

TKperf Test Report

Contents

1 Setup Information	2
2 General Information	2
2.1 Steady State	2
3 IOPS	2
3.1 Steady State Information	3
3.2 Measurement Plots	3
3.3 Measurement Window Summary Table	5
4 Throughput	5
4.1 Steady State Information	6
4.2 Measurement Plots	6
4.3 Measurement Window Summary Table	8
5 Latency	8
5.1 Steady State Information	8
5.2 Measurement Plots	8
5.3 Measurement Window Summary Table	10
6 Write Saturation	10
6.1 Measurement Plots	11

1 Setup Information

Tested Device:

- Intel SSD DC P4800X 375GB

Used command line:

- `/usr/bin/tkperf ssd IntelOptaneP4800X /dev/nvme0n1p4 -nj 2 -iod 16 -rfb -i nvme -dsc IntelOptaneP4800X`

Performance System:

- TKperf Version: 2.1
- Fio Version: fio-2.2.10
- Date of test run: 2017-09-19
- Number of jobs: 2
- Number of outstanding IOs (iodepth): 16

Operating System:

- Kernel Version: 4.4.0-93-generic
- Description: Ubuntu 16.04 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', '512k']
```

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])

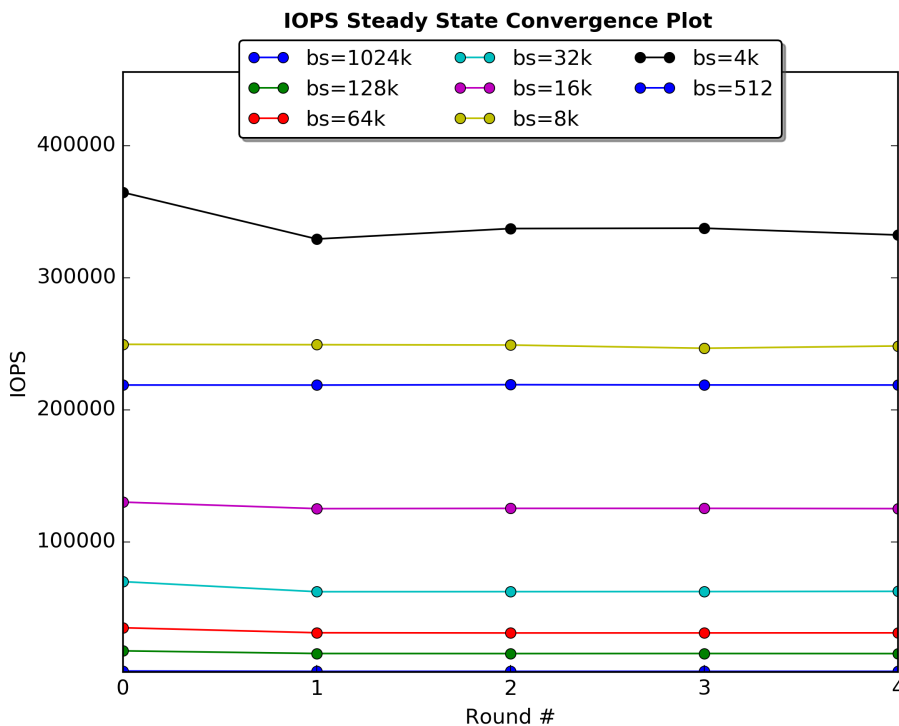
