A Virtual Comet

HTCondor Week 2017

May 3 2017

Edgar Fajardo
On behalf of OSG Software and Technology

Working in Comet

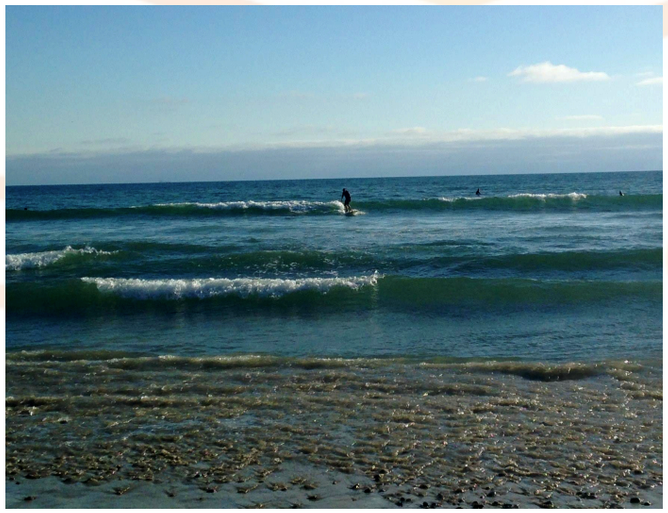

What my friends think I do

What Instagram thinks I do

What I think I do

What my boss thinks I do

SDSC SAN DIEGO SUPERCOMPUTER CENTER

UC San Diego

What is Comet?

"HPC for the long tail of Science"

Where is Comet?

- Comet is located at the San Diego Super Computing Center in La Jolla, California. On the UCSD Campus.

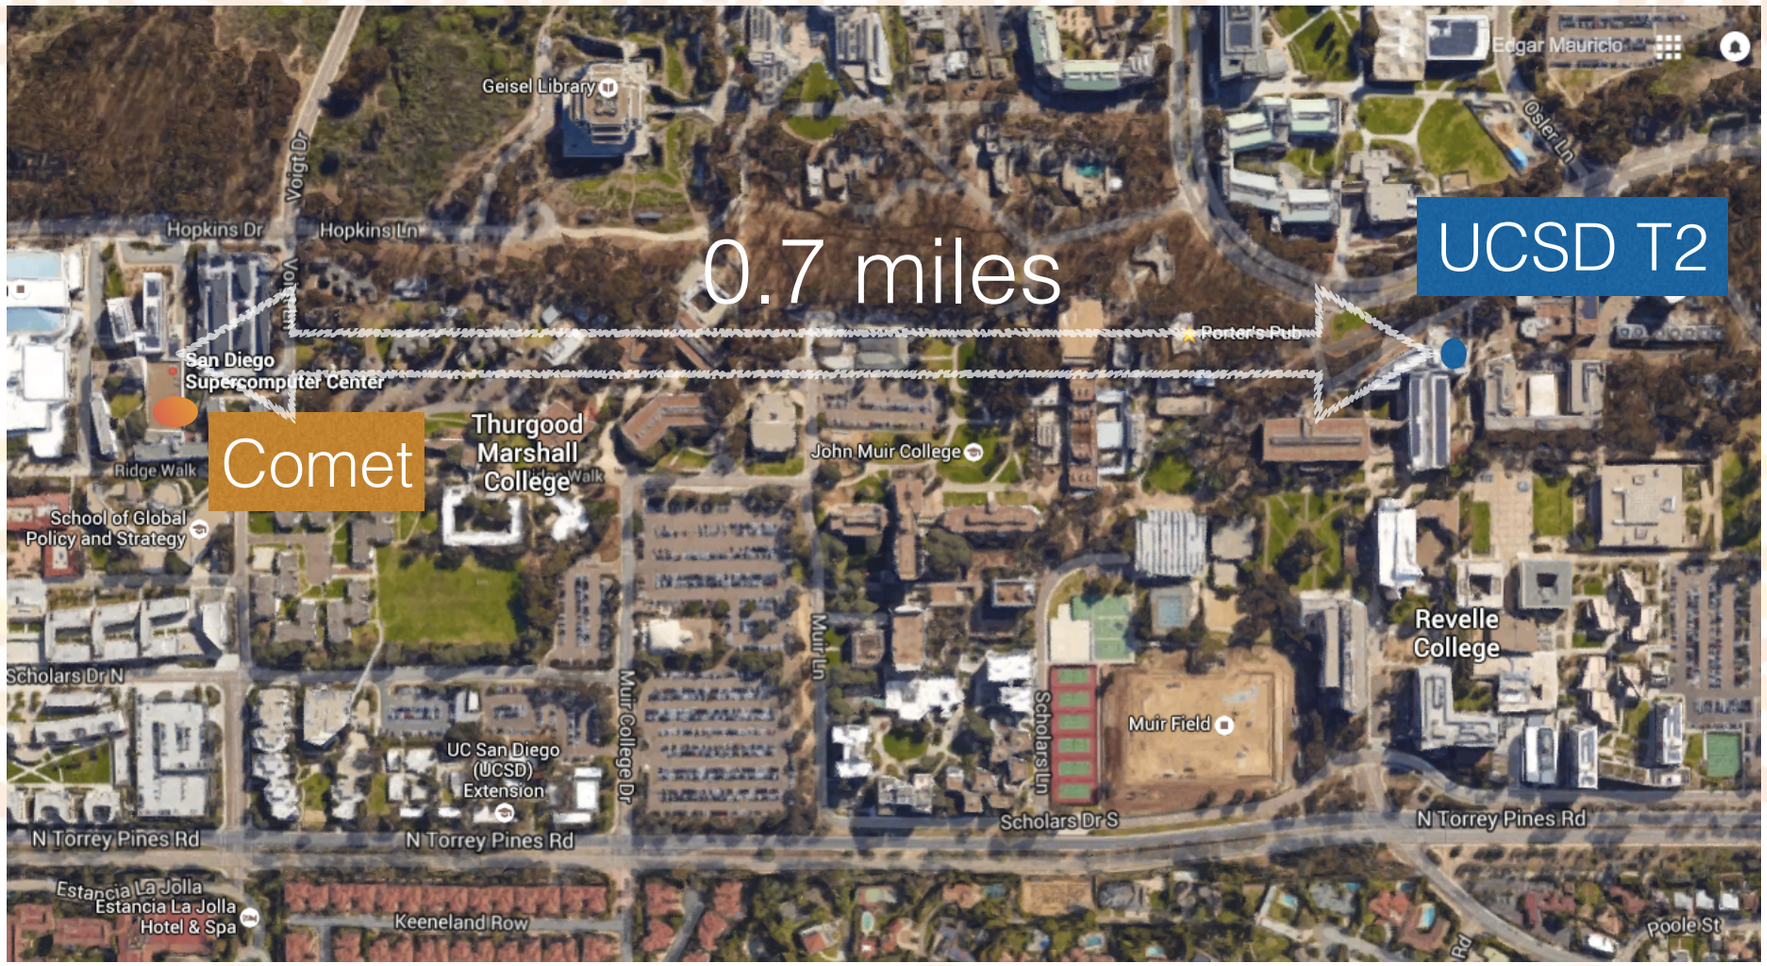

Comet by the numbers

System Component	Specs
Number of Racks	27
Computes nodes per rack	72
Cores per Node	24 x Dual socket Haswell 12 core @ 2.5GHz
Ram per Node	128GB
Total number of Cores	47000

Three ways of accessing a Comet

1. Usual "old" way submitting jobs to SLURM batch system
2. Virtual Cluster Interface
3. Science Gateways (web portals for science domains)

Virtual Cluster Interface

- A "cloud" like API to get resources and monitor them.
- For example how to request a VM:

```
cm comet start osg --count=1 --walltime=2d --allocation=csd428
```

Cluster Name

Comet allocation

XSEDE Allocation cannot be used, yet.

