

# THOMAS KRENN®

server.hosting.customized.



This work is licensed under a [Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-nc-sa/4.0/).

# TKperf Test Report

## Contents

|                                      |           |
|--------------------------------------|-----------|
| <b>1 Setup Information</b>           | <b>2</b>  |
| <b>2 General Information</b>         | <b>2</b>  |
| 2.1 Steady State                     | 2         |
| <b>3 IOPS</b>                        | <b>3</b>  |
| 3.1 Steady State Information         | 3         |
| 3.2 Measurement Plots                | 3         |
| 3.3 Measurement Window Summary Table | 5         |
| <b>4 Throughput</b>                  | <b>6</b>  |
| 4.1 Steady State Information         | 6         |
| 4.2 Measurement Plots                | 6         |
| 4.3 Measurement Window Summary Table | 8         |
| <b>5 Latency</b>                     | <b>8</b>  |
| 5.1 Steady State Information         | 9         |
| 5.2 Measurement Plots                | 9         |
| 5.3 Measurement Window Summary Table | 11        |
| <b>6 Write Saturation</b>            | <b>11</b> |
| 6.1 Measurement Plots                | 11        |

# 1 Setup Information

## Tested Device:

- Model Number: Samsung SSD 840 PRO Series
- Serial Number: S1ATNSAD744454L
- Firmware Revision: DXM05B0Q
- device size with M = 1000\*1000: 256060 MBytes (256 GB)
- write-caching = 1 (on)

## Used command line:

- `/usr/local/bin/tkperf ssd Samsung840PRO-256GB /dev/sda -nj 2 -iod 16 -rfb`

## Performance System:

- TKperf Version: 1.1
- Fio Version: fio-2.0.14
- Date of test run: 2013-10-01
- Number of jobs: 2
- Number of outstanding IOs (iodepth): 16

## Operating System:

- Kernel Version: 3.2.0-38-generic
- Description: Ubuntu 12.04.2 LTS

# 2 General Information

- *workloads*: The percentage of read operations in the random mixed workload. In the plots the term "100/00" means 100% read and 0% write, "95/5" 95% read and 5% write, and so on.
- *block sizes*: The block size of Fio to be used for IO operations.
- *measurement window*: Those rounds, where the dependence variable became stable.
- *dependence variable*: A specific type of test variable to determine the steady state.

## 2.1 Steady State

The Steady State is to determine if a test has reached a steady performance level. Each test has a different dependence variable to check if the state has already been reached. To check for the steady state the performance values of a test measurement window are taken (the last 5 rounds). The steady state is reached if:

- The maximum data excursion is less than 20% of the average in the measurement window.
- The slope of the linear best fit line is less than 10% of the average in the measurement window

If these two conditions are met the steady state has been reached for the specific dependence variable. Therefore the test can be stopped and the performance values of the measurement window can be taken for the measurement plots. If the steady state has not been reached after a maximum number of rounds the test can be stopped as well. The numbers for these two variables are:

- Measurement Window: 4
- Max. number of rounds: 25

## 3 IOPS

The IOPS test consists of looping over the following parameters:

```
Make Secure Erase
Workload Ind. Preconditioning
While not Steady State
  For workloads [100, 95, 65, 50, 35, 5, 0]
    For block sizes ['1024k', '128k', '64k', '32k', '16k', '8k', '4k', '512']
```

Each combination of workload and block size is carried out for 60 seconds using direct IO. The average number of read and write IOPS is measured and summed up, therefore 56 values are the result of the two loops. After these loops are finished one test round has been carried out. To detect the steady state the IOPS of 4k random write are taken.

- Dependent Variable: 4k block size, random write

### 3.1 Steady State Information

**Steady State has been reached:**

- True

**Steady State has been reached in rounds :**

- deque([0, 1, 2, 3, 4])


**Values in stdy measurement window:**

- deque([4747, 4199, 4410, 4647, 5064])


**Average in stdy measurement window:**

- 4613.4

### 3.2 Measurement Plots




The Steady State Convergence Plot shows the reached IOPS for all block sizes of random writes over all rounds.




The Steady State Verification Plot shows the measured IOPS of 4k random writes, the 20% average window and the slope of the linear best fit line in the measurement window.

**IOPS Measurement Plot**



The Measurement Plot shows the average of IOPS in the measurement window. For every workload the IOPS of all block sizes are plotted.

**IOPS 3D Measurement Plot**



The Measurement 3D Plot shows the average of IOPS in the measurement window. For every workload the IOPS of all block sizes are plotted.

