Comparison of PCIe SLC Flash cards – Virident FlashMAX and Texas Memory Systems RamSan-70

Hussein N. El-Harake and Thomas Schoenemeyer
Swiss National Supercomputing Centre (CSCS), Manno, Switzerland
hussein@cscs.ch, schoenemeyer@cscs.ch

Abstract – We evaluated two PCIe attached SSD devices manufactured by Texas Memory Systems and Virident in terms of I/O bandwidth and IOPs. One of them is the "RamSan-70 Gorilla" card designed by Texas Memory Systems. The other device is the Virident FlashMax, formerly called TachIOn. Both devices are based on SLC technology, provide similar usable storage capacity and are connected via PCIe Gen2 x8 to a dual-socket node. Fio was used as benchmark tool. The performance characteristics were collected with various block size and different number of threads using the cards as raw devices and with xfs file system. Also degradation tests with filled cards have been carried out.

1 Introduction

In our papers from June 2011 [1] and September 2011 [2] we presented performance numbers for several PCIe SLC Flash devices built by TMS, FusionIO and Virident.

We repeated the fio benchmarks using two new devices, the Virident FlashMAX with 800 GB of storage capacity and the Texas Memory System RamSan-70 card with 900GB of capacity.

The devices were installed in a dual-socket server manufactured by Supermicro. It deploys two Intel Xeon X5690 processors with 6 cores running at 3.46 GHz with 48GB main memory. The server offers four x16 PCIe Gen2 slots. We used SLES 11 SP1.

Two enterprise class PCIe flash storage devices were evaluated: the first storage device was the RamSan-70 from Texas Memory Systems (TMS) offering 900 GB of usable storage space. The datasheet [3] reports a maximum performance of 2.5 GB/s read bandwidth and up to 600K IOPs at 4 KB blocksize.

	Reads		Writes	
	512 B	4 KB	512 B	4 KB
Bandwidth *MB/s*	600	2,500 1,250	125	1,800 900
IOPS	1.2M	600K 300K	250K	440K 220K

900 GB model provides maximum listed performance.

Table 1: TMS RamSan-70 datasheet

The second storage device is the Virident FlashMax with 800 GB [4] usable capacity. From the datasheet provided by Virident (table 2) a user can expect a read performance of up to 1.4 GB/s and 340K IOPS at 4KB blocksize.

Both devices use SLC technology and are connected via PCI-Express Gen2 8x Interface.

FlashMAXSCM Specifications		
	MLC	SLC
Capacity (GB)	1000, 1400	300, 400, 800
Form Factor	Low Profile	Low Profile
Read Performance	1.3 GB/s (4 KB) 325 Thousand IOPS (4KB) 1 Million IOPS (512 byte)	1.4 GB/s (4 KB) 340 Thousand IOPS (4KB) 1.4 Million IOPS (512 byte)
Write Performance	600 MB/s	1100 MB/s
Sustained Mixed Performance (75/25 r/w) at full capacity	850 MB/s (4KB) 220 Thousand IOPS (4KB)	920 MB/s 235 Thousand IOPS (4KB)
Write Access Latency	19µs	16µs
Read Access Latency	62µs	47µs

Table 2: Virident FlashMAX datasheet

2 Methods

All experiments described below were performed using the cards as raw devices and with xfs file system. We also used ext3, with similar results. Therefore we present only results based on xfs.

We used fio 1.58 [5] for our tests and used the latest drivers provided by the vendors in December 2011 (Virident 2.1.1 and Texas Memory Systems 3.3.1).

3 Results

3.1 Bandwidth measurements

We first discuss the random write bandwidth measurements on raw and xfs for both devices (figure 1, 2, 3 and 4).

The write performance to a raw device achieves a maximum of more than 1.6 GB/s for the RamSan device (figure 1) and 1.1 GB/s for the Virident FlashMAX (figure 2) and is nearly independent of block size and the number of threads.