Values in stdy measurement window:

- deque([364642, 329344, 337257, 337525, 332385])

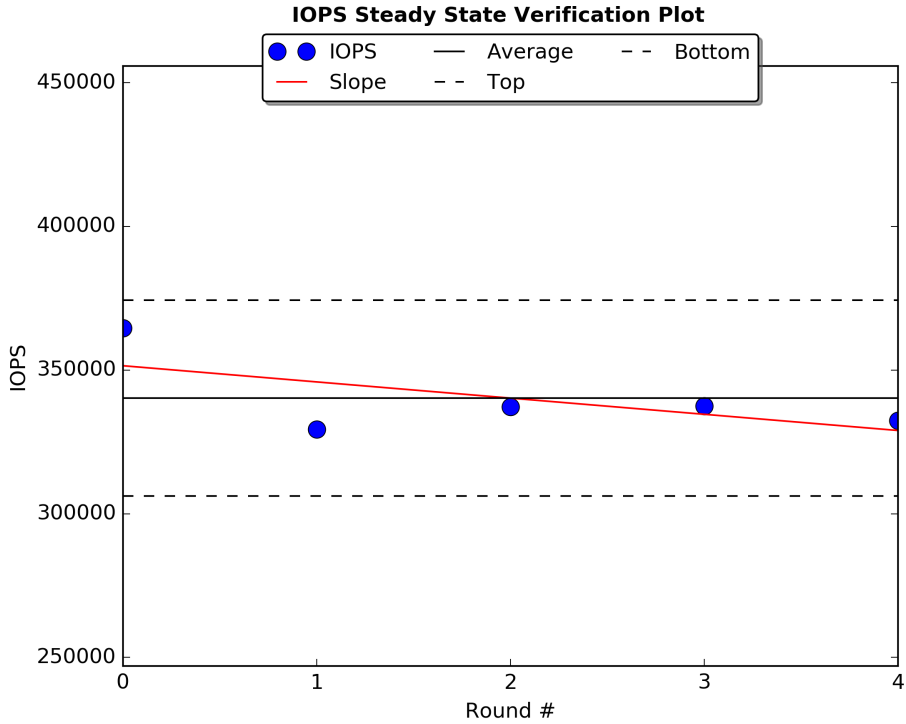
Average in stdy measurement window:

- 340230


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.



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 ↓	Wld. → 100/0	95/5	65/35	50/50	35/65	5/95	0/100
512	373494.6	403730.2	346639.2	305543.2	273896.0	228751.6	218819.8
4k	378150.0	386801.4	354212.2	355044.4	343331.8	336686.8	340230.6
8k	272830.2	265424.2	233625.2	228562.8	227707.8	247247.8	248575.2
16k	141172.4	138260.0	122912.8	120209.4	116607.4	125643.6	126228.4
32k	73952.8	72229.2	64015.0	62554.4	61060.6	62942.8	63859.2
64k	36375.2	35397.2	29033.2	28231.2	27874.4	31214.0	31848.8
128k	13825.0	16765.0	13613.2	13219.0	12361.8	15185.2	15831.8
1024k	2063.6	2092.4	1916.8	1881.0	1870.6	1944.4	1978.0

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])

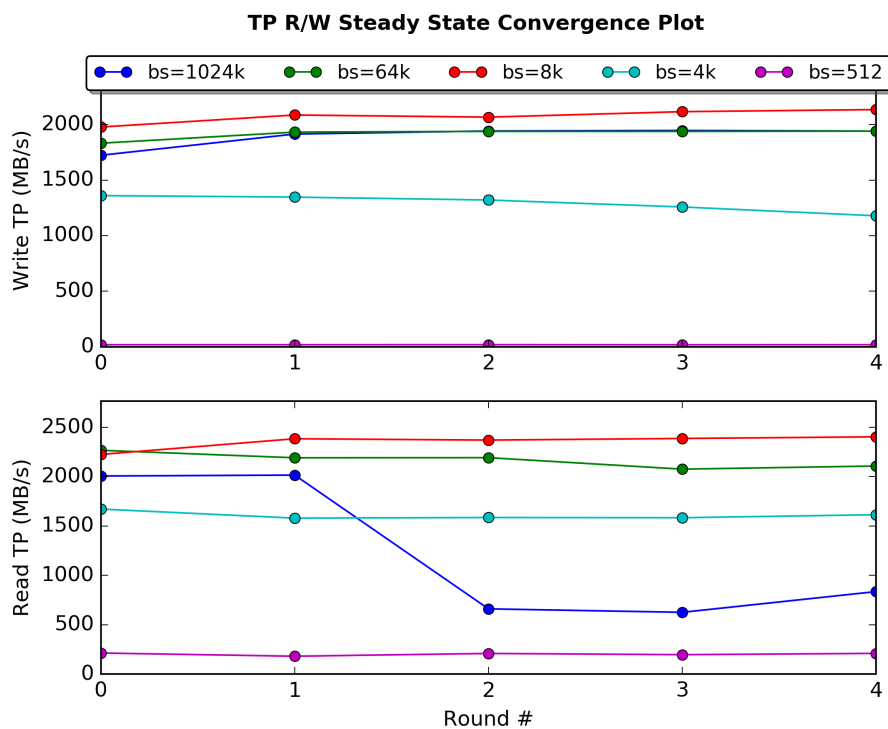
Values in stdy measurement window:

- deque([1765187, 1960440, 1988762, 1993421, 1987805])


Average in stdy measurement window:

- 1939123


