



# Production Support


Ken Herner

OSG Staff Retreat

7 November 2018


# Review of the last 12 months

▼ By VO



# Review of the last 12 months

WallHoursByVO



	total
cms	751 Mil
atlas	655 Mil
osg	80.3 Mil
dosar	48.3 Mil
glow	26.0 Mil
fermilab	15.5 Mil
belle	11.77 Mil
hcc	8.45 Mil
ligo	7.25 Mil
gridunesp	5.94 Mil
icecube	3.06 Mil
ciconnect	2.703 Mil
N/A	2.300 Mil
nova	2.118 Mil
gluex	2.086 Mil
nanohub	656 K
OSG-Connect	505 K

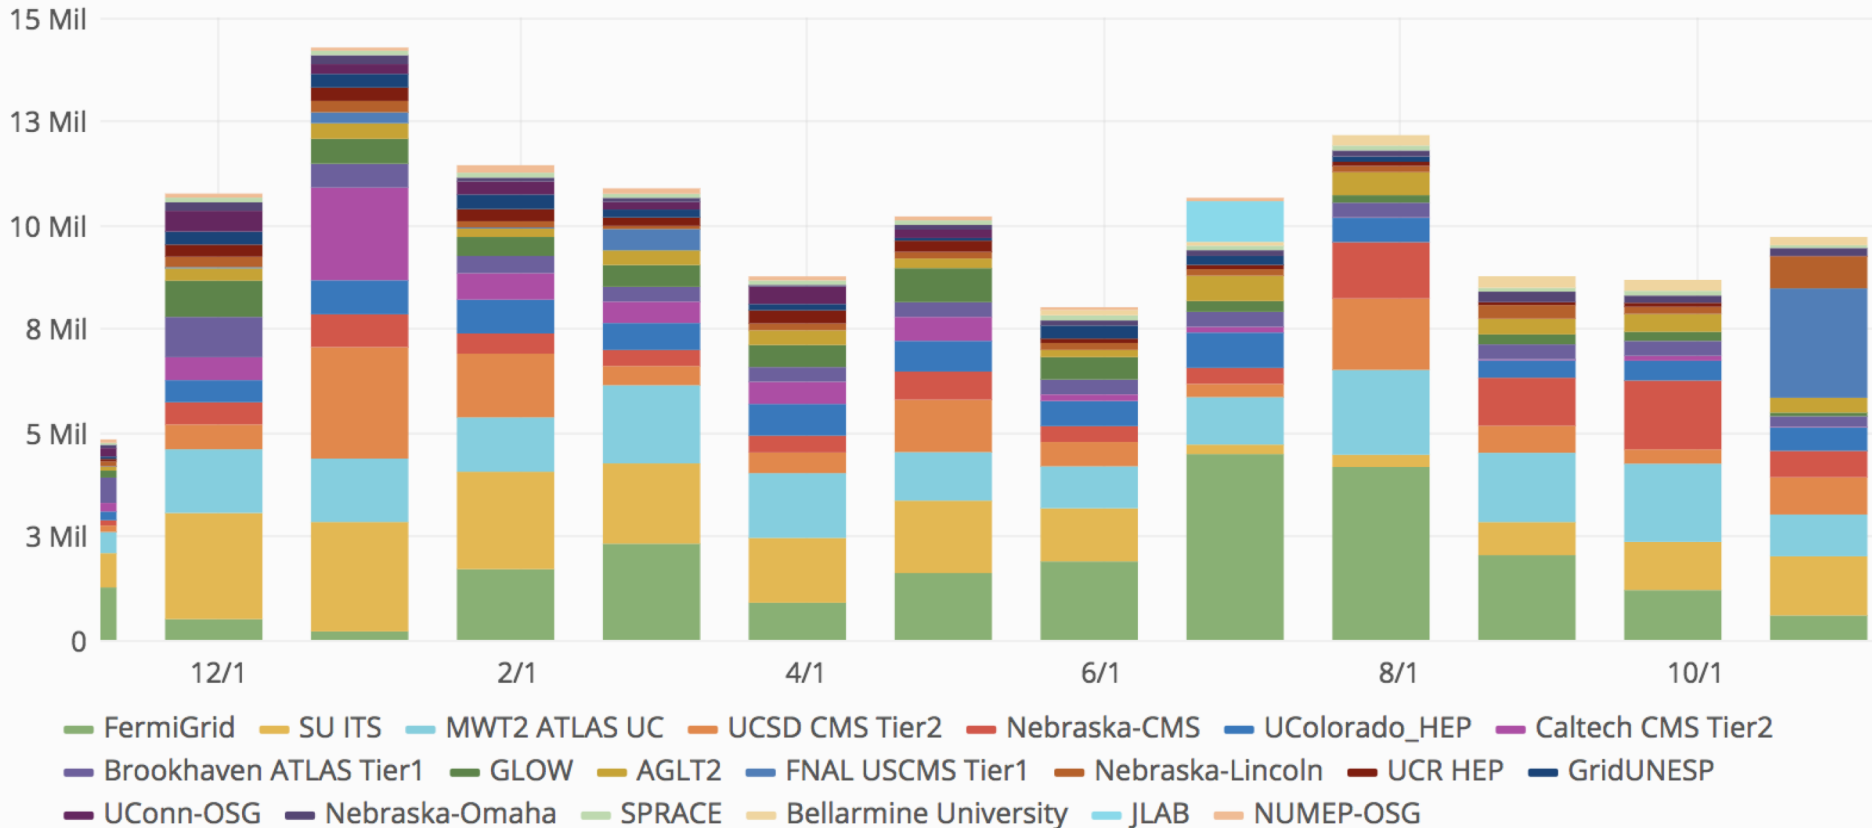
**Unofficial Total:  
1.62 B pilot hours**

