



# Vac and Vcycle

Luis Villazon Esteban

On behalf of  
**Andrew McNab**

University of Manchester  
LHCb and GridPP



# Overview

- Vac and Vcycle
- Vacuum model
- Machine/job features
- Accounting
- Usage
- Future plans



# Vac and Vcycle

- Both VM Lifecycle Managers
- Vac is a standalone daemon you run on each worker node machine to create its VMs
- Vcycle manages VMs on IaaS Clouds like OpenStack
  - Can be run at the site, by the experiment, or by regional groups like GridPP
- Both developed at Manchester as part of GridPP Clouds/VMs effort
  - With help from Lancaster, Oxford, IC, CERN, LHCb and ATLAS
- Both make very similar assumptions about how the VMs behave
  - The same LHCb and ATLAS VMs working in production on Vac and Vcycle

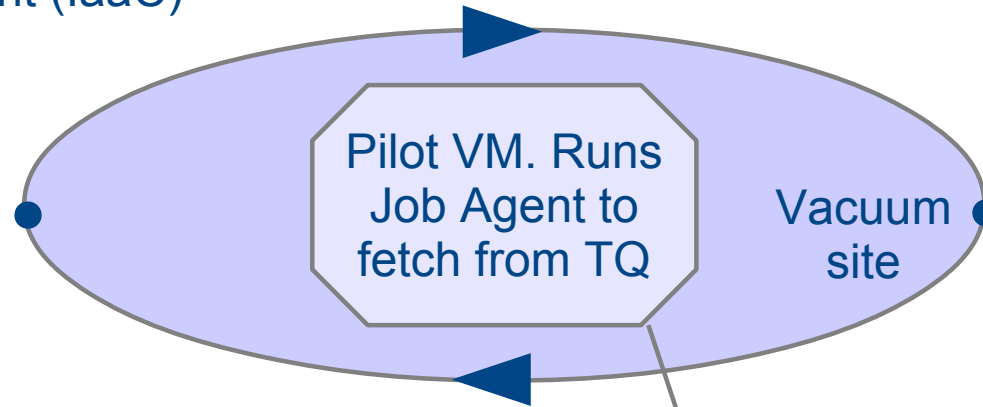


# “Pilot VM” lifecycle

- Vac and Vcycle assume the VMs have a defined lifecycle
- Need a boot image and user\_data file with contextualisation
  - Experiment provides procedure to make a site-wide user\_data file
- Virtual disks and boot media defined and VM started
- machinefeatures and jobfeatures directories may be used by the VM to get wall time limits, number of CPUs etc
- The VM runs and its state is monitored
- VM executes shutdown -h when finished or if no more work available
  - Maybe also update a heartbeat file and so stalled or overrunning VMs are killed
- Log files to /etc/machineoutputs which are saved
  - shutdown\_message file can be used to say why the VM shut down

# Vac's Vacuum Model

Infrastructure-as-a-Client (IaaSC)

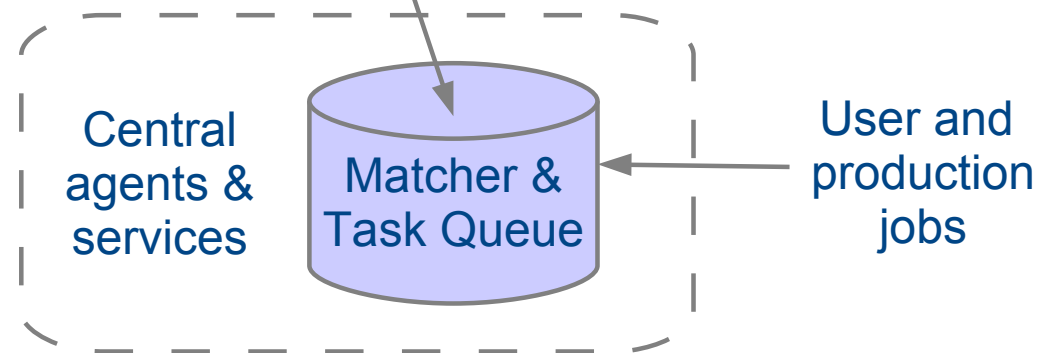


Since we have the pilot framework, we can do something much simpler.

