APRIL 2-4, 2013

COMPUTERWORLD SNIA

SNW

Intelligent Architecture for the Data-Driven Business

Analyst Perspective:
Test Lab Report – 16 Gb Fibre Channel Performance and Recommendations

Dennis Martin
President, Demartek

Demartek®

The original version of this presentation is available here:
http://www.demartek.com/Demartek_Presenting_SNWUSA_2013-04.html

Agenda

- About Demartek
- Fibre Channel – 16 Gigabit and futures
- Cabling considerations and recommendations
- Demartek performance lab test results
- Demartek free resources

About Demartek

- Industry analysis with on-site test lab
- Lab includes servers, networking and storage infrastructure
  - Fibre Channel – 4, 8 & 16 Gbps
  - Ethernet – 1 & 10 Gbps: NFS, SMB (CIFS), iSCSI & FCoE
  - Servers – 8+ cores, large RAM
  - Virtualization – VMware, Hyper-V, Xen
- We prefer to run real-world applications to test servers and storage solutions
  - Currently testing SSD, 10GbE, 16GFC and other technologies
- Website: www.demartek.com

Fibre Channel

► Marketplace*

- 2012 was the year of 10-10-10
  - 10 Million Fibre Channel ports (switches and adapters)
  - $10 Billion of Fibre Channel Enterprise Storage Systems
  - 10 Exabytes (EB) of external storage shipped with Fibre Channel
- Datacenter standard for storage area networks
  - Vast majority of enterprise storage uses Fibre Channel as the host interface

* Source: Fibre Channel Industry Association (FCIA)

Fibre Channel

► SAN Interface

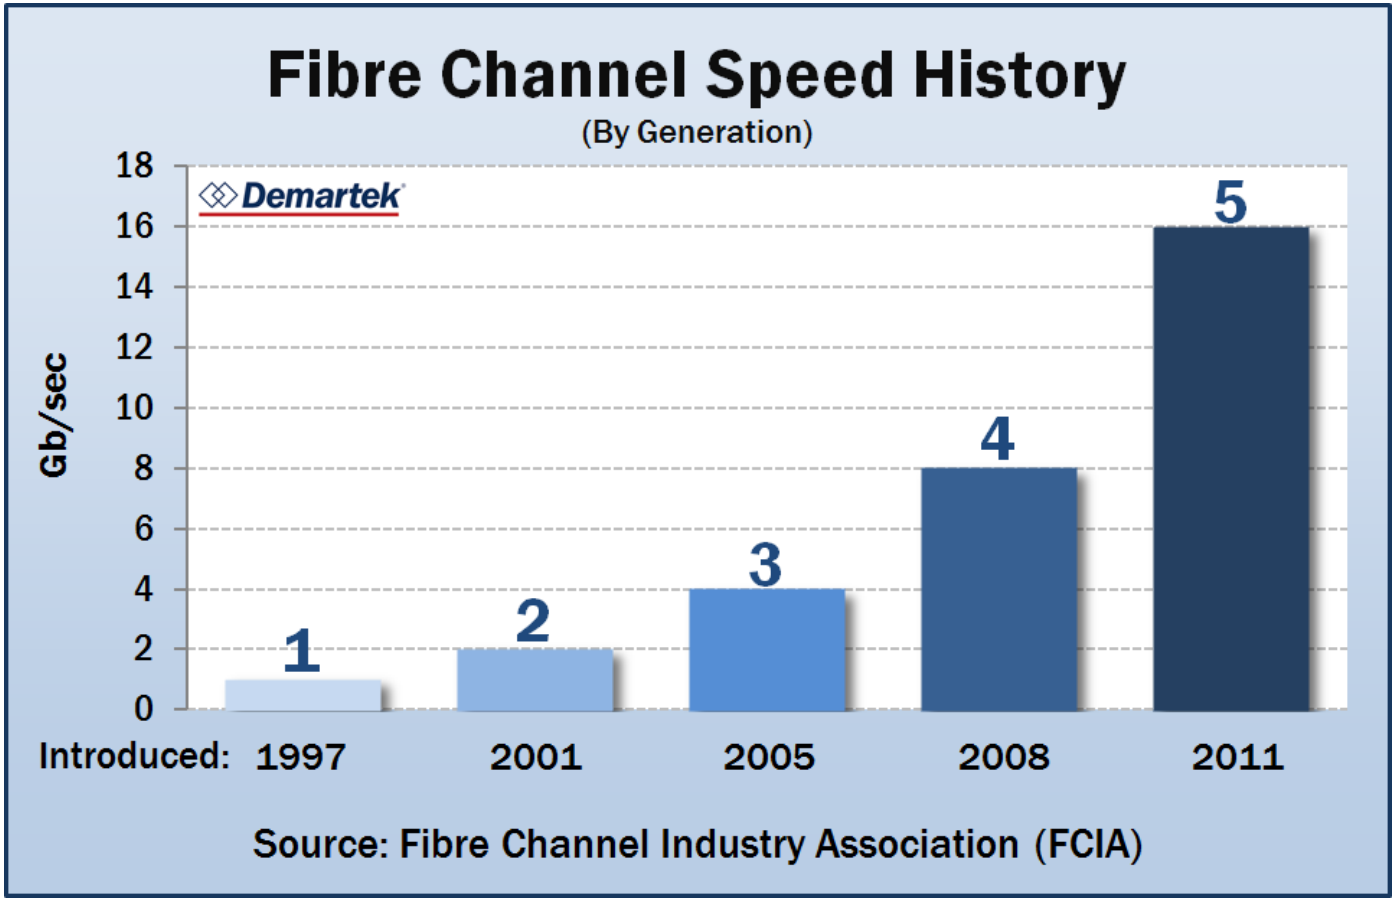

Fibre Channel

Doubles in speed every 3-4 years

FCoE (10 Gb) introduced in 2009

Fibre Channel

► Disk Drive Interface

- 4 Gb/s was the highest speed for FC interface disk drives

- HDD/SSD vendors have moved to 6 Gb/s SAS for enterprise drives
  - First 12 Gb/s SAS drives announced in Spring 2012

