

All About M.2 SSDs



SNIA WEBCAST



HOSTED BY THE  
SOLID STATE STORAGE INITIATIVE



Tom Coughlin,  
President  
Coughlin Associates  
[www.tomcoughlin.com](http://www.tomcoughlin.com)

Tom Coughlin, President, Coughlin Associates is a widely respected storage analyst and consultant. He has over 30 years in the data storage industry with multiple engineering and management positions at high profile companies.

Dr. Coughlin has many publications and six patents to his credit. Coughlin Associates provides market and technology analysis as well as Data Storage Technical Consulting services. Tom publishes the *Digital Storage Technology Newsletter*, the *Media and Entertainment Storage Report*, and other industry reports.

Tom is active with SMPTE, SNIA, the IEEE (he is Director Elect for IEEE Region 6 and active in the Consumer Electronics Society) and other professional organizations. Tom is the founder and organizer of the Annual Storage Visions Conference ([www.storagevisions.com](http://www.storagevisions.com)), a partner to the International Consumer Electronics Show, as well as the Creative Storage Conference ([www.creativestorage.org](http://www.creativestorage.org)). He is the general chairman of the annual Flash Memory Summit. For more information on Tom Coughlin and his publications. go to [www.tomcoughlin.com](http://www.tomcoughlin.com).

- M.2 Market Overview – Jim Handy, Objective Analysis
- M.2 Cards – Jon Tanguy, Micron
- M.2 Connection Schemes – Jaren May, TE Connectivity
- NVM Express – David Akerson, Intel
- M.2 Performance – Eden Kim, Calypso
- Wrap-up – Tom Coughlin, Coughlin Associates
- Q&A

# M.2 SSDs' Bright Future



# SNIA WEBCAST

Presented by:  
Jim Handy  
Objective Analysis.



HOSTED BY THE  
SOLID STATE STORAGE INITIATIVE



Jim Handy, Director  
Objective Analysis  
Los Gatos, California  
[Objective-Analysis.com](http://Objective-Analysis.com)

Jim Handy of Objective Analysis has over 35 years in the electronics industry including 20 years as a leading semiconductor and SSD industry analyst. Early in his career he held marketing and design positions at leading semiconductor suppliers including Intel, National Semiconductor, and Infineon. A frequent presenter at trade shows, Mr. Handy is known for his technical depth, accurate forecasts, widespread industry presence and volume of publication. He has written hundreds of market reports, articles for trade journals, and white papers, and is frequently interviewed and quoted in the electronics trade press and other media. He writes the [Chip Talk blog for Forbes](#), [TheMemoryGuy.com](#), and [TheSSDguy.com](#).

Both major SSD markets can use this format:

## ➤ Enterprise

- ◆ PCIe – Up to 4 lanes
  - › Preferred interface in the enterprise
  - › High bandwidth
  - › Mounts flush with motherboard, good for blades
- ◆ Can be used as SATA boot disk

