NAS for Server Virtualization

Dennis Chapman
Senior Technical Director
NetApp

Agenda

- The Landscape has Changed
- New Customer Requirements
- The Market has Begun to Move
- Comparing Performance Results
- Storage Management for Virtualized Environments
- Performance with Storage Efficiency and Protection
- A few words on SMB 2.2

The Landscape has Changed

- Yesterday before server virtualization
  - Select business critical apps
  - Small set of connected hosts
- Today with server virtualization
  - Hosts every application
  - Connected by every virtualized host
  - High level of data redundancy

New Customer Requirements

- High performance
  - Large pools of shared storage for the mass of "non-descript" VMs
  - Individual VMs with I/O demanding applications
- Operational Simplicity
  - Physical to virtual transparency

The Market has Begun to Move

- Forrester Research Storage Choices for Server Virtualization
  - Independent, non-commissioned analyst report
  - Results obtained from CIO survey data
  - Only multi-year industry report on subject
    - Published in Q1, 2009 & Q1, 2011

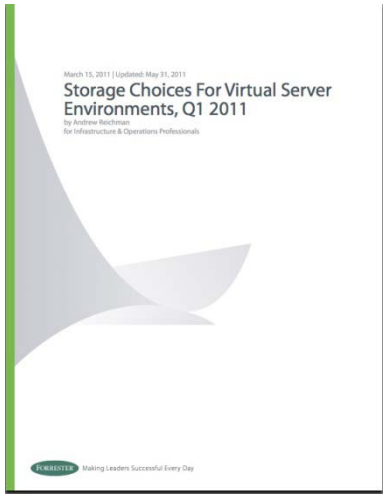

http://www.netapp.com/us/library/analyst-reports/ar-storage-choices-for-virtual-server-environments.html

The Market has Begun to Move

- 56% growth in reported use of NFS over 24 month period

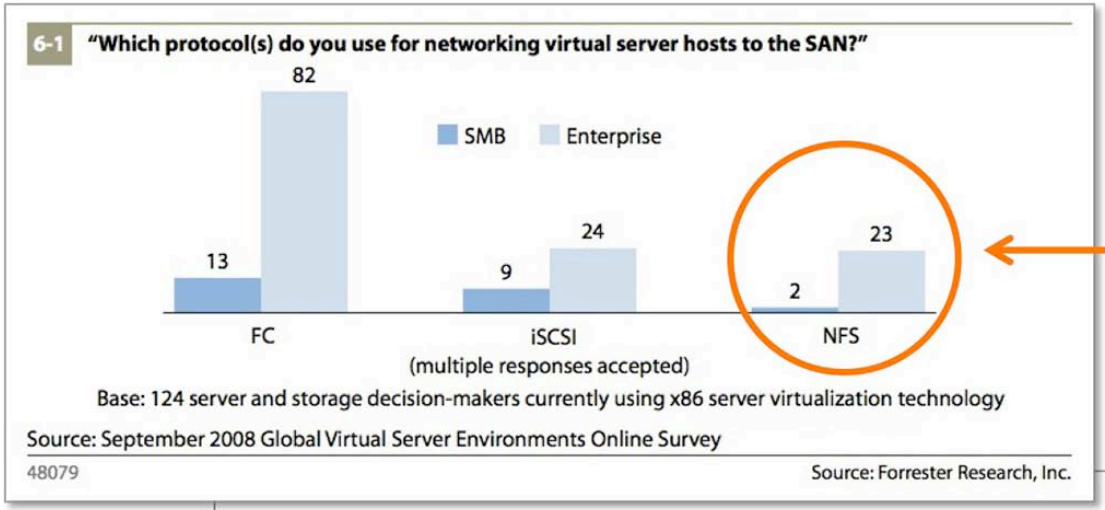

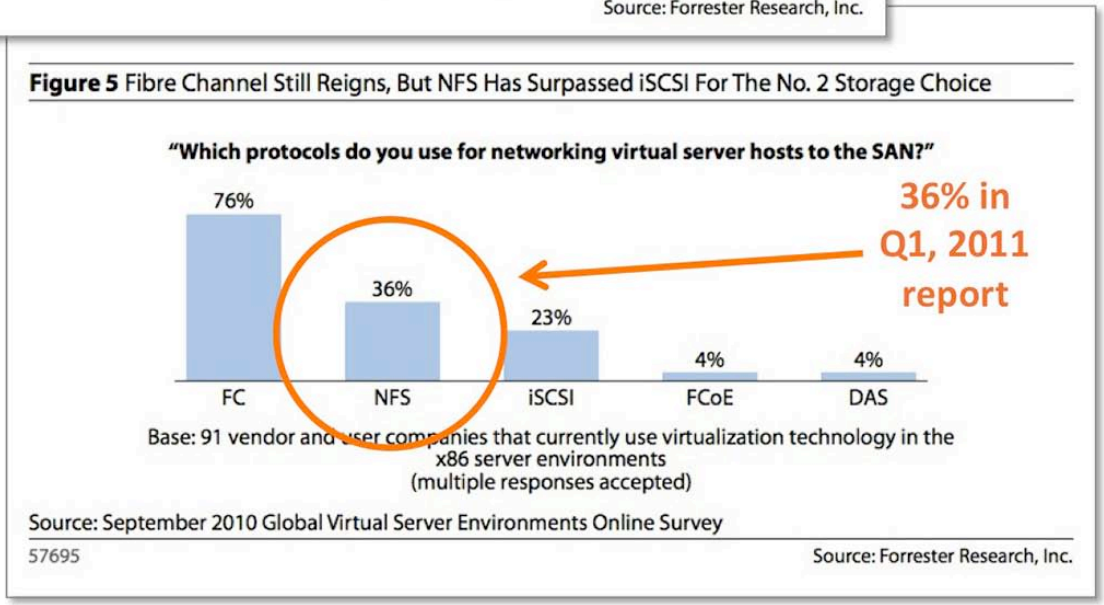

Comparing performance results

- Workload set A
  - Storage protocol performance using the same storage system and workload generator
  - Compares general application and active database workloads
- Workload set B
  - Compares performance with storage efficiency and protection using two different storage systems
  - Common general application workload generator

Testbed configuration A

Compute	Storage
8 Fujitsu RX200, 8-core Xeon, 48GB	NetApp FAS6210 active-active HA pair
VMware vSphere 4.1	192 x 300GB, 15k RPM SAS disk drives
Iometer controller & dynamos	Dual port 10Gb UTA (United Connect)

