



Vac and Vcycle status and plans

Andrew McNab

University of Manchester,
GridPP, and LHCb



Overview

- Vac vs Vcycle
- VacMon, Pipes, Multiprocessor
- Deployment status
- Docker containers in Vac
- Google Compute Engine



Vac vs Vcycle recap

- 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
 - Factory installation by Puppet etc or Vac-in-a-Box
- Vcycle - uses OpenStack, EC2, Google Cloud etc
 - VMs created via Cloud API in response to observed demand for each type of VM
 - Same VM definitions as Vac
- VMs are self-contained black boxes defined by experiments
 - Know how to pull in jobs to run from experiment HQ



New in the last six months

- VacMon - Ganglia-style monitoring at site, space, VM factory level
- Vac 2.0 deployed
 - Multiple VM sizes on the same VM factory: eg 8 and 1
 - Vacuum Pipes to reduce VO configuration to a URL
- New VM definitions, including ALICE VMs
 - Enabled Birmingham to start converting worker nodes to Vac
- VMCondor framework in production for ATLAS and ALICE, and available for generic VMs running HTCondor jobs
 - Should also work for CMS
- Google Compute Engine plugin for Vcycle


Deployment by site and experiment

		ATLAS	ALICE	LHCb	GridPP DIRAC
Vac	Birmingham	✓	✓	✓	
	Glasgow	✓		✓	✓
	Liverpool	✓	✓	✓	✓
	Manchester	✓	✓	✓	✓
	Oxford	✓	✓	✓	✓
	UCL	✓		✓	✓
Vcycle	Imperial			✓	✓
	CERN (LHCb)			✓	
	CERN (Dev)	✓	✓	✓	✓
	CC-IN2P3			✓	
	Yandex			✓	
	Datacentred			✓	✓
	INFN Naples				

