

# Xcache in ATLAS Distributed Computing Environment

A. Hanushevsky<sup>1</sup>, H. Ito<sup>2</sup>, M. Lassnig<sup>3</sup>, T. Li<sup>4</sup>, R. Popescu<sup>3</sup>, A. De Silva<sup>5</sup>, M. Simon<sup>3</sup>, R. Gardner<sup>6</sup>, V. Garonne<sup>7</sup>, J. De Stefano<sup>2</sup>, I.Vukotic<sup>6</sup>, A. Washbrook<sup>4</sup>, W. Yang<sup>1\*</sup> on behalf of the ATLAS Experiment


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



## What is Xcache ? - Xrootd Disk Proxy Cache

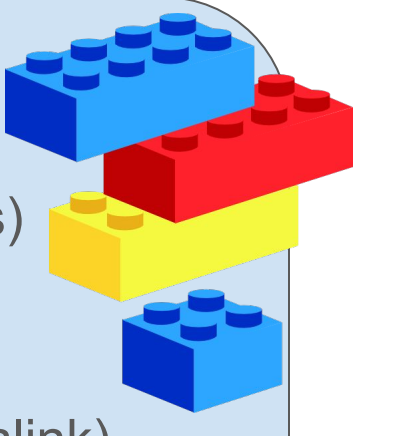
Squid-like Cache but speak the "root" protocol:

- 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

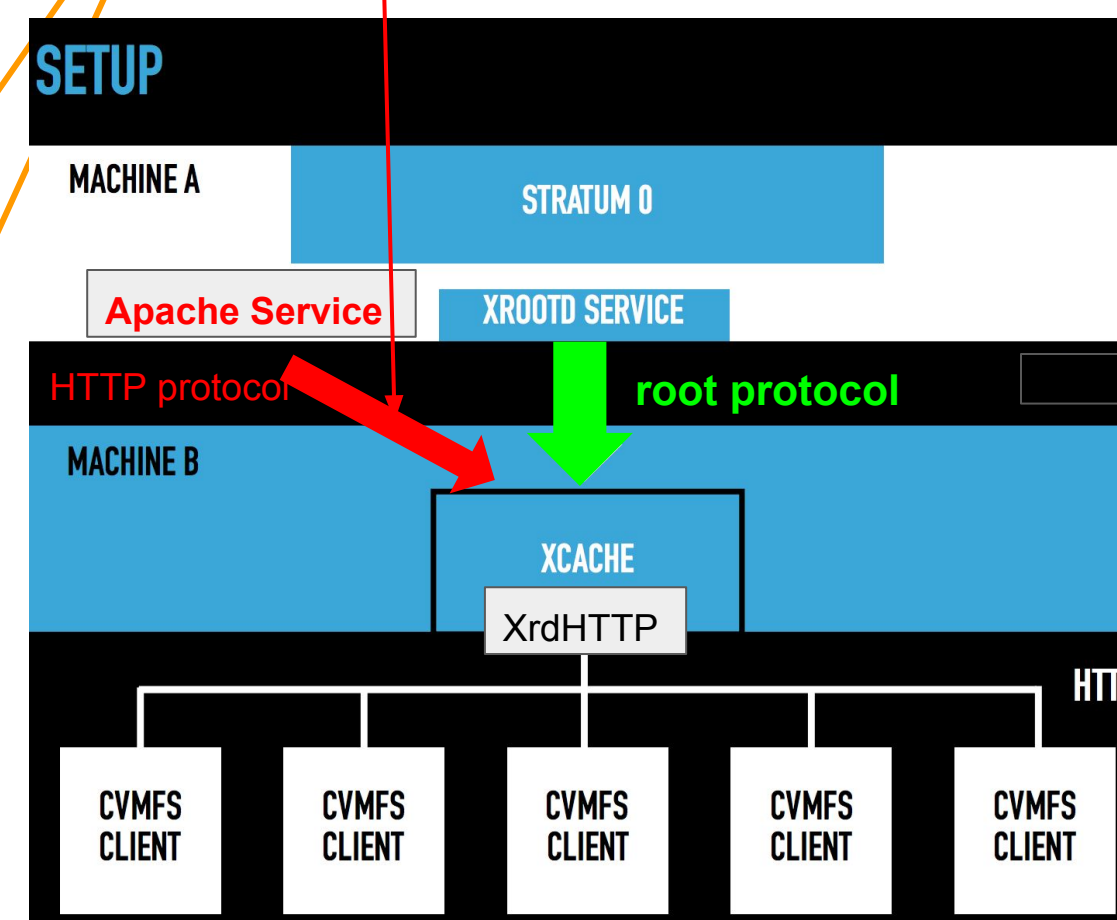
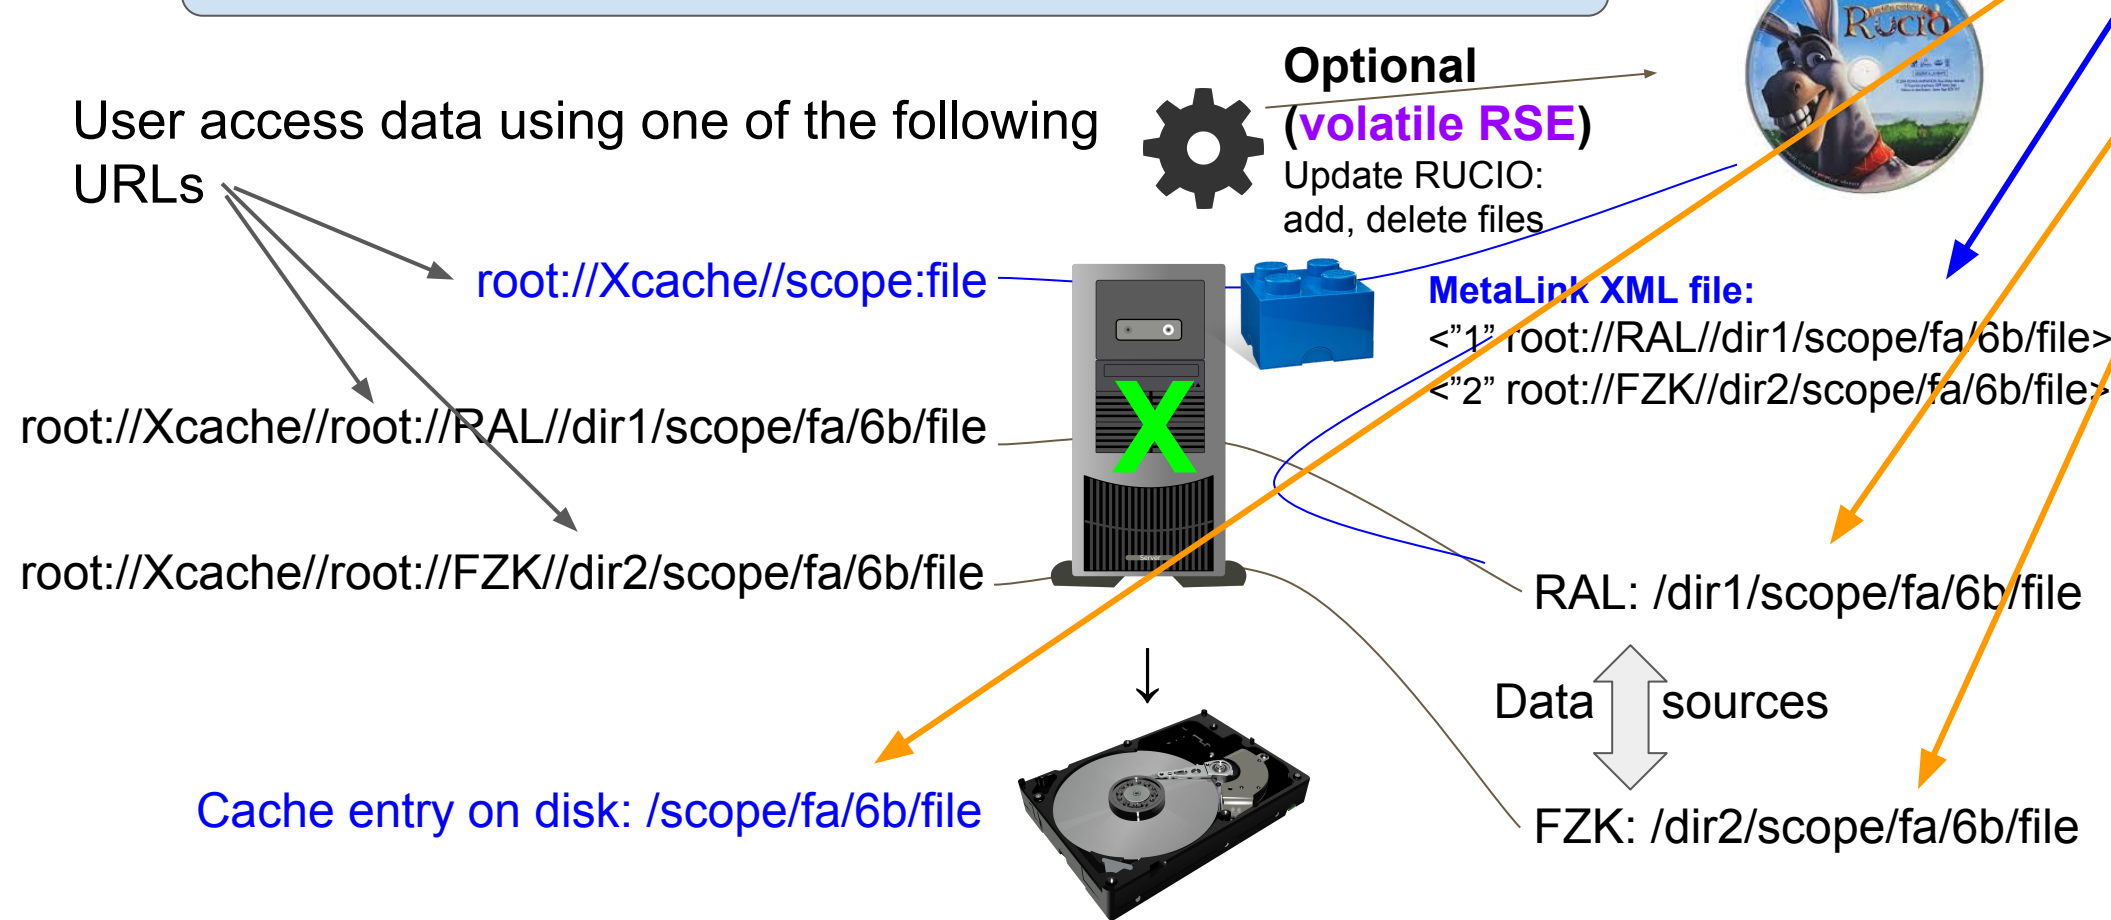


