FIFE Notes - February 2016 News

for

Distributed Computing at Fermilab

New sites for MicroBooNE

The MicroBooNE collaboration operates a 170 ton Liquid Argon Time Projection Chamber (LAr TPC) located on the Booster neutrino beam line at Fermilab. The experiment first started collecting neutrino data in October 2015. MicroBooNE measures low-energy neutrino cross sections and investigates the low-energy excess observed by the MiniBooNE experiment. The detector also serves as a next step in a phased program towards the construction of massive kiloton scale LAr TPC detectors for future long-baseline neutrino physics (DUNE) and is the first detector in the short-baseline neutrino program at Fermilab.

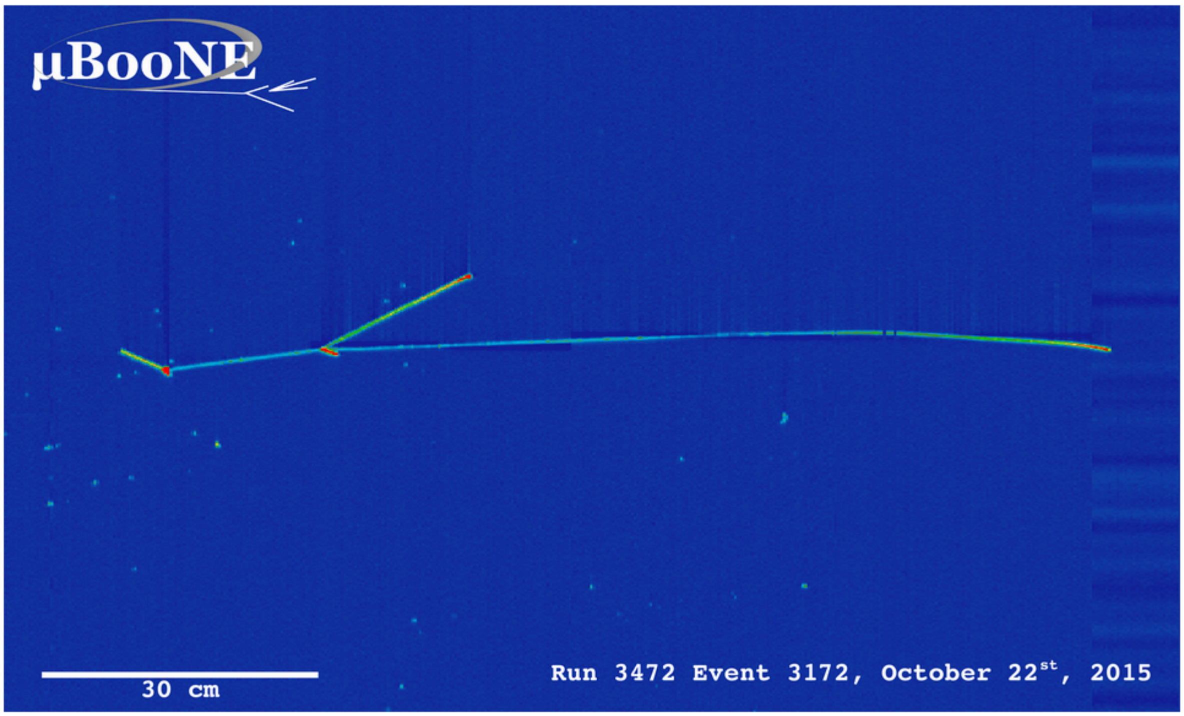

Photo from: http://www-microboone.fnal.gov/first-neutrinos/index.html

In past LAr TPC experiments, event selection was not fully automated. Candidate event selection required scientists and researchers to hand-scan the images depicting detector data. One of the critical steps in advancing LAr TPC based experiments is the development of pattern recognition algorithms and software to analyze the recorded data. In a November 2015 article titled 'A Neutrino in a Haystack: Brookhaven's Contributions to the MicroBooNE Neutrino

Experiment', Xin Qian said, “When you have an image, humans can actually process it pretty well, but to get a computer to achieve this level of pattern recognition is very difficult.” But with more than 30 million beam spills recorded to date, hand selection of events is untenable and fully automated event selection is needed to succeed.

Once algorithms are developed, they can be very CPU intensive. To accommodate this increased need for computing, MicroBooNE has been working to connect the computing at home institutions to the computing infrastructure located at Fermilab. With plans to incorporate computing from University of Manchester, University of Bern, Ohio Super Computing, and others, MicroBooNE will use these resources to reconstruct neutrino candidate events and increase the precision of theoretical predictions. These efforts have already paid dividends with a successful campaign at the University of Bern computing cluster to simulate an improved photon detector response after tuning with initial data. MicroBooNE is not limited to home institutions. It was the first Fermilab experiment to run applications at Clemson University’s OSG cluster through opportunistic access. The combination of Fermilab, university, and OSG opportunistic resources will play a critical role in the successful analysis of MicroBooNE’s data and the precision measurements of neutrino cross sections and searches for sterile neutrinos.

This is an excellent example of how combining computing clusters from numerous institutions can deliver the computational resources needed to make high precision measurements and drive future discoveries.