Storage Protocols	Network
Fibre Channel using both 4GB & 8GB	Cisco Nexus 5020 (storage network)
FCoE using 10Gb	Cisco 4948 switch (client network)
iSCSI using both GbE & 10GbE	Qlogic 8Gb or 4GB FC HBA
NFS using both GbE & 10GbE	Qlogic 10Gb CAN or Intel 1Gb NIC

Test topology A

- Joint testing by NetApp & VMware
  - Published as TR-3916
- Same hardware used in compute, network, and storage for testing of protocols
- Tested shared IO pools & independent high IO VMs

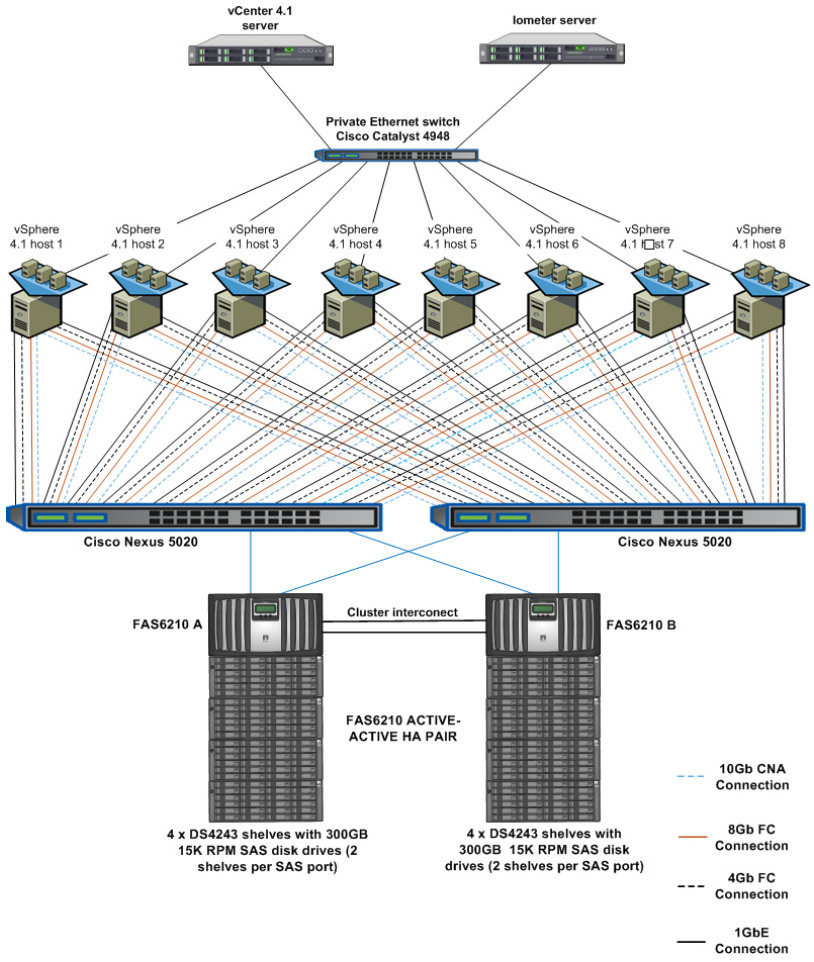

Workload characterization A-1

- VMs with general purpose applications
  - 75% reads / 25% writes
  - NTFS/EXT3 default block sizes (4k)
  - Relatively light I/O load from any particular VM, but in aggregate a large workload
  - Shared datastores
  - VM counts range 32, 64, 96, 128
  - Outstanding I/Os range 64, 128, 192, 256

Workload A-1 IOPS comparison

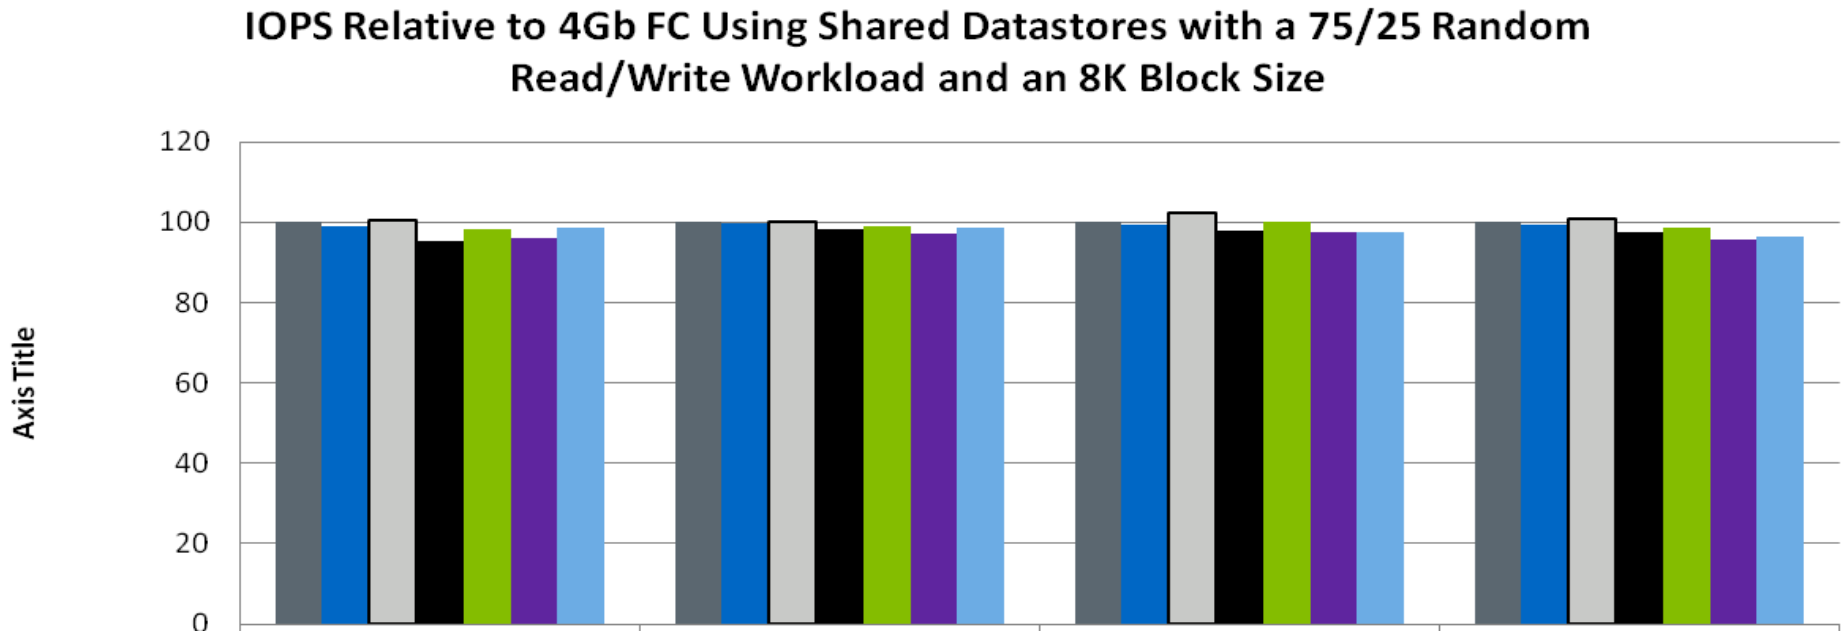

	32 VMs	64 VMs	96 VMs	128 VMs
FC - 4Gb	100	100	100	100
FC - 8Gb	99	100	99	99
FCoE - 10Gb	100	100	102	101
iSCSI - 1GbE	95	98	98	97
iSCSI - 10GbE	98	99	100	99
NFS - 1GbE	96	97	98	96
NFS - 10GbE	99	99	98	96