Something of an under-report:  
FIFE hours at EGI sits not all included

# Opportunistic Hours past 12 months

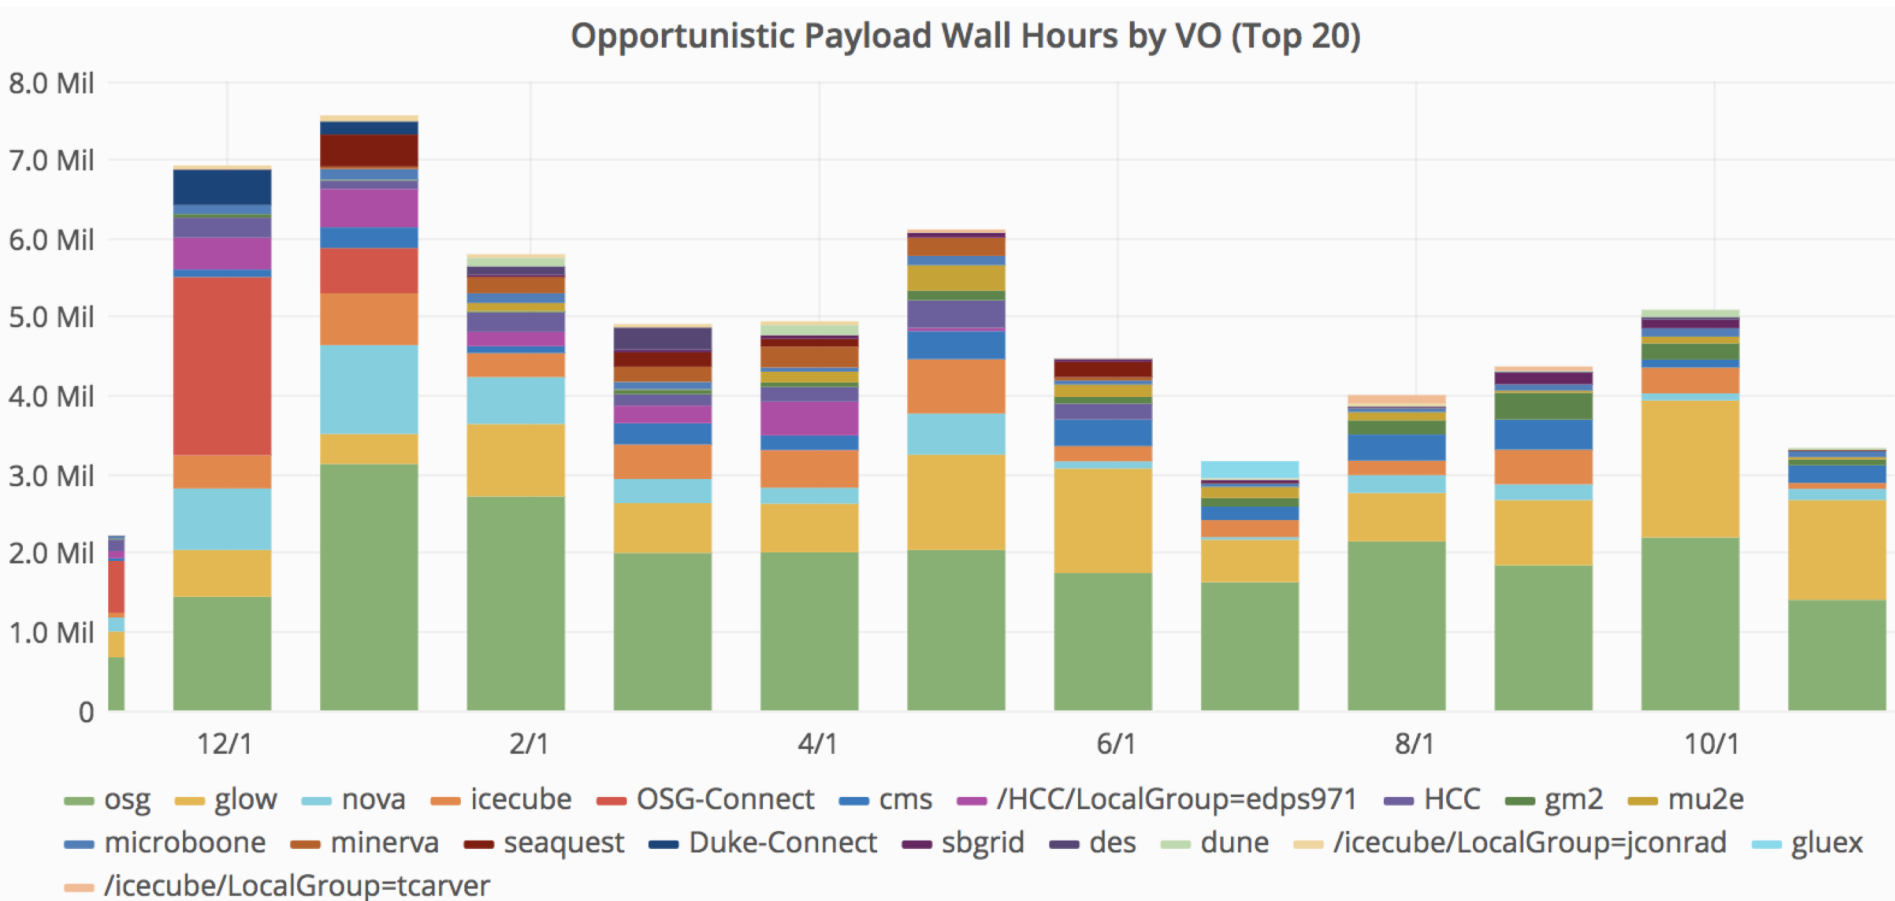
- Good summer availability at FermiGrid

Opportunistic Pilot Wall Hours by Site (Top 20)



