SMB 3.0 – New Opportunities for Windows Environments

Webcast Presenters

SNIA™
Ethernet Storage Forum

Linda Wu
Director of Product Management, Windows File Services
NetApp

Linda Wu is the Director of Product Management for the Windows File Services solution at NetApp. In this role, she is responsible for driving the file-based storage business for file sharing, BYOD, Microsoft Hyper-V and SQL, and technical applications. Her responsibility also includes defining the product direction for SMB protocol

Dennis Chapman
Senior Technical Director
NetApp

Dennis Chapman is Senior Technical Director for Microsoft Solutions at NetApp. Dennis joined NetApp in 1998 working on the Data ONTAP CIFS stack. Currently Dennis drives the architecture of NetApp's Windows product portfolio, providing tight integration between the Windows ecosystem and Data ONTAP's strong NAS, SAN and data management capabilities.

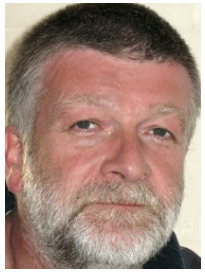

Alex McDonald
Office of the CTO
NetApp

Alex McDonald joined NetApp in 2005, after more than 30 years in a variety of roles with some of the best known names in the software industry . Alex is now part of NetApp’s Office of the CTO that supports industry activities and promotes technology & standards based solutions, and is co-chair of the SNIA Cloud Storage Initiative, and co-chair of the SNIA File Protocols Special Interest Group

SNIA Legal Notice

SNIA Ethernet Storage Forum

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Agenda

SNIA
Ethernet Storage Forum

- Market Trends for File-based storage
- Evolution of the SMB Protocol
- Overview of SMB 3.0
- SMB 3.0 in Action
- Key Takeaways

Market Trends

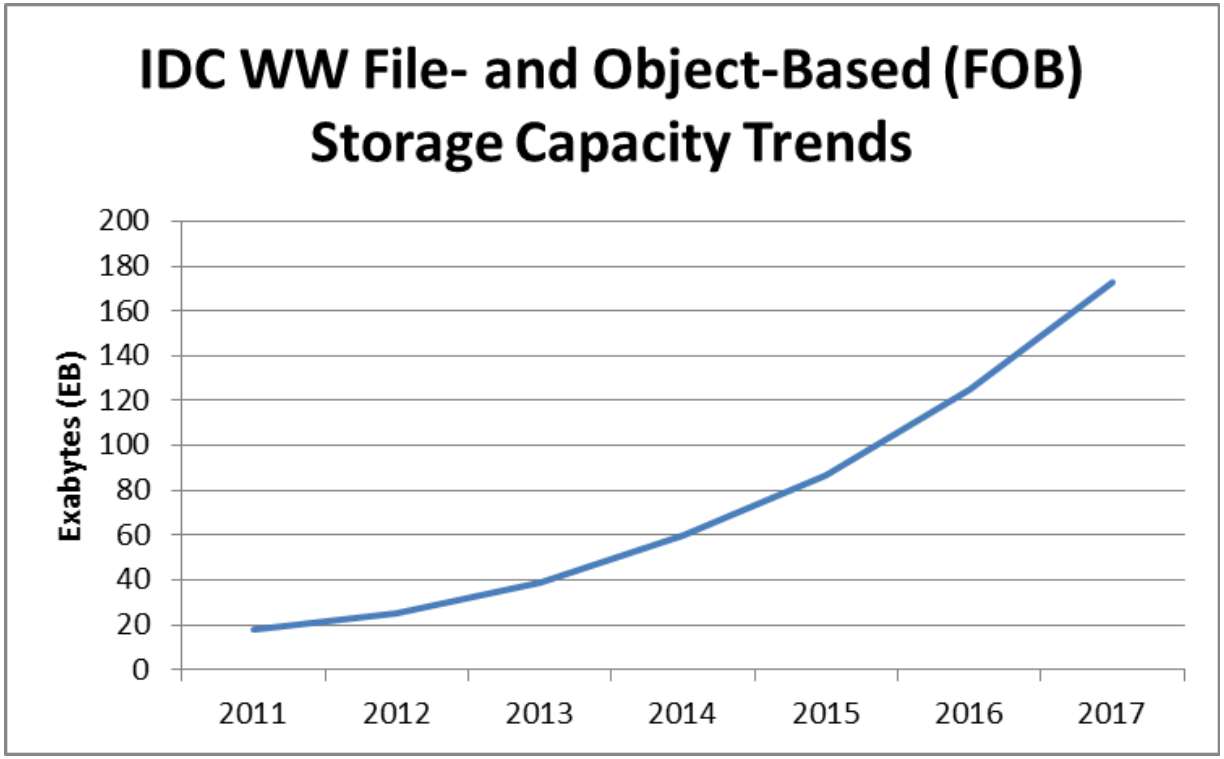

Source: IDC WW File- and Object-Based Storage 2013-2017 Forecast

- 2017: 4X data growth, 173 Exabytes, 47.2% CAGR
- File- and Object-based storage (FOB) is outpacing the overall storage growth

Key Drivers for File Storage

SNIA Ethernet Storage Forum

- Continued growth in traditional file sharing
- Traditional block-based apps moving to file
  - Hypervisors, VDI, Databases
- Hybrid application architecture - structured apps managing unstructured data
  - MS SharePoint
- Cloud and new apps
  - Archiving, Tiering, Mobile, Social Apps

Why Files?

SNIA
Ethernet Storage Forum

- Data Mobility & Granular Management
  - Easily copy, delete and move data around
  - Ability to restore single VMs, databases, or groups of VMs and databases
- Operational
  - Very little technical knowledge required - everyone knows how to manage shares and mountpoints
  - Less complex infrastructure, no HBA etc

Agenda

SNIA
Ethernet Storage Forum

- Market Trends for File-based storage
- Evolution of the SMB Protocol
- Overview of SMB 3.0
- SMB 3.0 in Action
- Key takeaways

SMB Timeline

Stylish SDC plugfest t-shirt provides a helpful timeline for SMB

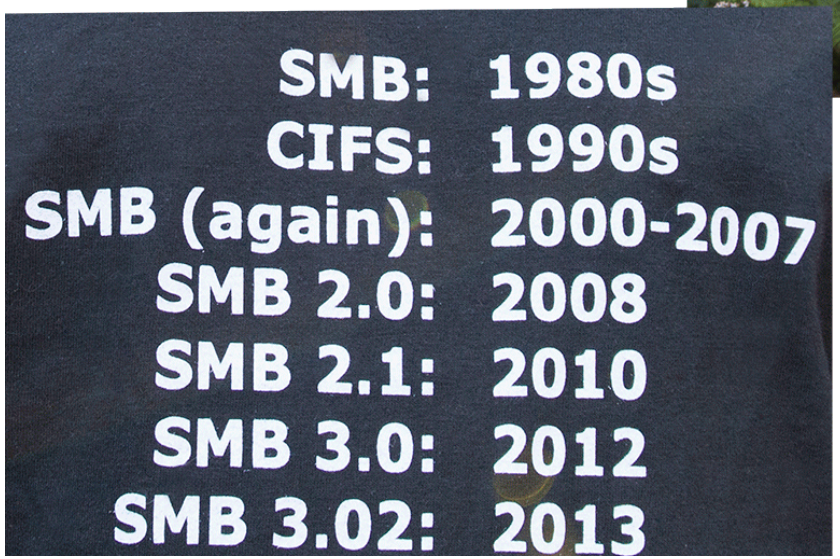

A long life with lots of activity over the last few years

In the beginning

SNIA
Ethernet Storage Forum

- SMB: 1980s
  - DOS 3.x - redirector & SMB remote file protocol for FAT
  - Simple and supported on a variety of OS platforms
  - Various docs e.g. X/Open (the ancient texts)
- CIFS: 1990s
  - Large number of operations added for NTFS with NT 4.0
  - More complex, harder to implement
  - IETF draft – Common Internet File System, 1997
  - SNIA Technical Specification, 1999

In the enterprise

SNIA
Ethernet Storage Forum

- SMB (again): 2000 – 2007
  - Active Directory, Kerberos and other ancillary protocols
    - A SMB server supports a suite of protocols beyond SMB
  - Increased file server complexity
    - Backward compatibility to FAT, e.g. 8.3 naming, code pages
    - Unevenly documented during the early years
  - Documentation access regulated by US DoJ & EC
    - Microsoft Communication Protocol Program – 2003
    - Large set of detailed reference documents developed
    - Licensed implementations from storage vendors

In the enterprise - reloaded

SNIA Ethernet Storage Forum

- SMB 2.0: 2008
  - Start with a clean sheet of paper
    - New, smaller command set greatly simplifies SMB
  - WS 2008, Vista
- SMB 2.1: 2010
  - WS 2008 R2, Windows 7
- SMB 3.0: 2012
  - SMB grows up
  - WS 2012, Windows 8

Question…

SNIA
Ethernet Storage Forum

- Do you have any issues about running hypervisor or database workloads over file-based protocols?

