Benefits of NVMe™ SSDs in Client Implementations

Jonmichael Hands, Product Marketing Manager, Intel SSDs

Cameron Brett, Director of Marketing, Outbound Marketing and Comms, Toshiba Memory America, Inc.

Dec 2018

Meet the Speakers

Jonmichael Hands
Intel

Cameron Brett
Toshiba

NVMe™ now Dominant Storage Interface for Client SSD

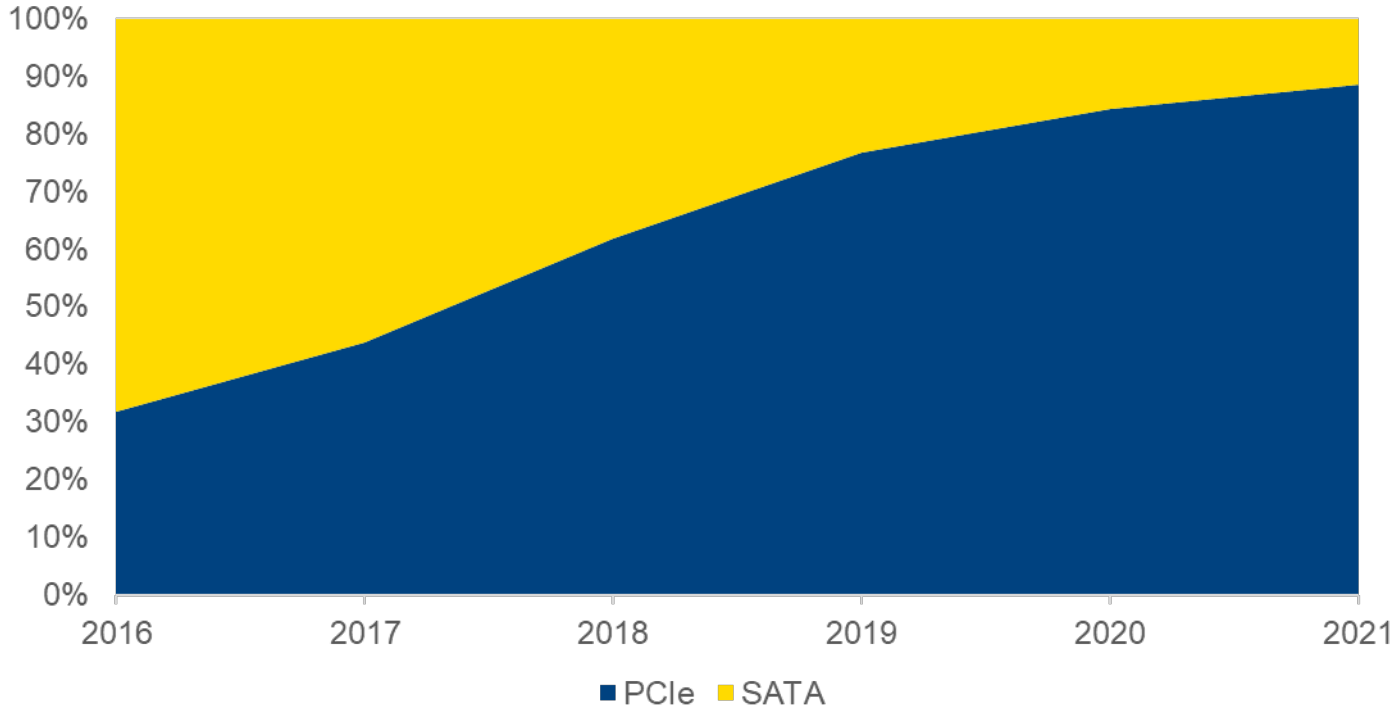

Source: Intel Forecasting, Q2'18

Client use cases for NVMe™

Gaming

Opens up the opportunity for unparalleled realism, with high quality textures and decreased load times

Content Creation

NVMe creates opportunity for new workflows for content creation when working with large data sets. Creators frequently move, backup, and duplicate storage

Workstation

Opportunity to accelerate any WS workload with large data requirements, reduce CPU idle time.

Speed up design, CAD, simulations

Client / Mobile

High performance is driving NVMe into client. Efficiency and features of NVMe lead to better battery life. Lower latency and better QoS delivers better application responsiveness

Media Creation

Rendering, high resolution (4k, 8k editing), audio production

Consumer product storage priorities

What are consumer storage needs

- Low cost
- Small form factor
- Optimal thermal and power management
- High performance
- Low active power usage
- Compatibility

Why is NVMe™ great for all consumer storage?

- Scalable streamlined storage stack
- Low latency
- Industry standard drivers in all OS
- Robust features to address power/thermals
- Scalability /w PCIe and next gen NVM
- Built in security and manageability features

Client Desktop PCIe Storage Form Factors

Add-in-card

M.2

U.2

Source: https://www.msi.com/Motherboard/X299-XPOWER-GAMING-AC.html

nvm EXPRESS®

M.2 mania!

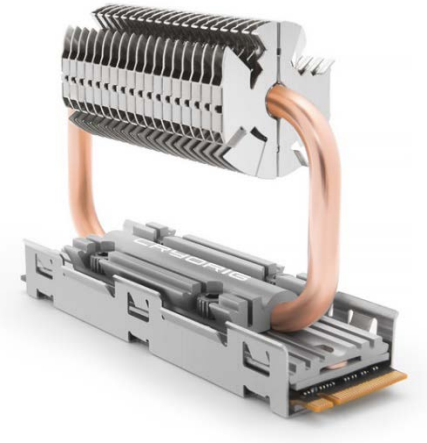

http://www.cryorig.com/news.php?id=80
https://www.asus.com/us/Motherboard-Accessory/HYPER-M-2-X16-CARD/gallery/
https://www.ekwb.com/shop/ek-m-2-nvme-heatsink-black

Choose the right laptop (hint…it needs NVMe™ SSD)

Choose the ultimate in form, function & style

Choosing a balance of performance, mobility & battery life in the right form factor is essential.

