



Cloud use by LHCb

Andrew McNab
University of Manchester
LHCb and GridPP



Overview

LHCb experiment at CERN

What I mean by “cloud”

DIRAC job submission

Vcycle cloud management


Usage in 2015-2016

LHCb VMs, CernVM


LHCb plans

Summary

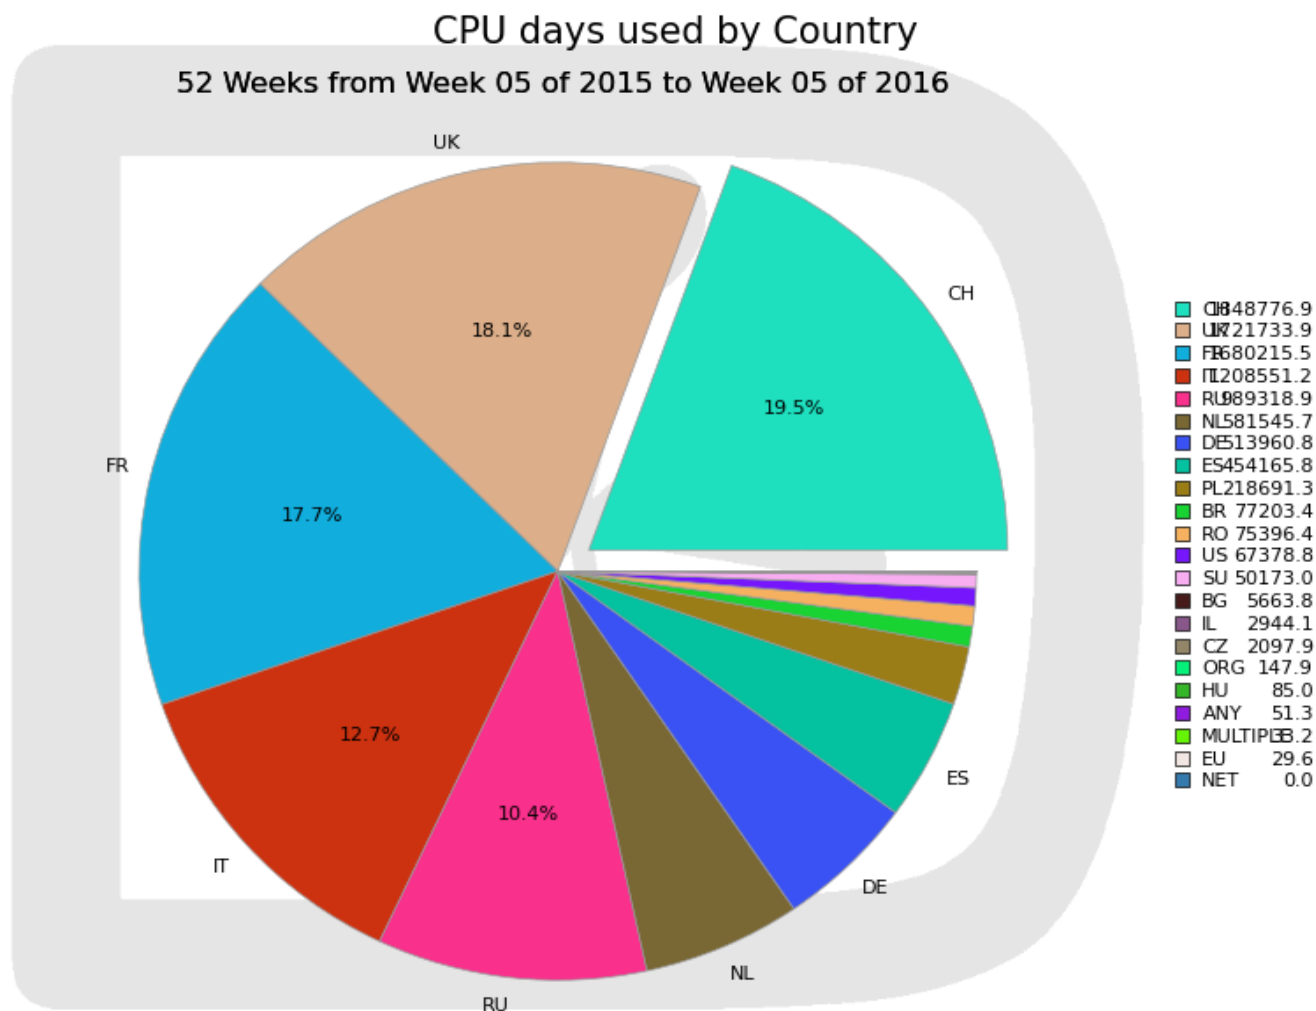
LHCb experiment at CERN



LHCb experiment at CERN



LHCb workload by country (mostly WLCG)



Generated on 2016-02-09 09:10:06 UTC

We need to use resources in the form presented by sites, now including Cloud



What do I mean by “Cloud”