Belle II

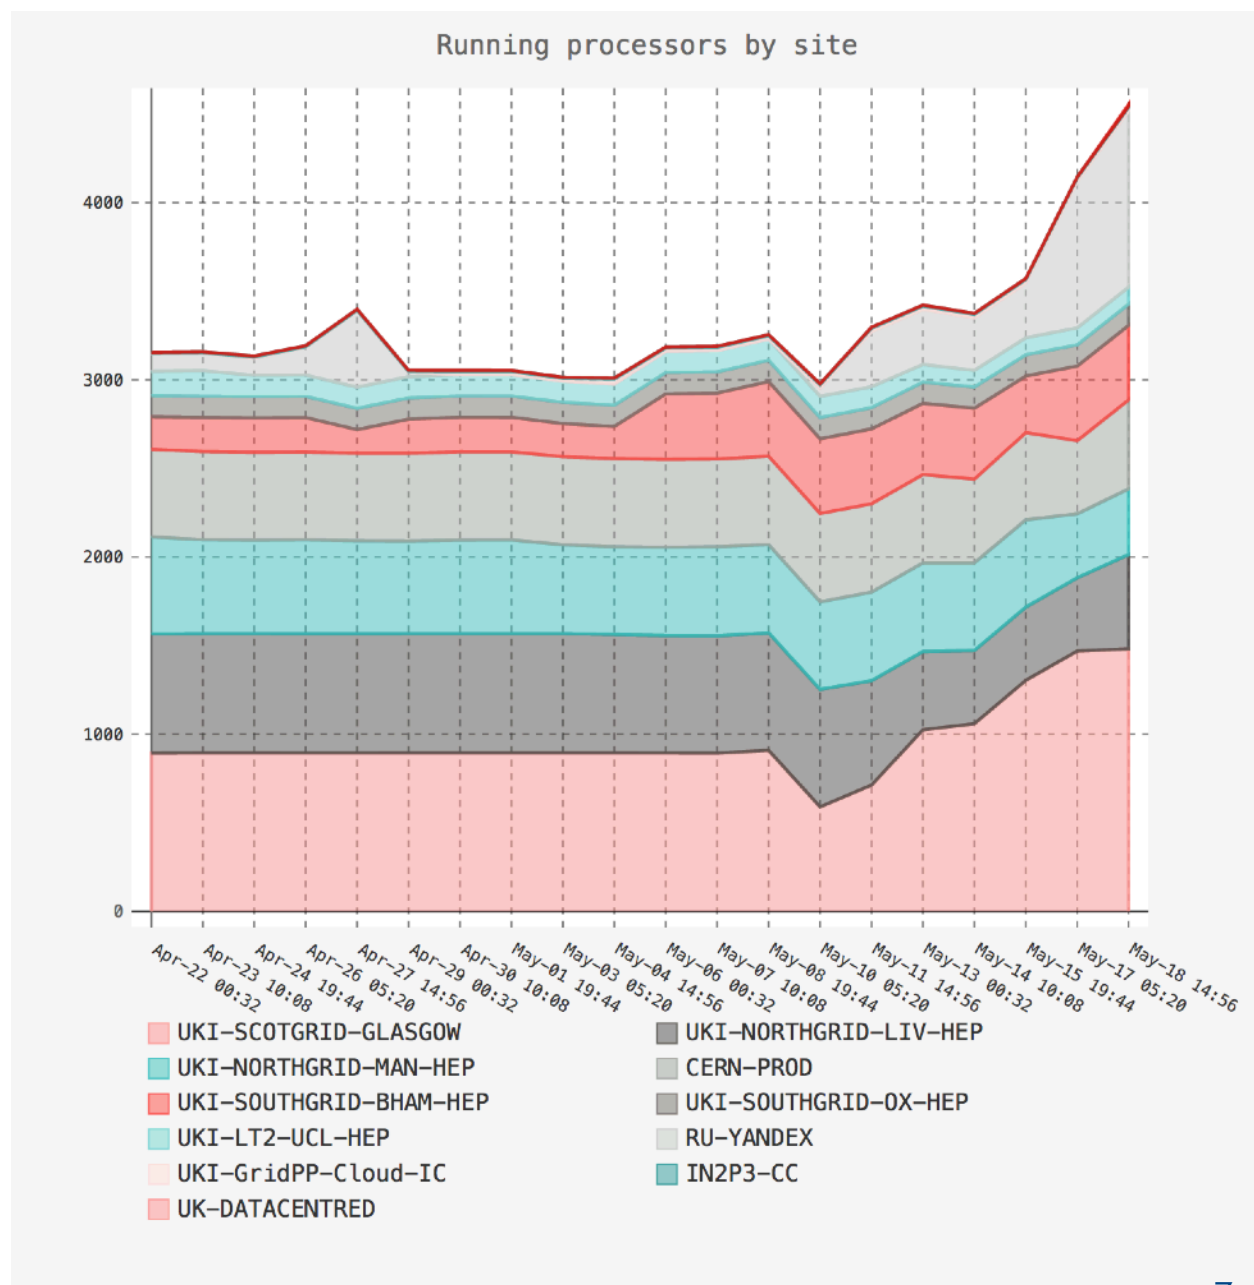
VacMon

- Ganglia-style monitoring at site, space, VM factory level
- Produces charts like this
- Uses Vac's internal JSON status message formats
- Sent over UDP to vacmon.gridpp.ac.uk
- Stored in ElasticSearch




Site status

- Glasgow has pushed Vac up to ~1400 processors as part of the acceptance tests of their new farm
- Yandex is moving resources from DIRAC.Yandex.ru to OpenStack, which we manage with Vcycle
- Other sites coming: Cambridge + RHUL