- Mike Kirby & Katherine Lato

Measuring the Universe -- one galaxy at a time --

The Dark Energy Survey (DES) is in its third year of gathering multi-colored digital images of large swaths of deep space with the 4m diameter mirror of the Blanco telescope on a mountaintop in Chile. The team of physicists, astrophysicists, engineers, computer professionals, technicians and managers who operate the survey are participants in a around-the-clock marathon consisting of a myriad of steps to bring the valuable sky data home and to get the images into shape for producing new insights into the nature of Dark Energy -- a phenomenon that is pushing distant galaxies within the visible universe apart from each other at ever accelerating rates contrary to our intuition of gravity as a force that pulls objects together.

The data arrive during the night from Chile in chunks called exposures, a digital package of 60 CCD images covering a 90-second snapshot of few square degrees on the sky. Sent on by scripts operated by team members at the National Center for Supercomputing Applications (NCSA) in Urbana, IL, the data are fed to grid nodes for initial processing which extends well into the daylight hours.

These fixed-sized exposures take a constant, predictable amount of time to 'detrend' (a processing step which recognizes defects and removes instrumental signatures from the digital CCD images). Detrending processing fits naturally into existing grid processing requirements (say, 2GB RAM and and 40 GB scratch disk per job running 6 hours with one CPU core). This stage of the processing can be run in a highly parallel fashion.

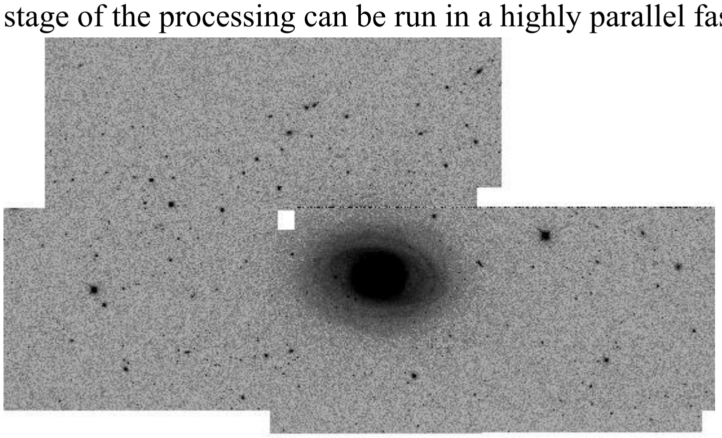

Figure 1 courtesy Nikolai Kuropatkin, Brian Nord, Dark Energy Survey

The first figure shows the results of processing a sample set of overlapping exposures through detrending. There will be several million individual rectangular regions processed by the end of the survey.

After several thousand exposures are gathered, they are grouped on an annual basis and sorted by position on the sky and by color. This is in preparation for the second stage of processing, where overlapping exposures are registered and combined into deep images of the sky, and where individual objects are carefully measured for shape, brightness and position. This stage is considerably more demanding of compute resources. These jobs, of which there are some 20,000 per year of observing, require up to 64 GB RAM, 2 TB of scratch disk and several linked cores, running occasionally for more than 24 hours/job -- although some regions of sky require significantly less than this -- the jobs vary in size.

Compute resources with such large requirements have not been widely available with multi-user grid-like common accessibility until very recently.

The FIFE group and Fermilab Computing division have played key roles in providing tools to make available efficient, flexible and widely distributed use of such dynamic systems. By building upon HT Condor-CE enhancements to provide access to dynamically configurable resources, GPgrid nodes can be reserved with these higher-capacity RAM, larger scratch disk, multi-core and extended run-time requirements as needed.

Figure 2, Courtesy Nikolai Kuropatkin, Brian Nord, Dark Energy Survey

The second figure shows a rendering of the combination of some 100 individual rectangular DES images (a superset of those in Figure 1) in blue, green and red filters to construct a coadded image of the barred spiral galaxy NGC 1398, which lies within the DES survey's footprint. This was done using a pipeline and high-capacity node similar to those now being developed for the DES coadd production campaign.

We are delighted to be able to use the configurable GPgrid resources for our DES production processing campaigns!

- Brian P. Yanny

OPOS helping MINERvA with offline production

MINERvA has been using offline production from OPOS for over a year. The following picture shows almost 200,000 jobs from MINERvA between June and December of 2015. (Note, DUNE's part of the graph only covers December since they started until OPOS on November 30, 2015.)

