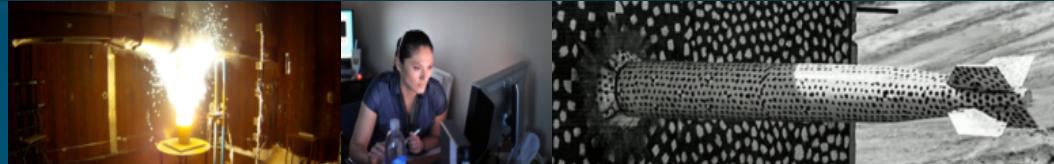




Sandia
National
Laboratories

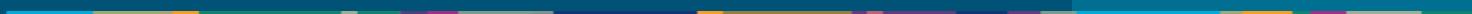
SAND2021-6014C

Migrating data from system to system is a piece of cake with Starfish



Presented by

Justin Wood, HPC Solutions Architect





- The HPC Systems department supports and manages multiple production clusters within several different network security domains
 - 4 systems on the Top500
 - 16 clusters, over 15,600 compute nodes
- Several global filesystems on each network, total capacity of ~48PB
 - Lustre (5)
 - GPFS (2)
 - NFS (6)
- Mixture of different storage hardware vendors (currently 5)

Our experiences so far



- Challenges

- Moving large amounts of data using traditional methods was labor intensive
- Tried a commercial product from another vendor, it couldn't handle the size/number of files, and was painful trying to get WOS data to migrate. Also had no lustre changelog/monitor support.

- Successes (using Starfish)