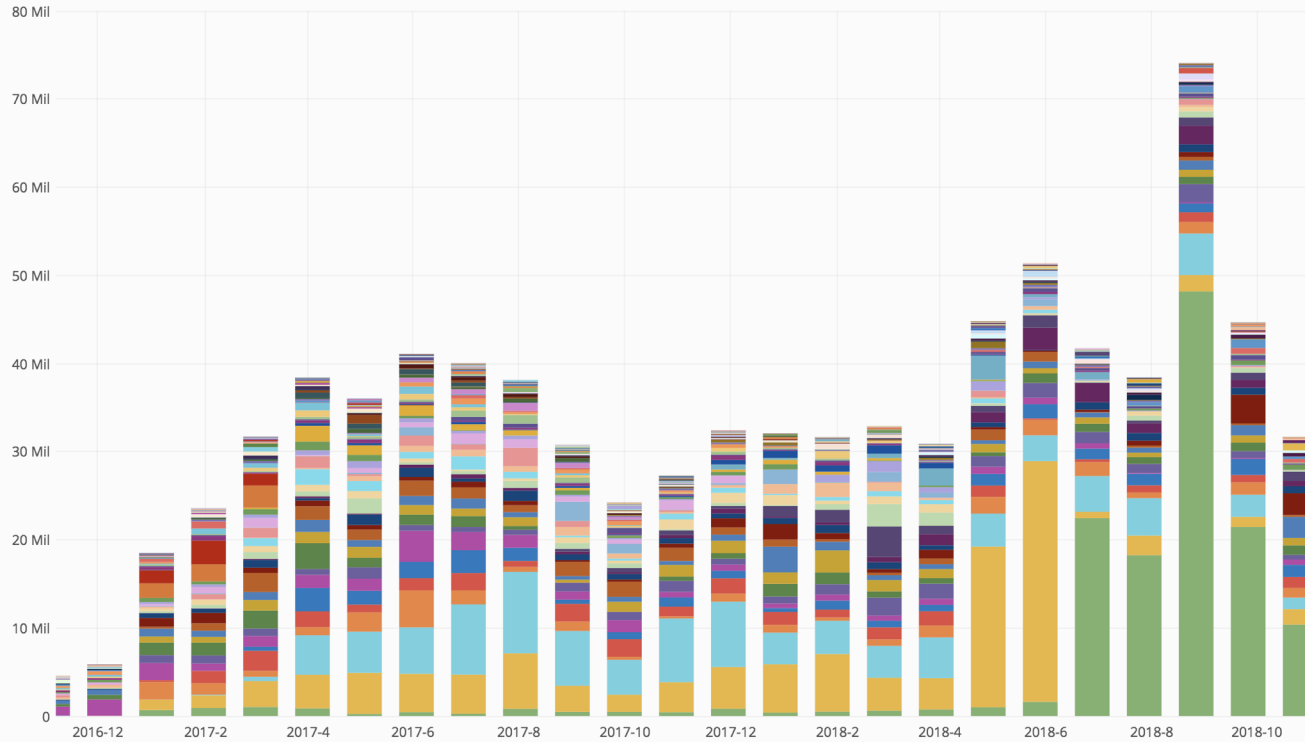
# Opportunistic Hours past 12 months

- Good summer availability at FermiGrid



# Projects last 2 years

Wall Hours per Project




	avg	total
cms	5.3 Mil	133.1 Mil
nova	4.6 Mil	114.7 Mil
IceCube	3.9 Mil	96.3 Mil
mu2e	1.2 Mil	29.0 Mil
IBN130001-Plus	1.1 Mil	28.5 Mil
GLOW	1.1 Mil	26.3 Mil
Duke-QGP	1.0 Mil	26.0 Mil
minerva	985 K	24.6 Mil
SPLINTER	958 K	23.9 Mil
microboone	937 K	23.4 Mil
TG-IBN130001	896 K	22.4 Mil
minos	833 K	20.8 Mil
dune	754 K	18.9 Mil
VERITAS	651 K	16.3 Mil
gm2	637 K	15.9 Mil
cms.org.cern	560 K	14.0 Mil
xenon1t	557 K	13.9 Mil
mars	477 K	11.9 Mil
UNLbcnf	444 K	11.1 Mil
BioGraph	411 K	10.3 Mil
LIGO	407 K	10.2 Mil
seaquest	370 K	9.3 Mil
TG-PHY150040	350 K	8.8 Mil
des	327 K	8.2 Mil
SourceCoding	307 K	7.7 Mil
macsSwigmodels	292 K	7.3 Mil
cms.org.fnal	287 K	7.2 Mil
icecube	245 K	6.1 Mil
glow	223 K	5.6 Mil


# Review of Goals for Year 6

- Weather the storm
  - Spring and summer 2018 will be very busy. IF experiments gearing up for Neutrino 2018 (July), LHC for Moriond (March) and ICHEP (July), NP for Quark Matter 2018 (May), protoDUNE coming online late July.
  - A **LOT** of mouths to feed over the next 12 months
  - Extremely important to keep growing the pie... user training part of this too (especially with certain experiments). Also, make sure ALL resources are accessible (containers, etc.)