How Comet/ Condor integration works

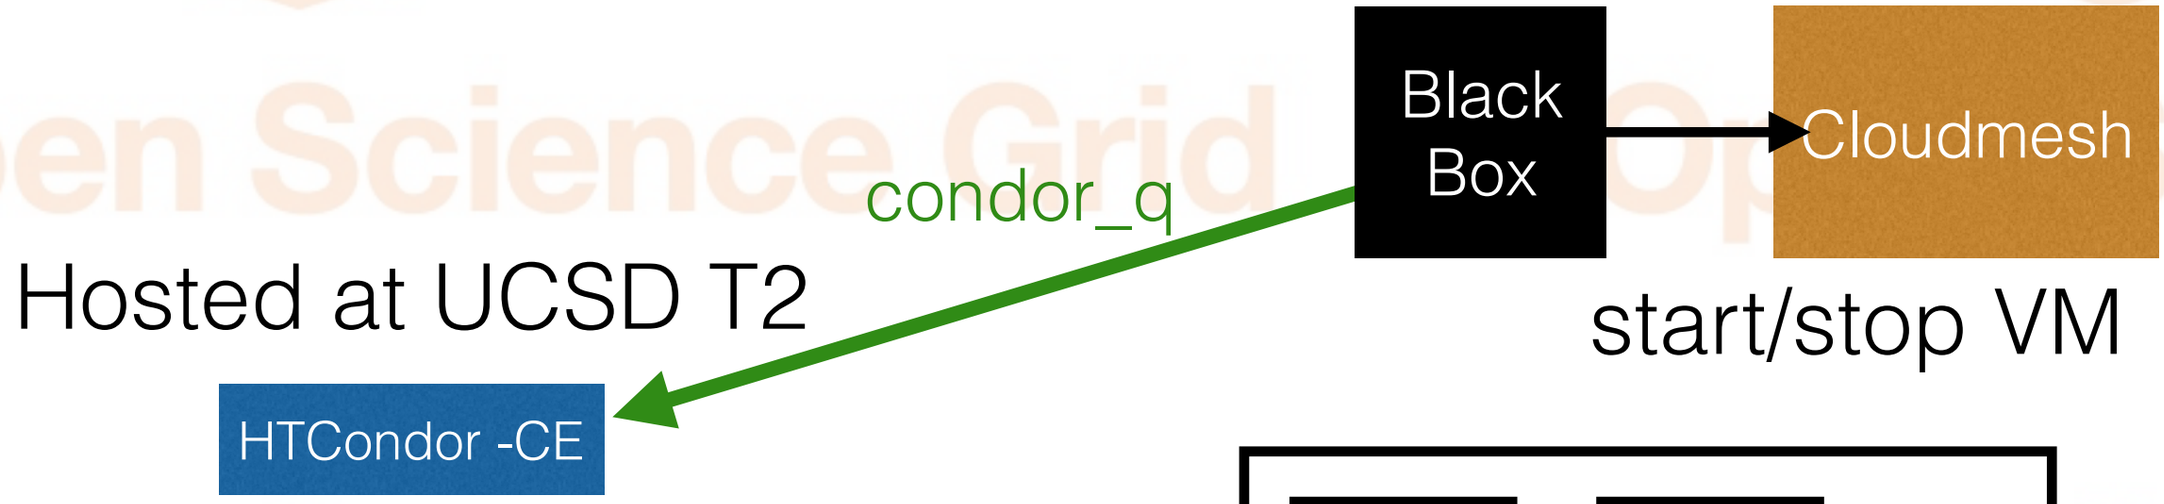

- job1: +project_Name="allocation1" +CometOnly=True
- job2: +project_Name="allocation1"+CometOnly=True
- job3: +project_Name="allocation1"+CometOnly=True

Where does OSG kick in?

Glideins can get into Comet using the already existing UCSD T2 grid infrastructure