## ➤ Client

- ◆ Small form factor, thin
- ◆ Upgrade path from SATA to PCIe for speed

## ❖ Hot Swap

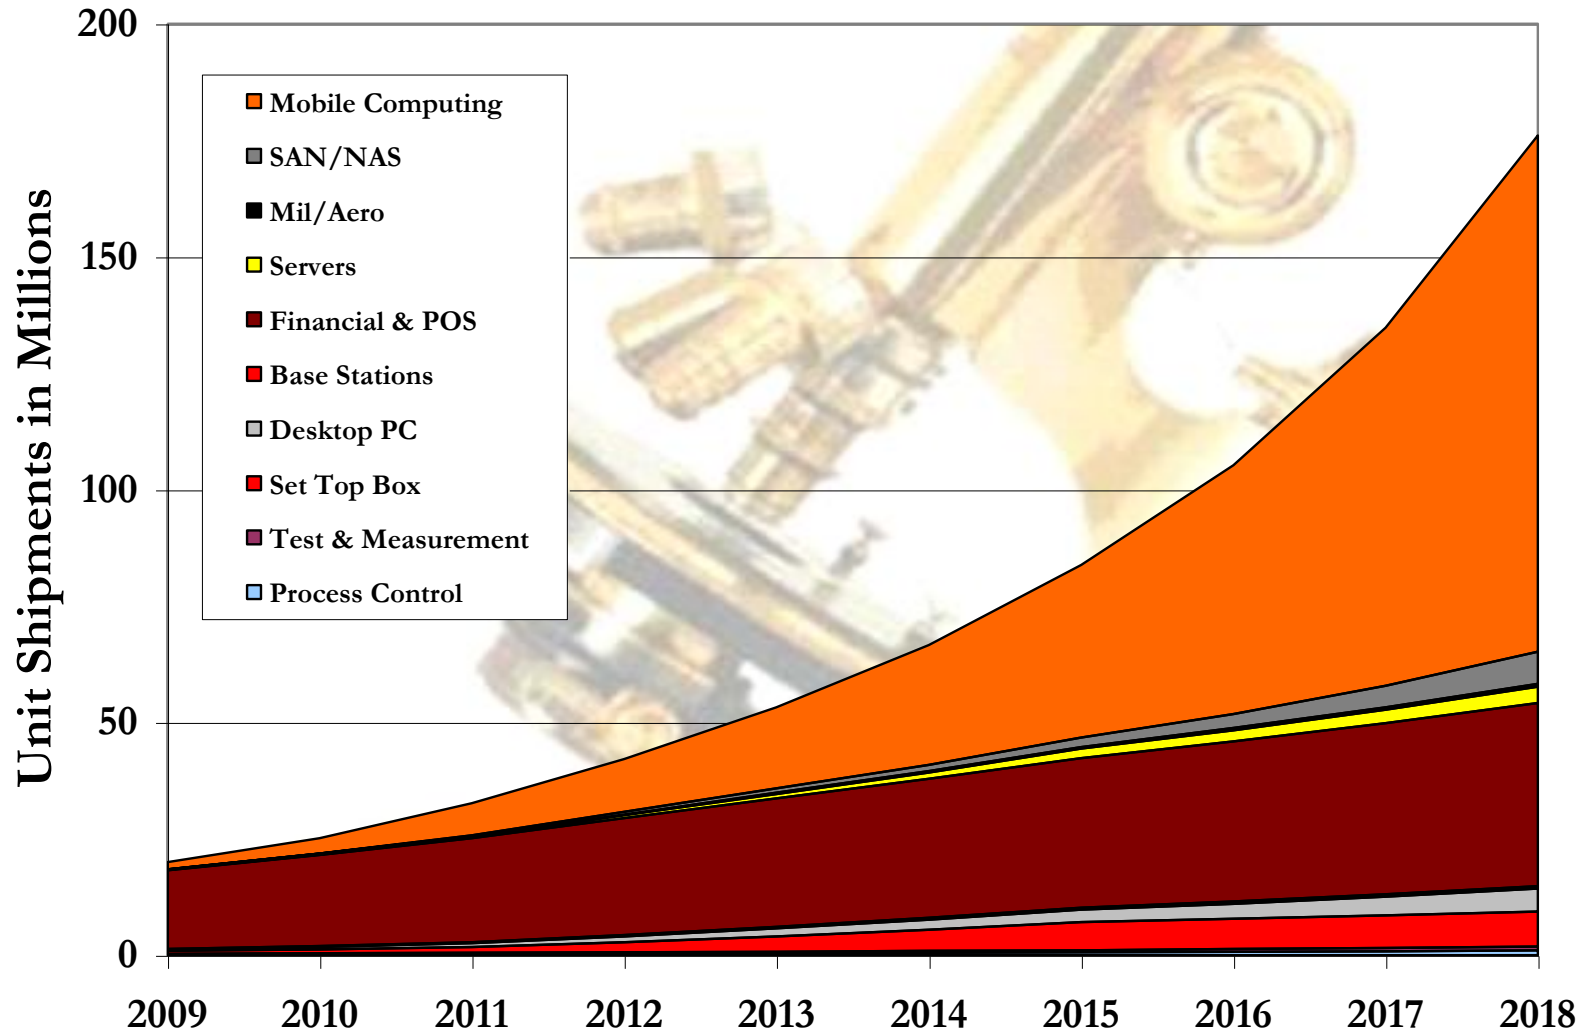
- ◆ Doesn't support standard HDD model of front-panel replacement
- ◆ An undesirable approach in these environments

## ❖ Legacy systems

- ◆ Older designs with IDE but not SATA or PCIe interfaces
- ◆ Slow to move to new form factors
- ◆ Many still use IDE/PATA DOM

Both of these are shrinking in importance

# SSD Unit Shipment Outlook





- M.2 will have broad market appeal
- SSD growth strong in both markets
  - ◆ Enterprise – Approaching 7M units by 2018
  - ◆ Client – 115M units by 2018
  - ◆ Embedded/Other – Over 50M units by 2018
- Good growth + broad appeal = bright outlook!

# M.2 SSD for Mobile Computing

## SNIA WEBCAST

Presented by:  
Jon Tanguy  
Micron Technology, Inc.



HOSTED BY THE  
SOLID STATE STORAGE INITIATIVE



Jon Tanguy  
Sr. Technical  
Marketing Engineer  
Micron Technology, Inc.  
Boise, Idaho  
[www.micron.com](http://www.micron.com)



Jon is a Sr. Technical Marketing Engineer in Micron's Storage Business Unit, which includes SSD and NAND Flash organizations, serving in this capacity for the past five years. Jon facilitates new product integration and customer qualifications for notebook and desktop applications, as well as SSD in the data center.. Jon plays a key role in product planning and development, with an eye toward market requirements.


Jon has more than 20 years of experience in the data storage industry, working with both magnetic media and solid state technologies.

Jon earned his Bachelor of Science degree in Electrical and Computer Engineering from the University of Colorado at Boulder.



# Overview of M.2

- Initially Introduced as “NGFF” for Next Generation Form Factor, soon became “M.2”; initial proposal in SATA-IO and PCI-SIG.
- Detailed M.2 specifications are included in the PCI-SIG M.2 spec; the SATA version of M.2 is described in the SATA v3.2 spec.
- Intended to resolve the extensibility issues with mSATA SSD.
- Brings superior throughput capability to “Ultra thin and light” computing, by leaping past the plateau of 6.0 Gbps SATA.
- Enables 2- or 4-lane transfer speeds ~ 900 MB/s (read) & 800 MB/s (write) for first generation drives. Significantly faster in the x4 options to come!



**mSATA evolves to  
M.2 PCIe Gen2 x2**




- ▶ An M.2 SSD supports up to 4 lanes of PCIe
  - ◆ 2 lanes PCIe 2.0 = 1GB/s
  - ◆ 2 lanes PCIe 3.0 = 2GB/s
  - ◆ 4 lanes PCIe 3.0 = 4GB/s
  
- ▶ Actual performance will be less

- Legacy SATA/AHCI. Allows form factor transition; no driver change for SATA 6 Gbps hosts.
- PCIe/AHCI. Allows backwards compatibility, but can't reach full performance potential.
- PCIe/NVMe. Allows full access to PCI Express lanes, with an interface designed to work most efficiently with new, faster storage technologies, e.g. Flash. Takes advantage of the ability of SSD to execute data transfers in parallel. Requires driver updates.
- Greater flexibility than mSATA in physical dimension.
- Like mSATA, low insertion rating; not intended for hot-plug!

