SNIA ESF | ETHERNET STORAGE

# What Is NFS?
# A Brief Introduction

Chad Hintz, Cisco
Alex McDonald, NetApp

**23 March 2016**

# SNIA Legal Notice

SNIA ESF | ETHERNET STORAGE

- 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.

# Today's Presenters

SNIA ESF | ETHERNET STORAGE

**Chad Hintz**
**Board Member**
**SNIA Ethernet Storage Forum**
**Cisco**

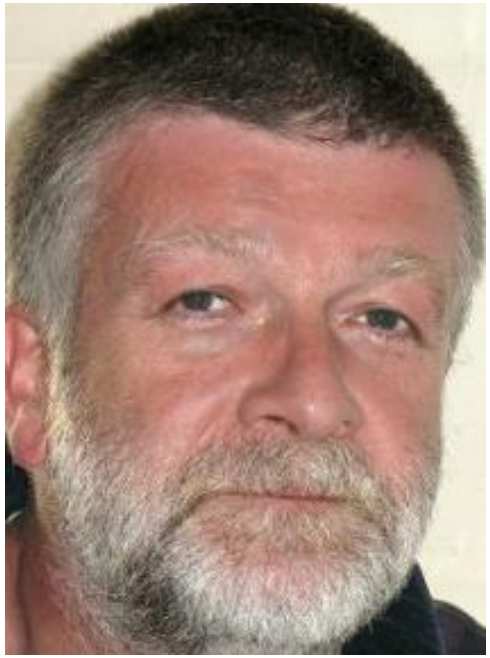

**Alex McDonald**
**Vice Chair, SNIA Ethernet Storage Forum**
**Chair, SNIA Cloud Storage Initiative**
**NetApp**

# A Polling Question

SNIA ESF | ETHERNET STORAGE

- How would you rate your NFS knowledge?
  - Frankly, it's a bit of a mystery
  - I know a little
  - I'm comfortable with NFS & I use it pretty regularly
  - I'm an NFS expert
  - I'm an NFS developer

# Overview



## About this presentation

- This is a **101-level introduction; it will be light on detail**
- NFS features discussed are NFS version 4.0+

- What Is a File System?
- A Brief NFS History; Distributed File Systems
- The Components of NFS
- Use Cases for NFS
- More Information & Where to Start

# What Is a File System?



- Characteristic of a file
  - Byte addressable
  - Randomly accessible
  - Named
    - But IO operations through file handle
- Operations
  - Opening and closing
  - Reading, updating and writing
  - Creating, renaming and deleting

