

# Xrootd Demonstrator

Progress Update  
WLCG Collaboration Meeting

# The Data Access Problem

- Physicists need a simple and quick mechanism to access data from outside grid sites.
- The operating cost of our data management tools are too high for marginally staffed sites (T3s).
- Jobs should only degrade, not fail, when encountering data access issues.

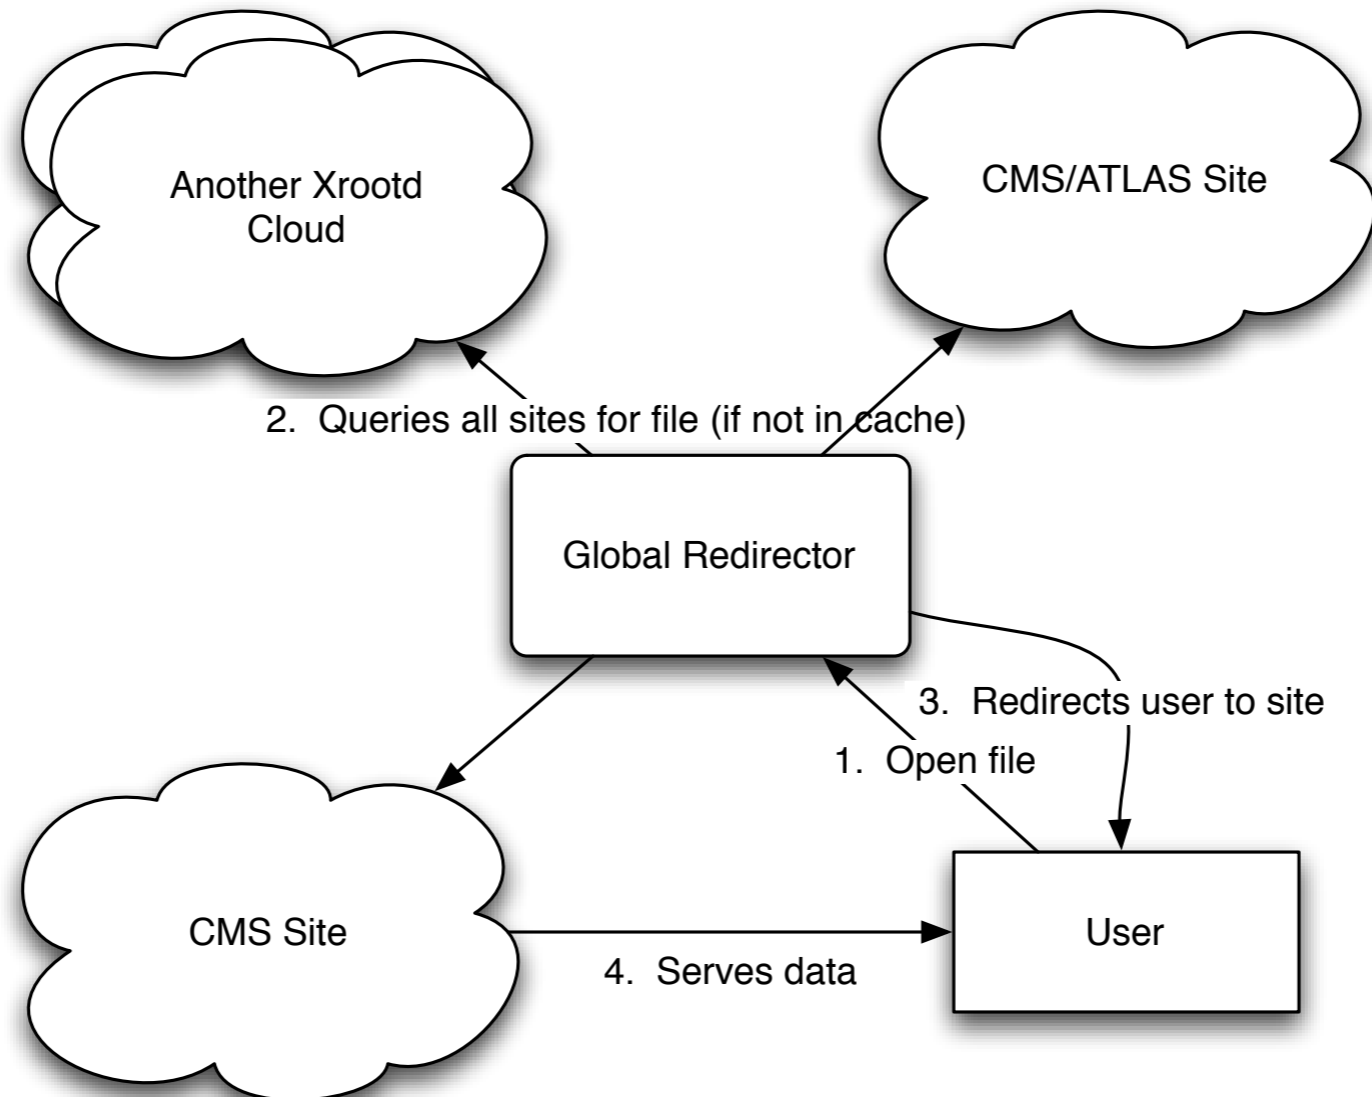
# Goals

- Develop a global data access architecture that is simple, reliable, and transparent.
- Data downloaded on demand, not pre-staged.
- Each file is represented by one URL; discovering sources is done by the system.
- Reliability!

# Xrootd Architecture

- We decided to look at a Xrootd-based global data service, because:
  - Xrootd protocol is ubiquitous in this field
  - Allows for a multi-level caching architecture (and deployed already in the field by ALICE!)
  - Can easily integrate non-Xrootd sites