Strip the system right down and have each physical host at the site create the VMs itself.

Instead of being created by the experiments, the virtual machines appear spontaneously “out of the vacuum” at sites.

Requests  
for real jobs



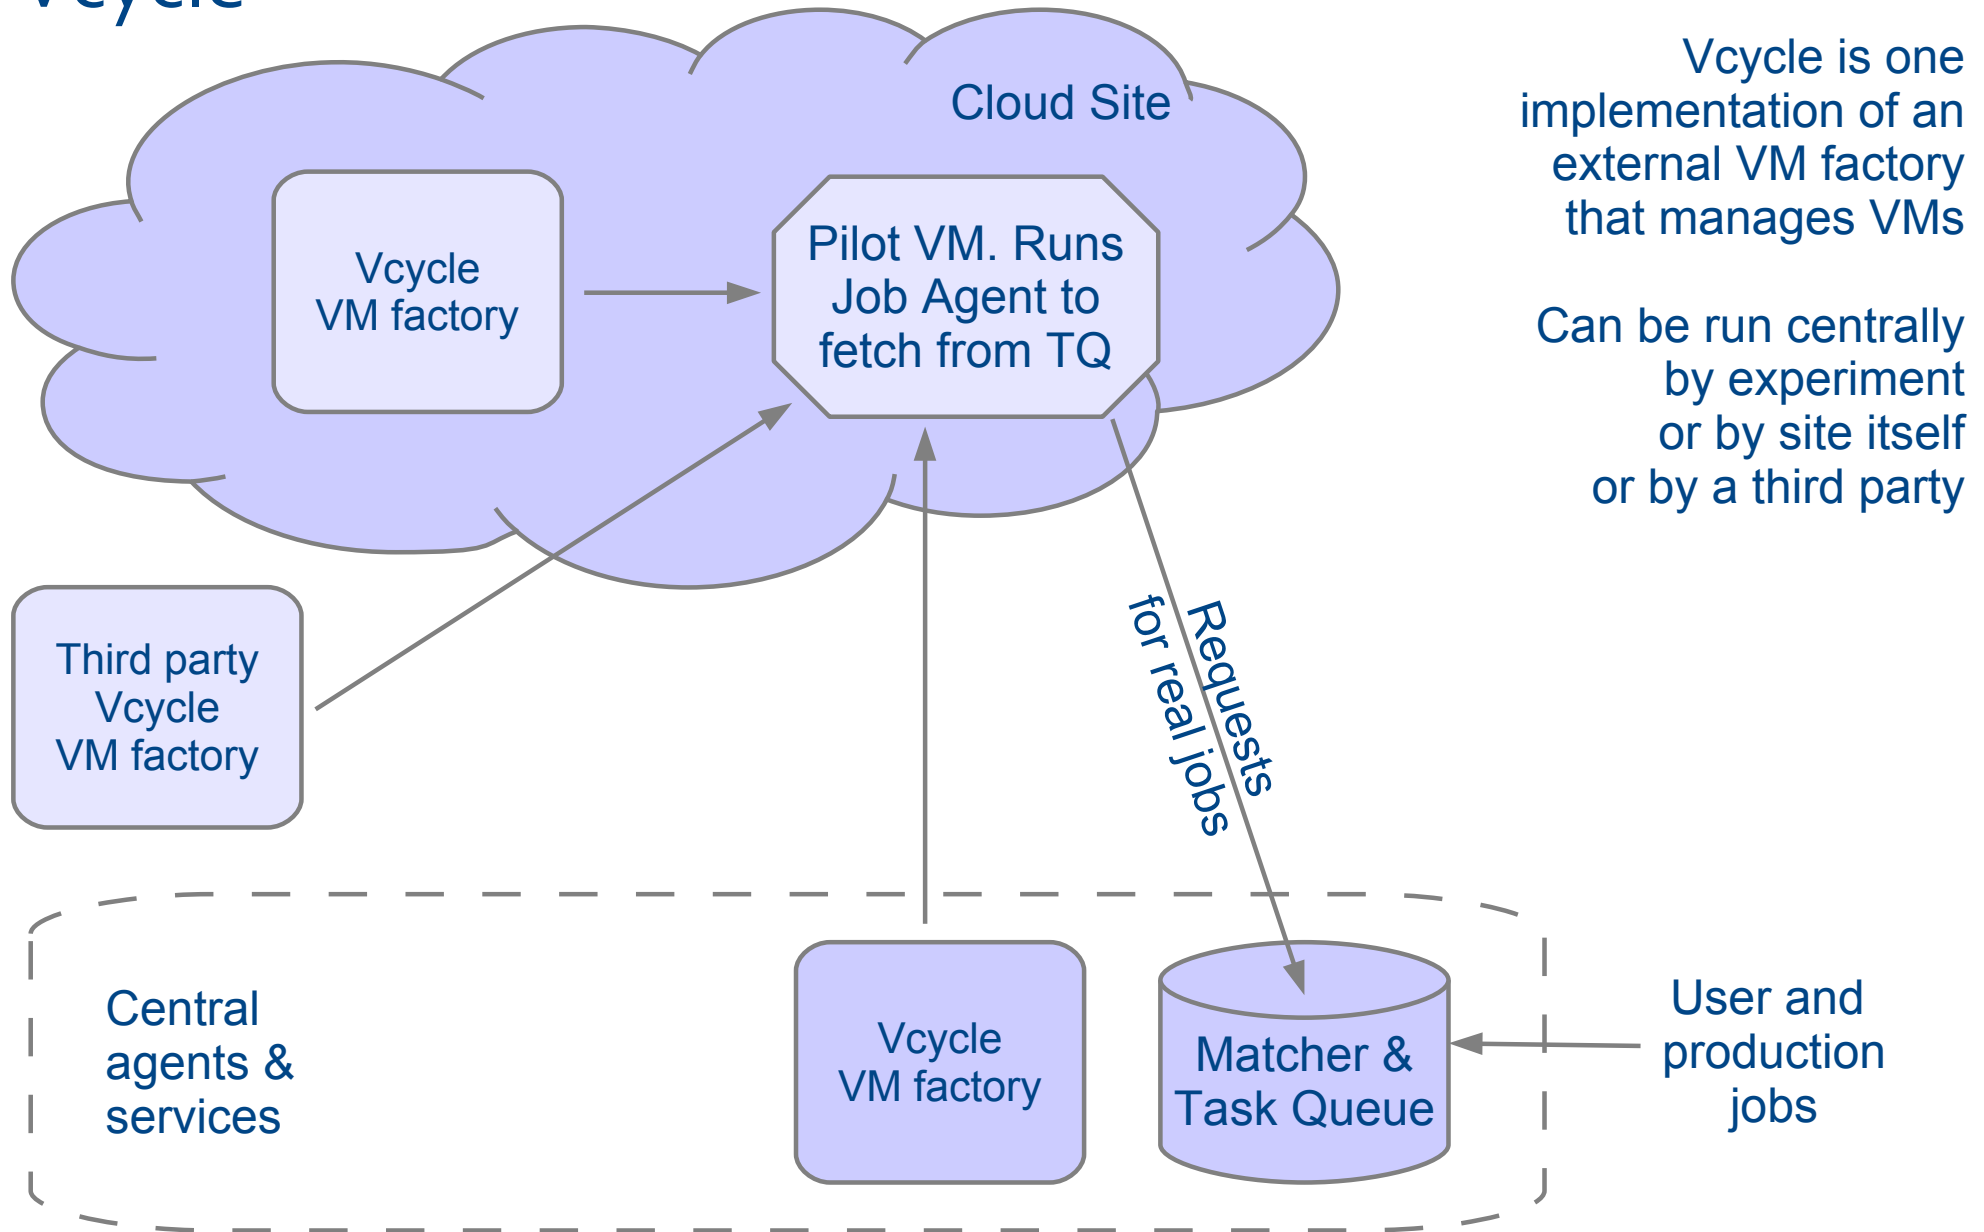
Ideally use same VMs as with IaaS clouds