Fibre Channel

► 16 Gigabit (16GFC)

- 16 GFC is backward compatible with 4 GFC & 8 GFC
- Uses 14 Gbps single-lane connectors
  - Doubles speed of 8 GFC due to newer 64b/66b encoding
- The first 16 GFC switches and HBAs shipped in 2011
  - Some of these HBAs can also function as 10 Gb NICs
- 16 GFC storage targets becoming available
- Fibre Channel speeds and server slots (dual-port)
  - 4 Gb: PCI-X 2.0, PCIe 1.0
  - 8 Gb: PCIe 2.0 x4 or PCIe 1.0 x8
  - 16 Gb: PCIe 3.0 x4 or PCIe 2.0 x8

Fibre Channel

► 32 Gigabit and 64 Gigabit

- Formal statement of direction:

  "The INCITS Technical Committee T11 is currently working on the 32 GFC Fibre Channel specifications. The 32 GFC specifications are going to letter ballot in April and should be complete this year. TheT11 committee is also investigating a multi-lane 128 GFC interface that is based on the 32GFC work. Work has not yet begun in T11 for developing the 64 GFC specifications, but 64 GFC is on the FCIA Speed roadmap."

  Steve Wilson, Director of Technology and Standards, Brocade and INCITS Technical Committee T11 Chairman

- 32 GFC will use 28 Gbps connectors (25/28G), and will double the speed of 16 GFC

Network Virtualization

► NPIV and Virtual Fibre Channel

- Available with Windows Server 2012 Hyper-V
  - Similar in concept to SR-IOV but for Fibre Channel
  - Supported by most FC HBAs (requires NPIV support)
    - NPIV is enabled by default in some FC HBAs, not in others
  - Requires support by the guest O.S. (Windows 2008R2/2012)
- Procedure
  - Assign a pair of virtual WWPNs to a guest O.S.
  - Add the WWPNs to the zoning and storage LUN masking
- Benefits
  - VMs have their own virtual FC HBAs, just like physical servers
  - VMs can be moved and take their FC storage with them

Fibre Channel Adapter Specifications

- Fibre Channel can run in full-duplex mode, but storage protocols generally operate in half-duplex mode
  - Throughput numbers below are half-duplex (one-way)
- Host Adapter Requirements below are for dual-port cards

Demartek	Throughput (MBps)	Encoding	Line Rate (Gbaud)	Host Adapter Requirements
1GFC	100	8b/10b	1.0625	PCI-X
2GFC	200	8b/10b	2.125	PCI-X
4GFC	400	8b/10b	4.25	PCI-X 2.0 or PCIe 1.0 (x4)
8GFC	800	8b/10b	8.5	PCIe 1.0 (x8) or PCIe 2.0 (x4)
16GFC	1600	64b/66b	14.025	PCIe 2.0 (x8) or PCIe 3.0 (x4)

Encoding Schemes

- 8b/10b
  - For every 8 bits, adds 2 bits for command and control
  - 20% overhead = (10-8)/10
- 64b/66b
  - Used by 10 GigE and 16 GFC
  - For every 64 bits, adds 2 bits for command and control
  - 3% overhead = (66-64)/66
- 128b/130b
  - Used by PCIe 3.0
  - For every 128 bits, adds 2 bits for command and control
  - 1.5% overhead = (130-128)/130

Cabling Recommendations

► Fiber Optic Cables

- Fiber optic cabling service life – 15 to 20 years
- ***Recommendation*** – OM4 cables for current and future
  - OM4 will support 40/100 GigE and higher speeds of FC

Demartek	OM1	OM2	OM3	OM4
Jacket color	Orange	Orange	Aqua	Aqua
1 Gb/s	300m	500m	860m	–
2 Gb/s	150m	300m	500m	–
4 Gb/s	70m	150m	380m	400m
8 Gb/s	21m	50m	150m	190m
10 Gb/s	33m	82m	Up to 300m	Up to 400m
16 Gb/s	15m	35m	100m	125m

