

Security

I noticed quite a bit of interest in federated identity and in not requiring users to manage their own certificates.

One security topic of particular interest to me came not in one of the talks but in a conversation: Predrag Buncic suggested that to avoid having to send certificates to jobs, create a randomly named directory on the storage system that is writable by everyone but only readable by the real user. Send the name of the directory to only that job, so the only privilege the job has is to write to that directory; if the job is compromised, there is not much lost. Then after the job is finished, another process comes along and renames the output files to their real location and deletes the random directory name.

Talks:

- Enabling Federated Access for HEP
 - Hannah Short's talk about the AARC project, Europe's equivalent to what we're doing with cigetcert, but far more complex, primarily because they don't want to depend on Identity Providers supporting ECP.
<https://indico.cern.ch/event/505613/contributions/2227690/>
- Access to WLCG resources: The X509-free pilot
 - Hannah Short also gave this talk. It is not about avoiding X509 as the title implies, it's about managing the certs for users. It is about setting up a CILogon-equivalent federated X.509 certificate CA. The talk includes discussion about modifying VOMS to allow registrations without having a certificate, something we also need for DCAFI phase 2.
 - <https://indico.cern.ch/event/505613/contributions/2227687/>
- dCache, towards Federated Identities and Anonymized Delegation
 - Paul Millar's talk about OAuth2 integration into dCache, very interesting.
 - <https://indico.cern.ch/event/505613/contributions/2227700/>
- Internal security consulting, reviews and penetration testing at CERN
 - <https://indico.cern.ch/event/505613/contributions/2227686/>
- The future of academic computing security
 - <https://indico.cern.ch/event/505613/contributions/2227689/>

Posters:

- Grid Access with Federated Identities
 - This is my cigetcert/DCAFI poster
 - <https://indico.cern.ch/event/505613/contributions/2259539/>
- Developing the Traceability Model to meet the Requirements of an Evolving Distributed Computing Infrastructure
 - This is about the traceability & isolation working group, including singularity, presented by Ian Collier.

- <https://indico.cern.ch/event/505613/contributions/2227727/>
- BelleII@home: Integrate volunteer computing resources into DIRAC in a secure way
 - They use an intermediate gateway server that holds the user certificate, to move the job output from the BOINC server where a BOINC job writes its output.
 - <https://indico.cern.ch/event/505613/contributions/2227710/>
- Web technology detection - for asset inventory and vulnerability management
 - <https://indico.cern.ch/event/505613/contributions/2227711/>
- A Security Monitoring Framework For Container Based HEP Computing Infrastructures
 - <https://indico.cern.ch/event/505613/contributions/2227720/>
- IPv6 Security
 - <https://indico.cern.ch/event/505613/contributions/2227722/>
- The INDIGO-DataCloud Authentication and Authorisation Infrastructure
 - <https://indico.cern.ch/event/505613/contributions/2227724/>
- A lightweight access control solution for Openstack
 - <https://indico.cern.ch/event/505613/contributions/2227726/>
- Use of a hardware token for Grid authentication by the MICE data distribution framework
 - <https://indico.cern.ch/event/505613/contributions/2227728/>

CVMFS

Talks:

- Accessing Data Federations with CVMFS
 - Derek Weitzel's talk about the CVMFS enhancements to support osgstorage.org repositories.
 - <https://indico.cern.ch/event/505613/contributions/2230923/>
- Using Shifter to bring Containerized CVMFS to HPC
 - Lisa Gerhardt's discussion of how they get CVMFS into the NERSC supercomputer: they cheat. They take a snapshot of the entire repository, compress it with squashfs, and mount the whole huge filesystem image inside of a container on each worker node. The worker node then opens the one file on Lustre read-only and reads the pieces that it needs. It performs well, but the CVMFS repositories are huge because they're assumed to never need to all be read, and generating them is very resource intensive.
 - <https://indico.cern.ch/event/505613/contributions/2227429/>

Posters:

- New Directions in the CernVM File System
 - Jakob Blomer's poster on recent and upcoming CVMFS developments.
 - <https://indico.cern.ch/event/505613/contributions/2230961/>

Frontier/Squid

Talks:

- Web Proxy Auto Discovery for WLCG
 - My talk about wlcg-wpad.cern.ch/wpad.dat
 - <https://indico.cern.ch/event/505613/contributions/2230709/>
- Conditions Database for Belle II experiment
 - Lynn Wood's talk about the Belle II conditions system. They have a system much like the one used for FIFE, with a REST-based web api. They use IOVs and payloads like the CMS conditions. He did not talk about caching but said they're still planning to add it. I talked with him about Frontier last year, and again after his talk. They store their payloads outside of the database, probably in CVMFS (although he also didn't mention that in the talk), so depending on the number of parallel jobs they want to support they might be able to get away with just scaling up their central server. I advised that if that is not enough scaling and they want to use site squids, they should insert Frontier in the path rather than reimplementing many of the features in the frontier client.
 - <https://indico.cern.ch/event/505613/contributions/2230845/>
- Functional tests of a prototype for the CMS-ATLAS common non-event data handling framework
 - Roland Sipos' talk about the new exploratory framework for Run3 conditions data for both CMS & ATLAS. It is much like the Belle II design, except they aren't at this point anyway planning on storing payloads outside the database. They intend to use Frontier as the caching layer between the clients and their server, but hadn't yet involved me much. During the conference I set up a Frontier servlet to read their test server and discussed the requirements that Frontier will need on their server in order to do caching well. The long term plan for operations would be to add their server software to the Frontier launchpad machines.
 - <https://indico.cern.ch/event/505613/contributions/2230853/>