### 3.3 Measurement Window Summary Table

### Average IOPS vs. Block Size and R/W Mix %

| Block Size<br>↓ | Wld. →<br>100/0 | 95/5    | 65/35   | 50/50   | 35/65  | 5/95   | 0/100  |
|-----------------|-----------------|---------|---------|---------|--------|--------|--------|
| 512             | 53012.4         | 26836.2 | 18039.0 | 10540.6 | 7759.2 | 5354.0 | 4405.0 |
| 4k              | 44527.4         | 24394.0 | 10889.2 | 10038.0 | 7386.8 | 5100.6 | 4613.4 |
| 8k              | 34992.2         | 14613.6 | 6777.4  | 5007.6  | 3607.2 | 2517.2 | 2220.8 |
| 16k             | 24941.2         | 7532.2  | 3401.6  | 2493.8  | 1890.4 | 1347.4 | 1172.4 |
| 32k             | 15287.8         | 9030.4  | 1839.2  | 1349.6  | 981.2  | 713.0  | 633.4  |
| 64k             | 7604.2          | 4449.2  | 899.2   | 608.2   | 469.0  | 332.4  | 279.4  |
| 128k            | 3849.2          | 2244.8  | 440.6   | 291.2   | 221.8  | 155.6  | 146.8  |
| 1024k           | 443.4           | 360.2   | 78.4    | 44.2    | 28.0   | 19.0   | 18.4   |

## 4 Throughput

The throughput test consists of looping over the following parameters:

```
For block sizes ['1024k', '64k', '8k', '4k', '512']
  Make Secure Erase
  While not Steady State
    Sequential read
    Sequential write
```

For each block size sequential read and write is carried out for 60 seconds using direct IO. The number of kilobytes for read and write is measured, therefore 2 values are the result of one round. To detect the steady state the throughput of 1024k sequential write is taken.

- Dependent Variable: 1024k block size, sequential write

### 4.1 Steady State Information

**Steady State has been reached:**

- True

**Steady State has been reached in rounds :**

- deque([0, 1, 2, 3, 4])


**Values in stdy measurement window:**

- deque([501736, 503393, 518103, 518248, 518248])


**Average in stdy measurement window:**

- 511945.6

### 4.2 Measurement Plots




The Read/Write Steady State Convergence Plot shows the bandwidth for all block sizes of seq. reads over all rounds. On the top the write throughput is plotted, below the throughput for read.



The Steady State Verification Plot shows the bandwidth of 1024k seq. writes, the 20% average window and the slope of the linear best fit line in the measurement window.





The Measurement Plot shows the average bandwidth of reads and writes in the measurement window. For all block sizes the seq. read and write bandwidth is plotted.

### 4.3 Measurement Window Summary Table

Average MB/s vs. Block Size and R/W

| Block Size ↓ | Read    | Write   |
|--------------|---------|---------|
| 1024k        | 532.078 | 499.947 |
| 64k          | 482.567 | 452.868 |
| 8k           | 461.128 | 348.757 |
| 4k           | 387.883 | 336.656 |
| 512          | 25.19   | 52.277  |

## 5 Latency

The latency test consists of looping over the following parameters:

```

Make Secure Erase
Workload Ind. Preconditioning
While not Steady State
    For workloads [100, 65, 0]
        For block sizes ['8k', '4k', '512']

```

For all block sizes random read, a 65/35 read/write mixed workload and random write is carried out for 60 seconds using direct IO. For every combination the Min, Max and Mean Latency is measured. After these loops are finished one test round has been carried out. To detect the steady state the mean latency of 4k

random write is taken.

- Dependent Variable: 4k block size, random write mean latency

## 5.1 Steady State Information

Steady State has been reached:

- True

Steady State has been reached in rounds :

- deque([1, 2, 3, 4, 5])