Workload A-1 Latency comparison

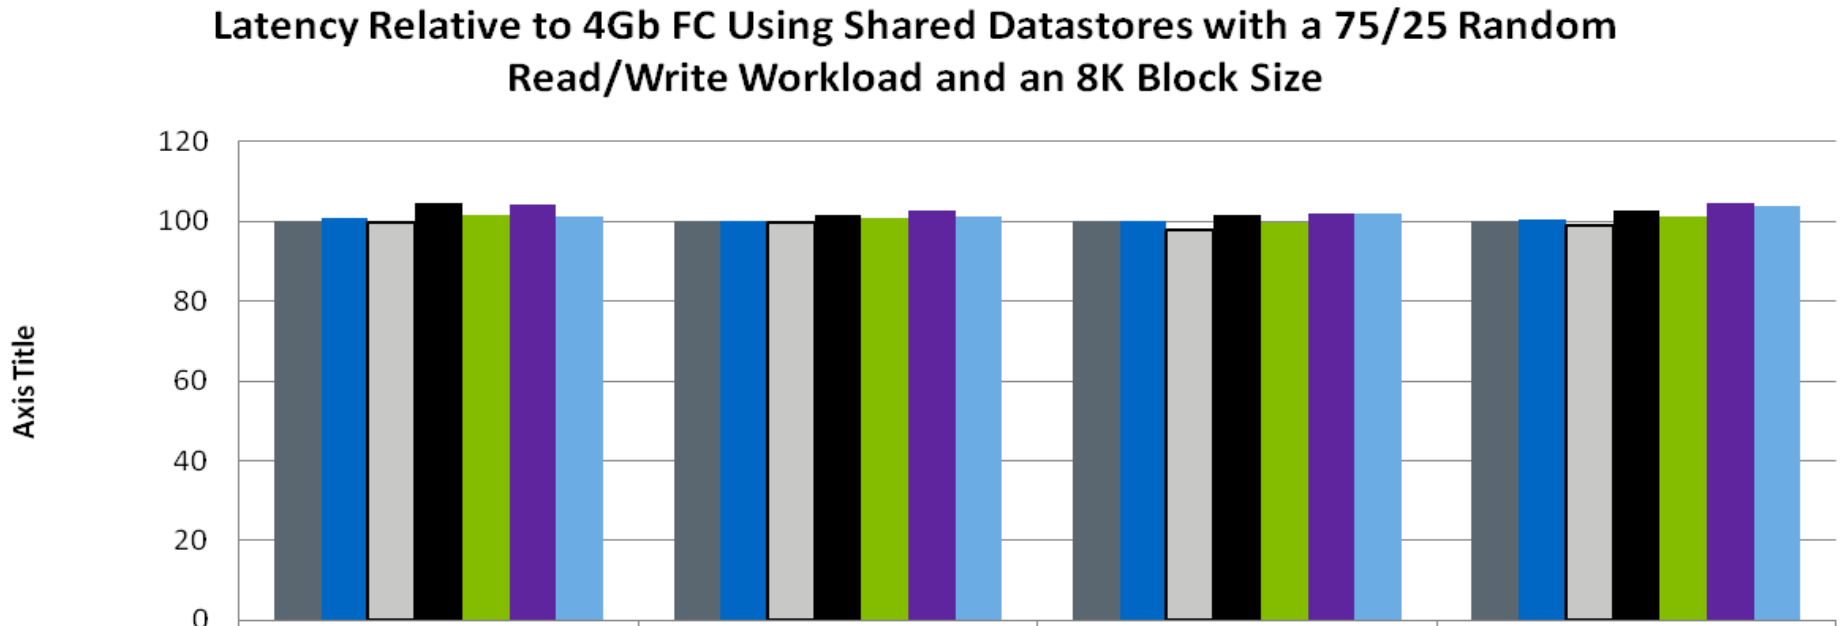

	32 VMs	64 VMs	96 VMs	128 VMs
FC - 4Gb	100	100	100	100
FC - 8Gb	101	100	100	101
FCoE - 10Gb	100	100	98	99
iSCSI - 1GbE	105	102	102	103
iSCSI - 10GbE	102	101	100	101
NFS - 1GbE	104	103	102	105
NFS - 10GbE	101	101	102	104

Workload characterization A-2

- VMs hosting an active database
  - 128 concurrent I/Os with 60% reads / 40% writes
  - Fully random access pattern
  - 8k block size
- Emulates an I/O-intensive applications like OLTP databases hosted in VMs
- Uses ESX Paravirtual SCSI (PVSCSI) driver
  - Improves throughput reduces CPU utilization