Connectors

Single-lane – SFP, SFP+

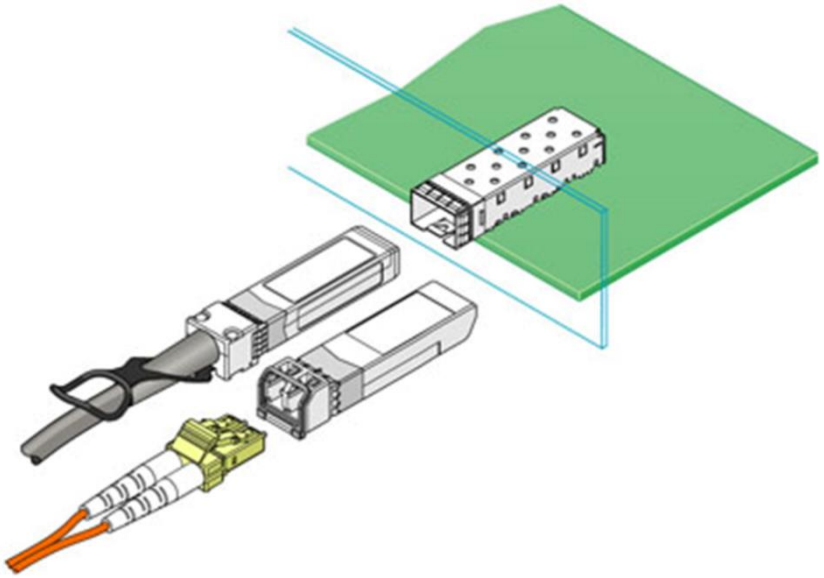

Four-lane – QSFP, QSFP+

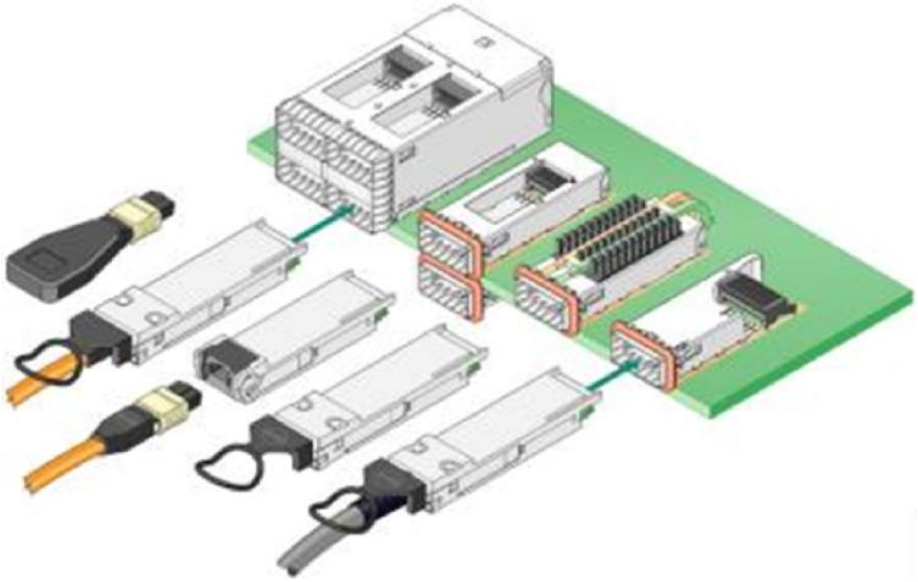

Demartek	SFP	SFP+	QSFP+
Ethernet	1GbE	10GbE	40GbE
Fibre Channel	1GFC, 2GFC, 4GFC	8GFC, 16GFC	–
Infiniband	–	–	QDR, FDR

