



# Lightweight UK sites using VMs

**Andrew McNab**  
University of Manchester,  
GridPP, and LHCb



# Overview

- Simplification with VMs
- Vac and Vcycle
- Vacuum Platform
- Experiment VMs
- Vac-in-a-Box
- Next steps



# Simplification with VMs

- GridPP needs to run the WLCG infrastructure in the UK with less people
- Experiments developed VMs for running at Cloud sites
- Put the two together: we developed simple daemons for managing those VMs
- This removes the need to run the “zoo” of grid middleware at sites:
  - CREAM/ARC, Batch, BDII, ARGUS, APEL, ...
- The experiments’ own frameworks (DIRAC, PanDA, ...) inside the VMs are sufficient



# Vac vs Vcycle

- Two GridPP systems aimed at running VMs
- Vac - autonomous hypervisors
  - Each VM factory machine creates VMs in response to observed demand for each type of VM
  - More mature of the two, better documentation
- Vcycle - uses OpenStack etc
  - Factories created via Cloud API in response to observed demand for each type of VM
  - Code is solid, but docs are minimal: just man pages

# Vacuum Platform

- Drafting an HSF technical note describing the interfaces between VMs and Vac/Vcycle
- For VM-authors and authors of other Vac/Vcycle-like systems
- Proposing this to EGI as the basis of a “community platform”
- VacQuery / VacMon
- VacUserData
- \$JOBOUTPUTS

THE HEP SOFTWARE FOUNDATION (HSF)