## Plugin Architecture is Powerful

- **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



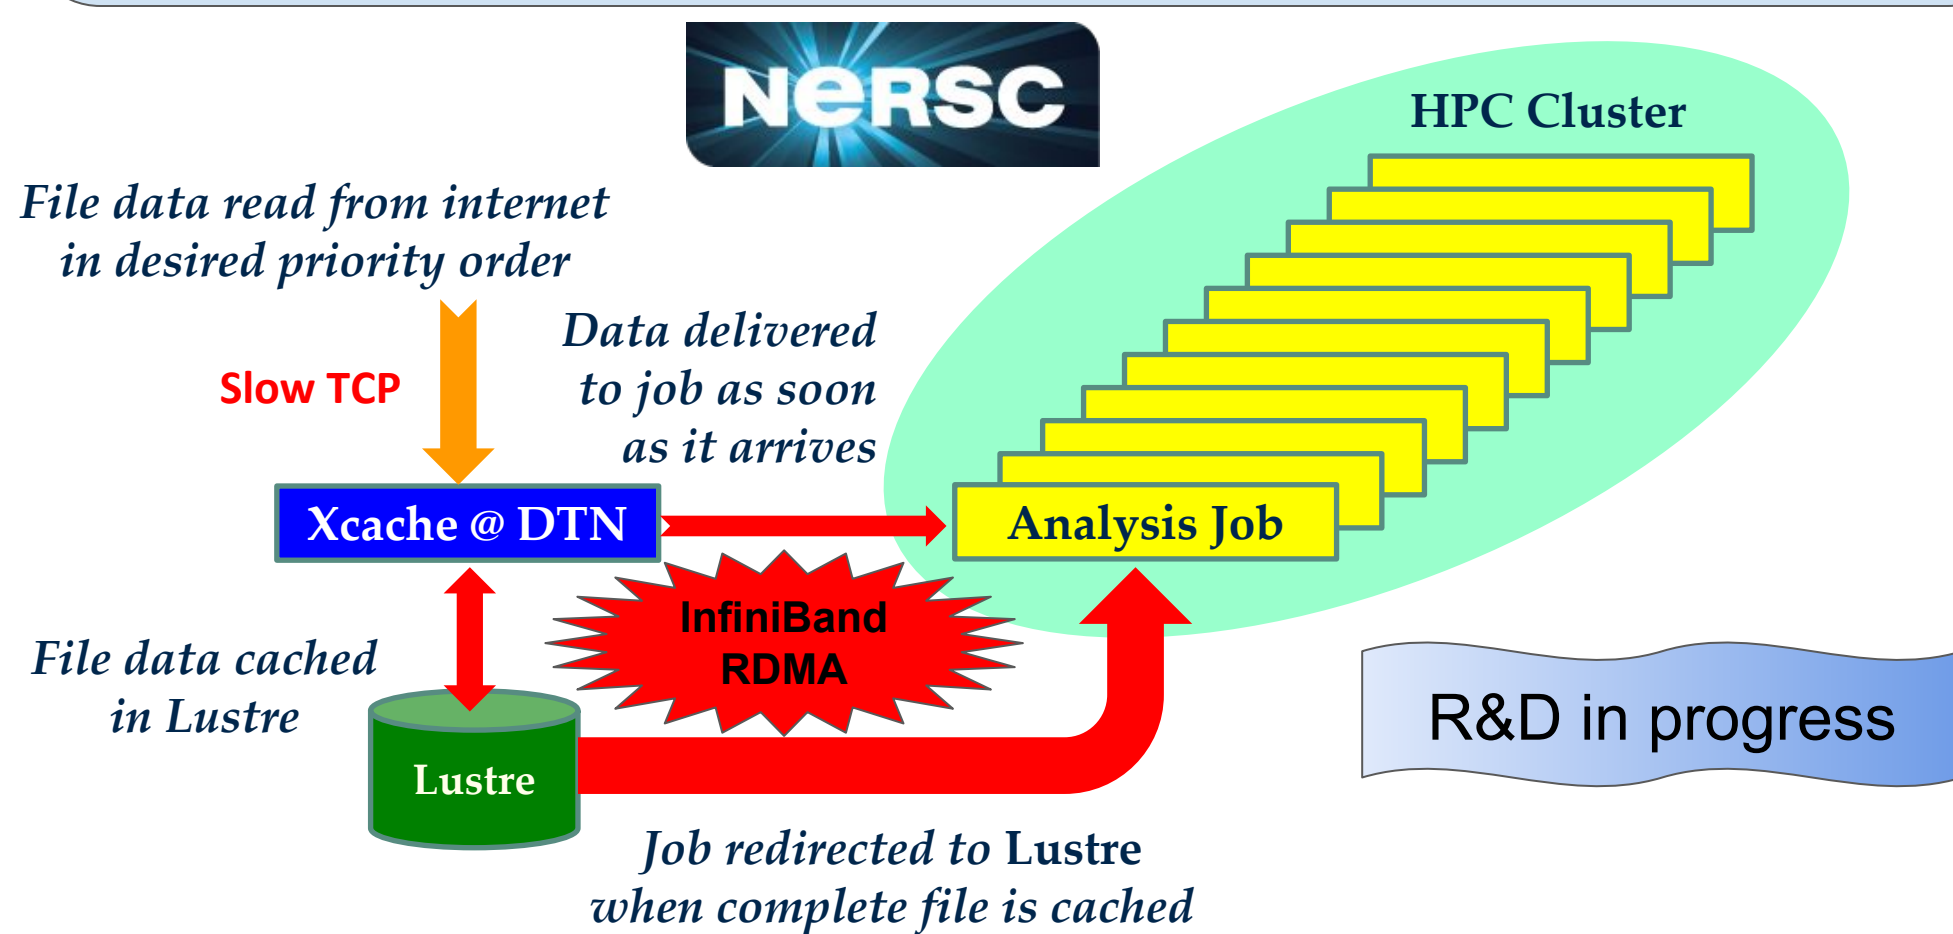
## Xcache for CVMFS

- 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
  - ...