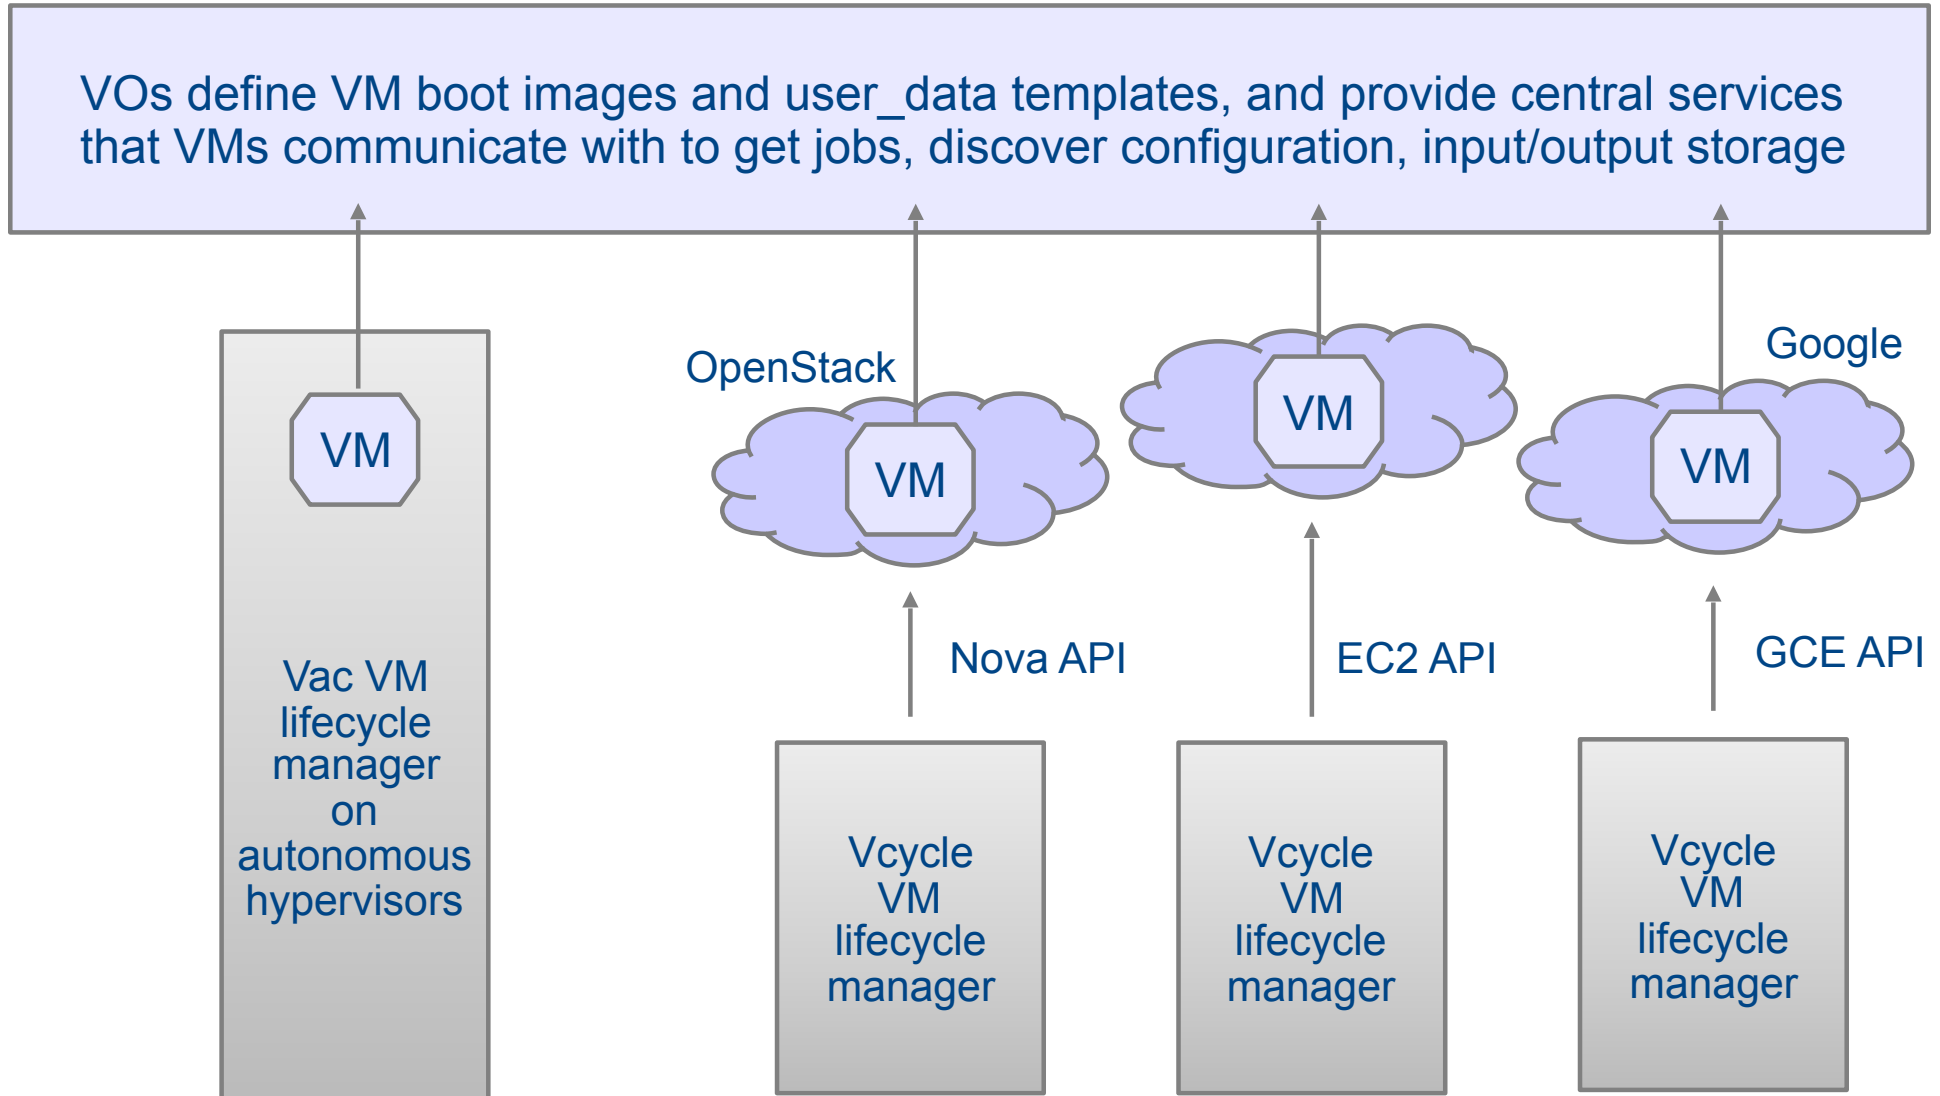


Use by VOs

- Still mainly LHCb
- Yandex and LHCb CERN Cloud are only LHCb
- Significant ATLAS fraction too at Vac sites
- ALICE notable and growing since 2016
- Smaller VOs via GridPP DIRAC service



Vacuum platform



Docker Containers

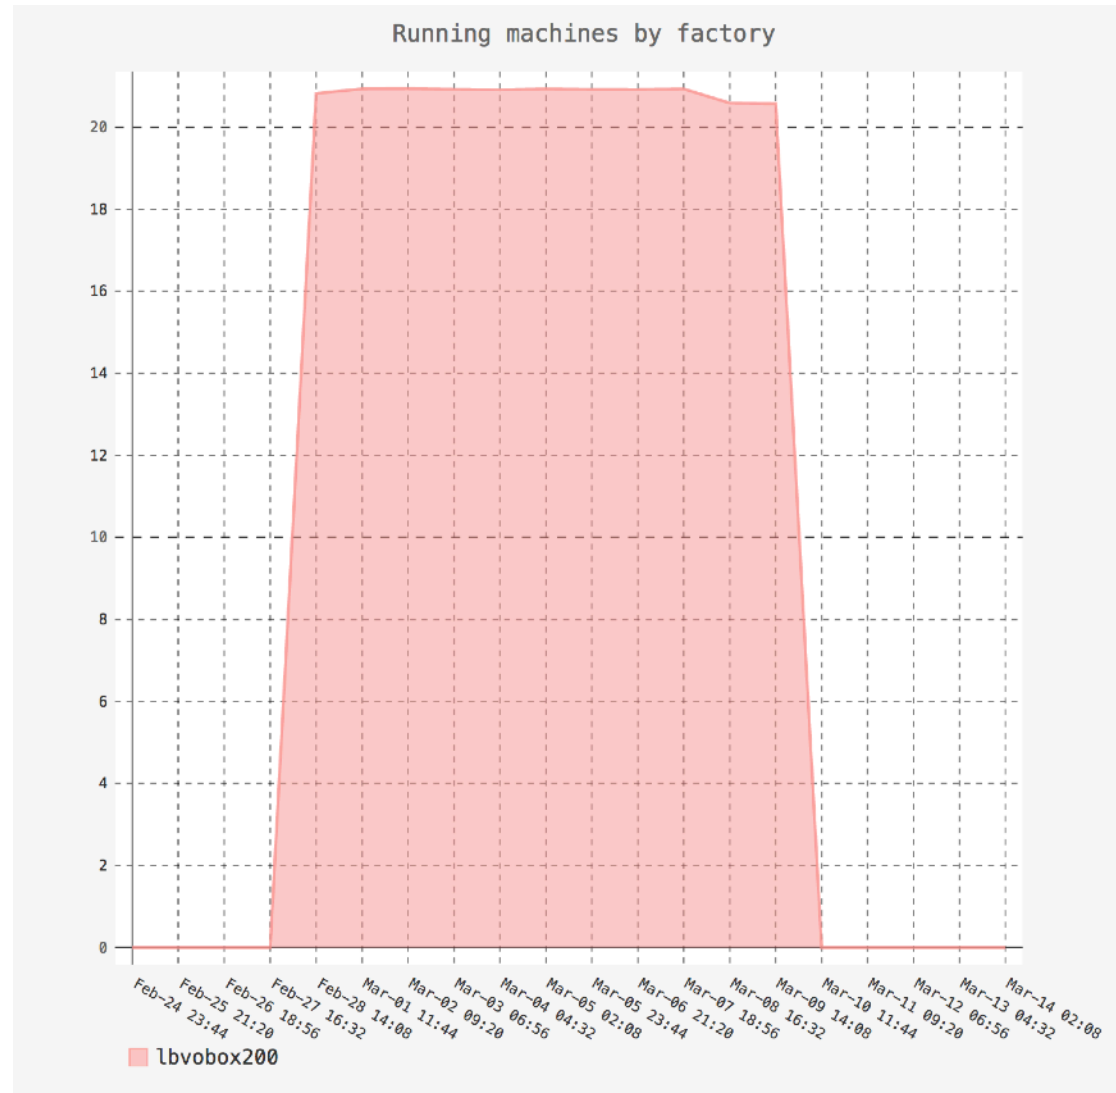
- Current Vac development is to add support for Docker Containers as another logical machine model alongside VMs
- Will be able to run arbitrary Docker images, or Vacuum Containers
 - Extending Vacuum Platform API to define how to provide CernVM-FS to unprivileged containers, init script as a volume etc
- So Vac factories will be able to run a mix of VMs and DCs alongside each other, using target share mechanism etc to decide what to start next
- Using LHCb container definition first, but will extend VMCondor framework too (so available to ATLAS + ALICE)