Agenda

SNIA
Ethernet Storage Forum

- Market Trends for File-based storage
- Evolution of the SMB Protocol
- Overview of SMB 3.0
- SMB 3.0 in Action
- Key takeaways

SMB 3 – Key features

SNIA Ethernet Storage Forum

- SAN-level Availability and Resilience
  - High availability applications can use SMB
- SAN-level Performance
  - Scales well compared to Fibre Channel
- Integrated Data Protection for Applications
  - Remote VSS for data stored on a SMB share
- Optimized Data Transfer (ODX)
  - Server-side copies using SCSI and/or SMB
- Security & Management
  - Encrypt messages over the wire
  - PowerShell and SMI-S

Highly Available SMB 3 configuration

- SMB 3 clients
  - Physical or virtual
  - View managed namespace
  - If permitted can connect to one or more SMB shares
- SMB 3 servers
  - Typically physical
  - Two or more nodes for failover handling
  - May include scale out clustering with a private interconnect

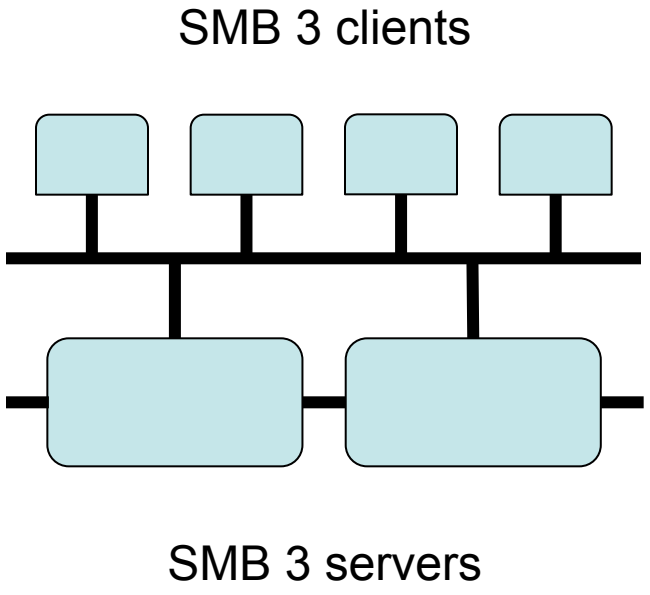

Cluster Client Failover - 1

- Clustered application runs on two client nodes

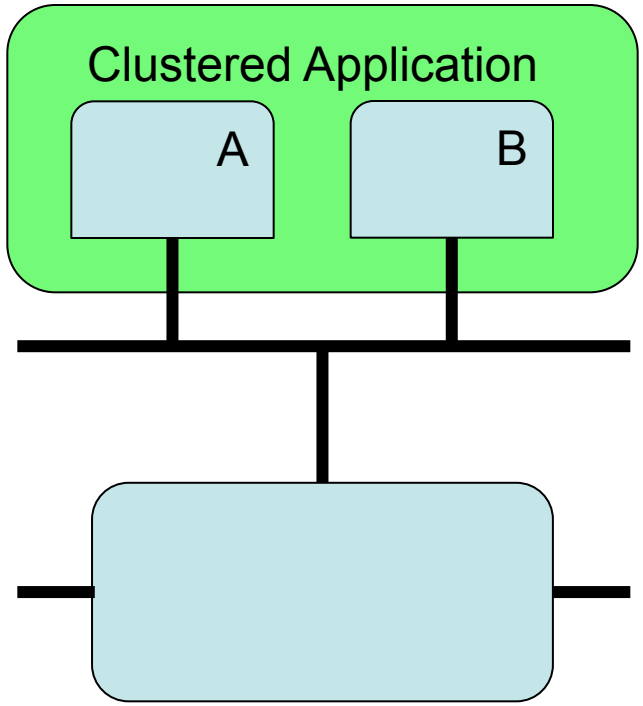

Cluster Client Failover - 2

- Clustered application runs on two client nodes
- Client A establishes session with server 1 using an unique ID
- Server 1 associates with client state with the ID

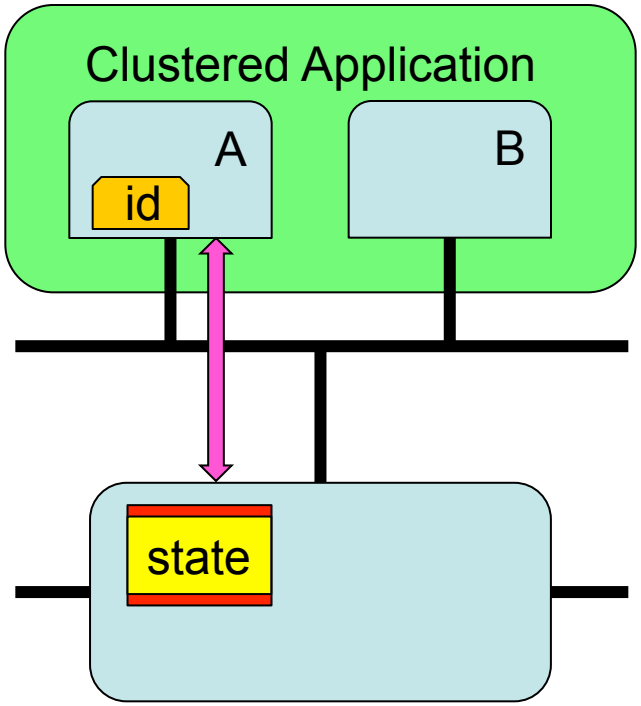

Cluster Client Failover - 3

- Clustered application runs on two client nodes
- Client A establishes session with server 1 using an unique ID
- Server 1 associates with client state with the ID
- Application work changes storage

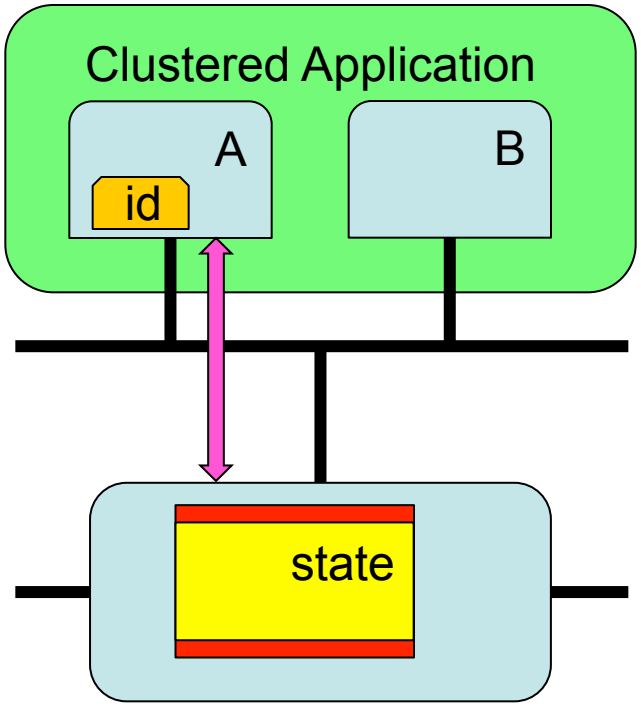

Cluster Client Failover - 4

- Clustered application runs on two client nodes
- Client A establishes session with server 1 using an unique ID
- Server 1 associates with client state with the ID
- Application work changes storage
- Client A fails

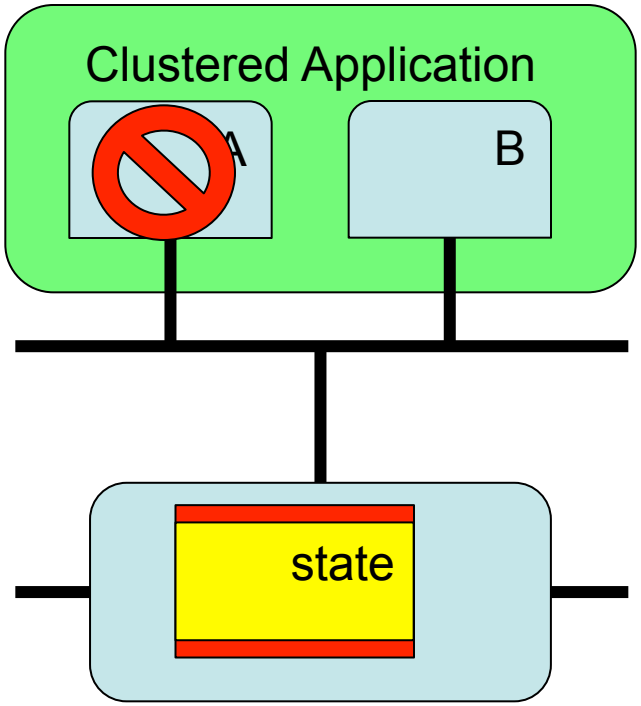

Cluster Client Failover - 5

- Application restarts. Client B establishes session using A's id and continues work

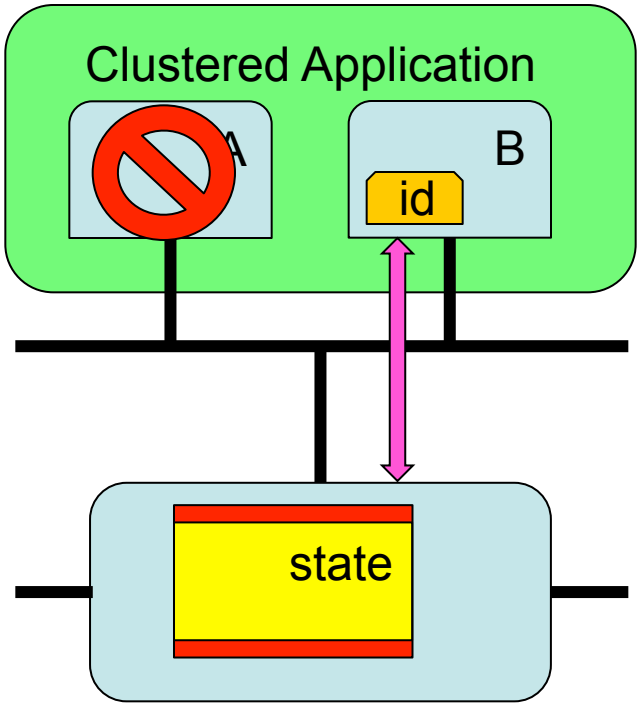

Transparent Server Failover - 1

- Server 1 offers a continuous availability (CA) share

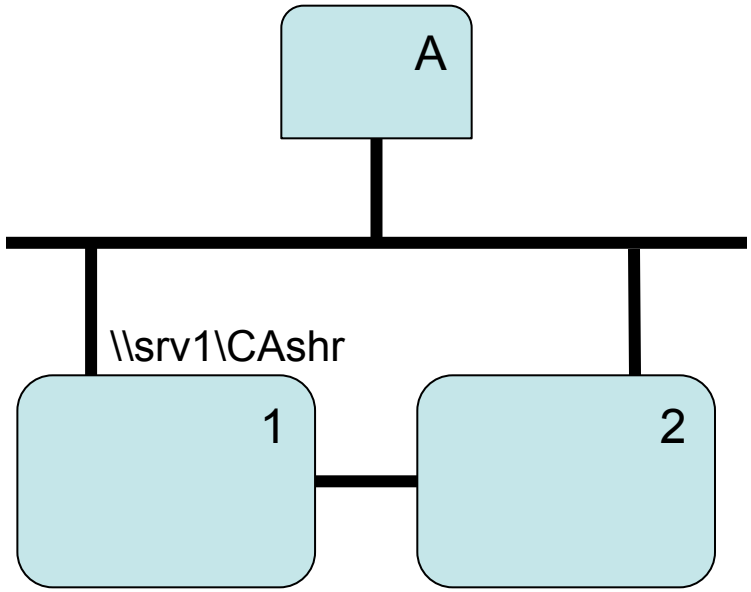

Transparent Server Failover - 2

- Server 1 offers a continuous availability (CA) share
- Client A mounts share and works
- Server 1 updates Server 2 with client A's state changes