Performance Results

Performance Results

► 16 GFC IOPS Test – Configuration

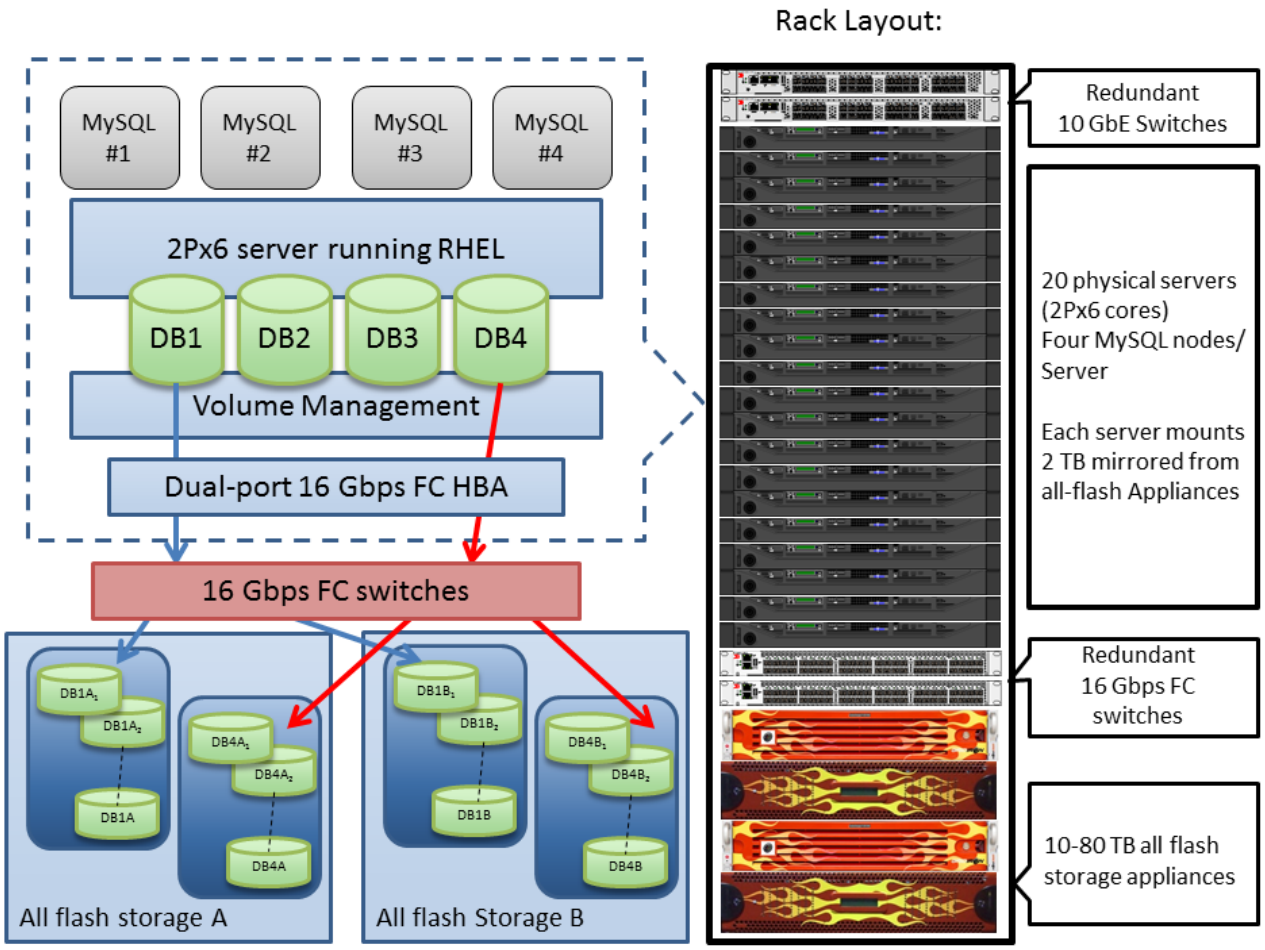

Performance Results

► 16 GFC IOPS Test – 1 Million IOPS

- **MySQL Clusters**
  - 10 servers, 4 instances of MySQL on each server
  - Databases – 256 GB each instance
  - 8 mirrored storage volumes, all flash