Google Compute Engine

- Test of running LHCb MC jobs on Google Compute Engine in the US
- Using standard LHCb VM definition, managed automatically by Vcycle as we do for CERN, IN2P3, YANDEX, DataCentred
 - GridPP wrote a Google plugin for Vcycle (took ~week including testing)
- Using all three US regions: us-east1, us-west1, us-central1
- One Squid cache in each region, managed by LHCb
 - All peered over Google internal network so they attempt to resolve cache misses from their peers
 - Only then do they get files from outside (usually from cvm.fnal.gov via Google peering with ESNet)
 - CernVM-FS in VMs can use any of the three caches, again over Google internal network
- Run with ~20 VMs in parallel
- (Retail pricing is ~\$25/month/VM, and $60,000 * 25 * 12$ is 18 million)

Google Compute Engine

- 20 VMs and jobs, in the three regions
 - Plus squid cache VMs
- This chart shows the production run once development and testing was finished
- Stable and ran unattended from the LHCb production Vcycle instance
 - lbvobox200
- Ready to repeat at larger scale if Google resources become available





Summary and next steps

- Vac 2.0 deployed
 - Better multiprocessor support and Pipes
- VacMon monitoring website
- Significant increase in resources
 - Across multiple sites
- Docker containers being added for Vac 2.1
- Google Compute Engine support in Vcycle