# M.2 Form Factor Options

- ◆ Denoted by a “Type.”
  - ◆ Specifically: 2280, 2260, 2242, 2230. Also, 3030, 3042, 1630, etc.!
- ◆ “Type” designates X-Y dimension: e.g. 2280 = 22mm x 80mm




- ◆ Above is shown an M.2 SATA option.
- ◆ Interface is keyed to denote interface and device type!
  - ◆ Described as a “Socket”, Socket 2 configuration is for SATA or PCIe x2
  - ◆ Socket 3 configuration is for PCIe in a x4 configuration

# M.2 Form Factor Options

## ◆ Height in Z-dimension also has options

S = single-sided; D = double-sided, as below:

Designator	Top side (mm)	Bottom side (mm)
S1	1.2	
S2	1.35	
S3	1.5	
D1	1.2	1.35
D2	1.35	1.35
D3	1.5	1.35





# M.2 Form Factor Options: Connector Keys

- M.2 provides great flexibility for system designers by providing key options to determine device types; storage, WiFi, WiGig, Bluetooth, et. al.
- Card specification can be very specific.
  - ◆ E.G.: “TYPE 2242-D2-B-M”
  - ◆ This indicates a 22x42mm card, double-sided, with keys in the B and M positions (can fit B or M connector!).

Key ID	Pin Location	Interface
A	8-15	2x PCIe x1 / USB 2.0 / I2C / DP x4
B	12-19	PCIe x2/SATA/USB 2.0/USB 3.0/HSIC/SSIC/Audio/UIM/I2C
C	16-23	Reserved for Future Use
D	20-27	Reserved for Future Use
E	24-31	2x PCIe x1 / USB 2.0 / I2C / SDIO / UART / PCM
F	28-35	Future Memory Interface (FMI)
G	39-46	Not Used for M.2; for Custom/Non-Standard Apps
H	43-50	Reserved for Future Use
J	47-54	Reserved for Future Use
K	51-58	Reserved for Future Use
L	55-62	Reserved for Future Use
M	59-66	PCIe x4 / SATA

# M.2 Connection Schemes



# SNIA WEBCAST

**Presented by:**  
**Jaren May**  
**TE Connectivity**



HOSTED BY THE  
SOLID STATE STORAGE INITIATIVE

# Webcast Presenters




Jaren May,  
Global Product Manager  
TE Connectivity  
Taipei, Taiwan  
[www.te.com](http://www.te.com)

Jaren May is the Global Product Manager of Internal Interconnects including the M.2 connector family. Jaren has worked in the electronic components industry for 5 years. He is currently based in Taipei, Taiwan to enhance his knowledge of the Asian business culture while studying Chinese. During his time outside the office you can find him exploring Asia, most recently trekking to Mt. Everest base camp in Nepal.

# Connector Overview

A natural transition from the Mini Card and Half Mini Card to a smaller form factor.

M.2 supports wireless and SSD module cards.



	Stack Height (z)	Card Centerline to PCB	Length (x)	Width (y)	Volume (xyz)
Mini Card	4.0	2.1	9.09	29.9	1087
M.2 (NGFF)	2.25	1.08	8.7	21.9	426
Savings	1.75	1.02	0.39	8.0	659

## Features

- Available in various heights
- 0.5mm pitch with 67 positions
- Designed for both single and double-sided modules
- Available in various keying options for module cards
- Support PCI Express 3.0, USB 3.0 and SATA 3.0

## Benefits

- Wide product offering to meet customers design needs
- Save more than 20% PCB real estate compared to PCIe Minicard
- Reduces connector height by 15%
- Ensures proper mating with various modules
- Supports higher data rates

## Wireless Applications


- Wi-Fi
- WWAN (2G, 3G, 4G)
- Bluetooth (BT)
- WiGig
- GPS
- Global Navigation Satellite Systems (GNSS)
- Near Field Communication (NFC)
- Hybrid Digital Radio (HDR)



# Module Card Nomenclature

## Module Nomenclature

Type    XX    XX    -    XX    -    X    -    X



Key ID	Pin	Interface
A	8-15	2x PCIe x1/USB 2.0/I2C/DP x4
B	12-19	PCIe x2/SATA/USB 2.0/USB 3.0/HSIC/SSIC/Audio/UIM/I2C
C	16-23	Reserved for Future Use
D	20-27	Reserved for Future Use
E	24-31	2x PCIe x1/USB 2.0/I2C/SDIO/UART/PCM
F	28-35	Future Memory Interface (FMI)
G	39-46	Generic (Not used for M.2)
H	43-50	Reserved for Future Use
J	47-54	Reserved for Future Use
K	51-58	Reserved for Future Use
L	55-62	Reserved for Future Use
M	59-66	PCIe x4/SATA


Note: M.2 Connectors have only one key even though module cards could have two keys

