Virident FlashMAX M1400
MySQL - Tpcc-Mysql Report

Percona Inc.
Vadim Tkachenko

November 2011

Contents

0 Introduction

The goal of this research is to compare the performance (throughput) of the Virident FlashMAX M1400 card with the Fusion-io ioDrive Duo 1.28TB card under a tpcc-mysql workload.

The research is sponsored by Virident Systems Inc.

1 Testing methodology

For the tests we used a `tpcc-mysql` benchmark on MySQL systems.

1. Server hardware: Cisco UCS C250; details in Appendix A.1
2. Storage: Virident FlashMAX M1400 MLC, driver 2.1(B5.33596)
3. Storage: Fusion-io ioDrive Duo 1.28TB MLC. Fusion-io driver version: 2.3.1 build 123; Firmware v5.0.6, rev 101583
4. Software: Percona Server 5.5.15; configuration in Appendix C
5. Client hardware: HP ProLiant DL380 G6; details in Appendix A.2

`tpcc-mysql` tests were run for following combinations:

- 5000 Warehouses (about 500GB datasize)
  - `innodb_buffer_pool_size`= [100GB]
- 2500 Warehouses (about 250GB datasize)
  - `innodb_buffer_pool_size`= [50GB]
  - `innodb_buffer_pool_size`= [150GB]

The purpose of the different combinations of data and memory sizes was to check how the data/memory ratio affected results.

We used a number of different user sessions and performed a 3600-sec run, gathering the throughput metric every 10 seconds.

That is, for each set of user sessions, we took 360 throughput measurements.

Based on these calculations, we constructed the following metric, which we used as a final result:

- Median Throughput for the last 1800 sec, to avoid a warm-up influence on the results.

2 Results

2.1 5000W

	Threads	Fusion-io ioDrive Duo	Virident FlashMAX	Ratio
1	1	486.00	550.00	1.13
2	4	1688.00	1709.00	1.01
3	16	3310.00	4662.00	1.41
4	24	3806.00	5000.00	1.31
5	64	4874.00	6057.00	1.24
6	256	3282.00	4724.00	1.44
7	1024	2170.00	2321.00	1.07

Table 1: Throughput, NOTP/10sec, more is better

2.1.1 5000W, Jitter graph

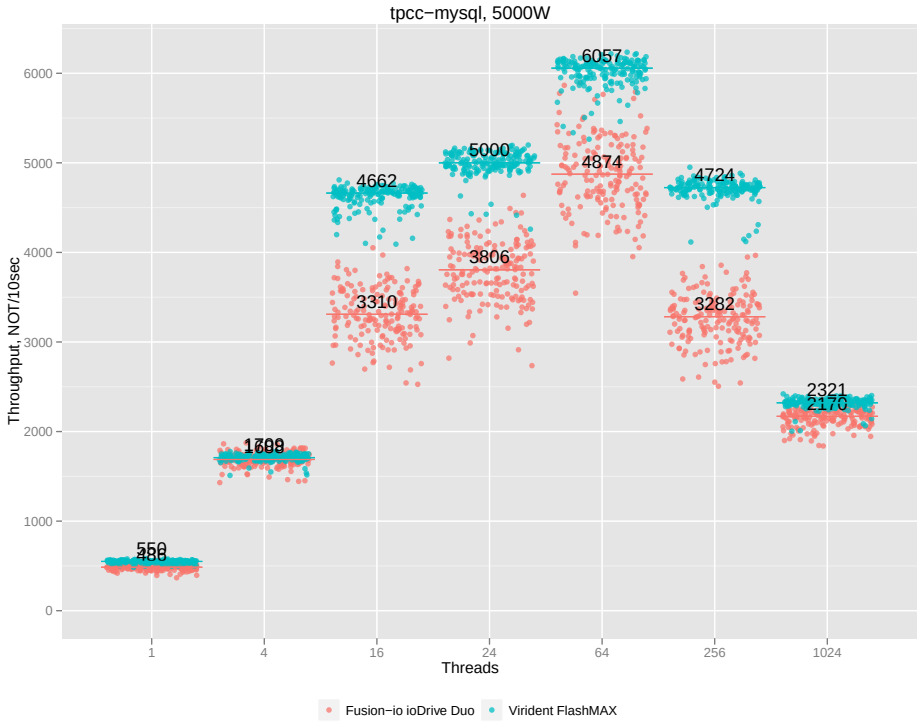

2.1.2 5000W, Stability graph

This graph compares the stability of throughput for both cards.

