DDM 3: Summary and Actions

- DDM as a global POSIX namespace
  - No, mapping does not really work well, but limited use into local storage for T3s works well enough
- Metadata searches are useful should be developed (dq2-find)
  - Separate from internal optimisations in CC
- Web interface to data demonstrator proposed
  - For now will need to be a proxy service, but support from storage is strongly desired
- Do not ask storage systems for exotic non-standard features
  - Limits our flexibility to adapt to future technologies
- Continuing need for a file transfer service is foreseen
  - Requirements: http(s) interface support (S3); probably no channels, only endponts
- For LAN access file:// is very desirable
- Tier-2s continue to need archival storage
  - Robust storage for unique outputs
  - Primary copies as T1s grow short on space
  - Cache can be catalogued (like now), but could use panda affinity or publication model
    - T3 'no registration' transfers to T2 cache
  - Push xrootd demonstrator to some T2s to improve data access robustness
- NoSQL databases under study to address scaling/performance issues and this is welcomed
  - Continue to inform rest of community, work together where possible
- Rule based engines for policy
  - Demonstrator for group data replication policy
- Merge of LFC functionality into DDM is agreed in principle
  - Needs some thought starting this year

ATLAS and WLCG Discussion