When using the xfs file system the write bandwidth for the TMS card decreases by up to 20% for less than 2 or more than 128 threads, in all other cases the numbers are very close to the raw performance (figure 3).

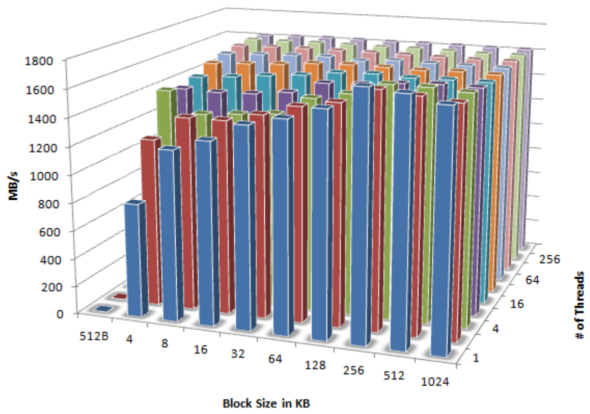

Figure 1: RamSan-70 Throughput Random Write (raw)

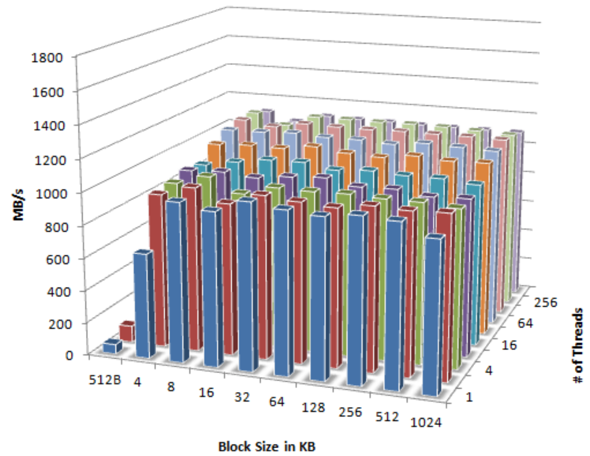

Figure 2: Virident FlashMAX Throughput Random Write (raw)

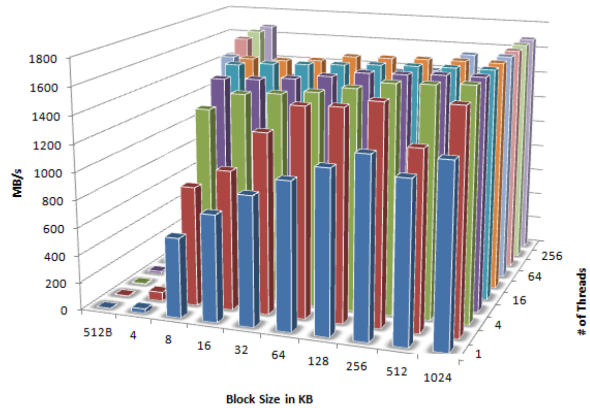

Figure 3: RamSan-70 Throughput Random Write (XFS)

The difference between raw and XFS write performance for the Virident device is more significant (figure 4) and appears for all block size numbers less or equal than 64KB. In this area the write performance drops down from 1.1 GB/s to 0.8 GB/s. For larger block size numbers the performance is similar to the raw performance.

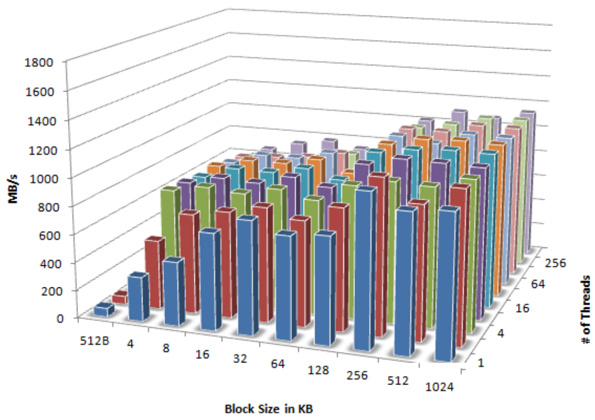

Figure 4: Virident FlashMAX Throughput Random Write (XFS)