# Key to Application Use

Key	Intended Use	Host Interfaces	Applications	Module Card Type
A	Connectivity Version 1-DP	2x PCIe x1/USB 2.0/I2C/ DP x4	Wireless Connectivity devices including combinations of Wi-Fi, BT, HFC, and/or WiGig.	1630, 2230, 3030
B	WWAN/SSD/ Others Primary Key	PCIe x2/SATA/USB 2.0/ USB 3.0/ HSIC/SSIC / Audio/UIM/ I2C	WWAN+GNSS, or SSD	3042, 2230, 2242, 2260, 2280, 22110
E	Connectivity Version 1-SD	2x PCIe x1/USB 2.0/ IC2/SDIO/UART/PCM	Wireless Connectivity devices including combinations of Wi-Fi, BT, NFC, and/ or GNSS.	1630, 2230, 3030
M	SSD 4 Lane PCIe	PCIe x4/ SATA	SSD devices and Host I/Fs supported are PCIe with up to four lanes or SATA.	2242, 2260, 2280, 22110

# Single vs. Double Sided Modules


## M.2 (NGFF) Single vs Double Sided Modules



The difference between single and dual sided modules is functionality. All soldered down type module cards are single sided and do not have a pin layout. Single sided modules are intended to be used in low profile applications.

There are two types of module cards, soldered down and connectorized

# Module Insertion



- Step 1:** Move the module against the housing's chamber
- Step 2:** Rotate the module to 25 degrees and insert it until the bottom of the module surface reaches the ramp
- Step 3:** Rotate the module to horizontal position by hand
- Step 4:** Fix the module with a PCB screw to secure the module

Please refer to the application specification for a complete description of both the insertion and removal method for modules.




# Standard Products

Base PN	Max. Height	Nominal Height Dim. A	Dim. B	Dim. C	Double Sided Modules	Key	Gold Plating		
							Gold Flash	15u"	30u"
2199125-x	2.25	2.15	1.08	0.63	No	A	-7	-8	-9
						B	-1	-3	-5
						E	-2	-4	-6
						M	-13	-14	-15
2199119-x	3.2	3.1	1.94	1.48	Yes	A	-7	-8	-9
						B	-1	-3	-5
						E	-2	-4	-6
						M	-13	-15	-16
2199230-x	4.2	4.1	2.94	2.48	Yes	A	-7	-8	-9
						B	-1	-3	-5
						E	-2	-4	-6
						M	-13	-15	-16

Dim. A: Max. connector height

Dim. B: Card Centerline to PCB

Dim. C: Bottom of Card to PCB



**NVM Express**



**SNIA** WEBCAST

**Presented by:  
David Akerson  
Intel Corporation**



**HOSTED BY THE  
SOLID STATE STORAGE INITIATIVE**



David Akerson,  
Strategic Marketing  
Engineer  
Intel Corporation  
[www.intel.com](http://www.intel.com)

David Akerson has worked in high tech for 20 years for such notable companies as Intel, Tektronix and Sequent Computer Systems.


David has extensive experience in Hardware and Software Product Management and Marketing.

In addition to his position at Intel, David works closely with NVM Express.

# What is NVM Express™?

➤ NVM Express (NVMe™) is a standardized high performance host controller interface for PCIe\* Storage, such as M.2 PCIe SSDs

- ◆ Standardizes register set, feature set, and command set
- ◆ Architected from the ground up for NAND and next generation NVM
- ◆ The specification was developed by an open industry consortium for Client and Enterprise platforms



# Why NVM Express™?

## ➤ High Performance


- ◆ Low latency (no HBA overhead or protocol translation)...6x latency reduction enables path to next generation NVM
- ◆ Full duplex, multiple outstanding requests, out of order processing

## ➤ Low System Cost

- ◆ Direct attach to CPU subsystem eliminates HBA cost

## ➤ Power Efficient

- ◆ Direct attach to CPU subsystem eliminates HBA power
- ◆ Link power management (dynamic link width/speed)

