Central NAS Service

Status update

Fermilab

Hepix 2007

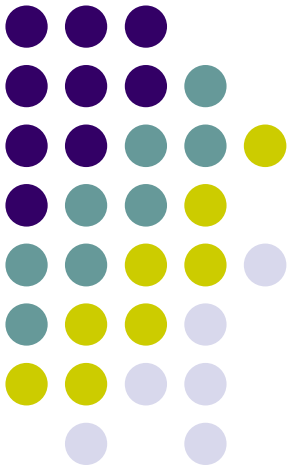

Status Summary

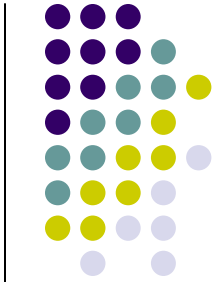

- Production 08/2006
  - 2x BlueArc Titan 21 heads (first generation)
  - 2 Hitachi AMS-500 storage arrays
  - 1x Nexsan SATAbeast (CMS)

Status Summary

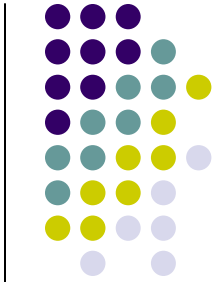

- Upgrade 11/2006
  - Bluearc upgraded to Titan 2200
  - Additional 64x 400GB HDS SATA disks
  - 3 additional Nexsan SATAbeasts (CMS)

Status Summary

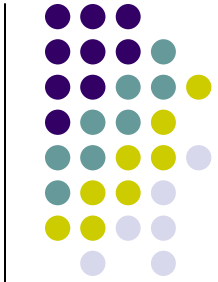

- Today
  - 2x Titan 2200 heads
    - 117+TB of storage in cluster (FC/SATA mix)
    - ~1200+ NFS/CIFS/FTP clients

