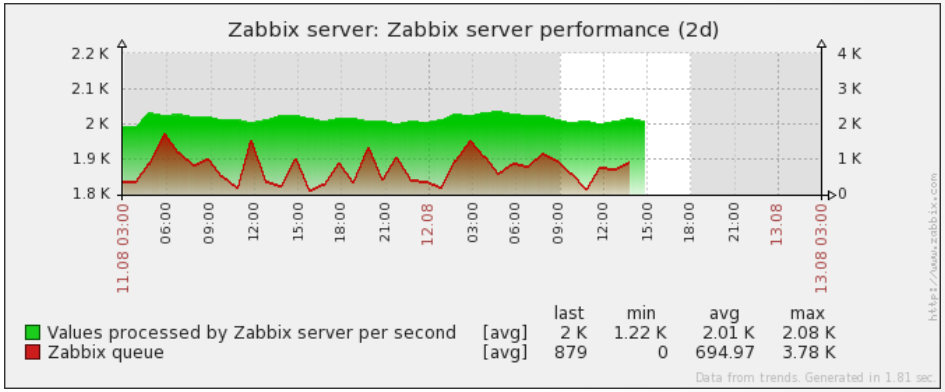

In spite of using more CPU, PG is still stable.

Test 4: CPU utilization

MySQL PG(C lang)

User System iowait

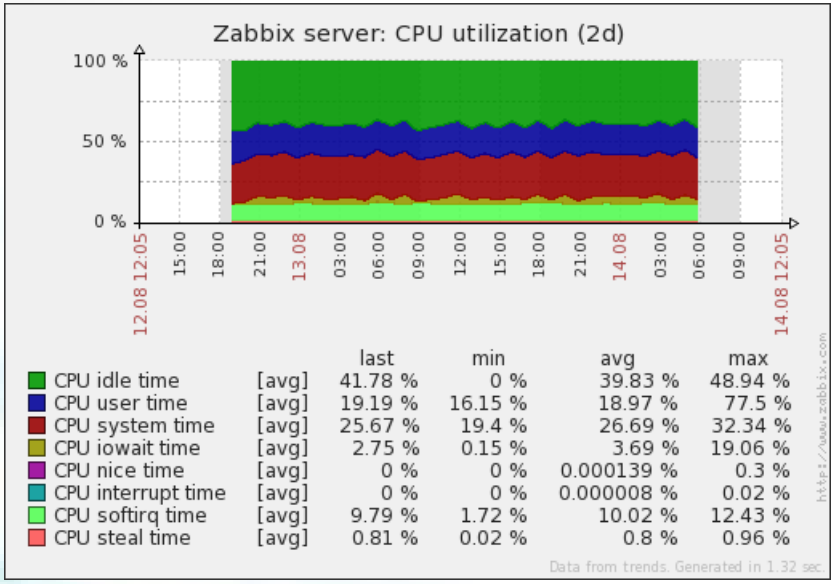

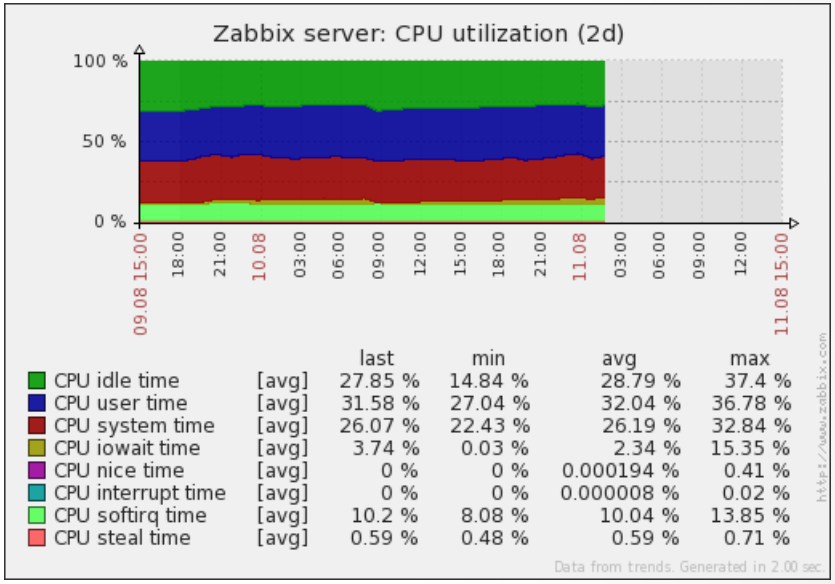

The cpu usage of PG(C lang) is better than PG(pl/pgsql), But it is worse than MySQL.

Test 4: Zabbix Performance

MySQL　　PG(C lang)

zabbix[wcache,values] :Values processed by Zabbix server per second

zabbix[queue]:Zabbix queue

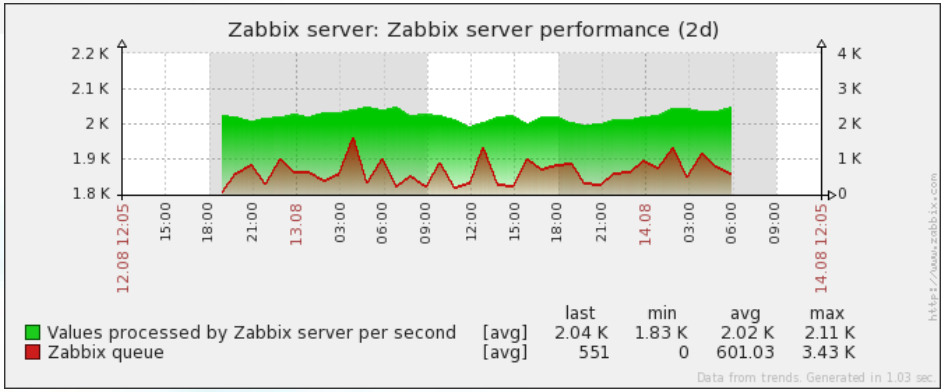

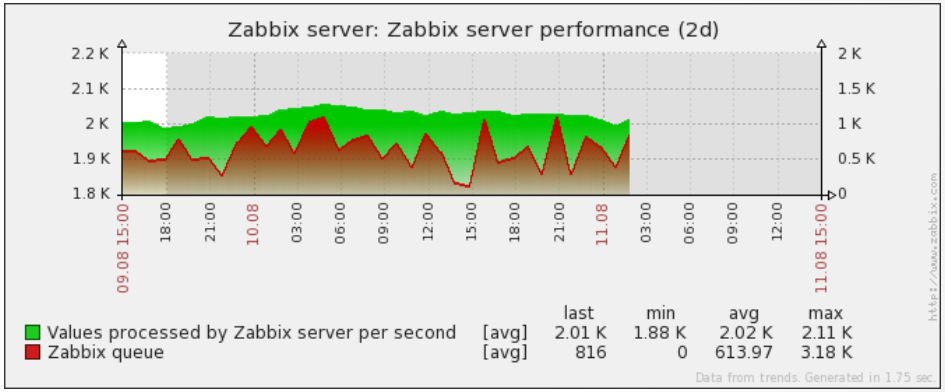

Almost same performance

Conclusion

- In our test
  - PostgreSQL and MySQL is almost same performance as Zabbix DB.
  - Tuning Checkpoint(buffer and transaction log) is important.
  - MySQL's partition is easy and better cpu usage than PostgreSQL.
    (PostgreSQL's partition supports foreign key)
  - PostgreSQL is more stable in heavy io situation.