FIFE Notes - April 2016 News

for

Distributed Computing at Fermilab

2016 Open Science Grid All-Hands Meeting

Every spring, the entire Open Science Grid (OSG) community–consisting of resource owners and operators, users, and staff–gathers at the annual OSG all-hands meeting. The 2016 OSG all-hands meeting was held between Monday, March 14 and Thursday, March 17 at Clemson University in Clemson, SC, thanks in large part to Jim Bottum, the CIO and vice provost for technology at Clemson. The OSG is, as befitting a vehicle for distributed high-throughput computing, a highly distributed organization, with a community spread out across the US. As such, the all-hands meeting offers one of the few opportunities for face-to-face interaction for this community. Some of the highlights of the past year in the OSG were noted, including the passage of the 1 billion CPU hour/year

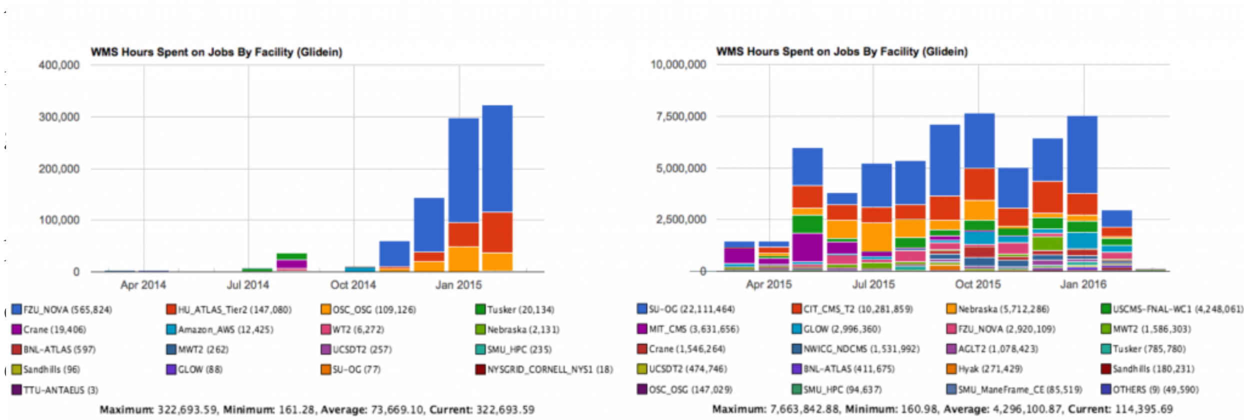

Opportunistic OSG usage by FIFE experiments in the 12-month period leading to the 2015 OSG AHM (left) and leading to the 2016 OSG AHM (right) shows a dramatic increase in usage.

At the main plenary session on Wednesday, Ray Culbertson of Fermilab discussed the mu2e experiment and their rapid ramp-up in using the OSG [1]. A parallel session on Monday dedicated to non-LHC physics experiments saw presentations by Ken Herner about the DES follow-up-to-LIGO campaign [2], Steve Timm about DUNE [3], and a more technical discussion of mu2e's OSG computing by Ray Culbertson [4]. Ken Herner also presented an overview of the FIFE project with an emphasis on recent progress using OSG resources at a plenary session on Tuesday [5].

Despite the "Grid" in Open Science Grid, recent years have shown the importance of getting OSG users access to “non-grid” resources such as commercial clouds and High Performance Computing (HPC) platforms. Burt Holzman presented progress on the Fermilab HEPCloud project, including successful use of Amazon Web Services resources by both CMS and NOvA [6]. Finally, the critical work of security in OSG operations continues; Jeny Teheran [7] and Dave Dykstra [8] reported on transitions to a federated identity system (CILogon).

1. The Mu2e Experiment at Fermilab : Experience with OSG Opportunistic

2. Dark Energy Survey on the OSG

Lunch at the OSG All-hands Meeting: Photo courtesy Ken Herner

— Bo Jayatilaka

HEP Cloud: How to add thousands of computers to your data center in a day

Throughout any given year, the need of the HEP community to consume computing resources is not constant. It follows cycles of peaks and valleys driven by holiday schedules, conference dates and other factors. Because of this, the classical method of provisioning these resources at providing facilities has drawbacks, such as potential over-provisioning. Grid federations like Open Science Grid offer opportunistic access to the excess capacity so that no cycle goes unused. However, as the appetite for computing increases, so does the need to maximize cost efficiency by developing a model for dynamically provisioning resources only when they're needed.