# Vacuum model

- For the experiments, VMs appear by “spontaneous production in the vacuum”
  - Like virtual particles in the physical vacuum: they appear, potentially interact, and then disappear
- Following the CHEP 2013 paper:
  - *“The Vacuum model can be defined as a scenario in which virtual machines are created and contextualized for experiments by the resource provider. The contextualization procedures are supplied in advance by the experiments and launch clients within the virtual machines to obtain work from the experiments' central queue of tasks.”*
- At many sites, 90% of the work is done by 2 or 3 experiments
  - So a simple, reliable way of running their “baseload” of jobs is worthwhile
- cvmfs and pilots mean a small user\_data file is all the site needs
  - Experiments can provide a script to create the site's user\_data

# Vcycle





# Vcycle implementation

- Applies Vac ideas to OpenStack etc IaaS resources
- Experiment-neutral, and can be run by experiment or site or 3rd party
- Vcycle daemon creates VMs using user\_data file
- Watches what they do
- Backs off if they are failing to stay running
  - No work? Fatal errors?
  - Can also use shutdown messages to make better decisions
- Provides machine/job features via HTTP
  - Also used to collect log files, shutdown messages, and heartbeat updates
- Currently uses OpenStack native nova API
  - OCCi fork has been done by WLCG team at CERN (Luis Villazon Esteban)
  - Reintegration and EC2 support in progress