	2 in 1 personal laptops equipped with Intel® Core™ Processor (Y-Series)	Versatile laptops equipped with Intel® Core™ Processor (U-Series)	Intel® Core™ Processor-based clam shell form factor laptops (H-Series)	Intel® Core™ Processor-based clam shell laptops supporting overclocking (HK-Series)*
				
	• BGA or M.2 NVMe	• M.2 NVMe	• M.2 NVMe and 2.5in SATA	• Dual M.2 NVMe slots and 2.5in SATA
Portability				
Performance				
Battery life				

*Altering clock frequency or voltage may damage or reduce the useful life of the processor and other system components, and may reduce system stability and performance. Product warranties may not apply if the processor is operated beyond its specifications. Check with the manufacturers of system and components for additional details.

NVMe™ Scales to Mobile and Removable Storage

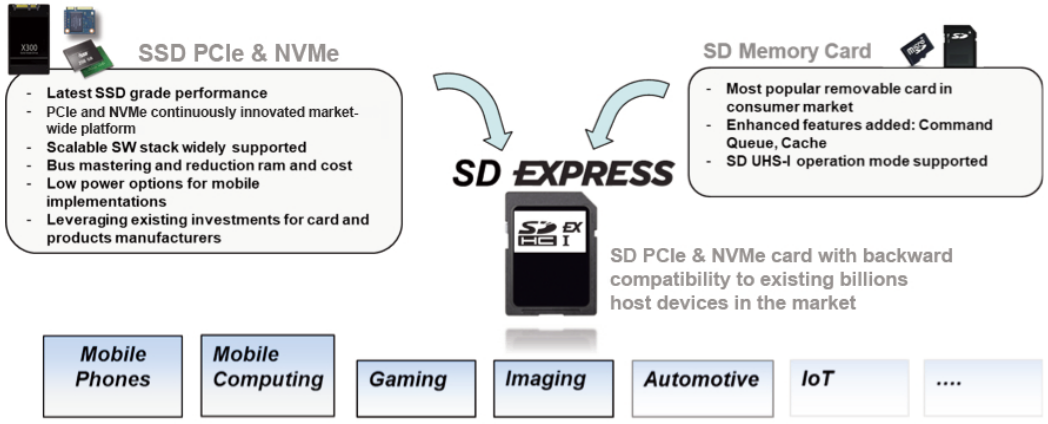

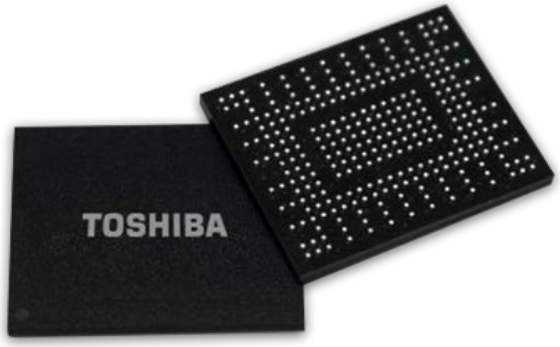

BGA 11.5x13mm

Source:
https://www.sdcard.org/downloads/pls/latest_whitepapers/SD_Express_Cards_with_PCIe_and_NVMe_Interfaces_White_Paper.pdf
https://business.toshiba-memory.com/en-us/product/storage-products/client-ssd.html

Google at the NVMe™ Developer Days 2018

Google

NVMe Storage for Consumer Product

Zhiping Yang, Ph.D.
Google LLC

1st Annual NVMe Developer Days
December 5–6, 2018, San Diego, California

2017 Pixelbook with the 1st 1113 BGA SSD in the world

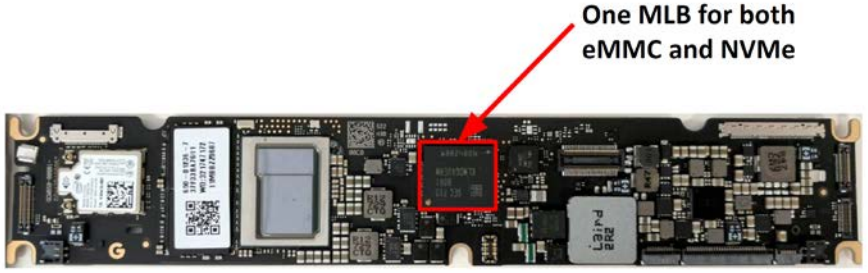

11.5x13 BGA SSD spec was finalized 2/17
Samsung samples in 7/17
Pixelbook with 512GB NVMe SSD was shipped in 12/17

WORLD RECORD

Google

Power Consumption

BAPCo MOBILEMARK 2014

55.32 mW

Device Idle Power

19.32 mW

4K Video Playback

112.19 mW

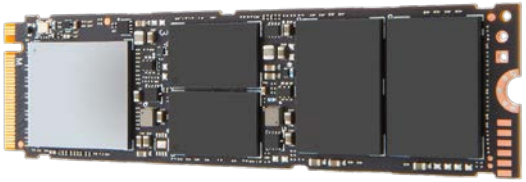

Intel® 760P SSD
Data is collected by Intel on Key sight 6705B* data logger by running Mobilemark* 2014 Office Productivity test for 2 hrs on Lenovo* Ideapad 720s. Windows* apps and other services are turned off for measurement consistency.
Data is collected by Intel on Key sight 6705B data logger by leaving the Lenovo Ideapad 720s for 10 mins and measuring the L1.2+PS3 power. Windows apps, radios, and other services are turned off for measurement consistency.
Data is collected by Intel on Key sight 6705B data logger by running 4K Video on the Lenovo Ideapad 720s for 1 hour and taking average of the measured power. Windows apps, radios, and other services are turned off for measurement consistency.
*Other names and brands may be claimed as the property of others.