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	1228.732	1893.675
64k	2167.191	1916.419
8k	2354.35	2076.307
4k	1607.256	1293.009
512	201.24	18.21

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([13, 14, 15, 16, 17])

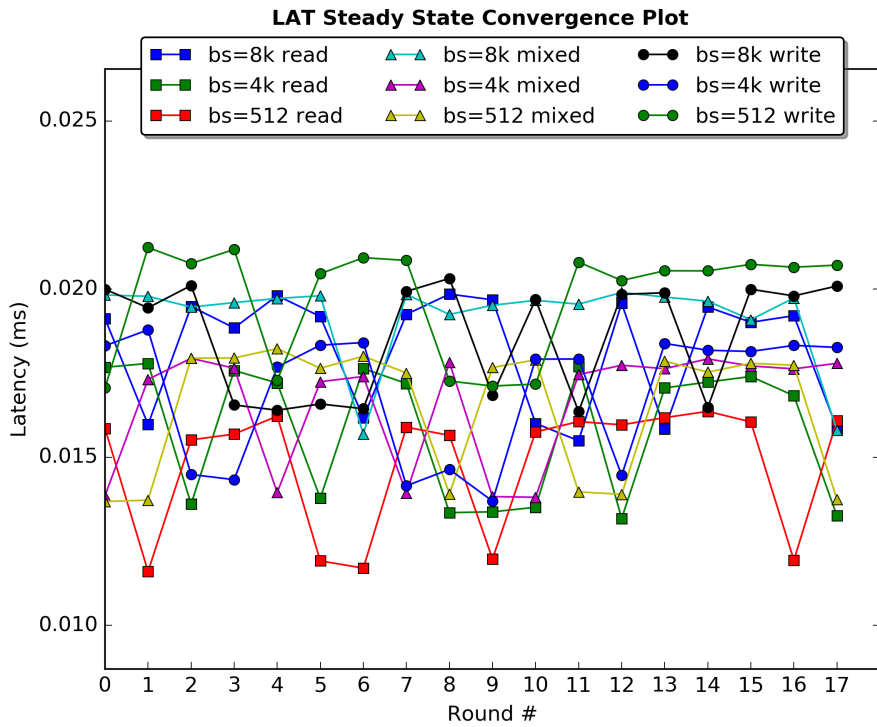
Values in stdy measurement window:

- deque([18.38568, 18.183962, 18.148539, 18.328561, 18.268658])


Average in stdy measurement window:

- 18.26308

5.2 Measurement Plots

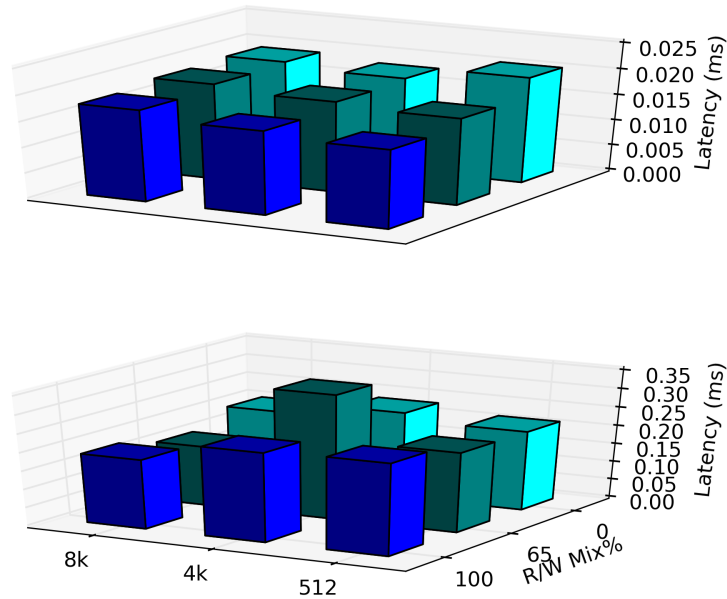


The Steady State Convergence Plot shows the mean latency for all block sizes of random read, mixed workload and write.



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	0.019	0.019	0.018
4k	0.018	0.018	0.016
512	0.021	0.017	0.015

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

Block Size ↓	Wld. → 0/100	65/35	100/0
8k	0.206	0.164	0.19
4k	0.236	0.338	0.243
512	0.217	0.218	0.25

6 Write Saturation

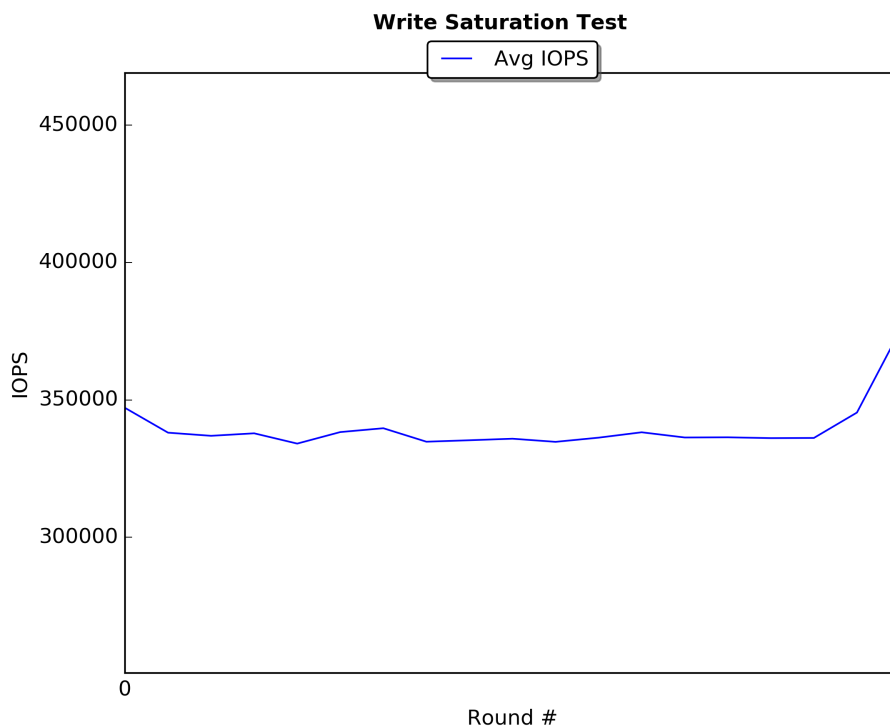
The write saturation test consists of looping over the following parameters:

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.