To address this issue, the HEP Cloud project was launched by the Scientific Computing Division in June 2015. Its goal is to develop a virtual facility that provides a common interface to access a variety of physical computing resources, including local clusters, grids, high-performance computers, and community and commercial clouds. Now in its first phase, the project is evaluating the use of the "elastic" provisioning model offered by commercial clouds such as Amazon Web Services. In this model, resources are rented and provisioned dynamically over the Internet as needed.

The HEP Cloud project team successfully demonstrated this elastic model for CMS in January and February using Amazon Web Services as discussed here and shown below:

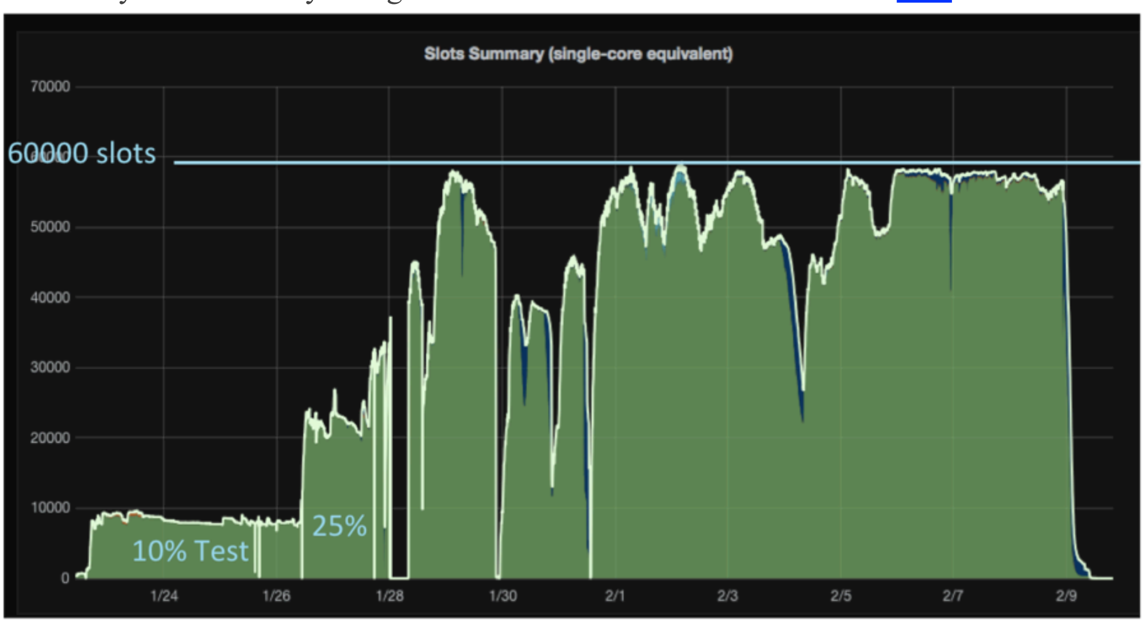

This plot shows the number of single core slots instantiated at Amazon over a period of just over two weeks, starting on Jan. 16. The slots have been obtained from Virtual Machines of different "types" (combinations of memory configuration, disk capacity and number of cores), instantiated at different Amazon Regions and Availability Zones. The figure shows the ability to sustain a plateau of 58,000 slots.

In March, the project team demonstrated that HEP Cloud is also a viable solution for the Intensity Frontier community. During the month, OPOS ran 3 productions activities for NOvA consisting of Monte Carlo and data event processing on HEP Cloud. The campaign contributed to the large computing "crunch" necessary to produce results for the Neutrino 2016 conference. It processed more than 90 TB of input data with 550,000 hours of computation, producing more than 150 TB of output. Data IO was handled efficiently by using S3, the highly scalable storage at Amazon. The team demonstrated that NOvA could sustain slot levels of 7,300 concurrent cores, a burst of almost 4 times over the slots allocated at Fermilab for NOvA. And thanks to the project's integration activities, NOvA is using the same familiar services they use for local computations, such as data handling and job submission.