NVMe™ removes the SATA performance bottleneck

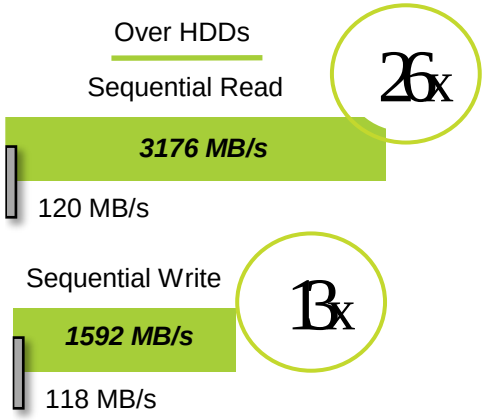

Intel®SSD 7 Series
versus
WD Blue* 5400RPM 500 GB HDD

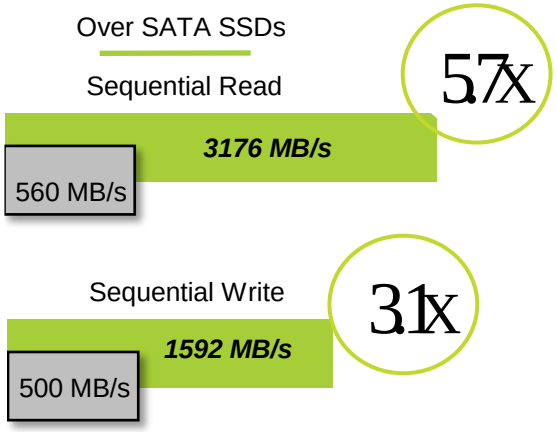

Intel®SSD 7 Series
versus
Intel®545 SATA-based SSD

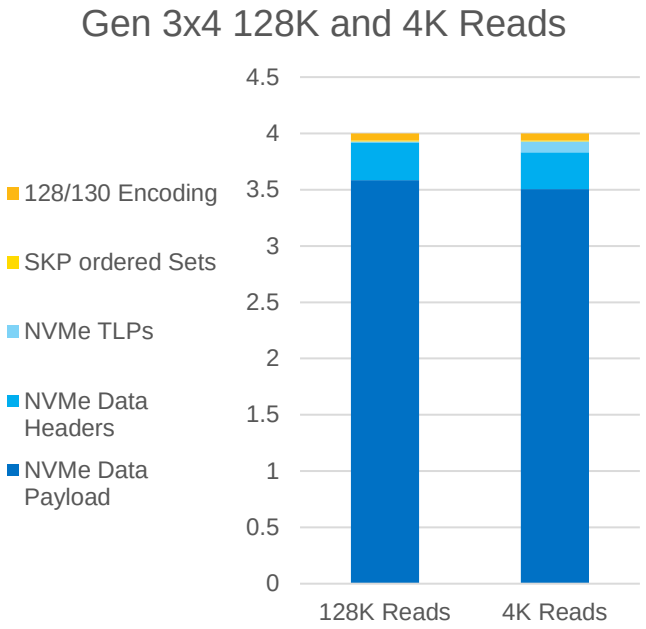

Storage performance comparison workload by Intel: CrystalDiskMark V5.2*. Drives being compared: Intel® 7600p vs Intel® 760p. System: processor: Intel® Core™ i7-7700K processor @ 4.5GHz Turbo Frequency, 8T/4C, 8MB cache, 91 W TDP, on motherboard: Asus Z270-A* Prime, memory: 2 X 4GB Corsair Vengeance DDR4* 3000MHz 8GB, operating system: Windows 10 Pro* (x64) OS RS2, storage: Intel® 600P and Intel® 760p 512GB. Drive under test is configured as a primary drive plugged into M.2 slot directly. System power profile set to performance mode. Data is collected at 500GB span size at Queue Depth 32 and thread 1 for sequential read and write. *Other names and brands may be claimed as the property of others.

Intel® SSD 660p Architecture: Dynamic SLC-Cache

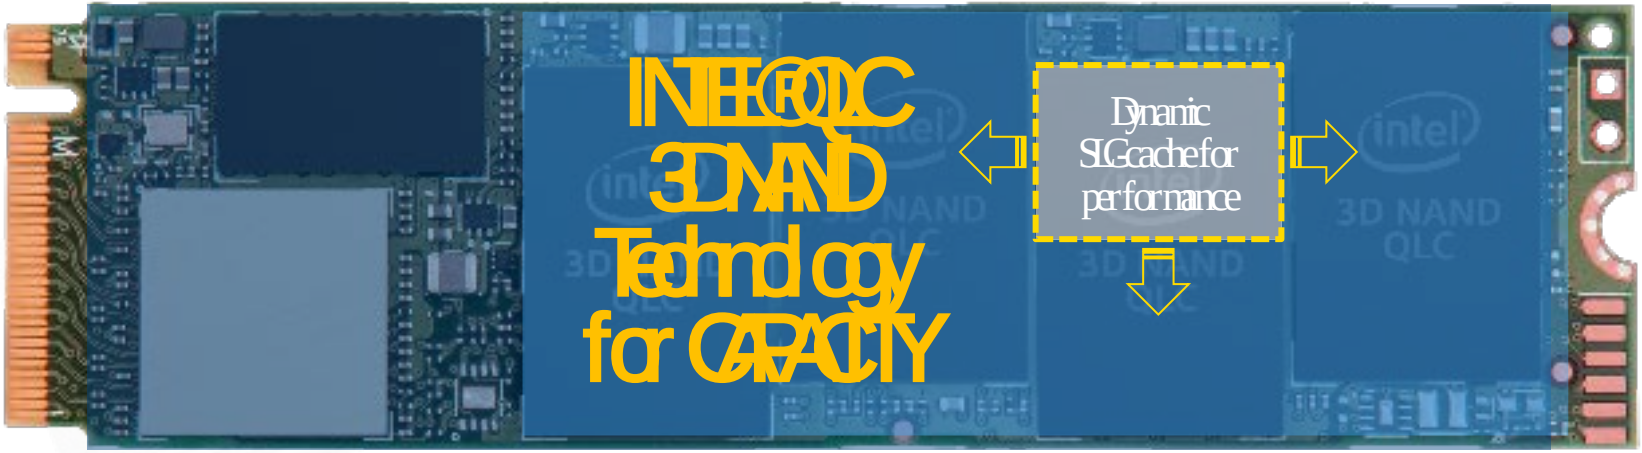

Intel® SSD 660p dynamic SLC-cache architecture:

- Dynamic SLC-cache is a **combination of static-SLC and QLC-mapped as SLC**
- **SLC-cache-first** algorithm ensures optimized performance[1]
- Intelligent firmware **configures unused QLC NAND cells as SLC** to dynamically expand SLC-cache
- **On-demand performance boost** option[2] to manually flush the SLC cache

[1] SSD performance within the SLC span.
[2] Performance Boost option available through Intel®SSD Toolbox version 3.5.3 within Intel®SSD optimizer feature, exclusively for Intel®SSD 660p Series.

Intel® SSD 660p. QLC Capacity and SLC Performance.

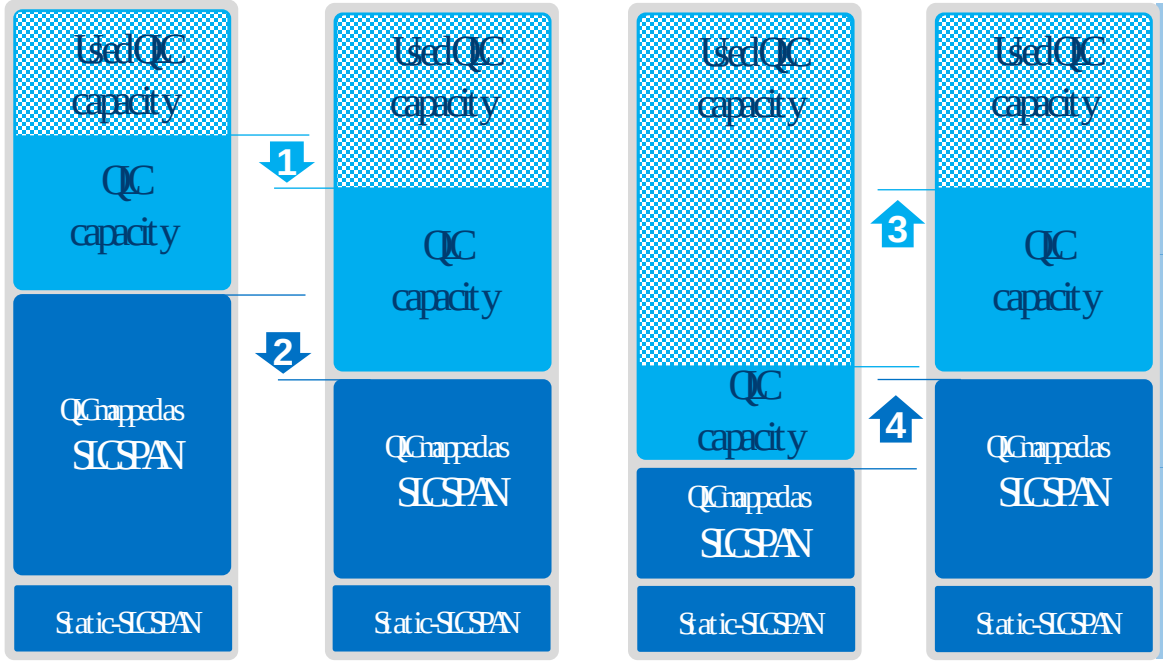

1. Increase in used capacity,
2. triggers SLC span decrease.
3. Decrease in used capacity,
4. triggers SLC span increase.

Intelligent architecture continually adjusts SLC spans up or down to boost performance over full life of SSD