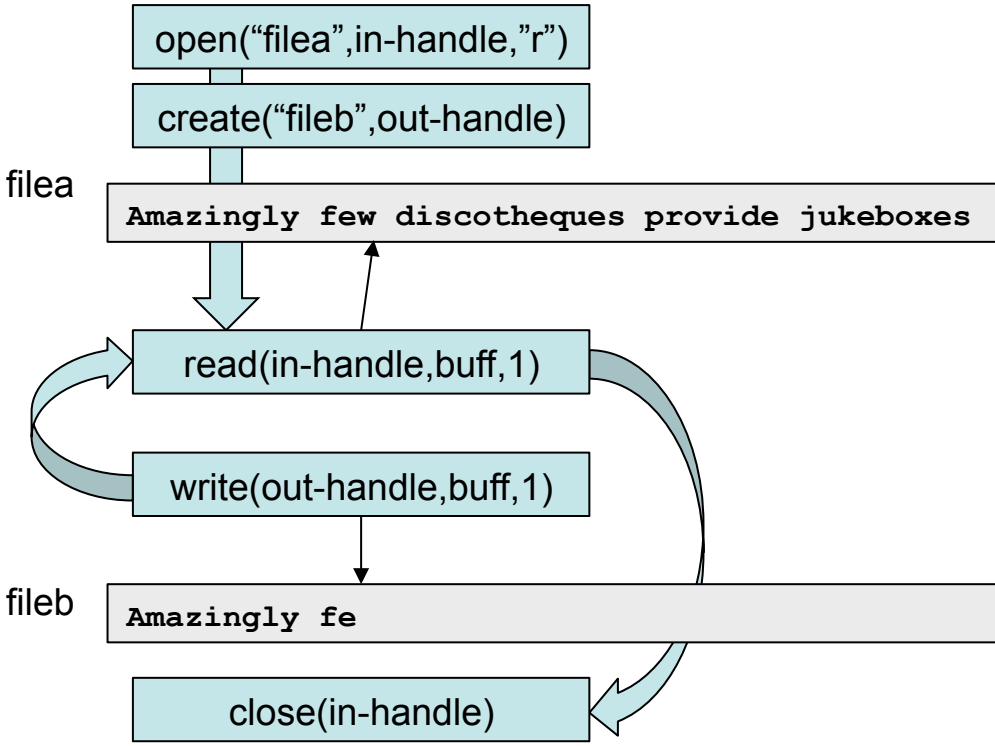

# What Is a File System?



- Organization
  - Rooted tree of directories and files
  - Names and paths
    - Fully qualified name: /home/bill/docs/a.txt

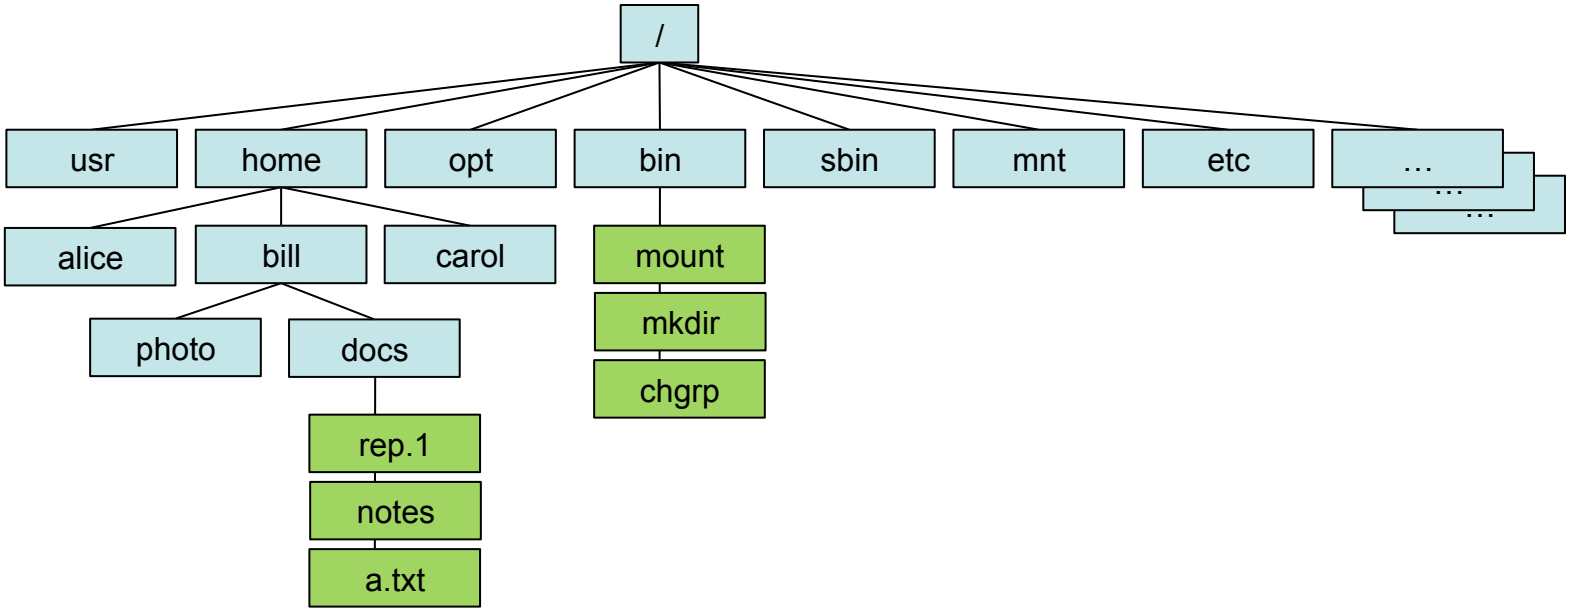

# The Application View



- POSIX or close to POSIX-like support in all major OSes
- Same system calls used allows source code portability
- Applications can treat files as
  - Streams or objects; a set of unstructured data
  - Structured data; sets of discrete units contained in a file
  - Block; set of randomly accessible blocks
  - …

