

# Chirp for User Output

- Rod Walker, LMU

- Current model for user output
- Chirp
- Status & experience
- Future Plans

# Currently in ATLAS

- Output(s) and logs written to local SRM
  - in SCRATCHDISK space token
    - 30 day cache turnover
  - registered in DDM dataset
- User can either
  - request subscription to permanent storage
  - more usually, dq2-get to desktop
- Most user outputs are tiny

# User Output File size




Fig 1: Snapshot of user output files at FZK from June 2010

# Problem with that

- Manual user step to get data
  - many jobs = many dq2-get and all SRMs need to work
- dq2-get is chaotic access to SRM
  - possible interference with controlled DDM transfers
- Moving files with DDM better for SRMs but ..
  - adding many small files to the system slows it down

# An Alternative

- Small files (<50MB) can be written directly from WN to place they will be analyzed
- Global File system is perfect for this
  - but what can we do today

# Requirements

- Destination should be 'file' access
  - users like posix(and so does Root)
- AAA via X509 proxy
  - destination need acs to share data with group
- Server must be easy to setup and maintain
  - user 2TB desktop, or HEP group service

# Technology choice

- Webdav, GridSite, NFS4 are candidates but I know Chirp has required features.
- Chirp spawned from Condor project
  - X509 and acls
  - FUSE module for 'file' access
  - easy to setup
  - good support (from Notre Dame now)

# Status

- Catch-all Chirp server at CERN
- Ganga(Panda backend) and Pathena have user switches to enable storing output to Chirp
- Wiki linked to FAQ (low-profile)
  - <https://twiki.cern.ch/twiki/bin/view/Atlas/ChirpForUserOutput>
- Few users tested, and 1 heavy user
  - 80k files on server
- Lxplus allows FUSE!
  - Chirp tools in ATLAS afs area
  - CVE reboot accidentally removed FUSE rights



# Experience

- Chirp server froze a few times under heavy load
  - possibly due authentication cpu, but did not investigate (did not reproduce)
  - holes in datasets recovered using dq2-get
    - i.e. must also store to local SRM, or have retries/failovers and require Chirp success for job success.
  - WN network store time is negligible
    - $1\text{MB/s} < 1\text{min}$ . SRM,LFC interaction can take longer on local store.

# Future Plans

- Encourage more test users, e.g. in tutorial
- See if they find it useful, or stop
- If useful then performance test and revisit technology choice