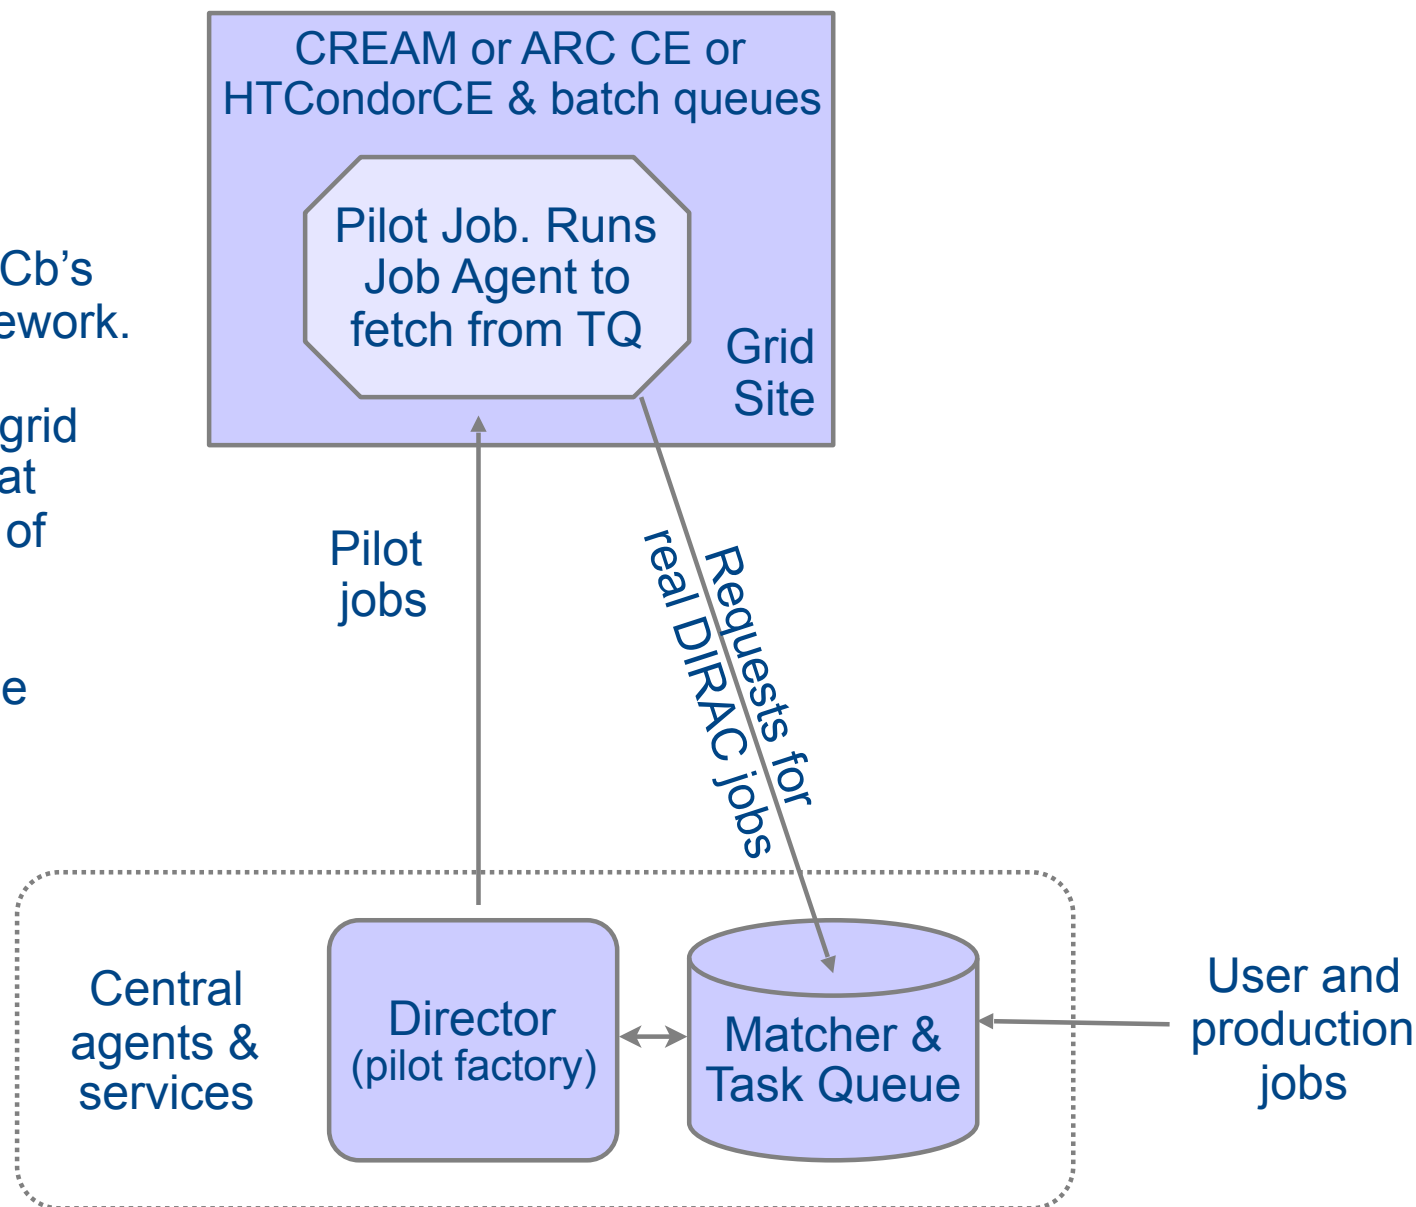
- In HEP we tend to say “Cloud” and just mean Infrastructure-as-a-Service (IaaS) Clouds
 - Arguably our Grid(s) are a kind of Platform-as-a-Service cloud
- IaaS Clouds give you a programmatic way to manage and use
 - Virtual Machines
 - Virtualised storage
 - Virtualised networking
- All this at a remote service
- “Cloud” terminology originally promoted by Amazon Elastic Compute Cloud in 2005
- OpenStack (as used by CC-IN2P3) has emerged as Open Source leader for IaaS implementations

