

# Closing Summary & Future Directions

Amber Huffman, Sr. Principal Engineer, Intel Corporation

August 11, 2015

# Legal Notices and Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at [intel.com](http://intel.com), or from the OEM or retailer.

No computer system can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit <http://www.intel.com/performance>.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

Statements in this document that refer to Intel's plans and expectations for the quarter, the year, and the future, are forward-looking statements that involve a number of risks and uncertainties. A detailed discussion of the factors that could affect Intel's results and plans is included in Intel's SEC filings, including the annual report on Form 10-K.

The products described 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.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel is a sponsor and member of the BenchmarkXPRT Development Community, and was the major developer of the XPRT family of benchmarks. Principled Technologies is the publisher of the XPRT family of benchmarks. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases.

Intel and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

\*Other names and brands may be claimed as the property of others. © 2015 Intel Corporation.



# Agenda


- **NVMe in the Real World Today**
- Continuous Innovation Underway in NVMe
- Looking Ahead and Resources to Learn More

# Lenovo ThinkServer RD650 with Intel® SSD DC P3700 Series Database TPC-H

7x

PERFORMANCE

Total database performance




New Server



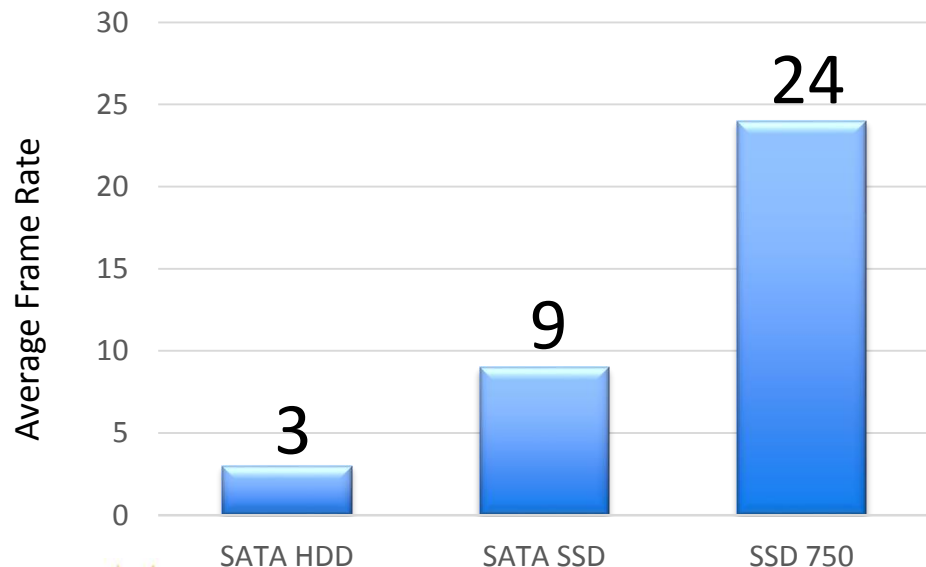
4 to 1 server consolidation

Legacy Servers



# Dual 4K Video Editing in Real Time with NVMe™

Uncompressed Video Playback of Two  
4K Streams in Adobe SpeedGrade



Real time 4K editing made possible

Design & build richer content with larger  
data sets, textures and assets

NVMe SSD = **~2.5x** (frames/sec) SATA SSD  
NVMe SSD = **~8x** (frames/sec) SATA HDD

# NVMe Driver Ecosystem is Strong

- NVMe drivers available on Windows\*, Linux\*, Solaris\*, VMware\*, UEFI
- Many are native / in-box drivers



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

# NVMe Available in a Variety of Form Factors

- NVMe products have been announced in form factors shown, depending on whether targeted for Data Center or Client


U.2 SSD  
(formally SFF-8639)



PCIe add-in card (CEM)



M.2



BGA SSD



# Robust Interop Program in Place

- The University of New Hampshire Interoperability Lab (UNH-IOL) has collaborated with NVMe to deliver a robust interop program
- Four plugfests have been held – populating robust NVMe Integrator’s List