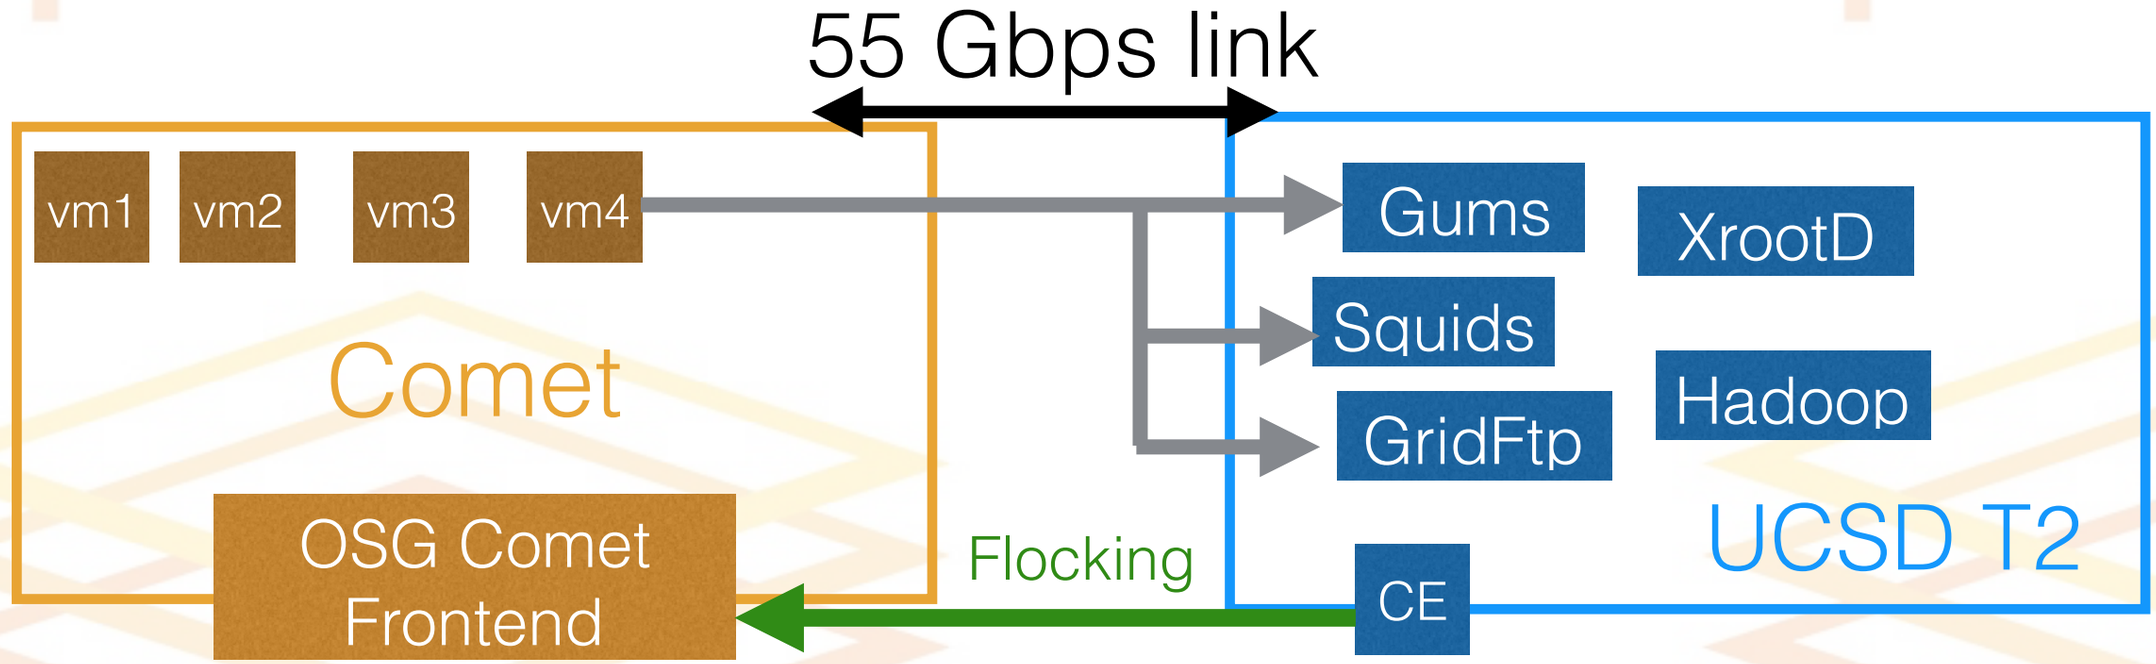

From the Comet Virtual Cluster Head Node

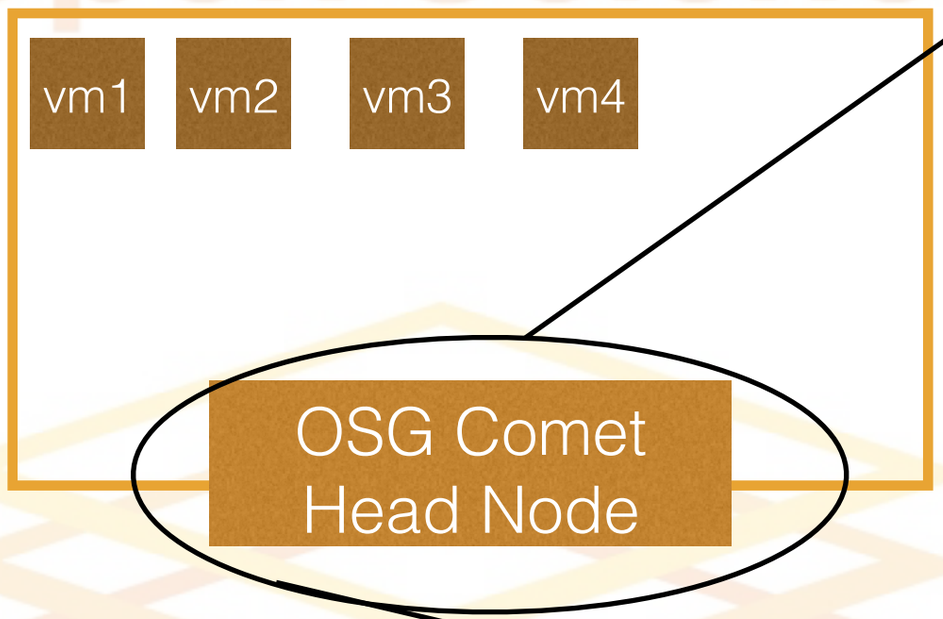

OSG Comet Head Node

- Puppet Master / Foreman
- DHCP
- HTCondor Central Manager
- Squid Proxy

~Same puppet config as the T2

From the Grid's perspective

It looks just as another UCSD entry in GlideinWMS:

```
<entry name="OSG_T2_US_UCSD_gw6_Comet_mcore" auth_method="grid_proxy+project_id" enabled="True" gatekeeper="osg-
gw-6.t2.ucsd.edu osg-g\
w-6.t2.ucsd.edu:9619" gridtype="condor" trust_domain="grid" verbosity="std" work_dir="Condor">
        <config>
            <max_jobs>
                <default_per_frontend glideins="5000" held="50" idle="100"/>
                <per_entry glideins="10000" held="1000" idle="4000"/>
                <per_frontends>
                </per_frontends>
            </max_jobs>
            <release max_per_cycle="20" sleep="0.2"/>
            <remove max_per_cycle="5" sleep="0.2"/>
            <restrictions require_glidein_glexec_use="False" require_voms_proxy="False"/>
            <submit cluster_size="10" max_per_cycle="25" sleep="2" slots_layout="fixed">
                <submit_attrs>
                    <submit_attr name="+CometOnly" value='"True"'/>
                    <submit_attr name="+maxMemory" value="98304"/>
                    <submit_attr name="+xcount" value="24"/>
                </submit_attrs>
            </submit>
```

Allocation base

Do not run at UCSD

24 cores per Pilot

Achievements

- Successfully ran LIGO, Xenon1T, IceCube, CMS Production and CMS UCSD user jobs in the Virtual