The random read bandwidth results for raw and xfs are shown in figures 5, 6, 7 and 8. Again the TMS RamSan achieves a higher peak random read bandwidth of nearly 2.2 GB/s to the raw device (figure 5), however this performance is only provided with large blocks if less than 16 threads are used. The Virident device achieves a peak read bandwidth of roughly 1.3 GB/s (figure 6) and compared to the TMS card, much less threads are needed to achieve the peak values.

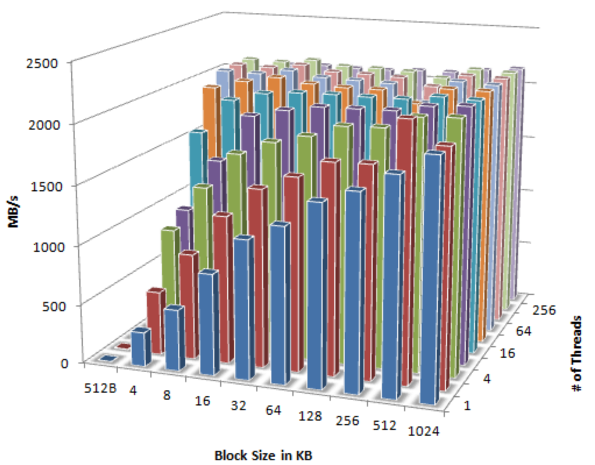

Figure 5: RamSan-70 Throughput Random Read (raw)

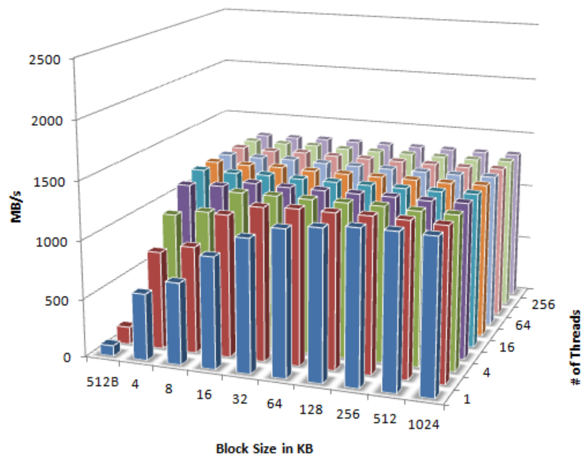

Figure 6: Virident FlashMAX Throughput Random Read (raw)

When using xfs, the read performance characteristics for both devices are not qualitatively different. Interestingly, xfs enables the TMS device to perform better especially for large blocks and high thread counts (figure 7). The peak read bandwidth achieves under these conditions up to 2.35GB/s. However the performance drops significantly for 4KB blocks (by 75% using 8 threads).

The Virident device delivers almost similar results when using XFS compared to raw (figure 8). Again this device provides a consistent performance profile over the whole range.

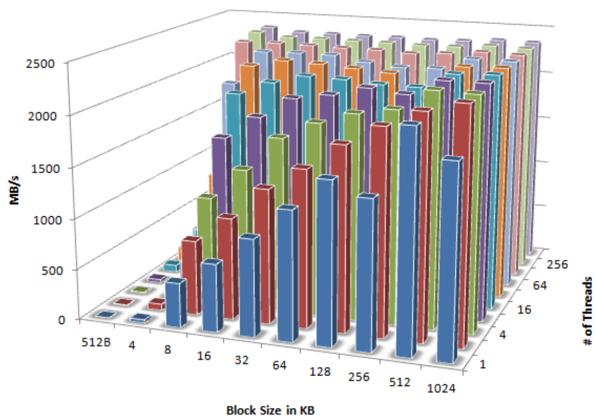

Figure 7: RamSan-70 Throughput Random Read (XFS)