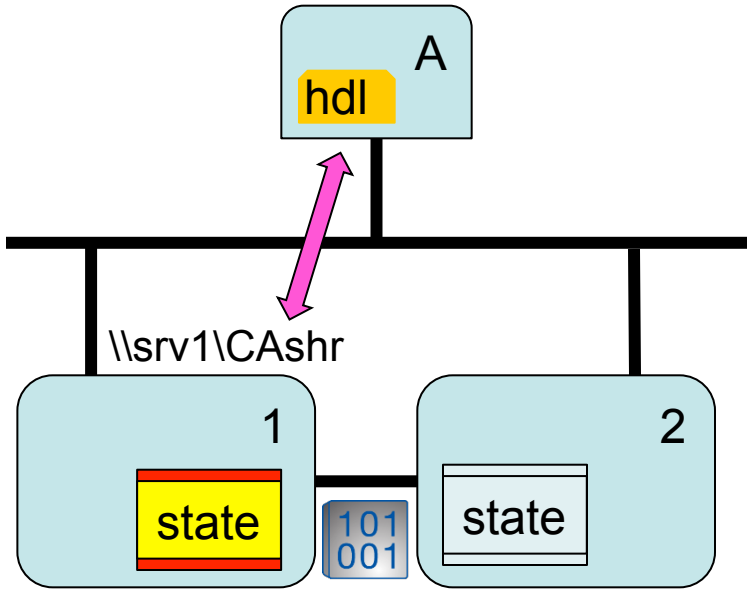

Transparent Server Failover - 3

- Server 1 offers a continuous availability (CA) share
- Client A mounts share and works
- Server 1 updates Server 2 with client A's state changes
- Server 1 fails

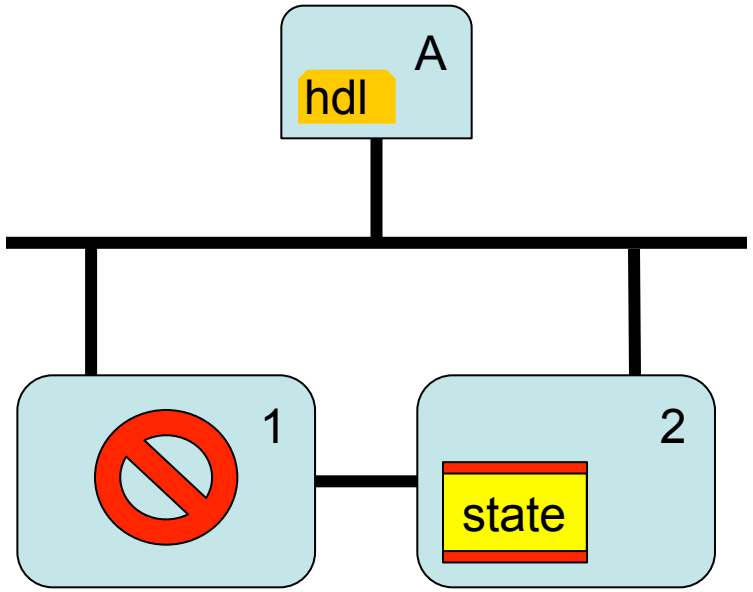

Transparent Server Failover - 4

- Client A and Server 2 auto-recover connections and handles
- No errors returned to the application/user

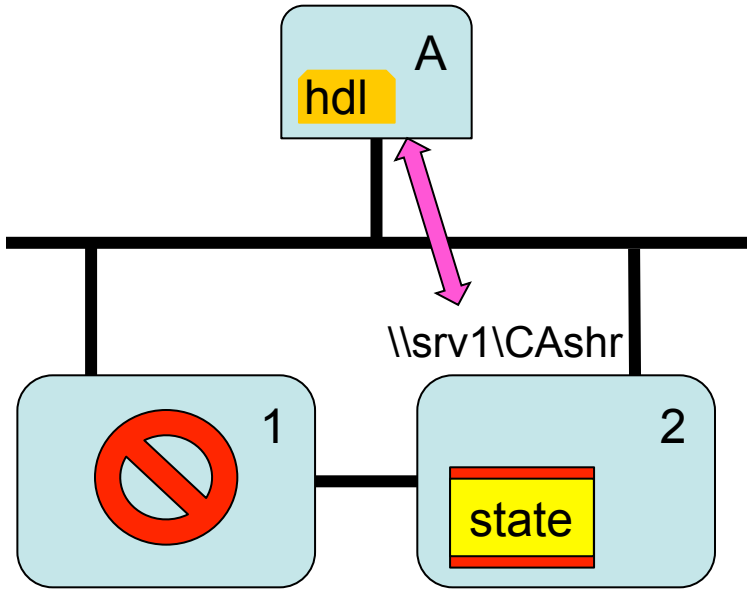

High Performance

SNIA
Ethernet Storage Forum

- Extensive LAN client caching
  - Sequence of leasing (oplocks) improvements since SMB 2.0
  - Cache both namespace and content
- Efficient use of network bandwidth and latency
  - Old SMB 1 chattiness is gone
- SMB Multichannel
  - Bandwidth aggregation with multiple NICs
  - Automatic session failover with NIC failure
- SMB Direct
  - SMB over RDMA
  - IB, iWarp & RoCE

Remote VSS overview

- Volume Shadowcopy Service coordinates an application with the OS for application consistent backups using snapshots
- Present since WS 2003 for SAN backup and restore scenarios
- Remote VSS enables this for SMB file shares

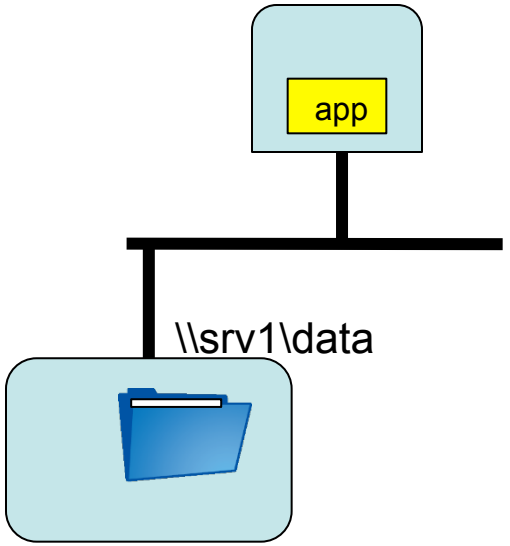

Remote VSS - 1

- Application is working modifying the storage on the server

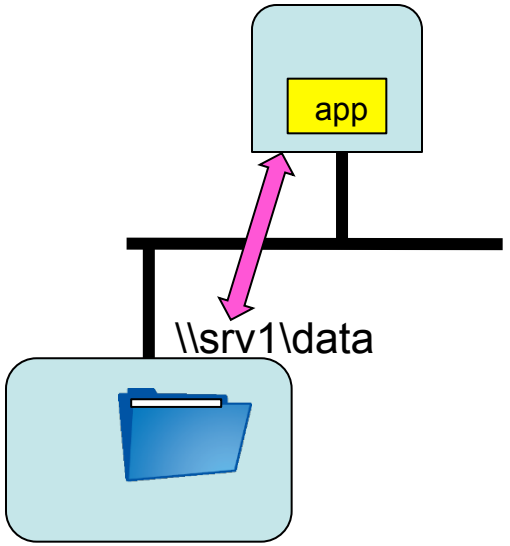

Remote VSS - 2

- Application is working modifying the storage on the server
- VSS backup sequence initiated

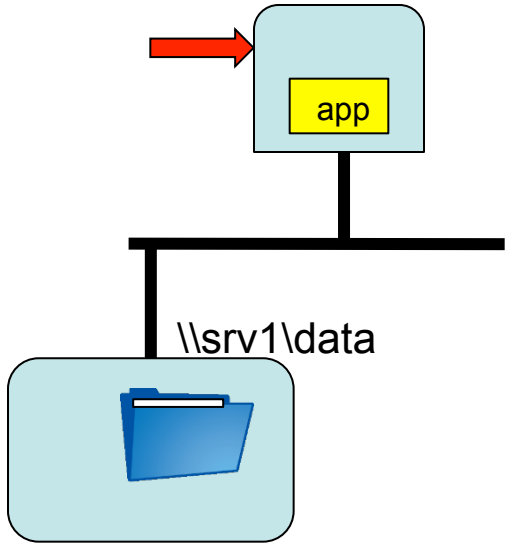

Remote VSS - 3

- Application is working modifying the storage on the server
- VSS backup sequence initiated
- Client coordinates and creates a shadow copy using a provider

