PCI Express* SSD Industry Standards

Jim Pappas
Director of Technology Initiatives
Intel Corporation

Legal Disclaimer

- INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL® PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. INTEL PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS.
- Intel may make changes to specifications and product descriptions at any time, without notice.
- All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.
- Intel, processors, chipsets, and desktop boards may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- and other code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user
- Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.
- Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. Go to: http://www.intel.com/products/processor_number
- Intel product plans in this presentation do not constitute Intel plan of record product roadmaps. Please contact your Intel representative to obtain Intel's current plan of record product roadmaps.
- Intel, Thunderbolt, UDK 2010 and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

Industry Standards for PCIe SSD's

Applications

Hardware

Standards	Status	Benefits
SNIA: NVM Programming TWG	• Approved Working Group in SNIA	• Deliver absolute highest application performance • Establishes NVM Memory Access Models
SCSI Express*	• Under development in T10	Preserves SCSI software infrastructure
NVM Express* (NVMe)	• Spec 1.0 completed March,2011 • 80+ members of NVM Express working group • OS Drivers released (Windows and Linux) • Numerous products in development	Optimized NVM Storage interface Designed for Client and Server
PCI Express* 3.0 (PCIe 3.0)	•	Highest performance Lowest latency Lowest power
SSD FF WG (2.5"/3.5" FF & connector)	• 60 member companies • Spec 1.0 published Dec'11 • 1.0 specification submitted to SFF • SFF8639 broadly supported by industry	Serviceability Scalable capacity Supports SATA IO, SATA 3.0, SAS 3.0 & PCIe 3.0

Driving technological advances that both establish and maintain high levels of standards-based innovation

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

Industry Standards for PCIe SSD's

Applications

Hardware

Standards	Status	Benefits
SNIA: NVM Programming TWG	• Approved Working Group in SNIA	• Deliver absolute highest application performance • Establishes NVM Memory Access Models
SCSI Express*	• Under development in T10	Preserves SCSI software infrastructure
NVM Express* (NVMe)	• Spec 1.0 completed March,2011 • 80+ members of NVM Express working group • OS Drivers released (Windows and Linux) • Numerous products in development	Optimized NVM Storage interface Designed for Client and Server
PCI Express* 3.0 (PCIe 3.0)	•	Highest performance Lowest latency Lowest power
SSD FF WG (2.5"/3.5" FF & connector)	• 60 member companies • Spec 1.0 published Dec'11 • 1.0 specification submitted to SFF • SFF8639 broadly supported by industry	Serviceability Scalable capacity Supports SATA IO, SATA 3.0, SAS 3.0 & PCIe 3.0

Driving technological advances that both establish and maintain high levels of standards-based innovation

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

Industry Standards for PCIe SSD's

Applications

Hardware

Standards	Status	Benefits
SNIA: NVM Programming TWG	• Approved Working Group in SNIA	• Deliver absolute highest application performance • Establishes NVM Memory Access Models
SCSI Express*	• Under development in T10	Preserves SCSI software infrastructure
NVM Express* (NVMe)	• Spec 1.0 completed March,2011 • 80+ members of NVM Express working group • OS Drivers released (Windows and Linux) • Numerous products in development	Optimized NVM Storage interface Designed for Client and Server
PCI Express* 3.0 (PCIe 3.0)	•	Highest performance Lowest latency Lowest power
SSD FF WG (2.5"/3.5" FF & connector)	• 60 member companies • Spec 1.0 published Dec'11 • 1.0 specification submitted to SFF • SFF8639 broadly supported by industry	Serviceability Scalable capacity Supports SATA IO, SATA 3.0, SAS 3.0 & PCIe 3.0

Driving technological advances that both establish and maintain high levels of standards-based innovation

Industry Standards for PCIe SSD's

Applications

Hardware

Standards	Status	Benefits
SNIA: NVM Programming TWG	• Approved Working Group in SNIA	• Deliver absolute highest application performance • Establishes NVM Memory Access Models
SCSI Express*	• Under development in T10	Preserves SCSI software infrastructure
NVM Express* (NVMe)	• Spec 1.0 completed March,2011 • 80+ members of NVM Express working group • OS Drivers released (Windows and Linux) • Numerous products in development	Optimized NVM Storage interface Designed for Client and Server
PCI Express* 3.0 (PCIe 3.0)	•	Highest performance Lowest latency Lowest power
SSD FF WG (2.5"/3.5" FF & connector)	• 60 member companies • Spec 1.0 published Dec'11 • 1.0 specification submitted to SFF • SFF8639 broadly supported by industry	Serviceability Scalable capacity Supports SATA IO, SATA 3.0, SAS 3.0 & PCIe 3.0

