Four Reasons To Start Working With NFSv4.1 Now

SNIA™ WEBCAST

Presented by:
Alex McDonald

Hosted by:
Gilles Chekroun

HOSTED BY THE
ETHERNET STORAGE FORUM

Ethernet Storage Forum Members

The SNIA Ethernet Storage Forum (ESF) focuses on educating end-users about Ethernet-connected storage networking technologies.

SNIA Europe™

Webcast Presenter

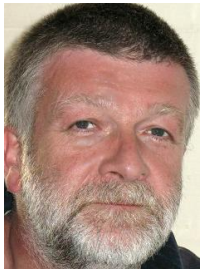

Alex McDonald
Office of the CTO
NetApp

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 .

With a background in software development, support, sales and a period as an independent consultant, 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 NFS Special Interest Group.

Gilles Chekroun
Distinguished Engineer
Cisco

Gilles joined Cisco 18 years ago. For the last ten years, Gilles' focus has been Storage & SAN extension technologies for designing and implementing Disaster Recovery Centers.

Gilles is now dedicated to Data Center Technologies like Unified Fabric, FCoE and Unified Computing System and is a member of the Cisco Europe Data Centre and Virtualisation Team as a Distinguished Engineer. He is a member of the Board of Directors of SNIA Europe (Storage Networking Industry Association) as Technical Chair.

SNIA Europe™

SNIA's NFS Special Interest Group

Education

- NFS SIG drives adoption and understanding of pNFS across vendors to constituents
  - Marketing, industry adoption, Open Source updates
- NetApp, EMC, Panasas and Sun founders
  - NetApp, EMC and Panasas act as co-chairs
- White paper on migration from NFSv3 to NFSv4
  - "Migrating from NFSv3 to NFSv4"

SNIA Ethernet Storage Forum

Learn more about us at: www.snia.org/forums/esf

SNIA Europe

NFS; Ubiquitous & Everywhere

- NFS is ubiquitous and everywhere
- NFS doesn't stand still
  - NFSv2 in 1983, through NFSv4.1 in 2010
  - NFSv4.2 to be agreed at IETF shortly
  - Faster pace for minor revisions
- NFSv3 very successful
  - Protocol adoption is over time, and there have been no big incentives to change

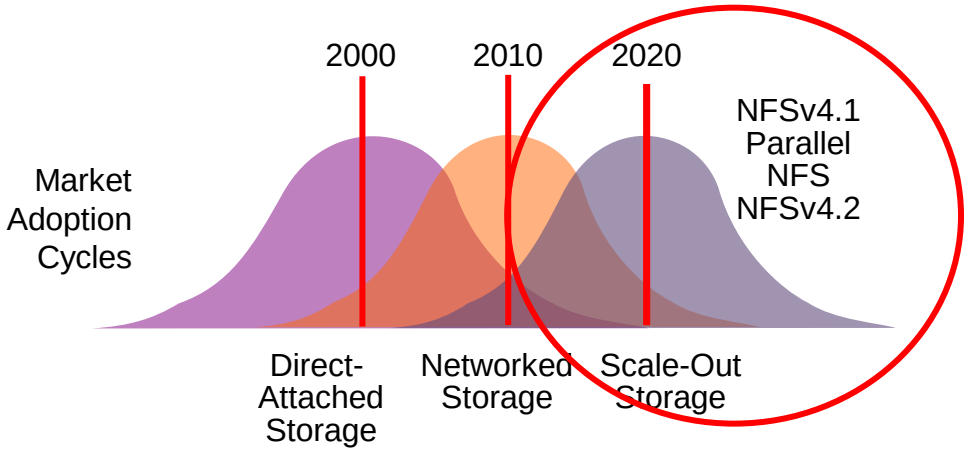

Evolving Requirements

- Economic Trends
  - Cheap and fast computing clusters
  - Cheap and fast network (1GbE to 10GbE, 40GbE and 100GbE in the datacenter)
  - Cost effective & high performance storage based on Flash & SATA
- Performance
  - Exposes single threaded bottlenecks in applications
  - Increased demands of compute parallelism and consequent data parallelism
- Powerful compute systems
  - Analysis begets more data, at exponential rates
  - Competitive edge (ops/sec)
- Business requirement to reduce solution times
  - Beyond performance; NFS 4.1 brings increased scale & flexibility
  - Outside of the datacenter; requires good security