QLC Offers Mainstream Performance and Price

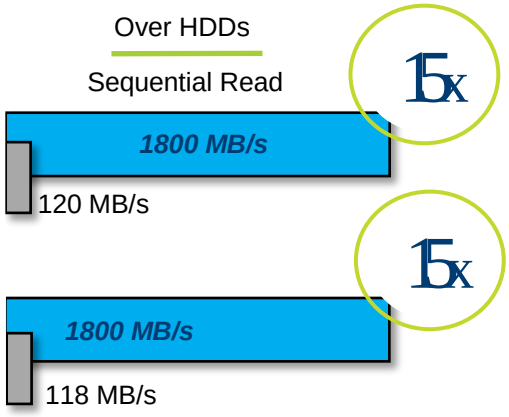

Intel® SSD 6 Series
versus
WD Blue* 5400RPM 500 GB HDD

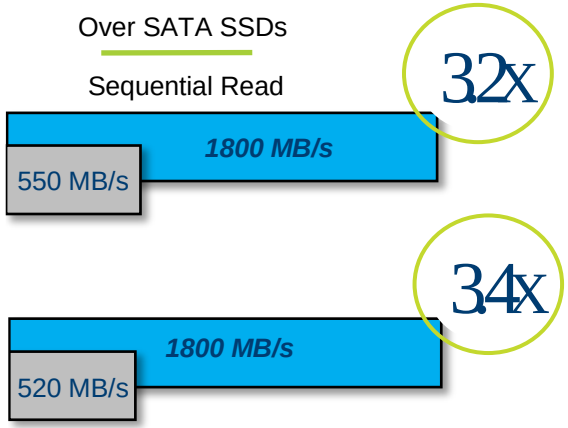

Intel® SSD 6 Series
versus
Samsung SSD 860 EVO*
SATA-based SSD

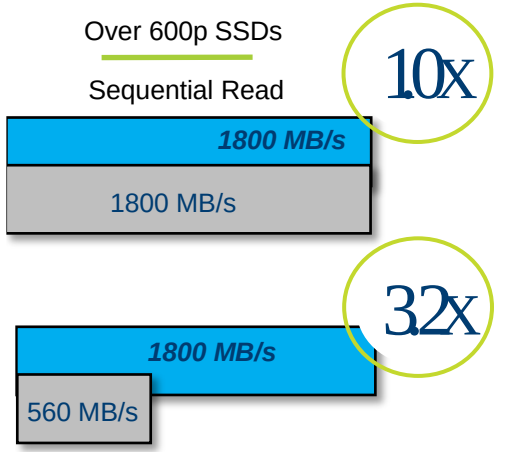

Intel® SSD 6 Series
versus
Previous generation Intel®
SSD 6 Series

Storage performance comparison workload by Intel: CrystalDiskMark V5.2.1*. Drives being compared: Intel® 660p. System: processor: Intel® Core™ i7-8700K @ 3.70GHz, Gigabyte Z370 AORUS Gaming 5 motherboard, EVGA GeForce GTX 1060 6GB SSC GAMING ACX 3.0, 6GB GDDR5 398.36, BIOS: American Megatrends Inc. F6 4/3/2018, Chipset: Intel® INF 10.1.1.42, Memory: 16GB (4X4GB) Crucial DDR4-2667*, Microsoft Windows 10* RS4 Enterprise 64-bit using native NVMe storage driver, Storage: Intel® 600P 1TB and Intel® 660p 2TB. Drive under test is configured as a primary drive plugged into M.2 slot directly. System power profile set to performance mode. Data is collected at 16GB span size at Queue Depth 32 and thread 1 for sequential read and write.

*Other names and brands may be claimed as the property of others.

QLC NVMe™ vs SATA Application Performance

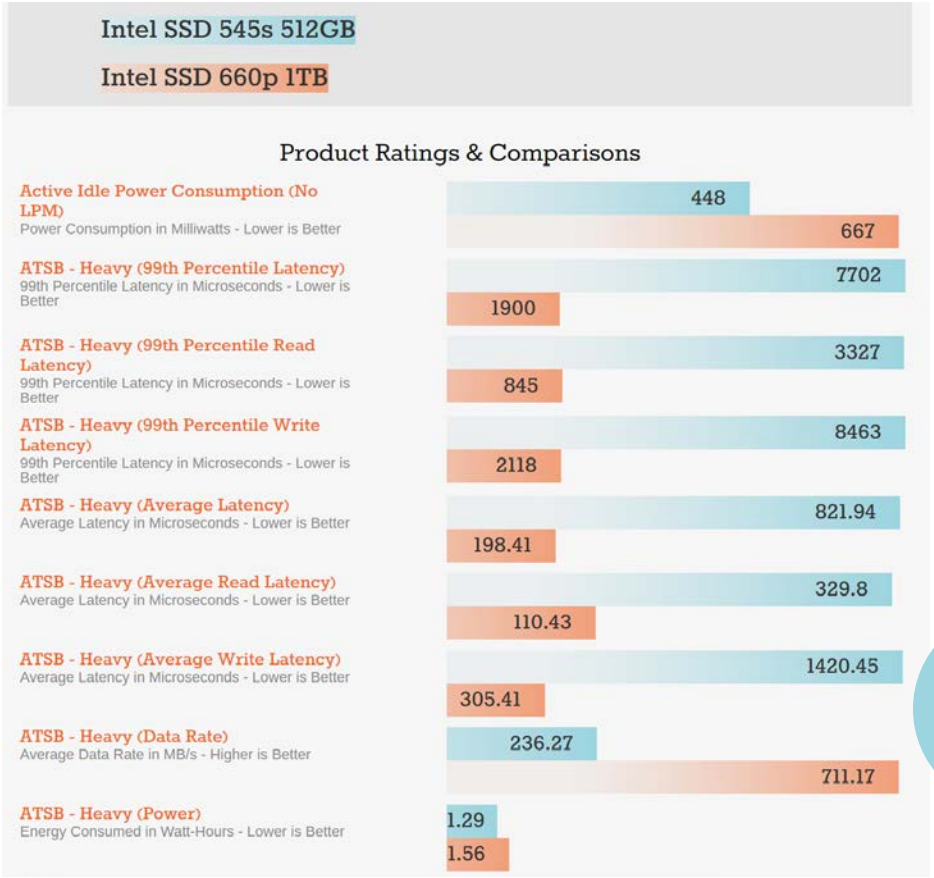

3x

*Other names and brands may be claimed as the property of others.

https://www.anandtech.com/bench/product/2191?vs=2166

Accelerate Video Editing With Intel® Optane™ SSDs

DAVINCI RESOLVE*

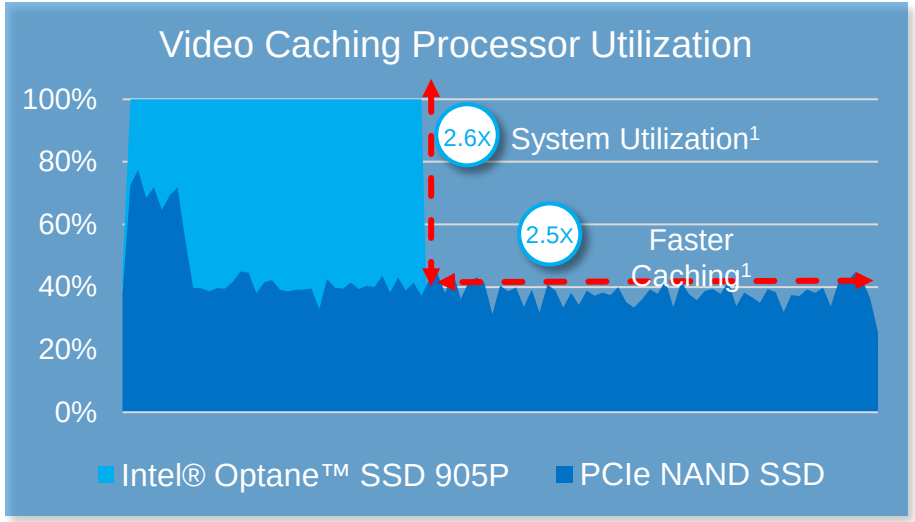

Up to **3.6X**
Faster renders[2]