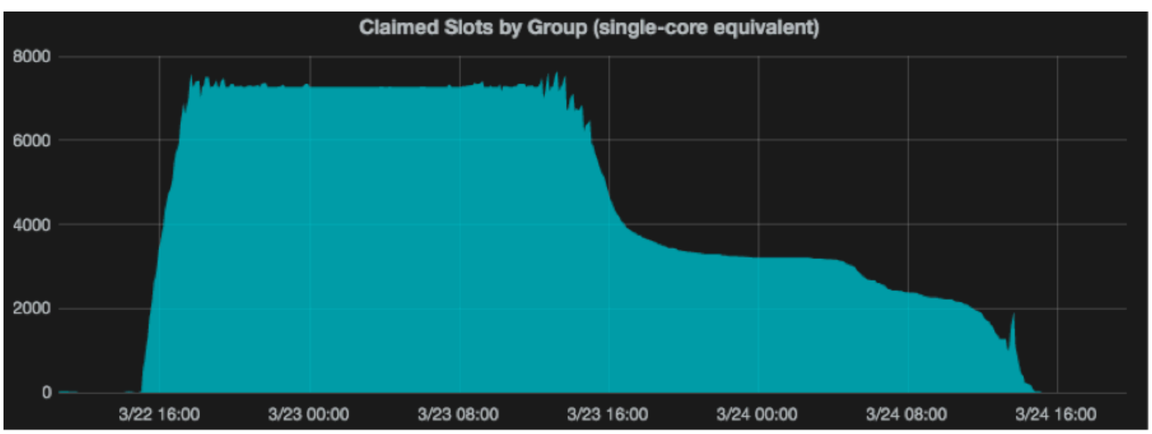

This plot shows the number of single core slots instantiated at Amazon for the NOvA experiment. Before the full number of jobs in the system (10,000) started to finish, the HEP Cloud sustained 7,300 cores for a day, almost a 4 times increase in the local resource *allocation for the experiment.*

A version of this article originally appeared in Feb 15, 2016 news.fnal.gov.

— Gabriele Garzoglio and Burt Holzman

DCAFI moving forward

The Distributed Computing Access with Federated Identities (DCAFI) project is moving forward on schedule and should be ready to start migrating the first experiment in June. For those of you who are unfamiliar with it, these are the motivations for the project:

- Dependency on Kerberos makes it difficult for non-Fermilab scientists to access our grid resources remotely, obstructing our lab´s goal of being an international laboratory.
- Fermilab's Kerberos Certificate Authority (KCA) server is losing its support starting September 2016, forcing us to find a replacement Certificate Authority for grid access.
- Asking users to manage their own certificates is a burden on them we avoided with KCA-based grid access, and we want to continue to avoid it.

Those motivations translated into these goals:

- Remove our dependency on Kerberos and KCA certificates.
- Continue to shield the users from the complexities of directly dealing with X.509 certificates.
- Integrate the grid authentication infrastructure with federated identities.
- In phase 1 we will use only Fermilab-based identities, but in phase 2 we will support other institutions. Phase 1 users need a Fermilab account but do not need to login to a Fermilab machine in order to submit jobs using the FIFE infrastructure.

From the end user's point of view, those who have Fermilab Kerberos credentials will be able to continue to submit jobs with no changes and no extra work for managing credentials. For those who want to submit from a remote site without Fermilab kerberos, the jobsub_submit command will tell them once a week when it is time to run a new "cigetcert" command to enter their Services password to get a new certificate. "cigetcert" will contact a Fermilab Identity Provider service to authorize the user, make use of the pre-existing CILogon Basic CA service to get a new certificate into /tmp, and place a longer-lived copy of the certificate on a MyProxy server where jobsub will be able to use it to get certificate proxies to run jobs.

For those more advanced users who had previously run kx509 to get a certificate, there will be a compatible replacement command provided that invokes cigetcert instead of contacting the KCA.

FIFE experiments will be transitioned to the new mechanism over the next few months in the following order:

Experiments	Transition Time
CDF, NUMI-X, GENIE:	6/6 – 6/20
g-2, ANNIE, CHIPS:	6/13 – 6/27
MINOS, Darkside:	6/20 – 7/4
DUNE:	6/27 – 7/11
Mu2e:	7/5 – 7/19
DES:	7/11 – 7/25
uboone:	7/18 – 8/1
Nova:	7/25 – 8/8
SBND:	8/1 – 8/15
Minerva:	8/15 – 8/29
Seaquest:	tentative 8/15 – 8/29

— Dave Dykstra

AFS transition

We will be turning off the Fermilab AFS servers in early May this year because the sort of worldwide file sharing once unique to AFS is now provided by the Web. Clickhere for details of the migration shutdown.