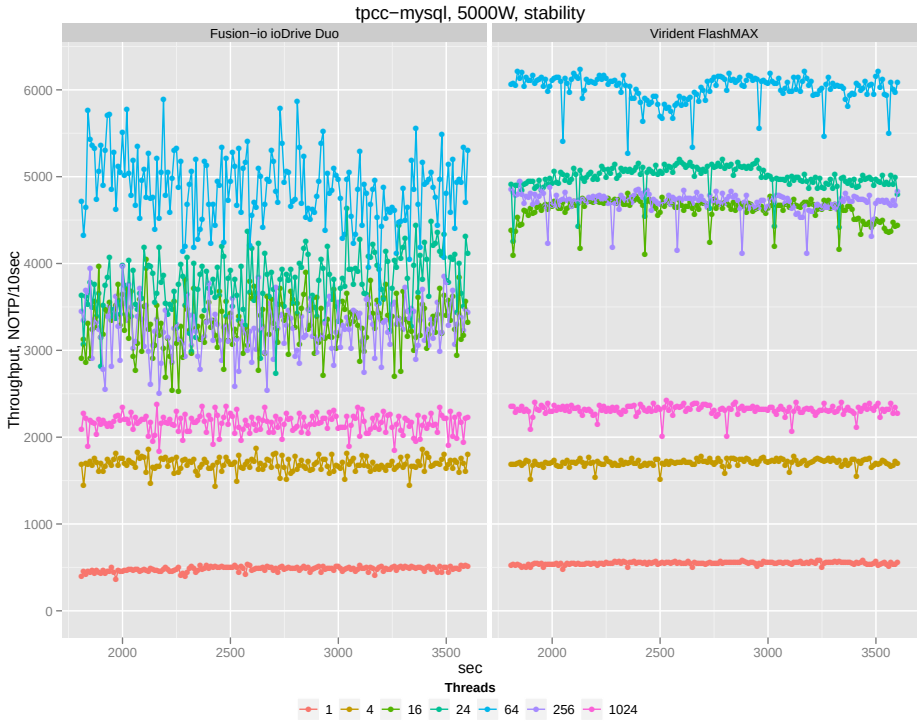

2.2 2500W

2.2.1 Buffer pool 50G

Threads	Fusion-io ioDrive Duo	Virident FlashMAX	Ratio
16	3398.00	4667.00	1.37
24	4294.00	5186.00	1.21
32	5382.00	5600.00	1.04
48	4760.00	6147.00	1.29
128	3222.00	5592.00	1.74

Table 2: Throughput, NOTP/10sec, more is better

2.2.2 Jitter graph

2.2.3 Stability graph

2.2.4 Buffer pool 150G

Threads	Fusion-io ioDrive Duo	Virident FlashMAX	Ratio
16	7490.00	8588.00	1.15
24	8508.00	10432.00	1.23
32	9911.00	12474.00	1.26
48	11321.00	14738.00	1.30
128	8171.00	10544.00	1.29

2.2.5 Jitter graph

2.2.6 Stability graph

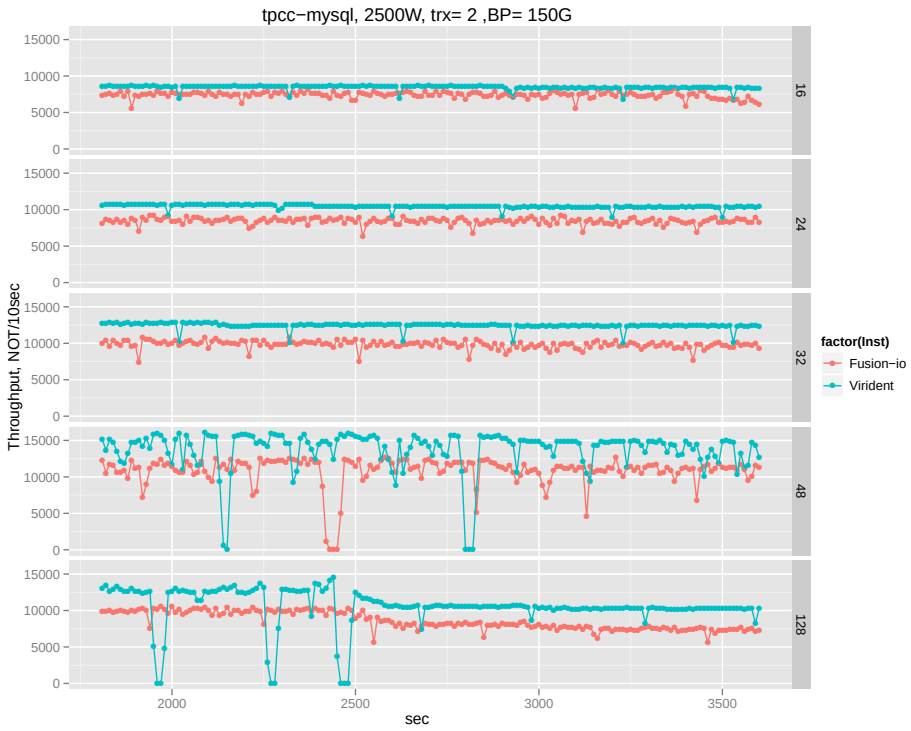

3 Conclusion

- Virident FlashMAX provides stability of performance and reveals a denser throughput.
- In addition to stability, in many cases there is also a better throughput in MySQL (up to 40%) using the Virident FlashMAX card.