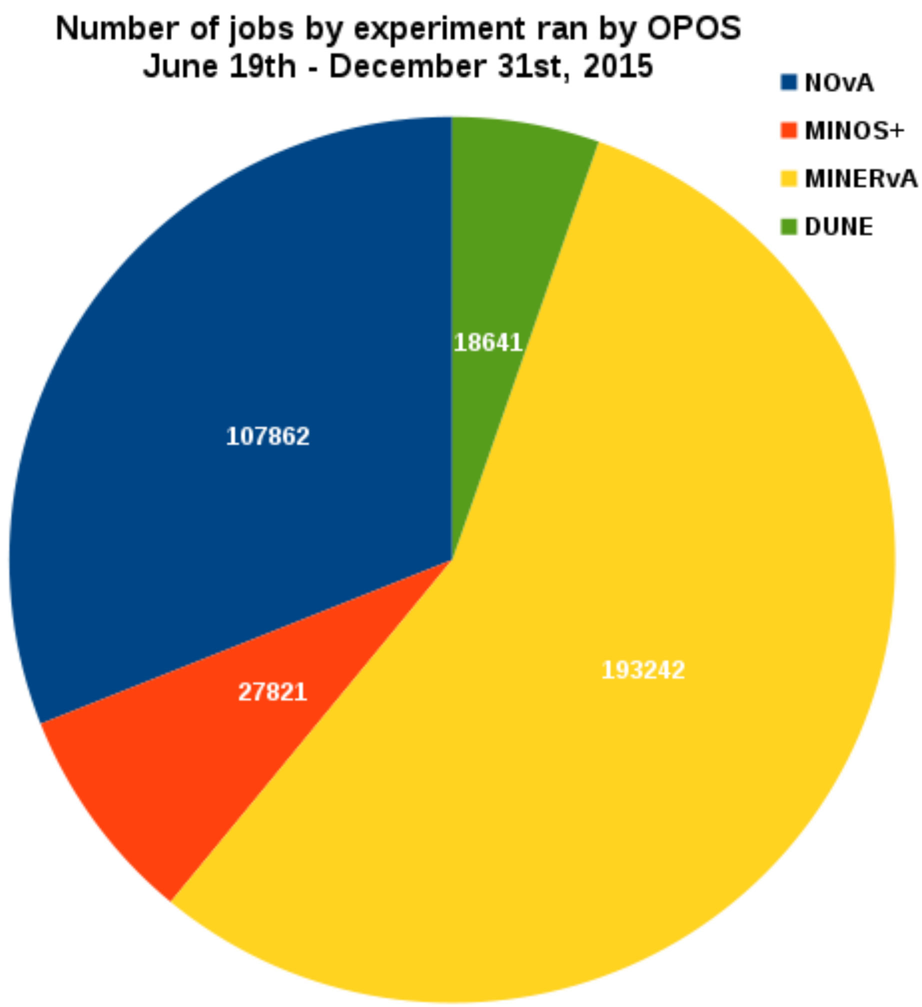

MINERvA started by using just monitoring services from OPOS, but now the team is running Monte Carlo (MC) simulation as well. It is anticipated that more services from OPOS will be used in the future.

The following quote is from the MINERvA spokespeople to the team. “OPOS takes on the task of submitting and monitoring grid jobs for MINERvA. This frees up the MINERvA personnel working on production, simulation and data analysis to spend time optimizing these processes to minimize resource use and maximize reliability. In addition, OPOS is able to identify common problems among experiments and to track statistics of analysis job performance and reliability. This data also proves essential for efficient processing and analysis of MINERvA data. Their efforts add up to more reliable and faster analysis of experimental data, which in turn allows MINERvA to continue to rapidly publish high quality physics papers."

One of the benefits of running production jobs for several experiments is that experience can be shared. For example:

- OPOS helped NOvA create procedures for running Keepup jobs. OPOS later helped MINERvA create their procedures using that work. Setting up these jobs was easier because of the experience of the OPOS team.
- OPOS helped MINOS change their workflow so they could use the sabweb tool. OPOS can transfer this successful work to MINERvA to enable them to begin to use samweb.

- Anna Mazzacane & Katherine Lato

Know before you go (on the OSG)

FIFE has a Wiki page to help match resource requirements and availability at remote sites

Previous editions of FIFE notes have shown glimpses of the tremendous computing resources available on the Open Science Grid. These resources come from a large number of remote sites, each of which has its own limitations and policies regarding opportunistic access. When users try to match their job requirements to sites, it can be a daunting task. In most situations, the FIFE Group recommends that users should not specify sites explicitly when submitting jobs to OSG locations in order to get the most resources. Instead, sit back, relax, and let the fifebatch system match your job to a free slot at an appropriate site. Occasionally, however, users want or need to send certain types of jobs only to a specific set of remote OSG sites.

The FIFE Group maintains a Wiki page containing information about each OSG site that supports FIFE experiments. For each site, the page lists the supported experiments, the proper site name to use for the jobsub --site option, as well as any known limitations on memory, job lifetime, or special details about the environment on the worker nodes. It is a good idea to consult this page before submitting your jobs to make sure that there are sites that can accept your resource requirements such as enough memory, scratch disk, run time limits, etc.

The Wiki page is not, however, a site status monitor. Just because a particular job's resource requests fit at a particular site, it does not guarantee that the job will start at that site. For example, there may not be any opportunistic slots available at that site at a given moment, or the opportunistic slots that are available may not match the job's requirements in terms of memory or job lifetime. Even so, this Wiki page can help distinguish between a situation where users simply need to wait a little bit for the job to start, or a situation in which the resource requests can never be met at the sites you requested. Many happy return codes on the OSG!

- Ken Herner

Coming soon to Fifemon: job resource monitoring

Wondering why your jobs have been put on hold? Want to better set your resource requests to make more efficient use of the grid (and to get your jobs starting faster)? This information, and more, will be available in the FIFE monitoring application, Fifemon, soon, and is already available for testing in pre-production (https://fifemon.fnal.gov/monitor-pp/).

