

Xcache in ATLAS Distributed Computing Environment

W. Yang^{1*}, A. Hanushevsky¹, H. Ito², M. Lassnig³, T. Li⁴, R. Popescu³, A. De Silva⁵, M. Simon³, R. Gardner⁶, V. Garonne⁷, J. De Stefano², I. Vukotic⁶, A. Washbrook⁴ on behalf of the ATLAS Experiment


¹ SLAC National Accelerator Laboratory, ² Brookhaven National Laboratory, ³ CERN, ⁴ University of Edinburgh, ⁵ TRIUMF, ⁶ University of Chicago, ⁷ University of Oslo, * Corresponding Author



What is Xcache ? - Xrootd Disk Proxy Cache


Squid-like Cache but speak the “root” protocol:

- Unix environment variable XROOT_PROXY - similar to “http_proxy”
- Caching at either sub-file level, or whole file
- Multi-thread, async data fetching
- Designed for both **large** and **small static** data files
- Clusterable for scaling up
- **ROOT and HTTP protocol between client and Xcache**
 - ROOT protocol between Xcache and data sources, will support HTTP

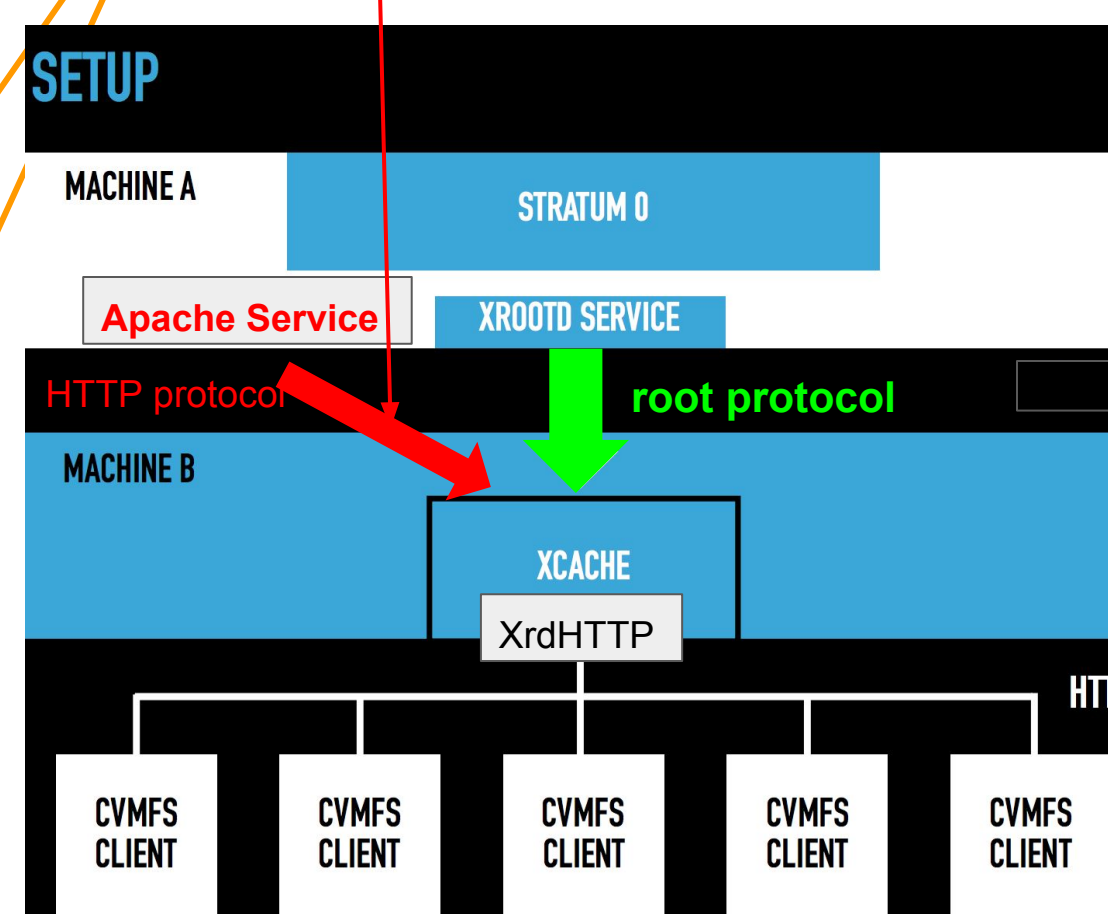
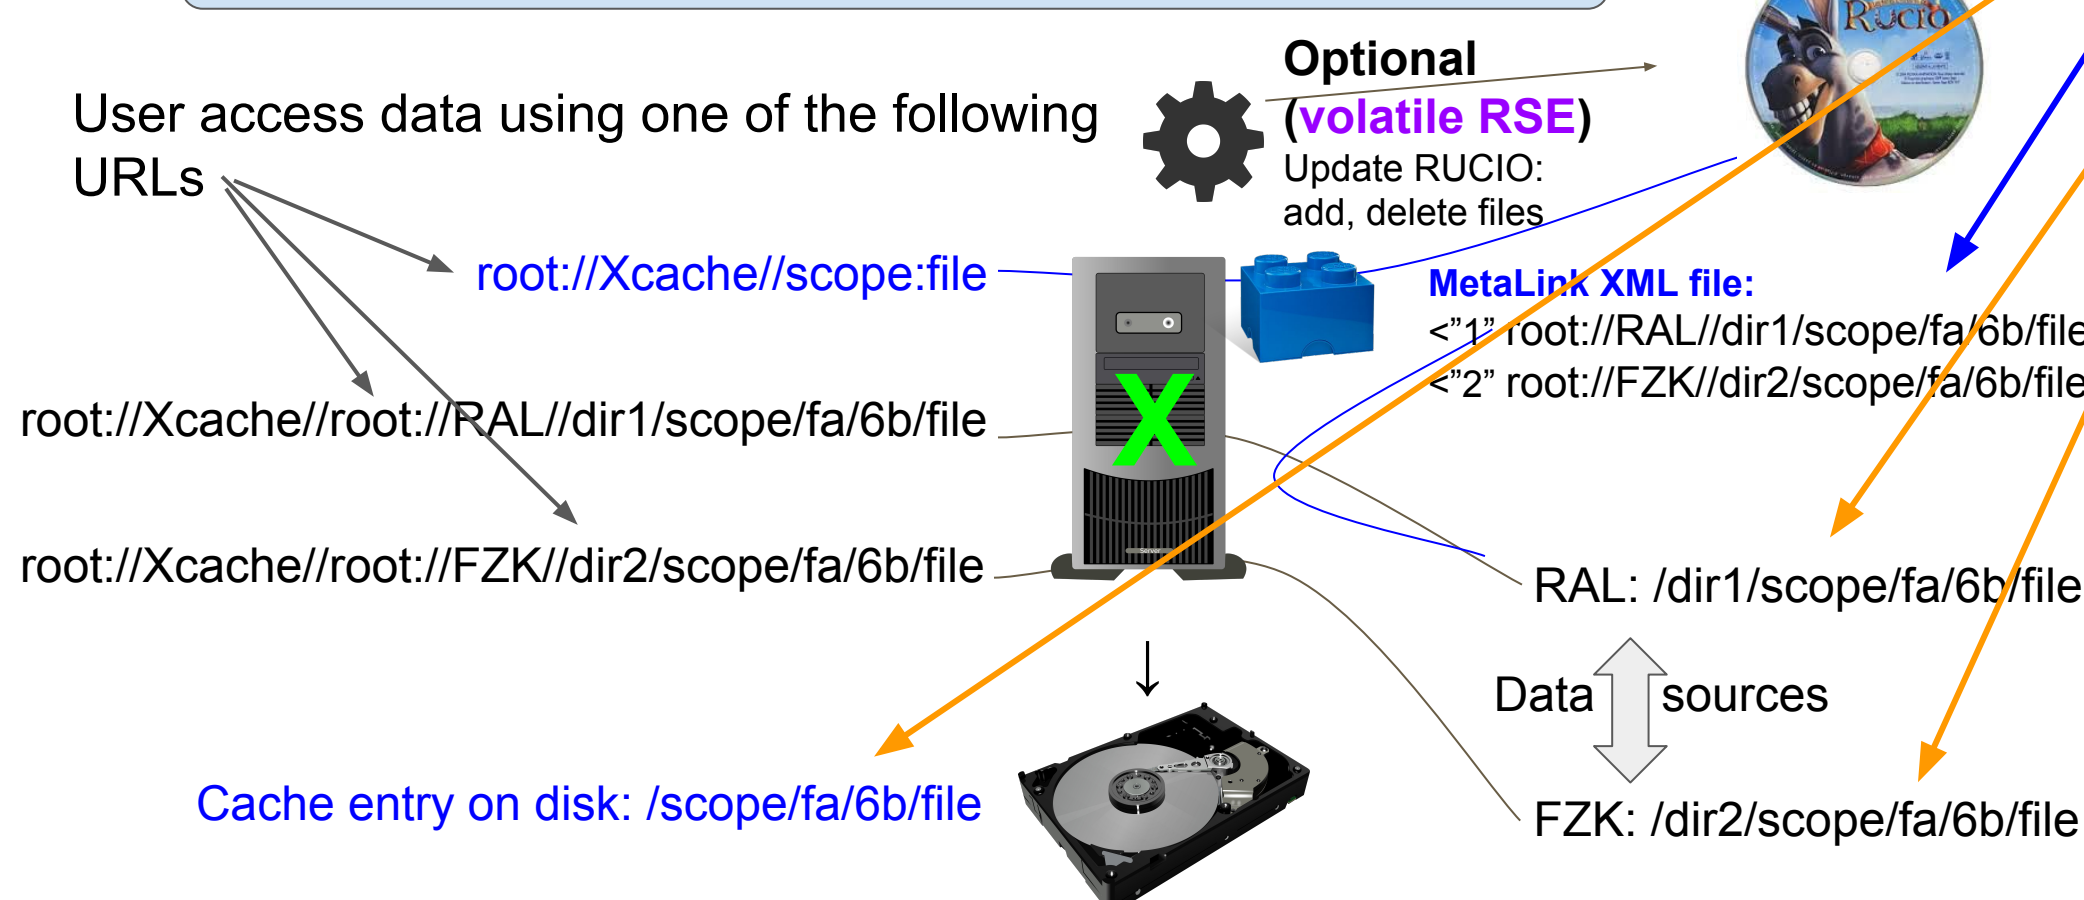


