NFS servers

Alexandr Mikula

FZU Fyzikální ústav Akademie věd České republiky

NFS servers overview

- Ten servers
  - One hosts the „home“ storage
    - Daily backups (except for the second day of given month)
    - VM based hosted on the enterprise storage array
    - Active Quotas 3GiB soft, 4GiB hard, one wee grace
  - Two for all users*
    - 28TiB and 91TiB
    - Planned upgrade of capacity to ~220%
  - Rest is astroparticle reserved*

* These are **NOT** backed

NFS service problems

- Servers are optimized for bulk transfers of big files
  - magnitude of gigabytes and larger
  - Small IO operations have huge overhead
  - Please try to avoid:
    - Direct IO from nfs in jobs
    - Cloning of whole git repositories (shallow clone should hel
    - Any other unnecessary IO
- Avoid running of 100+ jobs requiring data from one server
  - Limiting of running jobs on next slide

NFS concurrency limits

You can limit number of concurrently running jobs for single nfs
`concurrency_limits` condor expression:
Format is following:

`concurrency_limits = <nfsID>:<divisor>`

Example:

`concurrency_limits = nfs19:4`

Explanation:

Limit these jobs to ¼ of my running jobs limit for server nfs19
Limit is <UserJobLimit>/<divisor>

See the FAQ for more details

NFS discussion and future

- Moving to the transfer input, output files mechanics form cond
  - Unmounting of the nfs servers from worker nodes
  - More control and abuse prevention
- Distributed storage (possibly CEPH)
- Quotas?
- Upgrades and financing
- Hosting and purchases of new hardware

Thank you for your attention

FZU
Fyzikální ústav
Akademie věd
České republiky