# Review of Goals for Year 6

- Weather the storm 
  - Spring and summer 2018 will be very busy. IF experiments gearing up for Neutrino 2018 (July), LHC for Moriond (March) and ICHEP (July), NP for Quark Matter 2018 (May), protoDUNE coming online late July.
  - A **LOT** of mouths to feed over the next 12 months
  - Extremely important to keep growing the pie... user training part of this too (especially with certain experiments). Also, make sure ALL resources are accessible (containers, etc.)

Did get some help from NERSC...



One million cores!





# Review Goals for Year 6 (2)

- CPU resources
  - Continue to increase opportunistic availability across new and existing resources (continue to onboard new sites) and make sure all VOs can use as many sites as possible (Fermilab, LSST, GlueX, etc.)
  - Work with VOs to commission HPC resources (cross-pollination is key here! See CMS and MINOS+ at Stampede)
    - Several potential tech issues (CVMFS, etc.) here. We (I) need to get more VOs sharing information. It's better for *everyone* if we can make this an easy process!
- GPU resources
  - Continue to push this forward and try to make a (at least somewhat) standard prescription for accessing the resources. Common, simple approaches go a long way.

# Review Goals for Year 6 (2)

- CPU resources



- Continue to increase opportunistic availability across new and existing resources (continue to onboard new sites) and make sure all VOs can use as many sites as possible (Fermilab, LSST, GlueX, etc.)
- Work with VOs to commission HPC resources (cross-pollination is key here! See CMS and MINOS+ at Stampede)
  - Several potential tech issues (CVMFS, etc.) here. We (I) need to get more VOs sharing information. It's better for *everyone* if we can make this an easy process!

- GPU resources




- Continue to push this forward and try to make a (at least somewhat) standard prescription for accessing the resources. Common, simple approaches go a long way.

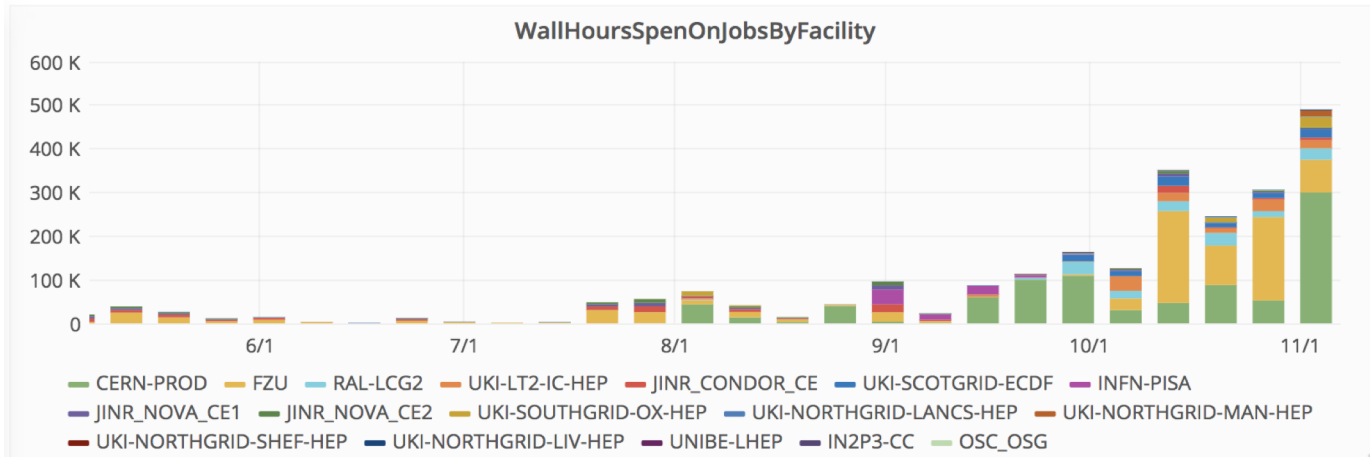
# Y6 Additional Sites

- Added a few new sites for OSG (Prod Support not the primary driver in most cases). International coordination is the way of the future...

By VO



By Facility




# Review of Goals for Year 6 (3)

- Software and Infrastructure

- Support VOs in using as much of the standard stack as possible.