The /afs/fnal.gov file system has served Fermilab well since about 1992.

The primary services were:

- web server content
- Unix account login areas
- shared code and data areas for interactive use

The advantages of AFS included:

- Kerberos authentication for flexible network sharing
- Access Control Lists (ACL's) for good control of web content
- worldwide file sharing (/afs/cern.ch and such) thanks to good security.
- compatibility with many systems (AIX, IRIX, Linux, OSF1, SunOS, Solaris, Windows)

Disadvantages, as deployed at Fermilab, included:

- nonstandard file access permissions due to the strong authentication
- very limited capacity, with a maximum volume size of 50 GB.
- slow file access
- lack of support, as the market has moved on to other tools
- AFS drivers not always available for the latest kernels
- did not scale for grid use (removed from Fermigrid years ago)

Since deployment of newer NFS 4 file servers supporting ACL's and strong authentication, Fermilab has moved all central web services from AFS to the central NFS servers.

/publicweb for personal web pages, and /web/sites for group web pages.

Unix account login areas areas are being moved to /nashome. If you are affected, you have been informed!

We are presently verifying that the shared code and data areas are no longer in use.

For details of the migration, see http://computing.fnal.gov/nasan/afs-mig

–Arthur Kreymer

Experiment with most opportunistic hours Feb – March 2016

The experiment with the most opportunistic hours on OSG between Feb. 1, 2016 and March. 31, 2016 was Mu2e with 4,804,996 hours.

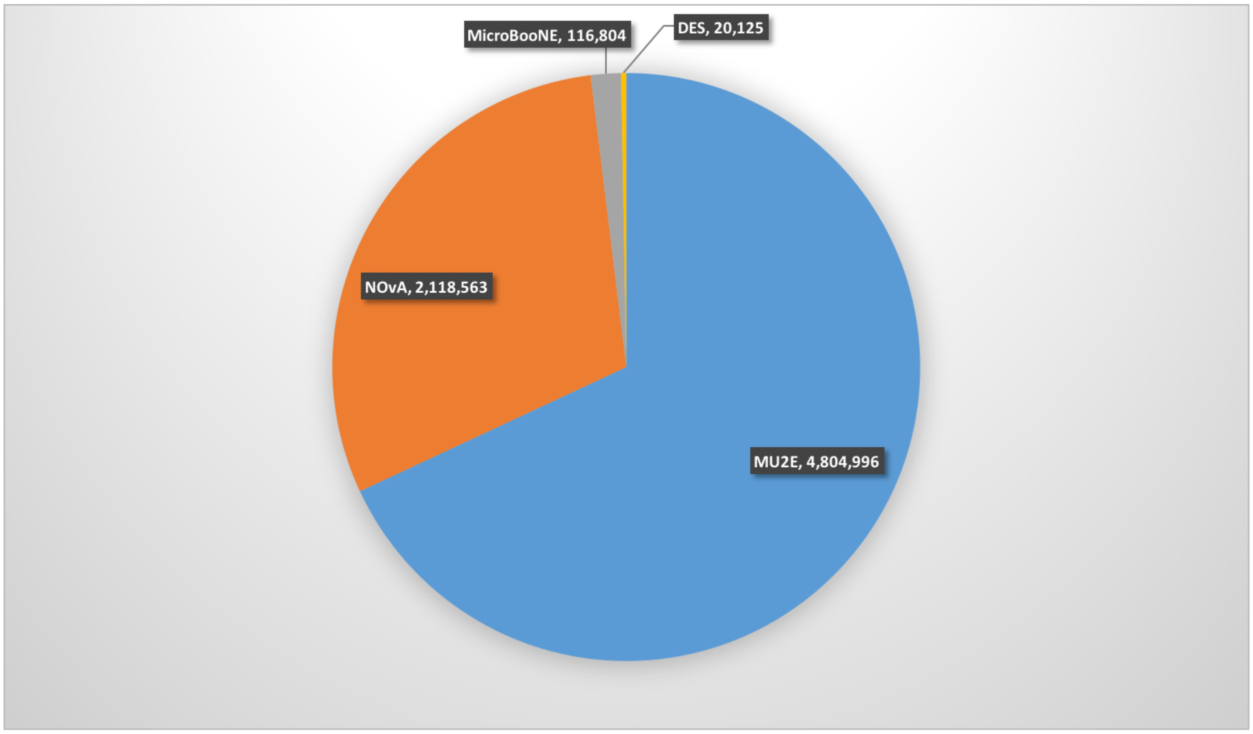

— Tanya Levshiva

Most efficient big non-production users – Feb – March 2016

Most efficient big non-production users on FermiGrid who used more than 100,000 hours since Feb. 1, 2015 was Willis K.Sakumoto with 100% efficiency.

