

XRootD cache sites

Alastair Dewhurst

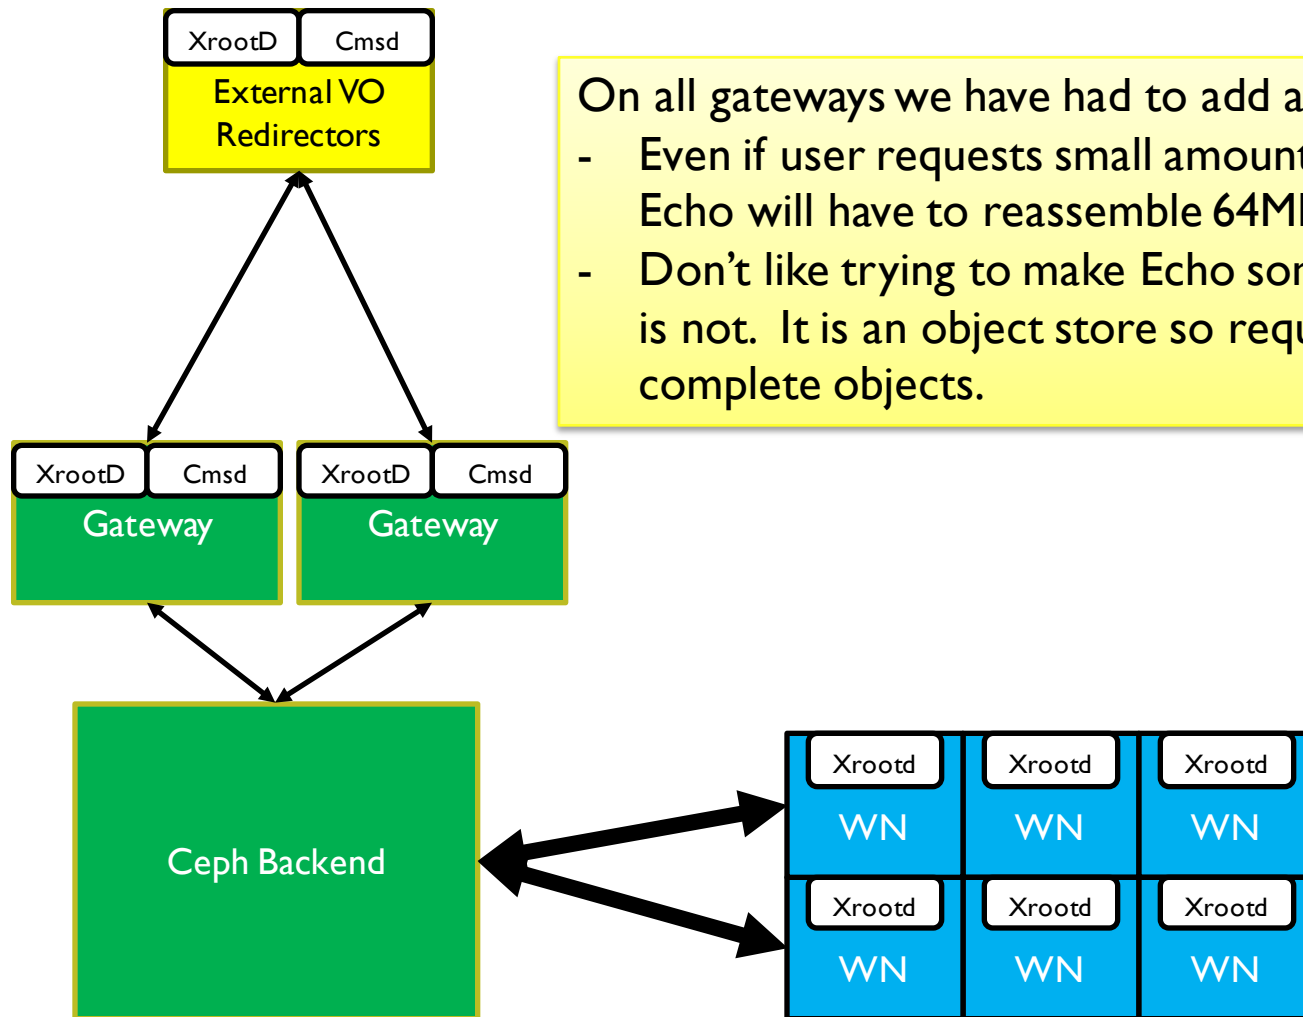


Motivation

- ATLAS are pushing forward with diskless sites:
 - https://indico.cern.ch/event/642836/contributions/2608398/attachments/1467335/2268911/Diskless_28May.pdf
 - ATLAS consider Birmingham, Cambridge, Sheffield to have small storage but don't know have a clear plan.
 - Leaving it to us to decide at the moment.
- Requirements:
 - No new hardware and can work with increasingly unreliable hardware.
 - Extremely easy to maintain for the site.
 - Personal opinion – If it is useful for local users, more likely to be maintained.
 - Transparent to ATLAS.



RAL XRootD Setup



‘ECDF’ XRootD Setup

- Setup:
 - 2 redirectors working in load balancing mode, one at ECDF and one in Glasgow.
 - 2 storage servers behind it right now, one in ECDF and one in Durham.
- More sites can easily connect a server to the redirectors.
- Instructions: <https://www.gridpp.ac.uk/wiki/XRootD>
- Main problem for this site was how to get data there (No GridFTP for FTS).
 - I have been informed it is ‘trivial’ to configure this to stage in data from an external source (i.e. act as a cache).



Start of a plan

- Add configuration to XRootD storage to allow it to function as a cache.
 - Ideally to pre-fetch complete files.
 - Point it at local storage first and test if it works.
 - Test it against Echo.
 - Setup 'UCL style' Panda queue.
- Wei Yang has written a N2N plugin that caches identical Rucio files (that have been given different scopes)
 - <https://github.com/wyang007/xrootd-N2N-rucio-pfn2glfn>
 - Could increase efficiency of cache (but not necessary for functionality).

