

NDGF Update

WLCG GDB, CERN, Geneva



norden

NordForsk



Nordic e-Infrastructure
Collaboration

Overview

- News
- Storage protocols
- New tape integration
 - teaser for next week
- IPv6



News

- External review
 - Efficiency, distributed vs centralized, etc
 - Interviewing candidates this month
- New resources, as usual
- New versions of dCache, as usual
 - But still on 2.12.9, not on 2.13 yet
- Finally migrated most services off the old blade center
 - Ubuntu 10.04 EOS finally pushing it hard
 - Production on the redundant Ganeti cluster

Storage protocols

- ATLAS ND production has experimented with SRM+http[s] in addition to SRM+gridftp
- Only for downloading
 - See next slide
- Side-effect of using IPv6 for much of the traffic since a decent share of ND ARC-CEs are already dualstacked
- Found some suboptimalities in ARC's http downloader, after upgrading to ARC 5.0 it might be worth to re-run
- SRM still protocol of choice, for bulk speed as well as negotiation and redirection

Difference between https and gridftp

- For downloading data from an SE https + http redirect has slightly lower overhead than gridftp + data connection
 - More round-trips for multiple gridftp commands in a normal download transaction
 - And http is just request-response (twice)
- For uploading http [and xrootd] is/was more dangerous than gridftp to dCache endpoints (and others?)
 - Late errors could be propagated through control channel
 - But the http/xrootd mover that got the redirect doesn't redirect back to the central door after finished upload

ENDIT - Efficient Nordic Dcache Interface to TSM

- Reimagining of ENDIT with native dCache HSM provider
- Much more scalable than a perl script per outstanding read request
 - 100k concurrent read requests per pool doable
 - Discovered
 - Great for tape drive efficiency, perhaps less so if some of the files requested from tape have priority
- Easier configuration with no “hidden” buffers when dCache manages all the space within the tape pools
- Production version some weeks away
- Talk at the dCache workshop in Amsterdam next week

IPv6

- In our long struggle to get ready for dualstacking our main production dCache `srm.ndgf.org`, progress!
- Last site of pools are getting them dualstacked ~now
 - Currently also means most http transfers from dualstacked clients will use IPv6
 - But not gridftp, since it uses IP+port for datachannel vs hostname+port, where hostname can resolve A & AAAA
- This means we can start looking at dualstacking the doors and headnodes

Changes this late in the game?

- Enabling IPv6 is probably doable, especially if we can back it out quickly
 - And is an area where there is plenty of testing work
 - Also helps that most others won't try to use IPv6 :)
- For interface/protocols between sites we're pretty much stuck with the current setup until LS2
 - NDGF: SRM+gridftp for ATLAS, xrootd for ALICE
- CEs can be a bit more flexible
 - Not affecting data taking
 - Independent from each other

Questions?