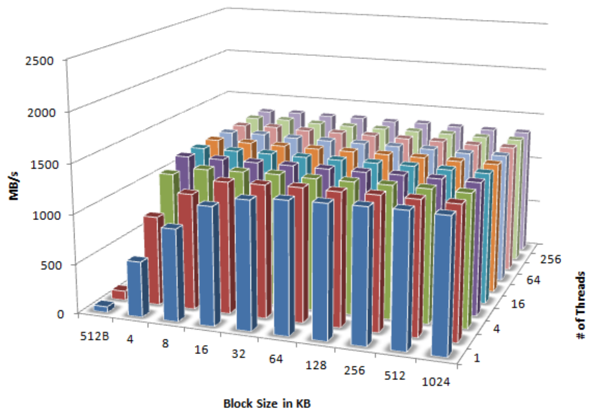

Figure 8: Virident FlashMAX Throughput Random Read (xfs)

3.2 IOPs measurements

Figures 9, 10, 11 and 12 summarize the results of the IOPS random write measurements for both flash devices using raw and xfs. From table 1 we can expect up to 440K IOPS using 4KB blocks for the TMS device. We measured peak values between 384K and 413K IOPS for more than 4 threads and 4KB blocks when writing to the raw device (figure 9).

The Virident behaves quite differently. For 4KB blocks, not more than 292K IOPS were achieved, however we observed surprisingly high numbers for 512B blocks up to 251K IOPS (figure 10).

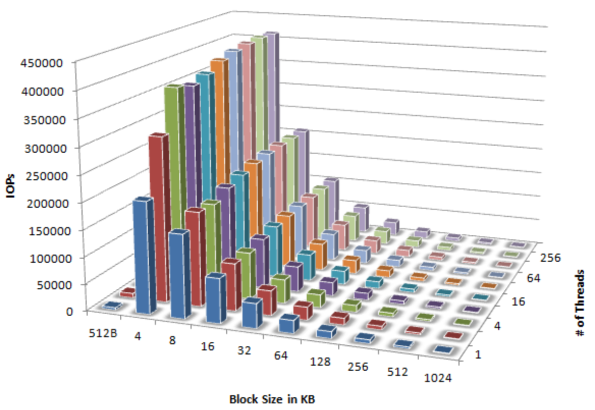

Figure 9: RamSan-70 IOPS random write (raw)

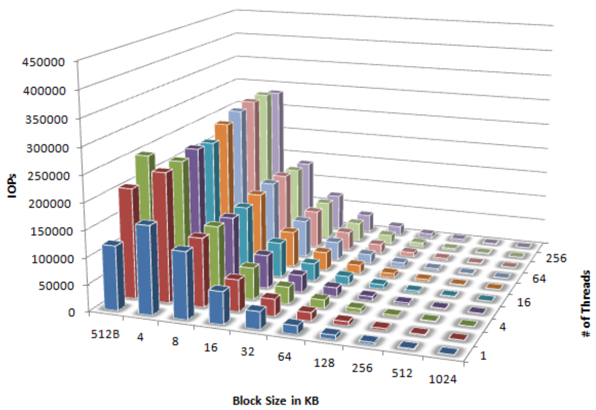

Figure 10: Virident FlashMAX IOPS random write (raw)

When using xfs the performance characteristics change significantly especially for the TMS device. At 4KB block size, the TMS device needs at least 64 threads to reach the 400K IOPs (figure 11).

The Virident device also performs much less using xfs, we observed at 4KB blocks a peak value of 203K IOPS, the device is saturated with 4 threads only (figure 12).