Starting with your User Batch Details dashboard, you can see what jobs have been put on hold and why, as well as a complete listing of job clusters currently in the system. Included in this table are the maximum resources used with how much was requested. If a cell is highlighted in red, it means the job has exceeded its request and has been put on hold, requiring intervention to either decrease the amount the job uses or increase the request. Please contact FIFE Support through Service Now if you need assistance.

Job Clusters

Cluster	Submit Time	I	R	C	X	H	Memory (MB)	Disk (MB)	Time (hr)	Max CPU Time	Starts
4234003	2016-01-27T14:29:52.000Z	0	1	0	0	1	464 / 1900	1918 / 1024	24 / 7	23.69 hr	7
4260124	2016-01-29T14:32:47.000Z	0	2	0	0	0	1118 / 1900	559 / 5120	10 / 12	10.09 hr	7
4265871	2016-01-29T22:40:24.000Z	0	0	0	0	3	1523 / 1900	13 / 1024	0 / 8	0.00 hr	11
4289942	2016-02-02T09:10:46.000Z	0	2	0	0	0	8 / 1900	0 / 5120	2 / 20	0.00 hr	1

Held Jobs

Last 5 minutes

jobid	hold_date	HoldReasonCode	HoldReasonSubcode	HoldReason
4265871.3610@fifebatch1.fnal.gov	2016-02-01 19:10:29	26	4	SYSTEM_PERIODIC_HOLD Starts/limit 11/10
4265871.3613@fifebatch1.fnal.gov	2016-02-01 16:32:52	26	4	SYSTEM_PERIODIC_HOLD Starts/limit 11/10
4265871.3590@fifebatch1.fnal.gov	2016-02-01 16:25:57	26	4	SYSTEM_PERIODIC_HOLD Starts/limit 11/10
4234003.9402@fifebatch1.fnal.gov	2016-02-01 09:36:46	26	2	SYSTEM_PERIODIC_HOLD Disk/limit 2000000/1.572864000000000E+06

https://fifemon.fnal.gov/monitor-pp/dashboard/db/user-batch-details?var-user=mu2epro

Click on a cluster number, and you'll be taken to the Job Cluster Summary dashboard showing a variety of information about that job cluster: request parameters and resources, number of processes in each state, resources used by running and completed processes, and a timeline of Condor events.