Conventional grid job submission with DIRAC

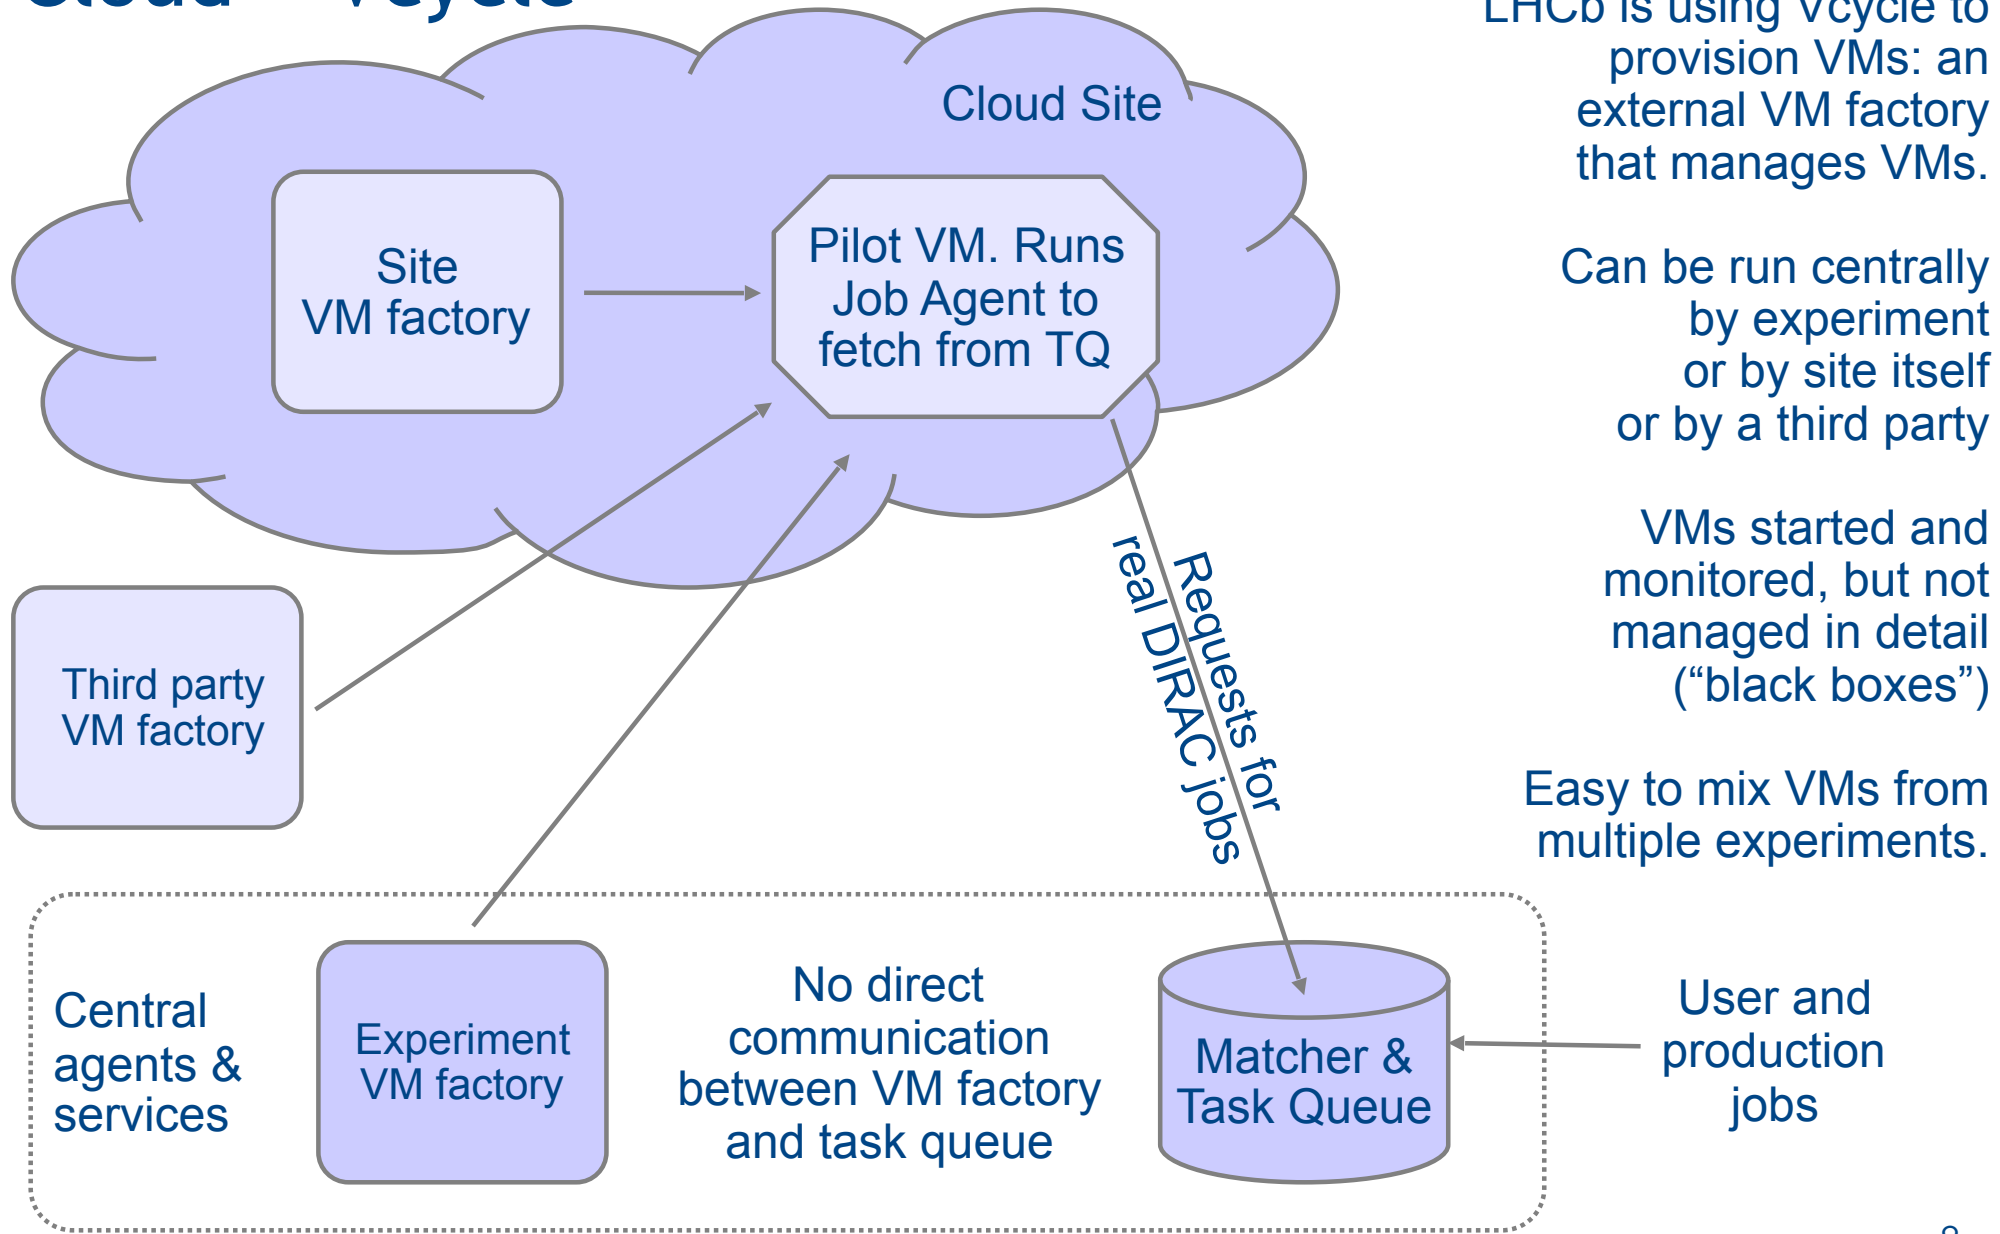
DIRAC began as LHCb's job submission framework.

Provides an overlay grid so that jobs can run at many different kinds of site.

Hides details from the user.



Cloud + Vcycle



LHCb is using Vcycle to provision VMs: an external VM factory that manages VMs.

Can be run centrally by experiment or by site itself or by a third party

VMs started and monitored, but not managed in detail ("black boxes")

Easy to mix VMs from multiple experiments.