Performance, Management and Reliability

Education

- Random I/O and Metadata intensive workloads
  - Memory and CPU are hot spots
  - Load balancing limited to pair of NFS heads; originally designed for HA
    - Not a limitation of the NFS 4.1 protocol
- Compute farms are growing larger in size
  - NFS head can handle a 1000+ NFS clients
  - NFS head hardware comparable to client CPU, I/O, Memory
  - NFS head requires more spindles to distribute the I/O
- Reliability and availability are challenging
  - Data striping limited to single head and disks
  - Non-disruptive upgrades affect dual-head configurations
  - Access and connectivity is typically limited to a pair of NFS server heads

NFSv4 Major Features; Security

- Strong security framework
- Access control lists (ACLs) for security and Windows® compatibility
- Mandatory security with Kerberos
  - Negotiated RPC security that depends on cryptography, RPCSEC_GSS

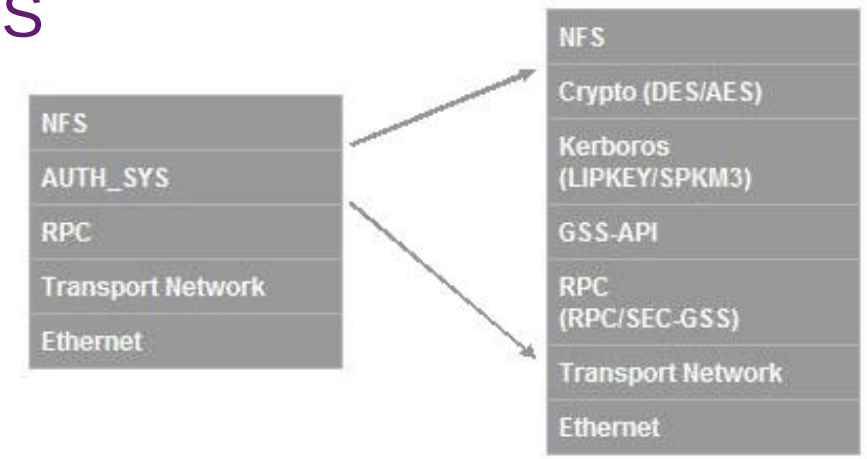

NFSv4 Major Features; Namespace

- Uniform and “infinite” namespace
  - Moving from user/home directories to datacenter & corporate use
  - Meets demands for “large scale” protocol
  - Unicode support for UTF-8 codepoints
- No automounter required
  - Simplifies administration

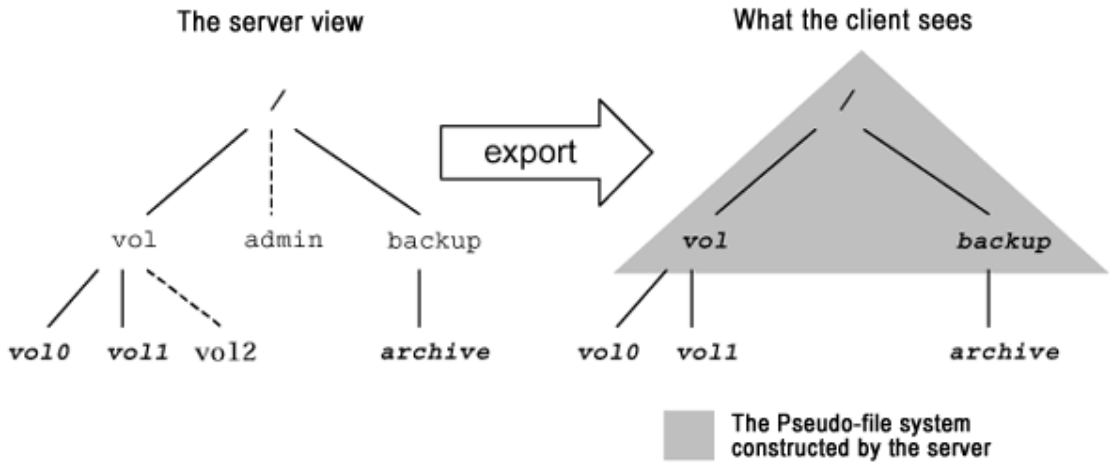

NFSv4 Major Features; Stateful Clients

- NFSv4 gives client independence
  - Previous model had “dumb” stateless client
  - Server had the smarts
- Pushes work out to client through delegations & caching
- Why?
  - Compute nodes work best with local data
  - NFSv4 eliminates the need for local storage
  - Exposes more of the backend storage functionality
    - Client can help make server smarter by providing hints