## *Portion of Integrator’s List*

NVMe Integrator's List v1.2					
• NVMe Integrator's List Policy v1.2					
NVMe Devices					
Product	Product Type	Firmware Version	Interop Program Revision	Date Listed	Further Info
Intel SSD DC P3608 Series	NVMe SSD	8DV10160	v1.2	7/6/15	<a href="http://www.intel.com">http://www.intel.com</a>
Memblaze Pblaze4 Card Series	NVMe SSD	v0.08.0060	v1.2	7/21/15	<a href="http://www.memblaze.com">http://www.memblaze.com</a>
Phison NVMe PS5007-E7	NVMe SSD Controller	v1.0	v1.2	7/17/15	<a href="http://www.phison.com">http://www.phison.com</a>
PMC FlashtecTM NVMe2016/2032	NVMe SSD	NA	v1.2	7/16/15	<a href="http://pmcs.com/products/">http://pmcs.com/products/</a>
Seagate SandForce SF3700 SSD Controller	NVMe SSD Controller Chip	G.0.1	v1.2	7/6/15	<a href="http://www.seagate.com">http://www.seagate.com</a>
SK Hynix SSD PE3000 Series	NVMe SSD	20050A00	v1.2	7/6/15	<a href="http://ssd.skhynix.com">ssd.skhynix.com</a>

More details at [www.iol.unh.edu/services/testing/NVMe](http://www.iol.unh.edu/services/testing/NVMe).



# Data Center Use Cases for NVMe



## Cloud computing

Better SLAs for CSPs, lower opx/capx, get developers to market faster, consumers services on demand



## Virtualization

Lowering enterprise IT by increasing system utilization and improving virtual machine scalability



## HPC

Eliminating bottlenecks in HPC workflows. NVMe keeps up with high bandwidth demands of HPC to speed up overall workflow times by an order of magnitude



## Database

High performance and great QoS shine in traditional database



## Big data

High bandwidth and low latency can provide business insights with real time analytics

# 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



## Workstation

Opportunity to accelerate any WS workload with large data sets  
Caching from backend SAN in large organizations



## Client / Mobile

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



## 4K

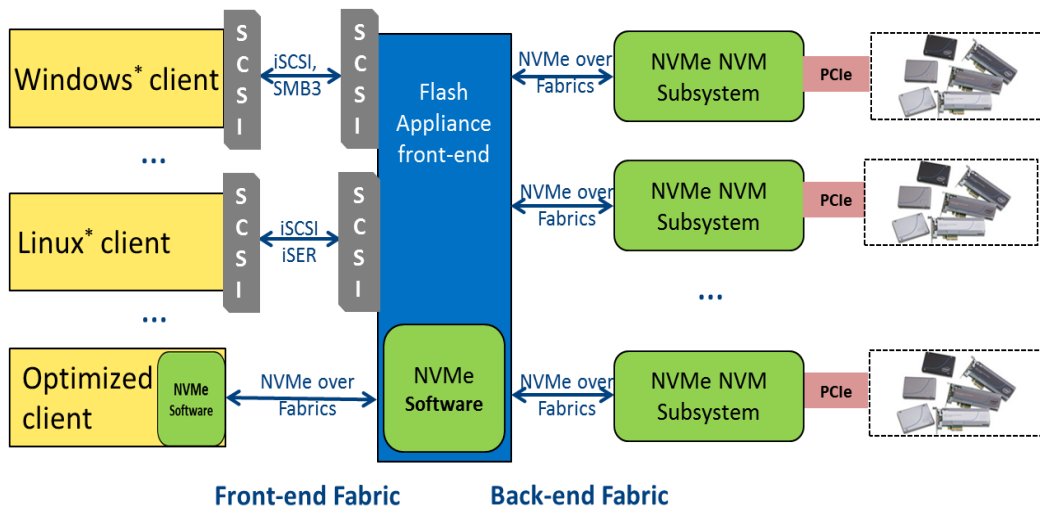
High bandwidth is required for real time 4K editing

# Agenda

- NVMe in the Real World Today
- **Continuous Innovation Underway in NVMe**
- Looking Ahead and Resources to Learn More