# Machine/job features

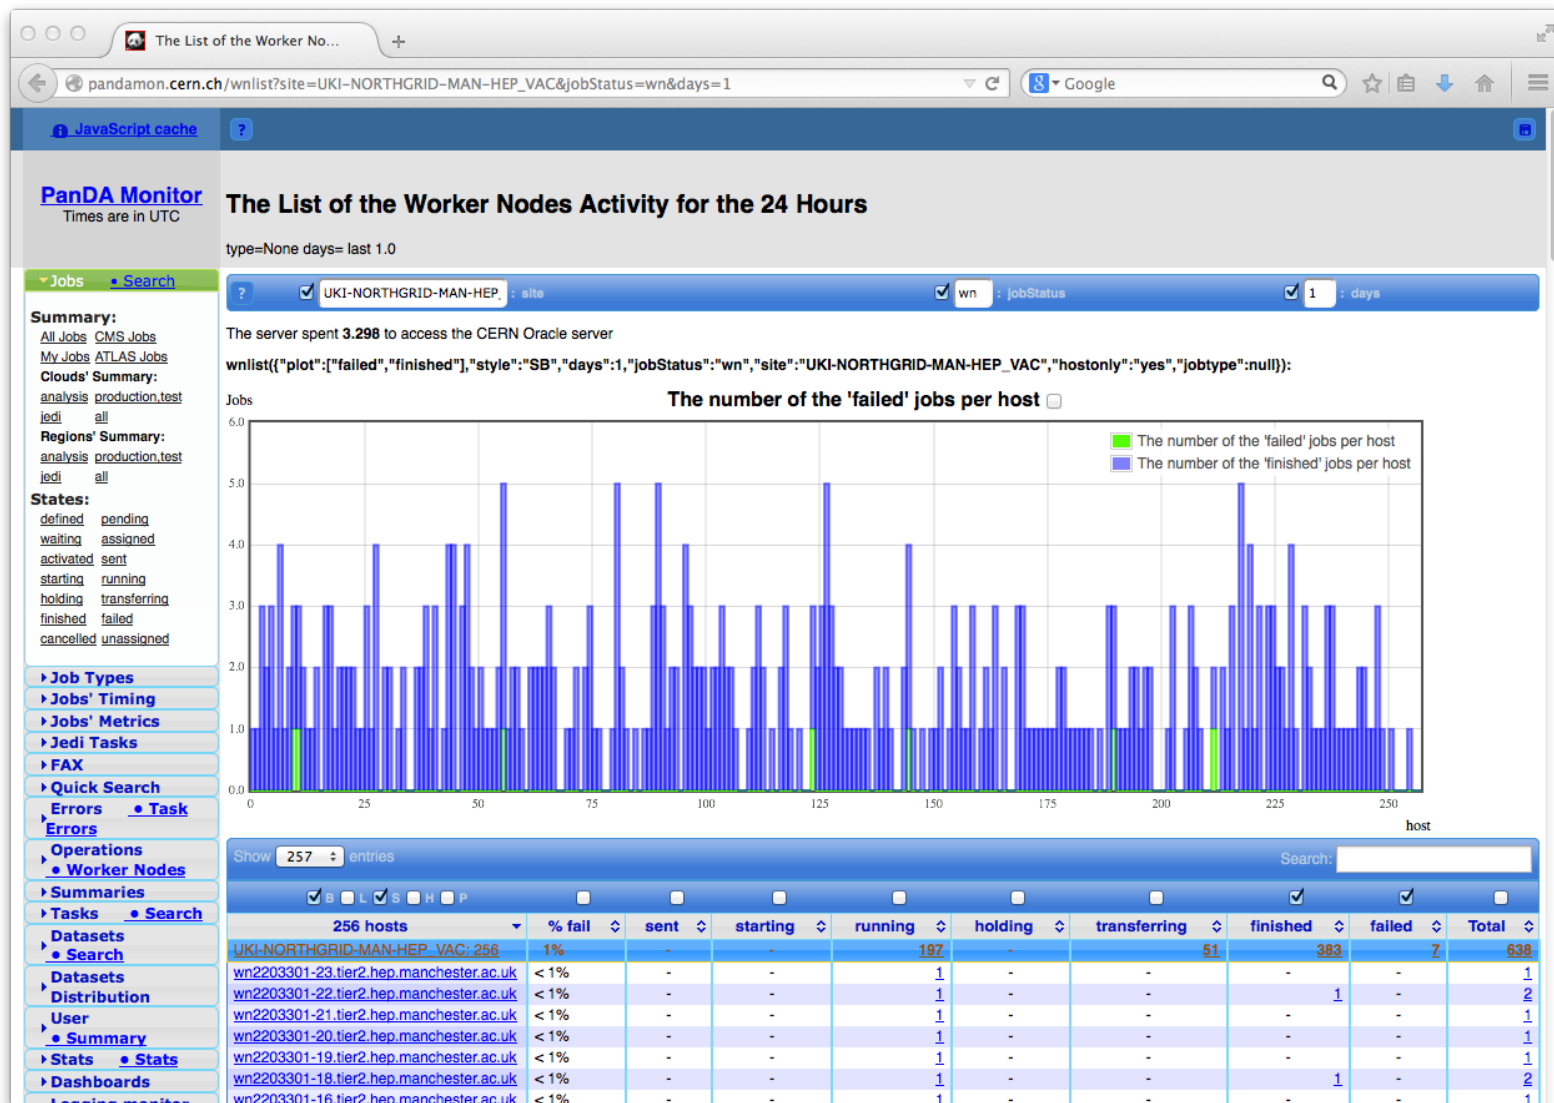
- Proposed HEPiX protocol and current WLCG task force
- Allows site/host to communicate details of machine and the job slot to the job or VM
  - HSo6, shutdown time for VM/host, CPU and memory limits, ...
- One key file per key/value pair, in one of two directories
- The Vac factory node offers these directories to its VMs via NFS over its internal private network
  - Also a writeable NFS directory for log files, shutdown reason, heartbeat files etc
- Vcycle does this use HTTP(S) web server on the Vcycle machine
- The basis for several scenarios for telling VMs and payload jobs what resources they have and how long they can run for
  - Want graceful termination of VMs to avoid disrupting payload jobs
  - `/etc/machinefeatures/shutdowntime` always set using `max_wallclock_seconds`