Driving technological advances that both establish and maintain high levels of standards-based innovation

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

NVM Express

- NVMe is defined to scale for future NVM
  - Host controller standards live for 10+ years
  - Future NVM may have sub microsecond latencies
- 1M IOPS needs highly efficient driver approach
  - Benefits from removing OS queues, IO scheduler, and SCSI layer while optimizing for NVMe
- Block layer attach reduces overhead > 50%
  - Block layer: 2.8 µs, 9100 cycles
  - Traditional: 6.0 µs, 19500 cycles

Linux * Storage Stack

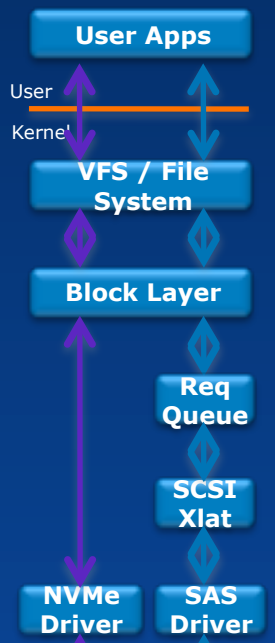

Chatham NVMe Prototype

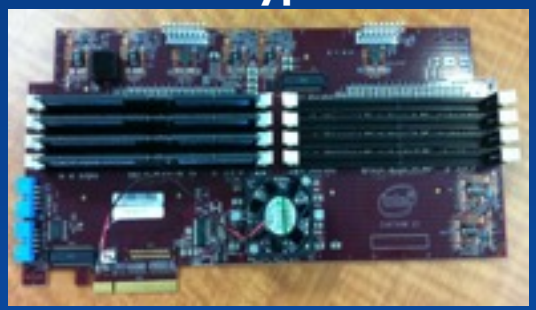

Prototype Measured IOPS

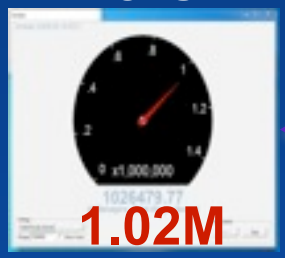

Cores Used for 1M IOPs

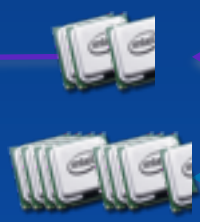

www.nvmexpress.org

Measurement taken on Intel Core i5-2500K Sandy Bridge 3.3GHz 6MB L3 Cache Quad-Core Desktop Processor using Linux RedHat EL6.0 2.6.32-71 Kernel

Industry Standards for PCIe SSD's

Applications

Hardware

Standards	Status	Benefits
SNIA: NVM Programming TWG	• Approved Working Group in SNIA	• Deliver absolute highest application performance • Establishes NVM Memory Access Models
SCSI Express*	• Under development in T10	Preserves SCSI software infrastructure
NVM Express* (NVMe)	• Spec 1.0 completed March,2011 • 80+ members of NVM Express working group • OS Drivers released (Windows and Linux) • Numerous products in development	Optimized NVM Storage interface Designed for Client and Server
PCI Express* 3.0 (PCIe 3.0)	•	Highest performance Lowest latency Lowest power
SSD FF WG (2.5"/3.5" FF & connector)	• 60 member companies • Spec 1.0 published Dec'11 • 1.0 specification submitted to SFF • SFF8639 broadly supported by industry	Serviceability Scalable capacity Supports SATA IO, SATA 3.0, SAS 3.0 & PCIe 3.0

Driving technological advances that both establish and maintain high levels of standards-based innovation