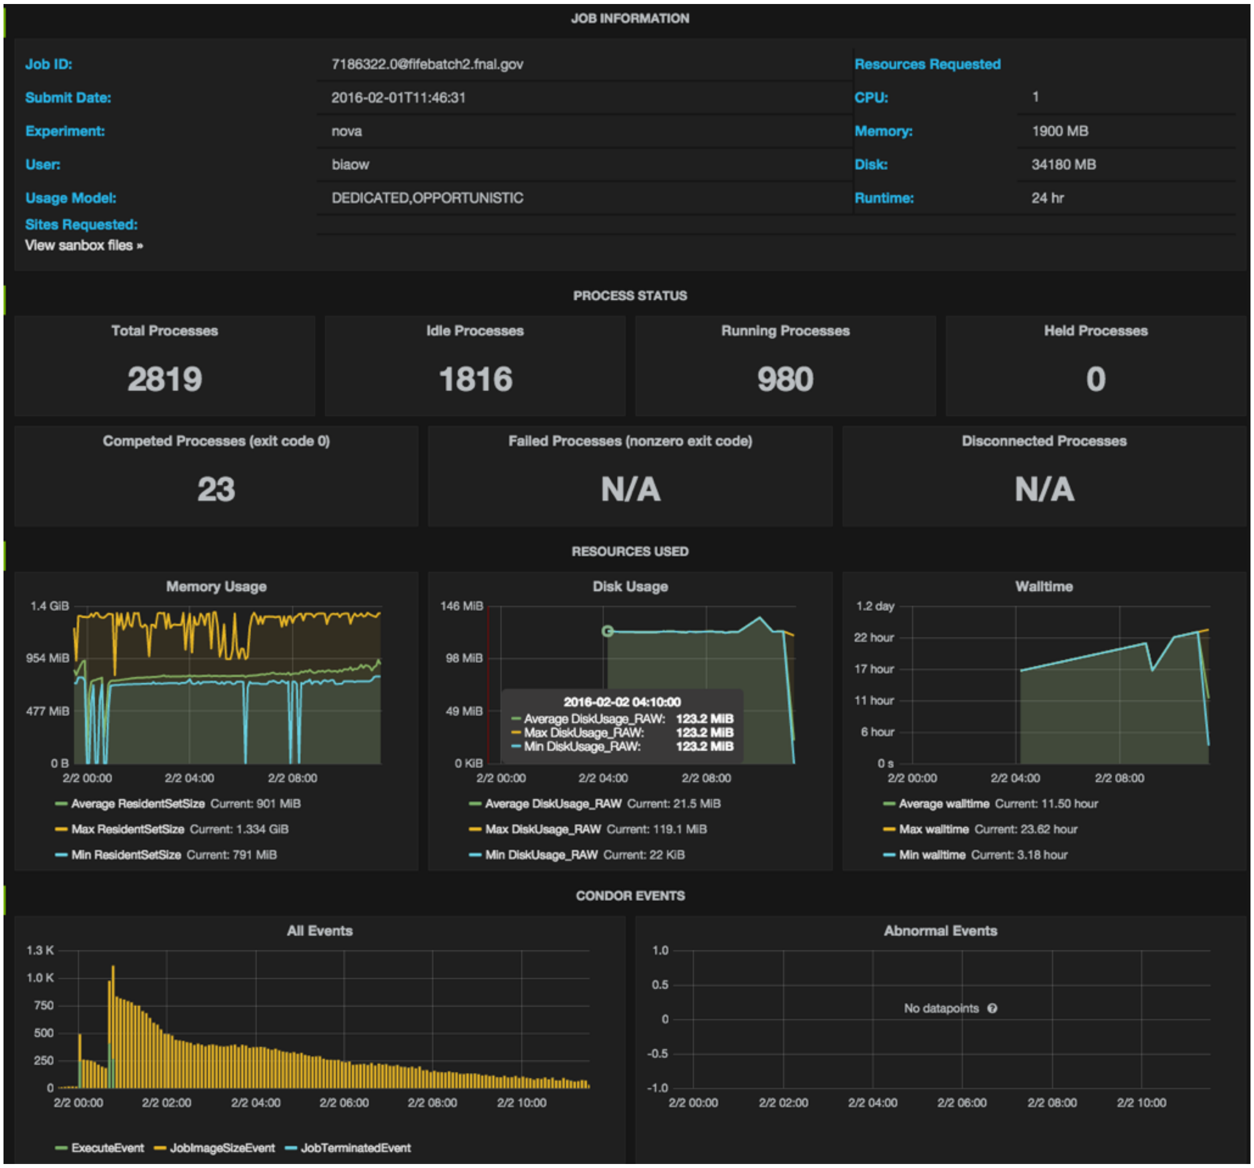

https://fifemon.fnal.gov/monitor-pp/dashboard/db/job-cluster-summary?var-cluster=7186322&from=1454391283934&to=1454434483934

We invite you to test these new features in pre-production and to provide feedback. We hope you find this information useful for tracking your job progress and for better understanding the resources your jobs are using. As a general rule of thumb, the higher your resource requests, the longer it will take your jobs to start running

- Kevin Retzke

Optimizing job submissions

Carefully tailoring your resource requests will increase your job throughput

With partitionable slots now the norm on GPGrid, it's important to have a good understanding of resource requirements. The memory, disk, and maximum runtime available in free job slots changes as users submit jobs. As a result, there may be free slots that have less resources available to them than the defaults, as they are leftovers from the way the cluster was partitioned at a given moment. There's nothing inherently wrong with these free slots, and any job that can fit into those requirements can run without problem.

When you submit jobs via jobsub_submit, the default memory, disk, and job runtime requests are 2 GB, 35 GB, and 23 hours 40 minutes, respectively. Unless you override any or all of these options with the –memory, --disk, and –expected-lifetime options, you'll get the defaults. But if the default is bigger than you need, your job can't take advantage of the smaller-than-usual free resources that we just described.

As an analogy, imagine it's lunchtime and you walk into a restaurant. If you ask for a table (the default), you may well be seated at a table for four. That is fine when the restaurant isn't busy. But if they're busy (lots of red dots) and you insist on a table, you may have to wait.

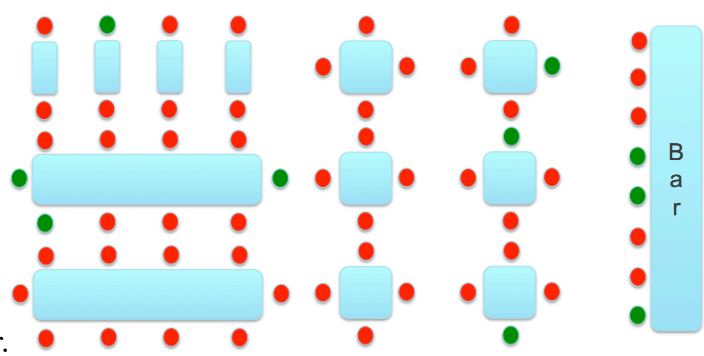

Unless you're willing to sit at the bar.

The same holds true in a computing environment. If you have a workflow that consistently uses fewer resources than the default requests, it makes sense to lower your resource request so that you can take advantage of small-sized slots that are unusable by the default request, just like sitting at the bar. By doing so your jobs will start faster.

In addition to the slots that have a smaller than usual amount of memory and/or disk assigned to them, there can be free resources available both on GPGrid and OSG that don't satisfy the default 24-hour run time request. The vast majority of jobs submitted by the FIFE experiments take less than twelve hours. By lowering your expected runtime via the --expected-lifetime option to more closely match your actual runtime, you can again get access to resources that aren't available to jobs submitted with the default request. If you do have longer jobs that take closer to 24 hours or more, a way to still utilize resources with shorter allowed run times might be to divide a large job into smaller pieces that can run at the same time. You could then lower your lifetime request to fit within resources available.

Consider the analogy of eight people going out to eat who don't care if they all sit together or not. Asking for a table for eight when the restaurant is busy may result in a long wait.