NFSv4.1 Major Features; Sessions

- NFSv3 server never knows if client got reply message
- NFSv4.1 introduces Sessions
  - Major protocol infrastructure change
  - Exactly Once Semantics (EOS)
  - Bounded size of reply cache
  - Unlimited parallelism
- A session maintains the server's state relative to the connections belonging to a client

NFSv4.1 Major Features; Layouts

- Layouts
  - Files, objects and block layouts
  - Provides flexibility for storage that underpins it
  - Location transparent
    - Striping and clustering
- Examples
  - Blocks, Object and Files layouts all available from various vendors

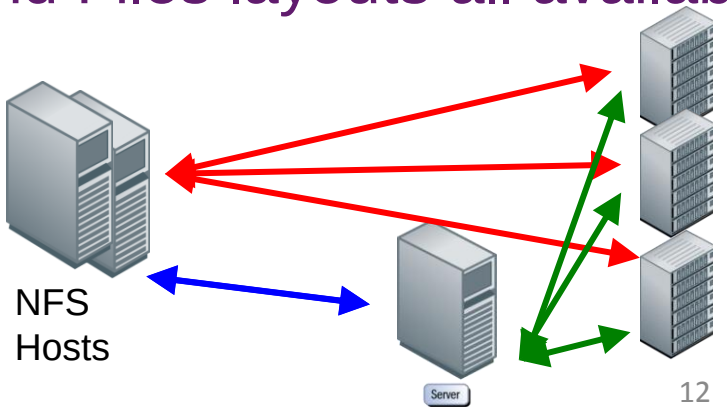

NFSv4.1 Major Features; pNFS

- NFSv4.1 (pNFS) can aggregate bandwidth
  - Modern approach; relieves issues associated with point-to-point connections

- pNFS Client
  - Client read/write a file
  - Server grants permission
  - File layout (stripe map) is given to the client
  - Client parallel R/W directly to data servers
- Removes IO Bottlenecks
  - No single storage node is a bottleneck
  - Improves large file performance
- Improves Management
  - Data and clients are load balanced
  - Single Namespace

