Developing Extremely Low-Latency NVMe SSDs

Young Paik

Director of Product Planning

Samsung Electronics

Flash Memory Summit

Disclaimer

- This presentation and/or accompanying oral statements by Samsung representatives collectively, the “Presentation”) is intended to provide information concerning the SSD and memory industry and Samsung Electronics Co., Ltd. and certain affiliates (collectively, “Samsung”). While Samsung strives to provide information that is accurate and up-to-date, this Presentation may nonetheless contain inaccuracies or omissions. As a consequence, Samsung does not in any way guarantee the accuracy or completeness of the information provided in this Presentation.
- This Presentation may include forward-looking statements, including, but not limited to, statements about any matter that is not a historical fact; statements regarding Samsung’s intentions, beliefs or current expectations concerning, among other things, market prospects, technological developments, growth, strategies, and the industry in which Samsung operates; and statements regarding products or features that are still in development. By their nature, forward-looking statements involve risks and uncertainties, because they relate to events and depend on circumstances that may or may not occur in the future. Samsung cautions you that forward looking statements are not guarantees of future performance and that the actual developments of Samsung, the market, or industry in which Samsung operates may differ materially from those made or suggested by the forward-looking statements in this Presentation. In addition, even if such forward-looking statements are shown to be accurate, those developments may not be indicative of developments in future periods.

Trend Is Toward Convergence

After years of development on NAND, products are beginning converge on increasingly similar numbers.

Are We At The Limit of Flash?

- Samsung looked at a redesign of NAND
  - Goal was to reduce media latency
  - Result is Z-NAND → Z-SSD

- What happens as media latency decreases?
  - Application performance/behavior
  - Limits

Benchmarks Tests Will Tell

Impact of lower latency?

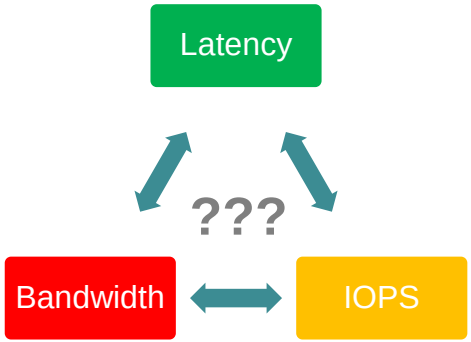

Use Case	Benchmark Tool	Description
Caching	Twemperf	Read/write tests with SSDs as a cache.
NoSQL Database	DB_Bench	Simple read/write tests of small objects.
SSD as Swap	YCSB	Read/write tests in memory with SSDs acting as swap.

Benchmark Hardware

	PM1725a	Z-SSD
Form factor	HHHL	HHHL
PCIe Lanes	Gen3 x8	Gen3 x4
Seq. Read BW (MB/sec)	6,400	3,200
Seq. Write BW (MB/sec)	3,000	3,200
4K Random Read (KIOPS)	1,080	750
4K Random Write (KIOPS)	170	170

Real Goal Was Lower Latency

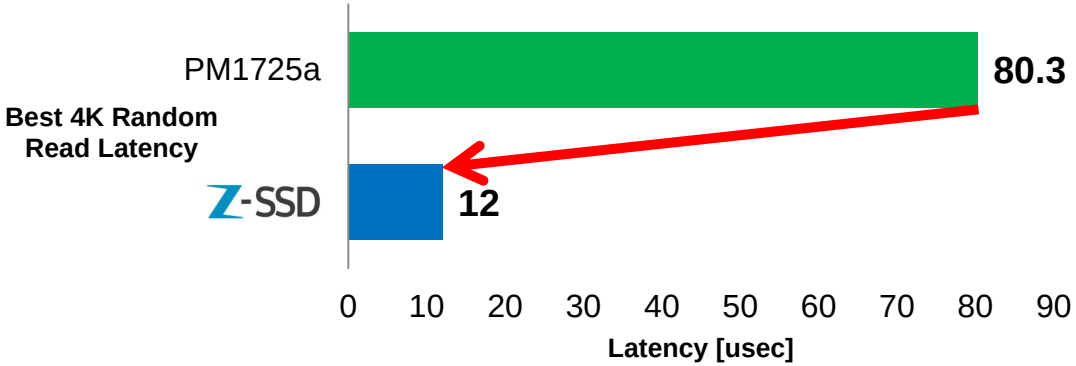

Ran Read Latency	12 - 20 us
Ran Write Latency (Typical)	16 us

+ software overhead

Caching

Used for:

