

LHCb experience running jobs in virtual machines

Andrew McNab,
University of Manchester

Federico Stagni
& Cinzia Luzzi,
CERN

on behalf of the LHCb collaboration



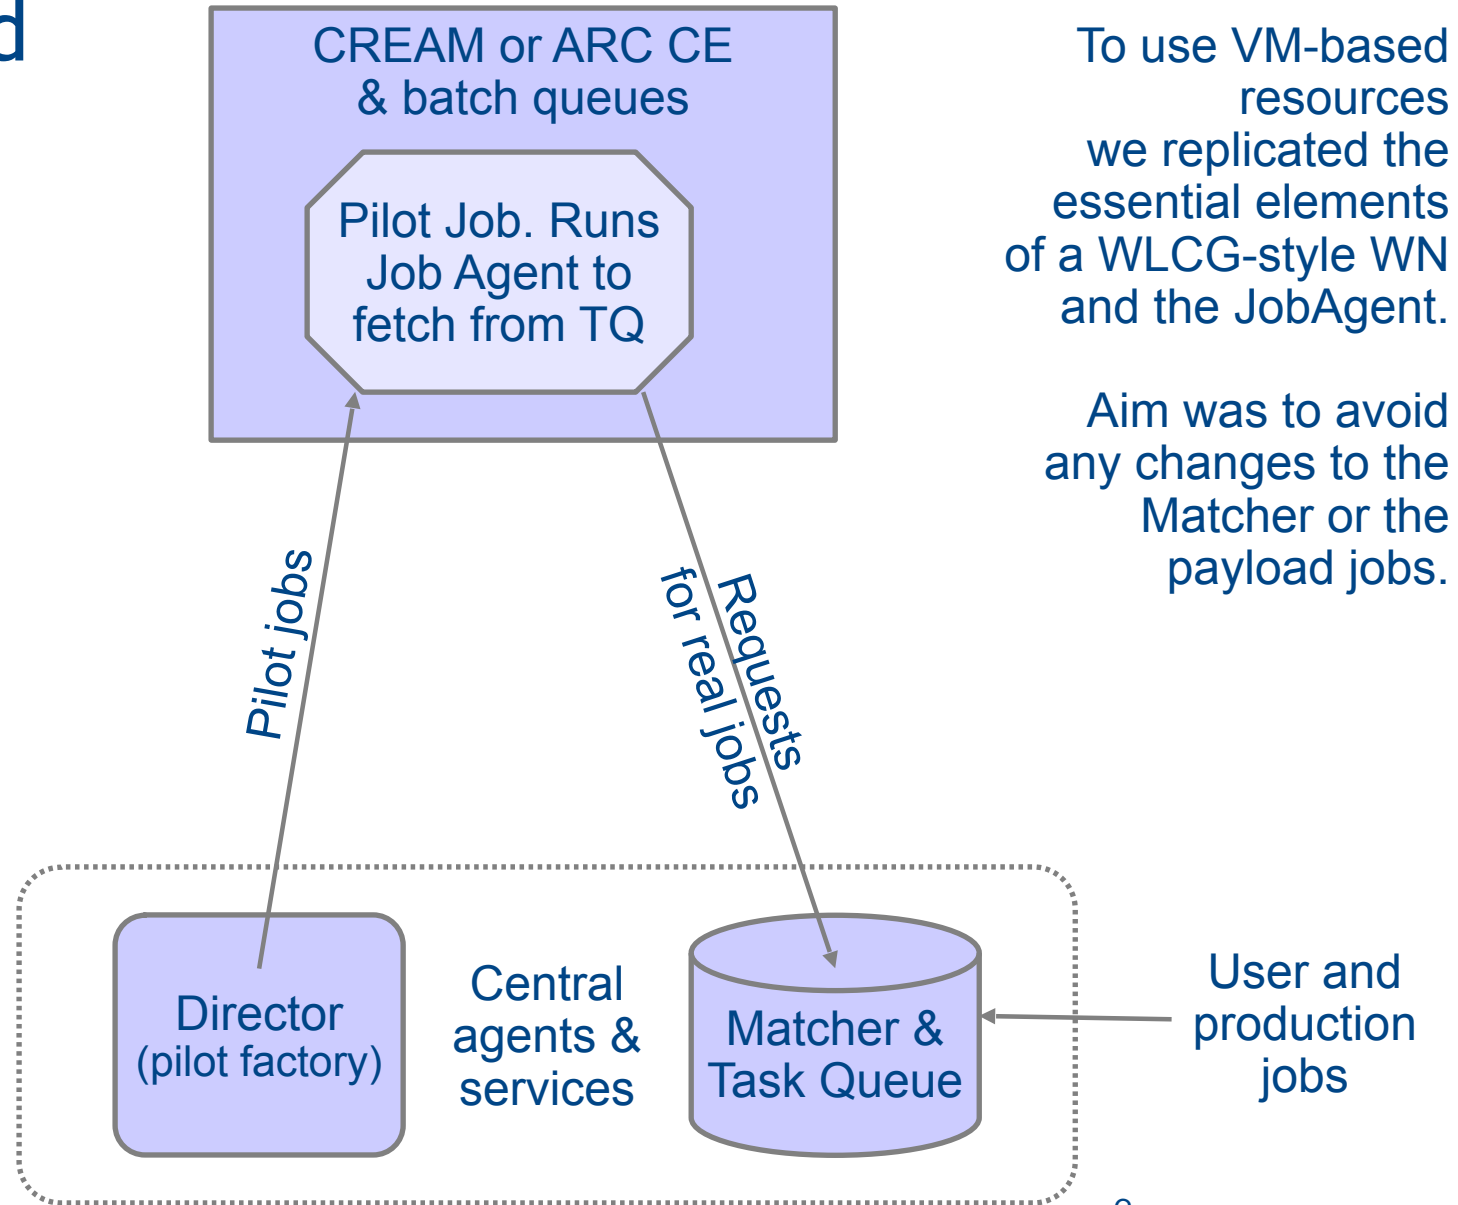
Overview

- Starting from DIRAC + Grid
- CernVM as a WLCG worker node
- Pilot VMs
- Machine/Job Features
