

DUNE Authentication Needs

Steven Timm

10 Sep 2019

For DUNE Software + Computing


At WLCG Pre-GDB AuthZ session

My Hats

An orange fedora-style hat with a black band.

DUNE Data Management
DUNE VO Security Contact

My comments today given with my DUNE hat, but informed by my experience in running Grid services at Fermilab from 2005-present.

A blue baseball cap.

FERMILAB VOMS
service operator
FERMILAB HEPCloud
service operator

Outline

- I will describe
 - Current layout of DUNE distributed computing and storage
 - Current authentication/authorization scheme and the problems we currently face.
 - Expected evolution and growth of our setup
 - Questions I have based on what I've seen presented of the suggested technologies to date.
 - I will purposely keep to logical descriptions as much as possible
- Important caveat:
 - DUNE experiment currently does not manage our own AuthN/AuthZ
 - The fact that I am a DUNE collaborator and also operator of VOMS at Fermilab is pure coincidence.
 - We rely on Fermilab site-provided utilities for group membership, certificates, Kerberos, and hope to continue to do so in new regime.

DUNE Web Services

- ServiceNow
- Wiki.dunescience.org (mediawiki)
- www.dunescience.org (wordpress)
- Slack.dunescience.com
- ZOOM
- Indico (cern.ch and fnal.gov)
- EDMS (CERN)
- Cdcvs.fnal.gov. (redmine)
- Github.com
- Github.com
- Docs.dunescience.org (docdb)
- Dune-data.fnal.gov
- Landscape.fnal.gov (monitoring)
- CERN Twiki (protodune)
- Cern E-log (protodune)
- Collaboration Database
- ECL Electronic Logbook

Job submission and file services

- (most of which use certificate auth right now)
- SAM*
- Jobsub*
- POMS*
- Rucio*
- dCache/Enstore
- EOS* / Castor* / CTA
- HEPCloud
- GlideinWMS
- Fermi FTS, CERN FTS-3
- Lots and lots of databases everywhere
- DIRAC*

State of DUNE Dist. Computing

- DUNE already computes at > 25 compute sites around the world
 - All possible grid middlewares, all possible docker/singularity config
 - Global pool unified through GlideinWMS following CMS model.
- DUNE has 12 storage elements declared in Rucio
 - All possible types of SE recognized by the WLCG plus some that aren't. (and all types of manual ad-hoc mechanisms for cert DN's)
 - Most can't 3rd-party-transfer between each other even now
 - When gridFTP goes away we lose our lowest common denominator.
- Need to *get* the data flowing in the current regime and then *keep* the data flowing in the post-cert regime.

DUNE Production AuthN/AuthZ

- DUNE production servers run on IGTF certified InCommon certs
- DUNE production job submission done with IGTF certified InCommon cert
- DUNE data movement CERN->FNAL->everywhere else done with IGTF certified InCommon Cert.
- Production jobs can run anywhere, works fine.
- Long proxy of service cert is stored in MyProxy and shorter-lived delegations used in running Grid jobs.

DUNE grid user AuthN/AuthZ

- Utility called “cigetcert” is used to contact CILogon on behalf of user and retrieve a certificate. (also Web UI available)
- Can use Fermi Kerberos credential or Fermi single sign on credential to authenticate
- Automatically done on behalf of user when they submit a job with jobsub.
- 1-month long cert stored in MyProxy
- Jobsub server is only machine that can get myproxy delegations of shorter proxies and get them pushed out to running jobs
- Most users never have to voms-proxy-init, openssl pkcs12, grid-proxy-init, or any of that—certificate is transparent to them
- Unfortunately cilogon-basic not IGTF classic accredited (see earlier Mine Altunay talk). Previous Kerberos CA (SLCS) was.

Current Challenges

- Initial Fermi ID and renewals taking much longer to process
- Initial CERN ID always took a long time to process.
- It is likely that DUNE will have some collaborators who are unbadgeable at Fermilab and others who are unbadgeable at CERN.
- Need ways to:
 - 1) Make the “How To Join DUNE and Get an ID” page public
 - 2) Give access to computing technical people who are not necessarily DUNE collaborators—even whose institution may not be part of DUNE.
 - 3) Make proxy service for file fetching

DUNE VO Membership

- Huge amount of work recently done at Fermilab to automate collaboration accounts, the FERRY project.
- User that wants to join DUNE has to get sign off from institutional board representative
- Request for DUNE computing account (and the associated Fermilab ID you need to get it) goes through DUNE collaboration office at Fermilab
- Once approved all account creation and VO population is automatic.
- Can add extra non CILogon-basic certs to your VO entry via a service desk form.
- VOMS-Admin interface is still there for now but it's read-only.

Questions of an old Grid guy

- All diagrams I've seen of Indigo-IAM include an input from CERN HR database as the membership information.
- What happens if our membership comes from FERRY instead.. Is hackery necessary? Can Indigo-IAM function at all without CERN HR DB? Will our European colleagues trust it if we don't use CERN HR DB?
- How do we build a trust relationship of which tokens we trust and which we don't?
 - The equivalent of CA and CRL files?
 - The equivalent of VOMS LSC files?
 - How is VO Trust expressed? How do we distribute all these files everywhere? Do I have to beg to get the DUNE key on each site?
- dCache Macaroons—they sound tasty and are supposed to be compatible but are they really?
- Token capabilities sound very handy, how can I leverage them to restrict permissions in rucio and rucio-admin?
- How do we define the interim state so I can move data from CERN to Fermilab and the rest of the world in the fall of 2021 which is our next beam run?