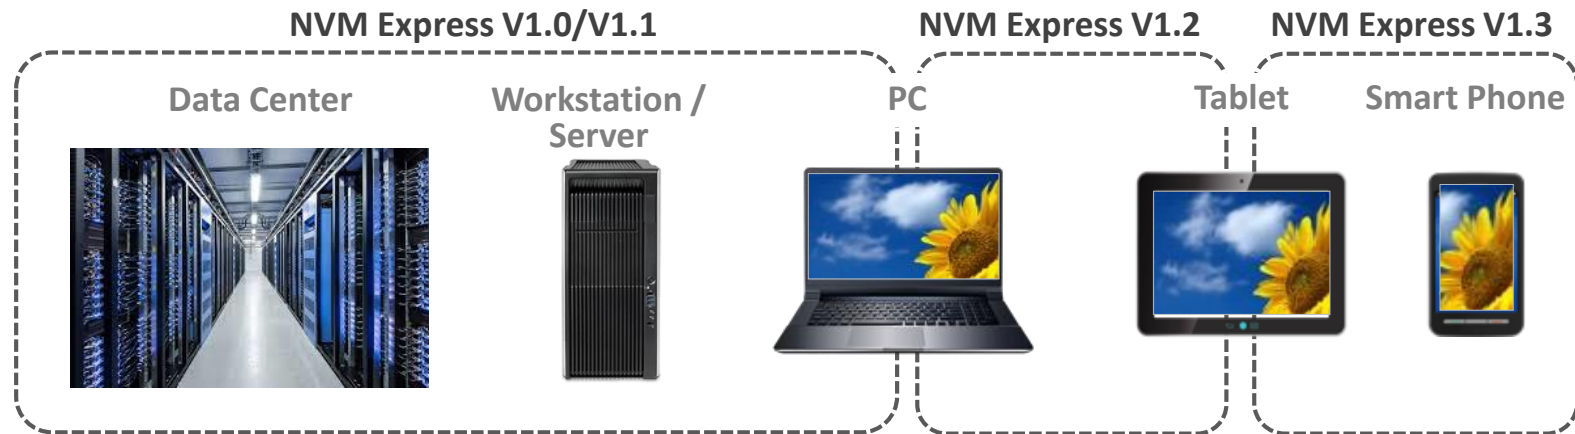
# NVMe over Fabrics

- Scales NVMe across the Datacenter over Fabrics like Ethernet, Fibre Channel and InfiniBand™
- Encapsulates NVMe commands to leverage ~ 90% of NVMe
- Three separate prototypes have shown < 10 μs adder between local and remote NVMe



*Extend efficiency of NVMe over front and back-end fabrics*

# Enabling NVMe in Mobile Segments



NVMe Express

- Optimized for NVMe
- Low Latency
- Exploits Parallelism
- Efficient SW stack

PCI Express

- Smart Device / Bus Master
- No HBA

- Host memory buffer
- Replay Memory (RPMB)
- Power Enhancements

- **Boot Partitions for non-BIOS boot**

- M.2
- Emerging BGA
- L1.2 Sub-states

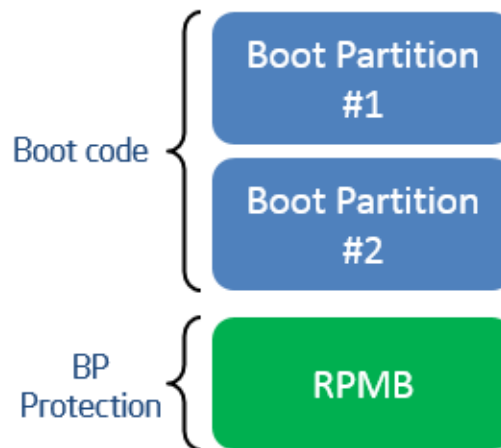
- **Smaller BGA**

NVMe/PCIe  
Data Center to Mobile

# Enhancements for Mobile

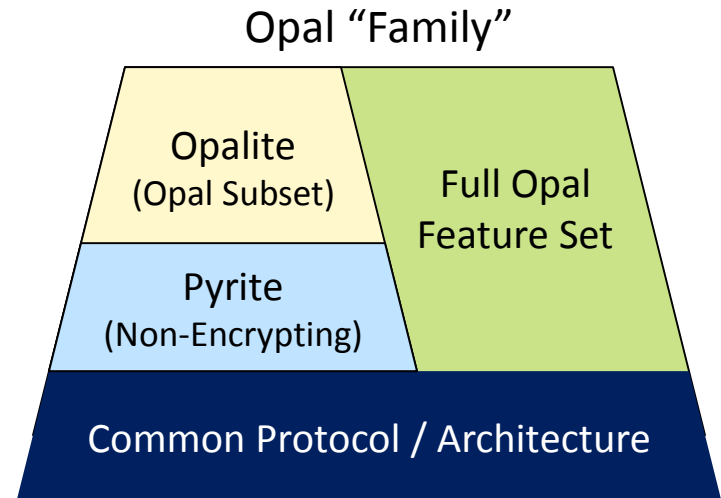
- PCIe is a low power interface
- Smaller BGA coming for mobile
- NVMe defined “Boot Partitions” to enable non-BIOS boot
  - Read via MMIO registers
  - Write via enhanced firmware download commands
  - Protect/lock with RPMB