Workload 2 IOPS comparison

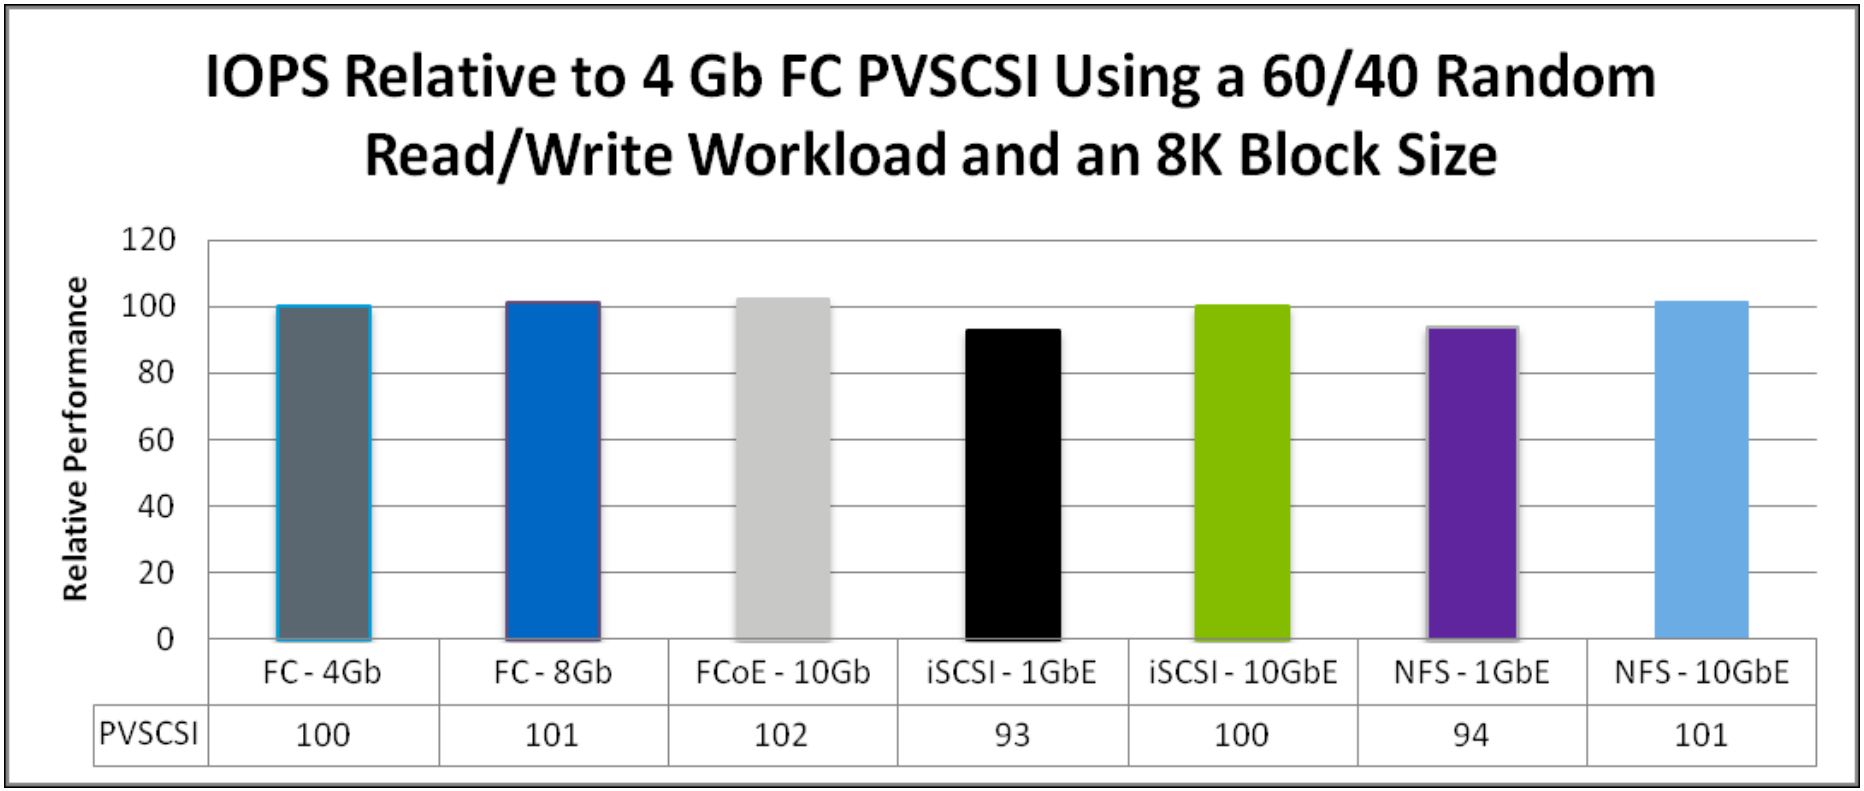

Workload 2 latency comparison

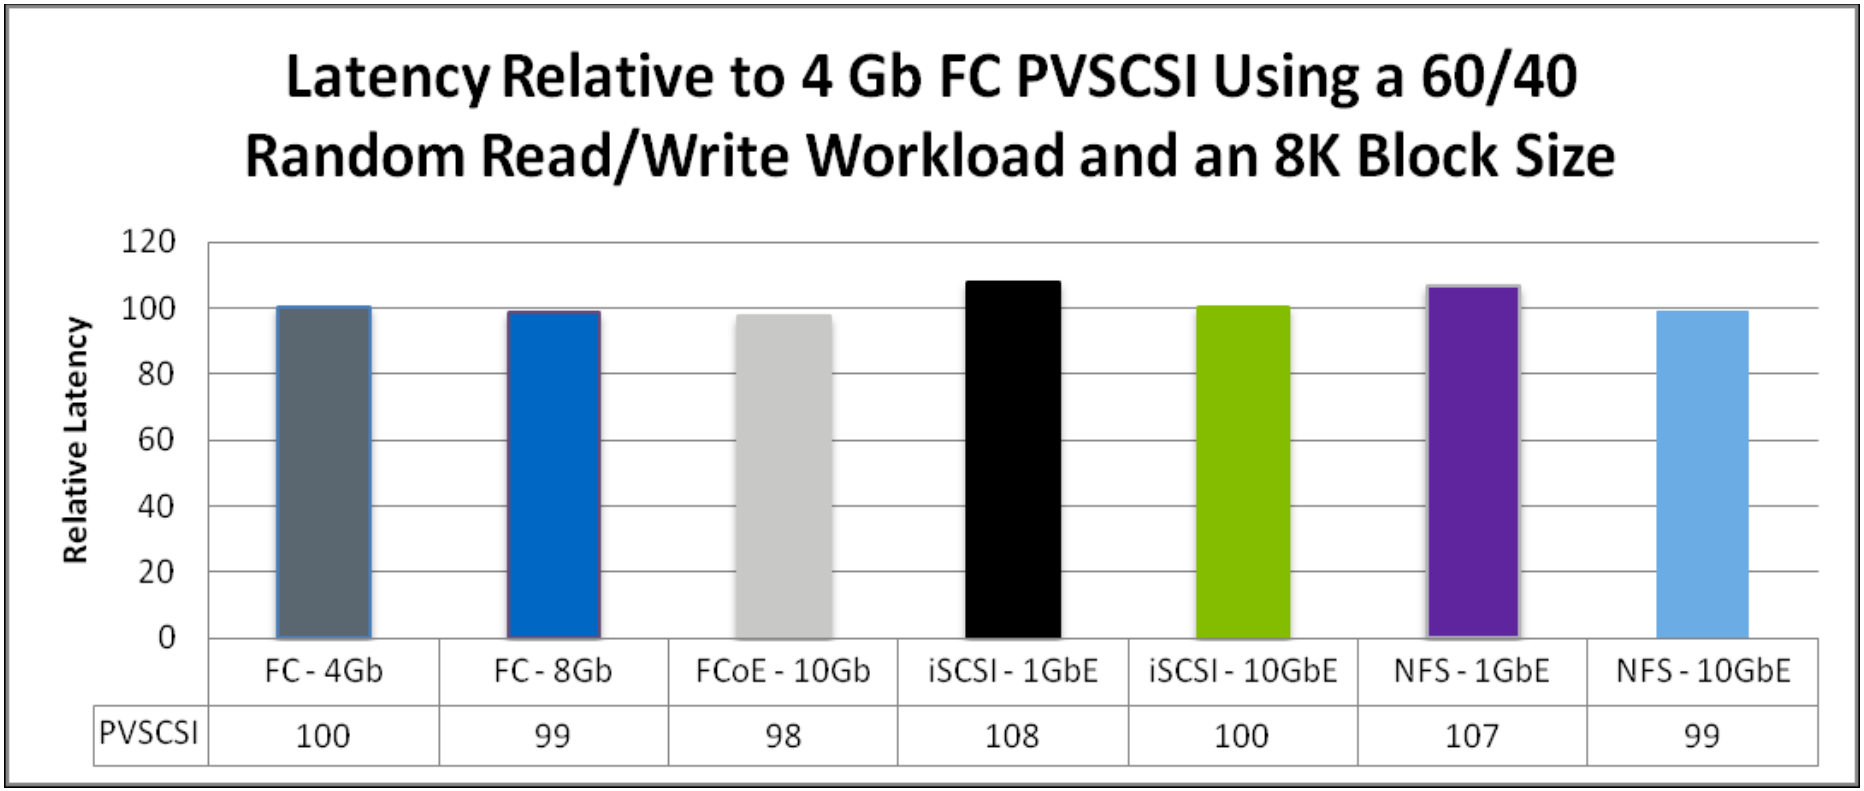

Storage Management

- With performance concerns addressed, customers focus shifts to management or reducing operational costs

Focus is Storage Management

- SAN
  - Integrations limited to software executing on hypervisor without consideration of array
  - Clustered file systems restrict data access to a single hypervisor platform
- NAS
  - Allows direct integration of file level virtualization in coordination with hypervisor
  - Open standard, provides storage as a network service to any vendor's hypervisor

Focus is Storage Management

- SAN
  - Hidden I/O bottlenecks
    - LUN queue limits
  - Reduced scaling in # VMs per shared storage pool
    - Common to see 6-12 VMs per shared pool
    - Increased provisioning and replication tasks based on VM to LUN ratio