- Increasing database performance
- Increase Web server performance

Fatcache is a cache that uses SSDs for storage of data. Twemperf was used to drive the benchmark.

Benchmark characteristics:

Very high read load

Small, 1 KB objects

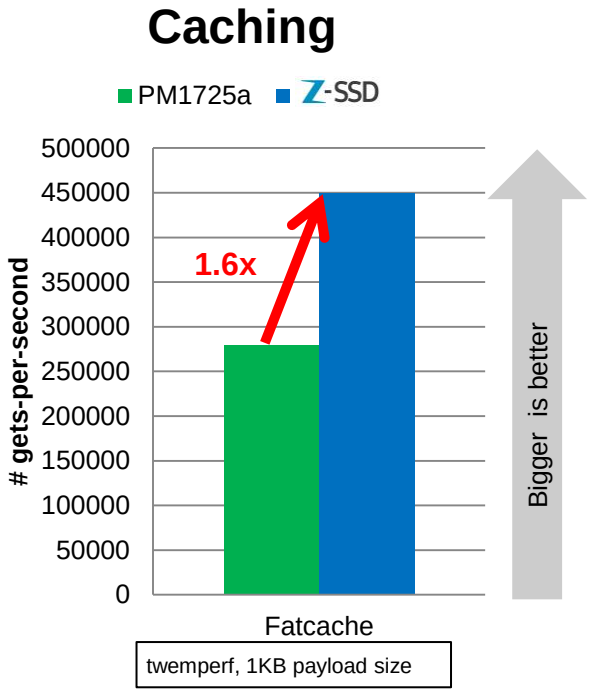

NoSQL Database

Used for:

- High speed, simple transactions
- Fraud prevention

RocksDB is a key-value store that can act as a back end for other databases such as MongoDB, Redis and MySQL.

Benchmark characteristics:

- Read while writing
- 800 byte objects

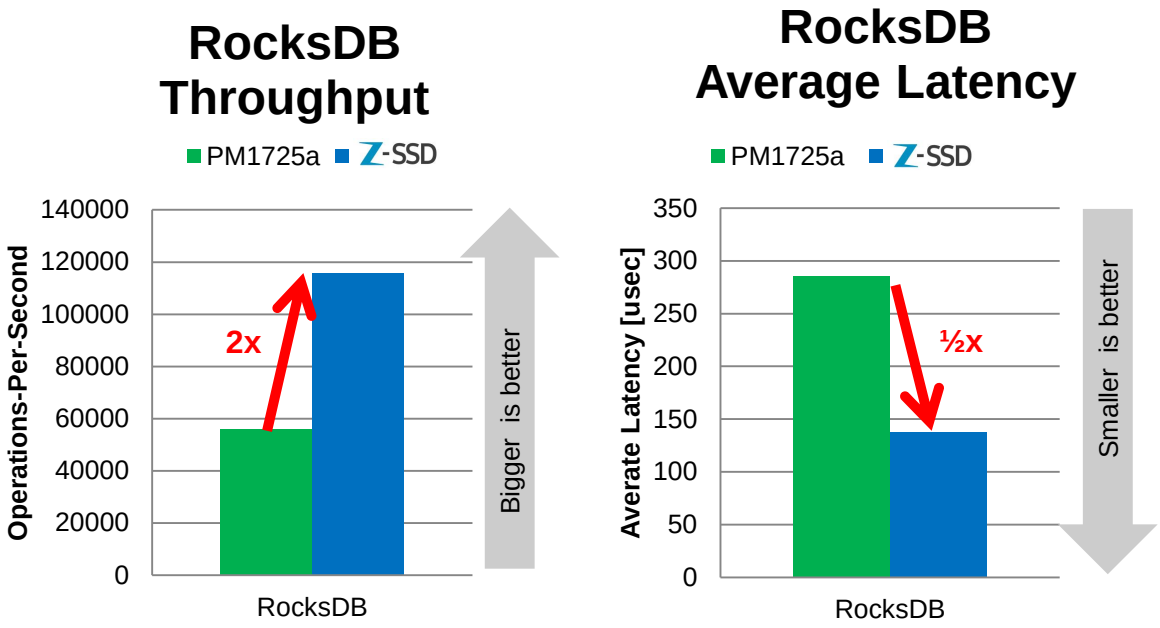

SSD as Swap

Used for:

- Extension to memory, when performance is less critical
- Overflow for memory if requirements are very dynamic

Memcached – Is an Open Source project and the most popular caching solution.

Benchmark characteristics:

- 90%read
- 10% update
- 1,000 B objects

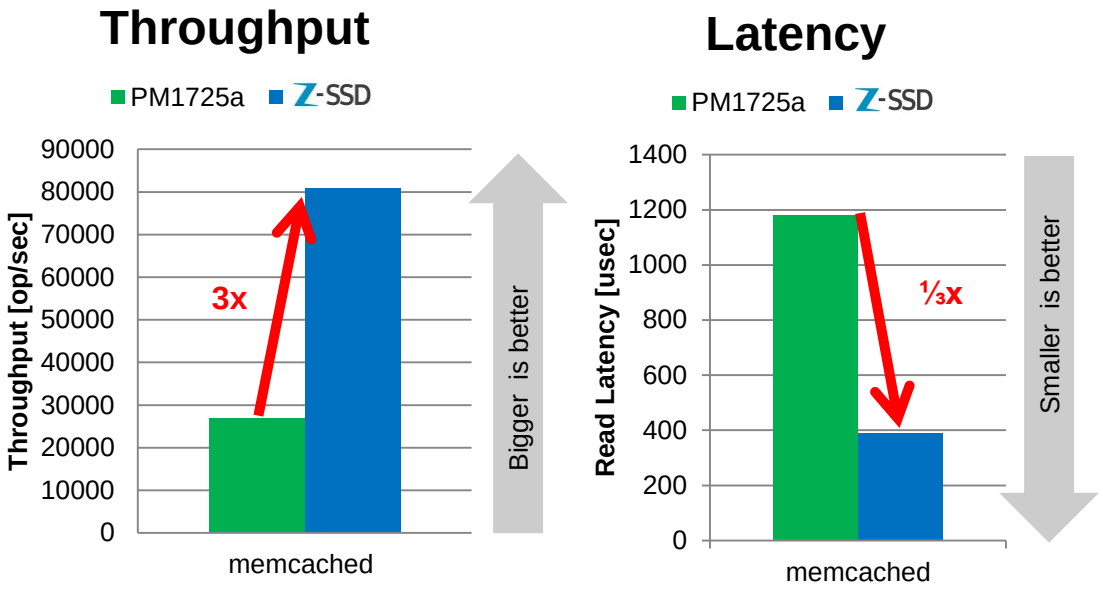

Memcached with 60GB dataset (45 million records, value size=1000B)
DRAM capacity=32GB (less than half of the used memory space)
Driven with YCSB, 90% read/ 10% update

Conclusions

- Lower latency does result in better performance in application benchmarks
- NVMe driver latency forms a barrier to SSD performance, so reducing latency more may not help.
- Flash improvements are not over – all flash is not equal

Would You Like To Learn More?

	Z-SSD
4K Random Read (KIOPS)	750
4K Random Write (KIOPS)	170
Seq. Read BW (MB/sec)	3,200
Seq. Write BW (MB/sec)	3,200
PCIe Lanes	Gen3 x4
Ran Read Latency	12 - 20 us
Ran Write Latency (Typical)	16 us
Endurance (DWPD)	**30**

Come see us at our booth.