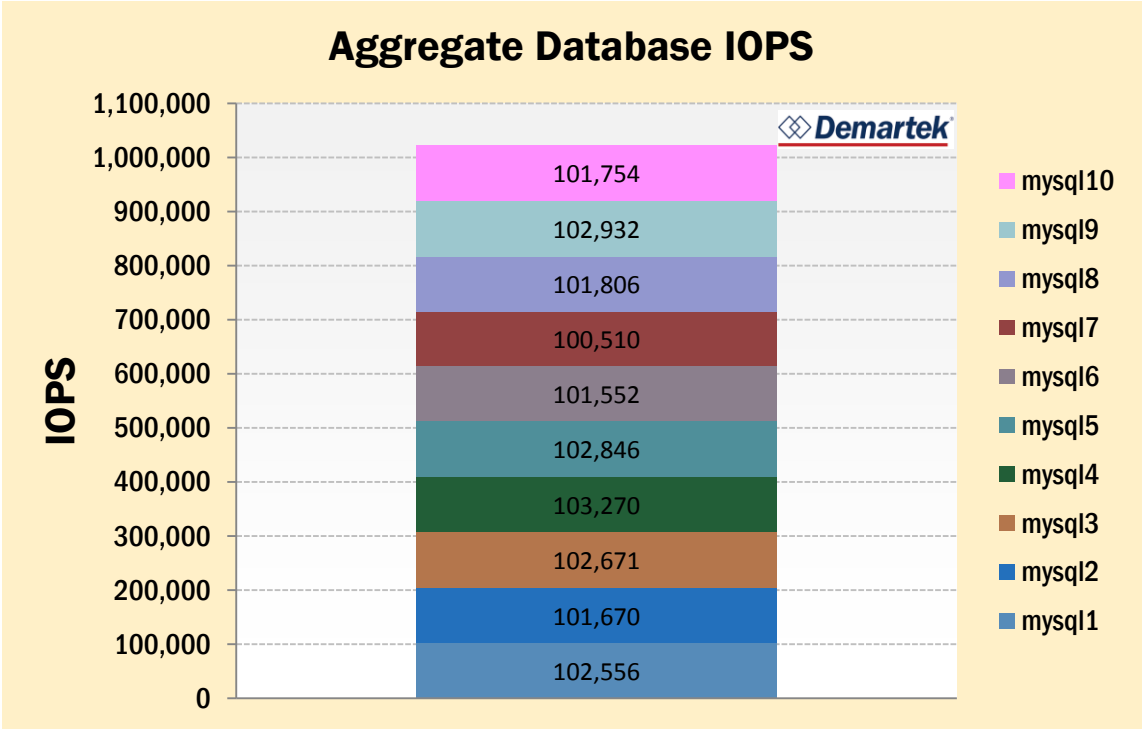

Source: Demartek Fibre Channel Zone – http://www.demartek.com/FC

Performance Results

► 16 GFC Bandwidth

Test – Configuration

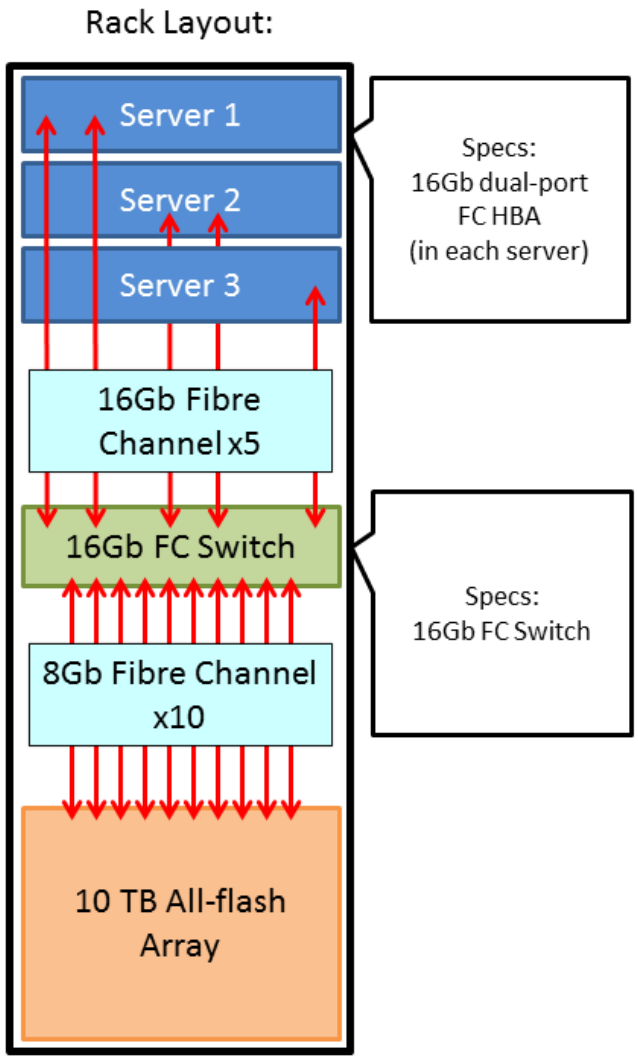

Performance Results

► 16 GFC Bandwidth Test – 7200 MB/sec

- Oracle RAC single large database cluster
  - "Select count(*)" all 10.2 billion rows completed in < 8 minutes