A Systems Specifications

A.1 Server configuration

```
# Percona Toolkit System Summary Report ######################
        Date | 2011-11-14 23:46:29 UTC (local TZ: PST -0800)
      System | Cisco Systems Inc; R250-2480805; v0
    Platform | Linux
     Release | Red Hat Enterprise Linux Server release 6.1 (Santiago)
      Kernel | 2.6.32-131.0.15.el6.x86_64
Architecture | CPU = 64-bit, OS = 64-bit
   Threading | NPTL 2.12
     SELinux | Disabled
 Virtualized | No virtualization detected
# Processor ##################################################
  Processors | physical = 2, cores = 12, virtual = 24,
               hyperthreading = yes
      Speeds | 24x2926.185
      Models | 24xIntel(R) Xeon(R) CPU X5670 @ 2.93GHz
      Caches | 24x12288 KB
# Memory #####################################################
       Total | 346.92G
```

A.2 Client configuration

```
        Date | 2011-11-14 23:56:12 UTC (local TZ: PST -0800)
      System | HP; ProLiant DL380 G6; vNot Specified
    Platform | Linux
     Release | Red Hat Enterprise Linux Server release 6.1 (Santiago)
      Kernel | 2.6.32-131.0.15.el6.x86_64
Architecture | CPU = 64-bit, OS = 64-bit
   Threading | NPTL 2.12
    Compiler | GNU CC version 4.4.4 20100726 (Red Hat 4.4.4-13).
     SELinux | Disabled
 Virtualized | No virtualization detected
# Processor ##################################################
  Processors | physical = 2, cores = 8, virtual = 16,
     | hyperthreading = yes
      Speeds | 16x1600.000
      Models | 16xIntel(R) Xeon(R) CPU L5520 @ 2.27GHz
      Caches | 16x8192 KB
# Memory #####################################################
```

```
       Total | 70.80G
```

B Storage configuration

Fusion-io partition was created as:

```
yes | fio-format /dev/fct0
yes | fio-format /dev/fct1
fio-attach /dev/fct0
fio-attach /dev/fct1
mdadm create verbose /dev/md127 level=0 raid-devices=2 /dev/fioa /dev/fiob
mkfs.xfs -s size=4096 /dev/md127
mount /dev/md127 /mnt/fio -o nobarrier
```

Virident partition was created as:

```
fdparm prepare -f /dev/vgca0
mkfs.xfs -s size=4096 /dev/vgca0
mount /dev/vgca0 /mnt/tachion -o nobarrier
```

C MySQL configuration

```
[mysqld]
gdb

datadir=/mnt/tachion

#for SSD
innodb_flush_neighbor_pages = 0
innodb_adaptive_flushing_method  = keep_average
innodb_buffer_pool_restore_at_startup=300


#####fixed innodb options
innodb_file_per_table = true
innodb_data_file_path = ibdata1:10M:autoextend
innodb_flush_log_at_trx_commit = 1
innodb_flush_method = O_DIRECT
innodb_log_buffer_size = 256M

innodb_buffer_pool_size = 174G
```

```
innodb_log_file_size = 4G
innodb_log_files_in_group = 2


#####plugin options
innodb_read_io_threads = 16
innodb_write_io_threads = 16
innodb_io_capacity = 10000


#not innodb options (fixed)

port = 3306
back_log = 50
max_connections = 2000
max_prepared_stmt_count=500000
max_connect_errors = 10
table_cache = 2048
max_allowed_packet = 16M
binlog_cache_size = 16M
max_heap_table_size = 64M
sort_buffer_size = 4M
join_buffer_size = 4M
thread_cache_size = 1000
query_cache_size = 0
query_cache_type = 0
query_cache_limit = 2M
ft_min_word_len = 4
memlock
#default_table_type = InnoDB
thread_stack = 192K
tmp_table_size = 64M

server-id = 10

#*** MyISAM Specific options
key_buffer_size = 8M
read_buffer_size = 1M
read_rnd_buffer_size = 4M
bulk_insert_buffer_size = 8M
myisam_sort_buffer_size = 8M
```

```
myisam_max_sort_file_size = 10G
#myisam_max_extra_sort_file_size = 10G
myisam_repair_threads = 1
myisam_recover

socket=/var/lib/mysql/mysql.sock
user=root
skip-grant-tables

[mysql]
no-auto-rehash
socket=/var/lib/mysql/mysql.sock

[client]
socket=/var/lib/mysql/mysql.sock
```