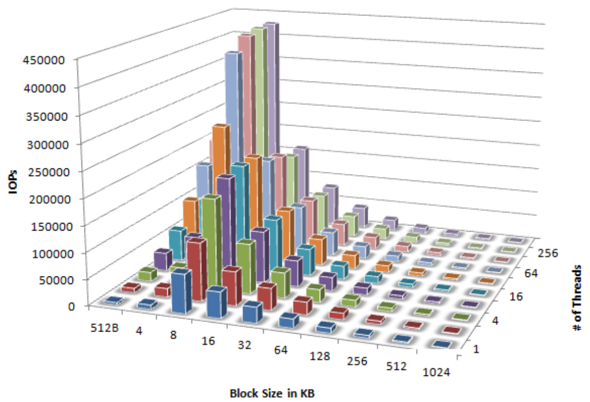

Figure 11: RamSan-70 IOPS Random Write (xfs)

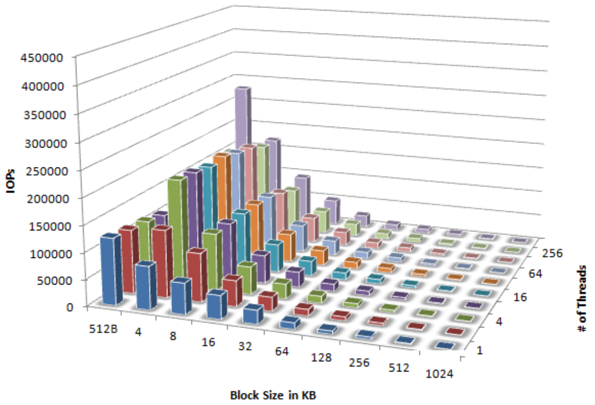

Figure 12: Virident FlashMAX IOPS random write (xfs)

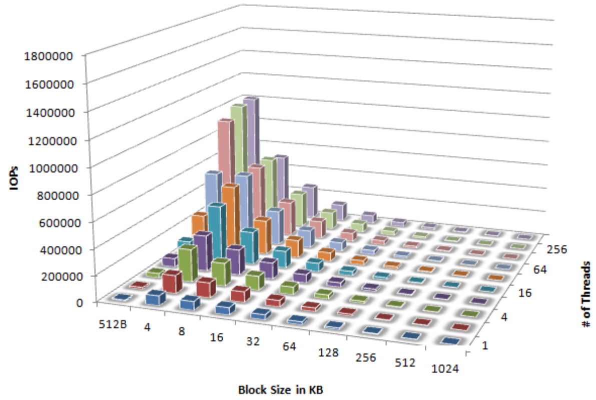

Figure 13: RamSan-70 IOPS Random Read (raw)

Figures 13, 14, 15 and 16 show the IOPS read performance numbers on raw and xfs. These results are very different from the write performance numbers. First of all, the peak values are achieved for both devices at 512B blocks. Secondly, the performance does not change regardless of using xfs or raw and thirdly, the Virident clearly outperformed the TMS device for 512B block size by a factor of nearly 1.7 (figure 14). A value of 1.68 MIOPS on a single SSD device has never been measured before at CSCS, also we could not find any hint in the literature about a similar measurement.

In our environment at CSCS the results achieved at 4K block size are more relevant, here the Virident card delivers between 146K and 339K, the TMS card provides between 74K and 506K (figures 13 and 15). All in all, the TMS device performs much better at 4K blocks, especially when more than 4 threads are used as shown in figure 17.

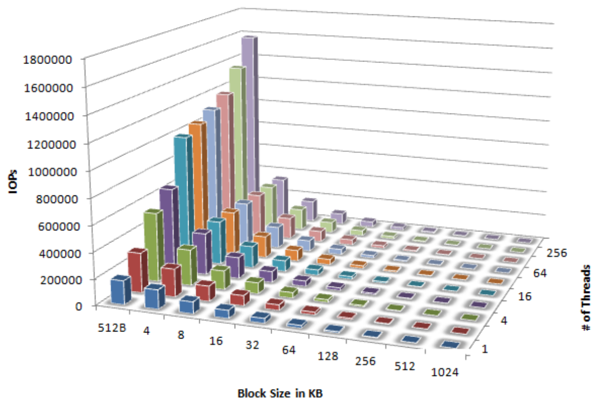

Figure 14: Virident FlashMAX IOPS Random Read (raw)