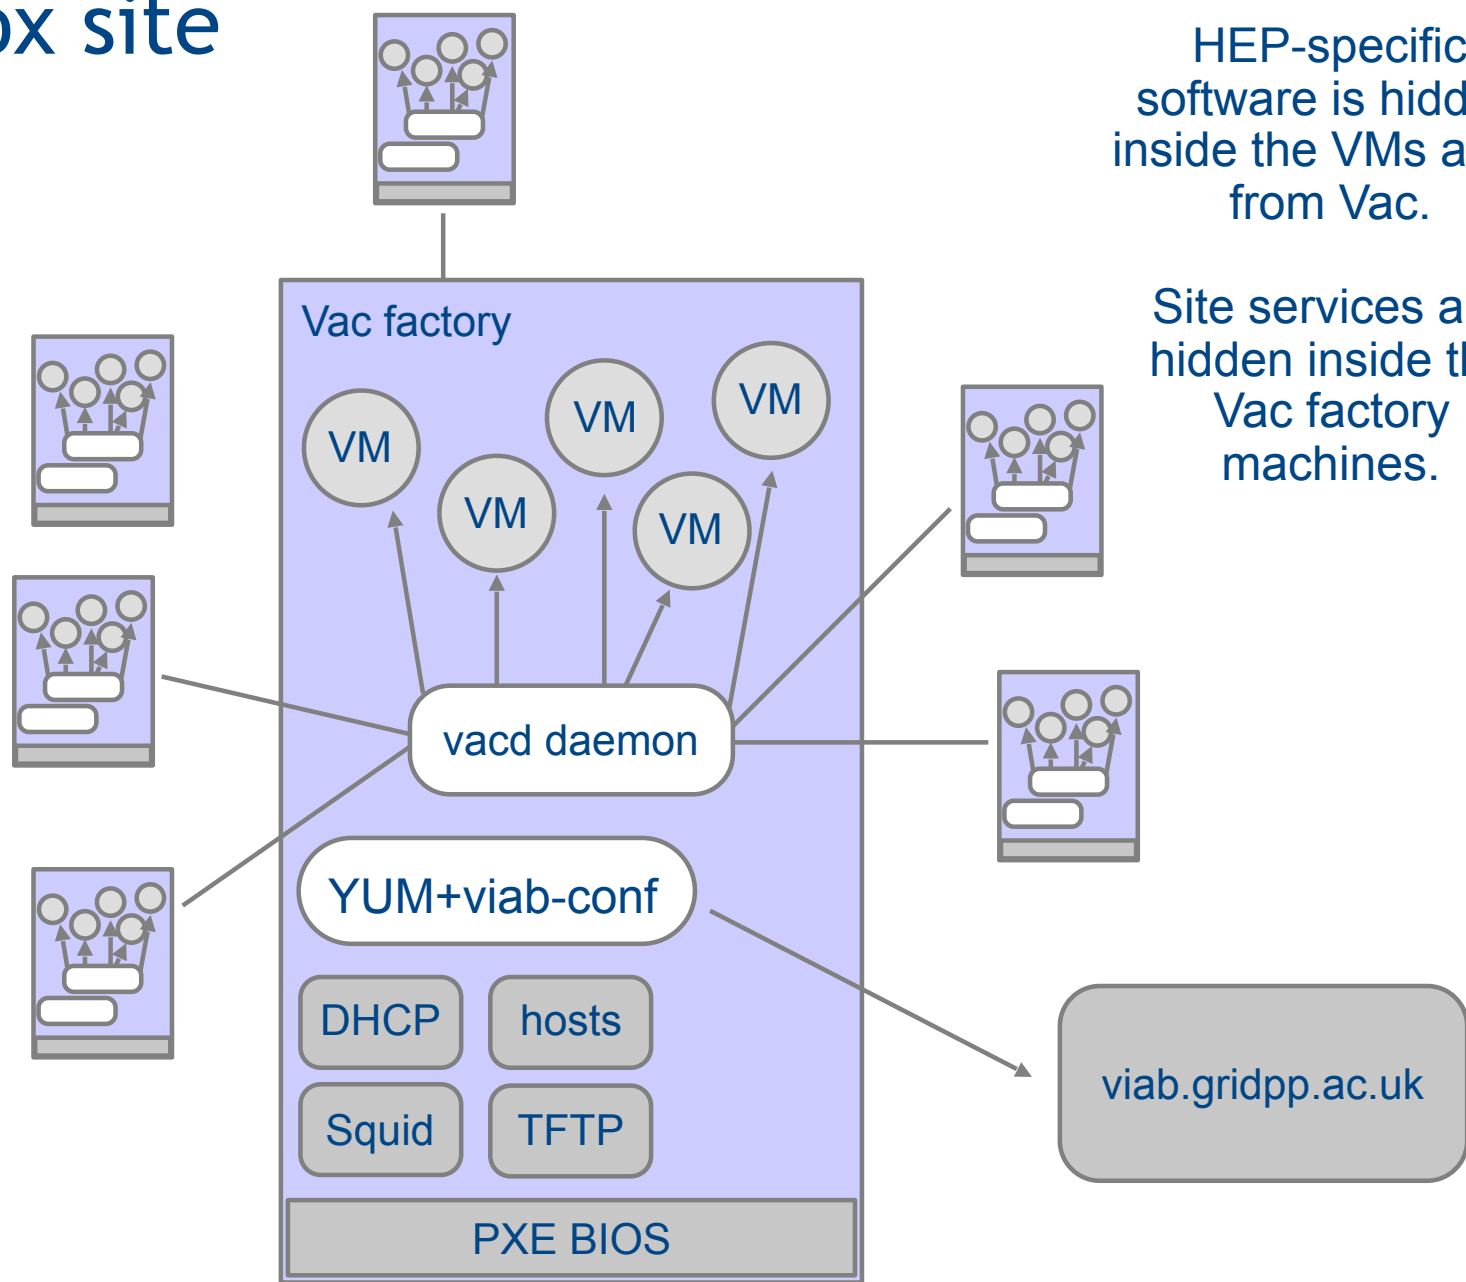
Vac-in-a-Box site

Simpler than installing via Puppet, Ansible etc.

Per-site dashboard at 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 vBrIWROMYNLaTt/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 vBrIWROMYNLaTt/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