Rendering a 3.5min 4K video went from **5.3mins** to **1.4mins**[2]

Up to **2.6X**
System Utilization[1]

System utilization improved **39%** to **100%**[1]

Up to **2.5X**
Faster Video Cache[1]

Caching a 3.5min 4K video went from **1.6mins** to **39secs**[1]

Performance results are based on testing as of July 21st 2018 and may not reflect the publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.
1.Test: Blackmagic DaVinci resolve 14* Video Caching of a 3.5mins @4K by using the command "media optimization." Test done by Intel in Sep 2018. System Configurations: Intel® Core™ i9-7900X, Asus X299 motherboard BIOS version 1401, NVIDIA* GeForce GTX1080, Memory 64GB (4X16GB) DDR4-2133, OS Win 10 version 1803, Storage 1TB Intel® SSD 760p vs. 480GB Intel® Optane™ SSD 905P.
2.Test: Blackmagic DaVinci resolve 14* Video Rendering of a 3.5mins @4K by rendering it to DPX file format at 4K/24FPs/10b. Test done by Intel in Sep 2018. System Configurations: Intel® Core™ i9-7900X, Asus* X299 motherboard BIOS version 1401, NVIDIA* GeForce GTX1080, Memory 64GB (4X16GB) DDR4-2133, OS Win 10 version 1803, Storage 1TB Intel® SSD 760p vs. 4800GB Intel® Optane™ SSD 905P.
*Other names and brands may be claimed as the property of others.