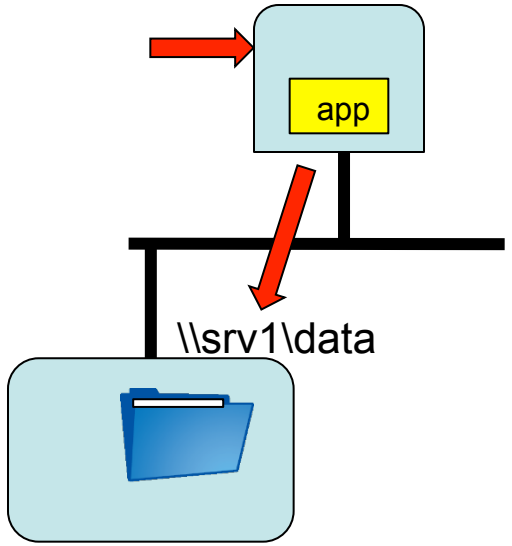

Remote VSS - 4

- Application is working modifying the storage on the server
- VSS backup sequence initiated
- Client coordinates and creates a shadow copy using a provider
- Server snapshots share contents and roots a new share on that snapshot

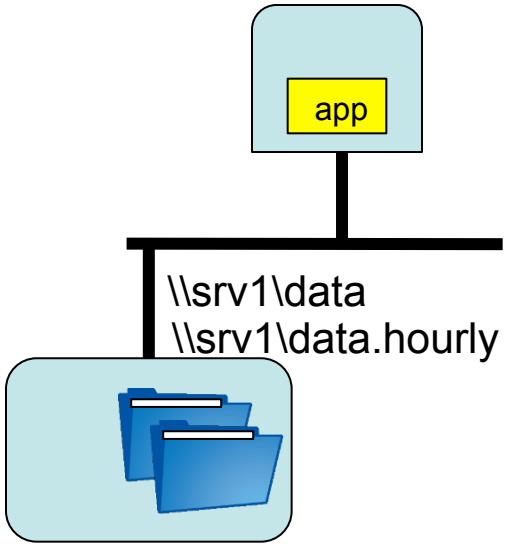

Remote VSS - 5

- Application is working modifying the storage on the server
- VSS backup sequence initiated
- Client coordinates and creates a shadow copy using a provider
- Server snapshots share contents and roots a new share on that snapshot
- Application resumes work

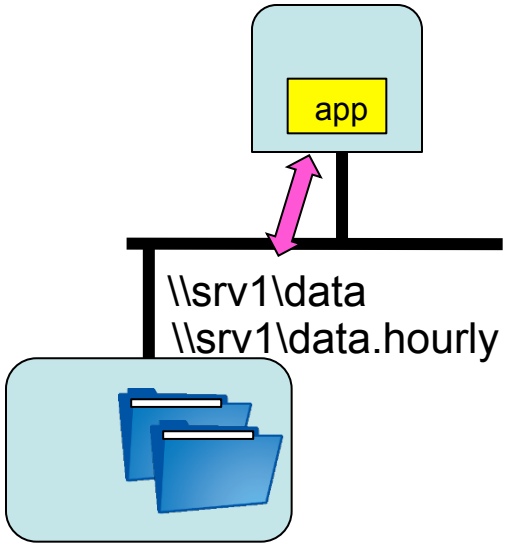

Optimized Data Transfer - ODX

- Enables server-side copies of clients' files
- Token authorization mechanism supported by iSCSI & SMB 3
- Copy files from
  - LUN to LUN
  - Share to Share
  - Share to/from LUN

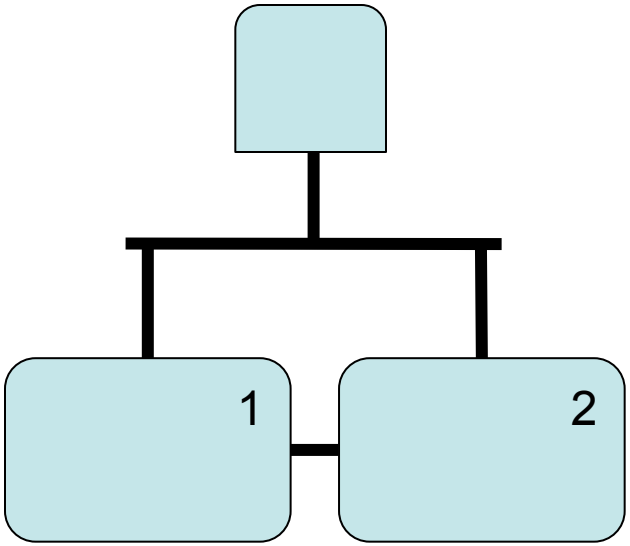

SMB ODX - 1

- Client wishes to copy the file abc from one SMB share to another

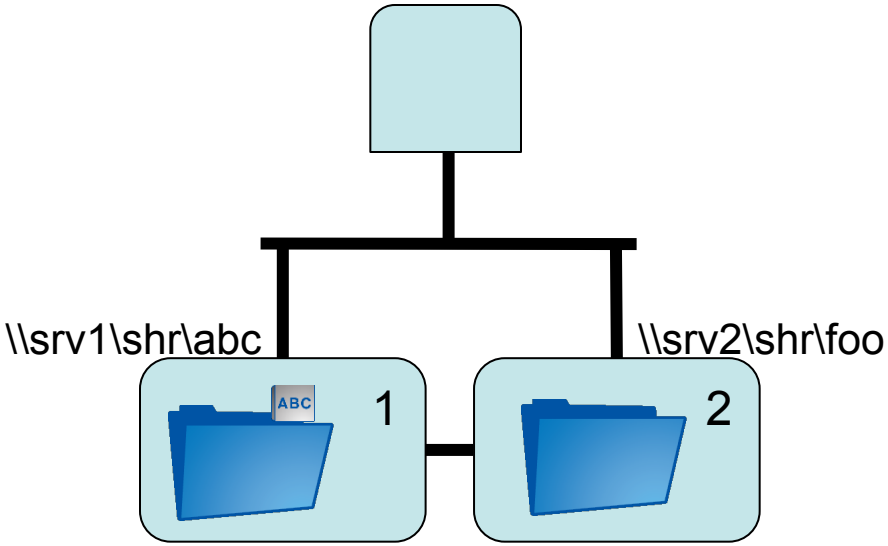

Non-Optimized Data Transfer

- Client wishes to copy the file abc from one SMB share to another
- Without ODX, file content passes through client adding latency

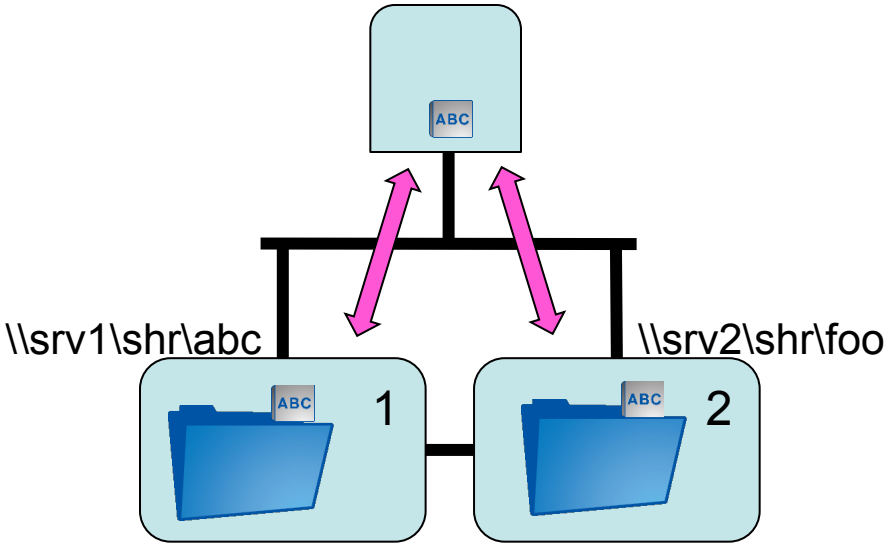

SMB ODX – 2

- Client wishes to copy the file abc from one SMB share to another
- With ODX client gets a token after file open on Server 1

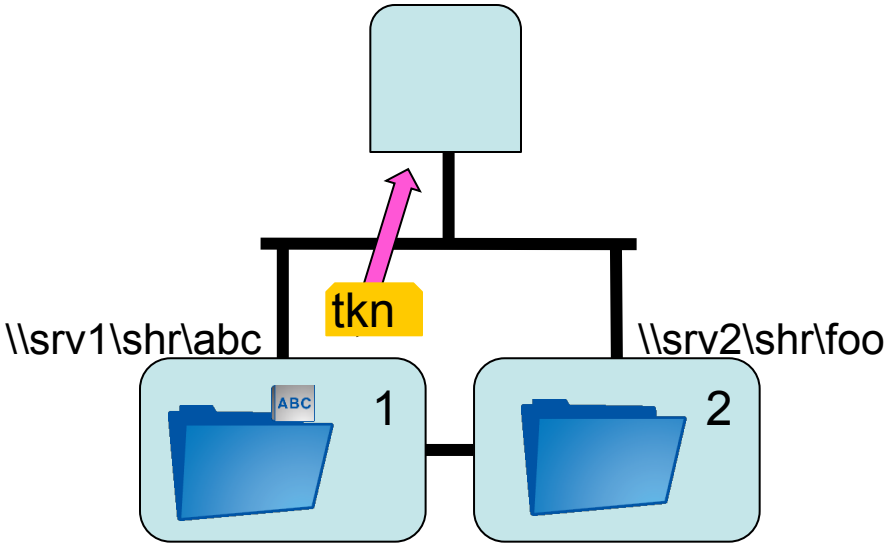

SMB ODX – 3

- Client wishes to copy the file abc from one SMB share to another
- With ODX client gets a token after file open on Server 1
- Client passes token to Server 2