Item	PCIe Gen3	PCIe Gen2	M-PHY Gear3
Line Speed [Gbps]	8	5	5.83
PHY overhead	128/130, 1[GB/s]	8/10, 500[MB/s]	8/10, 583[MB/s]
Active Power [mW]	60 (L0)	46 (L0)	58 (HS)
Standby Power [mW]	0.11 (L1.2)	0.11 (L1.2)	0.2 (Hibern8)
MB/mJ (higher better)	14-18	8-12	8-12



# Security – The NVMe and TCG Partnership

- NVMe is leveraging the security expertise of the Trusted Computing Group (TCG)
- TCG has developed a “family” of specifications to scale across the needs of NVMe in different Client and Enterprise solutions
- NVMe and TCG plan to continue collaborating on future security features for NVMe

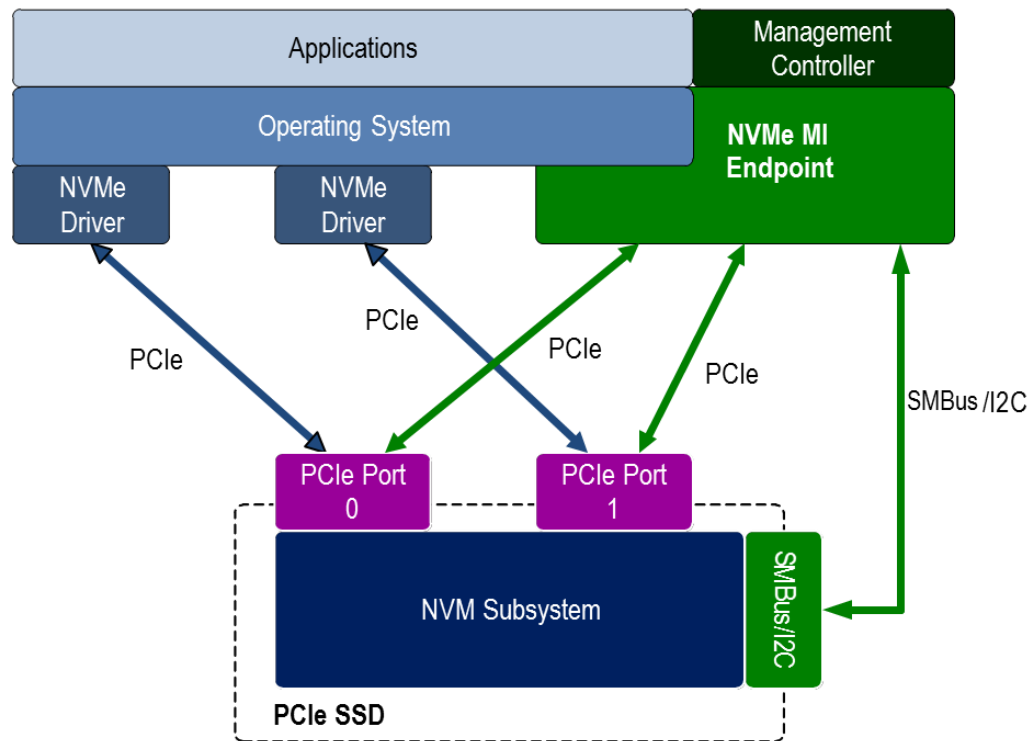


# Management Interface

NVMe Specification

NVMe Device Management Specification

- Enables pre-boot and out-of-band management during run-time
- Power budgeting, inventory, health monitoring, firmware update, etc
- Standardization benefits include:
  - Reduces cost & broadens adoption
  - Common feature set
  - Industry ecosystem (including compliance testing)