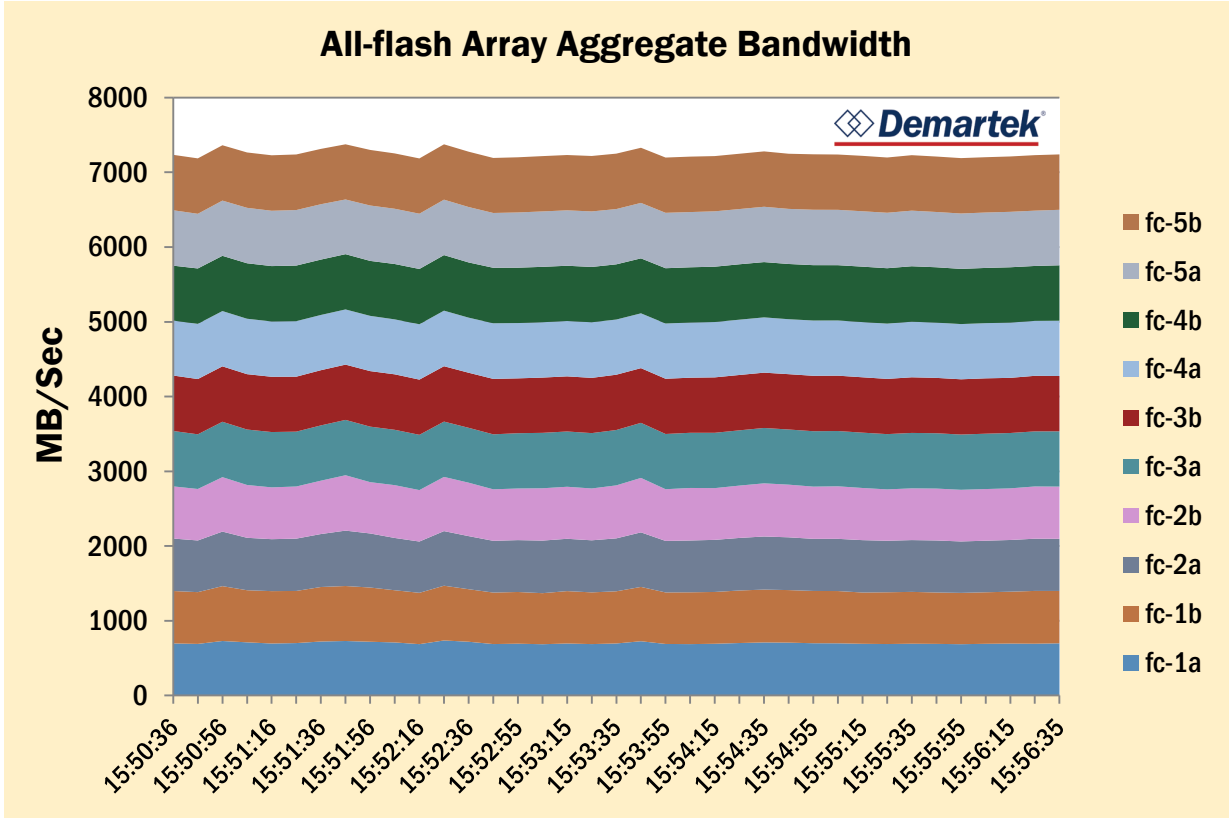

Source: Demartek Fibre Channel Zone – http://www.demartek.com/FC

Performance Results

► 16 GFC vs 8 GFC – Configuration

Demartek

Windows Server 2012
SQL Server 2008

16Gb/s 8Gb/s

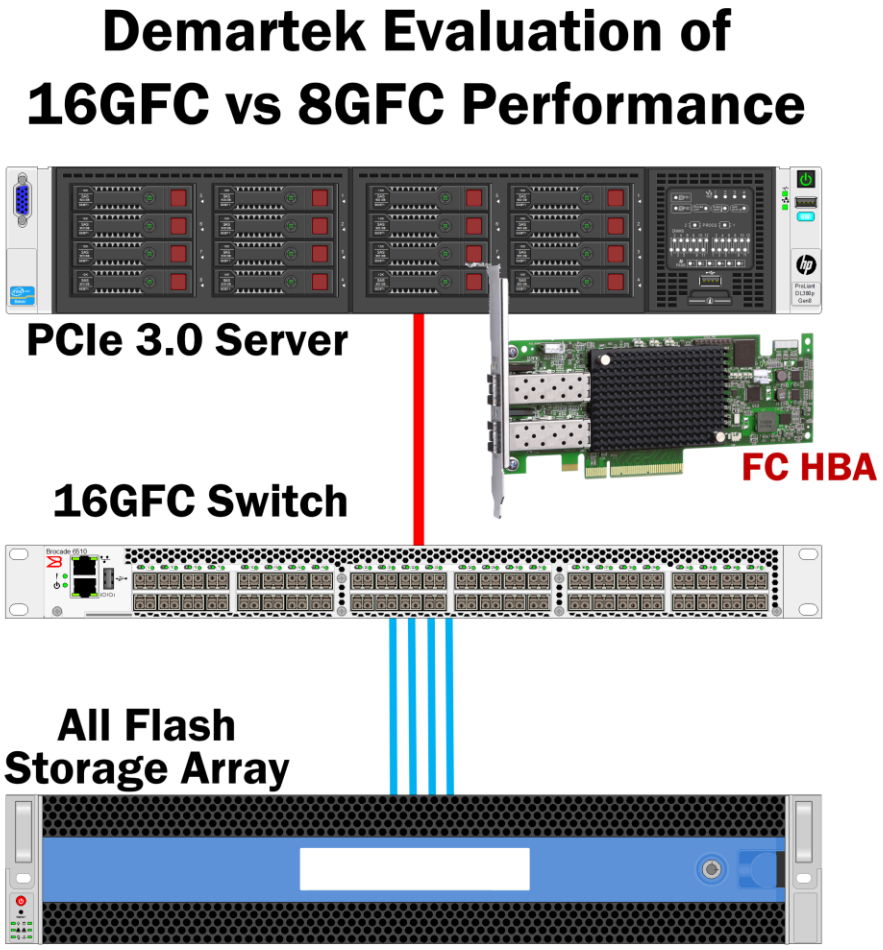

Performance Results

► 16 GFC vs 8 GFC – Test Results

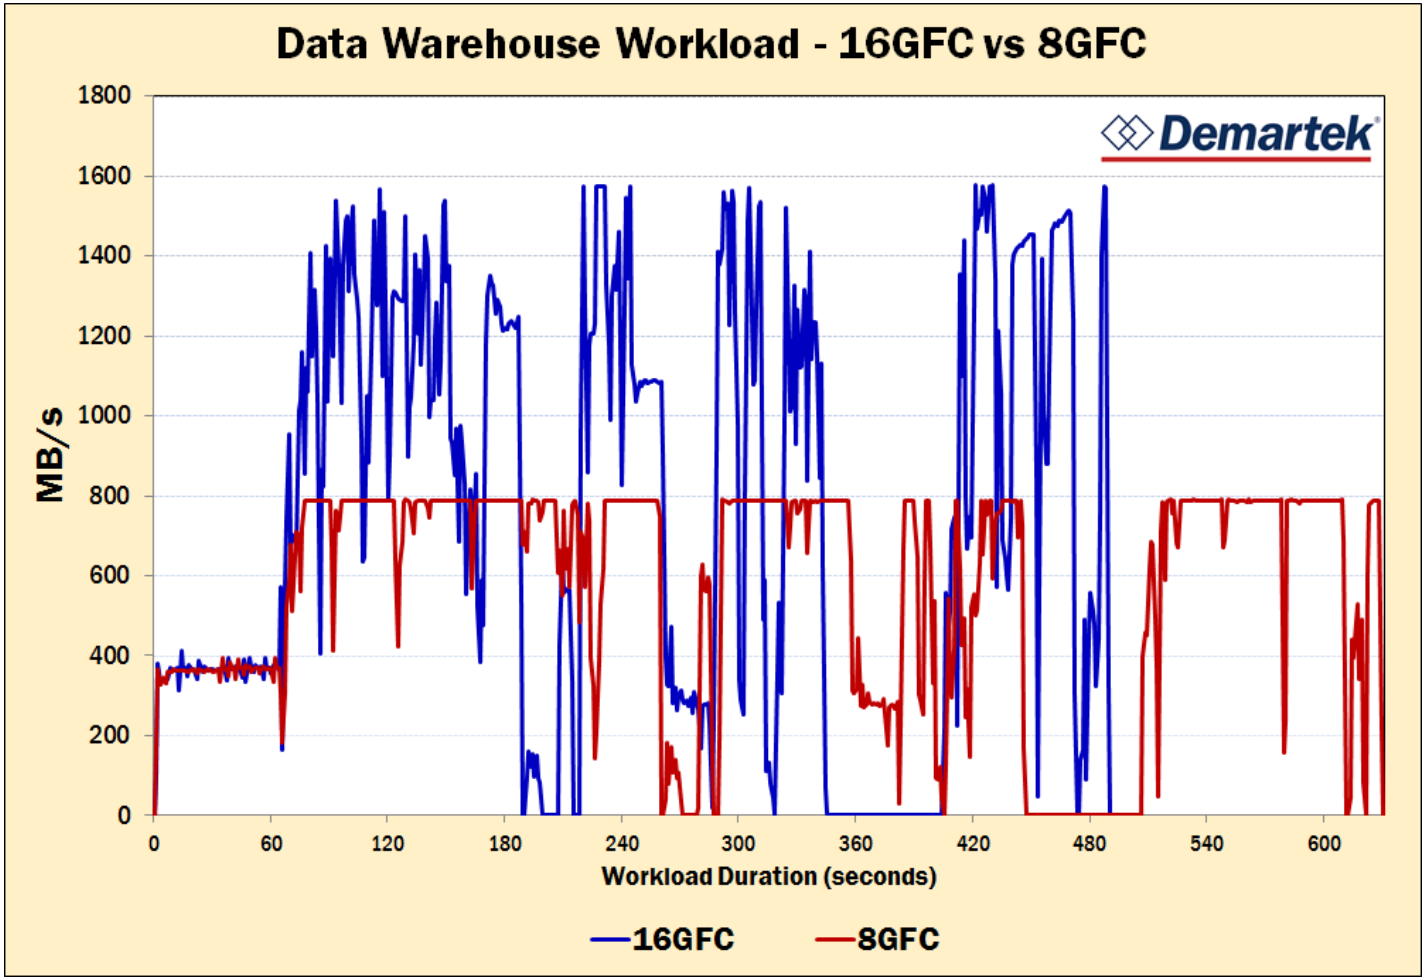

Source: Demartek Fibre Channel Zone – http://www.demartek.com/FC

Performance Results

► Virtual Fibre Channel – Configuration

Virtual Fibre Channel Server-Cluster VM Moves

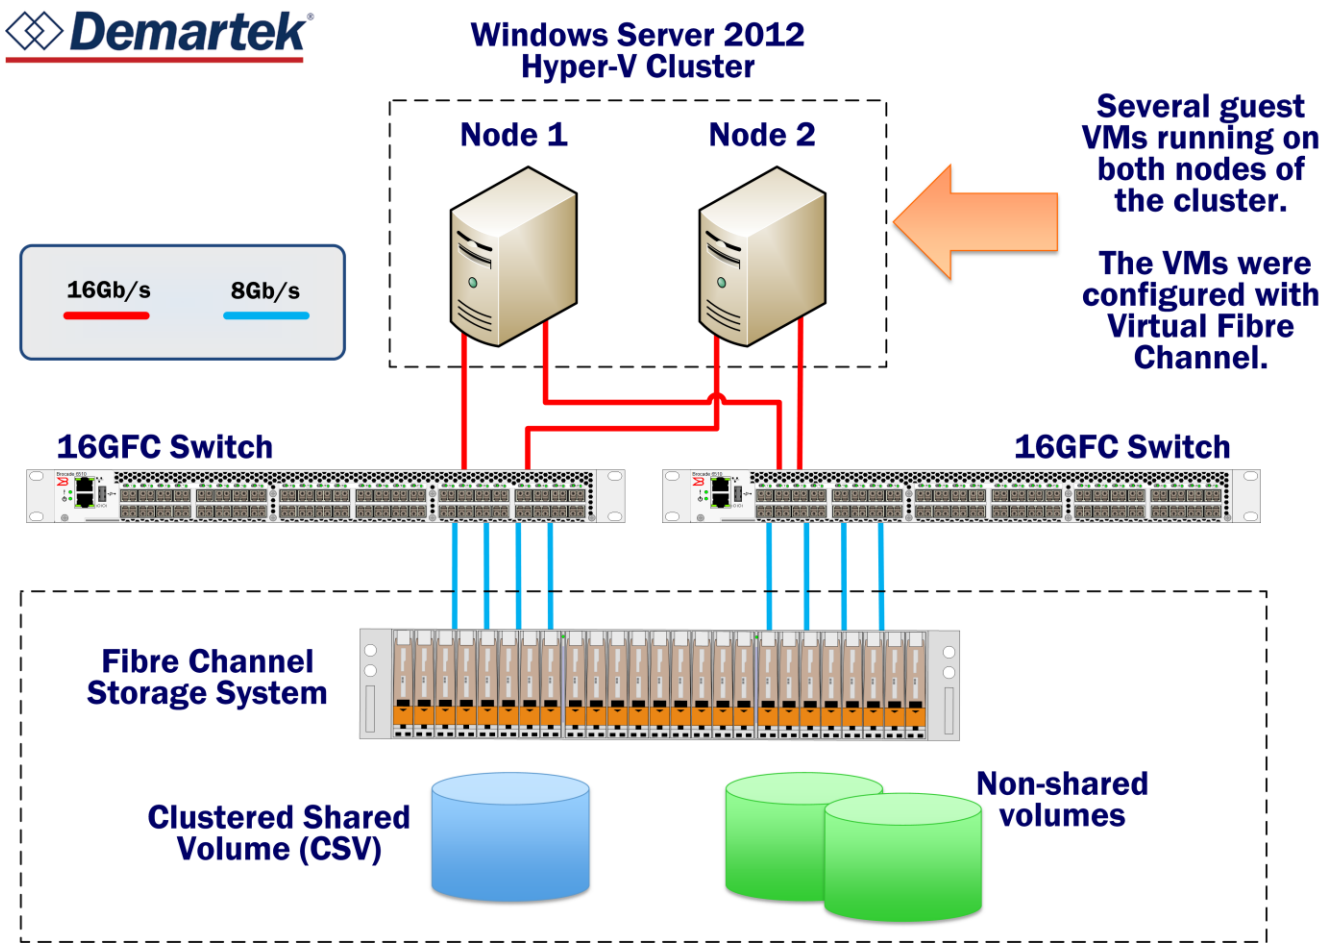

Performance Results

