



Storage Grid Backup & NDMP

Rudy Nedved


Connectathon 2006

- ▶ **Grid Computing**
- ▶ **Who cares?**
- ▶ **New problems.**
- ▶ **Some ideas.**

- ▶ **“Too Short” Version**
 - 1000s of cheap computers
 - Similar production done by many.
 - “High point” demands can shift daily.

- ▶ **Longer Version**
 - Everything is a computer and we have more and more of them!
 - Different “rush hours” for different things.
 - New tricks done to solve old problems.
 - Can’t afford schedule downtime
 - Everything done in “bulk”

Storage Grid Picture



- ▶ **Replace “Grid” with “Farm”**
- ▶ **“Product” with hundreds of disks together**
- ▶ **Examples**
 - **Andrew File System (Carnegie Mellon University)**
 - **Spinnaker Networks**
 - **Panasas**
 - **Network Appliance (based on Spinnaker)**


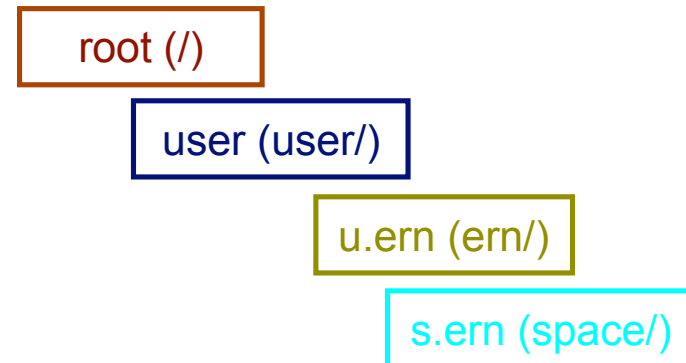
- ▶ **Storage Grooming**
 - Volumes separated from disks
 - Lots of volumes not 1 or 2
 - Oversubscribing disk space
 - Non-Disruptive Volume movement
- ▶ **Non-Disruptive Head Replacement**
 - Add a new head, move volumes
- ▶ **Single System Image**
 - One “command” to bind them all!
- ▶ **Add CPU, Network, Disks & Bandwidth Easily**

- ▶ **Instead of backing up /vol/vol0 and /vol/vol1**
- ▶ **We now have 400 to 1000 or more volumes**
 - **Do users really want to specify them all?**
 - **What happens with the 1 or 2 that fail out of the 1000? Do all 1000 get retried?**
 - **How does different backup service get implemented?**
 - **If a storage pool (aggregate) must be restored, are the volumes required spread out through out the tape set or collated together in the smallest number of tapes?**

- ▶ **Why does a volume move?**
 - The ethernet connection fails.
 - The NAS head fails.
 - The SAN connection fails.
 - The administrator moves hot volumes off of hot disks.
 - The administrator moves small volumes off of low space disks.

- ▶ **Backup sees a volume disappear on host A and a volume appear on host B**
 - Middle of the week full backup?

- ▶ 20,000 volumes on 20 NAS heads
- ▶ Administrator talks about volumes like “vol0”
- ▶ Users talk about paths like “/user/ern”
- ▶ How do you find where to look in the DMA for “/user/ern/space/work/p4/src/fred.cc”?



- ▶ **NDMP support of Unix-like “mount table”**
 - **DMAAs can use associate pathnames with volumes?**
- ▶ **Add signature to volumes to detect volume moves.**
- ▶ **Augment NDMP_GET_FS_INFO to provide standard attributes for disk storage location and others to aid in sorting.**
- ▶ **Augment DMA file system selection for backups.**
- ▶ **Other ideas?**