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

SCSI Express*

- Marketing name for T10 SCSI over PCIe (SOP) project
- Utilizes the decades of existing SCSI software infrastructure

www.T10.org and www.scsita.org

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

Industry Standards for PCIe SSD's

Applications

Hardware

Standards	Status	Benefits
SNIA: NVM Programming TWG	• Approved Working Group in SNIA	• Deliver absolute highest application performance • Establishes NVM Memory Access Models
SCSI Express*	• Under development in T10	Preserves SCSI software infrastructure
NVM Express* (NVMe)	• Spec 1.0 completed March,2011 • 80+ members of NVM Express working group • OS Drivers released (Windows and Linux) • Numerous products in development	Optimized NVM Storage interface Designed for Client and Server
PCI Express* 3.0 (PCIe 3.0)	•	Highest performance Lowest latency Lowest power
SSD FF WG (2.5"/3.5" FF & connector)	• 60 member companies • Spec 1.0 published Dec'11 • 1.0 specification submitted to SFF • SFF8639 broadly supported by industry	Serviceability Scalable capacity Supports SATA IO, SATA 3.0, SAS 3.0 & PCIe 3.0

Driving technological advances that both establish and maintain high levels of standards-based innovation

SNIA: NVM Programming TWG Status

- Founding members
  - Dell, EMC, Fujitsu, HP, IBM, Intel, NetApp, Oracle, QLogic, Symantec
- Charter: Develop specifications for new software “programming models” as NVM becomes a standard feature of platforms
  - Scope:
    - In-kernal NVM programming models
    - Kernal-to-application programming models
  - Programming models specify the exact technical behavior, up to (but not including) the OS specific API semantics
- APIs
  - Each OSV codes the programming models to specific to OS
  - Linux Open Source project underway to provide the Linux