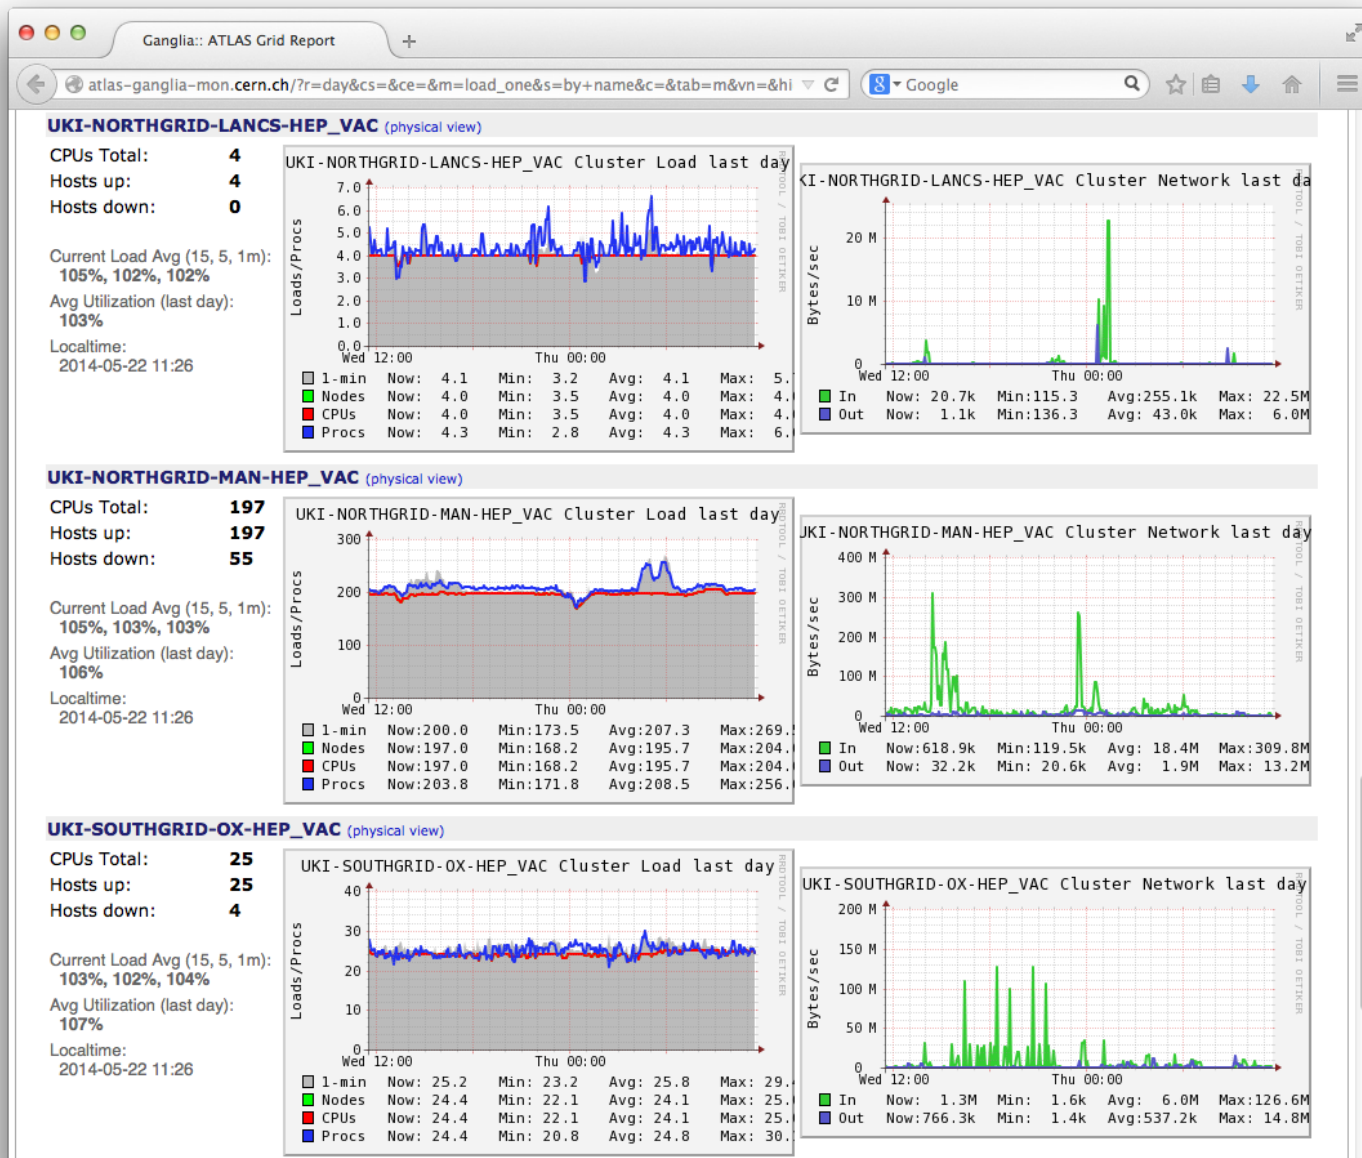
# Vac and Vcycle at sites

- Manchester
  - Vac running ATLAS and LHCb VMs
  - Vcycle instance managing Imperial (LHCb/ATLAS) and CERN (ATLAS) VMs
- Lancaster
  - Vac running ATLAS and LHCb VMs
  - Vac being added to old farm now
- Oxford
  - Vac running ATLAS and LHCb VMs
- Imperial
  - GridPP OpenStack tenancy “gridpp-vcycle” has LHCb and ATLAS VMs
- CERN
  - LHCb tenancy managed by Vcycle on LHCb vbox at CERN
  - (A few) ATLAS VMs managed by Vcycle at Manchester