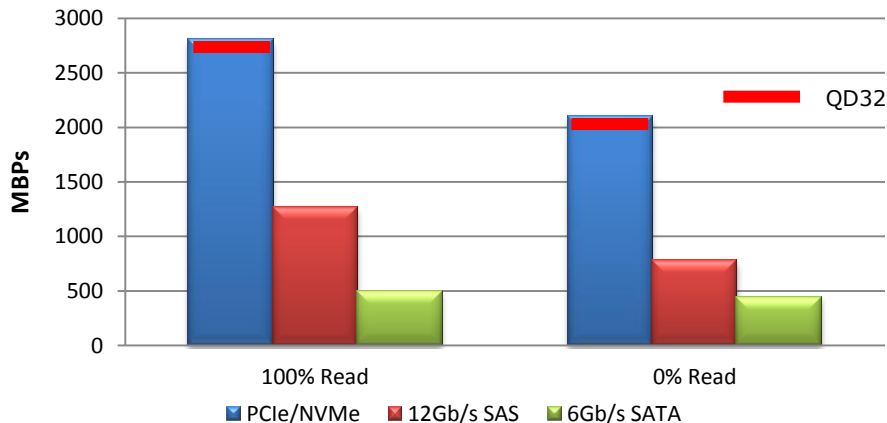


# Why NVM Express? Performance Leadership

## NVM Express delivers versus leadership SAS/SATA products

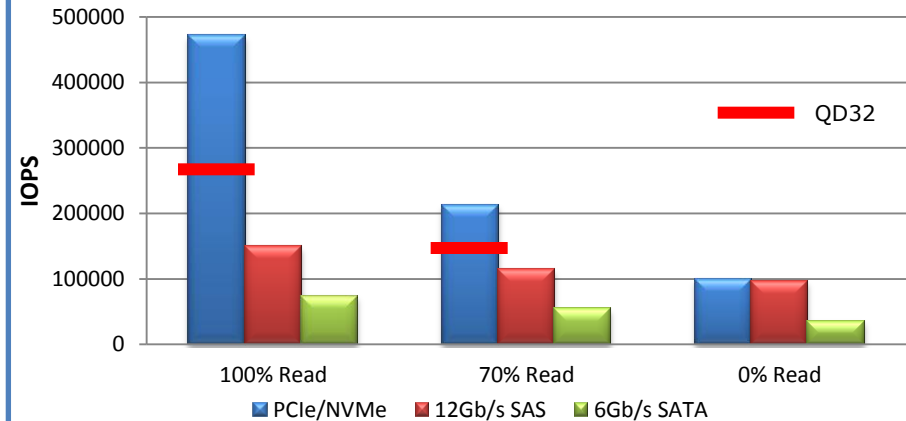
- ◆ Random Workloads
  - › > 2X performance of SAS 12Gbps
  - › 4-6X performance of SATA 6Gbps

**Sequential Workloads, QD = 128**



PCIe/NVMe 128K @ QD 32/128 ; SATA 128K @ QD32; SAS 64K @ QD32

**4K Random Workloads, QD = 128**



PCIe/NVMe @ QD 32/128 ; SAS/SATA @ QD32

## For sequential workloads, realize close to 3 GB/s reads

- ◆ Sequential Workloads
  - › > 2X performance of SAS 12Gbps
  - › > 4X performance of SATA 6Gbps

# Driver Development on Major OS'

## Windows\*

- Windows\* 8.1 and Windows\* Server 2012 R2 include inbox driver
- Open source driver in collaboration with OFA

## Linux\*

- Native OS driver since Linux\* 3.3 (Jan 2012)

## Unix

- FreeBSD driver released

## Solaris\*

- Delivered to S12 and S11 Update2
- Compliant with 1.0e

## VMware\*

- vmklinux driver targeted for Q2 '14 release

## UEFI

- Open source driver available on SourceForge

***Native OS drivers already available, with more coming!***

# NVM Express™: Unlock the Potential



Cloud & Big Data Computing



Financial Analysis



Datacenter Refresh

▶ Performance for Client, Enterprise and Data Center applications



Digital Content



Web 2.0



Gaming



Energy Exploration



# M.2 Performance



# SNIA WEBCAST

**Presented by:**  
**Eden Kim**  
**Calypso Systems**



HOSTED BY THE  
SOLID STATE STORAGE INITIATIVE



Eden Kim,  
CEO

Calypso Systems, Inc.

[www.calypsotesters.com](http://www.calypsotesters.com)

Eden, CEO of Calypso Systems, Inc., is Chair of the SNIA Solid State Storage Technical Working Group (SSS TWG) which has published the industry standard solid state storage Performance Test Specification (PTS) for both client and enterprise SSD performance testing. The PTS sets forth standardized test methodologies for drive preparation, test and reporting.

Eden is also Chair of the SNIA SSSI PCIe SSD and TechDev Committees and is a member of the SNIA SSSI Governing Board.

Eden has authored several white papers and presented talks at industry events on the PTS, SSD test methodologies, PCIe solid state storage, SSD Performance evaluation and Enterprise Synthetic Workload testing.

Calypso Systems is an SSD test house that provide test equipment and test services for SSD Performance testing. Calypso is also the developer and manufacturer of the SNIA PTS Reference Test Platform SSD tester.