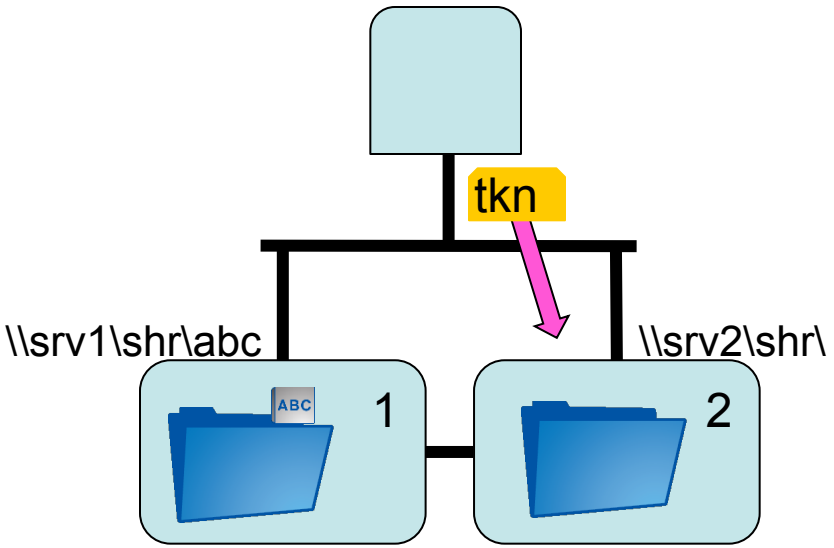

SMB ODX – 4

- Client wishes to copy the file abc from one SMB share to another
- With ODX client gets a token after file open on Server 1
- Client passes token to Server 2
- Server moderated optimized copy between shares

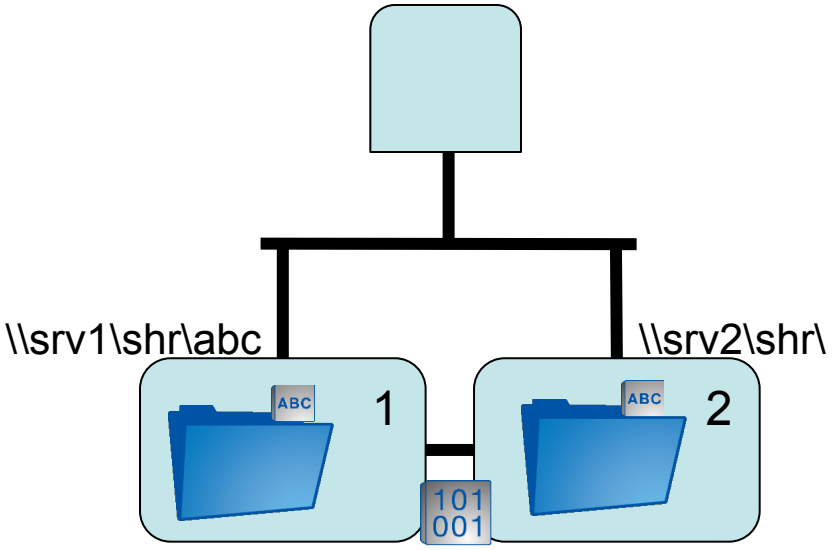

SMB SCSI ODX - 1

- Client wishes to copy the file abc from a SMB share to a mounted LUN

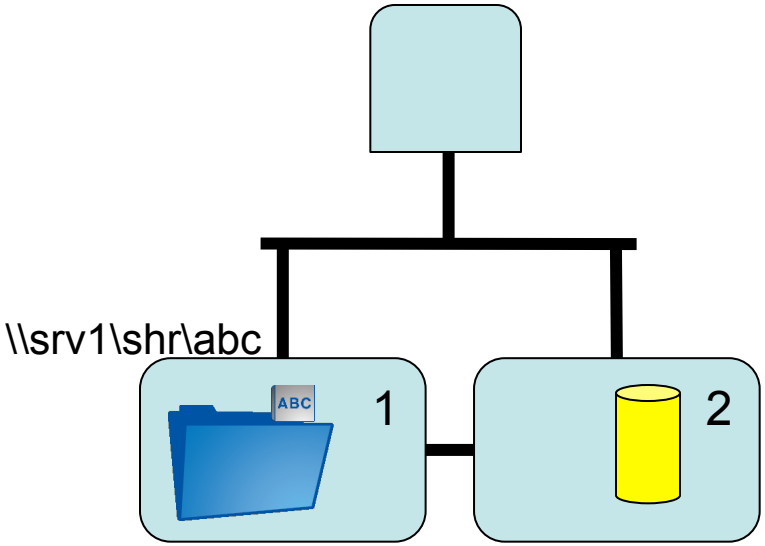

SMB SCSI ODX – 2

- Client wishes to copy the file abc from a SMB share to a mounted LUN
- Client gets a token after file open on Server 1

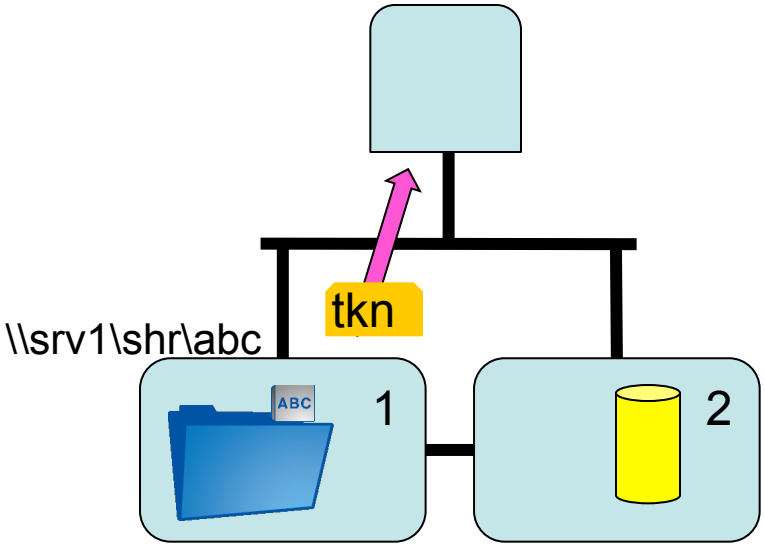

SMB SCSI ODX – 3

- Client wishes to copy the file abc from a SMB share to a mounted LUN
- Client gets a token after file open on Server 1
- Client passes token to Server 2

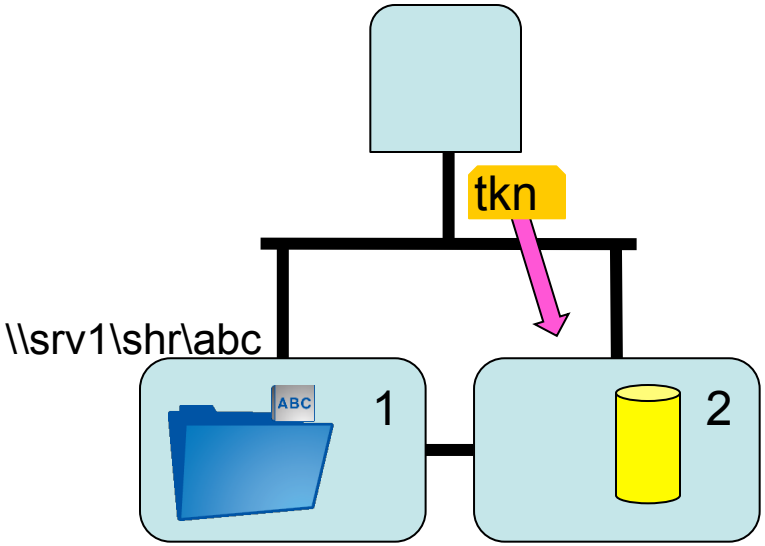

SMB SCSI ODX – 4

- Client wishes to copy the file abc from a SMB share to a mounted LUN
- Client gets a token after file open on Server 1
- Client passes token to Server 2
- Server moderated optimized copy between share and LUN

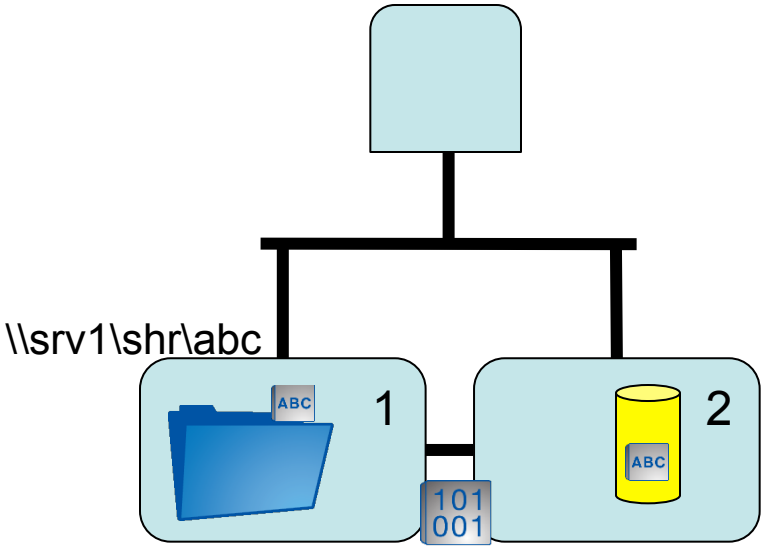

Security and Management

SNIA Ethernet Storage Forum

- Signing
  - AES-CMAC
- SMB encryption
  - End-to-end encryption of data in flight
  - AES CCM 128 bit
- PowerShell
  - WMI objects for SMB management in Windows
  - Integration with non-Windows SMB servers
- SMI-S
  - WMI objects mapped to SMI-File object model

Question…

SNIA
Ethernet Storage Forum

- Which of the following workloads are you considering deploying over SMB 3.0 in the next 12-18 months?