Plugin Architecture is Powerful

- **Talking to HTTP data source** (XrdCIHttp package - work in progress)
- Make decision on what to cache, and what to NOT cache
- Map a logical file name to the best data source
 - For example, querying RUCIO for a list of data sources (in a metalink)
- Where to place the cached data file
 - Identify the same files at different data sources and share a cache entry



Integration with RUCIO



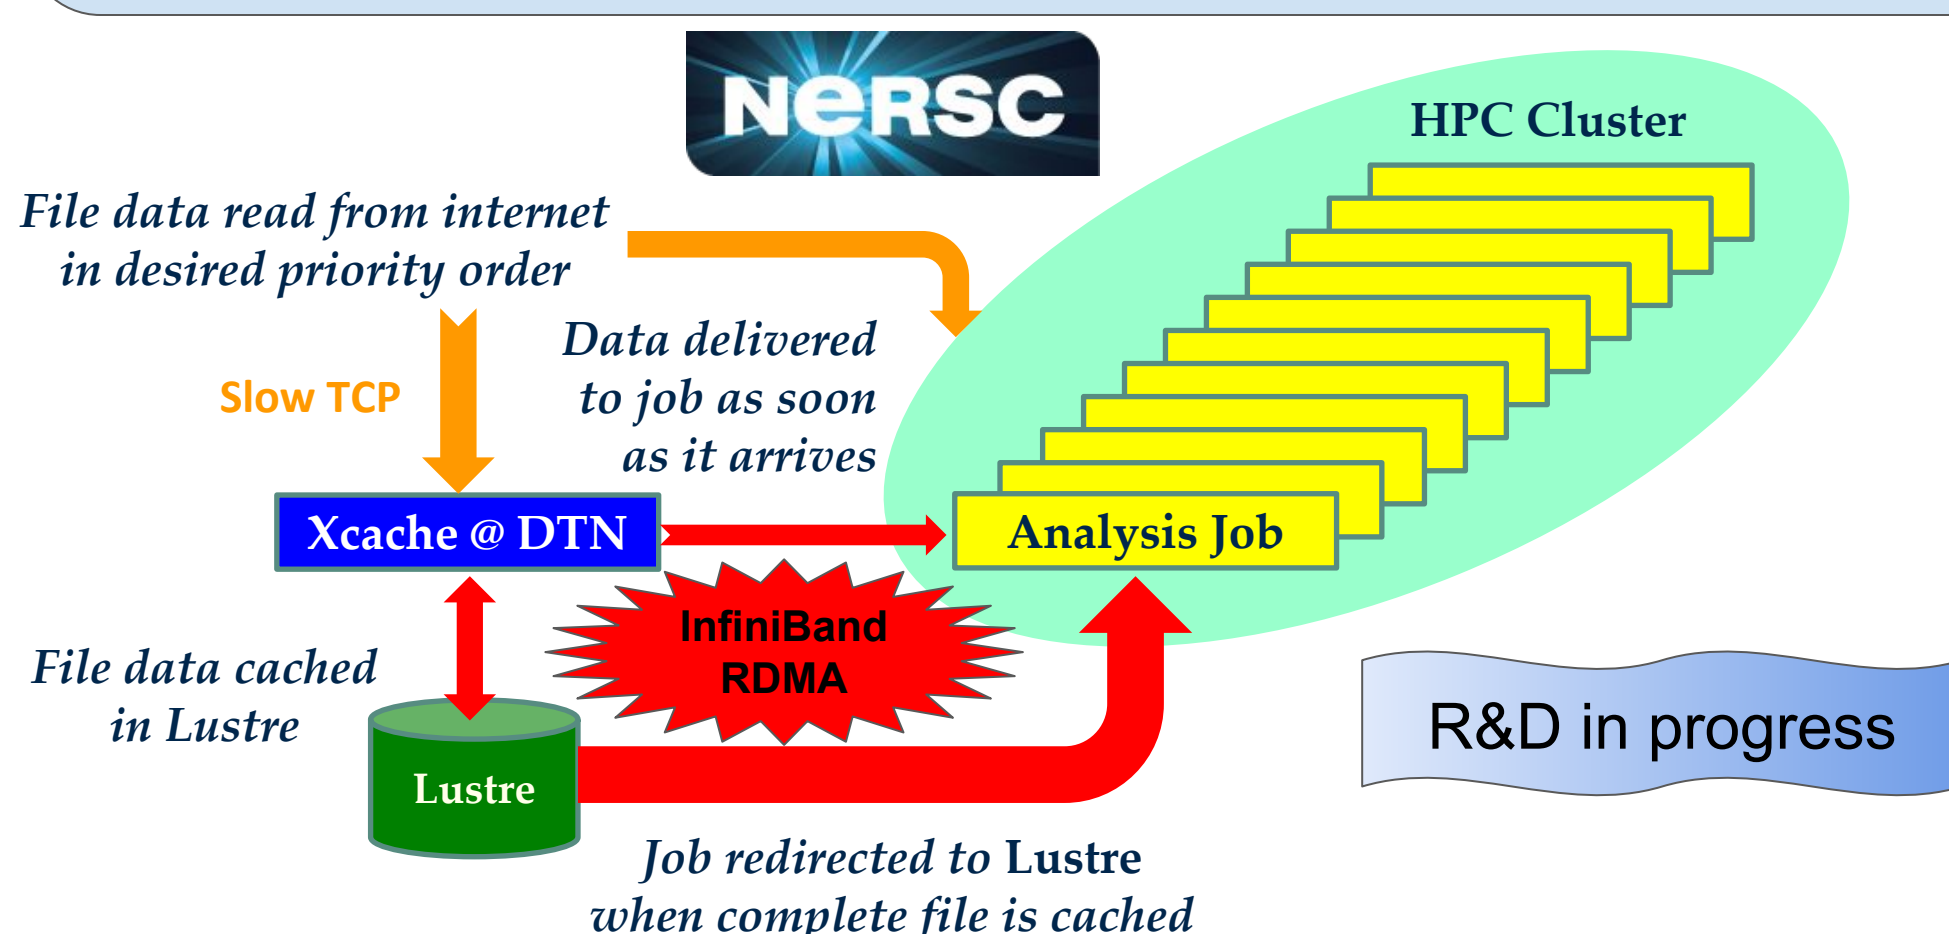
Xcache for CVMFS

- Export BNL Stratum-1 via Xrootd
 - Xcache Instance at SLAC
 - CVMFS client on SLAC batch nodes use Xcache instead of Squid
- OR
- (with XrdCIHttp)
 - Export Stratum-1 via Apache/HTTP
 - Xcache Instance at SLAC
 - ...