LHCb use of CC-IN2P3 cloud

Small 10-VM
OpenStack
tenancy to
establish we
can use
resources.

Managed by
Vcycle at
CERN.


Running real
work.

Extremely
useful for
validating
operations
(thanks!)



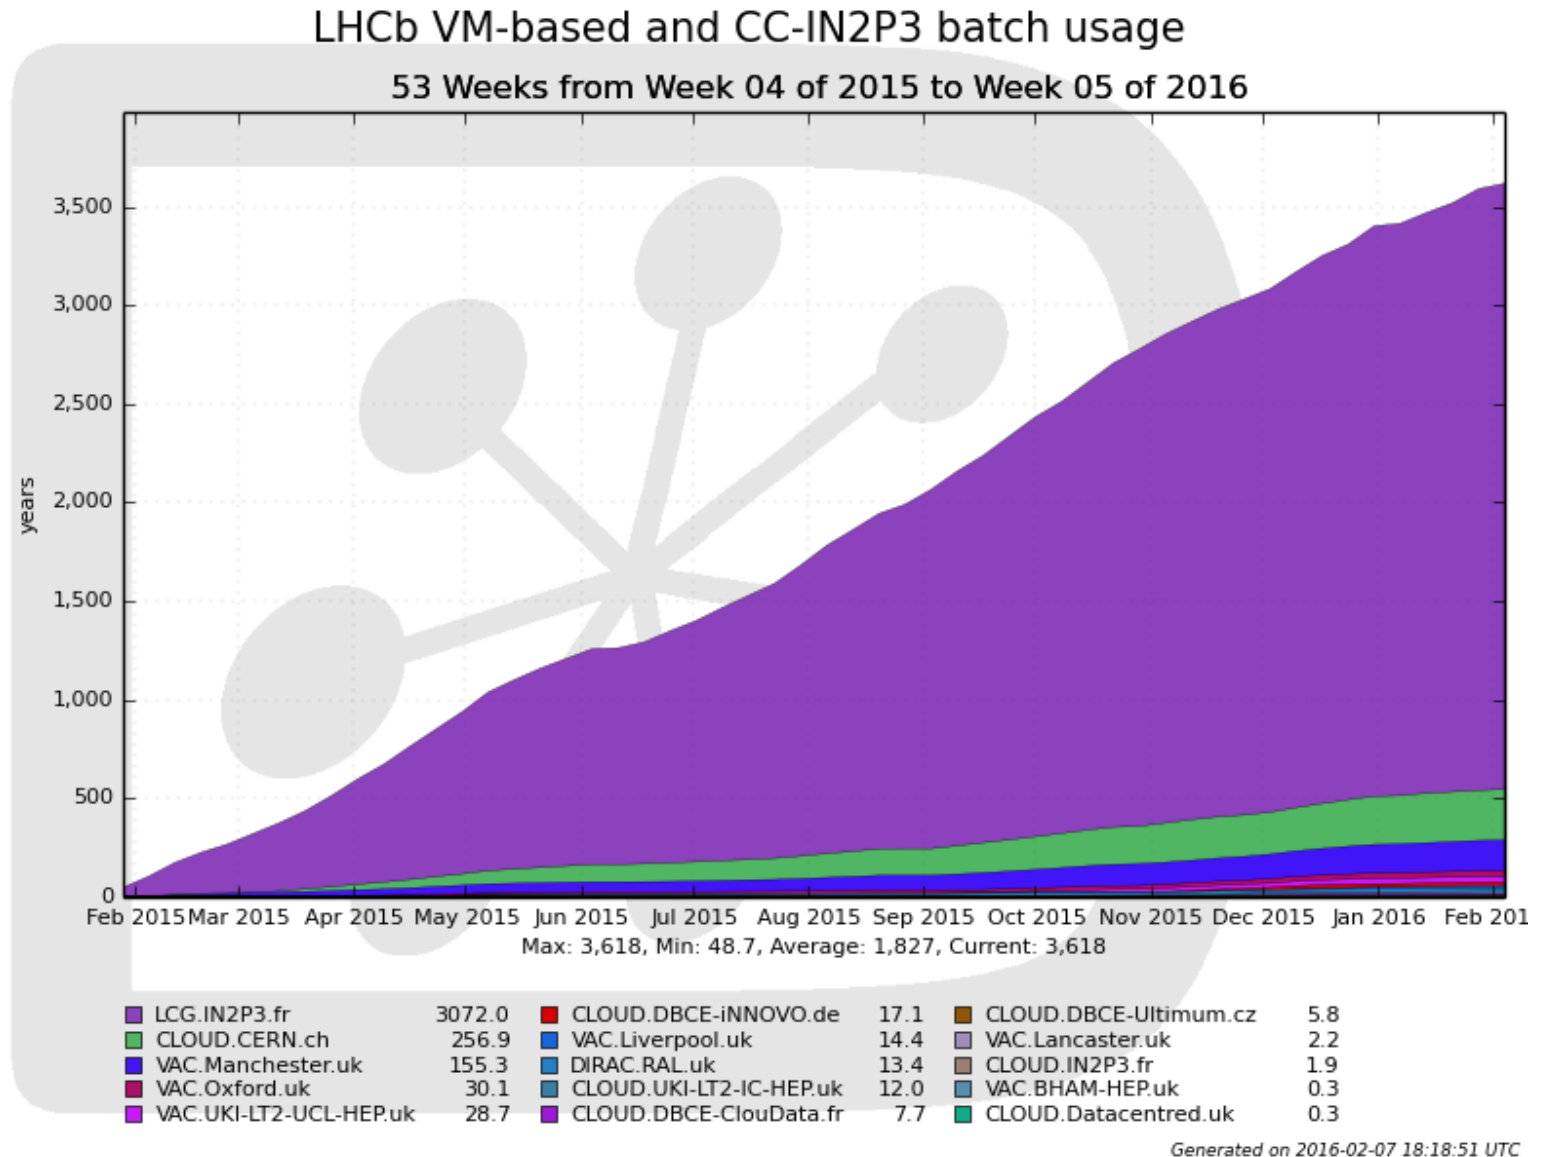
Generated on 2016-02-07 18:16:43 UTC 9

LHCb use of VM-based resources (Cloud+Vac)