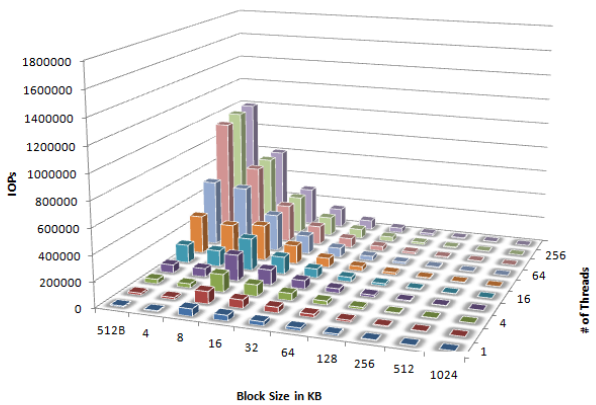

Figure 15: RamSan-70 IOPS Random Read (XFS)

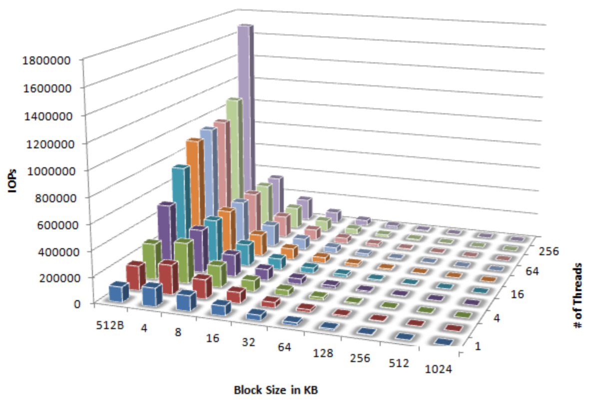

Figure 16: Viridant FlashMAX IOPS Random Read (XFS)

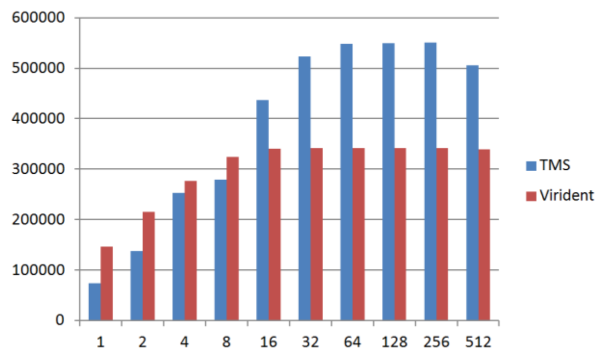

Figure 17: IOPS for random read on raw for TMS and Virident at 4KB block size and various number of threads.

3.3 Bandwidth degradation tests

All tests described above were carried out with not more than 70% of the device capacity. Our interest was to analyze the performance impact caused by filling an SSD close to its full capacity.

We filled both devices to 95% of the available net capacity and repeated the read and write throughput tests using xfs as file system.

For the experiment we used a block size of 1024KB and varied only the number of threads from one to 512.

Our results are shown in figures 18, 19, 20 and 21. All in all, the performance degradation is much less than what is reported by other studies, where the measured bandwidth can drop down to 15% of the peak values. Obviously both vendors found an efficient way to achieve a high steady state performance under these conditions.