Cluster Performance

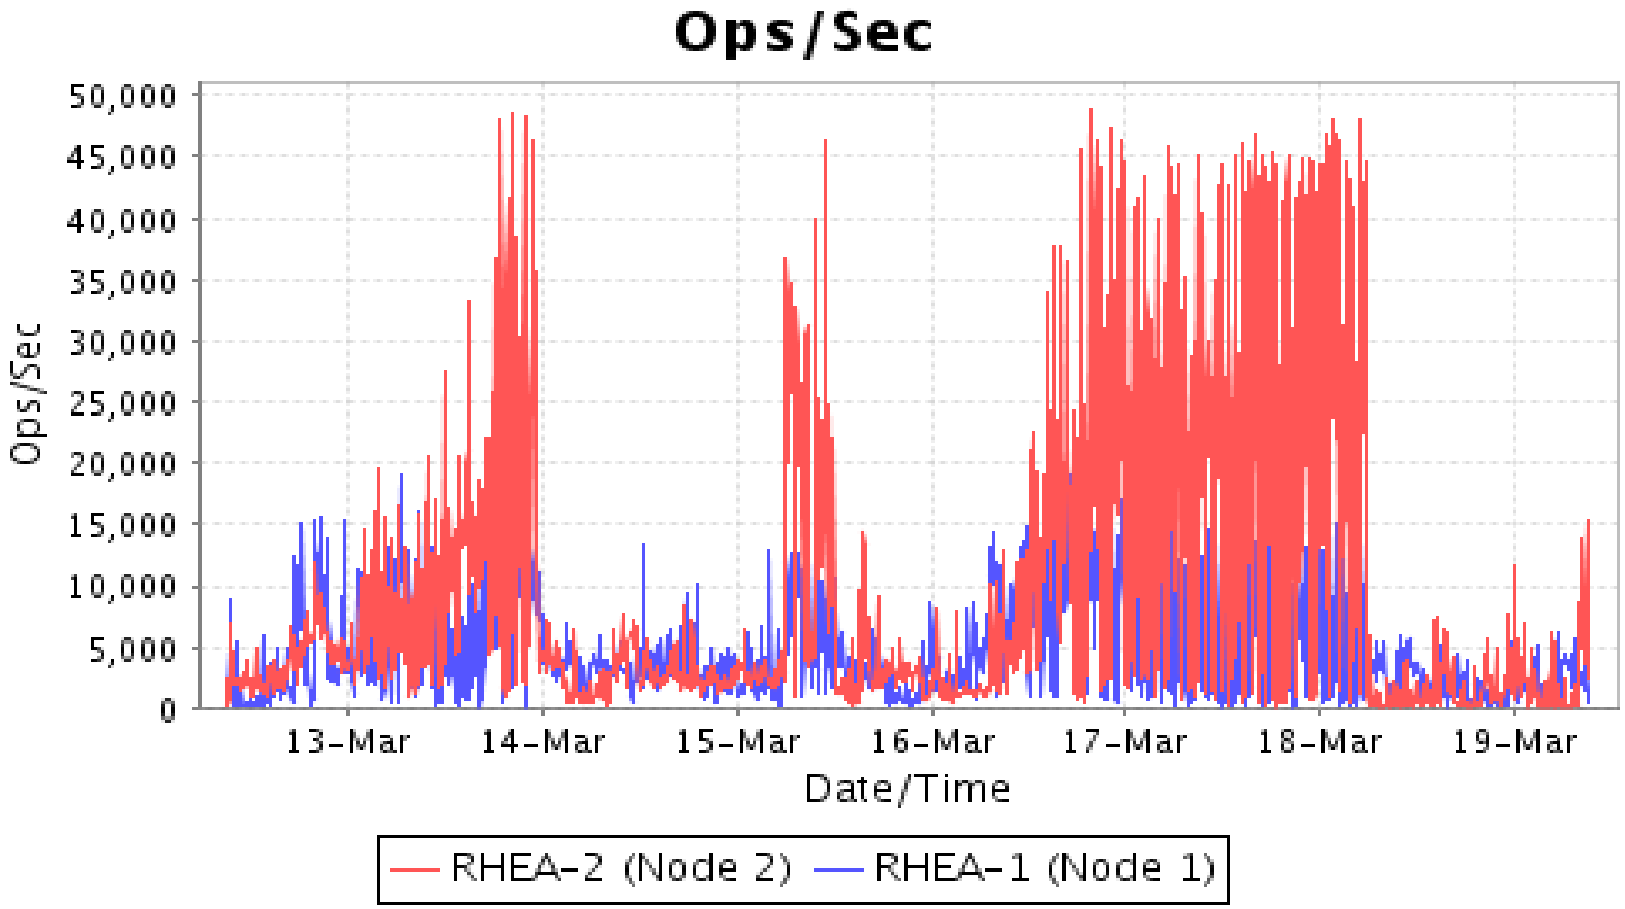

FermiGrid

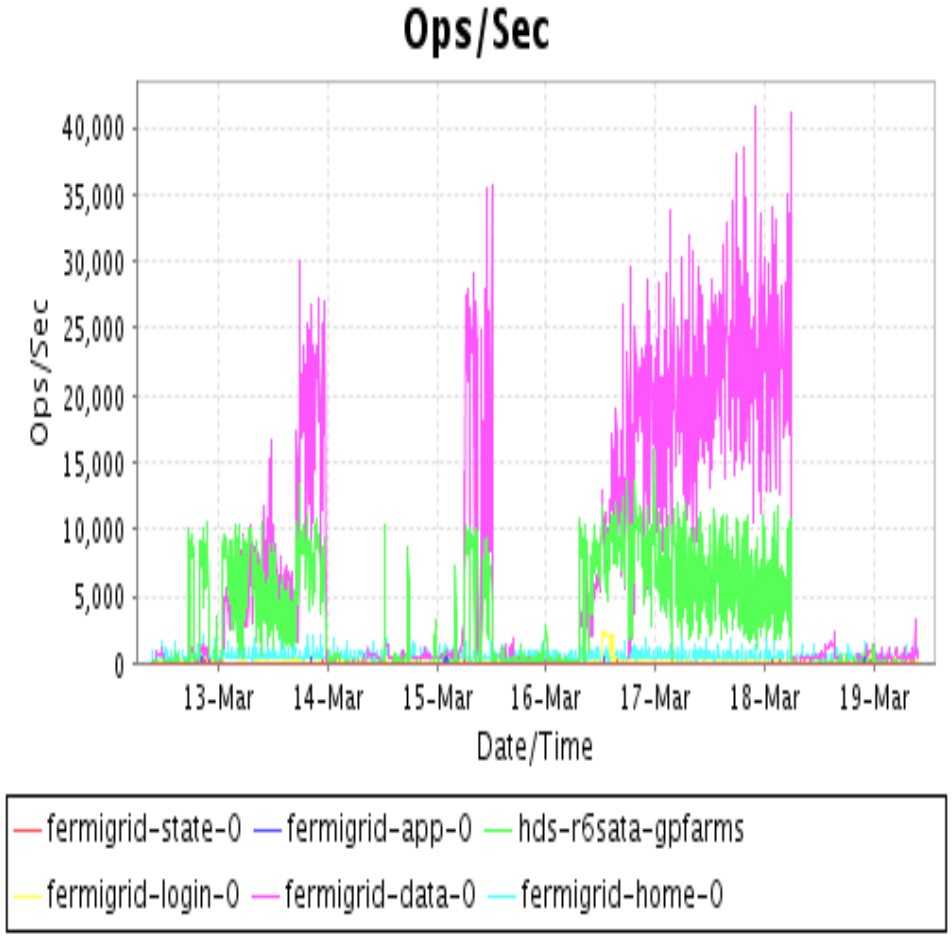

- ~520 Linux nodes
- Hitachi RAID6 SATA disks

CMS

