



Manchester status and plans

Andrew McNab
University of
Manchester
LHCb and GridPP

Context: Blakett Cluster

- Particle Physics Group (PPG) has run a grid service for High Throughput Computing (HTC) since 2001
- Blakett Cluster currently provides two services:
 - A WLCG Tier-2, funded through GridPP, mainly used by ATLAS and LHCb but also ALICE, skatelescope.eu, ILC, LSST, Icecube and about ten others
 - Manchester Analysis Facility for local users, funded by STFC consolidated and project grants, and University money
- Split between off-campus space and Schuster Laboratory
 - Several generations of Intel CPU, as bare metal and virtual machines: ~6000 processors
 - 3 PB of storage with GridFTP, xroot and WebDav interfaces
 - A small subcluster with K40 and V100 Nvidia GPUs

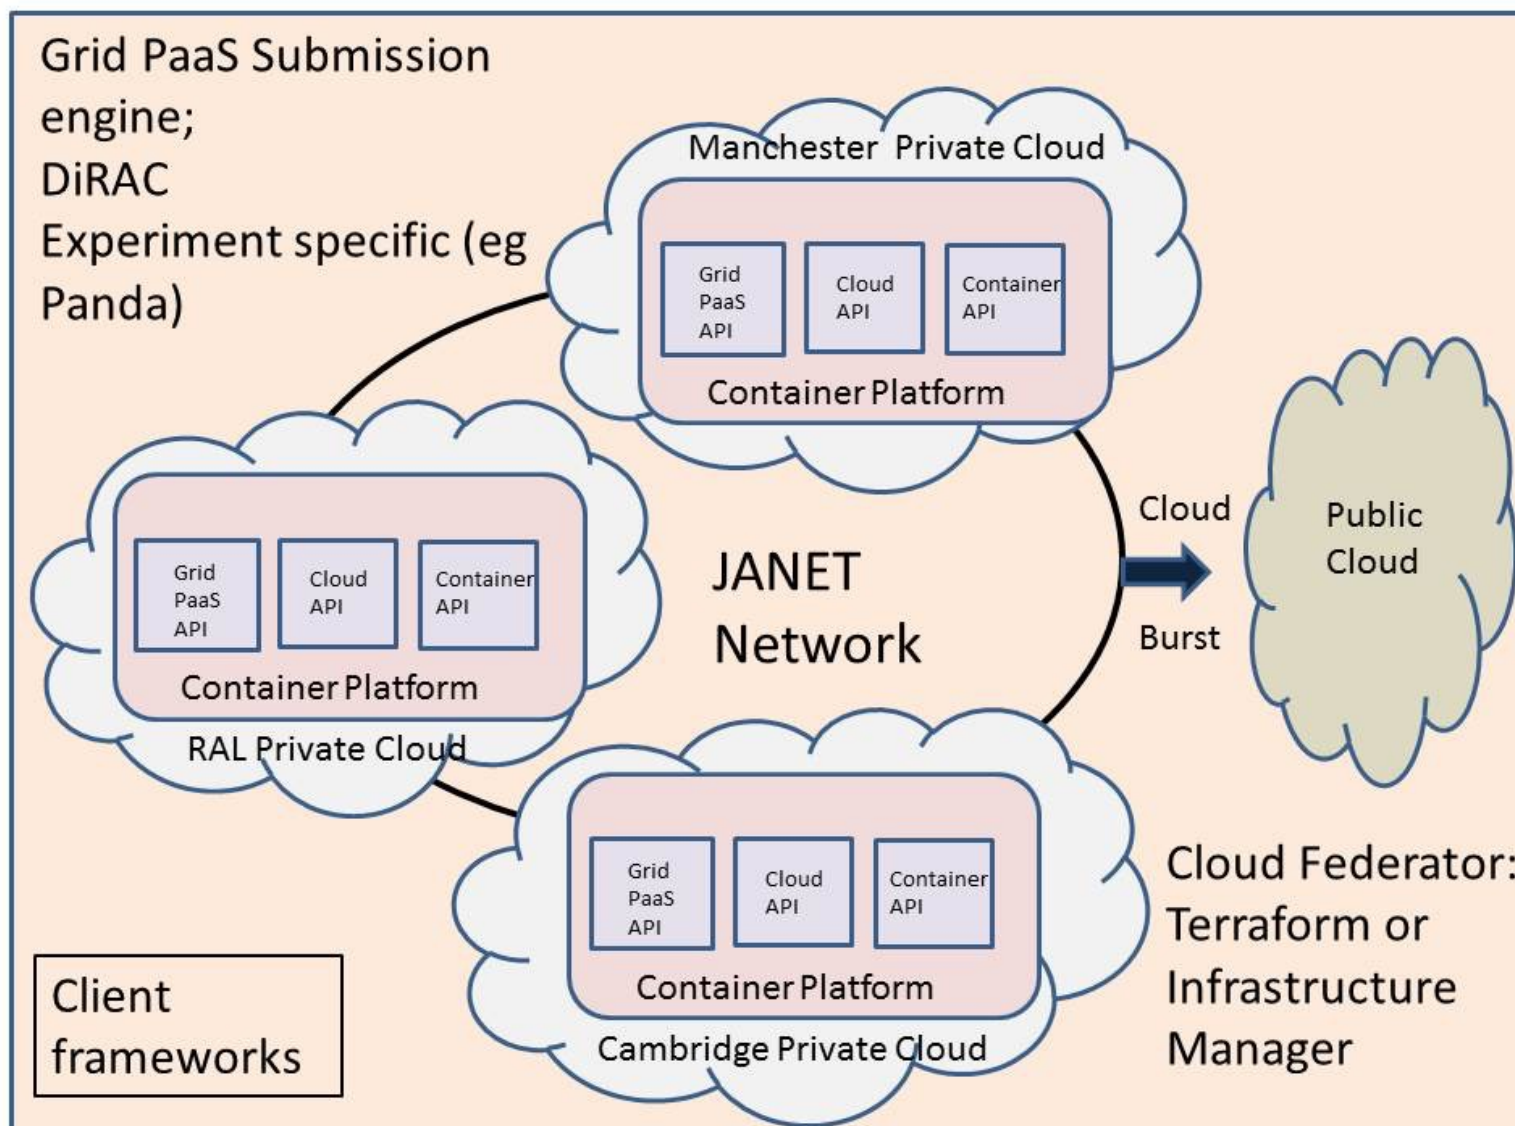
Context: other services

- Particle Physics Group
 - Provide DNS, web server and email for GridPP
 - Develops the Vac/Vcycle VM management systems and hosts associated repositories and services (eg VacMon monitoring service)
 - Hosts and manages the GridPP VOMS service, including many non-PP user communities (eg skatelescope.eu)
 - Has a development OpenStack service as part of a CASE PhD
- Jodrell Bank Centre for Astrophysics has several HPC and HTC clusters, with uses including:
 - Specific types of analysis
 - The operation of the e-Merlin network of telescopes
 - Participation in European VLBI network

2017 UKT0 funding

- £350,000 from the 2017 capital grant
- Used a tender process within the Northwest Universities Framework - fast and straightforward to use
 - Asked for 1 PB storage and as much CPU as possible
- Placed order with Dell
 - CPU: 48 C6420 @ 64 logical processors, 192GB memory (3GB/processor) giving 3072 processors (1536 cores vs 1250 target)
 - Storage: 10 R540 @ 12 10TB disks to provide 1PB usable target
- All items received ok
- All machines now installed in racks (including removing some older machines to make space)
- Next: finish first configuration / installation; start running jobs; start using storage on a provisional basis, subject to future RRB