Agenda

SNIA
Ethernet Storage Forum

- Market Trends for File-based storage
- Evolution of the SMB Protocol
- Overview of SMB 3.0
- SMB 3.0 in action (through the lens of Hyper-V)
- Key takeaways

Quick Provisioning

- Provision via Hyper-V Manager or SCVMM
  - Point to SMB shares
- Leverage PowerShell cmdlets for automation

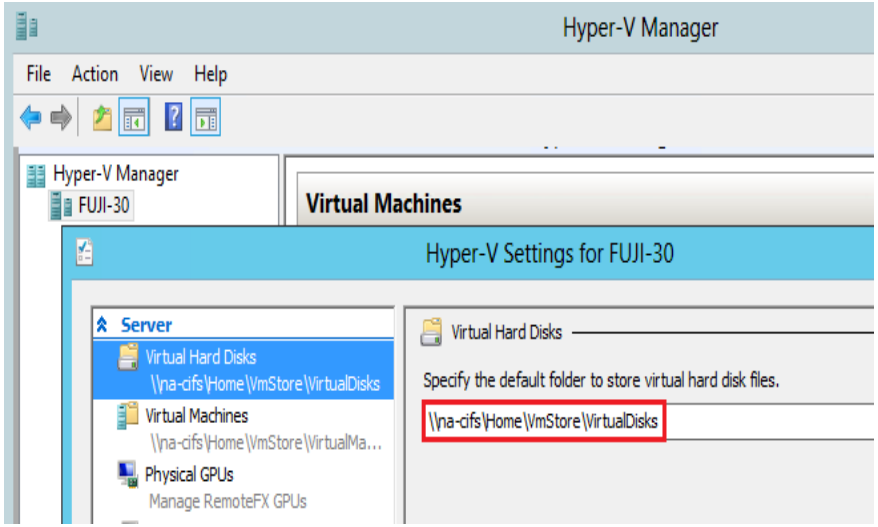

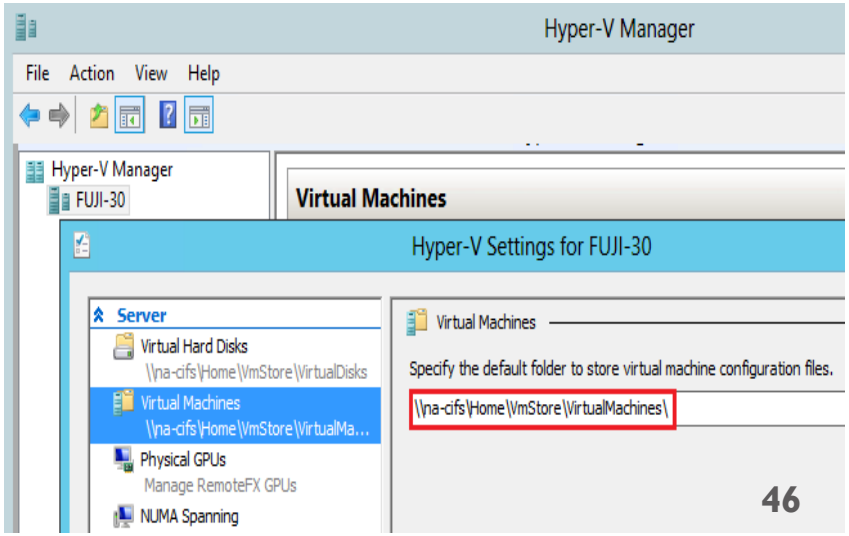

Transparency of Files – From Hyper-V Manager

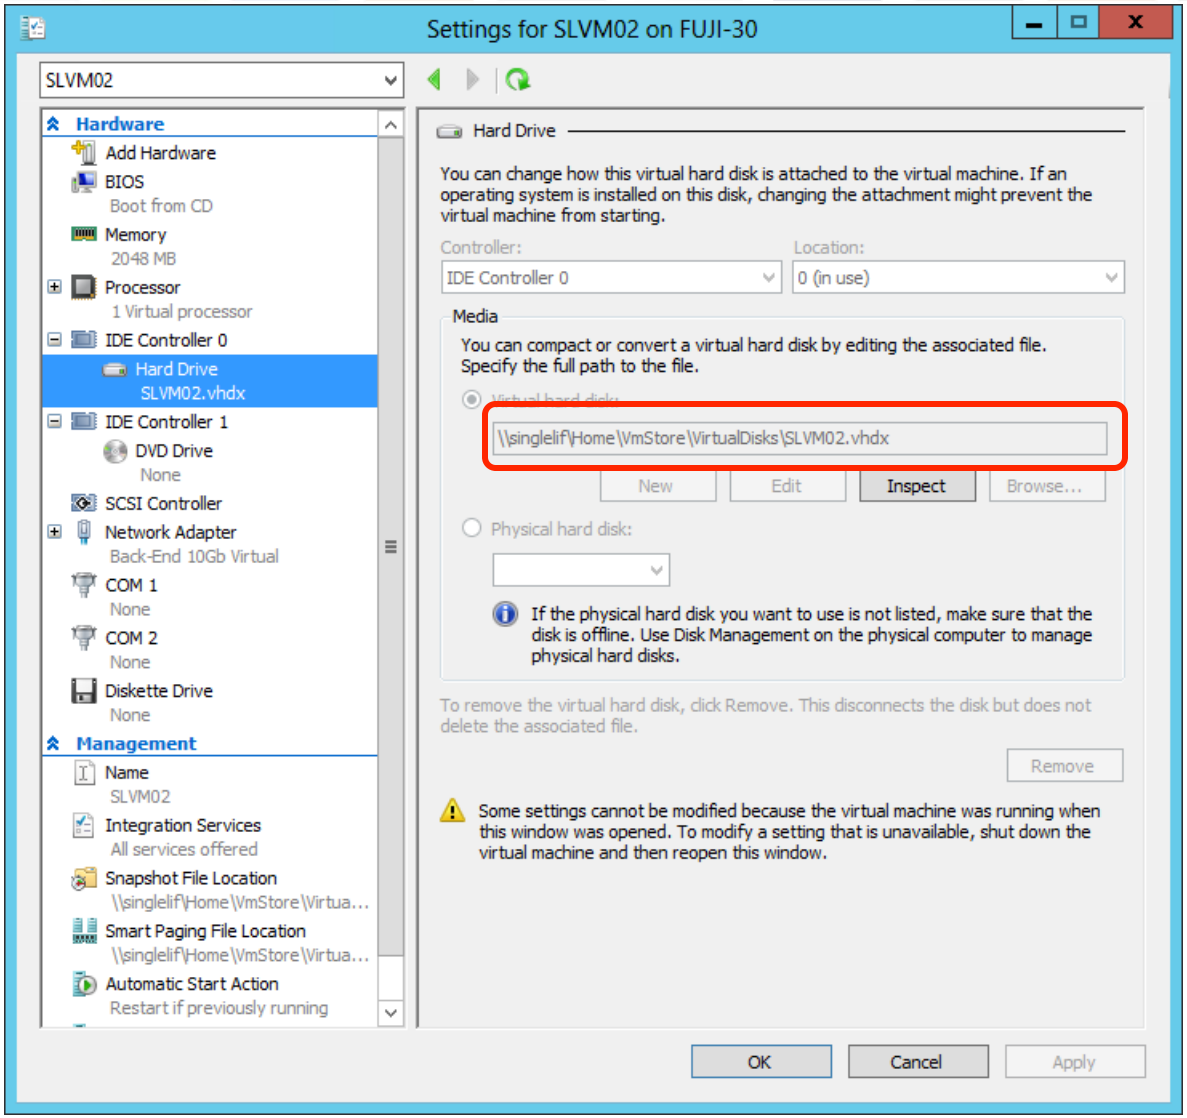

Transparency of Files – From Windows Explorer

SNIA Ethernet Storage Forum

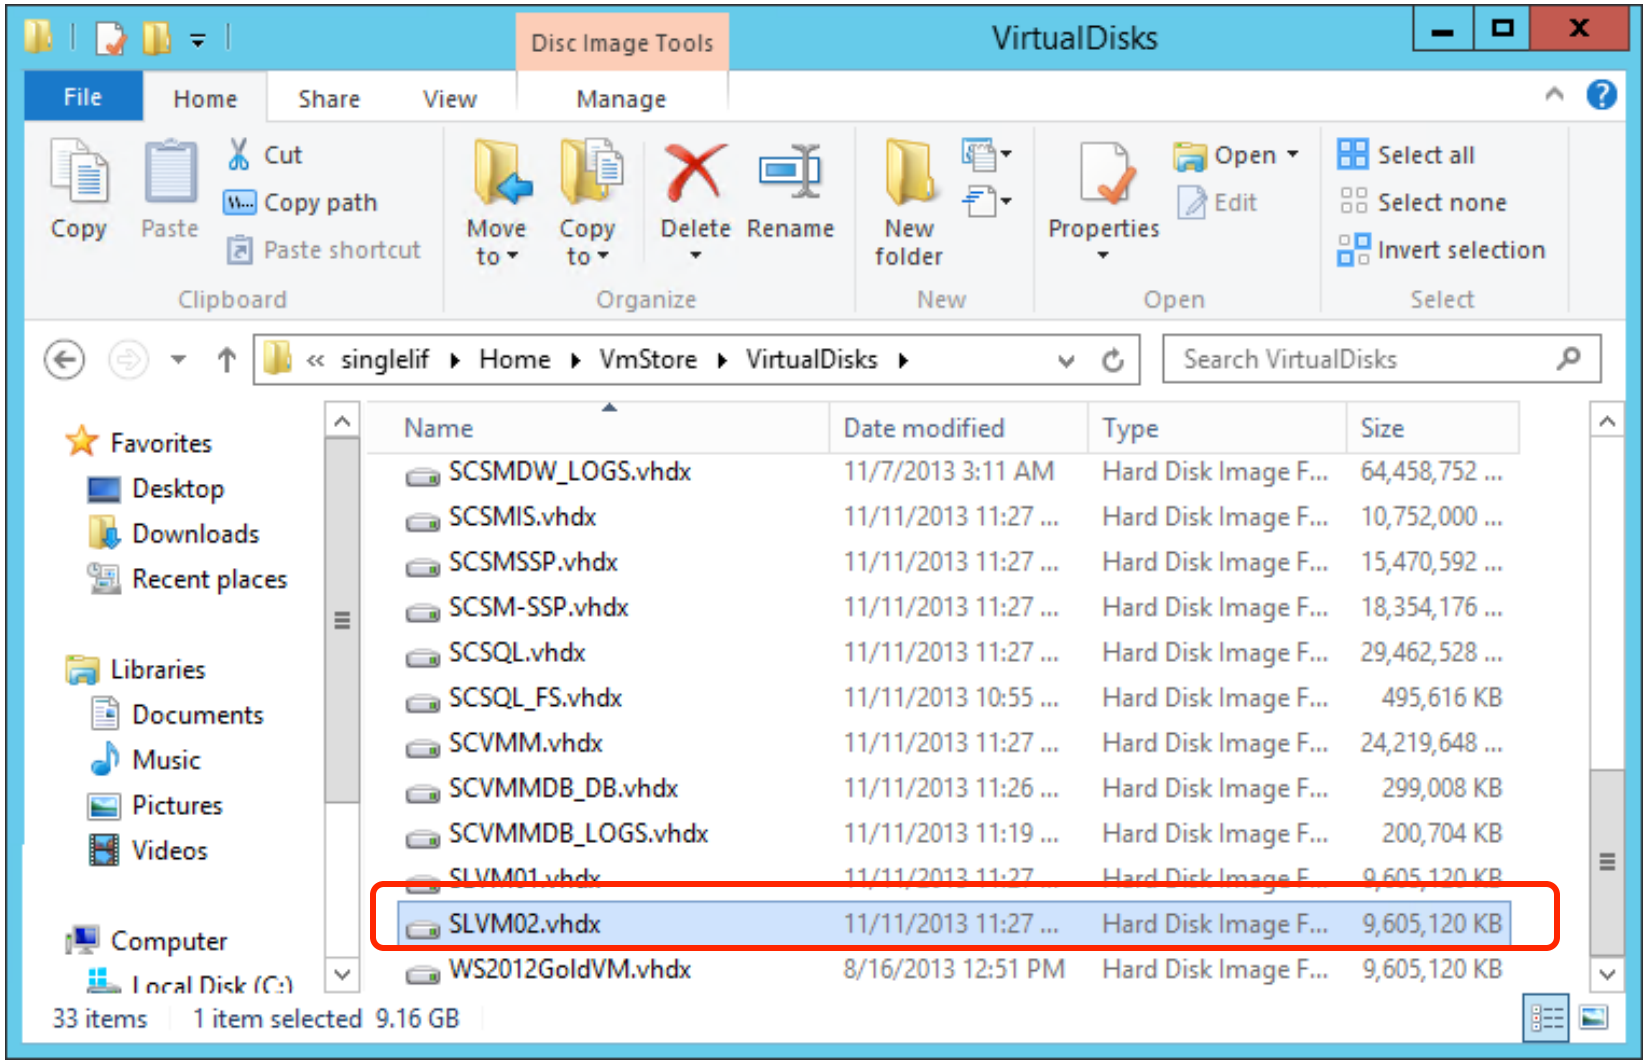

Easy Data Protection

SNIA
Ethernet Storage Forum

- Backup VM(s)
- Restore VM(s)
- Many storage vendors have dedicated applications to further streamline operations