POSIX™ defines a standard operating system interface and environment … to support applications portability at the source code level
http://pubs.opengroup.org/onlinepubs/9699919799/

# Overview

SNIA ESF | ETHERNET STORAGE

- What Is a File System?
- A Brief NFS History; Distributed File Systems
- The Components of NFS
- Use Cases for NFS
- More Information & Where to Start

# Distributed File Systems



- Characteristics of Distributed File Systems
  - Make distributed look exactly like local file system
- Key is transparency
  - Access & location
  - Consistency & concurrency
  - High level of tolerance to failure
  - Heterogeneous, scalable, replicatable, migratable
- Uses RPCs (Remote Procedure Calls) & network protocol
- No knowledge required of underlying structures; nothing “pokes through” to the end user

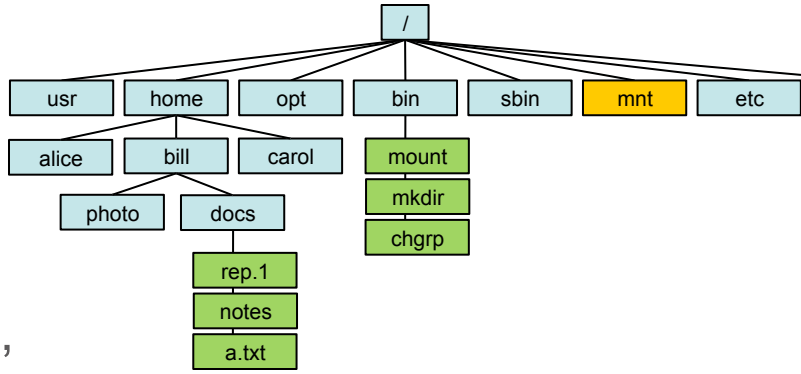

# A (Very!) Brief NFS History



- Plenty of "bake time"
  - 20+ years of NFS
  - NFSv3 - 1995, NFSv4 - 2000, NFSv4.1 – 2010, NFSv4.2 - ~2016
- Internet Engineering Task Force (IETF) standardization
  - Standardization = no surprises, no special interoperability rules to follow
    - RFC 1810 – NFSv3
    - RFC 3530 – NFSv4
    - RFC 5661 – NFSv4.1
    - NFSv4.2 RFC still in draft - https://tools.ietf.org/html/draft-ietf-nfsv4-minorversion2-41

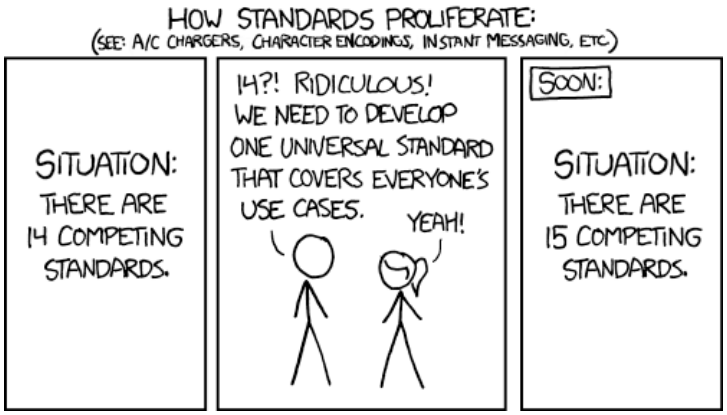

# SMB vs NFS



- Similar design goals but different starting points
- Different enough to make it confusing
  - Different file handle semantics
  - Permission structures and ACLs
    - SIDs/GUIDs vs UID/GIDs
  - Share access modes (Shares vs. Exports)
    - NFS allows client-based fencing
  - User and group identifiers
  - Path name construction (eg case sensitivity)
  - Standardization process; “de jure”
  - Some commercial solutions support dual access via SMB and NFS
- SMB has influenced NFS security implementation
  - Particularly authorization with ACLs

# Overview

SNIA ESF | ETHERNET STORAGE

- What Is a File System?
- A Brief NFS History; Distributed File Systems
- The Components of NFS
- Use Cases for NFS
- More Information & Where to Start

# NFSv4 Namespace

- NFSv4 can “stitch together” directories into pseudo-file system
  - Uniform and “infinite” namespace