# Xrootd Architecture

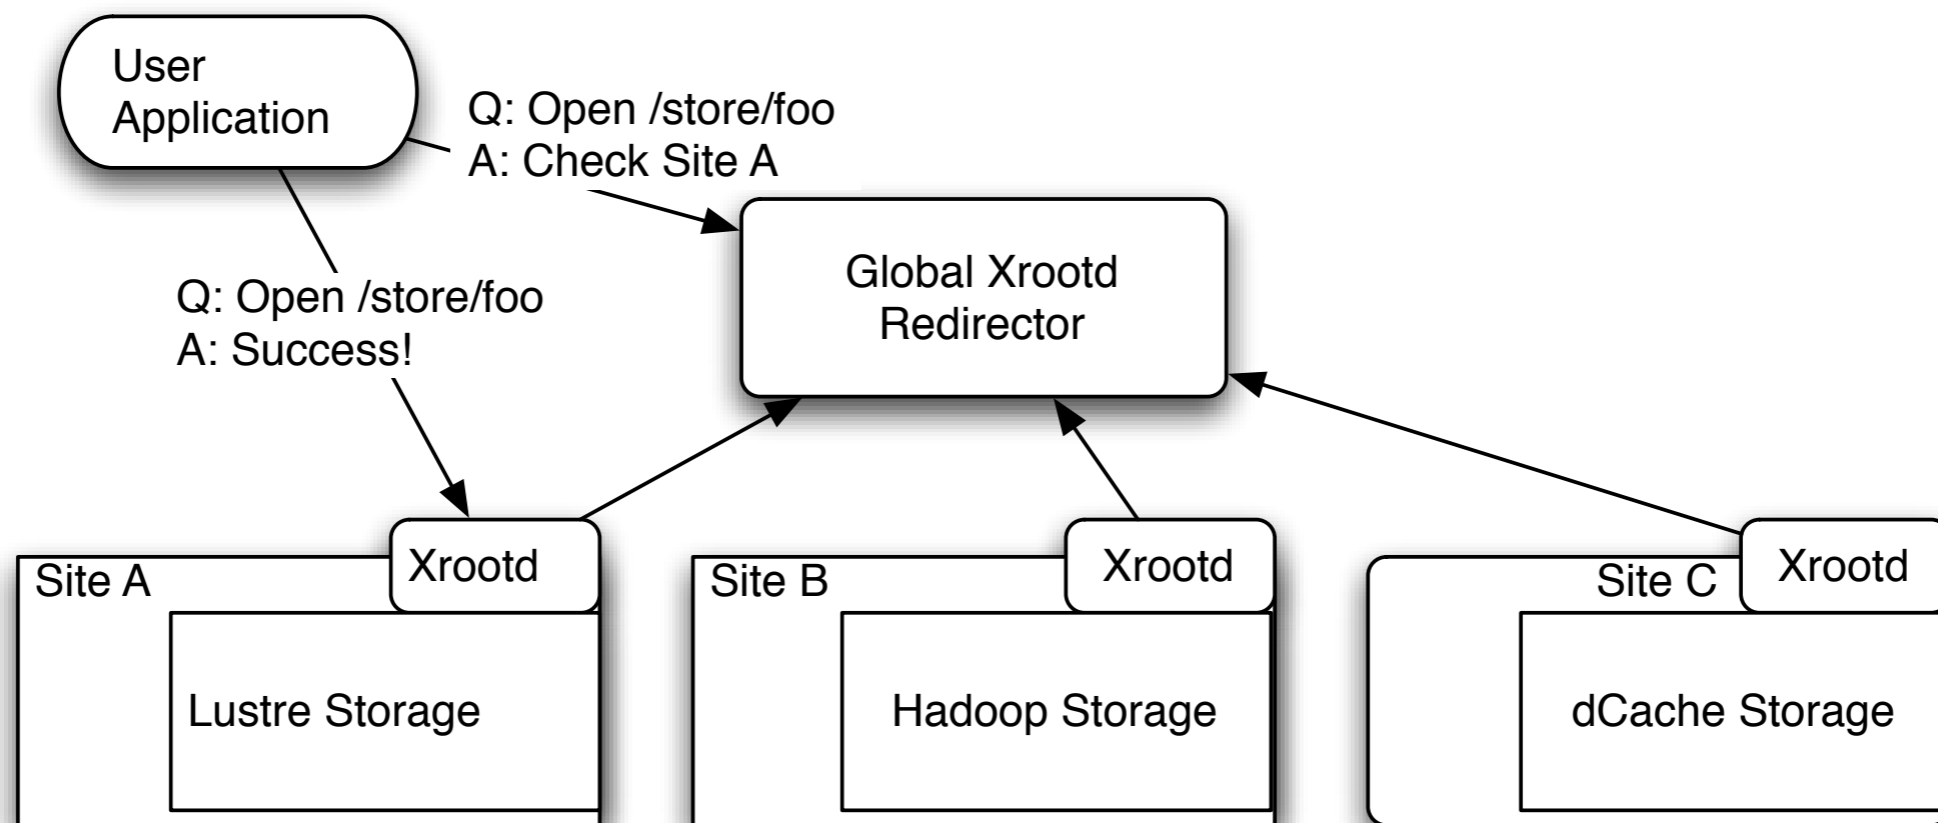


- Notes:

- “Global redirector” can be up to 16 actual hosts (highly available)
- Sites need to run at least 1 xrootd host, but can keep dCache/Lustre/HDFS/DPM/etc.
- Each site exports according to their capacity - no distinction in terms of T0 vs T1 vs T2.
- T3 is a special case; more later.