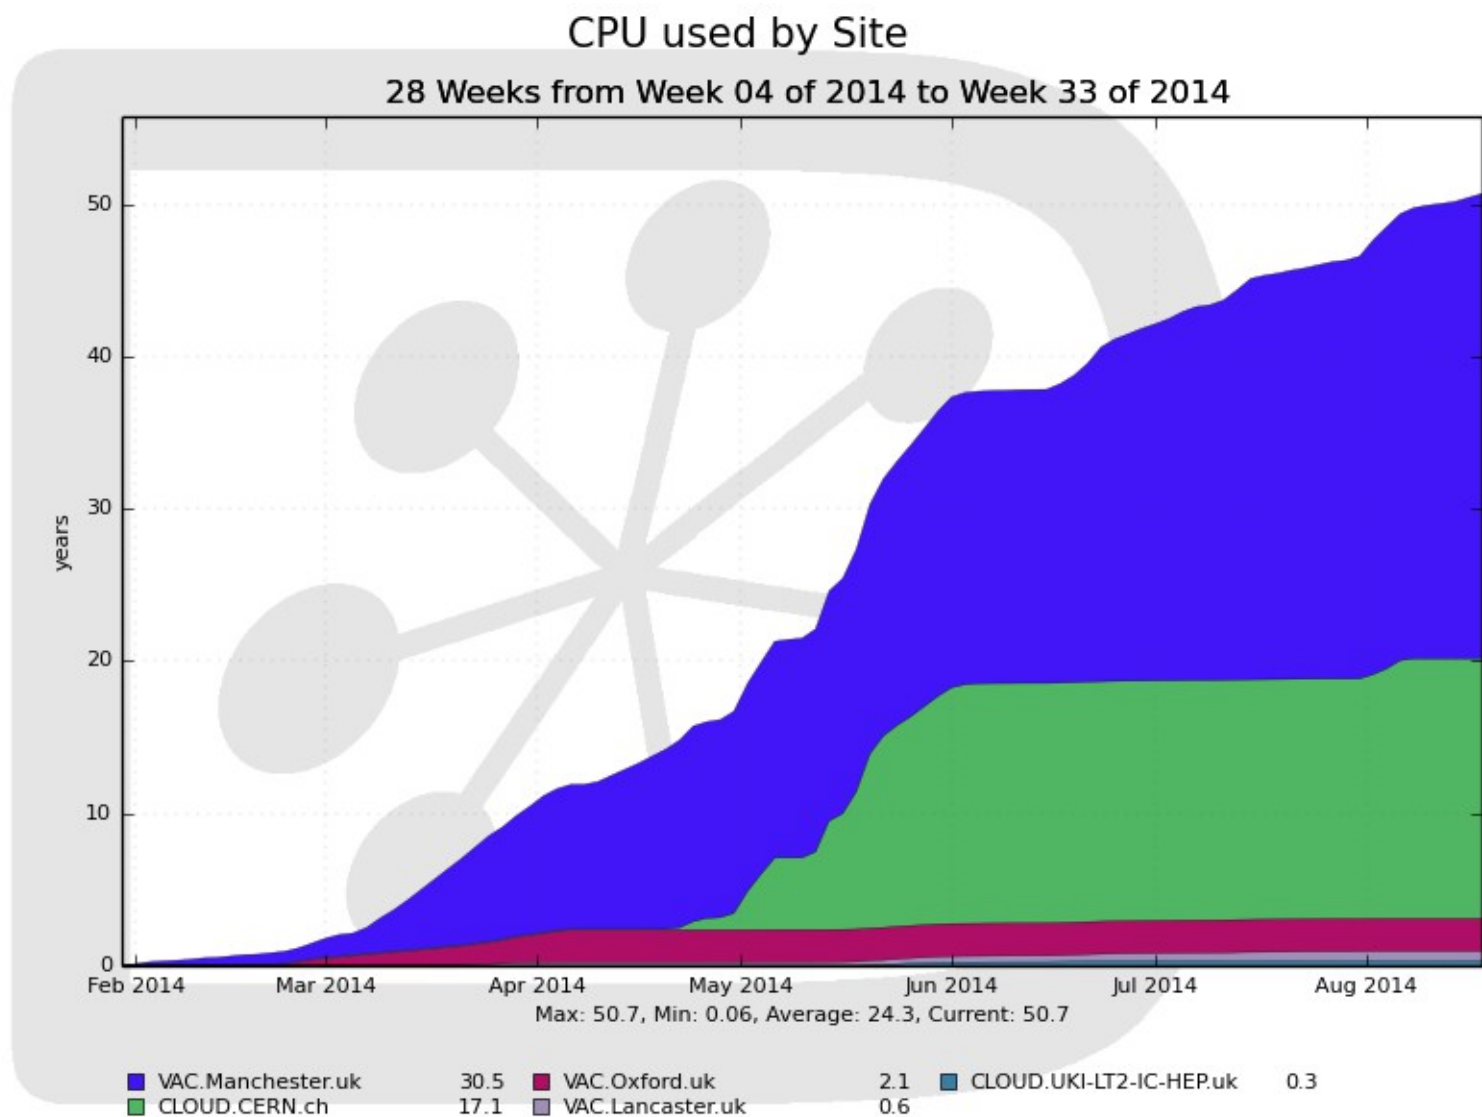
# ATLAS Panda jobs running on Vac VMs



# Ganglia monitoring of ATLAS Vac sites



# LHCb CPU time from Vac and Vcycle VMs



Generated on 2014-08-20 09:29:54 UTC



# Vcycle work at CERN

- Using Vcycle to run VMs in tenancies available through EGI
  - See the presentation on the WLCG use case.
- Thinking of the problem in terms of wanting to provide a “steady pressure” of job slots as long as there are payloads waiting to be run
- Vcycle is one way of doing this, as it monitors whether VMs find payloads to run or not
- Have produced a version of Vcycle that uses OCCI API rather than OpenStack's Nova API
- Using this for ATLAS, LHCb and looking at CMS. Would be interested in collaboration with ALICE too.
- Of interest as EGI aims to take operational responsibility for the cloud infrastructure and to take care of APEL accounting etc.



# Summary

- Vac provides a simple way for sites to run VMs
- Vcycle provides a simple way to manage VMs on OpenStack
- Both demonstrated with ATLAS and LHCb production jobs
  - Measured VM efficiency for MC is very good now
- GridPP is in a very good position in this field
- They are able to run production work in VMs if/when the experiments or projects start pushing in that direction
- See <http://www.gridpp.ac.uk/vac/> for RPMs, Yum repo, links to GitHub, docs, man pages etc