HSF-TN-2016-VACPLAT  
April 2, 2016

## Vacuum Platform

A. McNab<sup>1</sup>

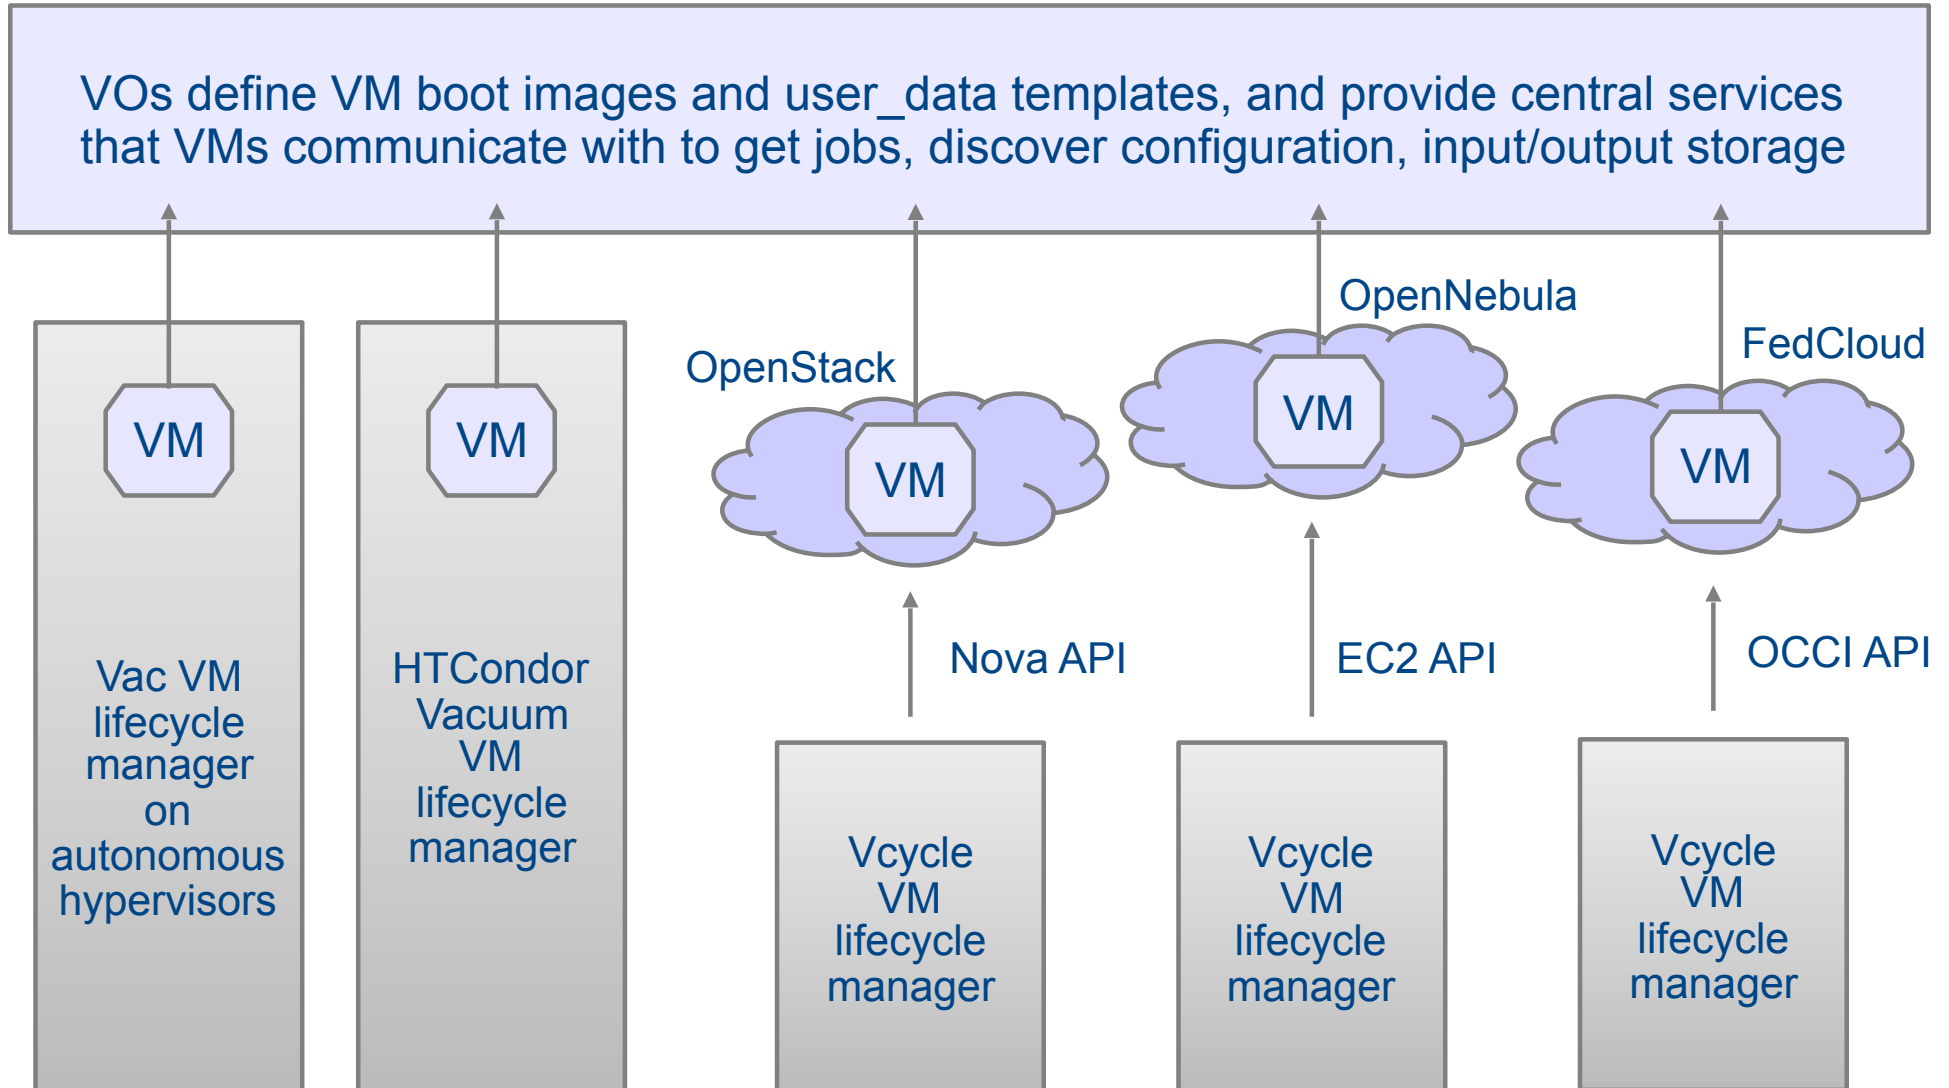
<sup>1</sup>*University of Manchester*

### Abstract

This technical note describes components of the Vacuum Platform developed by GridPP for managing VMs, including the \$JOBOUTPUTS, VacQuery, and VacUserData interfaces.

