



---

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

---

# Running on the Open Science Grid- Best Practices

Ken Herner

2015 FIFE Workshop

1 June 2015

# Introduction and assumptions

---

- Once you're running on Fermigrid, running on other Open Science Grid (OSG) sites is the next logical step
- Resources are plentiful, but some conveniences will not be there
- To save time I will make the following assumptions in this talk:
  - You have accounts with your experiment and can log in to your experiment's interactive VMs
  - You are familiar with using `jobsub_client` to submit jobs
  - You have heard of CVMFS, UPS, BlueArc, dCache, gridftp, Virtual Organizations
- Please consult the FIFE Tutorials or other talks in the workshop for more details on those points


# Motivation for running offsite

---

- Besides getting a lot more CPU hours?
- Fermigrid will not be able to meet projected future demand of all experiments
- Also provides a way to keep running if Fermigrid is down

I shouldn't be so skeptical but I thought maybe the schedd's would freak out today when all the local worker nodes went away. It's currently happily running all the jobs that have requested offsite resources. Too bad more people haven't included OFFSITE in their list.

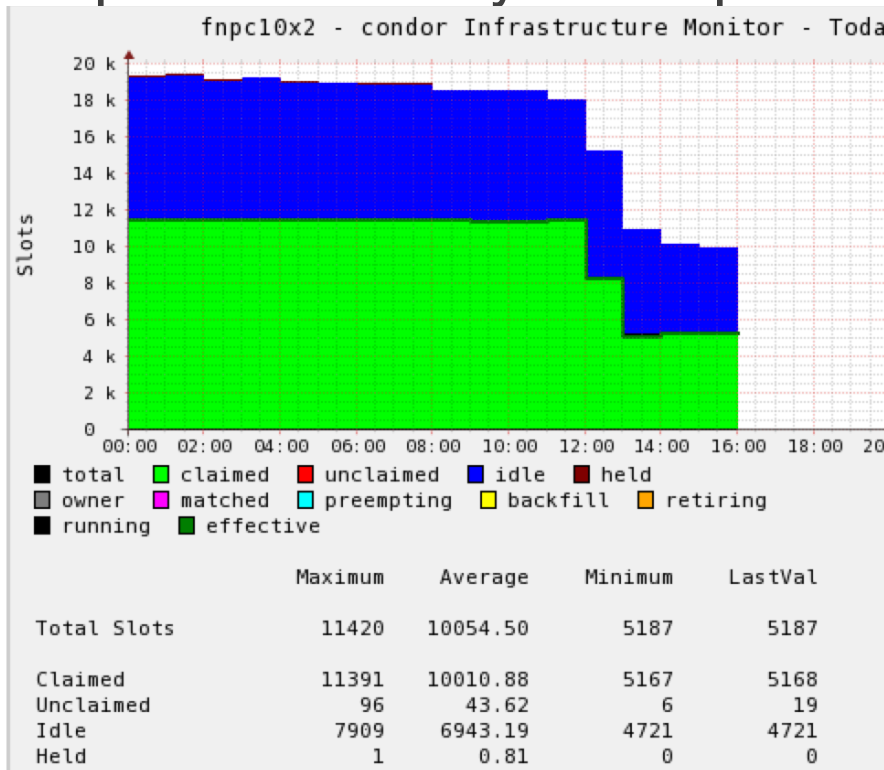
```
[root@fifebatch1 ~]# condor_status -format "%s\n" glidein_site | sort | uniq -c
  32 BNL
 265 Caltech
  118 FZU
  215 MIT
  149 MWT2
 106 Michigan
   32 Omaha
    4 SMU
  245 SU-OG
   16 UChicago
   19 Wisconsin
```



Fermigrid outage on 8 Apr

# Motivation for running offsite


- Besides getting a lot more CPU hours?
- Fermigrid will not be able to meet projected future demand of all experiments
- Also provides a way to keep running if Fermigrid is down



Two days ago....