► Virtual Fibre Channel – Test Results

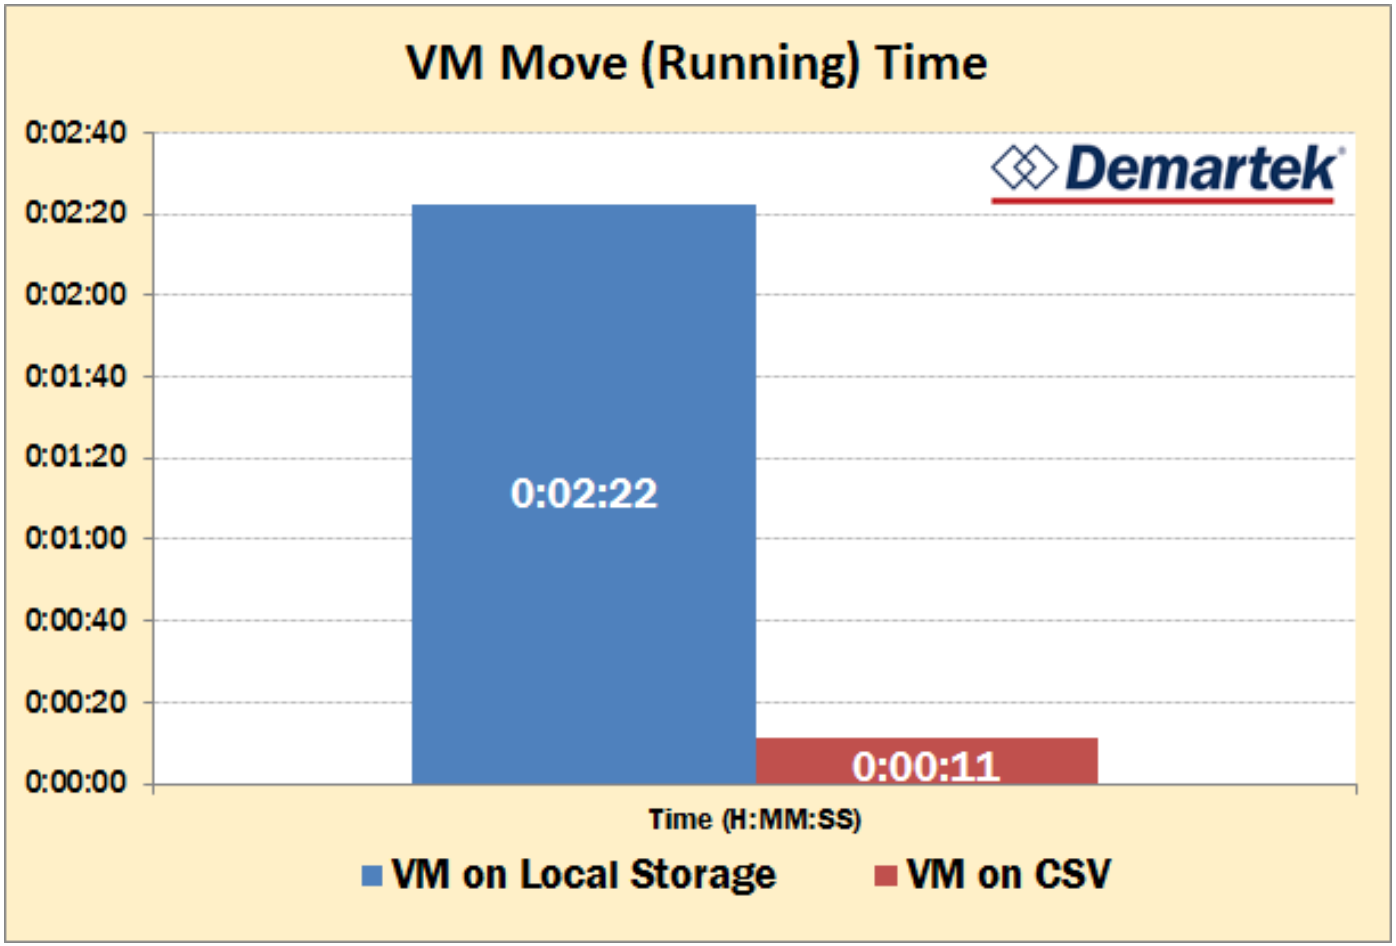

Demartek Free Resources

- Demartek FC Zone
  - www.demartek.com/FC
- Demartek FCoE Zone
  - www.demartek.com/FCoE
- Demartek iSCSI Zone
  - www.demartek.com/iSCSI
- Demartek SSD Zone
  - www.demartek.com/SSD
- Demartek SSD Deployment Guides
  - iSCSI and SSD available now, 16GFC coming…

Performance reports, Deployment Guides and commentary available for free download.

- www.demartek.com/Demartek_Interface_Comparison.html
- Or search for "storage interface comparison" in your favorite search engine

Free Monthly Newsletter

Demartek publishes a free monthly newsletter, ***Demartek Lab Notes***, highlighting recent reports, articles and commentary.

Look for the newsletter sign-up at:
www.demartek.com/Newsletter

Thank You!

Dennis Martin, President
dennis@demartek.com
www.linkedin.com/in/dennismartin

Demartek

(303) 940-7575
www.demartek.com
http://twitter.com/Demartek
www.youtube.com/Demartek
Skype: Demartek

To learn more about Demartek:

- Download the Aurasma App (Android/iPhone)
- Search and follow "Demartek"
- View image below with viewfinder.

*also on the back of Dennis' business card

Powered by: AURASMA