CPU in the e-Infrastructure bid

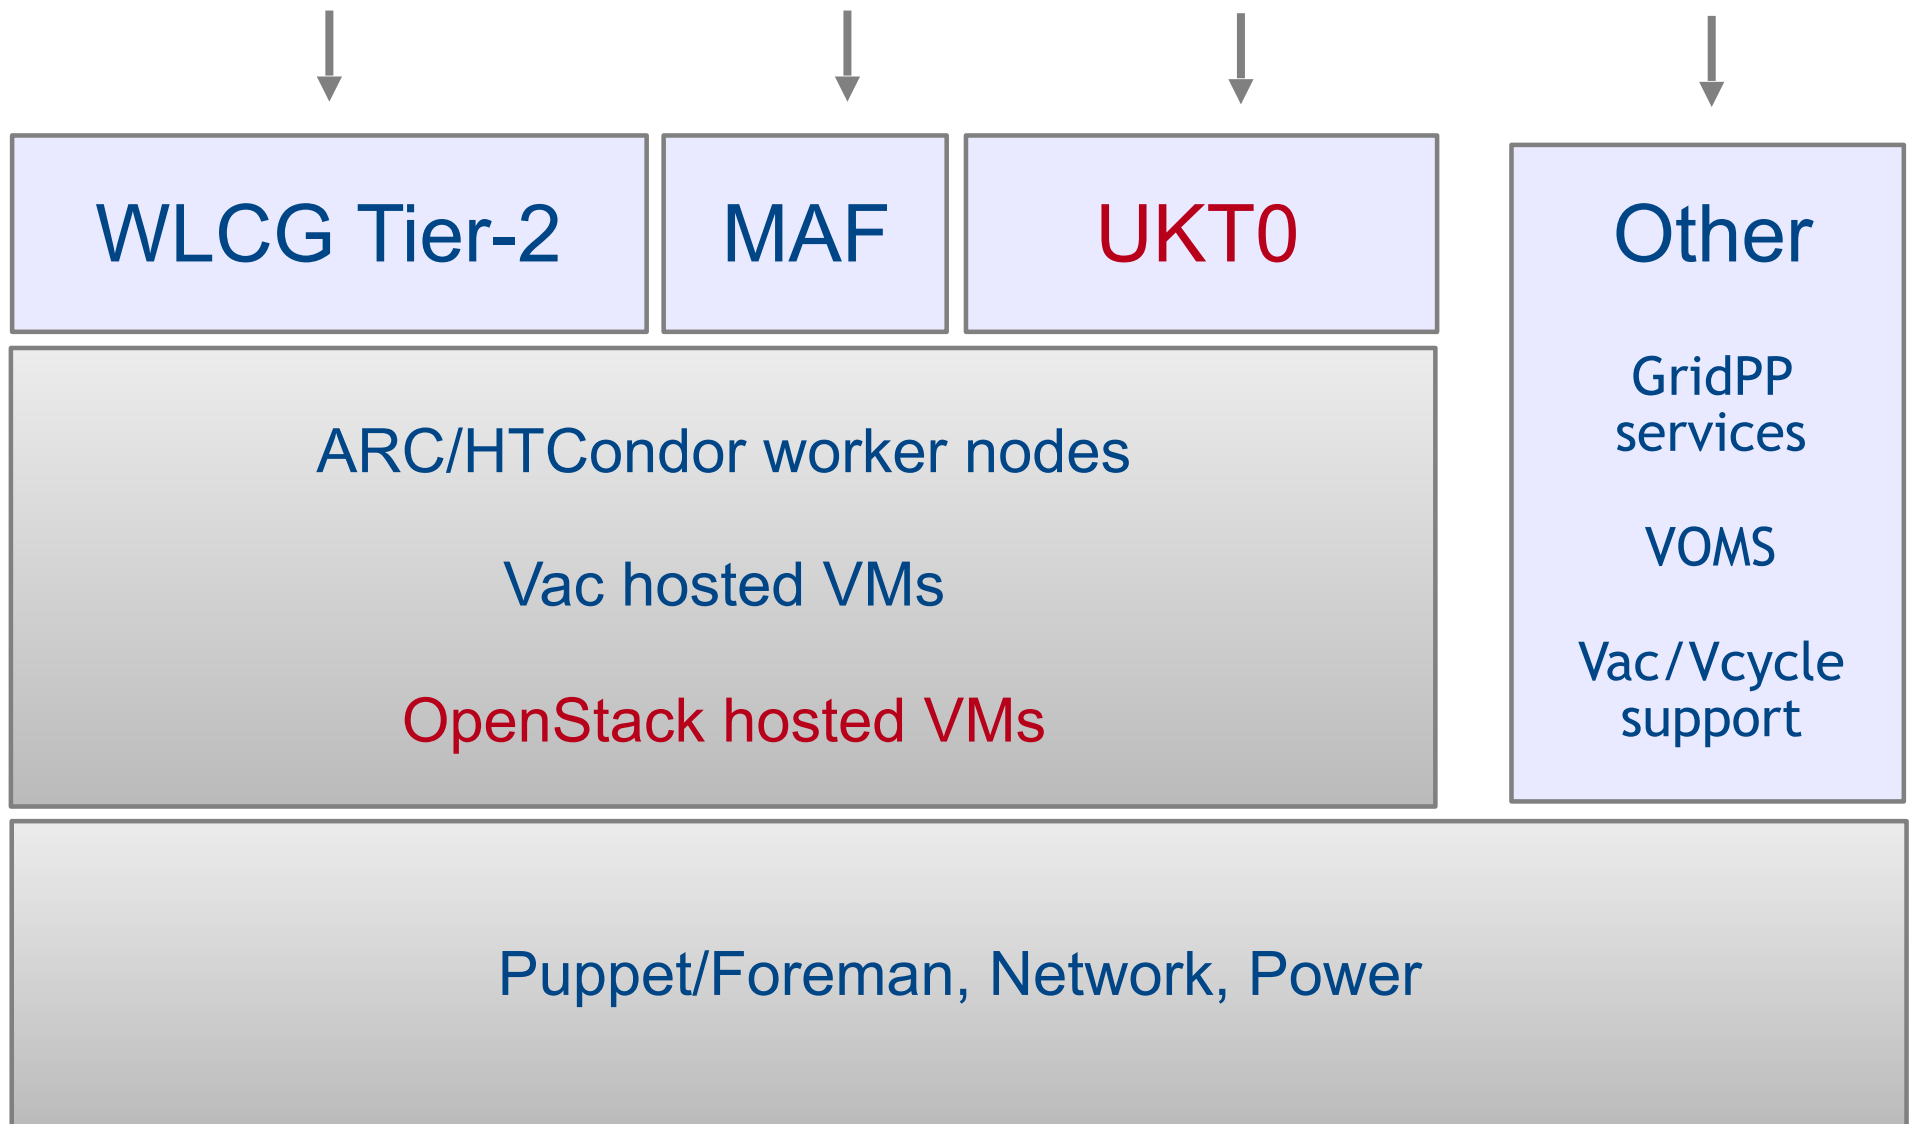




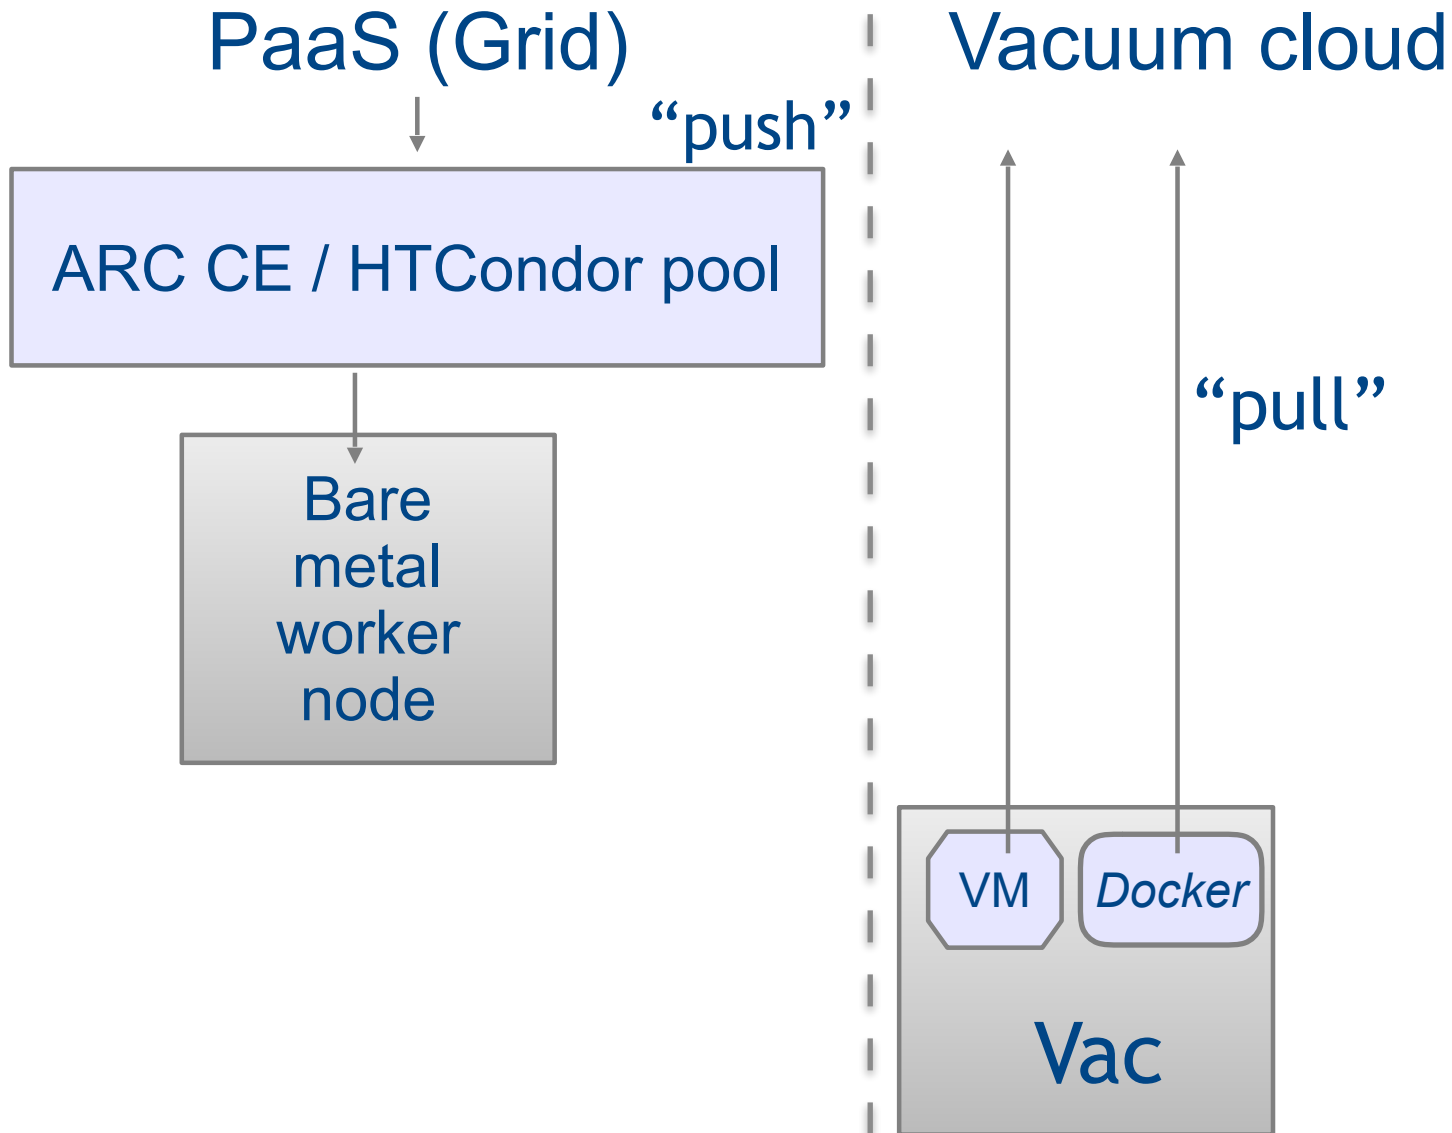
Aims of Manchester UKT0 service

- Support user communities across STFC science, not just LHC
- Provide a (Docker) container API
 - Use Vac's Docker support for this
- Provide OpenStack API for creating VMs, including teams of VMs that work together
- Be able to provide much larger slots than current single processor and eight processor HTCondor job slots and Vac VMs
 - Always 2-3 GB of physical memory per processor
 - Requirements e.g. for SKA-AENEAS to have 10s of GB for a job
 - Straightforward with Vac VMs/Containers, and possible to do efficiently with OpenStack VMs
- Work on controlling the effort required to operate services, and share experiences with other sites

Blakett cluster and its (new) services