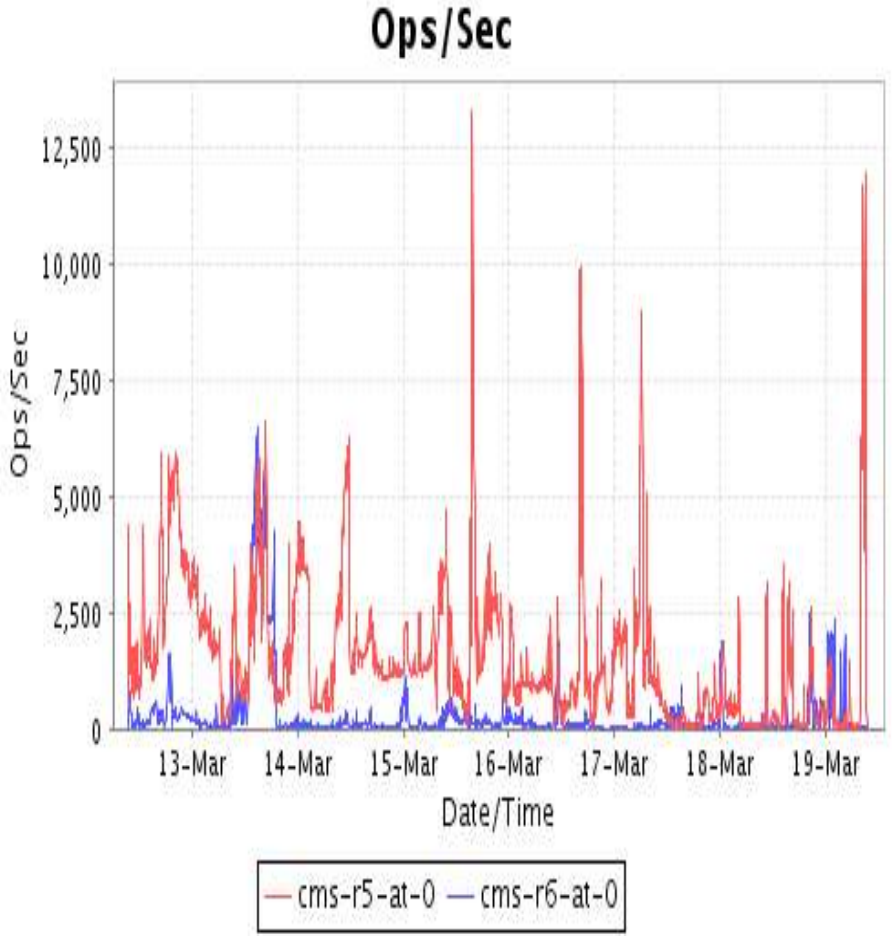

- ~600 Linux nodes
- 4 Nexsan Satabeasts
- Data/scratch areas- RAID5 across 3 storage arrays
- Home areas- RAID6 on single storage array