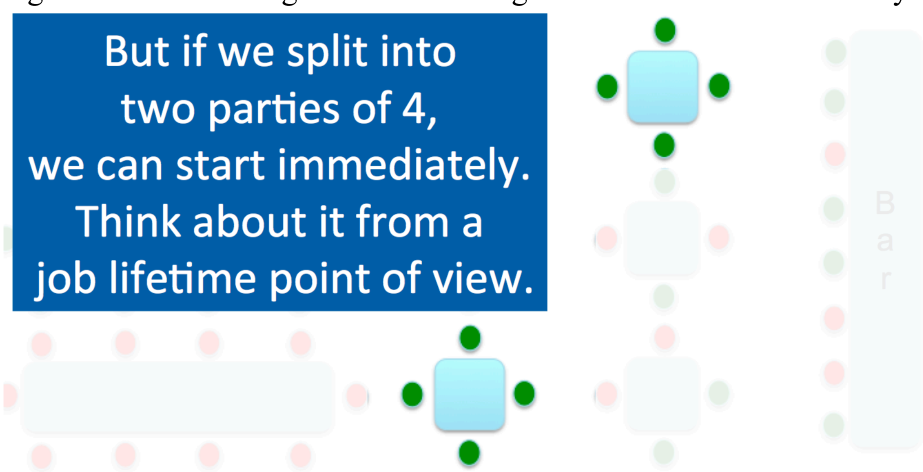

Suppose you have a 8-core 16 GB glidein running on a worker node with 23 hours of time remaining for jobs. It is running one job that has requested one CPU and 2 GB of memory. We could fit another 7 such jobs into this glidein, but jobs that request the default lifetime (24h) can't run on this machine, because the glidein only has 23 hours of time remaining. This can lead to significant overall inefficiency, as this very thing could be playing out on multiple worker nodes at the same time.

Please contact the FIFE group at **fife-group@fnal.gov** if you need help determining what resources your job requires.

You shouldn't ask for **less** than you need, but tailoring your request to match your resource needs as closely as possible will result in higher throughput for you and for your collaborators, as it will make the cluster as efficient as possible. Faster job completion enables faster analysis completion, faster publication, and more time to enjoy your lunch.

- Ken Herner & Katherine Lato

Recent Open Science Grid milestones

The Open Science Grid (OSG) has recently achieved a number of milestones and continues to provide distributed grid computing resources to scientists around the world at a record scale. 2015 marked the first year since the OSG's inception a decade earlier that over 1 billion computational hours were consumed by OSG users. Additionally, November 2015 marked the first calendar month where OSG usage exceeded 100 million computational hours putting the OSG on track to breaking last year's records in 2016. While the LHC experiments of ATLAS and CMS continue to be the cornerstones of both usage and provided resources on the OSG, FIFE experiments and Fermilab resources played a larger role on the OSG than ever before. In particular, mu2e was the 4th largest Virtual Organization (VO) to use the OSG in 2015 with 49 million computational hours logged. A feat that's particularly impressive given that over 35 million of those hours were harvested opportunistically across the OSG and that prior to March, the Mu2e experiment had never used OSG resources outside of Fermilab.

Fermilab continues to be one of the largest resource providers on the OSG. In 2015, Fermigrid was the third largest provider of CPU to the OSG with approximately 100 million hours delivered.

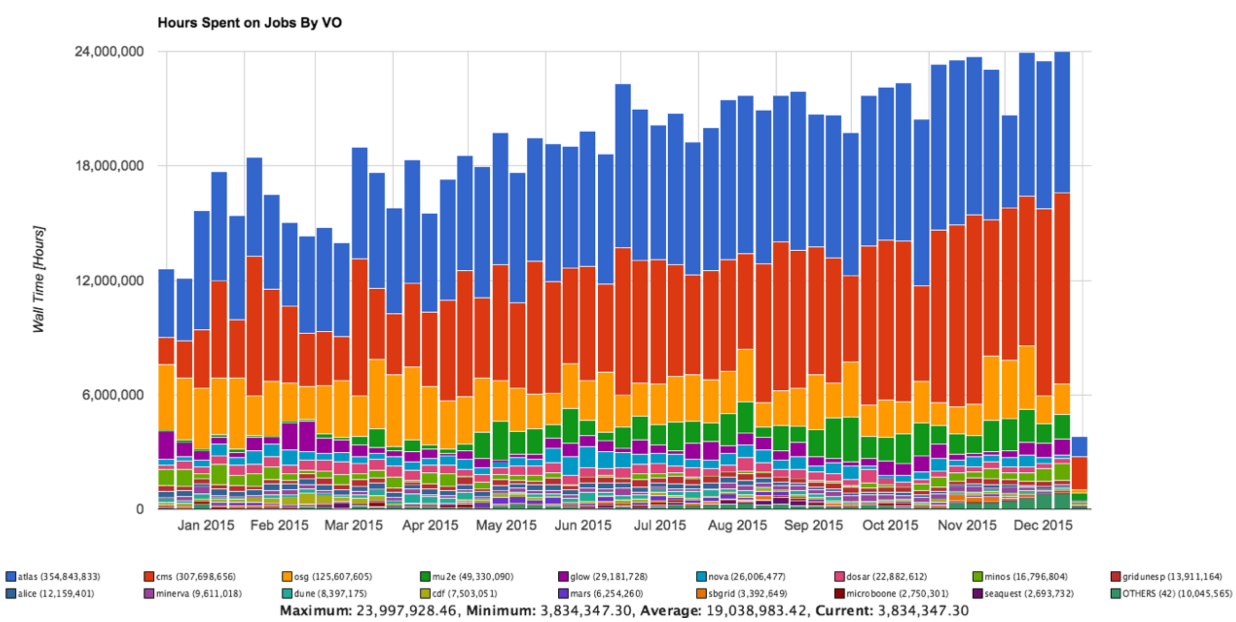

Wall hours summed to 1,009,066,075 hours from week 00 of 2015 to week 52 of 2015.