- Syslog and ssh
- Sites and platforms
- Plots
- Observations
- Summary

LHCb DIRAC and the Grid

Most LHCb jobs are run on CREAM or ARC sites, with batch queues.

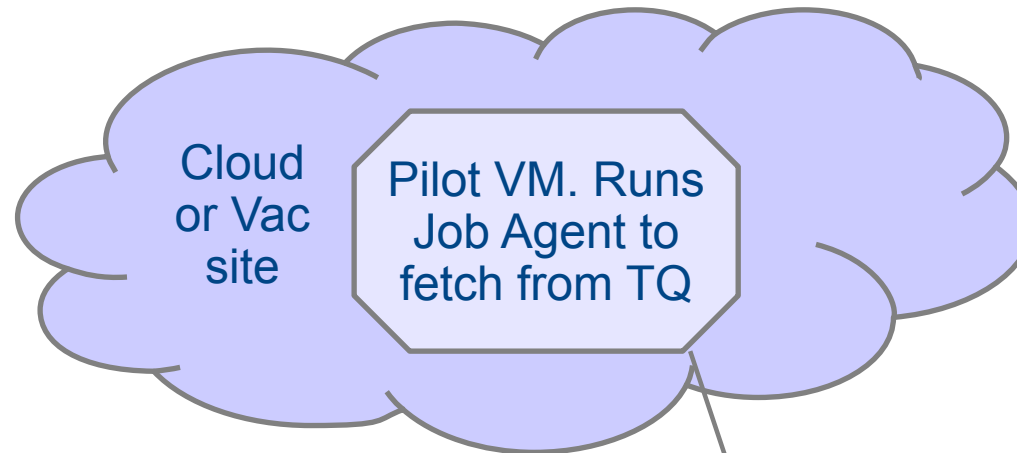
The core of this is the Job Agent on the Worker Node asking the Matcher for a payload job submitted by a user.