## What is available

- About 130M opportunistic hours for OSG sites last year (16% of all OSG hrs.)
  - That is NOT full capacity
  - 150M in past 12 months; on pace for 180M hours in 2015
- Let's do that math: 8760 hours per year means equivalent of 14,000 CPUs available continuously-- larger than Fermigrid
- Dozens of sites out there; 18 that support at least one FIFE expt. right now (most support all of them)



Total opp. hours  
in 2014 per VO

# How do I run on the OSG?

---

- If you are using `jobsub_client` this is trivial:
  - set `--resource-provides=usage_model=OFFSITE`
  - optionally specify the site(s) that you want to use with the `--site=siteA,siteB,siteC` option if you wish (no `--site` option gives you all available sites)
- No need for anything special with proxies, certs, etc.
- If you want things to run “anywhere”, do `--resource-provides=usage_model=DEDICATED,OPPORTUNISTIC,OFFSITE`
  - If you do this and only want to run at certain OSG sites, put Fermigrid as one of the possible sites
- Note that “OPPORTUNISTIC” above means opportunistic *Fermigrid* slots, not OSG. By definition almost everything on the OSG is opportunistic anyway

## What should I expect on OSG worker nodes

---

- Standard **SL5/6** environment (most places SL6 now)
  - requested in the usual way with the --OS jobsub option
- **NO BLUEARC**. Any reads/writes there **MUST USE IFDH**
  - e.g. doing /experiment/app/mypath/my\_executable in your script will not work
  - offsite bluearc reads/writes go through BestMan server, 1 Gb/s shared by all experiments. Use dCache whenever possible!
- All oasis.opensciencegrid.org/expt and expt.opensciencegrid.org **CVMFS** repositories will be available; can set up all UPS products and run code from there
- Some sites do not install various X11 libraries on the systems. Check if your application expects them. Solutions include adding them in CVMFS, copying at runtime, compiling without them (statically linking them)

# My job isn't starting!!!

---

- Some questions to ask before pushing the panic button:
  - How long has it been? **Several hours on busy sites is not unusual!!!**
  - Do my jobsub options make sense? Did I put contradictory options (eg DEDICATED, OPPORTUNISTIC and --site=MIT?)
  - Do my resource requests make sense for the site? By default jobsub asks for 2 GB RAM and 35 GB disk. Some sites cannot match that (especially the disk part. Use the --memory and --disk jobsub options.)
  - Can my experiment run at the requested site(s)?
- If that all seems fine, check FIFEMON to see if others from your experiment are running at any of your sites
- Could also be that the fifebatch server has hit its max number of running jobs (onsite and offsite count for that)
- If that all looks fine, open a ticket and include the job ID number



# My job isn't starting!!!

---

- Some questions to ask before pushing the panic button:
  - How long has it been? **Several hours on busy sites is not unusual!!!**
  - Do my jobsub options make sense? Did I put contradictory options (eg DEDICATED, OPPORTUNISTIC and --site=MIT?)
  - Do my resource requests make sense for the site? By default jobsub asks for 2 GB RAM and 35 GB disk. Some sites cannot match that (especially the disk part. Use the --memory and --disk jobsub options.)
  - Can my experiment run at the requested site(s)?
- If that all seems fine, check FIFEMON to see if others from your experiment are running at any of your sites
- Could also be that the fifebatch server has hit its max number of running jobs (onsite and offsite count for that)
- If that all looks fine, open a ticket and include the job ID number

# My job isn't starting!!!

---

- Some questions to ask before pushing the panic button:
  - How long has it been? **Several hours on busy sites is not unusual!!!**
  - Do my jobsub options make sense? Did I put contradictory options (eg DEDICATED, OPPORTUNISTIC and --site=MIT?)
  - Do my resource requests make sense for the site? By default jobsub asks for 2 GB RAM and 35 GB disk. Some sites cannot match that (especially the disk part. Use the --memory and --disk jobsub options.)
  - Can my experiment run at the requested site(s)?
- If that all seems fine, check FIFEMON to see if others from your experiment are running at any of your sites
- Could also be that the fifebatch server has hit its max number of running jobs (onsite and offsite count for that)
- If that all looks fine, open a ticket and include the job ID number

# My job isn't starting!!!

---

- Some questions to ask before pushing the panic button:
  - How long has it been? **Several hours on busy sites is not unusual!!!**
  - Do my jobsub options make sense? Did I put contradictory options (eg DEDICATED, OPPORTUNISTIC and --site=MIT?)
  - Do my resource requests make sense for the site? By default jobsub asks for 2 GB RAM and 35 GB disk. Some sites cannot match that (especially the disk part. Use the --memory and --disk jobsub options.)
  - Can my experiment run at the requested site(s)?
- If that all seems fine, check FIFEMON to see if others from your experiment are running at any of your sites
- Could also be that the fifebatch server has hit its max number of running jobs (onsite and offsite count for that)
- If that all looks fine, open a ticket and include the job ID number