© Named authors on behalf of the HSF, licence CC-BY-4.0.

# Vacuum platform



# Vcycle status

- VMs created via Cloud API in response to observed demand for each type of VM
- Default is OpenStack plugin
  - EC2 plugin added at the start of 2016
- OCCI (EGI), DBCE, and Azure (MS) plugins contributed by CERN
- Vcycle is used to manage LHCb OpenStack tenancy at CERN (500 VMs)
- Also LHCb tenancy at CC-IN2P3 and GridPP at Imperial
- Code and man pages good, but no admin guide etc
- Machine/Job Features done by HTTPS server



# Vac status


- Vac 01.00 reached last month
- Provides an OpenStack-compatible environment to VMs, but using autonomous hypervisors (VM factories)
  - Much simpler implementation: 3500 lines of Python vs 1,000,000 for OpenStack
- VM factories communicate per-experiment load via UDP to achieve desired target shares
  - No single point of failure
- Puppet module and Admin Guide for installations
- Generates APEL accounting records for ssm send
- In production today at Manchester, Liverpool, Oxford, UCL with about 1000 VMs in total



# Experiment VMs status

- Need VMs that comply with the “Vacuum” model
  - Mainly that they shutdown if no work is available
  - Ideally they should use Machine/Job Features to discover maximum lifetime etc
  - Same VMs work for Vac and Vcycle
- Several iterations of production quality VMs for LHCb, ATLAS, CMS and GridPP DIRAC service (for smaller VOs)
- Emphasis since late 2015 on getting VMs integrated into experiments’ operations properly
- Currently LHCb, ATLAS, GridPP DIRAC in production
  - Suitable CMS VMs are being reworked

# LHCb DIRAC VM in action



# Vac-in-a-Box

- Make it even easier to install Vac
- No local site services required. Not even DNS.
- ViaB website run by GridPP has per-site dashboards
  - Configure list of hypervisors, IP addresses, experiments, ssh keys, certificates etc
- USB boot image to download, which gets Kickstart file
- First hypervisor has full stack: DHCP, TFTP, Squid, /etc/hosts, Vac
- Second hypervisor PXE-boots off the first, and gets a full stack too. Etc etc.
- Configuration changes via hourly YUM updates from ViaB webserver
- In production at UCL, Oxford, and half of Manchester Vac nodes