# Management Interface Roll-Out

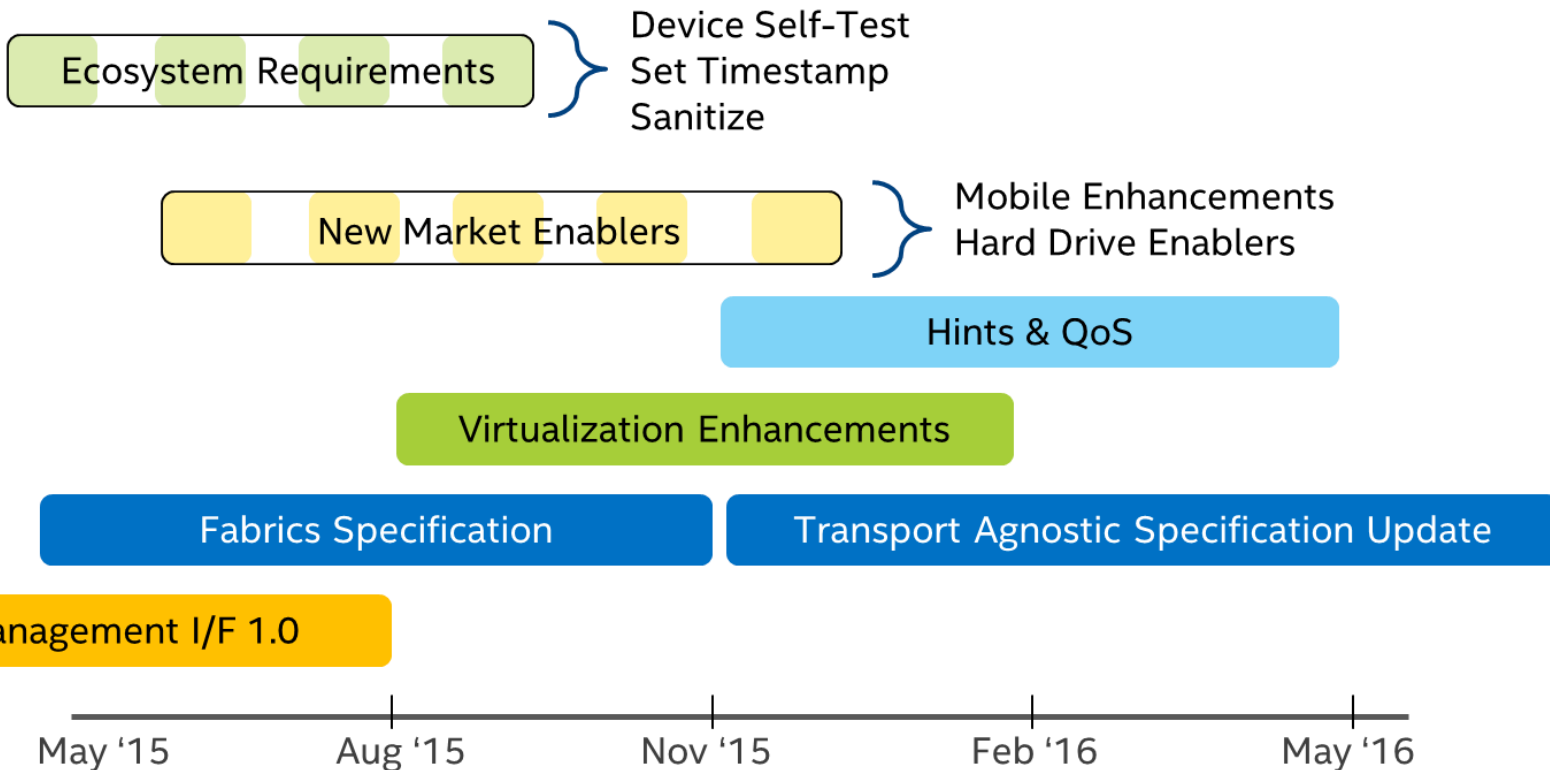
- Basic Management Command published in February 2015
  - Standardized way to poll NVMe devices for basic health status over SMBus
- Full Management Specification will be published in ~ October 2015
  - Full specification has started ratification, includes in-band and out-of-band built on MCTP
  - The Basic Management command is an optional feature – and will not be enhanced