Blackett cluster: now



Blakett cluster: with UKT0 service

PaaS (Grid)

“push”

ARC CE / HTCondor pool

Bare metal worker node

VM

OpenStack

VM

Docker

Vac

Vacuum cloud

“pull”

VM

Docker

Vac

VM

Open Stack

Vcycle

IaaS Cloud

“push”

Keystone + Nova

VM

OpenStack



Summary

- Manchester already supports user communities outside PP
 - UKT0 funding allows us to formalise that
- We've achieved the storage numbers in the 2017 bid and significantly surpassed the CPU figure
- Machines are racked and about to be put into service
- We're deploying new platforms (OpenStack, Containers) in addition to the extra capacity
- No significant development required for the baseline presented in this talk
 - Deployment of hardware
 - Integration of existing software components with the underlying infrastructure (Puppet etc)



Extras



OpenStack PhD Project

- We have a current PhD student with a CASE project to do OpenStack, Vcycle, DIRAC development
- The project aims to improve OpenStack's ability to deal with “eviction” scenarios where lower priority work is stopped in favour of higher priority users
- Also extensions to Vcycle to manage this
- And to (LHCb) DIRAC to handle jobs being asked to stop
- Different time scales: no notice, a few seconds (log what happened), a few minutes (finish processing and save cleanly)
- We intend to apply the products of the project to OpenStack at Manchester (and use UKT0 experience as input of course.)
- It will allow us to handle switchover from backfilling LHC jobs keeping the machines busy to new user communities
- Also more efficient provisioning of large geometry VMs