

Pre-GDB, GDB, and IRIS-HEP Retreat: Updates from Recent Meetings

**Brian Bockelman
OSG Technology Area Coordinator
Associate Scientist,
Morgridge Institute for Research**





Recent Meetings Involving OSG



- The week of Sept 9 featured three separate meetings involving OSG (and all at FNAL!):
 - GDB: WLCG “Grid Deployment Board”. The technical coordination body of the WLCG.
 - Meets once a month to discuss a broad array of technical topics.
 - First time in recent memory having the meeting in the US.
 - Pre-GDB: A day-long meeting to discuss a specific technical topic at length.
 - This meeting was about **authorization and authentication on the WLCG**.
 - IRIS-HEP Retreat: (Likely) Annual planning meeting for the IRIS-HEP institute.
 - Broad across the whole institute, but had a few specific items for OSG-LHC.

- This session was organized by the WLCG Authentication and Authorization Working Group.
- Morning session was spent finalizing the “WLCG JWT Profile”.
 - Since then, we have published this document with DOI <https://doi.org/10.5281/zenodo.3460258>.
 - Caps off about 18 months of work.
 - This documents the format and interpretation of future security tokens for WLCG.
 - Heavily influenced by the work done in the SciTokens project.
- Afternoon session was presentations from various technologies teams:
 - **IAM**: Identity and Access Management server, from INFN. Allows VO to manage group membership and access permissions.
 - **SciTokens**: Working on end-to-end of distributed capability tokens; talk was about interoperability with WLCG tokens.
 - **FNAL** and **DUNE**: Initial plans on converting infrastructure to token-based auth.



- See <https://indico.fnal.gov/event/21485/>
- Organized as part of the WLCG; driven by the SLATE team who are promoting remotely managed services at the edge.
 - Raises significant security & trust questions that need to be addressed.
 - Pulling in a good cross-section of the security community, including WLCG, EGI, OSG, and CTSC.
- Most of this meeting focused on drafting the charge and started working on policy language to describe the mdoel.

Next Hill To Climb

Where to next?

- Both the storage (XRootD) and compute (HTCondor-CE) have initial support for SciTokens; we need to demonstrate the end-to-end.
- I propose four challenges that, theoretically, require almost no development - but would demonstrate broad coordination!
 - **Challenge 1:** Acquire a token from IAM via OAuth2 and use it to upload files to dCache and XRootD.
 - **Challenge 2:** Acquire a token from IAM via OAuth2 and use it to submit a pilot job.
 - **Challenge 3:** Have the HTCondor “credmon” acting as an OAuth2 client acquire a token from IAM, send the token along with a job, and have the job stage out to dCache.
 - **Challenge 4:** Author a whitepaper describing how our community plans to use tokens for data management – including Rucio, FTS, IAM, XRootD, dCache, and others.

I suspect (1) is nearly done if you trawl through private email threads. The important aspects of an interoperability challenge would be to create knowledge and documentation for others to follow!

Why do we care?



- Why do we care about this work?
 - Some of the most significant investment in technology transformation is the migration from GSI.
 - We have been using the SciTokens technology as a key plank in the replacement strategy for GSI.
 - This document confirms WLCG commitment to head in the same direction.
 - WLCG JWT and SciTokens profiles are similar enough to share a client library.
 - By the end of the meeting, we were able to show the IAM server can produce SciTokens-compatible tokens (enough to send jobs to the HTCondor-CE).



- <https://indico.cern.ch/event/739882/>
- Topics included:
 - DUNE [computing outlook](#).
 - IceCube [computing outlook](#).
 - Rucio: [News & Outlook](#), work toward [multi-VO instance](#), and [DUNE plans](#).
 - Networking: [MULTI-One](#), [Network Virtualization](#), [SAND project status](#).
 - OSG Coordination topics: [OSG Overview](#), [global VO configuration](#).
- A few select highlights follow.