# My job isn't starting!!!

---

- Some questions to ask before pushing the panic button:
  - How long has it been? **Several hours on busy sites is not unusual!!!**
  - Do my jobsub options make sense? Did I put contradictory options (eg DEDICATED, OPPORTUNISTIC and --site=MIT?)
  - Do my resource requests make sense for the site? By default jobsub asks for 2 GB RAM and 35 GB disk. Some sites cannot match that (especially the disk part. Use the --memory and --disk jobsub options.)
  - **Can my experiment run at the requested site(s)?**
- If that all seems fine, check FIFEMON to see if others from your experiment are running at any of your sites
- Could also be that the fifebatch server has hit its max number of running jobs (onsite and offsite count for that)
- If that all looks fine, open a ticket and include the job ID number

# My job isn't starting!!!

---

- Some questions to ask before pushing the panic button:
  - How long has it been? **Several hours on busy sites is not unusual!!!**
  - Do my jobsub options make sense? Did I put contradictory options (eg DEDICATED, OPPORTUNISTIC and --site=MIT?)
  - Do my resource requests make sense for the site? By default jobsub asks for 2 GB RAM and 35 GB disk. Some sites cannot match that (especially the disk part. Use the --memory and --disk jobsub options.)
  - Can my experiment run at the requested site(s)?
- If that all seems fine, check FIFEMON to see if others from your experiment are running at any of your sites
- Could also be that the fifebatch server has hit its max number of running jobs (onsite and offsite count for that)
- If that all looks fine, open a ticket and include the job ID number

# My job isn't starting!!!

---

- Some questions to ask before pushing the panic button:
  - How long has it been? **Several hours on busy sites is not unusual!!!**
  - Do my jobsub options make sense? Did I put contradictory options (eg DEDICATED, OPPORTUNISTIC and --site=MIT?)
  - Do my resource requests make sense for the site? By default jobsub asks for 2 GB RAM and 35 GB disk. Some sites cannot match that (especially the disk part. Use the --memory and --disk jobsub options.)
  - Can my experiment run at the requested site(s)?
- If that all seems fine, check FIFEMON to see if others from your experiment are running at any of your sites
- Could also be that the fifebatch server has hit its max number of running jobs (onsite and offsite count for that)
- If that all looks fine, open a ticket and include the job ID number

# My job isn't starting!!!

---

- Some questions to ask before pushing the panic button:
  - How long has it been? **Several hours on busy sites is not unusual!!!**
  - Do my jobsub options make sense? Did I put contradictory options (eg
  - [unclear] jobsub
  - [unclear] at
  - [unclear] (e.g. [unclear] options.)
  - [unclear] ur
- If the exper
- Could also be that the fifebatch server has hit its max number of running jobs (onsite and offsite count for that)
- If that all looks fine, open a ticket and include the job ID number

**Remember:**  
**We have no way of boosting your priority or getting you a larger allocation at remote sites**

# How do I know where I can run?

- Current as of today (names you should put in --site option)

- **All Fermilab VO experiments:** CMS T1, not Fermigrid

- BNL, Caltech, Cornell, FNAL, Fermigrid, Michigan, MIT, MWT2, Nebraska, Omaha, SU-OG, TTU, UChicago, UCSD, Wisconsin

- **NOvA only:** FZU, Harvard, OSC, SMU

- **μboone only:** WT2

- **LBNE/DUNE (standalone VO):**

- BNL, Fermigrid, MWT2, Nebraska, Omaha, UChicago

- **Additional sites in the works**

- For most up to date information you can do some basic

condor\_status command:

MYVO=Fermilab, LBNE, DES, etc.