- NAS
  - Transparent I/O limits
    - Network bandwidth, array CPU, disk IOPs, etc
  - Greater scaling in # VMs per shared storage pool
    - 60-200 VMs per shared pool
    - Fewer provisioning and replication tasks

Focus is Storage Management

- SAN
  - Complex architecture
    - Troubleshooting challenges as shared LUNs don't map to VMs
  - Vendor APIs enable integration
- NAS
  - Simple architecture
    - Easy as files map to VMs
  - Native functionality and integration across multiple vendors and toolsets

Understanding Storage I/O

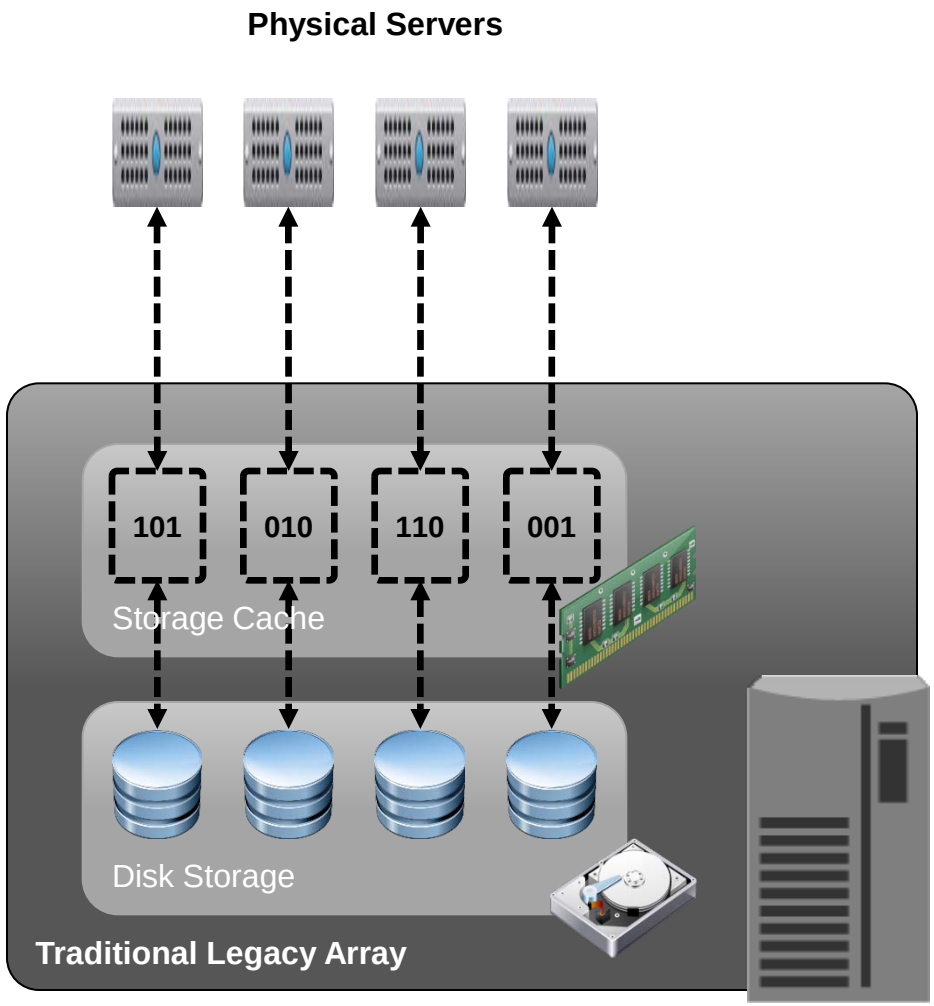

- Storage array Cache
  - Increases I/O performance
  - Copies data previously requested
  - Cache intelligence prefetches data in anticipation of future requests
- Legacy architecture
  - Caching is implemented on a per-server basis
  - Ideal design when shared arrays were exclusive to I/O intensive applications

