

OSG Rides a Comet

15 March 2016

Edgar Fajardo
on behalf of OSG Software and Technology

Acknowledgments

Although I am the one presenting. This work is a product of a collaborative effort from:

- Terrence Martin (UCSD) for setting up puppet and the PXE boot.
- Rick Wagner and the whole Comet Team for their support and help on making this happen.

Look up in the grid

- Dude 1: It's a new grid site?
- Dude 2: It's a cloud service?


No

It's a COMET



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

Two ways of accessing a Comet

1. Usual “old” way submitting jobs to SLURM batch system

2. Virtual Cluster Interface

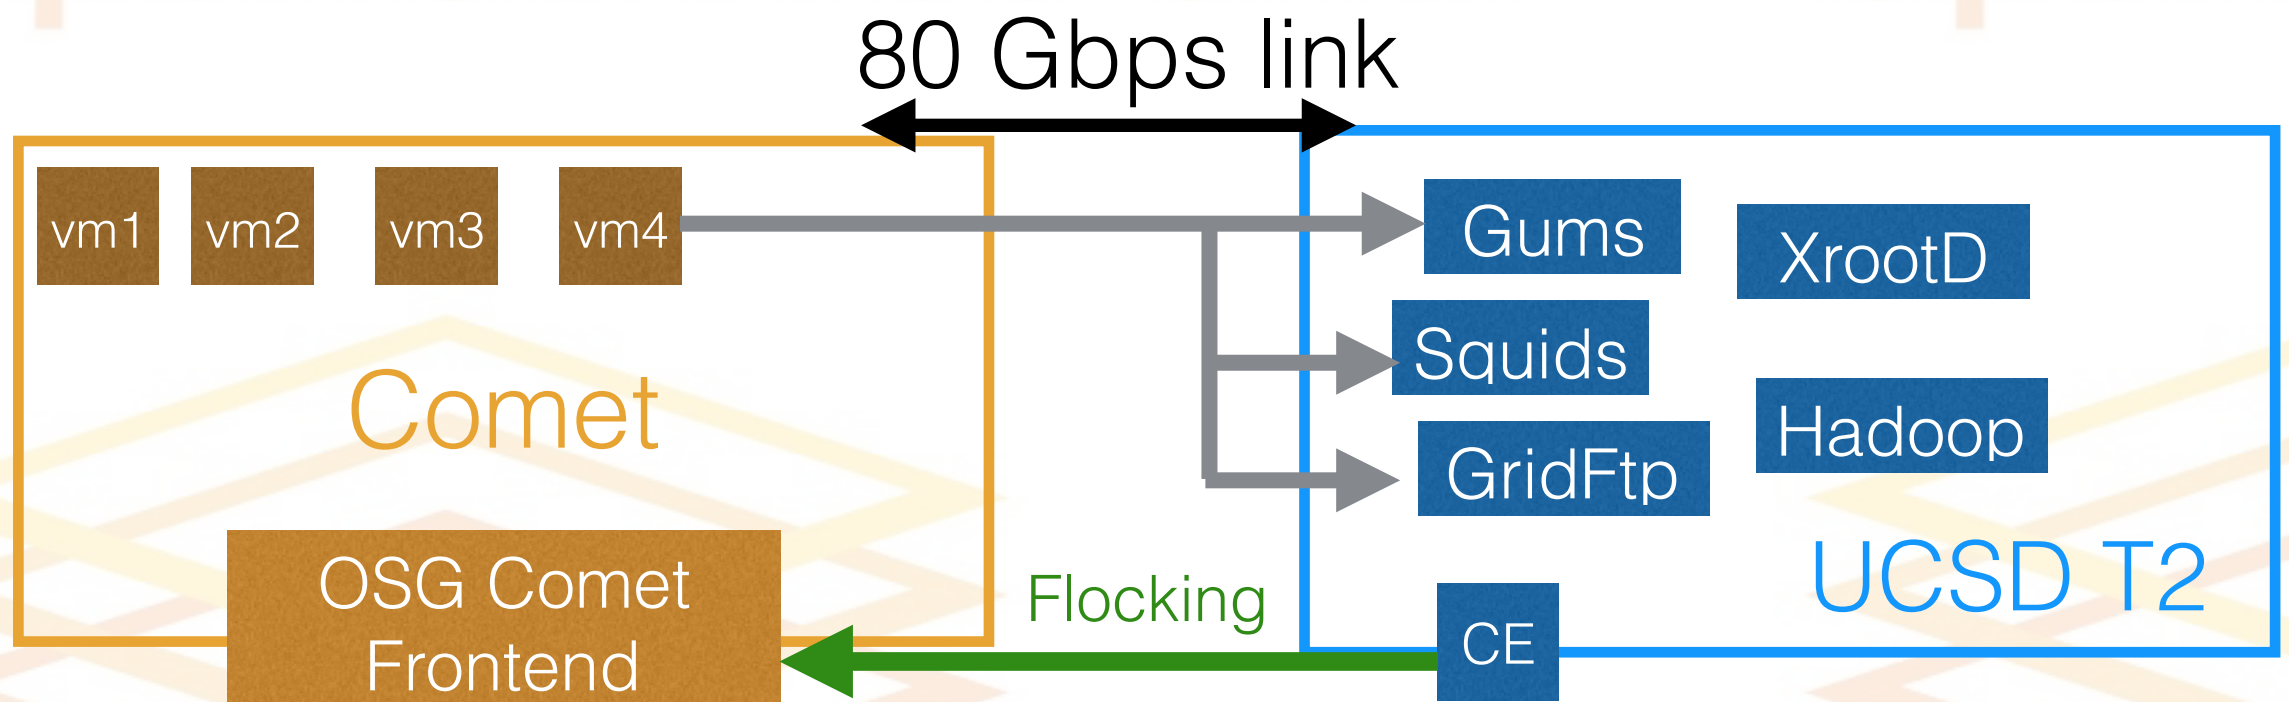
3. Science Gateways (web portals for science domains)