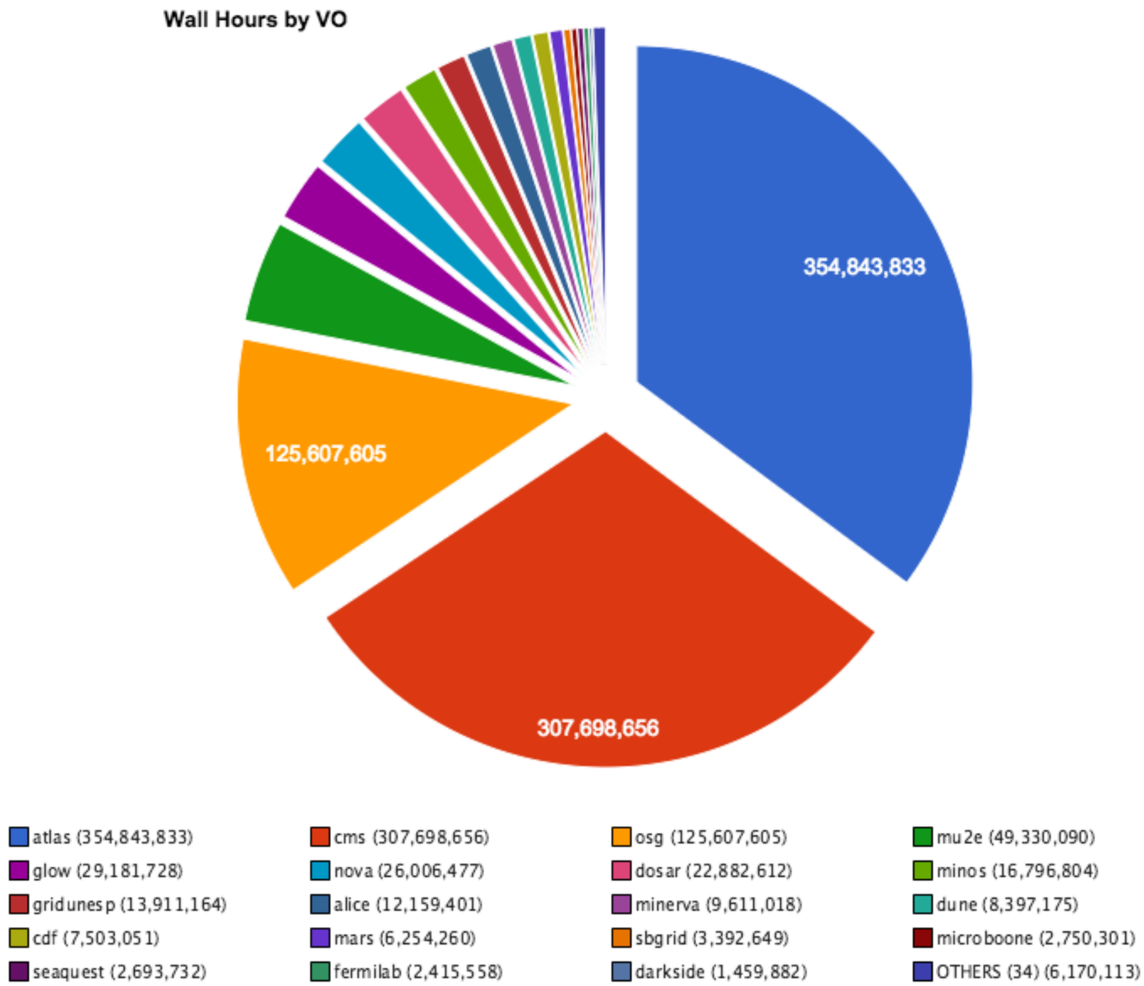

- Bo Jayatilaka & Tanya Levshina

Most efficient experiment

Most Efficient Experiment on FermiGrid that used more than 500,000 hours since December 1st - Minos (98.72%) and MINERvA (85.80%)