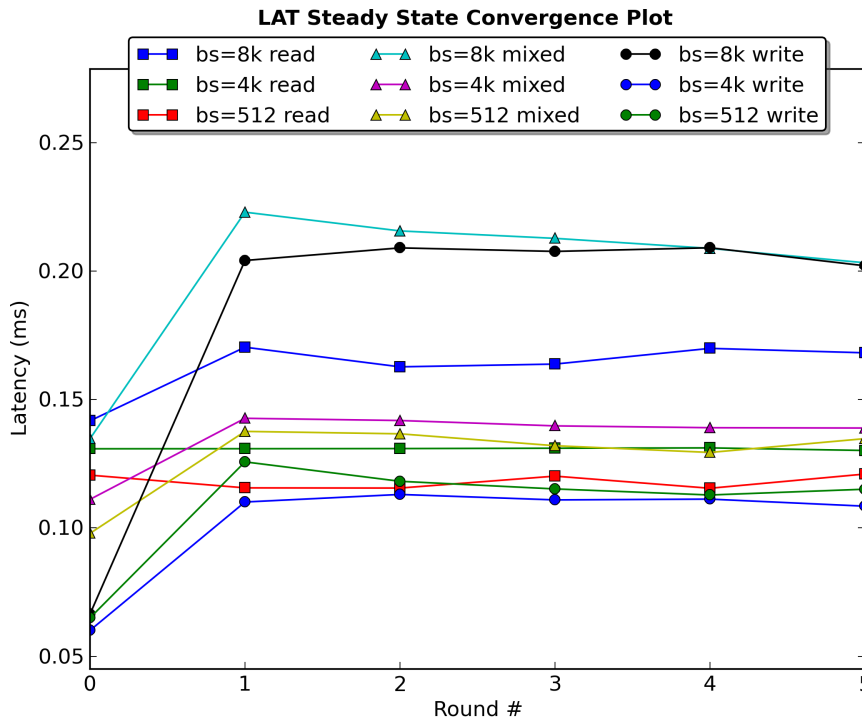
Values in stdy measurement window:

- deque([110.090464, 113.033035, 110.908359, 111.209767, 108.465582])


Average in stdy measurement window:

- 110.7414414

## 5.2 Measurement Plots

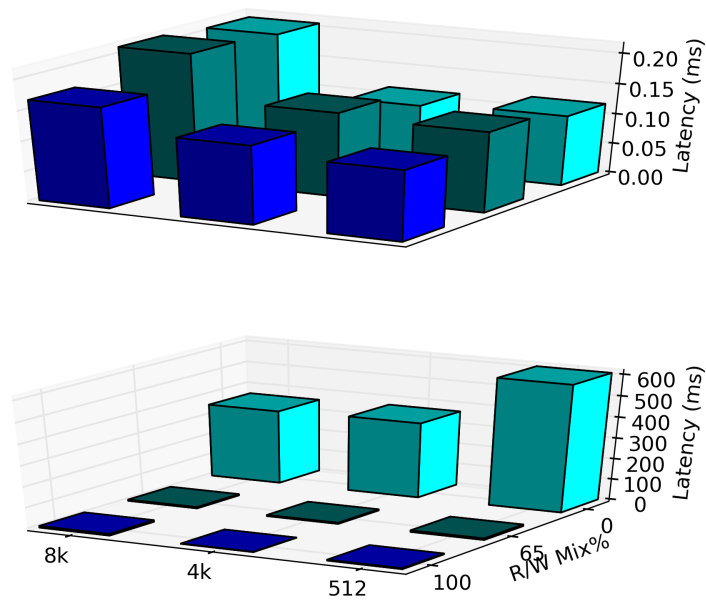


*The Steady State Convergence Plot shows the mean latency for all block sizes of random read, mixed workload and write.*



The Steady State Verification Plot shows the mean latency of 4k random writes, the 20% average window and the slope of the linear best fit line in the measurement window.

**LAT 3D Measurement Plot**



The Latency Measurement 3D Plot shows the average latency on top and the max latency below it. For the measurement window every workload including all block sizes is plotted.

## 5.3 Measurement Window Summary Table

*Average Latency (ms) vs. Block Size and R/W Mix %*

| Block Size ↓ | Wld. → 0/100 | 65/35 | 100/0 |
|--------------|--------------|-------|-------|
| 8k           | 0.206        | 0.213 | 0.167 |
| 4k           | 0.111        | 0.14  | 0.131 |
| 512          | 0.117        | 0.134 | 0.118 |

*Max Latency (ms) vs. Block Size and R/W Mix %*

| Block Size ↓ | Wld. → 0/100 | 65/35 | 100/0 |
|--------------|--------------|-------|-------|
| 8k           | 354.211      | 6.054 | 8.037 |
| 4k           | 362.909      | 5.484 | 1.115 |
| 512          | 607.793      | 7.47  | 4.315 |