*Take advantage of standardized enclosure management*

# Agenda

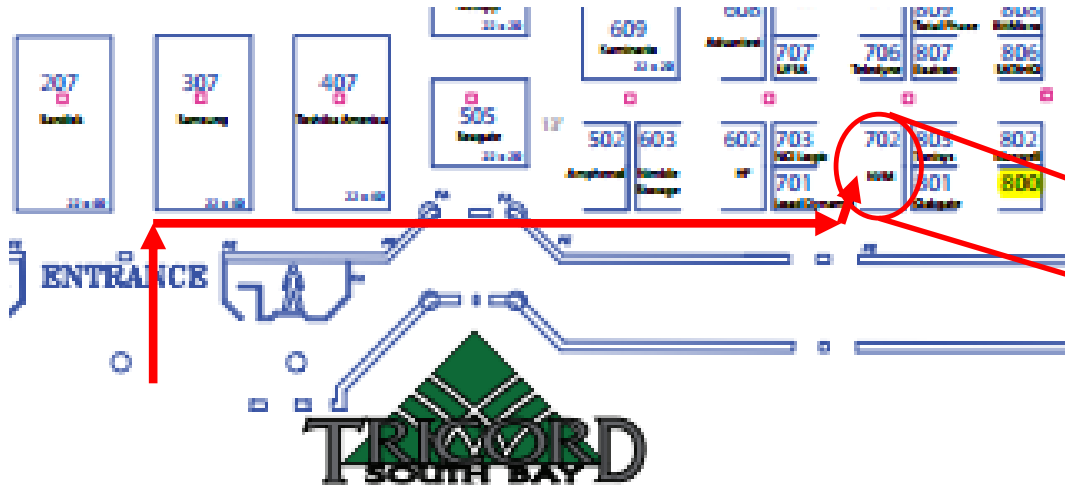
- NVMe in the Real World Today
- Continuous Innovation Underway in NVMe
- **Looking Ahead and Resources to Learn More**

# NVMe Technology Roadmap



# See NVMe in Action at FMS

Visit our FMS Exhibitor Booth #702



# Learn More at Intel Developer Forum



- IDF is August 18-20 at the Moscone Center in San Francisco
- Check out some of the great sessions on NVM Express

Session ID	Session Title	Presenter(s)
SSDS001	NVM Express™: The Data Center and Client Storage Transformation	Amber Huffman and Mike Shapiro
SSDS003	What You Need to Know to Win the Storage Transition – Preparing for NVM Express™ in the Data Center	Jonmichael Hands and Michael Hall
SSDC001	Tech Chat: Benchmarking Data Center Solid-State Drives – Insights Into Industry-Leading NVM Express* SSD Performance Metrics	Pallavi Pandit
SSDC003	Tech Chat: NVM Express* Features for High Availability and Storage Eco-System	Tahmid Rahman



# More Demos at IDF

- The NVM Express Community at IDF shows off technology from 16 companies
  - Check out **today's** NVMe PCIe SSD products
  - Preview **tomorrow's** early prototypes from several IHVs of NVMe over Fabrics



Company	Booth #
Aperion Data Systems	873
EMC	887
HGST	886
Intel	871 & 881
JDSU	874
Kazan Networks	880
Keysight Technologies	876
Microsoft	879
PMC-Sierra	882
QLogic Corporation	883
Samsung Semiconductor	884
Seagate Technology	878
SK Hynix	885
Storage Networking Industry Association (SNIA)	888
Super Micro Computer	877
Teledyne LeCroy	872
Viking Technology	875

# Keep Up to Date on the Latest

- NVMe Blog: <http://www.nvmexpress.org/blog>
- Wikipedia: [https://en.wikipedia.org/wiki/NVM\\_Express](https://en.wikipedia.org/wiki/NVM_Express)
- Twitter @NVMExpress: <https://twitter.com/NVMexpress>
- LinkedIn: <https://www.linkedin.com/grp/home?gid=4307826>



# Summary

- NVMe is available today on PCI Express in Data Center and Client
- New features and innovation are coming, including:
  - NVMe over Fabrics
  - Mobile Enhancements
  - Management Interface
  - Security


*Get involved – Join NVMe at <http://nvmexpress.org/join-nvme/>*



# SSD Give Away at the NVMe Booth!

- Giving away 7 SSDs
- Wed: 1:30, 4:30, 6:30
- Thursday: 12:30

## *NVM Express Booth #702*



*Stop by booth #702 to enter drawing (must be present to win)*

# *Thank You*





*Architected for Performance*