BGA Form Factor – Single-Package SSD

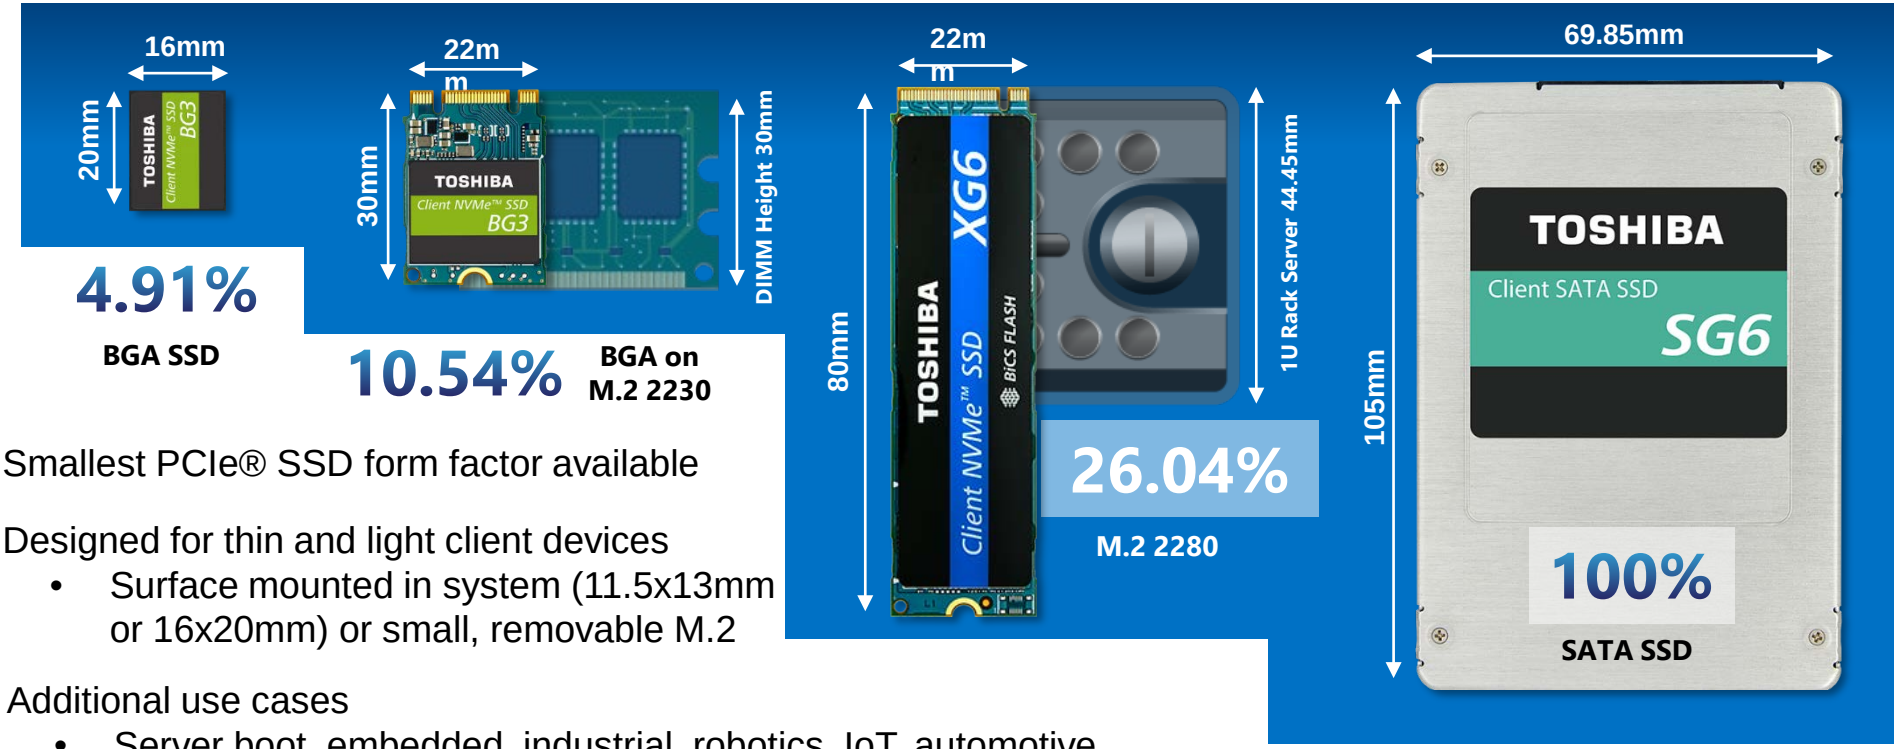

Smallest PCIe® SSD form factor available

Designed for thin and light client devices

- Surface mounted in system (11.5x13mm or 16x20mm) or small, removable M.2

Additional use cases

- Server boot, embedded, industrial, robotics, IoT, automotive

Benefits

- Smaller, lighter and low power
- Extends battery life and better user experience

Host Memory Buffer (HMB)

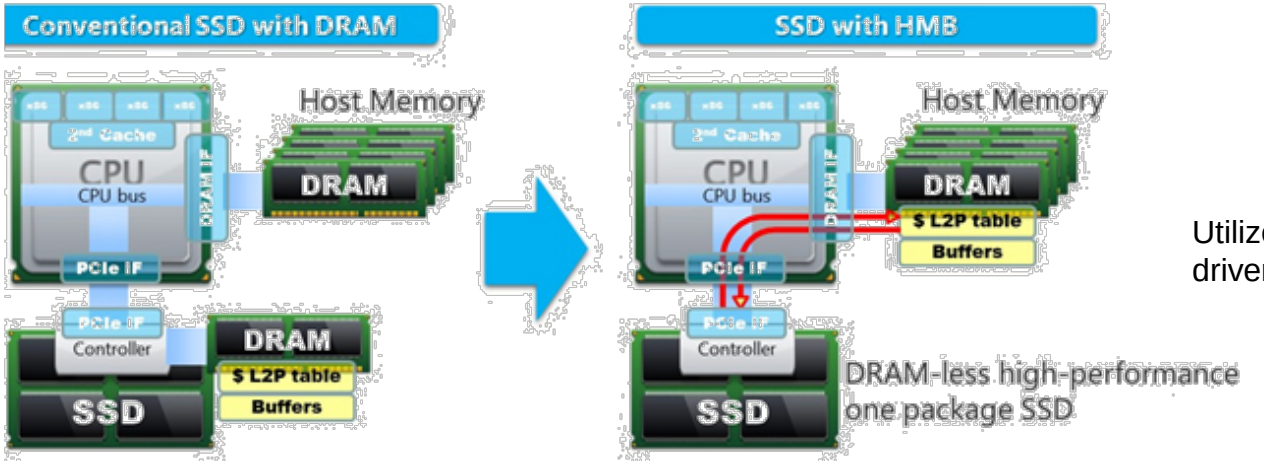

Utilizes HMB enabled device driver (i.e. Windows, Linux)

Standardized method (NVMe™ spec 1.2) to utilize host DRAM to hold LUT or other buffer data

Remove DRAM from SSD:

- Lowers BOM cost
- Reduced power requirements

Equivalent or better performance than SSD with DRAM

Off Module Power or Off-module PLP

- Power loss protection (PLP) is a mechanism to save DRAM-cached data and gracefully shut down an SSD upon an unexpected power loss condition
- PLP typically uses capacitors on the SSD to provide hold-up power until data is flushed from the DRAM to the NAND flash
- Off Module Power utilizes the host system or module to provide hold up power, eliminating the need for capacitors on the SSDs; helps reduce cost

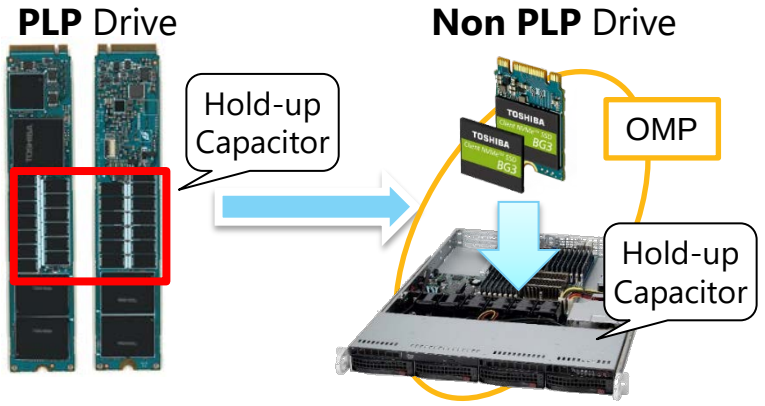

Standardization by PCI-SIG® and NVM Express®

New NVMe™ Features for Client SSDs

NVMe™ 1.2 Improvements for Client

RTD3

Allows safe shutdown to the storage to save platform power

Platform Value

- Enables safe shutdown of device
- Power savings

Specification Details:

- Spec provides registers for providing device details for entry/exit latencies.

Additional Power State Info

Provides host additional info to the power levels supported by the device

Platform Value

- Additional details of power states to assist in transitions.
- Power/thermal benefit

Specification Details

- Spec allocates details in SMART

NVMe innovations enable additional features for client to help manage power/thermals.