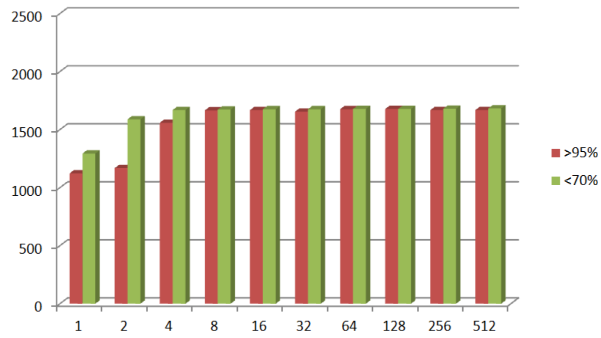

Figure 18: RamSan-70 Random Write Throughput (xfs)

As shown in figures 18 and 19, the write and read throughput performance of the RamSan-70 device is not affected, even at 95% fullness except for very low thread counts.

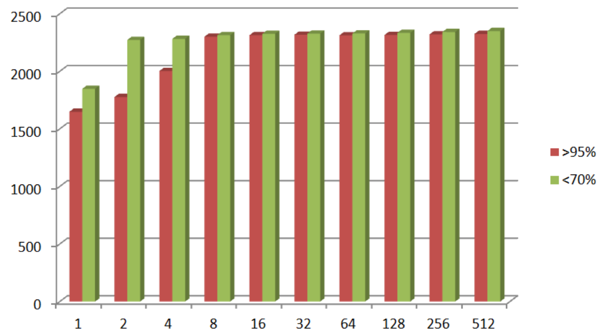

Figure 19: RamSan-70 Random Read Throughput (xfs)

The write performance of the Virident device is slightly lower under these conditions, especially at large thread counts, where we observed a bandwidth degradation to up to 30% (figure 20). The read performance degradation is similar to the RamSan-70 device (figure 21).

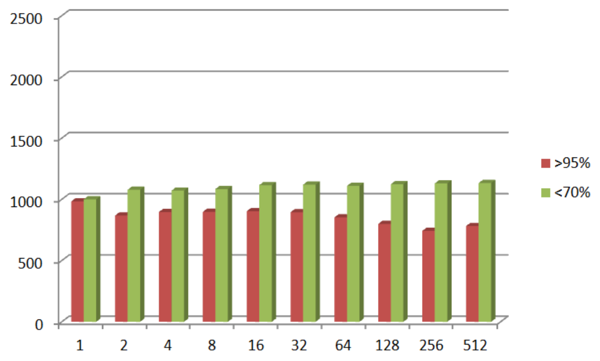

Figure 20: Viridint FlashMAX Random Write Throughput (xfs)

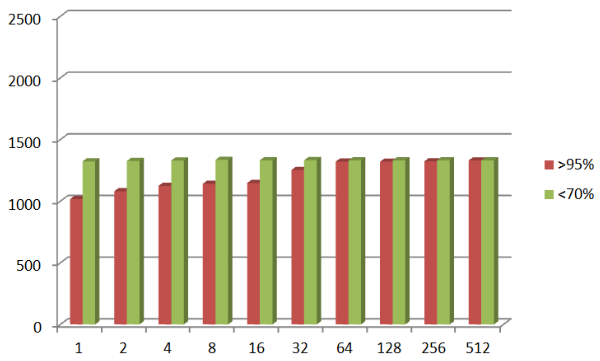

Figure 21: Virident FlashMAX Random Read Throughput (xfs)

3.4 Device dimensions

Apart from the performance measurements, it should be mentioned, that the FlashMAX is a low-profile card while the RamSan-70 is a full-height card. Therefore the Virident device having a low-profile card allows using it in high-density systems that cannot accommodate high-profile cards. For systems with vertical low-profile PCIe slots as the SuperMicro X8DTH-iF board, the Virident FlashMAX can deliver higher performance density and higher capacity density at the system level. In many real world application environments and server configurations they can be important metrics instead of just the performance or capacity per card.

4 Conclusion

We evaluated two enterprise class SSD devices based on SLC NAND technology using a PCIe connection. As test method we used fio 1.58 using the cards as raw devices and also with xfs file system.