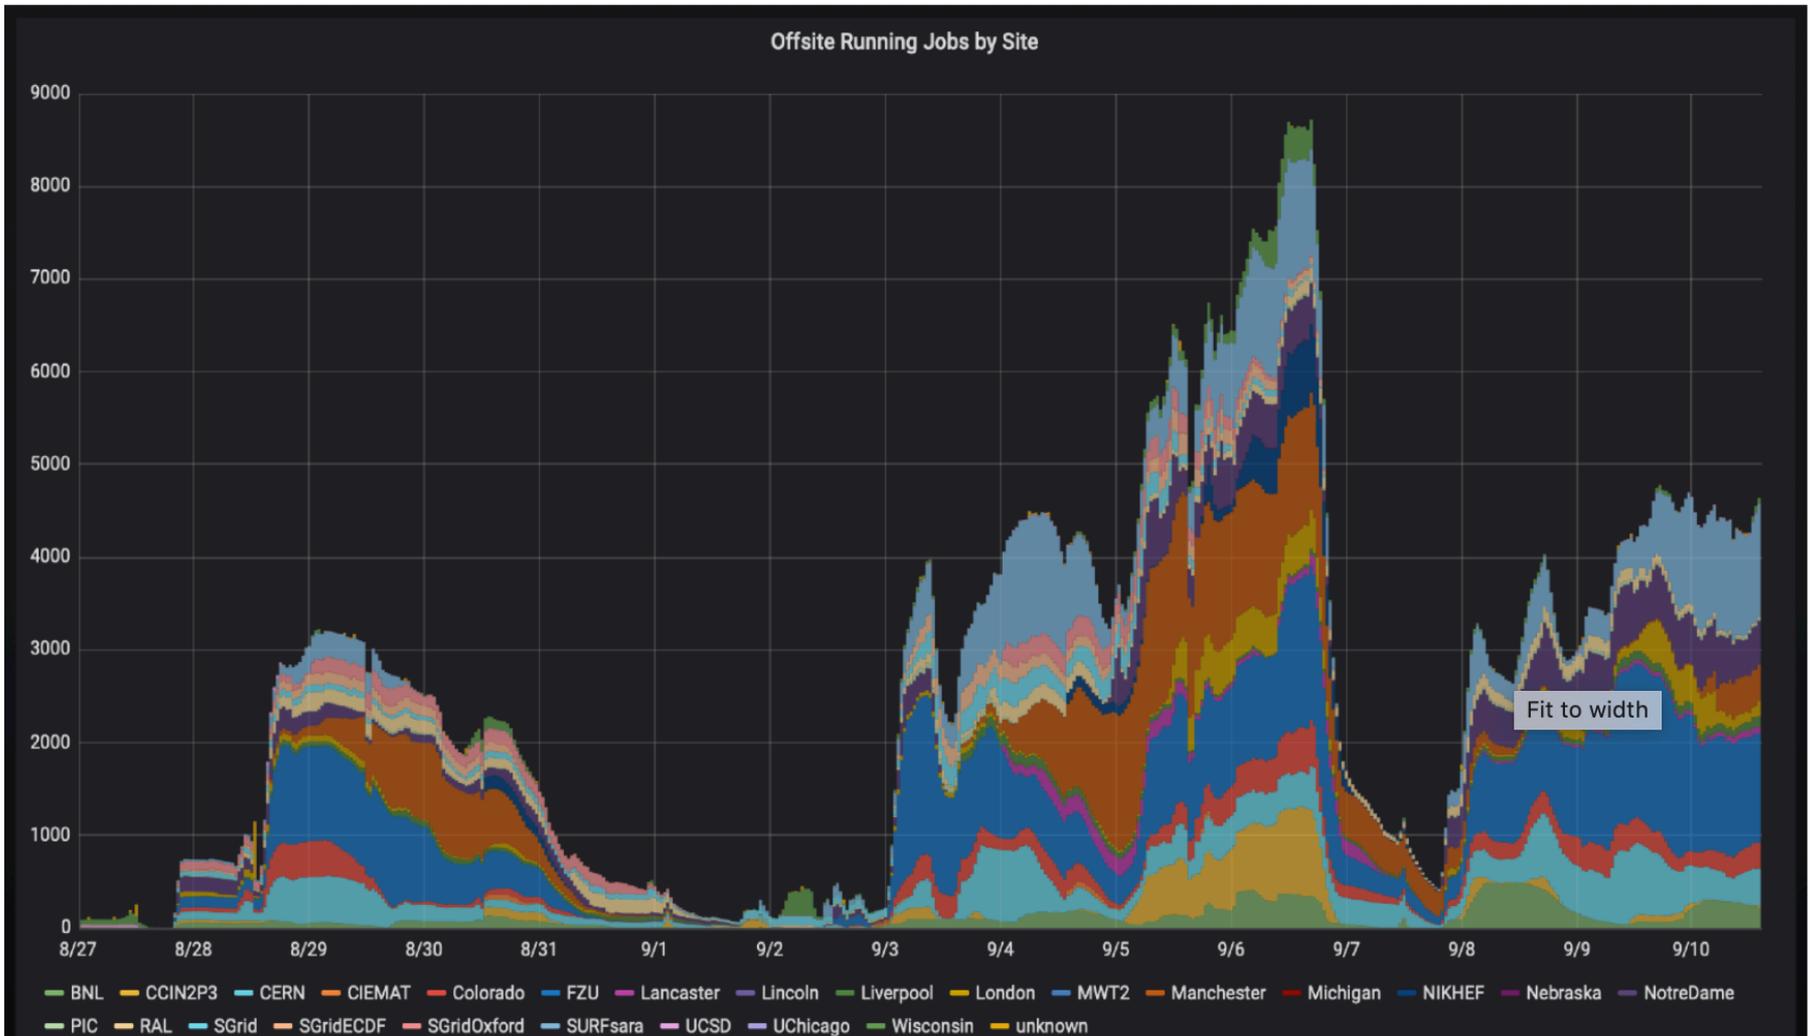


Open Science Grid

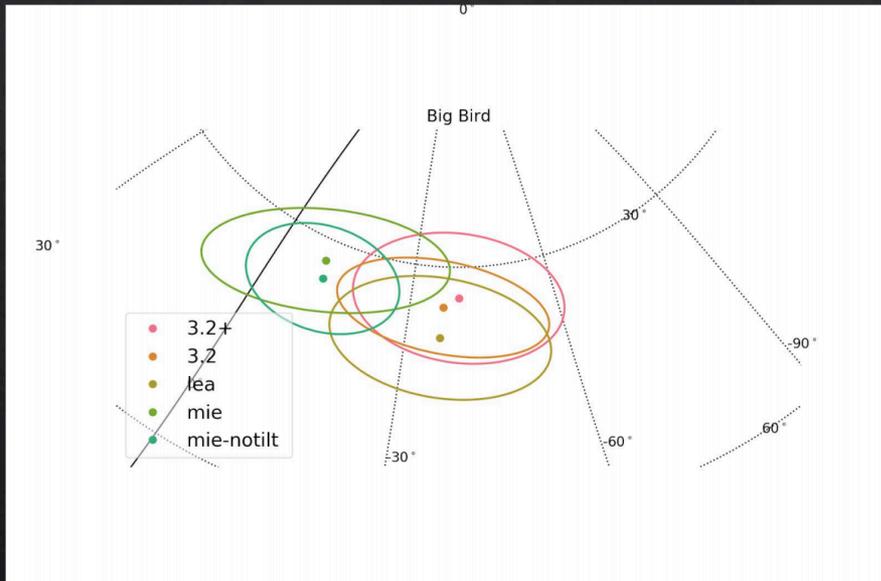
DUNE – Grid Activities



Last two weeks, excluding FNAL



Multi-Messenger Astrophysics



- ◇ Want to notify community as **soon as possible** of best-possible reconstructed location of interesting events
- ◇ Reconstructed by scanning across sky for best possible location
- ◇ Bursty resources - Requires massive ($O(10000)$ cores) offline (in the Northern Hemisphere) computing resources for short amount of time (30ish min)
- ◇ Fast ($O(\text{min})$) ramp-up is crucial
- ◇ Exploring cloud resources as means to get these kinds of resources - Funded through Internet2's E-CAS initiative

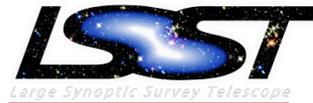
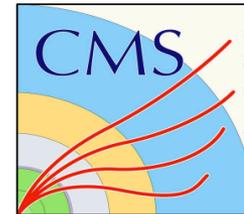