NVMe™ 1.2 Improvements for small form factors

Host Memory Buffer

Allows the host driver to allocate system memory for the SSD's exclusive use

Platform Value

- Enables DRAM savings & smaller BGA packages
- E.g., Allocate translation tables in host DRAM

Specification Details:

- Device indicates preferred HMB size
- Host enables/disables via Set Features

Composite Temperature

Allows host to monitor temperature of the SSD

Platform Value

- Platform has feedback to the device temperature.
- If the host believes the temperature is out of its limits, it can set a lower power state on the NVMe device

Specification Details

- Device indicates temperature in SMART
- Power State can be changed in power management

NVMe innovations enable scaling into smaller form factors delivering new differentiated platforms.

NVMe™ 1.3 - Boot Partitions

- Optional storage area that can be read with “fast” initialization method (not standard NVMe queues). Example: UEFI bootloader
- Saves cost and space by removing the need for another storage medium (like SPI flash, EPROM)
- Write using standard NVMe Firmware Download and Firmware Commit
- Can be protected with **Replay Protected Memory Block**

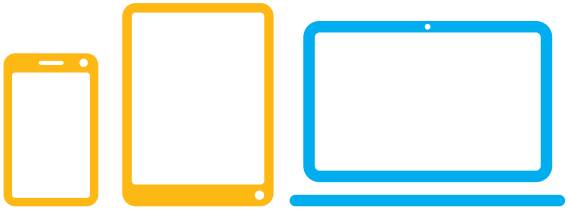

Makes NVMe more accessible for mobile and client form factors

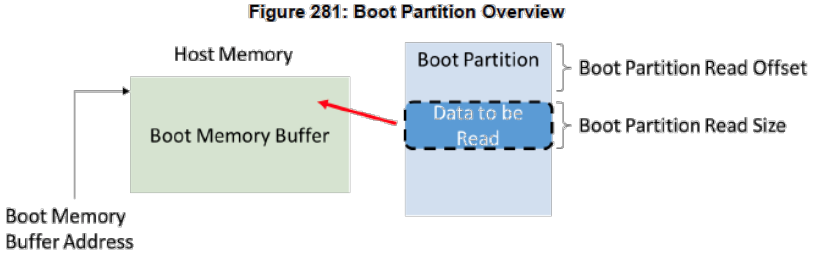

NVMe™ 1.3 - Host Controlled Thermal Management

Better thermal management in client systems like laptops and desktops.

Host can set **Thermal Management Temperature** at which a device should start going into a lower power state / throttling

- TMT1 – host tells SSD what temp in degrees K it should start throttling at
- TMT2 – threshold where the SSD should start heavy throttling regardless of impact to performance

Figure 264: HCTM Example

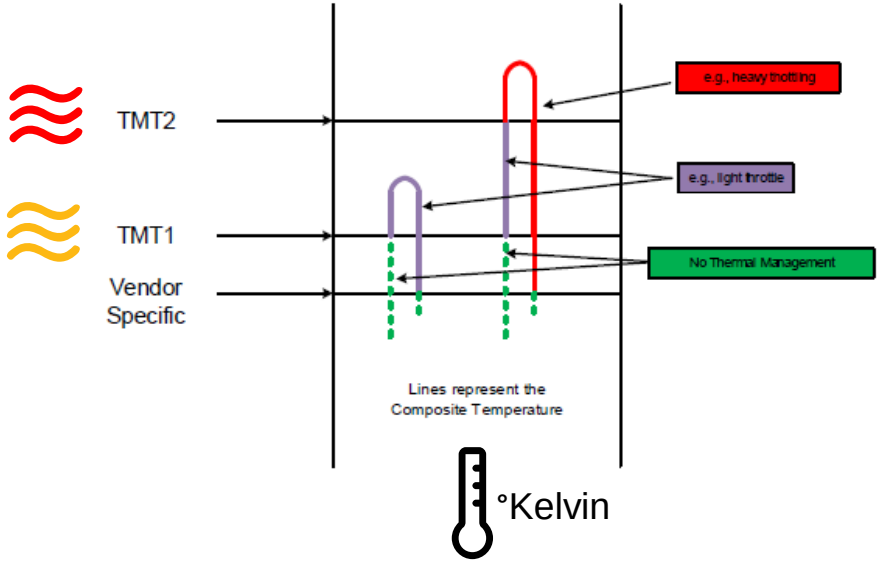

NVMe™ 1.4 –Namespace Write Protection

Namespace Write Protection is an optional configurable controller capability that enables the host to control the write protection state of a namespace.

(exactly what you think it does)

Could be used for secure space on drive, bootloader, backup image, important system files

Figure TBD1 – Namespace Write Protection State Machine Model

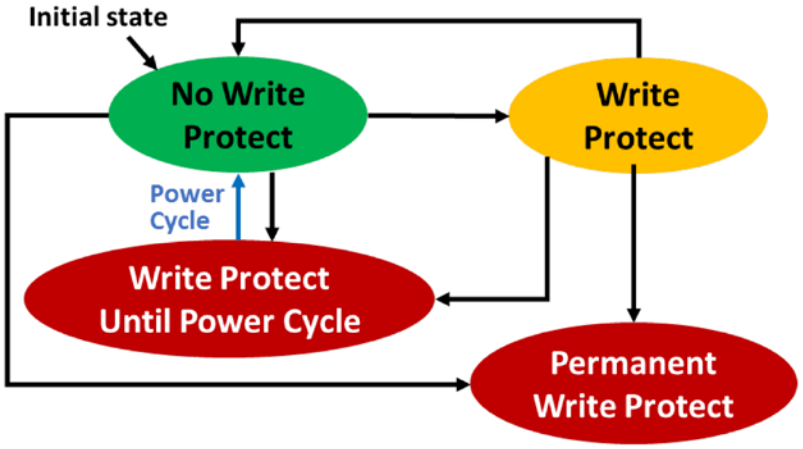


nvm
EXPRESS®
Architected for Performance