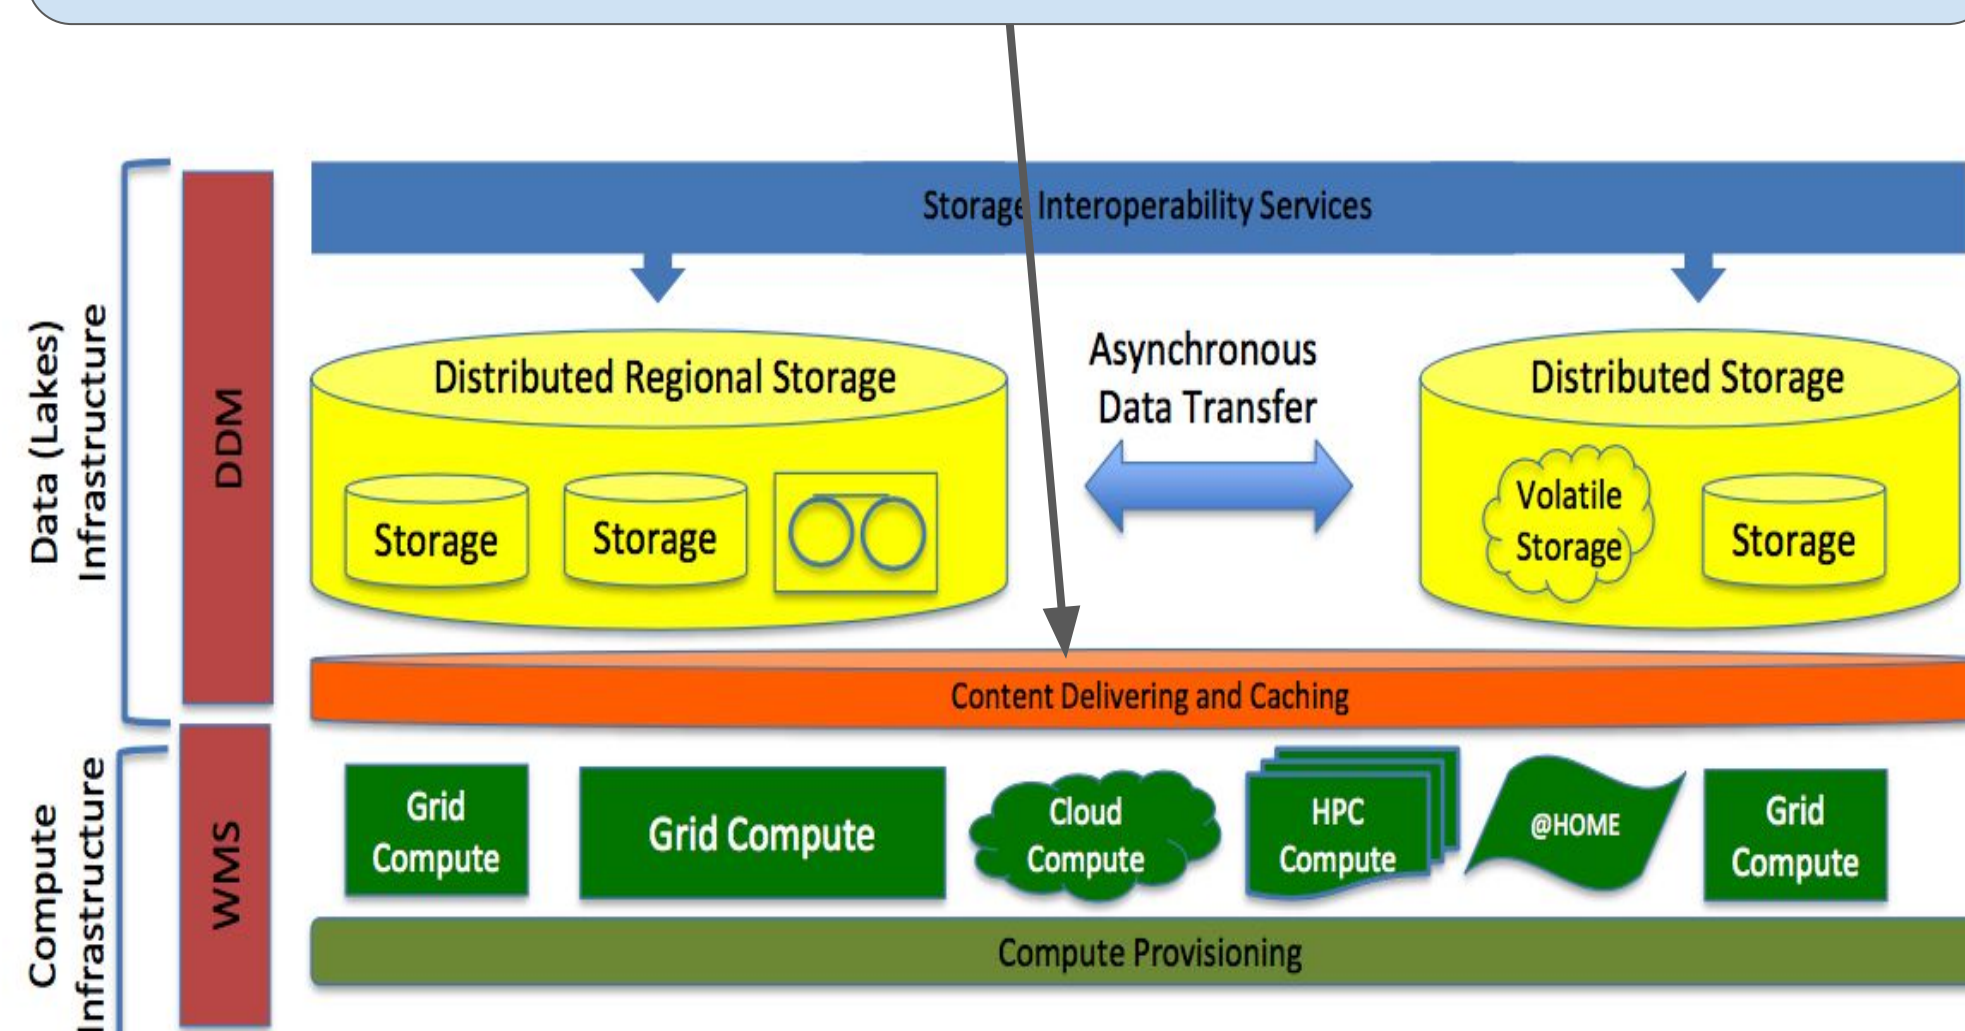
Xcache on HPCs - How to be Optimal

HPCs have fast interconnect and shared POSIX file systems

- Likely will put cache space on the shared file system
- Most efficient: user jobs open() completely cached file directly via shared file system
- Xroot over RDMA from Xcache to batch nodes - good for partially cached files.



Xcache Facilitates Data Delivery in the HL-LHC Data Lake Design



Xcache Packaging and Deployment

- Xcache is available in ATLAS CVMFS repository
 - nothing to install / deploy!
 - Run “lsetup xcache”, command “Xcache” will start a simple, preconfigured Xcache instance
- Xcache in containers
 - Singularity container and Docker container
 - Can run a single instance or a cluster of Xcache
 - Simple to use: run the container, supply a cache directory via bind mount - that is it!
- Orchestrated Deployment in US ATLAS
 - NSF R&D project SLATE will use Kubernetes to deploy Xcache to US ATLAS Tier 1/Tier 2 sites.

Other Activities in ATLAS Using Xcache

- Xcache with object store to improve POSIX IO performance ✓
 - Xcache in front of CEPH at RAL - **A Great Success!**
- Using Xcache as a **volatile** RUCIO storage element (RSE) for smaller sites
 - Since there is a RSE, these sites can run ATLAS Grid production
 - More tolerant to storage issues
 - Focus on providing CPU resources - minimum management for sites and for ATLAS
- Integrate Xcache with ATLAS production workflow
 - Lots of tidies hard works
 - Sorting out setting in ATLAS Grid Info System
 - Understand how ATLAS production uses Xcache and info in AGIS
- Exempt Tier 3 site from running a RSE to access data locally
 - Alternatively, they can just access data through Xcache.

ATLAS Sites Infrastructure Support Xcache

- Most storage systems provide Xrootd / HTTP doors
 - Need no special configuration to work with Xcache
- Availability and reliability of Xrootd and HTTP data source is important to the success of Xcache
- Behind firewall? No problem!
 - Xrootd Forwarding Proxy can help tunnelling through firewall
- Monitoring
 - Previous monitoring infrastructure for FAX can be reused
 - Collecting more info to the ATLAS analytics systems

Summary

- Many activities in ATLAS around Xcache
- The Xcache itself will evolve around a few areas (near term)
 - Optimal uses of Xcache on HPCs
 - Efficiently and reliably fetching data from HTTP data sources
 - Adapte the HL-LHC Data Lake model and facilitate the data delivery
- Modern containerized Xcache deployment is powerful and critical.