NVM Programming Models

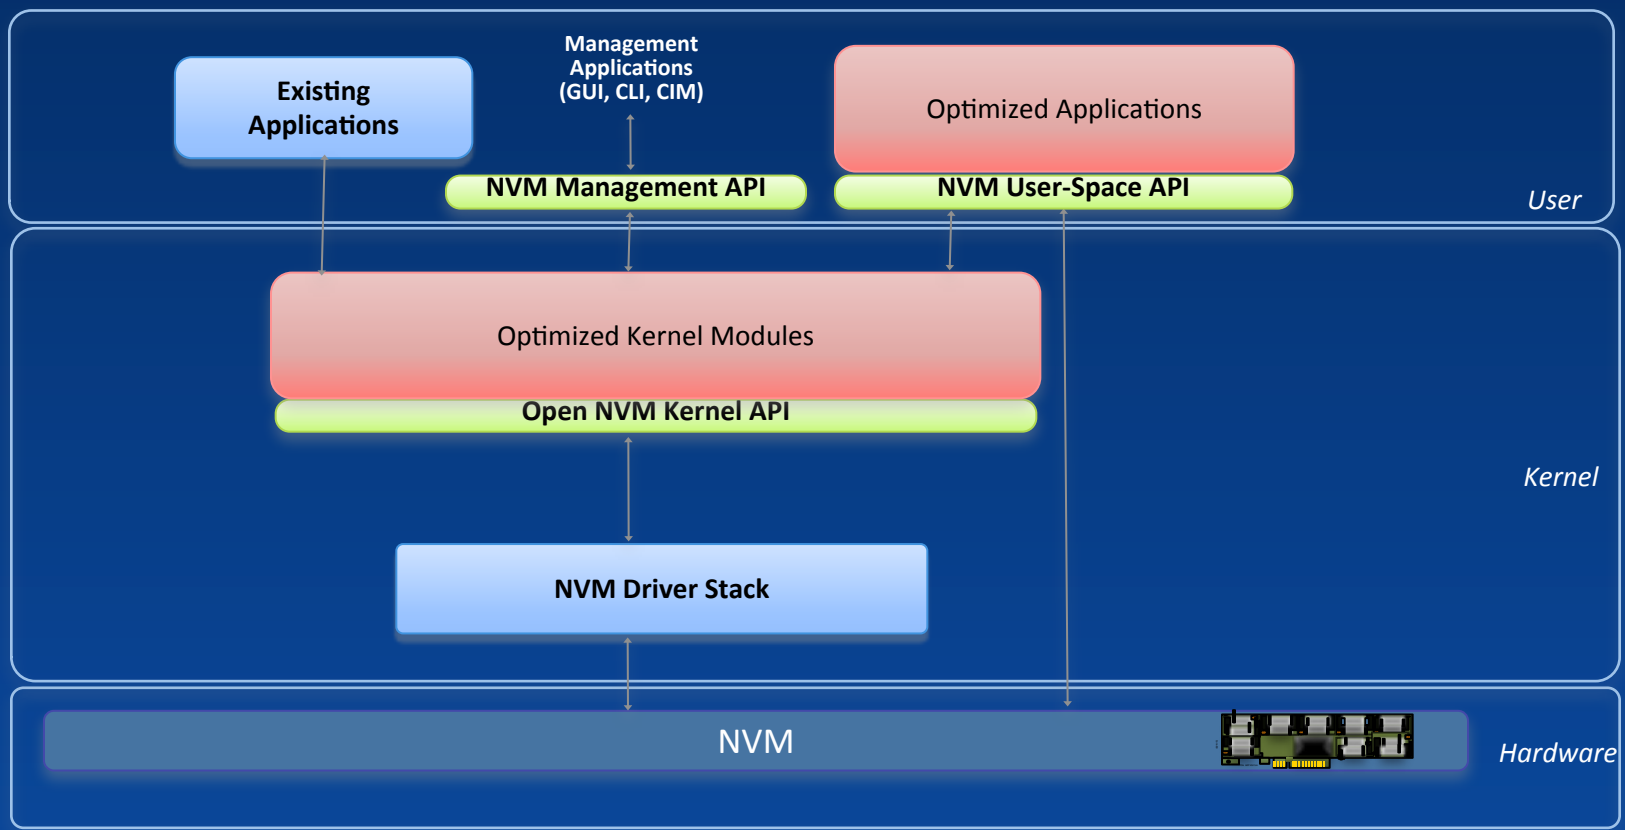

SNIA NVM Programming TWG

Linux Open Source Project -- or -- Microsoft

Existing/Unchanged Infrastructure

Summary

- PCIe* is becoming the primary SSD interconnect
  - Necessary standards in development or complete
  - Significant industry investment
- NVM Programming TWG will drive significant change to computer architecture
- Call to action: Join these organizations to participate in this movement
  - http://www.snia.org/
  - http://www.scsita.org/
  - http://www.nvmexpress.org/
  - http://www.pcisig.com/
  - http://www.ssdformfactor.org/

Risk Factors

- The above statements and any others in this document that refer to plans and expectations for the second quarter, the year and the future are forward-looking statements that involve a number of risks and uncertainties. Words such as “anticipates,” “expects,” “intends,” “plans,” “believes,” “seeks,” “estimates,” “may,” “will,” “should,” and their variations identify forward-looking statements. Statements that refer to or are based on projections, uncertain events or assumptions also identify forward-looking statements. Many factors could affect Intel’s actual results, and variances from Intel’s current expectations regarding such factors could cause actual results to differ materially from those expressed in these forward-looking statements. Intel presently considers the following to be the important factors that could cause actual results to differ materially from the company’s expectations. Demand could be different from Intel's expectations due to factors including changes in business and economic conditions, including supply constraints and other disruptions affecting customers; customer acceptance of Intel’s and competitors’ products; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Potential disruptions in the high technology supply chain resulting from the recent disaster in Japan could cause customer demand to be different from Intel’s expectations. Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term and product demand that is highly variable and difficult to forecast. Revenue and the gross margin percentage are affected by the timing of Intel product introductions and the demand for and market acceptance of Intel's products; actions taken by Intel's competitors, including product offerings and introductions, marketing programs and pricing pressures and Intel’s response to such actions; and Intel’s ability to respond quickly to technological developments and to incorporate new features into its products. The gross margin percentage could vary significantly from expectations based on capacity utilization; variations in inventory valuation, including variations related to the timing of qualifying products for sale; changes in revenue levels; product mix and pricing; the timing and execution of the manufacturing ramp and associated costs; start-up costs; excess or obsolete inventory; changes in unit costs; defects or disruptions in the supply of materials or resources; product manufacturing quality/yields; and impairments of long-lived assets, including manufacturing, assembly/test and intangible assets. Expenses, particularly certain marketing and compensation expenses, as well as restructuring and asset impairment charges, vary depending on the level of demand for Intel's products and the level of revenue and profits. The majority of Intel’s non-marketable equity investment portfolio balance is concentrated in companies in the flash memory market segment, and declines in this market segment or changes in management’s plans with respect to Intel’s investments in this market segment could result in significant impairment charges, impacting restructuring charges as well as gains/losses on equity investments and interest and other. Intel's results could be affected by adverse economic, social, political and physical/infrastructure conditions in countries where Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Intel’s results could be