- `condor_status -any -pool gfactory-1.t2.ucsd.edu -constraint 'glideinmytype=="glidefactory"' -format "%s\n" glidein_site -constraint 'stringlistmember("MYVO",GLIDEIN_Supported_VO)' | sort | uniq`



# Known site issues/features

---

- Small batch slot sizes
  - Jobs only have 9 GB scratch space at Syracuse (SU-OG) and 20 GB at FZU (Czech Republic)
  - Use the `--disk` option
- OS installations
  - Harvard and SMU are SL5 only right now
  - Jobs with `--OS=SL6` will not run there
- Linux 3.x kernels
  - MWT2/UChicago have some 3.x custom kernel worker nodes
  - Can cause problems with UPS setups
  - Need to set `UPS_OVERRIDE` environment variable

# Best Practices

---

- File I/O
  - Avoid BlueArc as much as possible in favor of dCache
  - Avoid excessive file copies in/out of input files, put small common files read by many/all jobs in CVMFS
  - Larger files (e.g. flux files) should not go in CVMFS. OSG StashCache may be a good option for you
- Job configuration (minimizing pre-emption chances)
  - Perhaps aim for shorter jobs and be sure to request min needed resources. Aim for a few hours' run time, not 12 for example
  - If worried about pre-emption and SAM, focus on workflows that don't need to copy files in, e.g. simulation
- Protect yourself against Black Holes (not the LHC kind)
  - See next slide

# Pre-emption

---

- Pre-emption: a higher-priority job comes in and kicks your job out (also called eviction)
- Fermigrid does not pre-empt any FNAL experiments
  - Other sites have their own policies; *we have no control over them*
- Pre-emption can and does happen on OSG sites
  - Jobs will eventually restart elsewhere
  - **This should not discourage you!!!**
- Each site has its own way of figuring out what jobs to pre-empt when high-priority jobs come in
  - Generally, jobs with lower resource requests and shorter running times fare better
- What can you do?
  - Lower your resource reqs. as much as possible; adjust jobs to aim for a few hours' running time instead of 8-12+ (e.g. run fewer evts.)

# Black Holes

- Black Hole: a bad worker node where jobs fail immediately
  - Can suck up hundreds of jobs before you know there's a problem
  - Cannot be anticipated; we have no way to “blacklist” such nodes at remote sites
  - GOC tickets should be opened with remote sites to fix them
- Some things can mitigate the problem somewhat



MU2E Success Rate on the OSG Resources (2015/05/13 - 2015/05/14).

SITE	TOTAL JOBS	FAILED JOBS	SUCCESS RATE	BAD HOST	EXIT CODE	FAILED JOBS ON HOST
UCSDT2	1360	0	100.0			
AGLT2	1	0	100.0			
CIT_CMS_T2	5443	0	100.0			
Nebraska	9339	7553	19.1			
				red-d19n20.unl.edu	1	7553
USCMS-FNAL-WC1	698	0	100.0			
Crane	9	0	100.0			
MIT_CMS	822	13	98.4			
				t2bat0212.cmsaf.mit.edu	2	13
MWT2	170	0	100.0			
SU-OG	999	6	99.4			
				LAW-D-TS440-S8-its-c6-osg-20150121	2	6
GLOW	1	0	100.0			
Total	18842	7572	59.8			

## Black Holes (2)

---

- Now have a script in place that checks for some cvmfs repos on the worker node; will not start job if they're not found
  - Will add other repos and system checks going forward
- The script will stop some problems *but not all*; you may want additional checks (look for other repos, versions of system libs, etc.) at the beginning of your scripts
  - Suggest putting in a sleep function before exiting if one of your checks fails-- prevents other jobs from hitting the same slot and dying for the same reason.
    - Amount of time up to you; suggest 15- 30 min
    - Will take longer to spot problems and you will lose some jobs, but losing 10 is better than losing 1000
- Other easy ways to help debugging
  - Print the hostname and OS (uname -a) at the beginning of each job
  - Carefully check exit status of different commands in your script; print error messages before exiting

# Testing the waters

---

- New fermicloud node available for testing workflows before submitting large batches of OSG jobs
- **fermicloud168**- configured with CVMFS but no bluearc (like an Amazon or OSG worker)
- login should be available to everyone, so try it out
- Copy your script here and run small test job interactively-- see if it works as you expect



# Summary

---

- Running on the OSG can give you a huge resource boost
- Trivial to get there via jobsub
- A few things to consider when designing workflows (biggest is no direct BlueArc reading)
- Best practices include using dCache wherever possible, aiming for shorter job length, requesting minimal resources, and adding checks to the beginning of your scripts for required resources, and sleep statements to minimize damage from black holes
- Bottom line: one experiment got **4 million hours in one week**-- and that was not using all sites.
  - That's **24,000 CPUs** running continuously. What's your experiment's Fermigrid allocation again?
- Feel free to **ask us** any questions about running offsite! We want to help you succeed!