Open Science Grid

Rucio – Community Building



Community





Ongoing/planned developments

- Globus Online as a transfertool in Rucio
 - Currently FTS is the only transfertool supported by Rucio
 - Development by BNL for Light source experiment and other usecases
 - Working Prototype implementation
- Token-based authentication (OIDC)
 - Extends authentication to OpenID Connect
 - Integration of IAM workflows to Rucio
 - Development done with funding from XDC project (H2020)
 - Authentication to Rucio, authentication from Rucio to Storage/FTS
- Community specific configuration, permissions and policies
 - Simplified way to enable community specific policies in Rucio



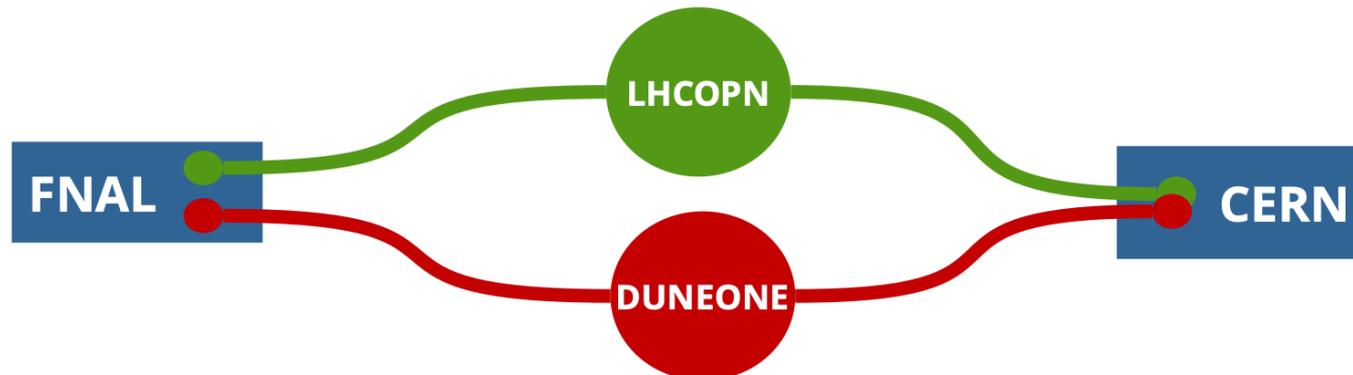
LHC-protoDUNE use case

Just agreed with FNAL to prototype the solution with protoDUNE between CERN and FNAL (protoDUNE is currently using the LHCOPN link of FNAL)

New VPN DUNEONE to be agreed with ESnet

No impact on existing protoDUNE traffic and other sites

Resources already distinct at FNAL. Mixed up at CERN





- See: <https://indico.cern.ch/event/840472/>
- Specific OSG topics:
 - XCache data integrity & requirements planning.
 - GridFTP / GSI migration: see earlier presentation for details.
 - New security policies: see [earlier presentation](#) for details.



XCache Development



- The XRootD/XCache software plays a central role in the OSG-LHC storage evolution.
 - Used by USATLAS, USCMS, and OSG (StashCache).
 - Coordinating the evolution and development priorities becomes important: small development team can't be stretched in too many ways.
- We covered a number of topics, including:
 - Packaging priorities (RPMs vs Docker vs k8s): current approach - focusing on RPM / Docker per VO – was kept.
 - Monitoring: ATLAS, CMS, and OSG will share a monitoring infrastructure run by OSG. OSG's focus will be to validate what we have, then help migrate to new cache monitoring infrastructure (Q2 2020).
 - Data Integrity:
 - In transit: will rely on XRootD-over-TLS, expected Q1 2020.
 - At rest: IRIS-HEP DOMA will develop a HEP-specific integrity checking mechanism for now. Longer-term (summer 2020), XRootD development team will develop a more generic mechanism.



Take-Home Message



- The OSG Community is leading across a number of lines of work:
 - Our Authentication and Authorization approach is being adopted across the WLCG community.
 - OSG plays a central role in the coordinating XCache community (even though we don't develop it).
 - Pushing forward security policy work for new models.
- Having the meeting at FNAL was fortuitously timed to increase engagement with DUNE, especially with the WLCG.