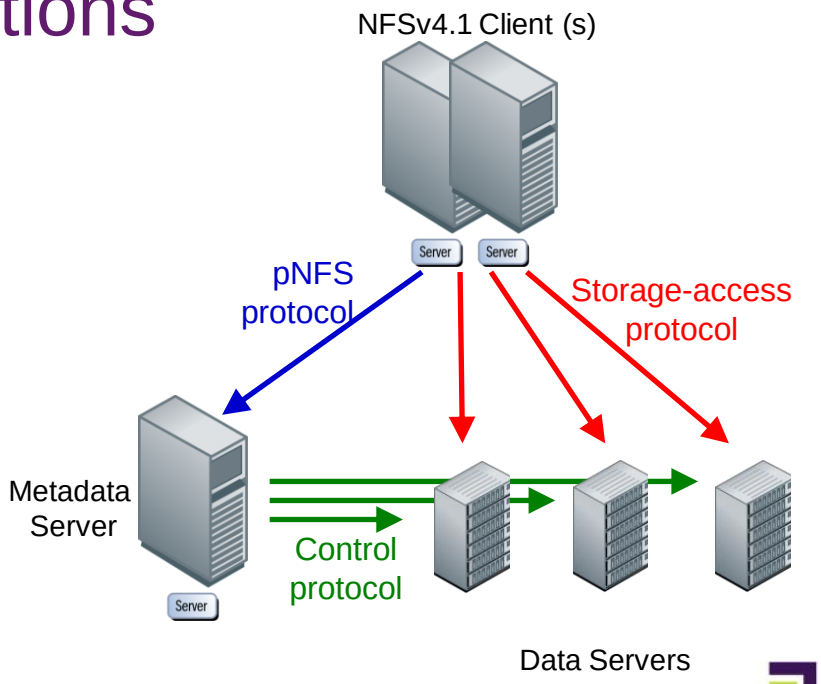

Linux Client and NFSv4.1

- Upstream (Linus) Linux NFSv4.1 client support
  - Basic client in Kernel 2.6.32
  - pNFS support (files layout type) in Kernel 2.6.39
  - Support for the 'objects' and 'blocks' layouts was merged in Kernel 3.0 and 3.1 respectively
- Full read and write support for all three layout types in the upstream kernel,
  - O_DIRECT reads and writes now supported.

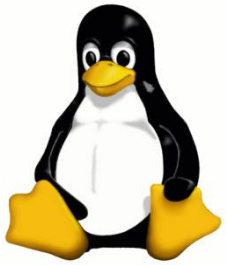

SNIA Europe

Linux Client and NFSv4.1

- pNFS client support in distributions
  - Fedora 15 was first for pNFS files
  - Kernel 2.6.40 (released August 2011)
- Red Hat Enterprise Linux version 6.2, 6.3
  - “Technical preview" support for NFSv4.1 and for the pNFS files layout type
  - Full support in RHEL6.4
- Other Open Source
  - Microsoft NFSv4.1 Windows client from CITI

Filesystem Implications

- Files, blocks, objects can co-exist in the same storage network
  - Can access the same file system; even the same file
- NFS flexible enough to support unlimited number of storage layout types
  - Three IETF standards, files, blocks, objects
  - Others evaluated experimentally
- NAS vs SAN; no-one cares any more
  - IETF process defines how you get to storage, not what your storage looks like
  - Underlying pNFS implementations will vary substantially between storage vendors

Virtualization; The Game Changer

- Server virtualization a major area of use
  - VMware, Citrix Xen
- Demands of 1000s of images on 100s of servers
- Requirements from a storage system
  - Single system image, resiliency, load balanced, transparent & non-disruptive upgrades…
- NFS a good fit in virtualized environments
  - Matches well datacenter use cases
  - NFS widely available and ubiquitous