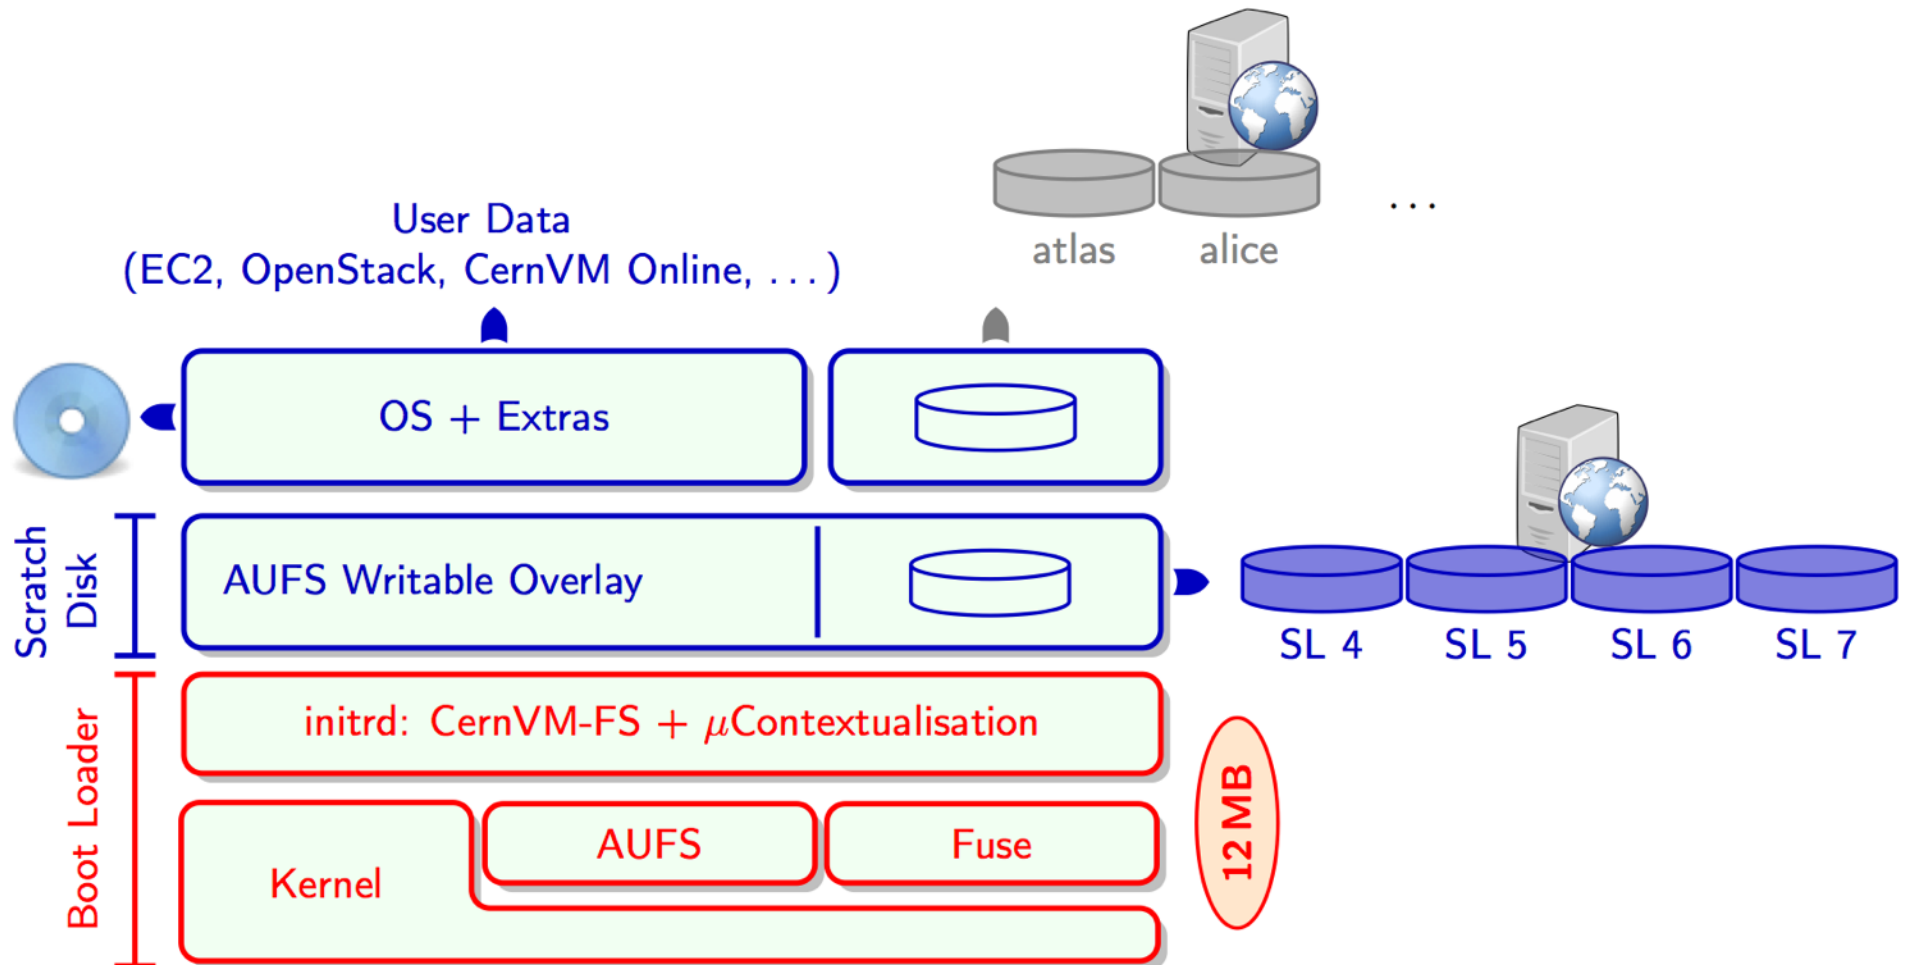
Generated on 2016-02-07 18:21:04 UTC

Add in CC-IN2P3 batch for comparison



LHCb VMs = CernVM + DIRAC

CernVM has operating system provided by WAN filesystem (CernVM-FS)





LHCb Cloud/VM plans

- Use the same VMs everywhere, on Vcycle/OpenStack and Vac sites
 - CernVM + DIRAC pilot client (no pilot job involved)
 - “Get it right once, should work everywhere”
- Now peaks at around 1200 mostly single-processor VMs (~3% of total)
 - Mostly Monte Carlo but also some production jobs
 - 500 VM LHCb tenancy at CERN runs with very little intervention
- Plan to request some pledged resources at Tier-1s as VMs
 - Will validate user analysis and production jobs with data access
 - Multiprocessor user jobs (in VMs, LHCb can test with cgroups etc)
- LHCb’s plan is to follow sites’ preferences for how the majority of resources are presented

Summary


- Clouds/VMs capacity is comparable to a large Tier-2 in LHCb
- Outside CERN, aim has been to verify :
 - Running at different types of VM-based site
 - Using different instances of a particular technology
- “Write one VM; run everywhere” strategy working well
- Will now request pledged capacity, starting at ~100 VMs per Tier-1
 - Also had spontaneous offers from interested T2-Ds
 - Want to verify use of Cloud/VM sites for user analysis and stripping jobs
- As always, LHCb policy is to follow sites preferences for the form that resources take
 - Subject to minimising the number of supported interfaces



Extra slides


Open source cloud implementations

- OpenStack is way ahead of other open source clouds in terms of commercial and academic adoption and number of developers
- Compare Linux taking off around 1995
- Particularly important for us as academic / research institutes tend to prefer open source
- OpenStack structure means institutes can contribute to features they need
- In particular, CERN has adopted OpenStack



Commercial providers

- Amazon (EC2), Microsoft (Azure), IBM (OpenStack,...), Google Cloud Platform, Rackspace (OpenStack/Azure/VMware), ...