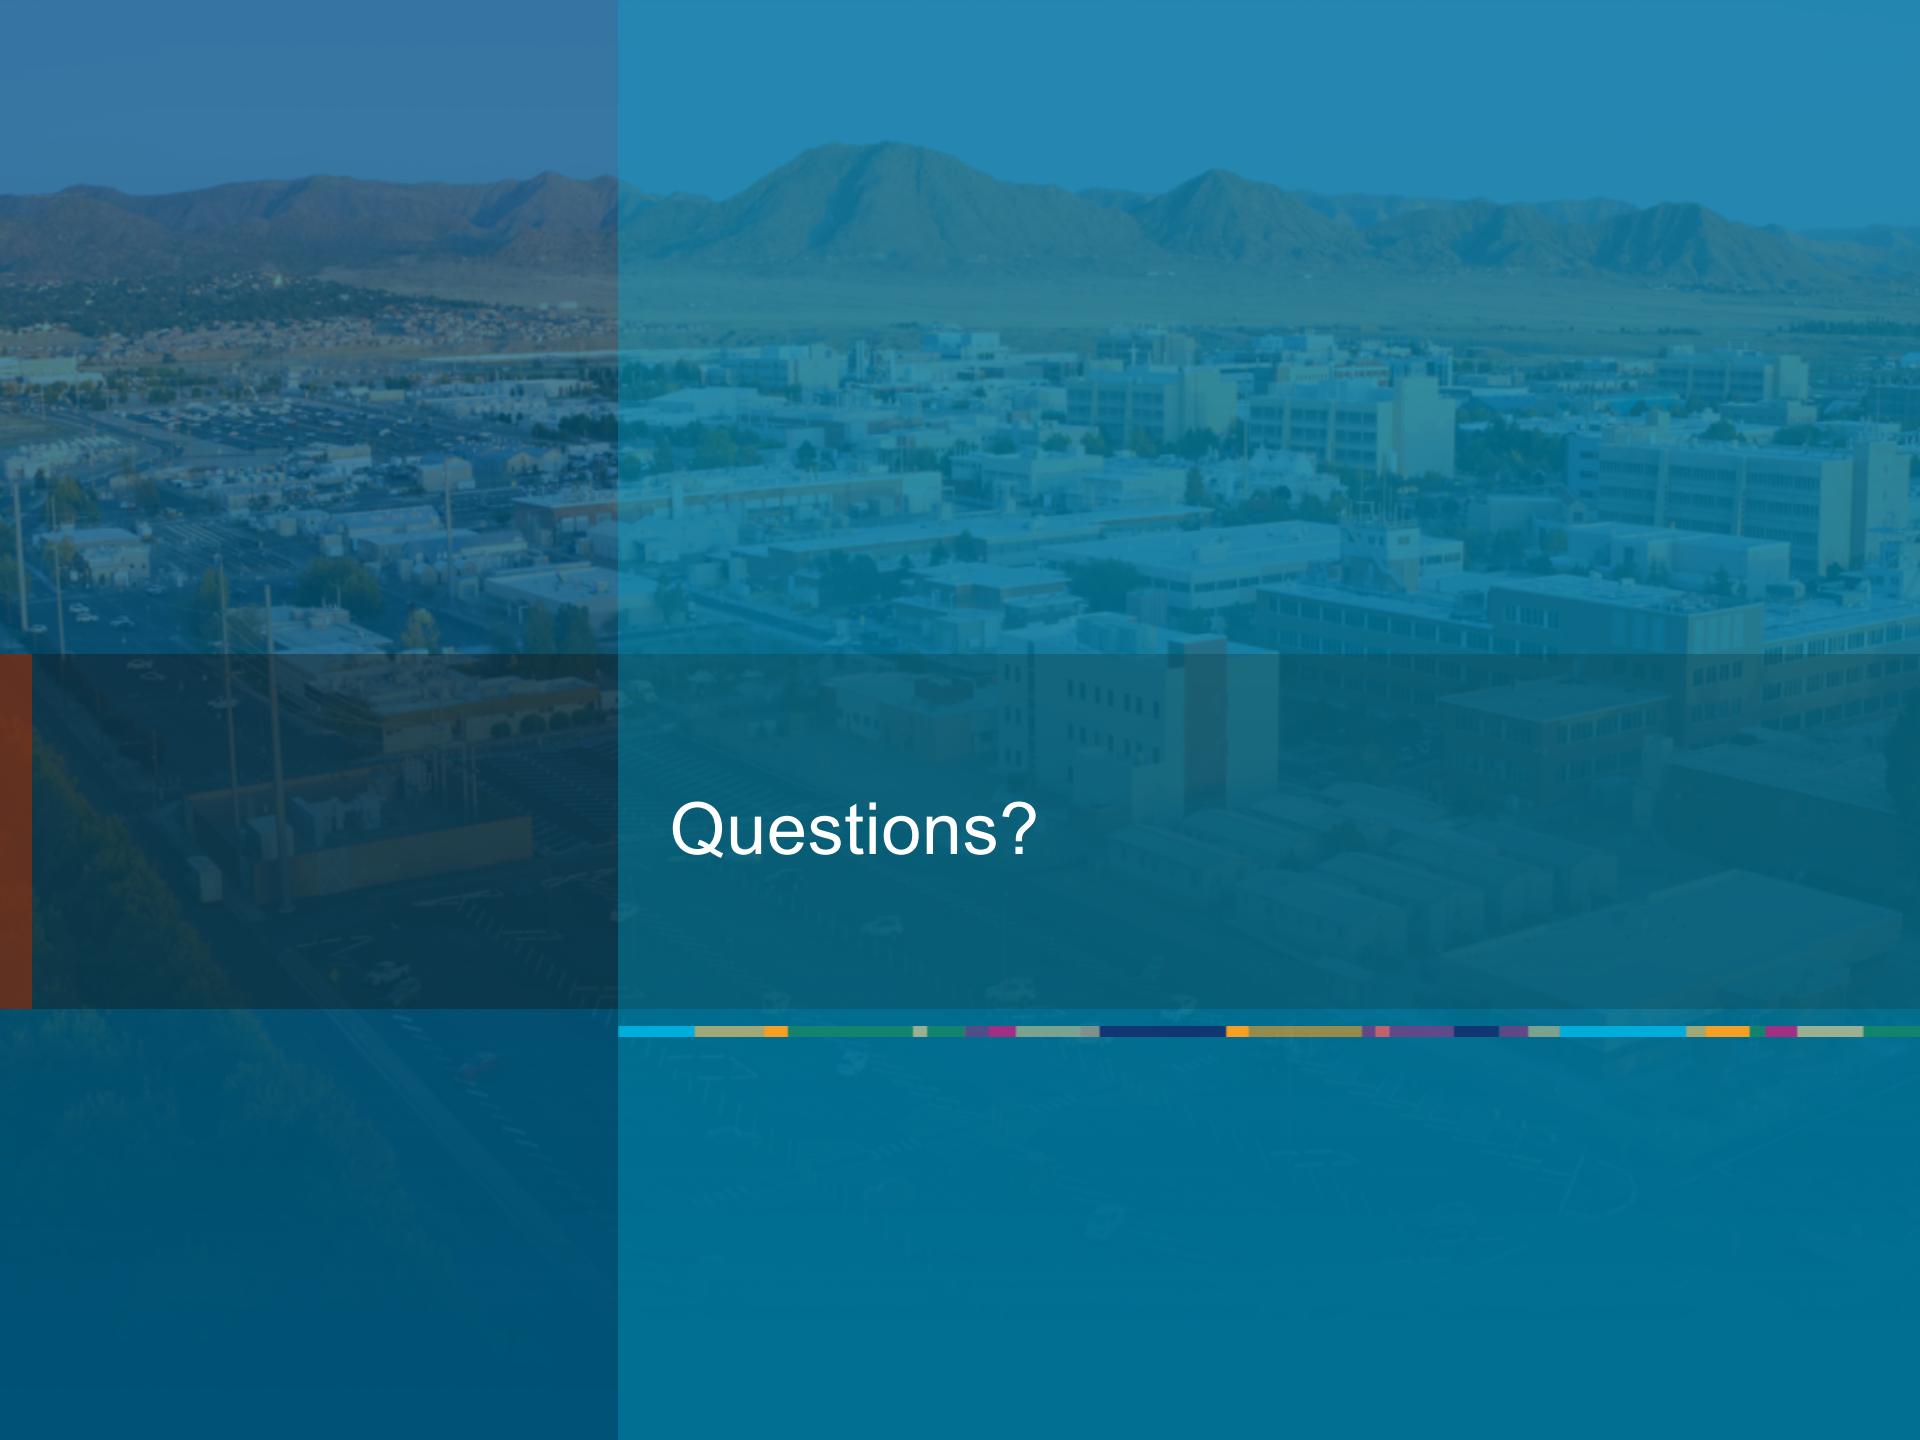
- Fully migrated 6 NFS filesystem so far (2 this week) – same filesystem, new hardware
- Total size for all filesystems ~115TB, largest being 45TB
- Primarily used sfsync-and-verify or scripted combination of rsync_wrapper and clean_target.py
- Received lots of help from Starfish support



- Use of a single script - `sfsync-and-verify`
 - The report from the script is very helpful – files that failed, time to complete each step
 - Ability/confidence to “set it and forget it”
- Watch current sync progress `sf job show xxxx`
- Historical information from `sf job list`
- Parallel rsync – several workers per agent node makes quick work of the migration
 - Currently using between 3 and 5 agent nodes, depending on network



- Currently working on migrating a 4-5PB gpfs filesystem, along with WOS object storage
 - Using copyit.sh script seems to work best. Thanks to Doug Hughes for the help.
- Exploring possibility of migrating data from lustre when we bring newer filesystems online
 - Historically we relied on users to move or delete their own data
- Would be nice to do clean up the data landfill prior to migration
 - Remove duplicates
 - Remove/archive older data



Questions?