NFS & The Virtualized Datacenter

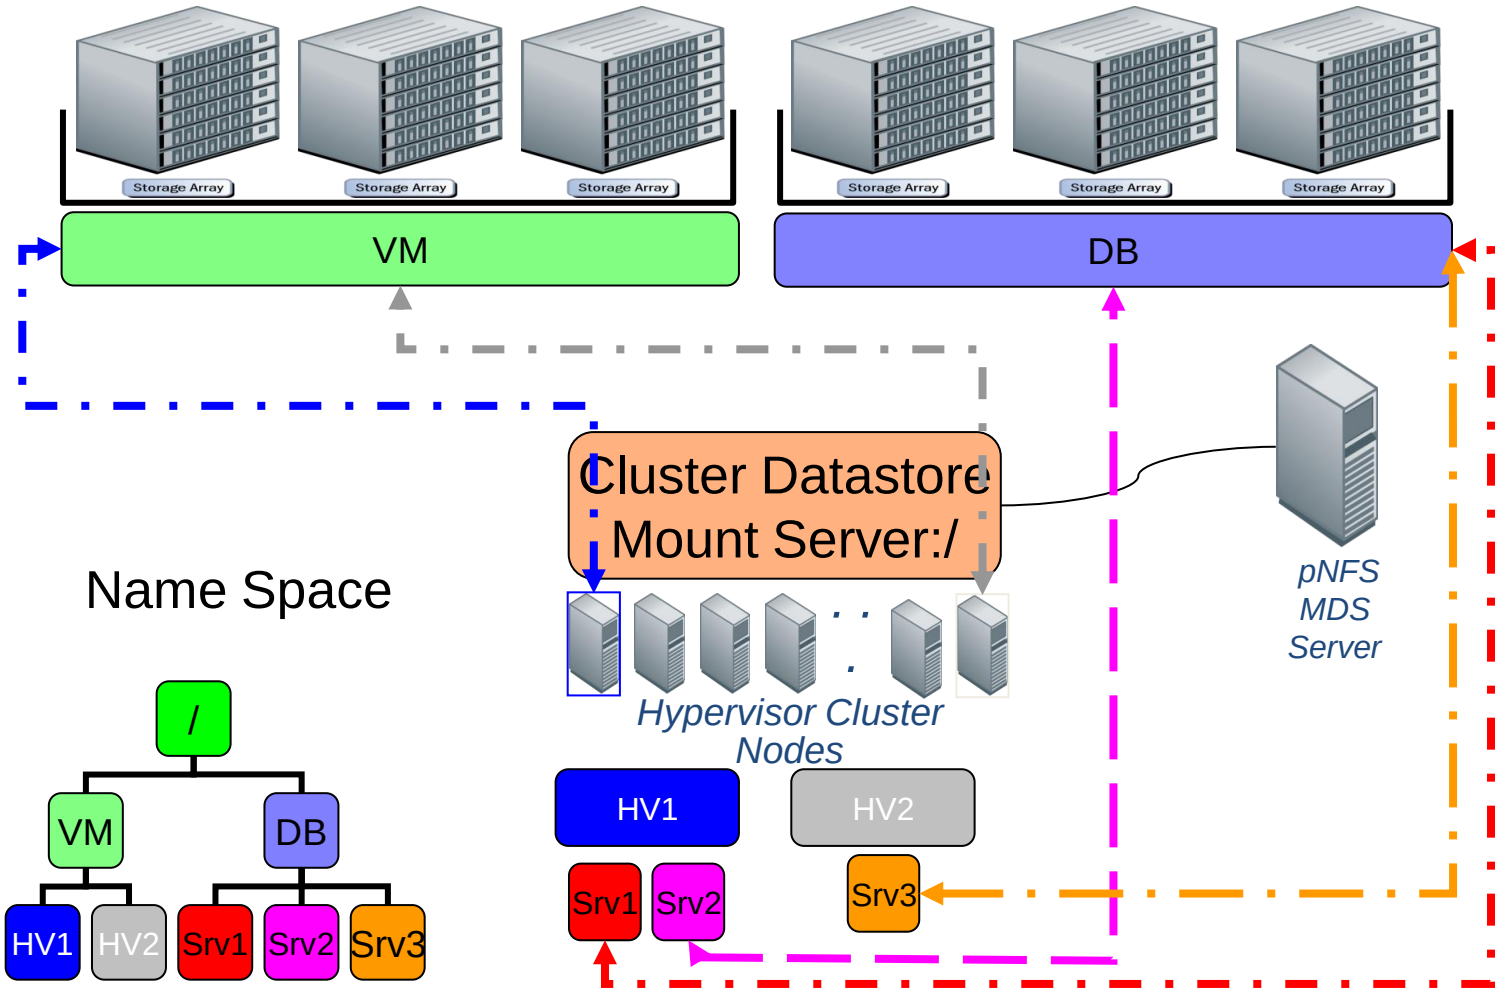

Conclusions

- NFS has more relevance today for commercial, HPC and other use cases than it ever did
  - Features for a virtualized data centers
- Developments driven by application requirements
- Adoption slow, but will continue to increase
  - NFSv4 support widely available
  - New NFSv4.1 with client & server support
  - NFS defines how you get to storage, not what your storage looks like

The Four Reasons for NFSv4.1

	Functional	Business Benefit
Security	ACLs for authorization Kerberos for authentication	Compliance, improved access, storage efficiency, WAN use
High availability	Client and server lease management with fail over	High Availability, Operations simplicity, cost containment
Single namespace	Pseudo directory system	Reduction in administration & management
Performance	Multiple read, write, delete operations per RPC call Delegate locks, read and write procedures to clients Parallelised I/O	Better network utilization for all NFS clients Leverage NFS client hardware for better I/O

Summary/Call to Action

- pNFS is the first open standard for parallel I/O across the network
  - Ask vendors to include NFSv4.1 support for client/servers
- pNFS has wide industry support
  - commercial implementations and open source
- Start using NFSv4.1 today
  - NFSv4.2 nearing approval
- Future BrightTalk Sessions
  - NFSv4.1 — Plan for a Smooth Migration
  - Advances in NFS; pNFS, FedFS and NFSv4.2

Question & Answer

To download this Webcast

after the presentation, go to

http://www.snia.org/about/socialmedia/