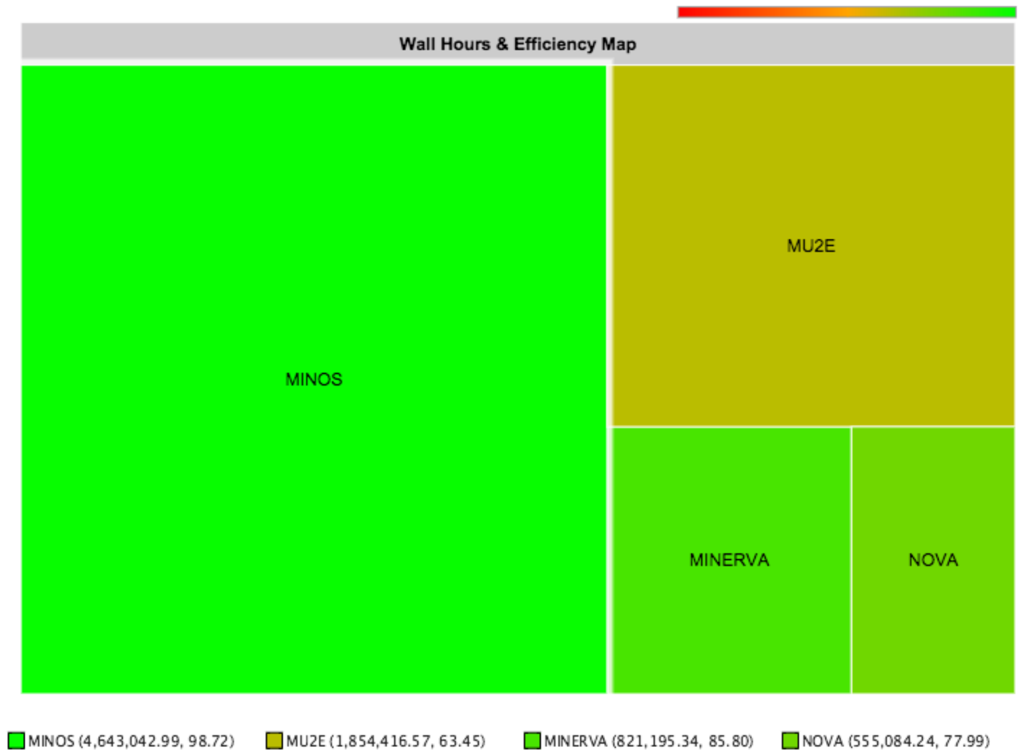

- Tanya Levshina

Most efficient big non-production user

Most efficient big, but not production, users on FermiGrid who used more than 100,000 hours since December 1st was Luri A. Oksuzian from Minos with 98.9% efficiency.

Experiment	User	Wall Hours	Efficiency
Mu2e	Luri A. Oksuzian	102,314	98.9%
Minos	Adam J. Aurisano	4,643,042	98.7%
Mu2e	Anthony Palladino	999,170	95.9%

- Tanya Levshina

Experiment with the most opportunistic hours

The experiment with the most opportunistic hours on OSG between December 1st and January 31 was Mu2e with 13,960,877.

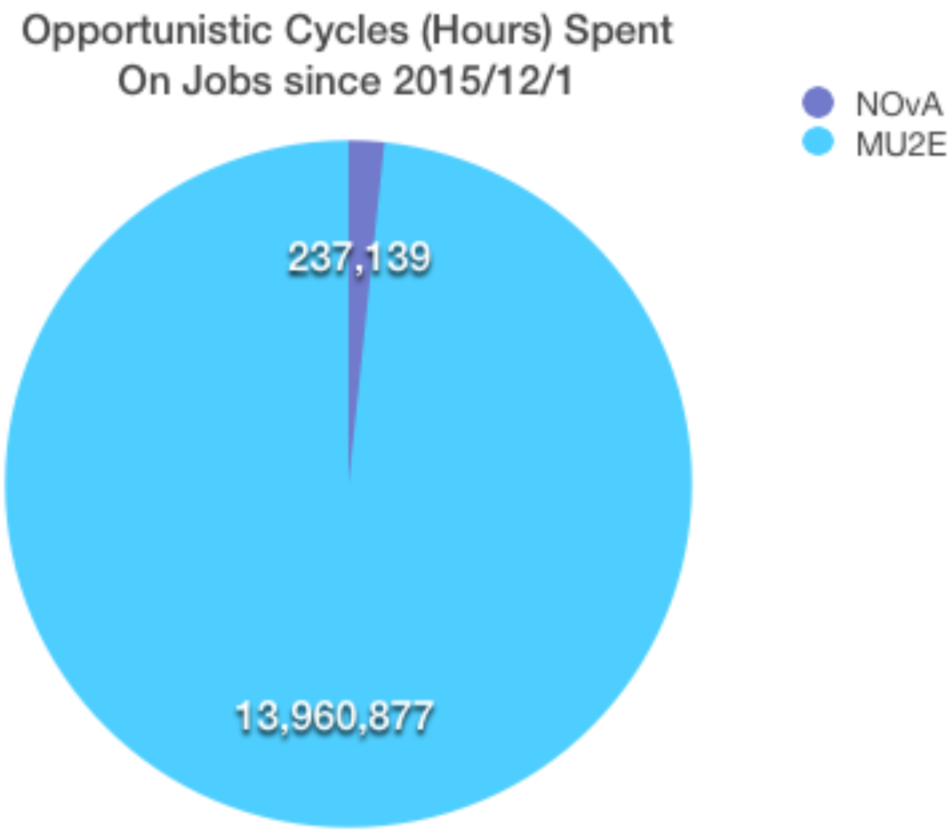

- Tanya Levshina

Brought to you by

This newsletter is brought to you by:

- Ken Herner
- Bo Jayatilaka
- Mike Kirby
- Katherine Lato
- Tanya Levshina
- Anna Mazzacane
- Kevin Retzke
- Brian P. Yanny

Feedback

We welcome articles you might want to submit. Please email fife-group@fnal.gov.