FIFE Workshop Data Handling

XROOTD

OLIVER GUTSCHE

Xrootd

- Provides POSIX-like access to files through the network with lowest possible latency
  - In plain English: you can open a file independent of your physical location through the network as if it would be local on your hard drive.
  - Was developed for the Root analysis framework to serve root files. However, the server is agnostic as to file type and provides byte level access to any type of file.
  - Can combine storage from different physical locations.
    - You open a file, the system figures out from where it serves the file.
    - Enables redundancy and load balancing
  - Supported authentication: GSI, Kerberos IV and V, as well as simple password authentication

- This solution provides
  - Access to experiment files independent of location on top of storage solutions
  - Enables off-site usage of files for processing and analysis
  - Suited for usage of GRID resources and other opportunistic resources for workflows that need access to input data

- Xrootd does not store your files, it provides access to them. It works on top of your storage and supports various storage systems (dCache, EOS, HDFS (Hadoop File System), etc.)

Xrootd

- Examples:
  - User stores ntuples for analysis at FNAL
    - He can access his files through bare root at CERN or any other place with network access
    - User is collaborating on analysis with off-site users and all are using common ntuples stored at FNAL -> collaborators have full access to ntuples from their home institute without having to copy files
  - Processing accesses files stored at FNAL
    - In case additional resources are available off-site (GRID, dedicated clusters, ..._) files can be accessed through Xrootd and the wide area network (WAN) and therefore enlarging the processing capabilities -> opportunistic resource usage
- Summary:
    - Data is accessible independent of the location through WAN
    - Xrootd enables load balancing and provides redundancy for files stored at multiple physical locations

Xrootd

- Xrootd setup schematics (taken from CMS)