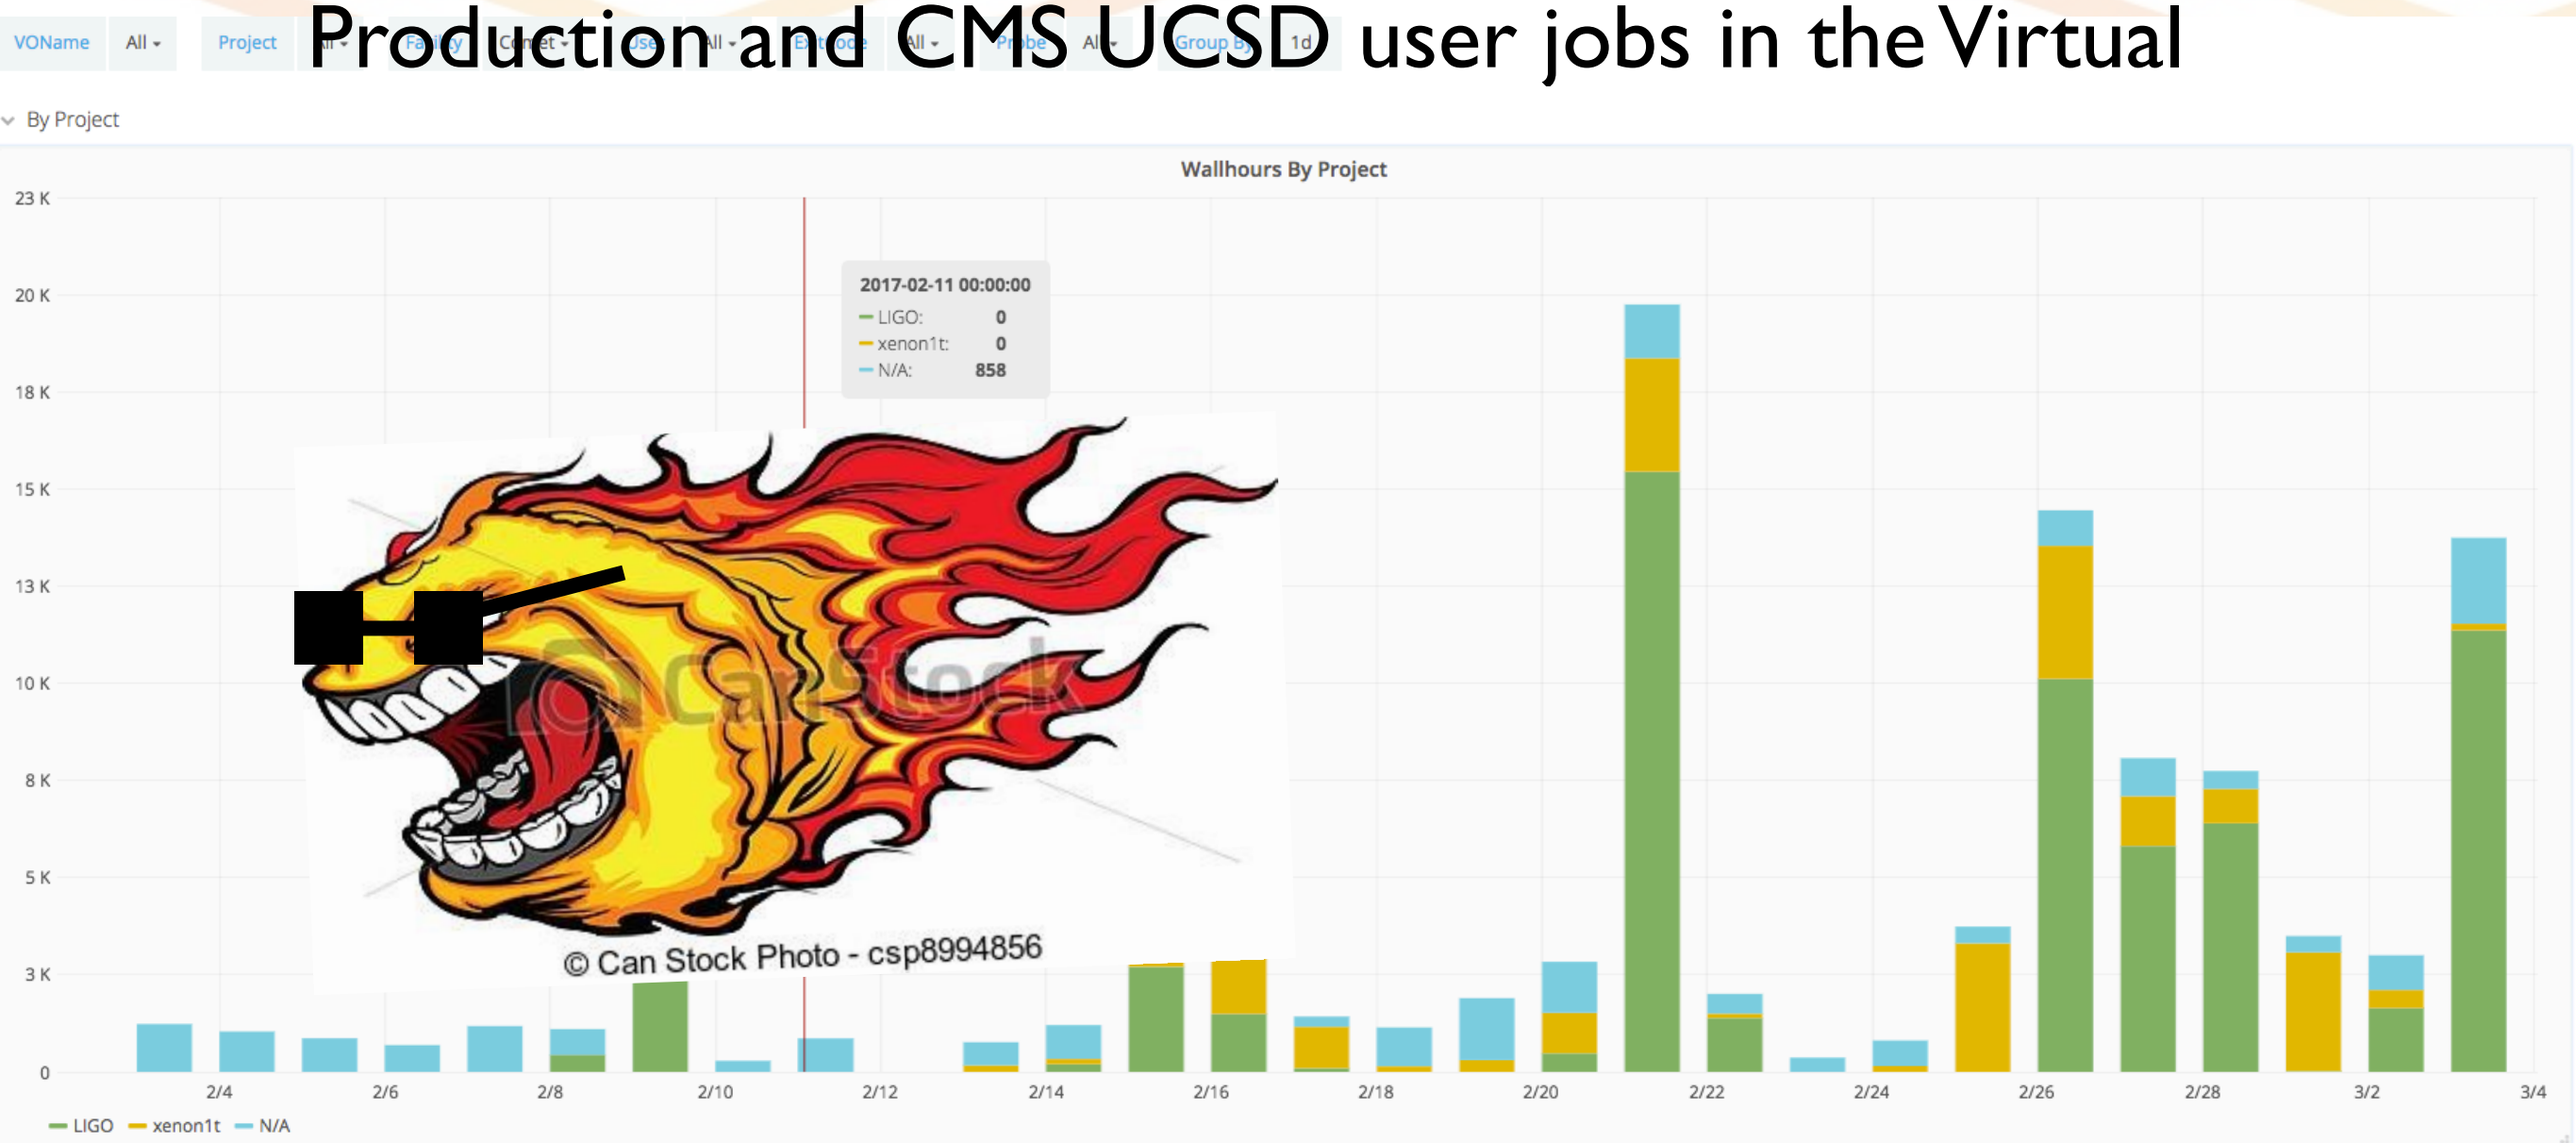

Achievements

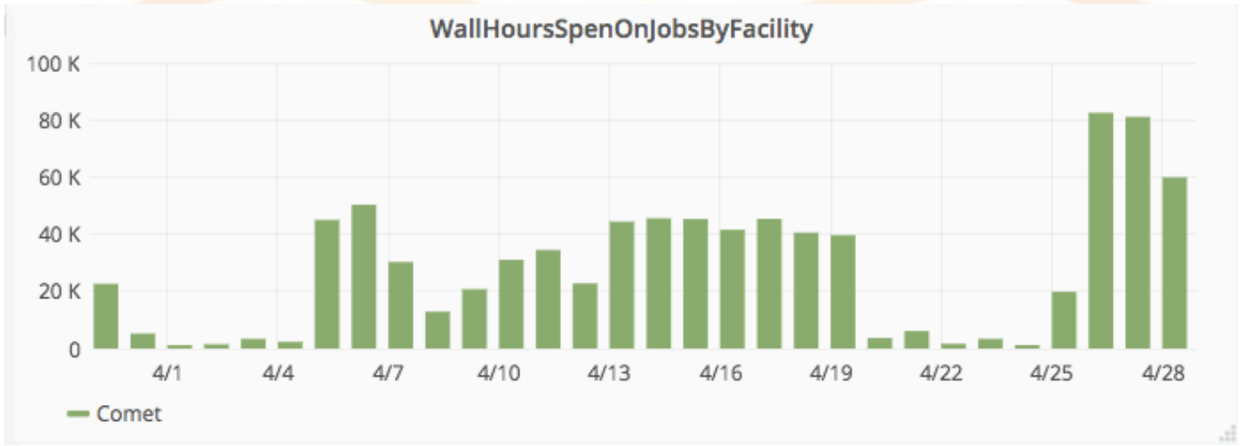

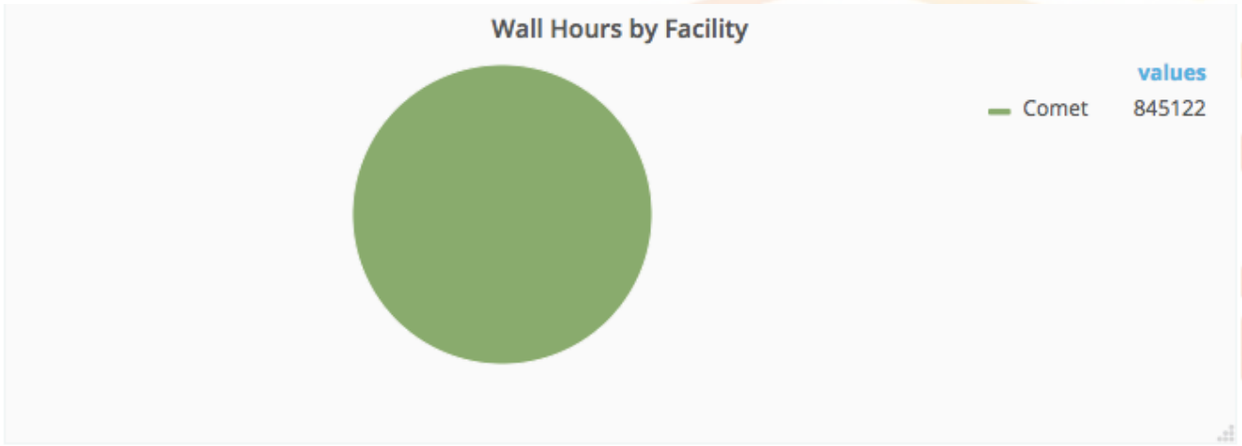

By Project

