

Questions for JLAB

XRootD

- **What is the backend storage for XRootD at Jefferson Lab?**
 - a) Most of our current xrootd storage is backed by zfs file servers that also operate as nfs file server for internal systems. -kjs
 - b) Our two newest xrootd system also export a /lustre file system, one in read only configuration, the osdf origin we just stood up offers writes as well (both systems require tokens to authenticate) -kjs
- **What is the authentication setup (comanage, vault, etc.)?**
 - a) Most xrootd systems use scitokens to authenticate for reads and writes. We do have one special case where reads are unauthenticated, but that system doesn't allow writes. -kjs
 - b) The full authentication chain involves using a comanage registry for user management, which goes through a htvault to get bearer / vault tokens. -kjs
- **What is the total amount of storage available at JeffersonLab (break down based on type)? Any decisions on what will be put on tape and how much tape storage will be available?**
 - a) Storage at jlab is, very broadly, broken down into three types. -kjs
 - i) The tape library, which acts as long term storage. The system stores over 50pb of data, and can take in more as more tapes are added.
 - ii) The Lustre file systems. The relevant lustre filesystem for the EIC project has a total usable capacity of ~10PB, but that is spread among various groups. It is also split into a "volatile" area and a "cache" area. The EIC presently has 100TB of cache space granted to it and 500TB of volatile space.
 - iii) Work disk areas. We also provide a work area to the EIC group. This a zfs backed system exported to local nodes via nfs, this is also the area that is available via xrootd. The total space here is changing as the the local EIC group is moving towards more use of the lustre file system.
 - iv) Users have control over what files get put to tape. The amount of tape a given group uses is determined through discussions that I'm not involved in.
- **What sort of network throughput does the JLAB server experience? Are there any monitoring metrics over the last few months?**
 - a) EIC uses two xrootd servers at present, both have 100Gbs connections out, looking at dtn2201 (the system that does reads and writes), the activity over the past week hasn't come close to maxing that out.

- **What's the process of data management and migration currently when the disk at JLAB fills up?**

-> This is a policy decision. Kurt can answer.

a) That depends on what area is being used. -kjs

- i) The /work areas are completely user managed, with quotas being assigned at the directory level (for example, /work/eic2 has a 500TB quota). Clean up of old data on that disk is left up to the user, as would be any migration to tape. The only data migration we do for users is the move to new underlying storage when a particular system reaches replacement age.
- ii) The /volatile area is managed by a deletion policy that automatically deletes older files (determined by modification time) based on given quotas, there is also a 6 month deletion policy that removes files based on their atime being > 6 months.
- iii) The /cache area was originally used to write data back to tape, but that system is being replaced by a read-only cache system. Files can be explicitly put to tape using internal jlab tools.

- **By when, should we expect RUCIO in production? What's the current status?**

-> AP:

A. The temporary EIC rucio policy that we agreed in a meeting is ready. One-to-one rucio did name to storage path is tested in test rucio.

B. Next step is current data migration. We are trying to mount the current work area which has EIC data into the Xrootd. (running into some permission issues).

Status: work ongoing.

C. After B, I think we are ready for a production "like" system.

-> Update : Nov 6. We are ready to register files to rucio. The EIC area file system at Jlab /work/eic2/EPIC is mounted to our xrootd instance.

- **Has there been tests conducted with a JLAB RUCIO instance that interfaces with a storage element at BNL? What were the conclusions?**

-> AP: Yes, the test was performed. We managed successful bi-directional transfers between JLAB and BNL. The test with X509 authentication.

-> More about test:

A.

- A file is upload to Jlab xrootd rse using rucio upload.

- A replication rule is created to have 1 copy at BNL xrootd rse.
- Rucio on the backend, submitted a transfer to FTS for that file.
- FTS initiated 3rd party pull copy to storages (BNL xrootd pull file from JLAB xrootd).
- File is copied to BNL xrootd and info is properly update

B. Repeated the same process but now file is uploaded using rucio upload to BNL xrootd and all the procedure is same. Replica is created at Jlab xrootd.

The test was performed using x509 auth.

Questions for ePIC Collaboration

- **ePIC officially chose XRootD as the method for data access. How flexible is ePIC in terms of the choice of backend storage? Does it need to be the same across all site?**

To first order, as long as the backend storage has token-based access via XRootD, the specific choice does not matter.

Questions for BNL

- **What is the status of the test storage element with which the bidirectional transfer with Jefferson Lab was tested?**

We are in the process of upgrading the integration/test instance. In principle it is expected to be fully functional in 2 weeks if not before.

- **What are the hurdles towards transitioning to token-based authentication from X.509?**

dcache supports token based authentication as a part of WLCG our configuration is based on this guidance

<https://wlcg-authz-wg.github.io/wlcg-authz-docs/token-based-authorization/configuration/dcache/>, you will need to provide the necessary token based configurations tailored to ePIC (oidc.provider,oidc.audience-targets)

Comment from Anil: Rucio/FTS3 is not capable of using cilogon token for transfer yet.

Contact List

Name	Relevant Expertise	Email
Sakib Rahman	ePIC Production Co-Convener	srahman1@bnl.gov
Kurt Strosahl	Jefferson Lab XRootD	strosahl@jlab.org
Anil Panta	RUCIO	panta@jlab.org
Xin Zhao	Setting up and testing OSG Access Point +	xzhao@bnl.gov

	XRootD at BNL	
--	---------------	--