Easy Data Protection – Backup All VMs

SNIA Ethernet Storage Forum

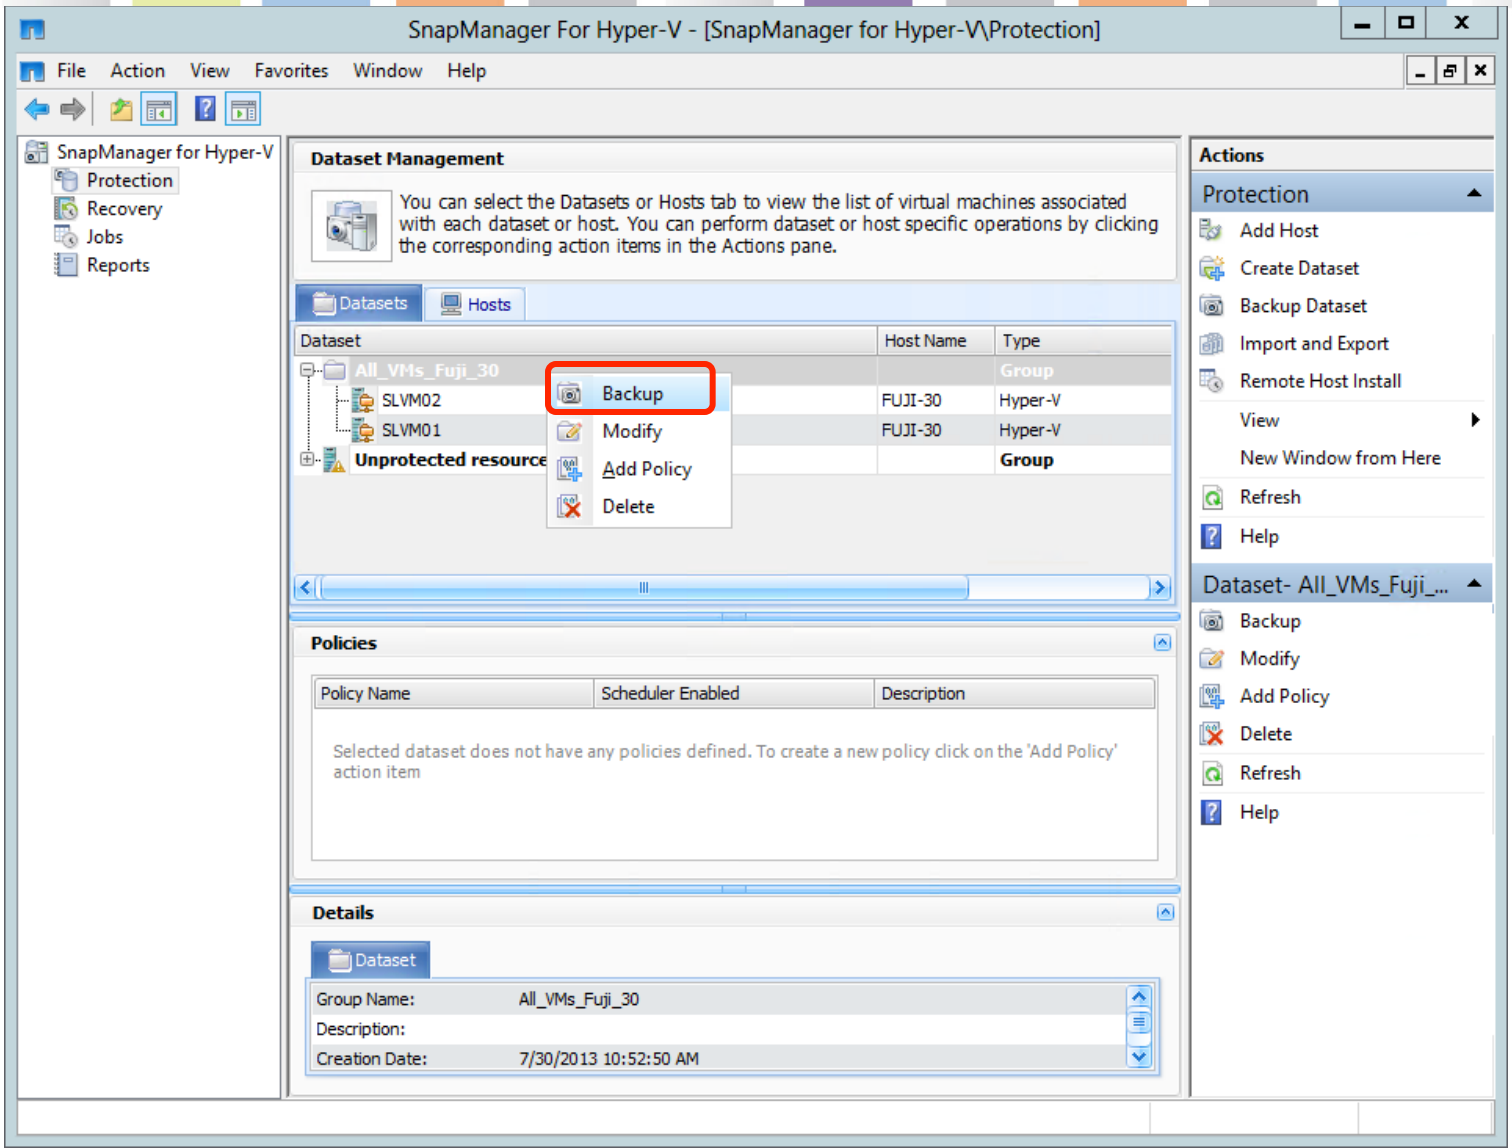

Easy Data Protection – Recover A Single VM

SNIA Ethernet Storage Forum

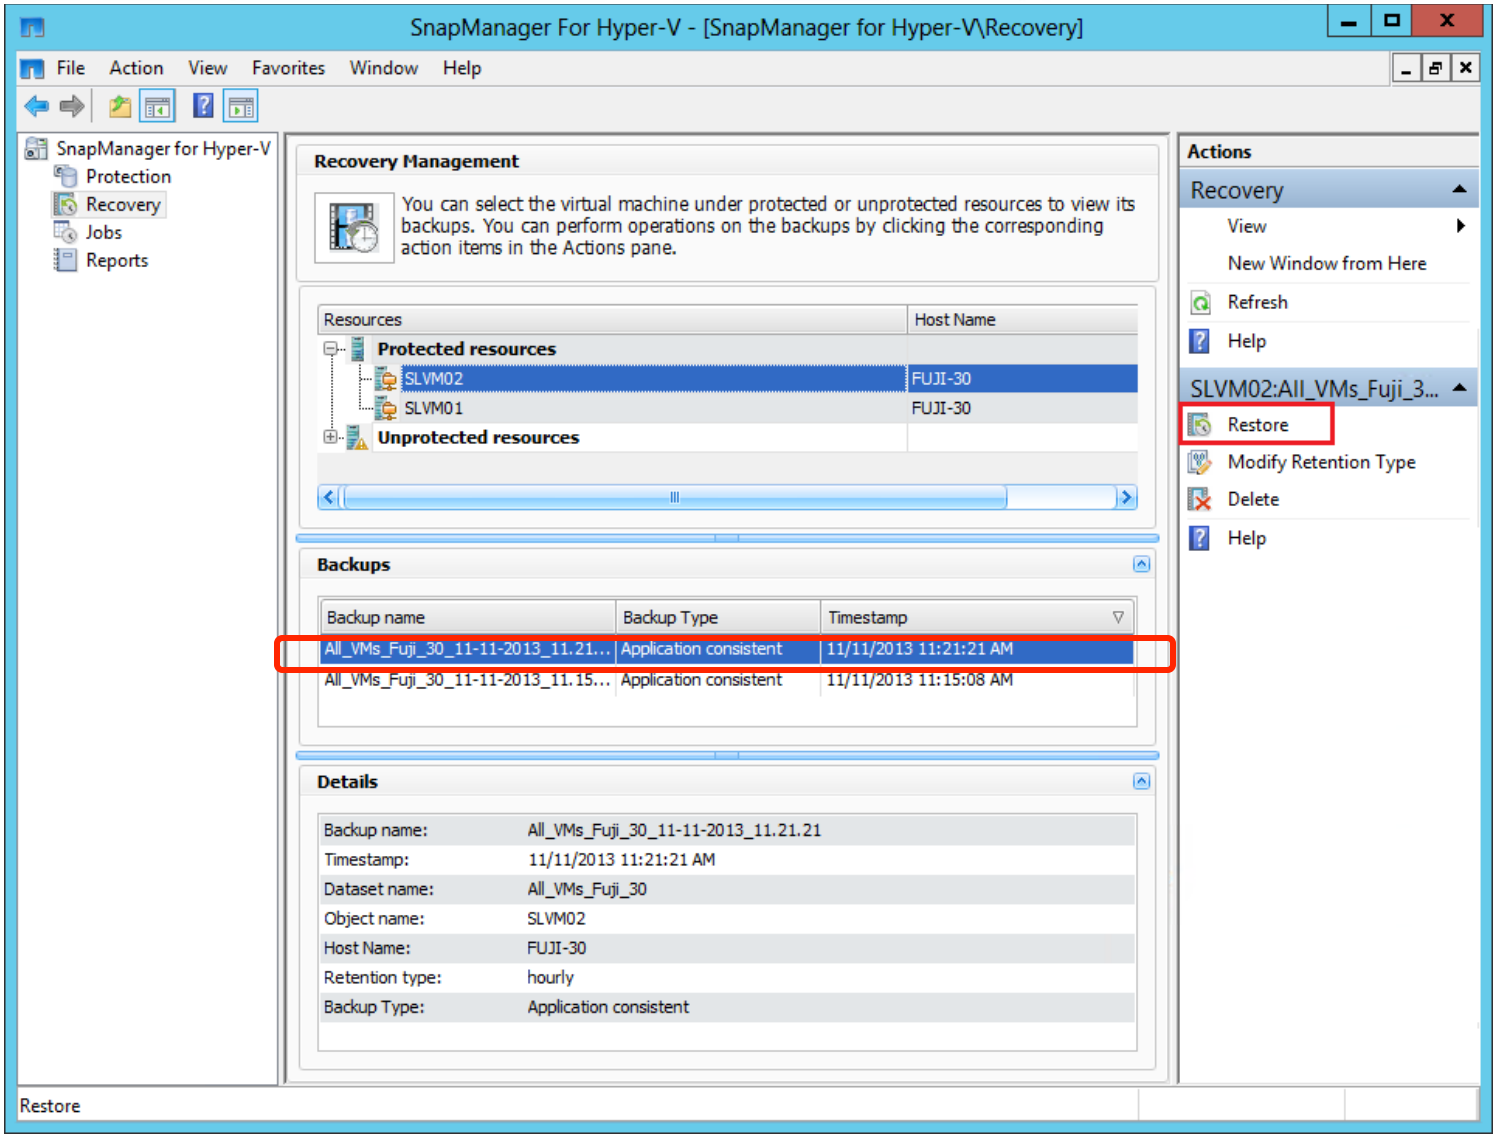

Performance Test Results - SMB 3.0 Is Comparable To FC Performance

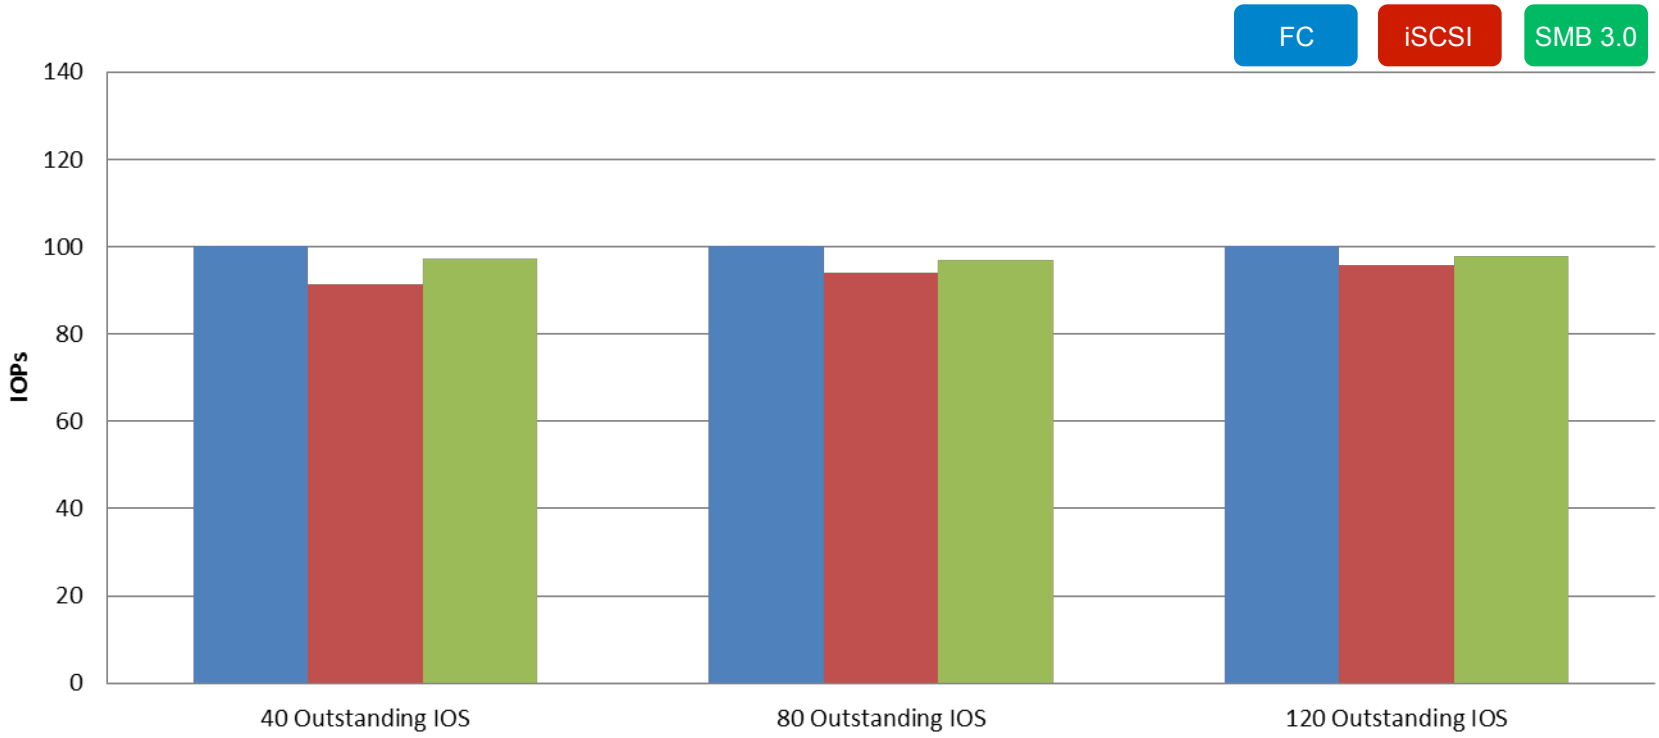

Hyper-V Protocol Performance Comparison, IOPs/sec (Normalized)
Based on a commercially available unified storage

- SMB3 throughput is 98% of FC
- SMB3 latency is 2% greater than FC

SMB 3.0 Enables Flexible IT

SNIA Ethernet Storage Forum

- Reduce complexity with additional design options
  - Ability to create different service offerings based on storage backend, ie file- based and block-based
- Reduce operational costs with heterogeneous data centers
  - Primary DC with block-based and the secondary DC with file-based

Key Takeaways

SNIA
Ethernet Storage Forum

- File-based storage will continue its upwards trajectory as a data store
- SMB 3.0 has the technology and capability to support this trend. It is real and ready.
- SMB 3.0 offers a new way of thinking in your designs for data centers and application storage

Additional Material

SNIA Ethernet Storage Forum

- Brought to you by SNIA ESF File Protocols Special Interest Group
- Additional Material
  - SNIA BrightTalk Channel
    - https://www.brighttalk.com/webcasts?q=SNIA
  - SNIA ESF (Ethernet Storage Forum)
    - http://www.snia.org/forums/esf
  - SNIA ESF Blog
    - http://sniaesfblog.org/
  - *"SMB remote file protocol (including SMB 3.0)"*
    - http://www.snia.org/sites/default/education/tutorials/2012/fall/file/JoseBarreto_SMB3_Remote_File_Protocol_revision.pdf
  - *"The Future of File Protocols: SMB3 Meets Linux"*
    - http://www.snia.org/sites/default/files2/SDC2012/presentations/Revisions/SteveFrench_Linux_CIFS-SMB2-year-in-review-revision.pdf
- Plugfests & Interoperability
  - SDC2013 SMB2/SMB3 Plugfest
    - https://www.snia.org/events/storage-developer2013/plugfest#smb