# Summary Performance Data By Storage Class

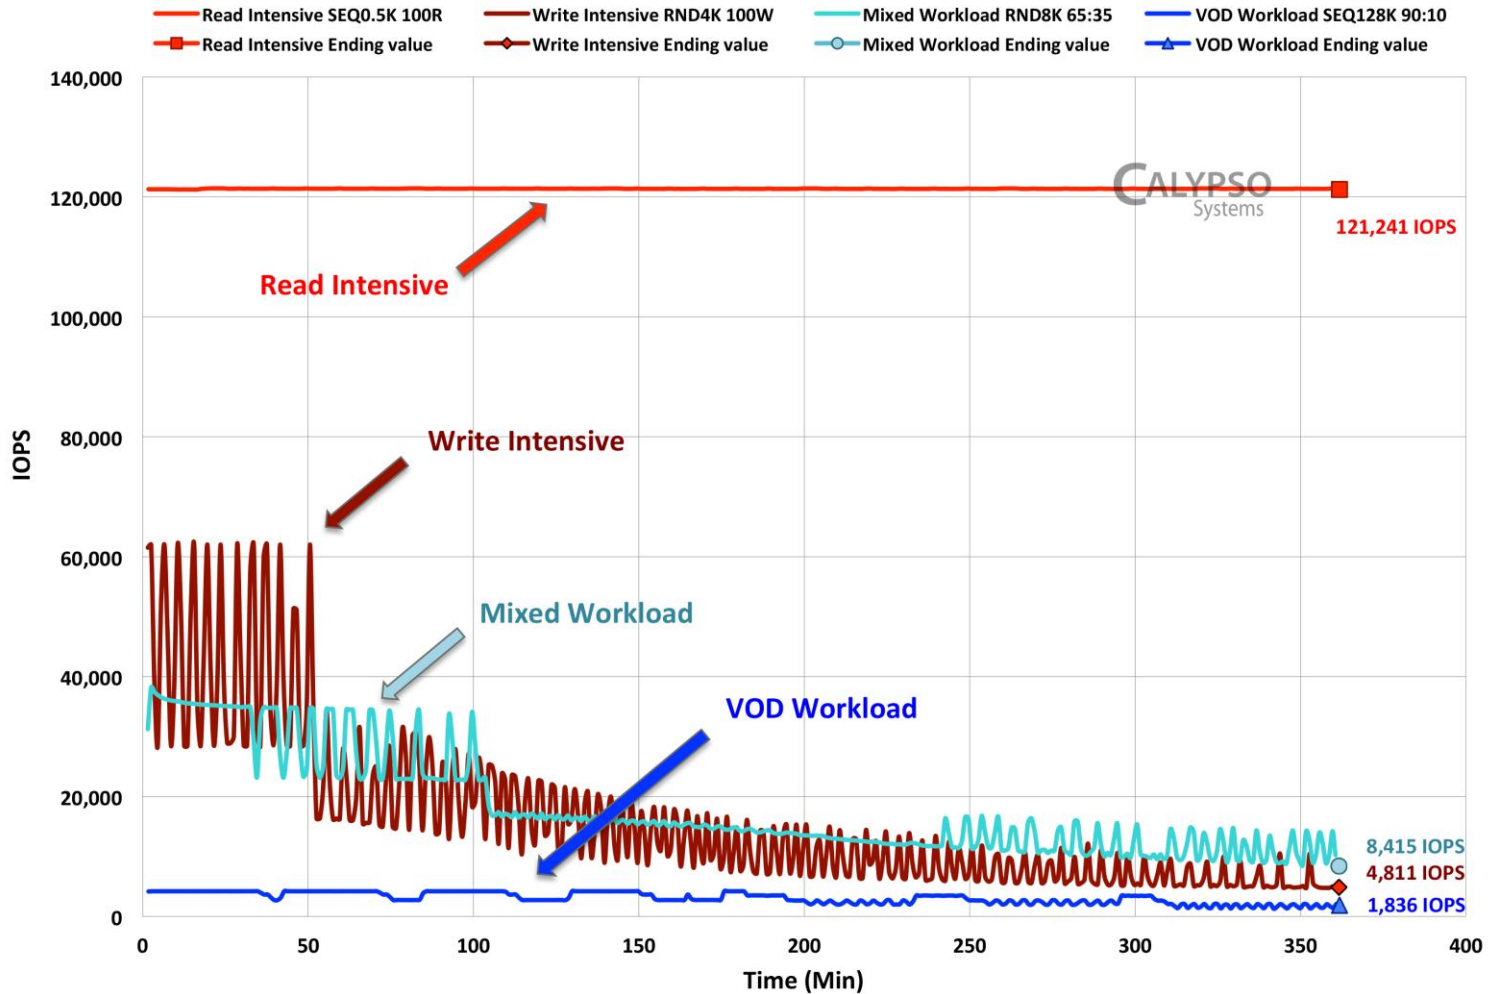
## Summary Performance Data by Storage Class

Storage Class		Type	FOB IOPS	IOPS (higher is better)			Throughput (larger is better)		Response Time (lower is better)	
Type	Device	Capacity & Write Cache	RND 4KiB 100% W	RND 4KiB 100% W	RND 4KiB 65:35 RW	RND 4KiB 100% R	SEQ 1024KiB 100% W	SEQ 1024KiB 100% R	RND 4KiB 100% W Ave	RND 4KiB 100% W Max
HDD & SSHD										
SSHD 1	7,200 RPM 2.5" SATA Hybrid	500 GB WCD	109	93	107	127	98 MB/s	102 MB/s	20.99 msec	138.62 msec
SAS HDD 2	15,000 RPM 3.5" SAS HDD	80 GB WCD	350	340	398	401	84 MB/s	90 MB/s	5.39 msec	97.28 msec
Client mSATA, M.2, SATA SSDs										
mSATA SSD 3	mSATA card 1.8" SSD	128 GB WCE	45,743	1,359	1,926	37,936	187 MB/s	533 MB/s	0.74 msec	543.41 msec
<b>M.2 SSD 4</b>	<b>M.2 x2 2280 SSD</b>	<b>512 GB WCE</b>	<b>61,506</b>	<b>4,185</b>	<b>9,532</b>	<b>71,282</b>	<b>455 MB/s</b>	<b>535 MB/s</b>	<b>0.29 msec</b>	<b>24.99 msec</b>
SATA SSD 5	SATA3 2.5" SSD	240 GB WCE	42,483	18,136	25,762	42,255	219 MB/s	478 MB/s	0.05 msec	22.60 msec
Enterprise SATA & SAS SSDs										
SATA SSD 6	SATA 6Gb/s 2.5" SSD	200 GB WCD	54,788	33,583	50,708	63,640	367 MB/s	480 MB/s	0.06 msec	11.95 msec
SAS SSD 7	SAS 6Gb/s 2.5" SSD	800 GB WCD	71,878	56,704	89,154	110,989	519 MB/s	536 MB/s	0.05 msec	3.88 msec