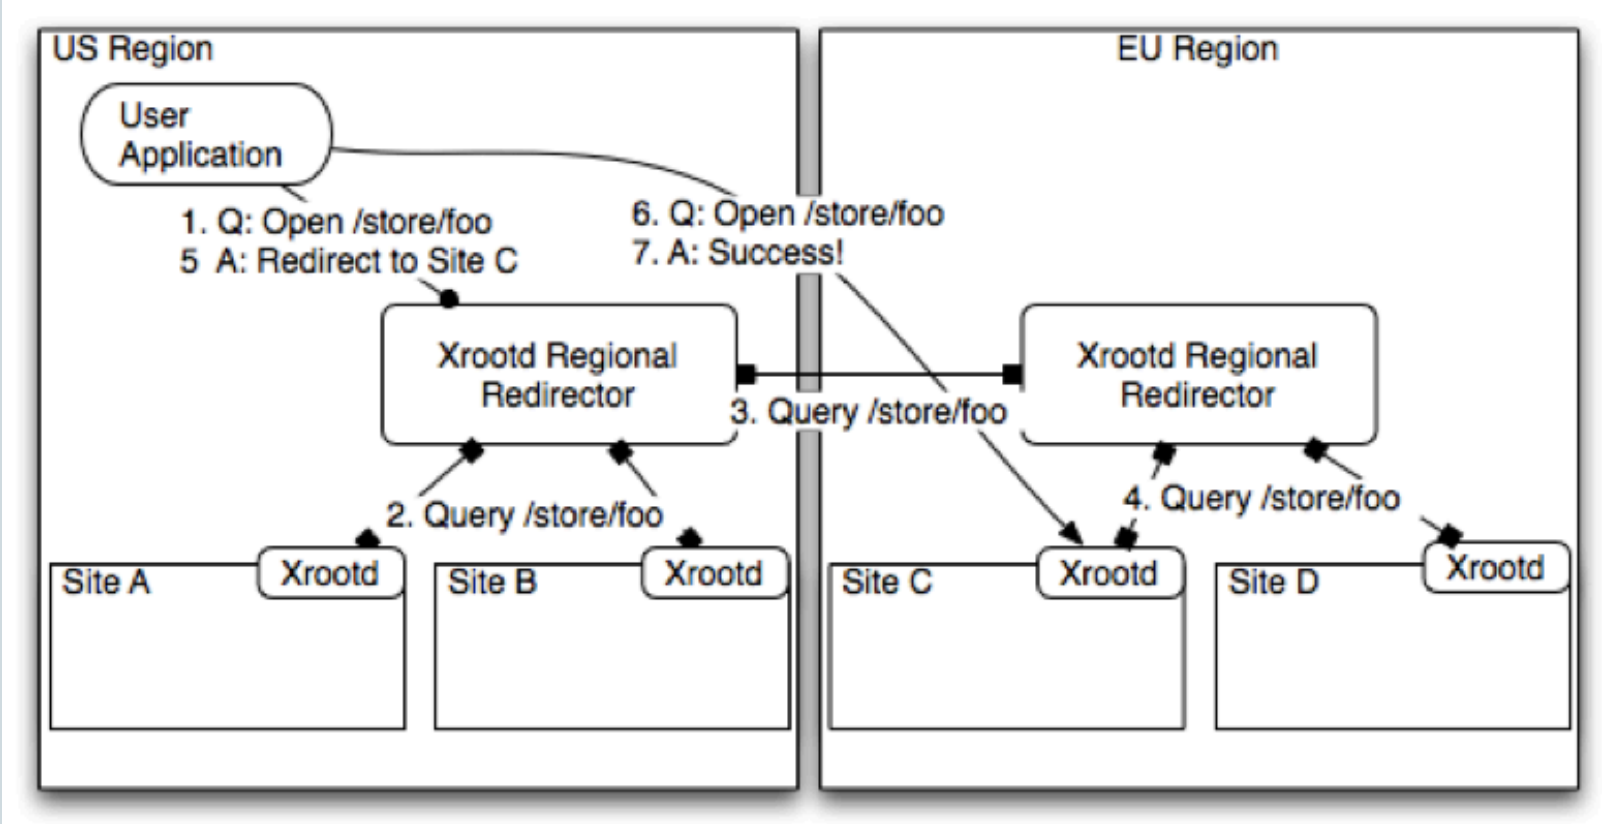

Xrootd

- configuration options for the service
  - Xrootd provides access to files stored on different storage systems (dCache, EOS, HDFS, …)
  - Xrootd can be configured per storage system and site, sharing between experiments possible
  - Xrootd can be enabled to do namespace translations to provide seamless access to logical file names
  - Xrootd allows authentication via GSI and Kerberos

Xrootd

- How is CMS using / planning to use Xrootd:
  - Merge all available disk storage worldwide (~80 PB) into one Xrootd federation
  - Combined with central dynamic file placement at sites, we can:
    - Redirect analysis jobs to sites with empty queues and access files through Xrootd as long as the sample hasn't been copied yet
    - Buffer file access problems by falling back to Xrootd access
    - Exploit the strong available networks (the US CMS T2 sites anticipate to have 100 Gbps connectivity in 2014, the dedicated LHCOPN between CERN and the T1 sites is at 120 Gbps) by
      - Running processing at any site (T1, T2, opportunistic) accessing input files through Xrootd -> more flexible and efficient use of CPU resources
  - Combined with global namespace of CMS (all files are accessed through Logical File Names (LFNs)), users have access to their files without having to deal with different storage technologies and different local paths
    - This also enables load balancing and redundancy.

Service Name

- Documentation
  - Xrootd project webpage: http://xrootd.slac.stanford.edu/
    - Introduction: http://xrootd.slac.stanford.edu/papers/Scalla-Intro.pdf
  - USCMS AAA project (Any Data, Anytime, Anywhere)
    - NSF proposal: http://osg-docdb.opensciencegrid.org/0010/001025/001/AnyDataAnyTimeAnyWhere.pdf
    - Project report during CENIC 2013: http://registration.cenic.org/cenic2013slides/AAA-Cenic2013.pdf