LHCb DIRAC and VMs

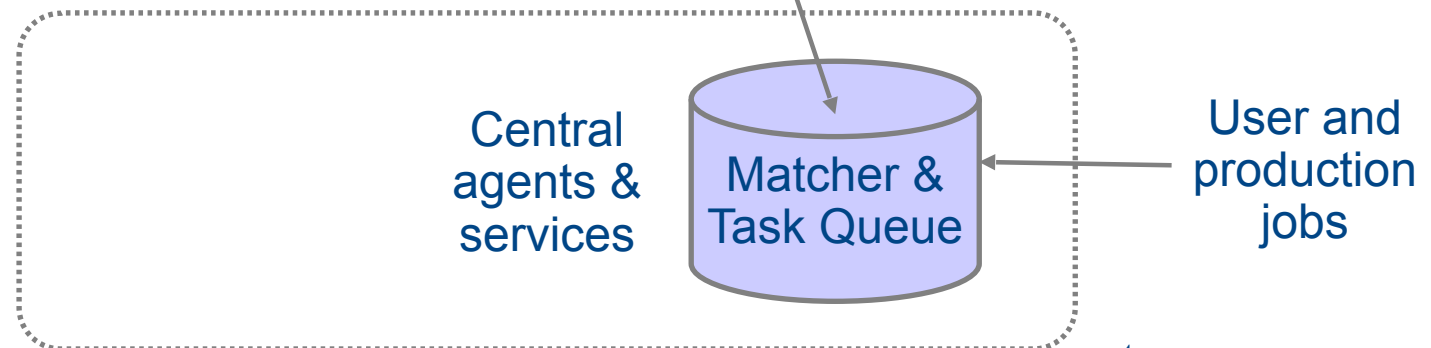
VMs are created by a VM lifecycle manager like Vac, Vcycle or HTCondor Vacuum.

Details of this don't matter once the VM is configured and ready to start DIRAC client.



To use VM-based resources we replicated the essential elements of a WLCG-style WN and the JobAgent.

Aim was to avoid any changes to the Matcher or the payload jobs.



CernVM as a WLCG worker node

- WLCG sites typically install and configure WNs with RPMs, Puppet, YAIM, local scripts
 - Frequently leads to job failures when an RPM is missing from WN
- VMs can allow us to have complete uniformity on WNs
- We use vm-bootstrap script, run by contextualization to set things up
- /cvmfs/grid.cern.ch is used to provide middleware
- Create /scratch from a 40GB+ partition provided by the hypervisor
 - Much faster than copy-on-write overlay root partition
- Increase the cvmfs cache to 11GB and put it on /scratch partition
- Create swap file on /scratch partition
- Set up /etc/machinefeatures, /etc/jobfeatures, Ganglia, etc
- Run the JobAgent to get the payload job

LHCb Pilot VMs

- Have a well-defined lifecycle compatible with Vac, Vcycle, and HTCondor Vacuum lifecycle managers
 - VM does shutdown -h when finished or if no more work available
 - shutdown_message reports why VM finished
- A heartbeat file is update every five minutes so zombie VMs can be killed
- We supply boot image and user_data template for contextualisation from our web server - new versions picked up automatically by Vac etc.
- Uses hostcert/key in contextualization to authenticate with DIRAC
- System setup done as root, then unix accounts used to partition VM
 - sudo dirac to run Pilot 2.0 client (Job Agent) now used on all platforms (See Federico's Track 4 talk, ID 113, from yesterday)
 - sudo diracuser to run the payload job
- Currently runs one payload per VM lifetime
 - Adding support for multiprocessor VMs with multiple concurrent and consecutive payloads

Machine/Job Features

- Proposed in HEPiX and by current (finishing) WLCG task force
- Allows site/host to communicate details of machine and the job slot to the job or VM
 - HS06, shutdown time for VM/host, CPU and memory limits, ...
- We obtain location of MJF files from contextualization (Vac, HTCondor Vacuum) or OpenStack metadata (Vcycle)
- Jobs use variables pointing to directories in the (NFS) filesystem or to URLs on a webserver
- VMs try these methods in turn so can set \$MACHINEFEATURES and \$JOBFEATURES for the jobs
- Further integration with LHCb elastic jobs (See Federico's talk, 112, on this in Track 4 yesterday)
 - Allow sites to better control VM lifetime
 - Deal with Masonry Problem of multiple payload jobs in a VM