# M.2 Benchmark Workloads

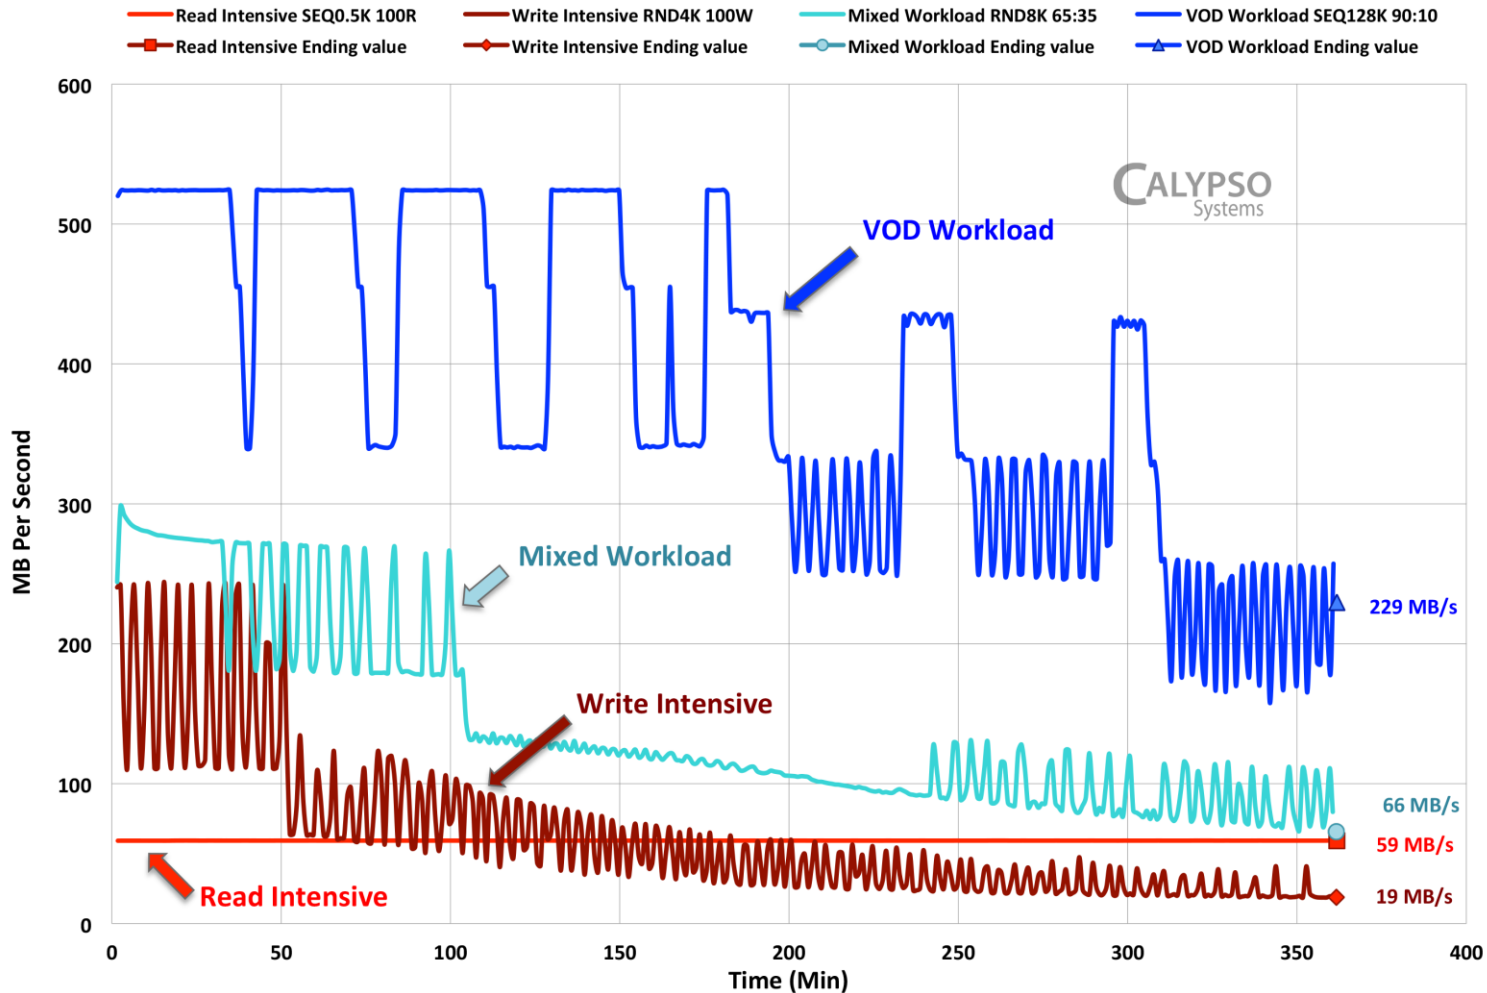
## IOPS versus Time

### M.2 x2 Benchmark Workloads - IOPS vs Time



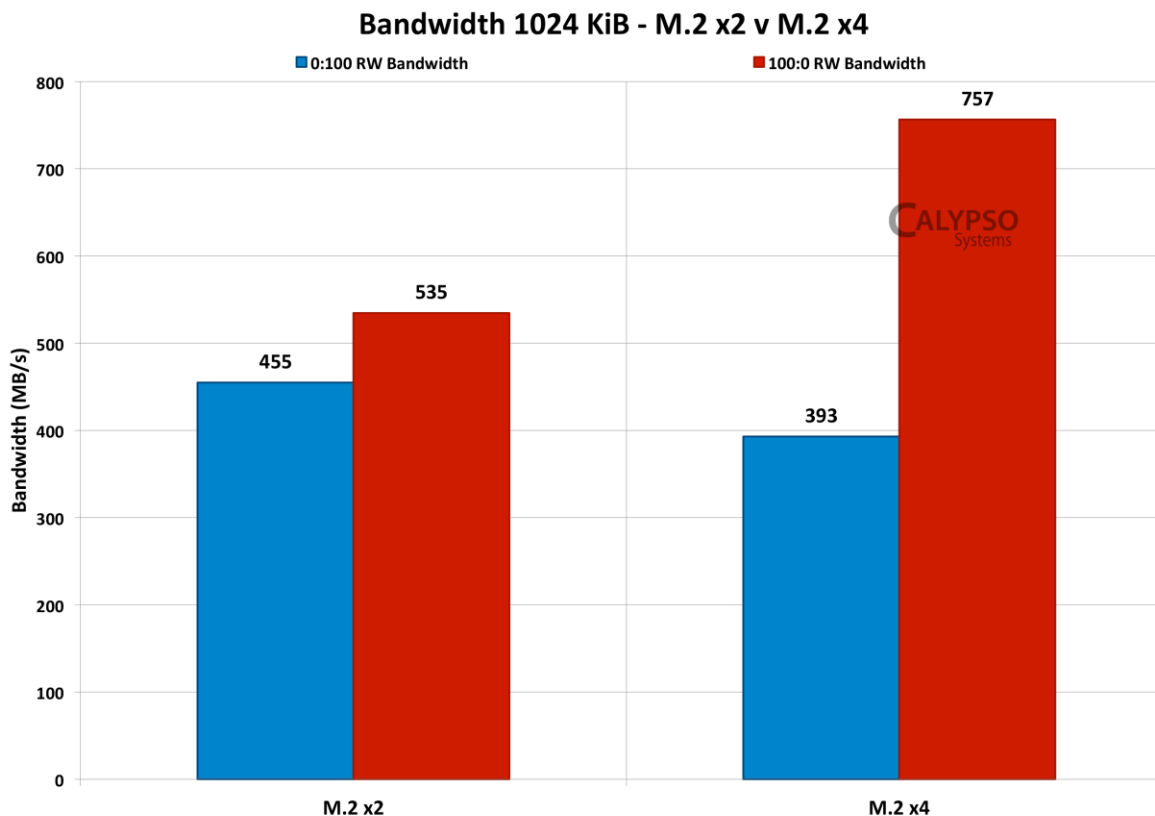
# M.2 Benchmark Workloads Bandwidth versus Time

## M.2 x2 Benchmark Workloads - Bandwidth vs Time



# M.2 Throughput 2 Lanes versus 4 Lanes

- Throughput (Bandwidth) shown for M.2 x2 vs M.2 x4 devices
- Bandwidth is in MB/s - shown is 100% Writes and 100% Reads
- M.2 x4 shows higher bandwidth, taking advantage of additional lanes



- **Performance Comparison by Storage Class Chart**
  - ◆ SSDs clearly have higher IOPS & Throughput and lower Response Times than HDDs and SSHDs
  - ◆ Different classes of SSDs have different performance ranges
  - ◆ SSDs can be designed and optimized for different workloads - hence one or more metrics may be enhanced or emphasized
  - ◆ Performance & cost generally increase as you go down the chart
- **M.2 Performance - IOPS v Time**
  - ◆ WSAT tests are a corner case test where a continuous workload is applied to the drive after a device PURGE
  - ◆ There is no workload independent pre-conditioning
  - ◆ The test workload itself is the pre-conditioning (i.e. workload dependent pre-conditioning)
  - ◆ 4 Corner case workloads are Write Intensive, Read Intensive, Mixed and Video-On-Demand workloads

## ➤ M.2 Performance - Bandwidth v Time

- ◆ Bandwidth measures the transfer rate of each workload in MB/sec
- ◆ Larger block size workloads generally show larger Bandwidth measurements
- ◆ Smaller block size workloads generally show smaller Bandwidth measurements
- ◆ Bandwidth often saturates the cache, resulting in dips in performance
- ◆ Mixed workloads can show oscillating performance depending on the access being performed