- Our observations are comparable with the specifications provided by vendor datasheets. However Virident provides fewer details. Tables 3 and 4 give a brief overview for the overall comparison. For that we picked the best numbers we measured regardless of the number of threads or raw/xfs. In a few cases our observations peaked the datasheet specs.
- The peak numbers for read bandwidth performance of the TMS RamSan-70 are up to 77% higher, the write bandwidth numbers are up to 56% higher than for the Virident FlashMAX device as shown in figures 22 and 23.
- The highest number of IOPS was measured on the Virident device with nearly 1.7 M IOPS using 512B blocks.
- Using xfs versus raw devices can affect read and write performance, especially for small blocks and also for small number of threads.
- Using 4KB blocks, the TMS RamSan outperforms the Virident device in IOPS performance.
- Both devices demonstrate good results even close to 95% of used capacity for 1M block size. Tests were finalized end of January 2012. After finalizing this report, new drivers were provided, (Virident 2.1.2 and Texas Memory Systems 3.4.0.X). We will rerun all tests and will publish an updated report in late March 2012.
- For large thread counts the performance results are directly proportional to the processor frequency.

- Both devices showed linear results in terms of scalability.

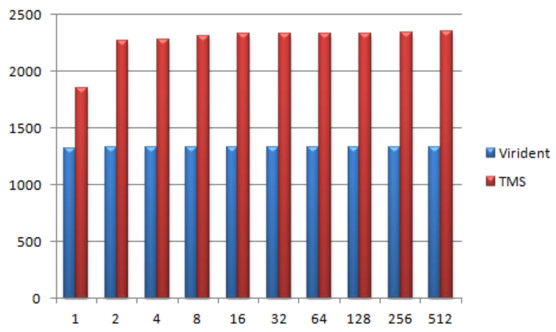

Figure 22: Comparison of random read performance (MB/s) for 1M versus number of threads using xfs.

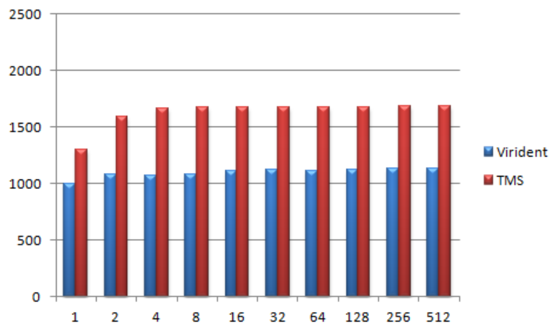

Figure 23: Comparison of random write performance (MB/s) for 1M versus number of threads using xfs.

TMS	Reads		Writes	
RamSan-70	512B	4KB	512B	4KB
Bandwidth Datasheet	600	2500	125	1800
Max. Bandwidth measured	490	2360	106	1690
IOPS Datasheet	1.2 M	600K	250K	440K
Max. IOPS measured	1.0 M	605K	225K	428K

Table 3: Device specs according to the datasheet compared to observations, TMS RamSan-70

Virident	Reads		Writes	
FlashMAX	512B	4KB	512B	4KB
Bandwidth Datasheet	NA	1400	NA	1100
Max. Bandwidth measured	824	1346	136	1143
IOPS Datasheet	1.4 M	340K	NA	NA
Max. IOPS measured	1.7M	344K	252K	293K

Table 4: Device specs according to the datasheet compared to observations, Virident FlashMAX

5 Literature

[1] Hussein N. Harake and Thomas Schoenemeyer: Detailed Analysis of Solid State Disks, Technical Paper, CSCS, July 2011.

[2] Hussein N. Harake and Thomas Schoenemeyer: Comparison of PCIe SLC Flash cards, Technical Paper, CSCS, September 2011.

[3] RamSan-70, Datasheet, Texas Memory Systems, 2011.

[4] Virident FlashMax PCIe Storage Class Memory, Datasheet, Virident, 2011.

[5] Fio I/O tool for benchmark and hardware verification, Source code, 2011