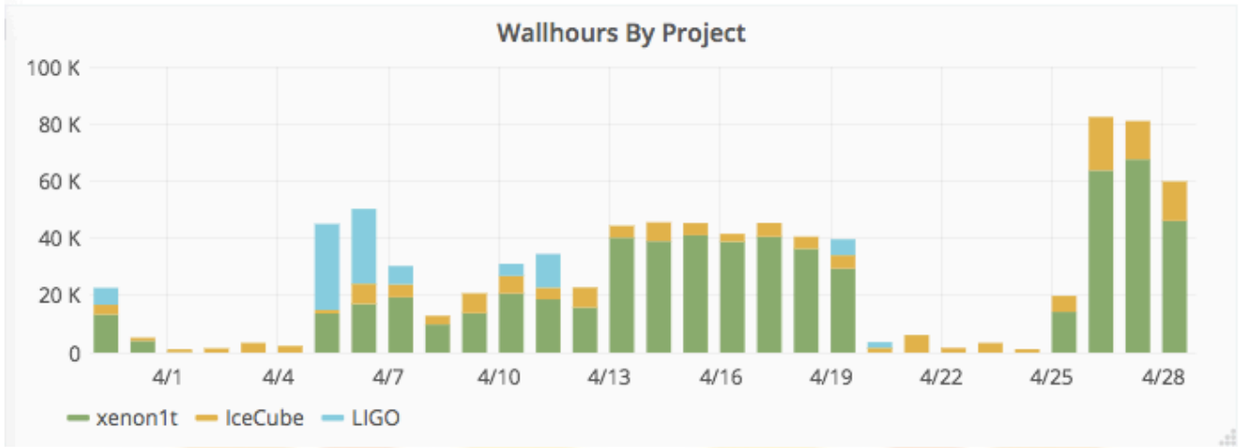

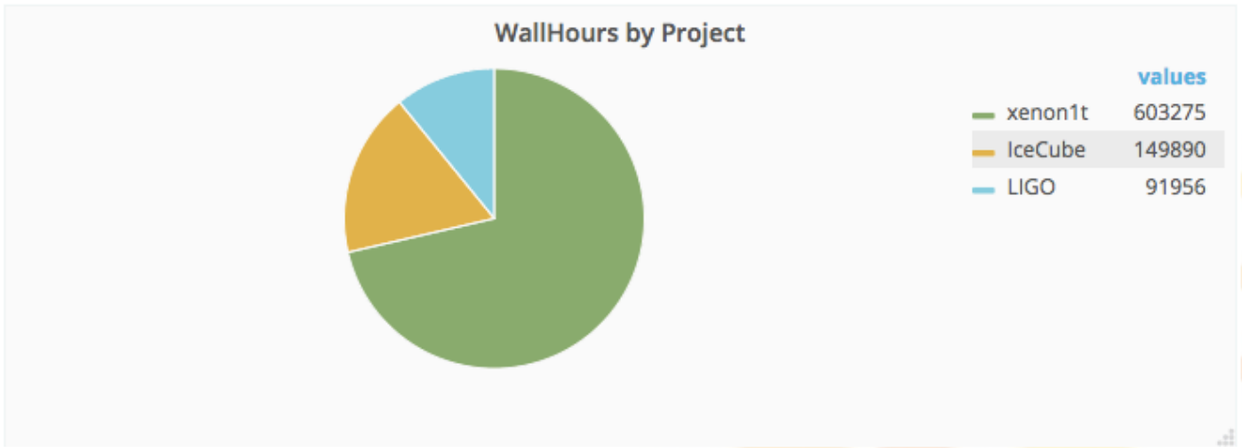

In last 30 days, It has delivered 845k CPU Hours with the OSG VC.