**Wrap-up**



**SNIA** WEBCAST

**Presented by:  
Tom Coughlin  
Coughlin Associates**



**HOSTED BY THE  
SOLID STATE STORAGE INITIATIVE**

## ➤ M.2 to Benefit from SSD's Bright Future

- ◆ Enterprise, Client, and Embedded/Other SSDs approaching 7 M, 115 M and over 50 M units respectively by 2018

## ➤ M.2 Cards

- ◆ NGFF renamed M.2 by PCI-SIG
- ◆ Allows extension beyond SATA 6 Gb/s and limits of mSATA
- ◆ Supports up to 4 lanes of PCIe at up to 1 GB/s per lane
- ◆ NVMe enables maximum PCIe performance
- ◆ Various card lengths from 30 to 110 mm, one-sided z-height of 1.2 to 1.5 mm

## ➤ M.2 Connection

- ◆ M.2 supports a variety of functions including wireless and SSD cards with SATA & PCIe
- ◆ Single and double-sided modules (soldered down and connectorized)

## ➤ NVM Express

- ◆ A standardized high performance host controller interface for PCIe storage (SAS, SATA, and PCIe)
- ◆ NVM Express was designed to support this and next generation Non-Volatile Memory in a multiple form factors such as M.2, 2.5" and Add-in Cards
- ◆ Faster random and sequential workloads possible with NVM Express
- ◆ Drivers available for major operating systems and client, enterprise and data center applications

## ➤ M.2 Performance

- ◆ Detailed summary performance shown for HDD, Hybrid HDDs and SSDs (client and enterprise)
- ◆ Most SSDs faster than HDDs, even hybrid HDDs
- ◆ Most SSDs better on throughput and latency than HDDs
- ◆ M.2 cards should provide better performance characteristics than mSATA cards
- ◆ Storage performance depends upon workloads

# Questions?

To replay this webcast, go to  
[www.snia.org/forums/sssi/knowledge/education](http://www.snia.org/forums/sssi/knowledge/education)

Scroll down to **Webcasts**