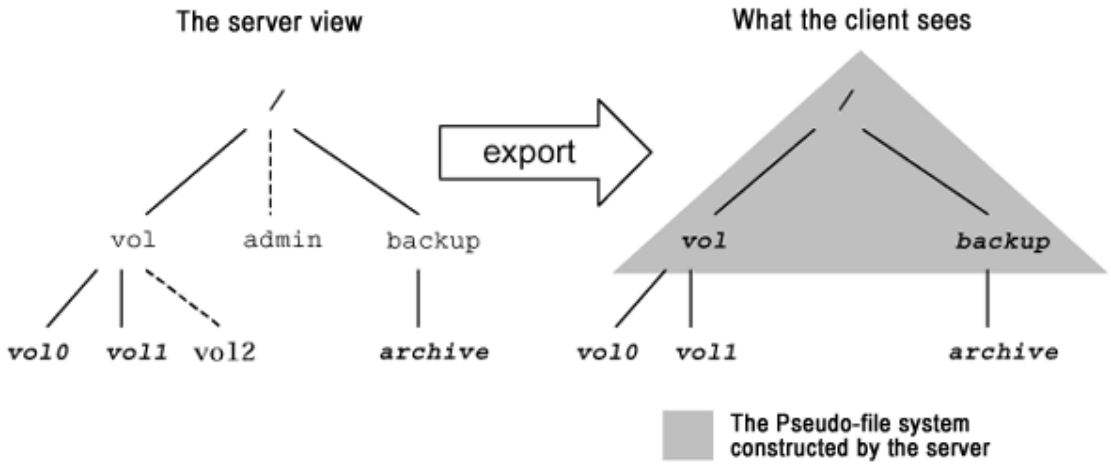

# Servers



- Servers run daemons
  - Several processes; one is the **nfsd** daemon
    - handles client requests for file system operations
  - **nfsd** uses entries in special file **/etc/exports**, or is contacted via **exportfs** command
  - **mount** command is used to “insert” a file system inside the file tree

```
mkdir /exports
mkdir /exports/opt
mkdir /exports/etc
mount --bind /usr/local/opt /exports/opt
mount --bind /usr/local/etc /exports/etc
exportfs -o fsid=0,insecure,no_subtree_check gss/krb5p:/exports
exportfs -o rw,nohide,insecure,no_subtree_check gss/krb5p:/exports/opt
exportfs -o rw,nohide,insecure,no_subtree_check gss/krb5p:/exports/etc
```

# Servers cont…



- Builds and exports a pseudo file system
  - For permanent exports, use **/etc/exports**
- Do all systems operate similarly?
  - Commercial NFS servers differ
    - See vendor doc
  - Linux servers vary for managing services
  - All Linux commands (like mount, exportfs and syntax of /etc/exports) are the same
  - At end of this presentation info for Linux servers

# Client Side

- Client mounts the file system

```
mount -t nfs4 -o sec=krb5 nfs-server:/ /mnt
```

- Now exported filesystem from nfs-server is visible under /mnt
  - /mnt contains files & directories from /exports; that is /mnt/opt and /mnt/etc
- Some applications have their own clients
  - Oracle, VMware etc

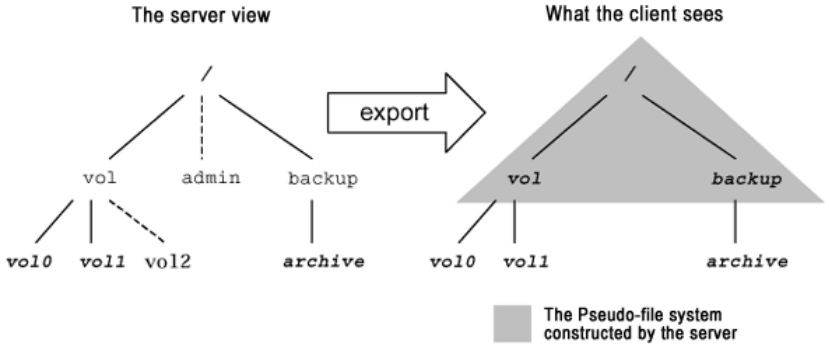

# Security



- Security a much misunderstood area
- NFS can use Kerberos
  - “network authentication protocol which works on the basis of 'tickets' to allow nodes communicating over a non-secure network to prove their identity to one another in a secure manner” https://en.wikipedia.org/wiki/Kerberos_%28protocol%29

| Authentication | Authorization |
| --- | --- |
| Proving to a system that you are who you claim to be<br>▪ Name mapping<br>▪ Kerberos tickets | What a user is allowed to do once in a system<br>▪ Group membership<br>▪ Access Control Lists (ACLs) |