Understanding Virtualized Storage

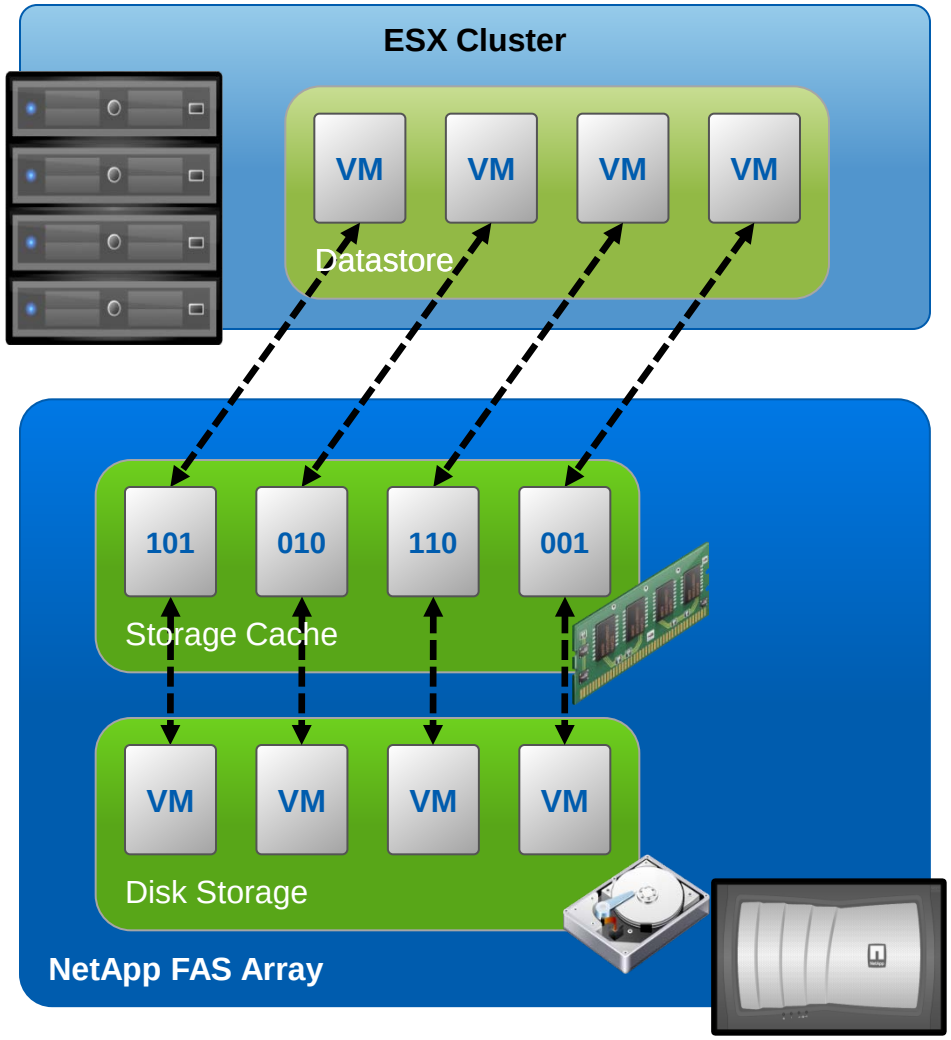

- Storage array Cache
  - Increases I/O performance
  - Copies data previously requested
  - Cache intelligence prefetches data
- Legacy architecture
  - Caching is implemented on a per-VM basis
  - Ideal design when shared arrays were exclusive to I/O intensive applications

Virtualized Storage I/O

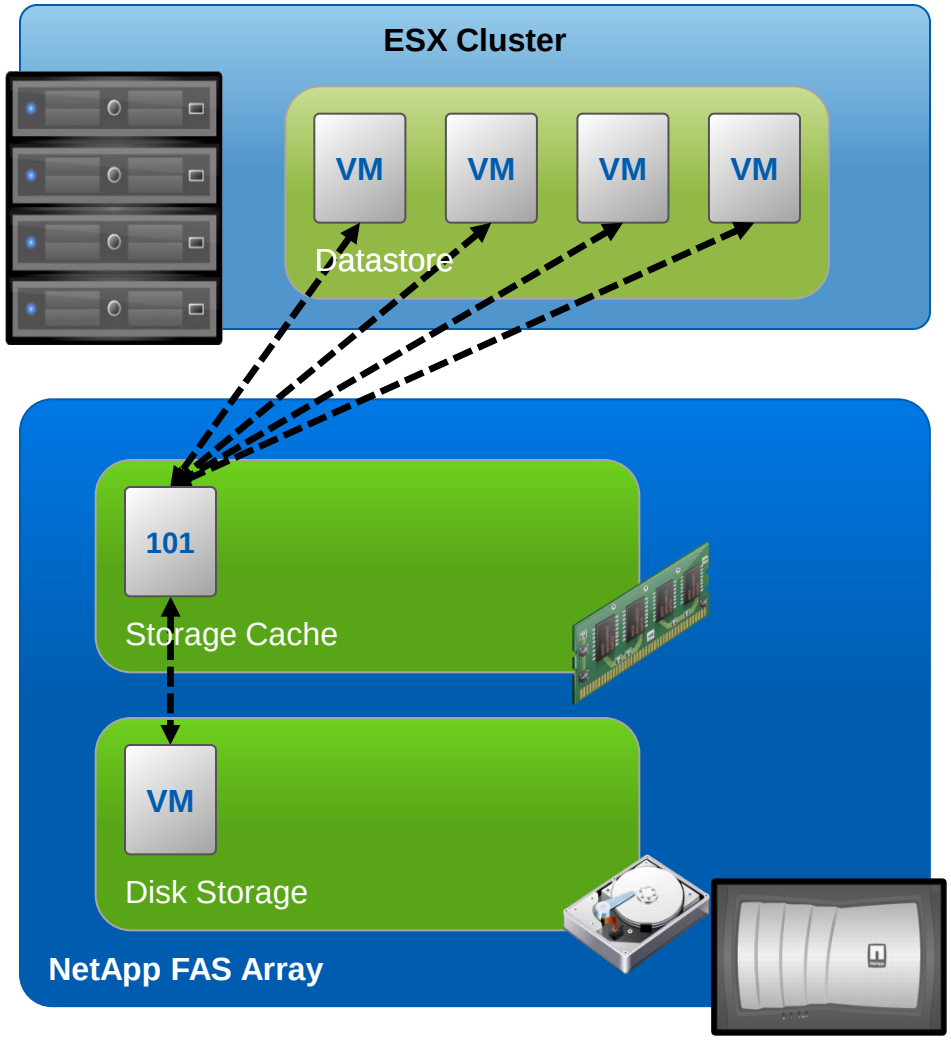

- Advanced Virtualized Cache
  - Dedupe aware caching
  - Increases I/O performance beyond non-deduped workloads
  - Shares previously requested data with multiple VMs
  - Shares prefetched data with multiple VMs
  - Ideal design for shared infrastructures