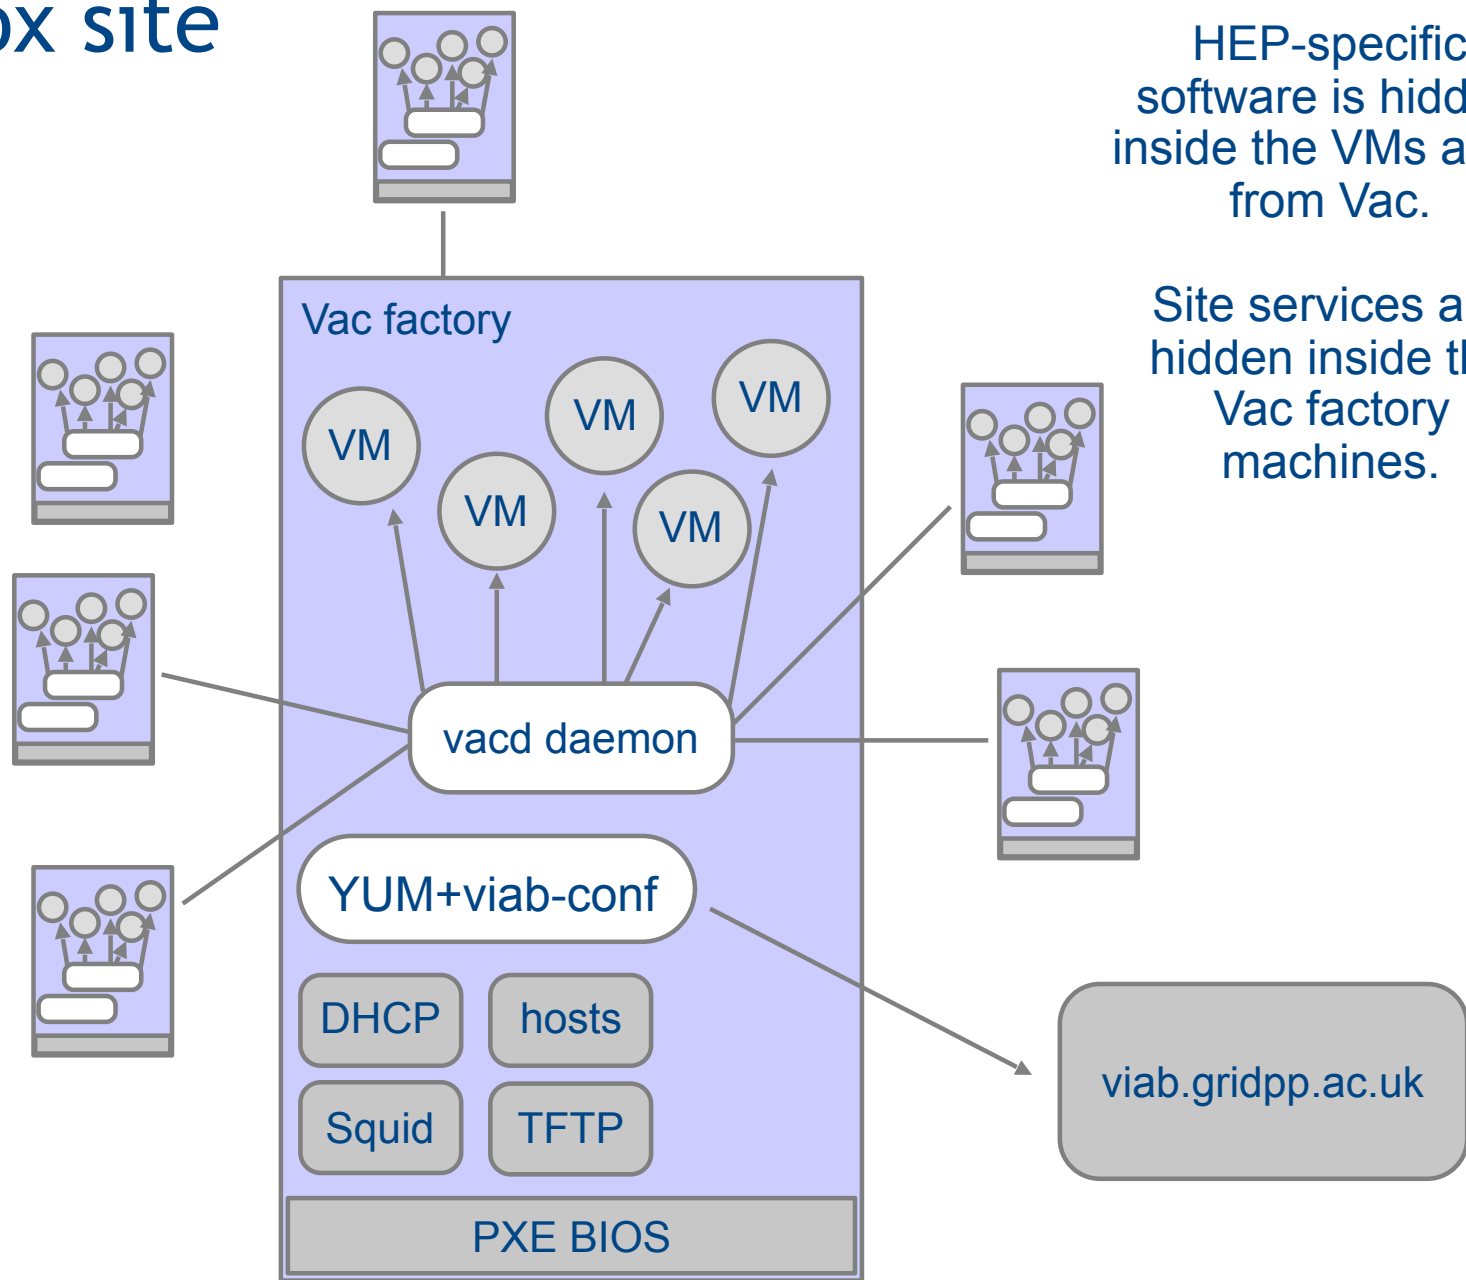
# Vac-in-a-Box site

Simpler than installing via Puppet etc.