- Containers, containers, containers 

- Monitor SW developments (especially as S2I2 evolves) and push for modular toolkits and middleware (Fermilab's ifdhc data handling client is a good example- under the hood data transfer protocols can be transparently changed) 

- HEPCloud

- Help onboard VOs as needed (likely Fermilab and CMS first)

HEPCloud Authorization To Operate granted!



# Review of Goals for Year 6 (3)

- Software and Infrastructure
  - Support VOs in using as much of the standard stack as possible.
  - Containers, containers, containers
  - Monitor SW developments (especially as S2I2 evolves) and push for modular toolkits and middleware (Fermilab's ifdhc data handling client is a good example- under the hood data transfer protocols can be transparently changed)
- HEP
  - HEP (first)  
**As far as software... Rucio seems to be converging as the standard tool for HEP data management (LHC, DUNE, ...)**

# Year 7 and Beyond

# Containers, containers, containers

- *This slide from June 2018 FIFE Roadmap talk*
- 18,693 jobs
- 5 users, 3 of whom were testing
- **Feedback requested:**
  - No demand?
  - Submission too hard?
  - Resources inadequate?



*Feedback: robust container construction too hard; lack way to seamlessly submit*


# Software Needs

- User analysis often has custom/modified code modules; not practical to store every possible iteration of code in CVMFS
  - How can users get custom code to jobs without overwhelming the experiment's/lab's SEs?
- FIFE stopgap solution: replicate user tarballs 20x on various dCache pools (works well but far from ideal)
- **Longer-term solution: rapid-turnaround CVMFS (FNAL: D. Dykstra)**
  - Publishing only allowed by job submission server (according to user-supplied tarball or directory); compares hash to avoid duplicate work
  - Updates every ~5 minutes
  - Don't let jobs start until CVMFS update finished



# Some Discussion Points

- Where does LSST fall in the support scheme? What is the overall OSG engagement strategy? They're doing mock data challenges and such already. 2021 is pretty soon...
- Don't forget about DESC...



## Some discussion points (2)


- What is the best model for container use in larger VOs? What is a sensible support model? Per experiment? Per user?
  - **Seems a logical place where all levels of user support can learn from and help each other!**
  - Especially high container interest for GPU resources, and HPC too.
- DOE labs and their experiments will be tightly coupled with HEPCloud. How can we cross-pollenate?
- Authentication/security: seems another place where everyone can/should work together. Potentially some land mines here wrt to larger VOs and their inherent visibility. They have to toe the ~~company~~ Department line very closely...
  - Close collaboration possible, but must proceed carefully.

# Some discussion points (3)

- Support for new compute resources pledged to one or more Large VOs (DUNE has been seeing this lately)
  - Currently I open a factory ops ticket, do tests, repeat as needed
  - Works very well for FIFE experiments; OK for Operations?

- **HPC/LCF**

- *“THOU SHALT UTILIZE LEADERSHIP COMPUTING”*
- Mid and large VOs will initially lead here
- What’s the best way for VOs to help each other navigate the new landscape? HEPCloud will somewhat abstract this away, but underlying structure affects many
- Some people talk already of course; is there more to do here?



# A concern

- Heraclitus: “The only constant is change.”
- Goes for the familiar, too (Linux, Python). SL6 is EOL in March 2020.
  - Personal view: second half of 2019 seems like a good time for lots of larger VOs to make the move. Containers reduce the urgency, but people are still better off changing IMO.
  - Longer-term: IBM-RedHat future
    - I imagine most of HEP will end up doing the same thing. IRIS-HEP will have lots to say
- Python 3 transition (<https://python3statement.org/>) will be lots of fun, too!
  - What, if any, is the OSG support staff’s role here (at all levels)?



# Backup

# HPC Successes

- CMS and MINOS+ successful on Stampede with startup XSEDE allocations (also CMS on Comet)
  - glidewms modifications to pass project/allocation IDs were key
- CMS recently successful at NERSC
- ATLAS has done well on Titan (90M CPU hours)
- Mu2e now testing at NERSC (via HEPCloud); other FIFE expts have expressed interest
- CVMFS tends to be a common stumbling block (Solved with Shifter at NERSC, for example)
- Are we collaborating as much as we can here?

## ATLAS at NERSC



CPU consumption Good Jobs in seconds (Sum: 45,529,858,476)