I/O gains like Cisco WAAS
or
VMware Transparent Page Sharing

Workload Set B

- Measure performance with storage efficiency & RAID protection
- General purpose workload
  - 100% random mix of 75% Read 25% Write using a 4k block size
  - 2 outstanding I/Os per VM
- Storage configuration
  - Protection provided by RAID
    - RAID-DP vs. RAID 5, RAID 6, & RAID 10
  - Storage efficiency
    - NTAP thin-provisioned data stores vs. thin VMDKs configured from a spanned VMFS partition
  - Workload generation using vSphere servers & IOmeter

Storage and Network configs

NetApp FAS3140 Mid-tier Unified Storage Array	Traditional Mid-tier FC Storage Array
2 x 4GB of storage controller cache	2 x 8GB of storage controller cache
40 x 450GB, 15k RPM SAS disk drives	152 x 146GB, 15k RPM FC disk drives
2 x 256GB Flash Cache expansion modules (used for perf run 2)	Dual storage controllers

NetApp FAS3140 Mid-tier Unified Storage Array	Traditional Mid-tier FC Storage Array
GbE Ethernet network	4Gb Fibre Channel network
Cisco 4948 Ethernet switch	Brocade 200E FC switch

Tested Storage Systems

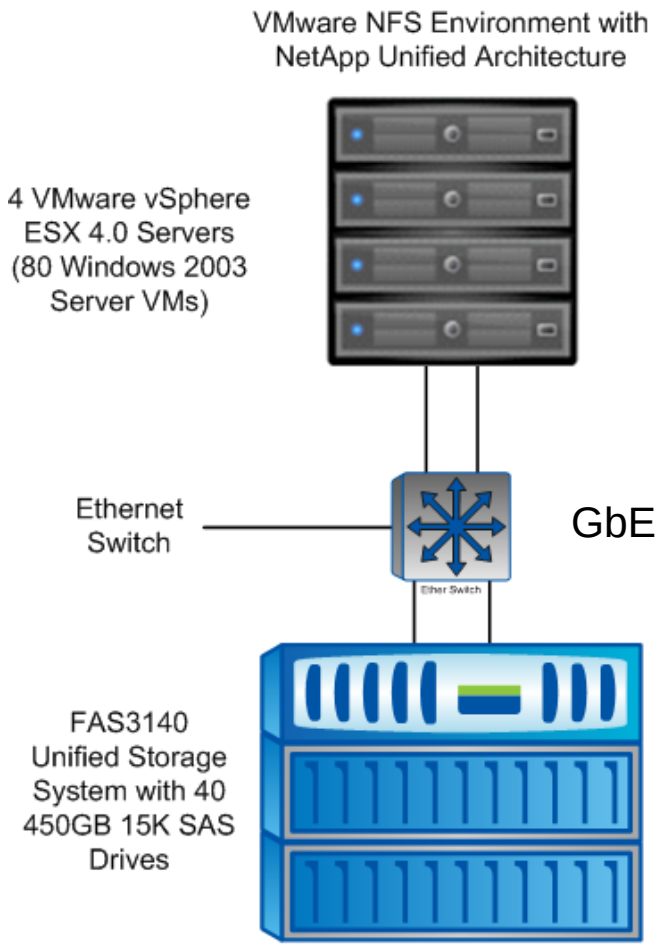

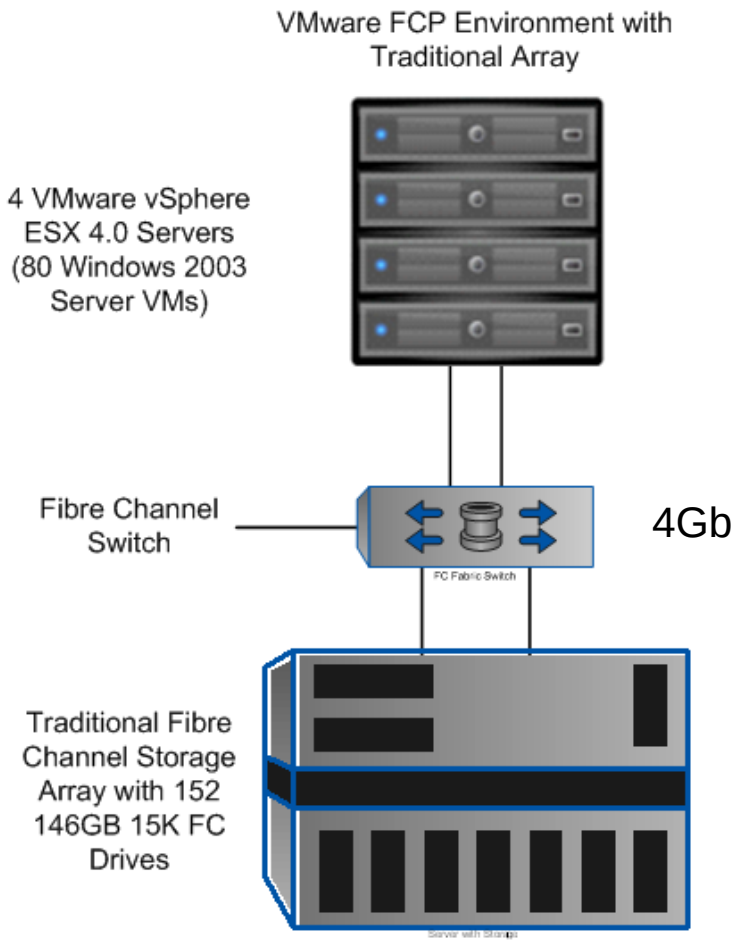

Workload IOPS comparison

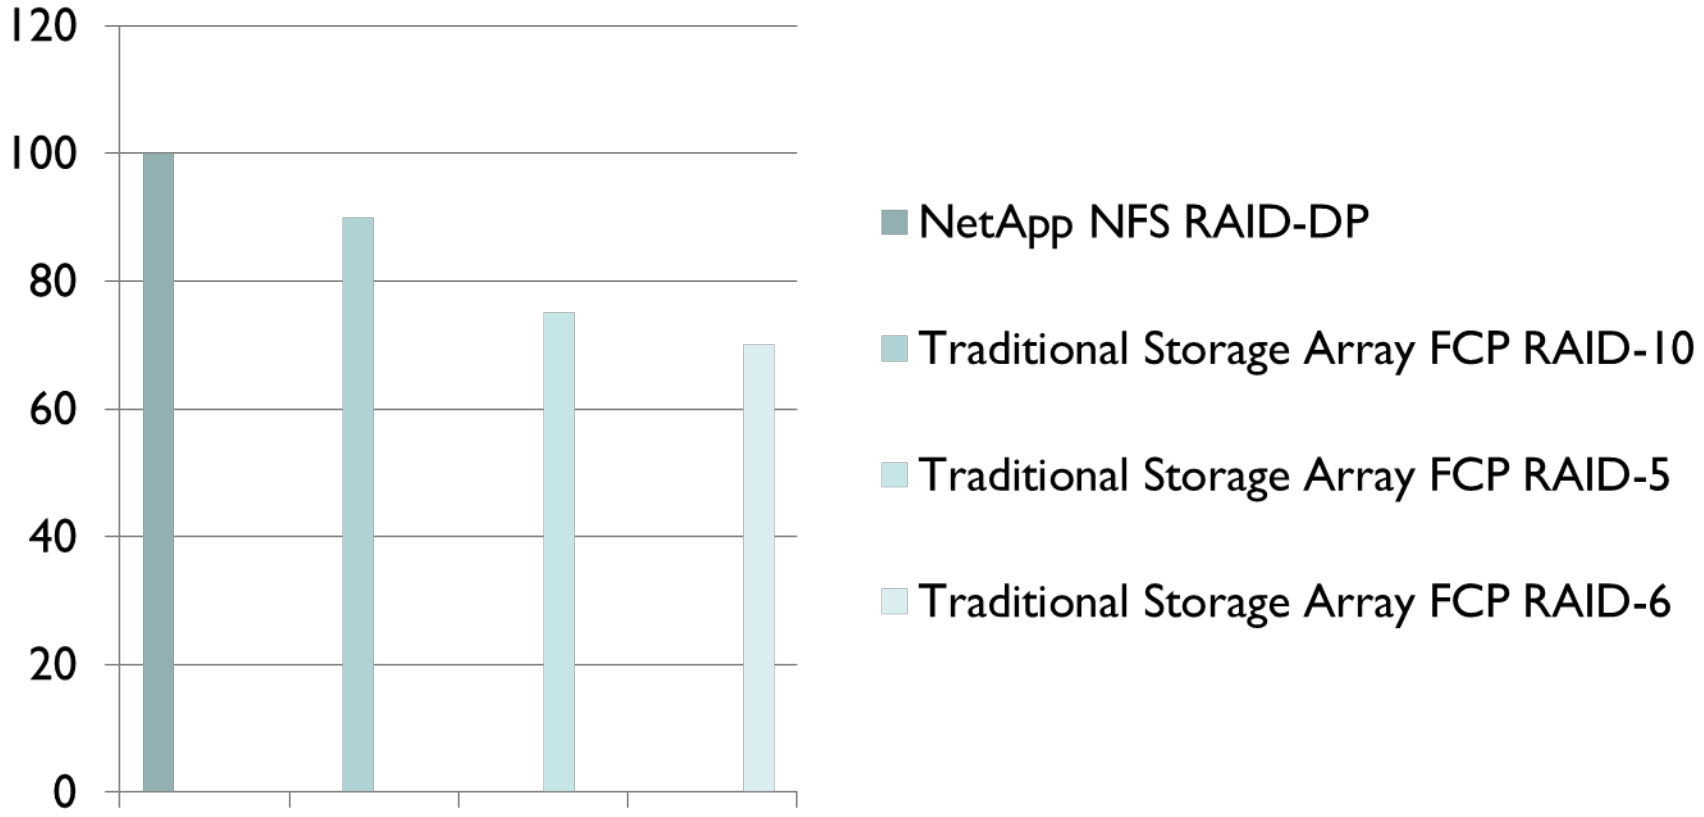

VMware Verified Performance

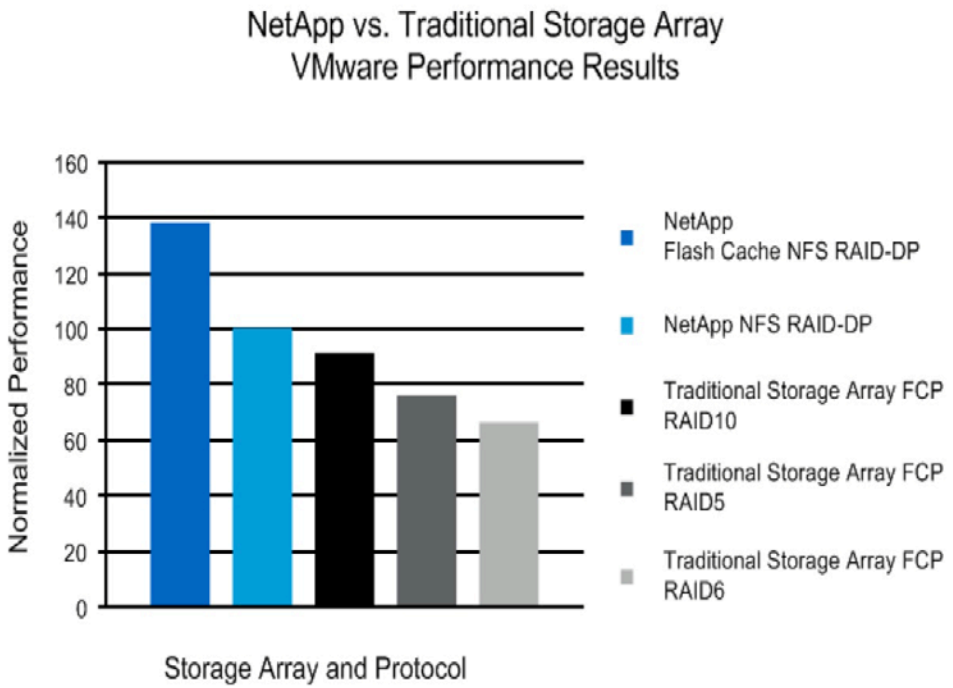

http://media.netapp.com/documents/tr-3856.pdf

- I/O intelligence from dedupe-aware array cache results in greater performance than traditional arrays

Closing

- Market trends show virtualization workloads moving to NAS
- Performance variances between FC, FCoE and NFS over 10GbE are minor and are overcome by storage management advantages with NAS
- Advanced caching enables superior performance with storage efficiency and RAID protection

NetApp and SMB 2.2

- Very important set of enhancements to the SMB2 protocol
- Microsoft – NetApp technical engagement during SMB information transfer has been excellent
- NetApp will ship support for SMB 2.2 as part of DataONTAP 8.2 Cluster-mode