General Public

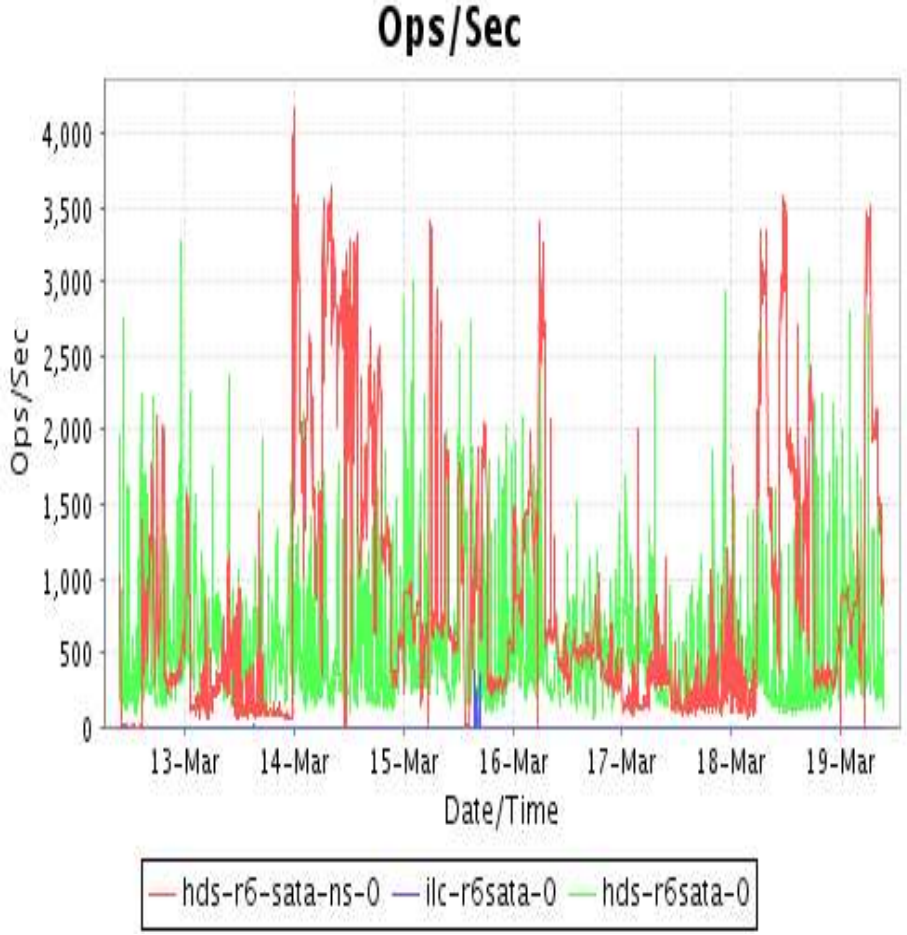

- SLF NFS/FTP install server
- Yum server
- Central CVS storage
- Oracle DB nearline backups
- SDSS
- LHC@fnal ROC
- Departmental fileserver
- Software distribution storage

Observations

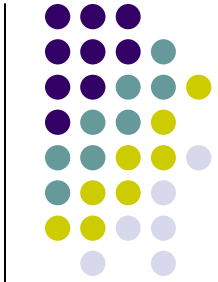

- Central Service provided by Computing Division
  - To share resources across experiments
    - CMS, Fermigrid, etc are users of service and are happy
  - Reduce overall costs and effort
- Support from BlueArc has been outstanding.
  - Quick to identify and resolve issues

Observations

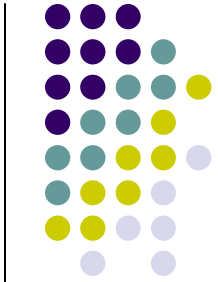

- In general, performance and reliability have been very good.
- Performance issues that have come up can be attributed to back-end storage or network bandwidth issues.
- Even with current load, we are only at ~35% NAS head capacity.

Observations

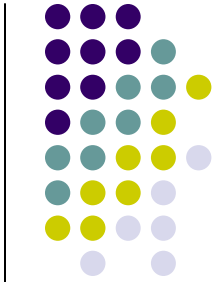

- New code to allow for > 2TB/lun support should allow us to have more than 512TB of storage behind NAS
- Reboots from upgrades are tolerated by NFS clients (~5minutes 10 seconds) with minimal disruption.
- A filesystem problem could take down entire cluster. New code (4.2.944b), currently in place, will now just place filesystem offline and not impact entire cluster.