- Ralph Finos, <http://wikibon.com/public-cloud-market-shares-2014-and-2015/>



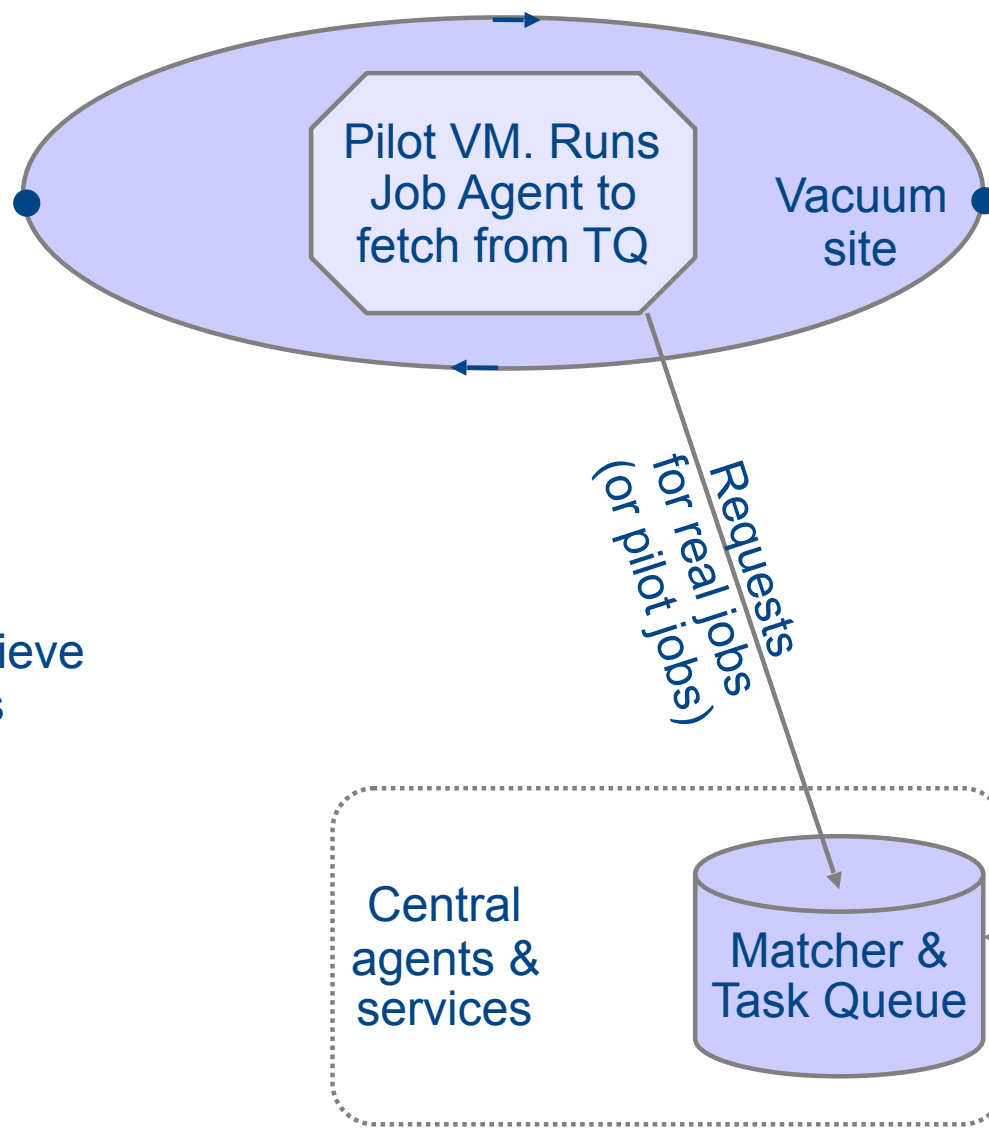
Guesses about the future

- Amazon and Microsoft will continue to dominate commercial provision, with closed source solutions
 - Only Amazon provides Amazon EC2; you can buy MS Azure software to run your own cloud service
 - Compare Apple MacOS vs Microsoft Windows
 - Maybe Google will catch up (Android did for phones)
 - Other players do not have the money/agility to catch up, but won't go away
- Academic/research providers will standardise on OpenStack for cloud services as they have on Linux
 - National/regional/institution HPC/HTC centres will increasingly offer cloud rather than batch
 - Big projects, e.g. SKA, looking at HPC on OpenStack, and at contributing required features to the codebase

Other types of “logical machine”

- Containers
 - Explosive growth of Docker bringing attention to this
 - “Heavyweight containers” approximate VMs, but share host’s kernel and native performance
 - Can be run by cloud systems (eg OpenStack)
 - Limitations of privileges model means need cvmfs from host too, or run something in userspace
- Unikernels
 - Build OS components into application, running directly on hypervisor or bare metal
 - Hypervisor takes care of hardware drivers, so one unikernel application can run everywhere
 - Application and kernel share address space so no context switching within unikernel logical machine

Vac: autonomous hypervisors



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

Use feedback from VM outcomes to decide which experiments' VMs to create as slots become free.

Easy to mix VMs from multiple experiments

Autonomous, but also inter-hypervisor communication to achieve desired targets shares between experiments.