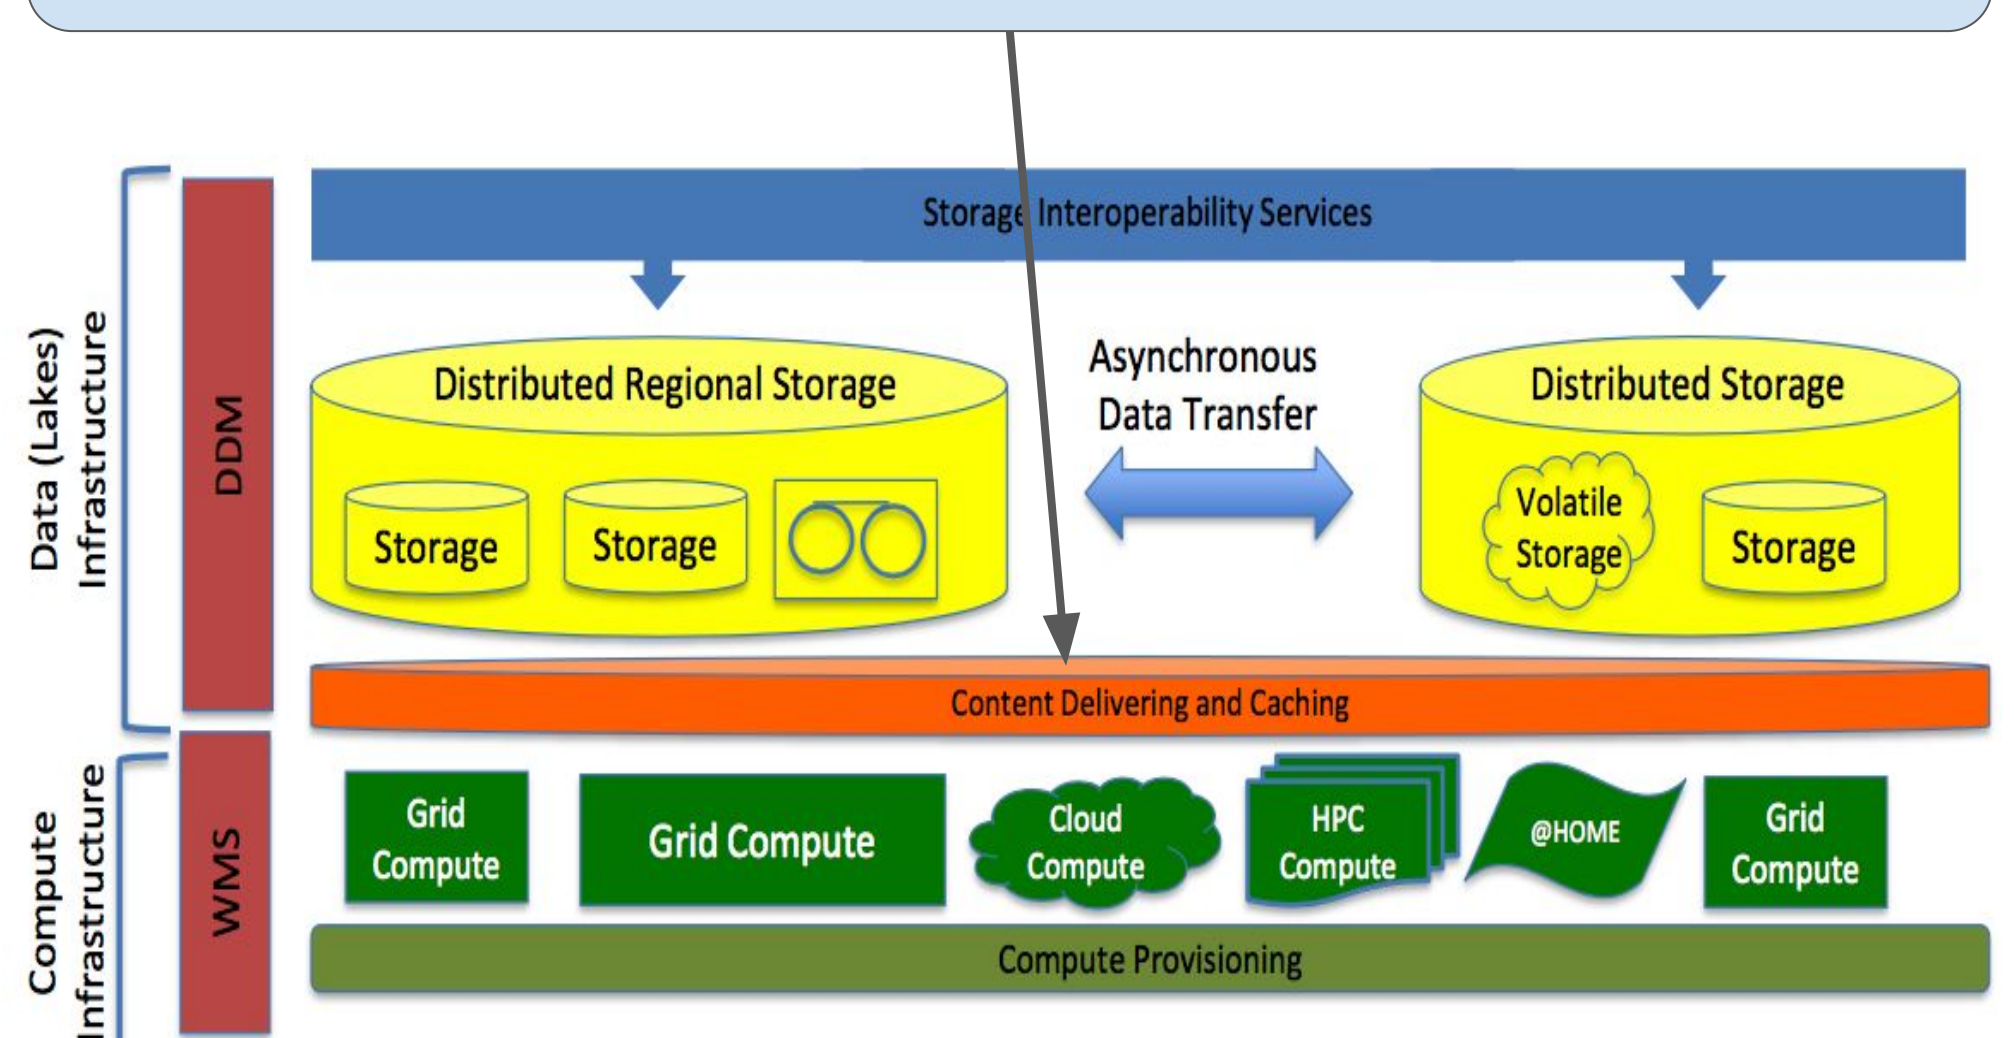
## Xcache on HPCs - How to be Optimal

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 "isetup 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
  - Adapt the HL-LHC Data Lake model and facilitate the data delivery
- Modern containerized Xcache deployment is powerful and critical.