Syslog and ssh support

- Done as part of WLCG Traceability Task Force
 - Address concerns from sites about logging and access to VMs
 - Included in LHCb VMs
- Check for `/etc/machinefeatures/syslog`
 - If it exists, use it to create configuration file `/etc/rsyslog.d/vm.conf`
 - Sends a copy of syslog messages to an external rsyslogd
- We intend to support whatever mechanism the task force recommends
- CernVM can supply authorized ssh public keys as part of contextualization
 - We don't stop this: we encourage it so site admins can help with debugging


Sites and platforms

- LHCb VMs run on all three Vacuum platforms at several sites
 - Vac
 - Manchester, Oxford, Lancaster
 - Vcycle + OpenStack
 - CERN, CC-IN2P3 (LHCb Vcycle instance at CERN)
 - Imperial College (GridPP Vcycle instance at Manchester)
 - HTCondor Vacuum
 - RAL
- As of this month, we are in discussion with three more Vac sites, and one more OpenStack site.
- The VMs are also the basis for LHCb's BOINC/@home work
- They were adapted to make the ATLAS, GridPP DIRAC, and CMS Pilot VMs

LHCb jobs in VMs

CPU used by Site

52 Weeks from Week 13 of 2014 to Week 13 of 2015



Routine production since May last year.

Increased capacity this year.

Periodic structure due to varying demand from other experiments and in availability of LHCb jobs



Observations

- We get a lot of uniformity from using VMs, and remaining problems tend to be associated with cvmfs use of site's squid cache(s)
- Our use of cvmfs is very sensitive to overloaded squids
 - Shows up as missing/corrupted scripts when starting jobs
 - Need to make sure resource provider is following site config properly
- For a new type of site, ssh access to VMs is very helpful
 - Logging to /etc/machineoutputs also invaluable
 - Need co-operation from resource provider at this stage
- If Vac or Vcycle is already working, tends to be trivial to add LHCb VMs to the site.
- Outcome: improving logging via LHCb-wide Pilot 2.0 to make it easier to debug problems remotely.
- Once the site is working, very little intervention required on either side.



Summary

- Adapted existing LHCb DIRAC framework to virtual machines
- Use CernVM as the basis of WLCG-style VMs
 - Particularly dependent on state of cvmfs squid caches
- Routine operation for production jobs
- Using resources comparable to a substantial Tier-2 at 7 sites
- Are able to make use of VM resources as they become available
- Straightforward to add LHCb to existing VM-based installations
- Very stable once set up



Extra slides

Example of Vac configuration

- Section of vac.conf used to enable LHCb VMs at Manchester
- They just need this and to create a hostcert/key.pem
- (vcycle.conf configuration is very similar)
- Compare what YAIM has to do to add a VO to a CE/Batch site

```
[vmtype lhcbprod]
vm_model = cernvm3
root_image = https://lhcbproject.web.cern.ch/lhcbproject/Operations/VM/cernvm3.iso
root_publickey = /root/.ssh/id_rsa.pub
backoff_seconds = 600
fizzle_seconds = 600
max_wallclock_seconds = 172800
log_machineoutputs = True
accounting_fqan=/lhcb/Role=NULL/Capability=NULL
heartbeat_file = vm-heartbeat
heartbeat_seconds = 600
user_data = https://lhcbproject.web.cern.ch/lhcbproject/Operations/VM/user_data
user_data_option_dirac_site = VAC.Manchester.uk
user_data_option_cvmfs_proxy = http://squid-cache.tier2.hep.manchester.ac.uk:3128
user_data_file_hostcert = hostcert.pem
user_data_file_hostkey = hostkey.pem
```


The Masonry Problem...



The Masonry Problem