Achievements

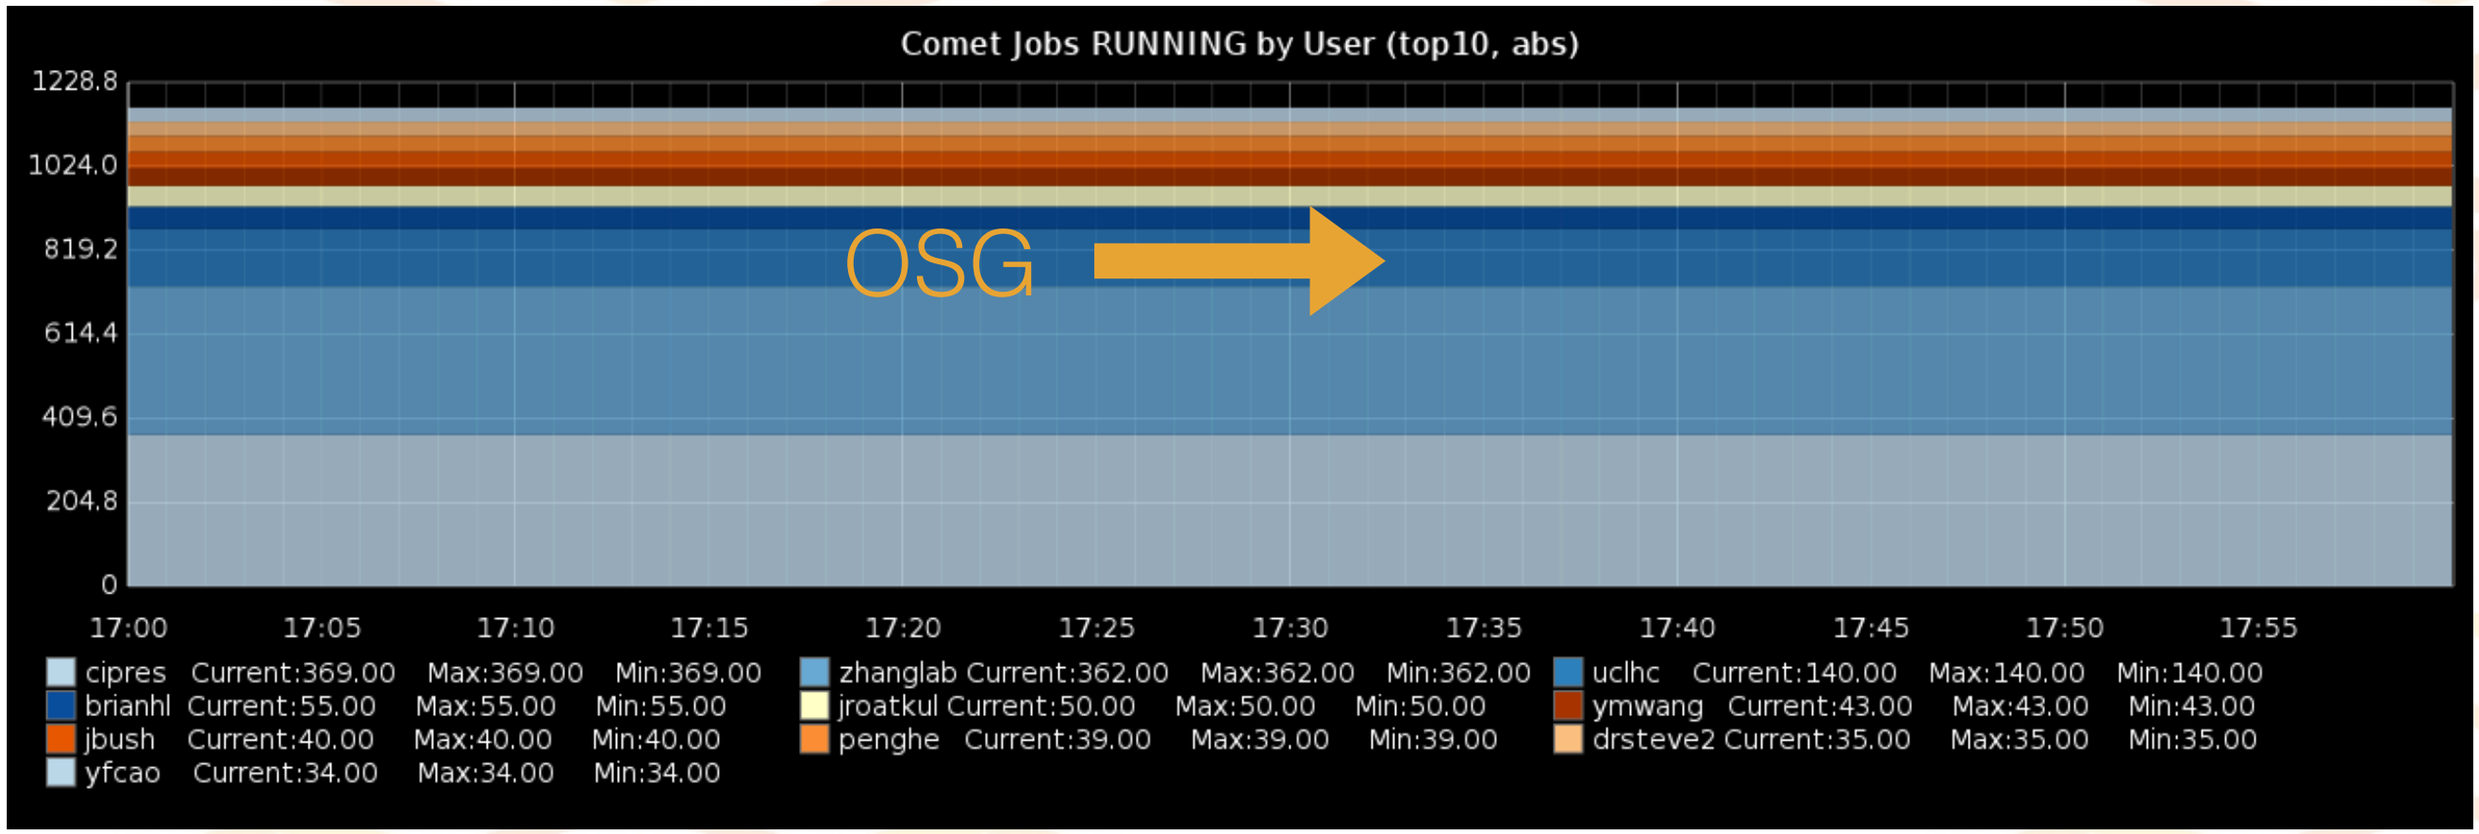

OSG: 3rd largest user of Comet

Scavenged Used Cycles

OSG Comet Virtual Cluster would like to make use of *unused cycles*…

free science

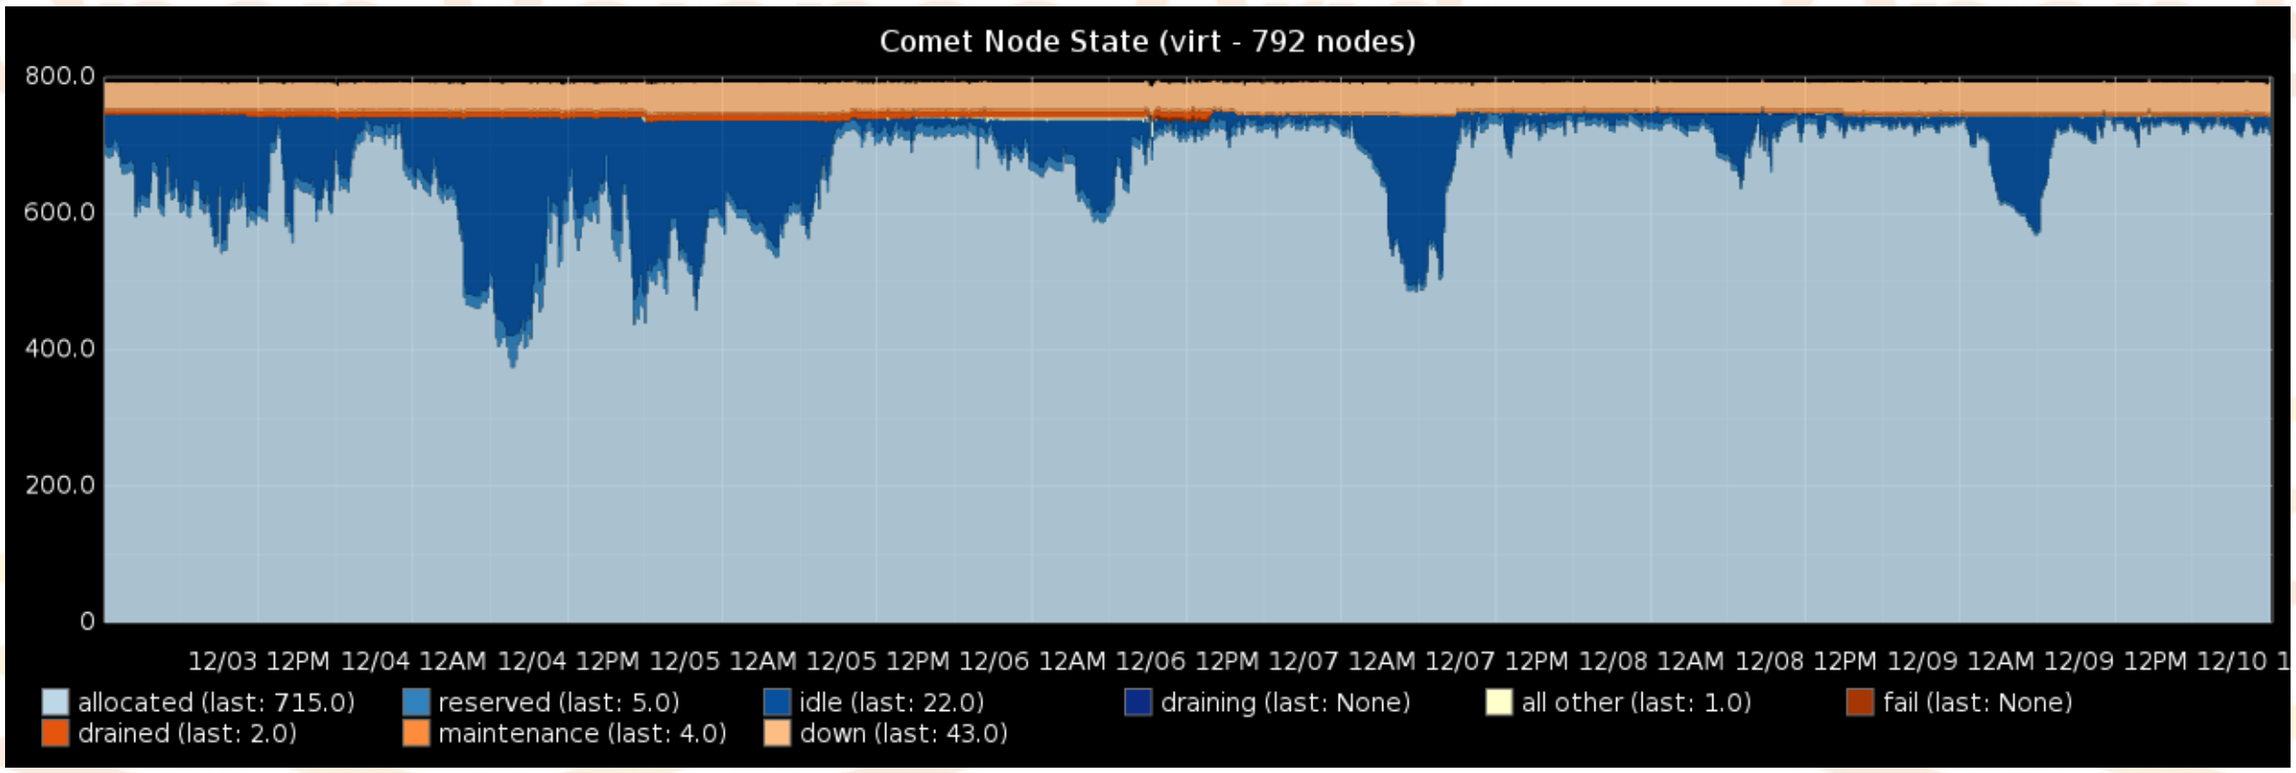

Comet available nodes shown in dark blue… 7 days in December 2016

Scavenged Used Cycles

OSG Comet Virtual Cluster would like to make use of *unused cycles*…

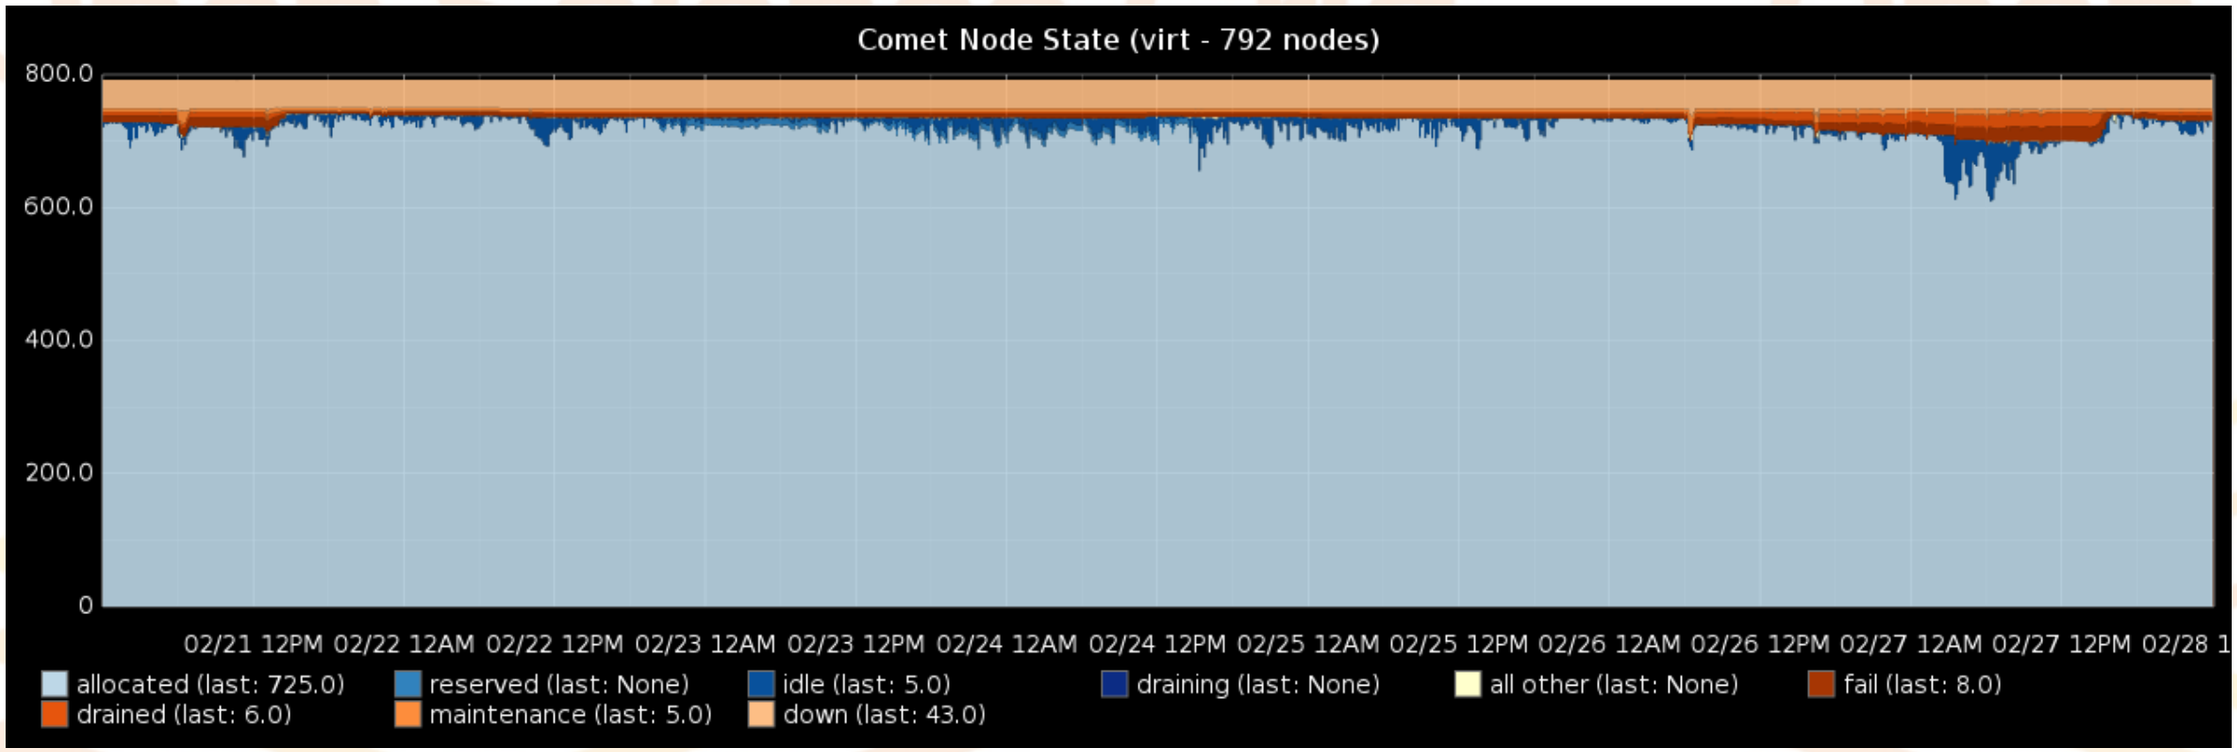

Comet available nodes shown in dark blue… 7 days in February 2017… where did they all go?

Acknowledgements

- Trevor Cooper, Dmitry Mishin (SDSC) and the whole Comet team.
- Fugang Wang and Gregor von Laszewski (Indiana University) for the troubleshooting in the Comet Cloudmesh.
- Terrence Martin (UCSD) for the full integration setup and help debugging the network infrastructure at Comet Virtual Cluster.
- HTCondor team. THANKS !!!!!

Edgar's Wishlist

These are things that would make my life easier:

- Support for Condor Metrics to InfluxDB
- Support for Python bindings in with python2.7 and pip install.
- Conditional flocking. (Even with just submitter ads)
- A job Ad that tells me the CPU efficiency of a job.
- Declare a resource uncontested (so its use does not count towards priorities).

Questions?

Contact us at:

1-900-OSG-Comet-Masters

Just Kidding

Contact us:

emfajard@ucsd.edu

Thank You

Additional Slides