Virtual Cluster Interface

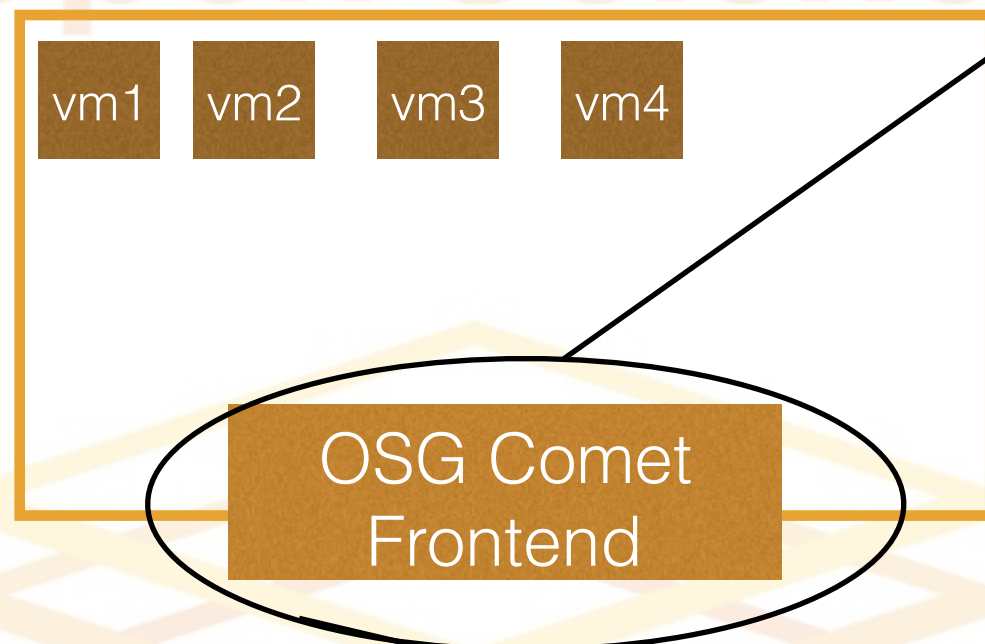
- A user U request a virtual machine with an allocation X for a walltime Y using the Nucleus API.
- The virtual machine PXE boots against a DHCP server (frontend) controlled by the user.
- Each virtual machine is registered as a SLURM job for accounting purposes. Agnostic to the users and to me :)
- User has full control of the environment (makes installation of specific libraries easier).

Where does OSG kick in?

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



From the Comet Frontend's Perspective



OSG Comet Frontend

- Puppet Master / Foreman
- DHCP
- HTCondor Central Manager
- Squid Proxy
- NAT (temporarily)

~Same puppet config as the T2


How does this look now

- OSG got access to the the dev rack of Comet ~1k cores (not hyper threaded)
- Dev Rack is limited to outside connectivity of 1 Gbps NATed to the outside.
- Around November 2015 LIGO and OSG VO's started using the virtual interface.
- The virtual interface will be in production around the end of march 2016.

Two ways of accessing the virtual cluster

	Allocation Based	Backfill
Wall Time Guaranteed	Yes	No
Only run at Comet	Yes	No
Pilot needs project_id	Yes	No
Who spins up the VM	OSG	Comet Admins

Where does OSG stand now in Comet?



There is a script that looks at the queues (gatekeepers) and requests VMs to Comet Nucleus API.

Who is running at Comet can be seen [here](#). Thanks to Tyson Jones for the plots

What is in the short term future

- Spin up VM's given an allocation. Making sure only glide ins with that allocation run there.
- Move to the production infrastructure (no longer behind a NAT).
- Try to backfill flock CMS glideins to Comet.
- Mount some lustre filesystem based on the allocation.

Long Term Future

- Move to MultiCore
- Offer the possibility of a glidein taking over a whole virtualized rack. Multinode pilot (like Stampede).
- GPU access via the virtual interface

In Summary



© Can Stock Photo - csp8994856

Questions? Comments?

Contact us at:

1-900-Comet-masters

Just Kidding

Contact us:

osg-software@opensciencegrid.org