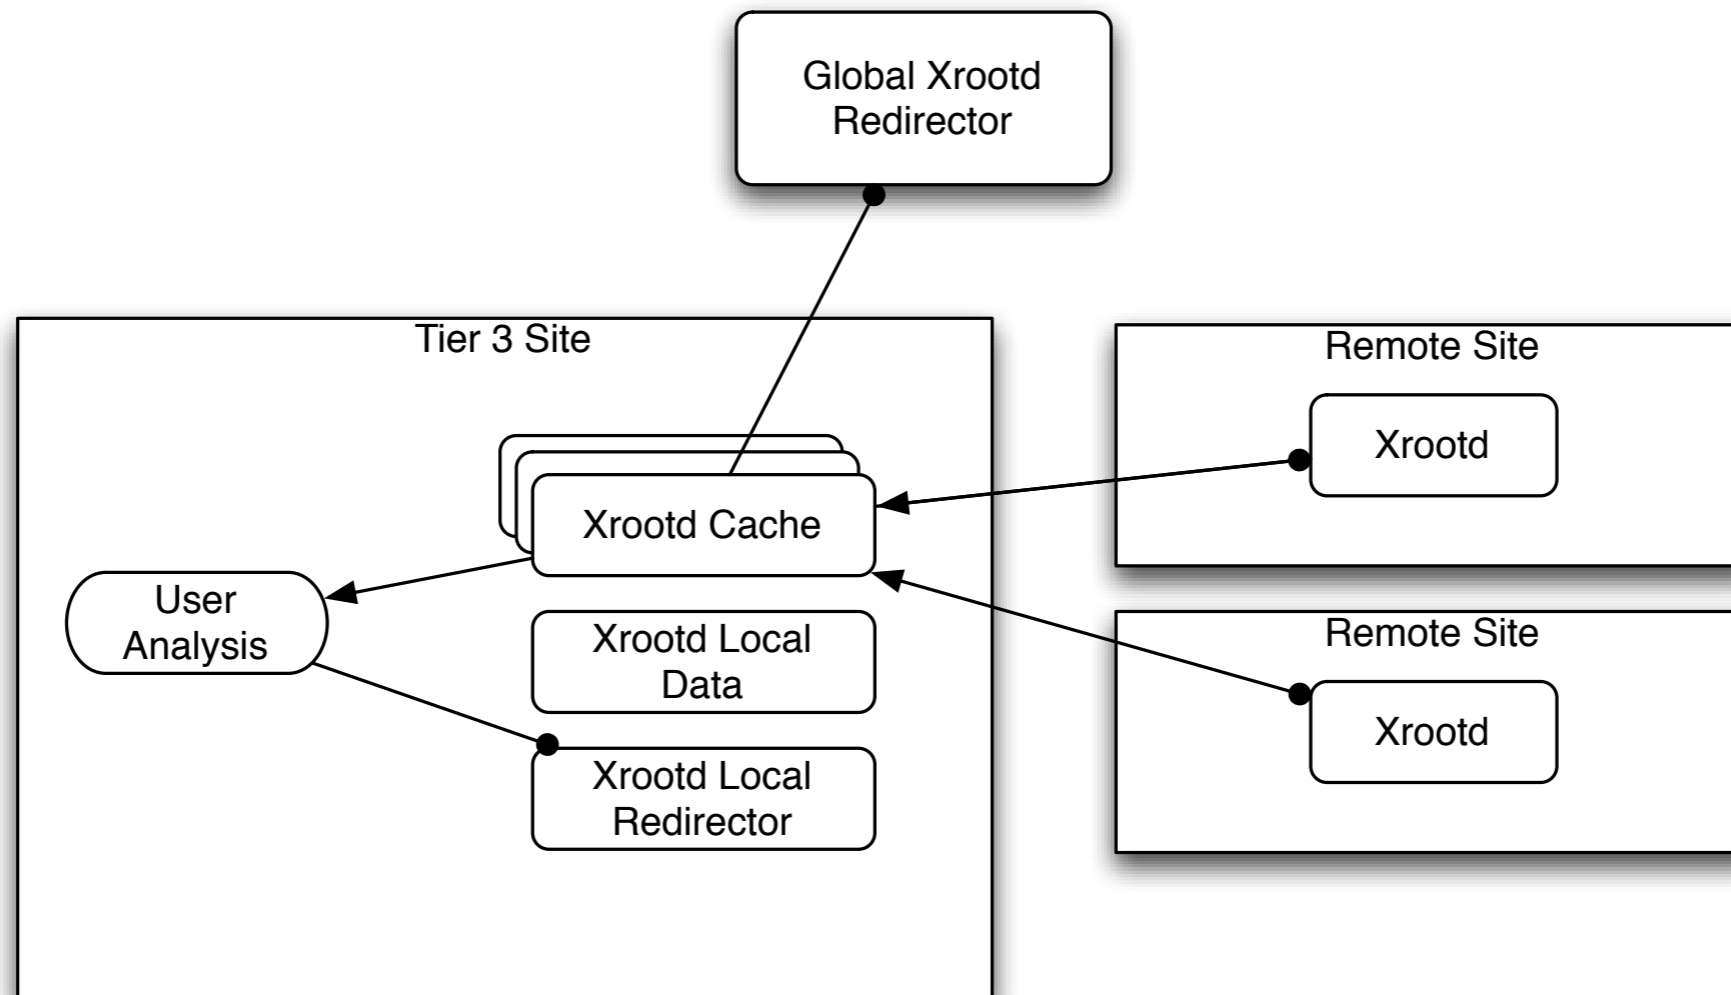
# User Use Cases

- Application running outside grid
- Incomplete dataset at a site



# Use Case

T3 Site - look! No data management responsibilities



The cache servers act as a client to the global system.  
Downloads from all possible sources as in bittorrent.

# Demonstrator Project

- A single global redirector, `xrootd.unl.edu`.
- Each site exports the same namespace (only need the LFN, no PFN/GUID stuff)
- GSI security (VOMS work being done by CERN).
- Have as many sites participating as possible.
- Show we can link it to experimental frameworks.
- “Flesh out” the operational issues and identify weaknesses.
  - Both globally and at the site level.



# Status

- Participating sites: Nebraska, Caltech, UCSD, Bari, FNAL.
- Upcoming sites: CERN, KIT
- Upcoming T3s: UCR, Omaha
- Demonstrated usage with CMSSW, ROOT, CMS event display.
- Up to 80% CPU efficiency on trans-atlantic link on experimental CMSSW versions!
- Submitted several rounds of CRAB-based analysis jobs demonstrating remote access.

# Metrics for Success

- *Not* in GB/s (we already can do that!)
- Being able to demonstrate promised capabilities to experiments.
- Understanding WAN impact of caching.
- Reduce costs of running a T3 site.

# Next Steps

- Need help!
- Looking for volunteer T1/T2/T3 sites.
  - Especially DPM sites (not yet represented!)
- Need users to abuse the system.
- Need feedback and concerns from sysadmins.