



---

# **Enterprise Flash Storage Demands Faster Interfaces**

Storage Valley Supper Club  
August 2015



# About Demartek


- ◆ Industry Analysis and ISO 17025 accredited test lab
- ◆ Lab includes enterprise servers, networking & storage (DAS, NAS, SAN, 10GbE, 40GbE, 16GFC)
- ◆ We prefer to run real-world applications to test servers and storage solutions (databases, Hadoop, etc.)
- ◆ Demartek is an EPA-recognized test lab for ENERGY STAR Data Center Storage testing
- ◆ Website: [www.demartek.com/TestLab](http://www.demartek.com/TestLab)



# Flash Storage


---

- ◆ Obviously faster than HDD storage
- ◆ Can drive up CPU utilization
  - ◆ Can alter your virtual-to-physical server ratio
- ◆ Can expose new bottlenecks such as network bottlenecks
- ◆ Flash storage and high speed networks were made for each other
- ◆ What new enterprise datacenter interfaces are coming that can take advantage of the speed of flash storage?



# Ethernet Today


- ◆ IEEE 802.3ba (40GigE & 100GigE) ratified June 2010
- ◆ The fastest Ethernet cables and connectors today are 10 Gbps per lane or channel
- ◆ Higher speeds today are achieved by bundling
  - ◆ 40GigE today = 4 x 10 Gbps together
  - ◆ 100GigE today = 10 x 10 Gbps together
- ◆ 40 Gbps NICs require PCIe 3.0 x8 or x16 slot in the server



# Ethernet Future

- ◆ 25 Gbps connectors will soon be available
  - ◆ 25Gbps over a single lane
  - ◆ 50GigE future = 2 x 25 Gbps together
  - ◆ 100GigE future = 4 x 25 Gbps together
  - ◆ 250GigE future = 10 x 25 Gbps together
- ◆ 25G Ethernet Consortium Announcement – July 1, 2014
  - ◆ [www.25GEthernet.org](http://www.25GEthernet.org)
- ◆ IEEE P802.3by 25 Gb/s Ethernet Task Force
  - ◆ <http://www.ieee802.org/3/by/index.html>

# Ethernet Public Roadmap



Development of four new speeds began in 2014:  
2.5 GbE, 5 GbE, 25 GbE, 400 GbE

<http://www.ethernetalliance.org/roadmap/>

# Fibre Channel

- ◆ **16GFC is current specification, first shipped in 2011**
  - ◆ Some 16GFC HBAs can function as 10 Gb NICs
- ◆ **In February 2014, “Gen 6” Fibre Channel was announced**
- ◆ **32 Gbps single-lane connection (“32GFC”)**
  - ◆ OM4 fiber-optic expected cable distance: 100m
- ◆ **128 Gbps parallel connection (4 x 32, “128GFCp”)**
- ◆ **Backward compatible with 16GFC and 8GFC**
- ◆ **32GFC products expected to be available in 2016**

# SAS

---

- ◆ Used for enterprise SSDs and HDDs and for connections to JBOD shelves and external array shelves
- ◆ Current specification is 12 Gbps SAS
- ◆ SCSI command protocol used in FC, FCoE, iSCSI and SAS
- ◆ Supports up to 16K devices on single “fabric”
- ◆ Roadmap for 24 Gbps SAS, at or near the same time as PCIe 4.0, approximately 2017





# SATA

---

- ◆ **Device (drive) types**
  - ◆ Common for client (consumer) SSDs & HDDs
  - ◆ Sometimes used for enterprise SSDs & HDDs
- ◆ **Mostly for inside the case connections**
  - ◆ eSATA allows for short external distances
- ◆ **SATA is point-to-point, single device per cable or connector**
- ◆ **Traditional SATA has no roadmap beyond 6 Gbps**
  - ◆ Some new enterprise features planned
  - ◆ Unclear if SATA Express will gain acceptance in the market

# Thunderbolt 3

- ◆ Thunderbolt 3 announced in June 2015
  - ◆ Increases speed to 40 Gbps
- ◆ Target audience is media creators and editors who use premium laptops, desktops, workstations and peripherals that connect to them.
  - ◆ Includes storage devices, especially SSDs
- ◆ Uses a USB type-C cable
  - ◆ Supports USB 3.1 (10 Gbps), Display Port (dual 4K displays), four lanes of PCI Express 3.0
  - ◆ Provides 15 watts for bus-powered devices and supports USB power delivery to charge laptop computers up to 100 watts.

# NVMe

---

- ◆ Scalable host controller interface designed for enterprise and client systems that use PCI Express SSDs
  - ◆ Designed with Flash memory and technologies coming after Flash memory in mind (non-volatile memory)
  - ◆ Much faster (lower latency) software stack than existing storage stacks such as SAS and SATA
- ◆ In-box drivers for Windows and Linux now
- ◆ Faster individual devices than other interfaces
  - ◆ PCIe card and drive form factor (SFF-8639 → U.2)
  - ◆ Not as well-established, but ramping up quickly

# Performance Results

---

- ◆ View my presentation notes from Tuesday's FMS session 104-C: "How Flash-Based Storage Performs on Real Applications"
  - ◆ [www.demartek.com/FlashMem](http://www.demartek.com/FlashMem)
- ◆ Demartek SSD Zone reports
  - ◆ [www.demartek.com/SSD](http://www.demartek.com/SSD)

# Storage Interface Comparison

- ◆ View the Demartek Storage Interface Comparison reference page
  - ◆ Search for “*Storage Interface Comparison*” in your favorite search engine



# Thank You!

---



Demartek public projects and materials are announced on a variety of social media outlets. Follow us on any of the above.



Sign-up for the Demartek monthly newsletter, *Demartek Lab Notes*.  
[www.demartek.com/newsletter](http://www.demartek.com/newsletter)