Experiment	User	Wall Hours	Efficiency
CDF	Willis K. Sakumoto	593,455	100%
MARS	Vitaly Pronskikh	274,953	96.93%
MU2E	Anthony Palladino Jr.	489,684	92.74%
MINERvA	Benjamin Messerly	358,353	91.81%
MARS	Nikolai Mokhov	105,269	91.41%

— Tanya Levshina

Most efficient experiments Feb – March 2016

Most efficient experiments on FermiGrid that used more than 100,000 hours since February 1st, 2016 – CDF (100%) and

MU2E (85.75%)

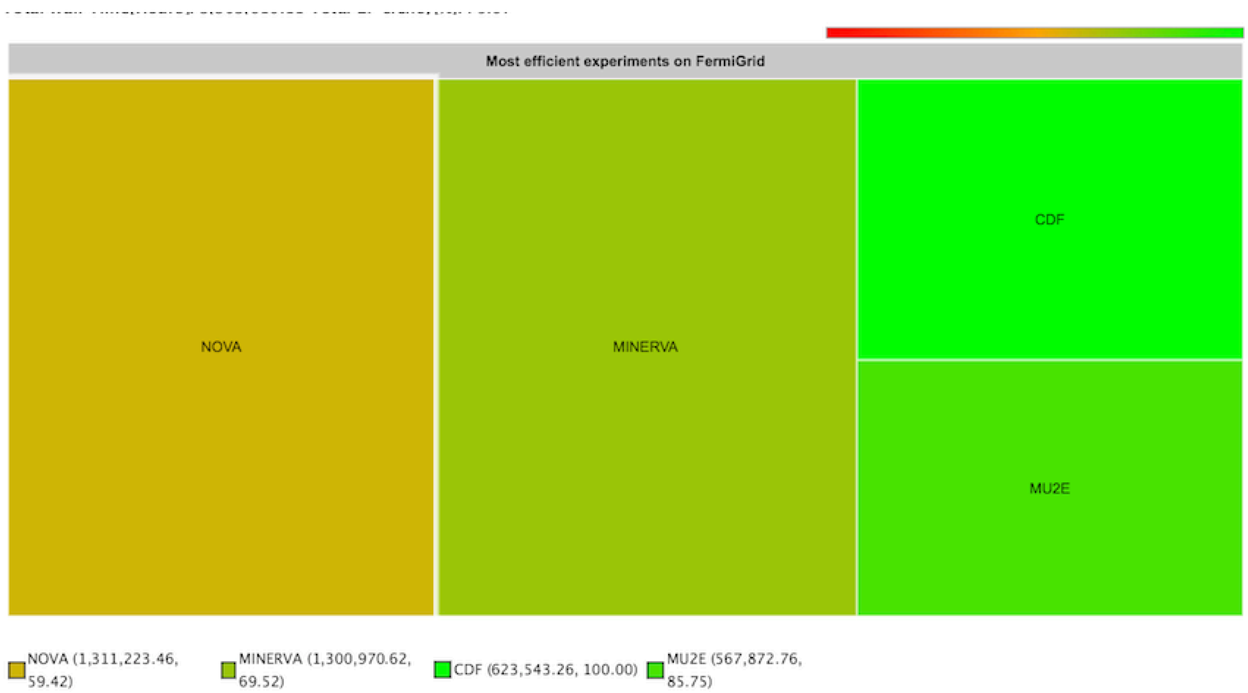

Va

— Tanya Levshi

Brought to you by

This newsletter is brought to you by:

- Dave Dykstra
- Gabriele Garzoglio
- Burt Holzman
- Bo Jayatilaka
- Mike Kirby
- Arthur Kreymer
- Katherine Lato
- Tanya Levshina

Feedback

To provide feedback on any of these articles, or the Fife notes in general, please email fife-support@fnal.gov.