## 6 Write Saturation

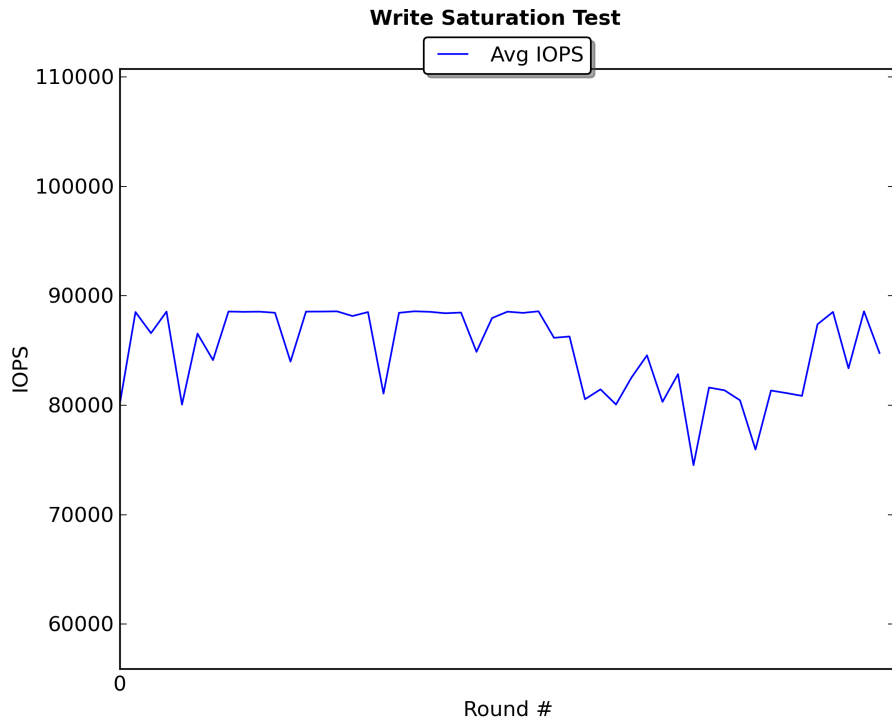
The write saturation test consists of looping over the following parameters:

```
Make Secure Erase
While not written 4x User Capacity or 24h
    Carry out random write, 4k block size for 1 minute.
```

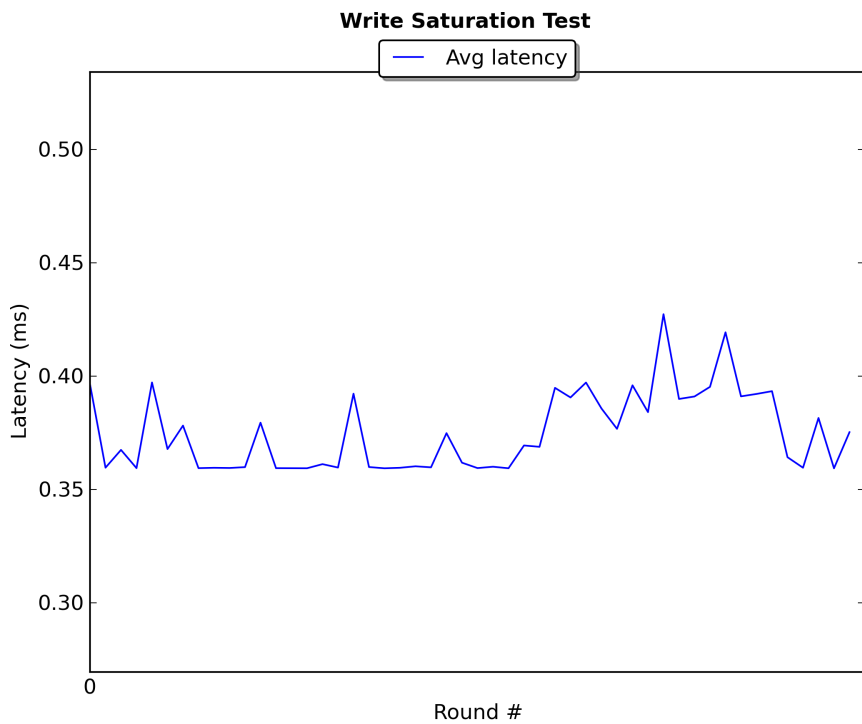
For 4k block size random write is carried out for 60 seconds using direct IO. For each round (60 second window) the write IOPS and latencies are measured. Also the total written IO is measured to check if 4x capacity has been written.

As no steady state detection is necessary there is no dependence variable.

### 6.1 Measurement Plots



The Write Saturation IOPS Plot shows the average IOPS of 4k random writes over all rounds.



The Write Saturation Latency Plot shows the mean latency of 4k random writes over all rounds.