Per-site dashboard at [viab.gridpp.ac.uk](http://viab.gridpp.ac.uk)

Kickstart from the website.

viab-conf RPM with configuration, via autoupdates from YUM repo.



HEP-specific software is hidden inside the VMs apart from Vac.

Site services are hidden inside the Vac factory machines.

# Vac-in-a-Box dashboard

viab.gridpp.ac.uk/admin/UKI-NORTHGRID-MAN-HEP

Vac-in-a-Box Sites admin Docs

All Sites / UKI-NORTHGRID-MAN-HEP

## Site UKI-NORTHGRID-MAN-HEP

### Spaces

Space	USB .iso	RPM published
testspace	-	Never
vac04.tier2.hep.manchester.ac.uk	Download	2015-08-20 16:15:01

**Add a space**

Space names should be in the DNS namespace controlled by the site, but they do not need to be registered in its name servers.

### SSH keys

Key	Type	Comment	Added
AAAAB3NzaC1yc2EAAAABIwAAAIEAuFxxq0w1gPEN Oxj6Uj4PhzomdVfJyBvWP9z8bWTYarErvqLQIZpU eBFW8sM+k/nnugUhYIn59nJHsZk7GhTdicZJ4YxJ F6mM3NMqisjYfuUdQXchTcKyy0yCdXv/P2xygvx0 vBrWROMYNLaTt/TdBeZQVC/JbWcJchrUSbpqec=	ssh- rsa	mcnab	2015- 08-08 22:18:45

### Add an RSA ssh key

The ssh keys will be installed on Vac factory machines to allow ssh access as root

Key:  Comment:

viab.gridpp.ac.uk/admin/UKI-NORTHGRID-MAN-HEP

Oxj6Uj4PhzomdVfJyBvWP9z8bWTYarErvqLQIZpU eBFW8sM+k/nnugUhYIn59nJHsZk7GhTdicZJ4YxJ F6mM3NMqisjYfuUdQXchTcKyy0yCdXv/P2xygvx0 vBrWROMYNLaTt/TdBeZQVC/JbWcJchrUSbpqec=	ssh- rsa	mcnab	2015- 08-08 22:18:45	<input type="checkbox"/>
---	-------------	-------	----------------------------	--------------------------

### Add an RSA ssh key

The ssh keys will be installed on Vac factory machines to allow ssh access as root

Key:  Comment:

### APEL certificate/key .p12 file

Uploading a valid cert/key will cause APEL accounting reports to be sent. The sitename UKI-NORTHGRID-MAN-HEP will be used when reporting to APEL.

.p12 file 2885 bytes, updated 2015-08-13 12:47:25

### Upload .p12 file

no file selected

### Site Admins

People with Vac-in-a-Box website admin rights are also able to update the site configuration.

X.509 DN	Added
/CN=Test Name	2015-08-20 15:50:42

### Add a site admin X.509 DN

X.509 DN:

© GridPP 2013-2015



## Next steps

- Exploit existing Vac to ElasticSearch reporting for Ganglia style monitoring
- Add VM size mixing: currently only one size of VM allowed per hypervisor (1, 8, 4, ...)
- “Vacuum Pipes” so new VOs can be added to sites with a single URL
- Manage mix of VMs and Containers on Vac hypervisors
- Continue to improve experiment/VO VMs, and:
  - Generic DIRAC VMs
  - HTCondor VMs for integration into local batch farms
- More sites ...



# Summary

- Vac and Vcycle both on a firm footing
  - Vac now at 01.00 release stage
- Vacuum Platform specification
- Clear progress/plans in getting VMs for LHCb, GridPP DIRAC, ATLAS, and CMS on to a production basis
- More things in the pipeline:
  - VacMon, Vacuum Pipes, VM size mixing, Containers

# New ATLAS VMs in action

- Using new VMs: one processor per VM