**Terms are often confused with each other; not interchangeable**

# Overview



- What Is a File System?
- A Brief NFS History; Distributed File Systems
- The Components of NFS
- Use Cases for NFS
- More Information & Where to Start

# Use Cases for NFS



- Application support across the board
  - Oracle, VMware ESXi, SAS Grid, SAP HANA, TIBCO, OpenStack, Docker, etc
  - VDI & database workloads
- Wide & readily available clients
- Simple, easy to manage
- No client OS file system needed
  - Easier to expand, shrink storage
  - Easier to move storage from client to client
- No special drivers, cards, switches required
  - Fewer moving parts to deal with
  - Runs over ubiquitous Ethernet
- Performance
  - Low latency
  - High throughput
  - Better performance than iSCSI in most cases
  - No expensive FC components needed

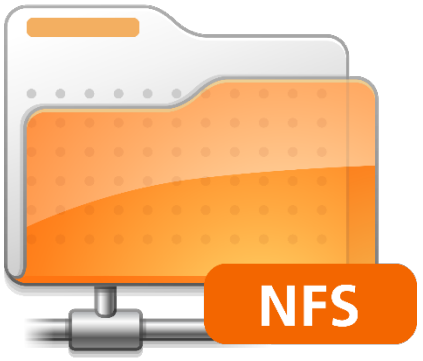

# Overview



- What Is a File System?
- A Brief NFS History; Distributed File Systems
- The Components of NFS
- Use Cases for NFS
- More Information & Where to Start

# More Information & Where to Start

SNIA ESF | ETHERNET STORAGE

- How I practiced on NFS
  - Laptop; 16GB and 4+ cores preferable
  - Install VMware Desktop Player on Windows
  - Install copies of Linux
    - I have used Ubuntu and Fedora
    - Desktop versions are OK,
      but check server packages installed
  - Assign an image as a server
    - Will also support Kerberos, but start without it (see links)
    - Build exportable directories with some useable data
  - Assign another as a client
    - Mount & inspect
  - Clone
    - Play with multiple client VMs, multiple server VMs
  - Install Kerberos on one server

# More Information & Where to Start

SNIA ESF | ETHERNET STORAGE

- VMware Workstation Player
  - http://www.vmware.com/uk/products/player
- Distributions
  - Fedora https://getfedora.org/
  - RedHat (eval) https://access.redhat.com/downloads
  - Ubuntu http://www.ubuntu.com/download/desktop
  - SUSE (eval) https://www.suse.com/products/server/download/amd64.html?gclid=CIKR1sv71MsCFUI_GwodblQIRQ
- NFS documentation
  - RedHat https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Storage_Administration_Guide/ch-nfs.html
  - SUSE https://www.suse.com/documentation/sles11/book_sle_admin/data/cha_nfs.html
  - Ubuntu https://help.ubuntu.com/lts/serverguide/network-file-system.html
- NFS Howtos
  - Oldie but good; Kereberos-free NFSv4 http://www.vanemery.com/Linux/NFSv4/NFSv4-no-rpcsec.html
  - Fedora, more up to date https://fedoraproject.org/wiki/User:Renich/HowTo/NFSv4
  - Ubuntu equivalent https://help.ubuntu.com/community/NFSv4Howto
- Kerberos
  - Many available; search on Google!

# After this Webcast

SNIA ESF | ETHERNET STORAGE

- Please rate this Webcast and provide us with feedback
- This Webcast, and a PDF of the slides, will be posted to the SNIA Ethernet Storage Forum (ESF) website and available on-demand http://www.snia.org/forums/esf/knowledge/webcasts
- A full Q&A from this webcast, including answers to questions we couldn't get to today, will be posted to the SNIA-ESF blog http://sniaesfblog.org/
- Follow us on Twitter @SNIAESF

SNIA ESF | ETHERNET STORAGE

# Thank You