

# Feedback from CMS

Andrew Lahiff

STFC Rutherford Appleton Laboratory

CernVM Users Workshop

6<sup>th</sup> June 2016

# Outline

- CernVM-FS
  - Content
  - Automated deployment
  - HPC
- CernVM
  - Virtual machines in CMS
  - CernVM usage by CMS

# CernVM-FS

# Software deployment in CMS

- Basically fully based on CVMFS
  - Recommended strategy by the WLCG Technical Evolution Group in 2011/12
- CVMFS used throughout CMS computing
  - Tier-0
  - Tier-1, Tier-2, Tier-3 grid sites
  - HLT cloud
  - cloud resources
  - (some) HPC resources
- 2 repositories
  - /cvmfs/cms.cern.ch
  - /cvmfs/cms-ib.cern.ch

# Content: CMSSW

- Main CMS software package
  - includes everything from generators, simulation to reconstruction
- Packaged as RPMs
  - externals shared among several CMSSW releases
    - ROOT, GEANT, GCC & various libraries
    - Also packaged as RPMs by CMS
  - Typical numbers for a CMSSW release
    - 150k+ files, 10k+ directories, ~6GB disk required
- Releases typically built for different architectures or GCC versions
  - e.g. MacOS X, AArch64, POWER8

# Content

- Additional input required for some workflows
  - Gridpacks
    - Generator files
  - Particle density functions
    - PDF sets from LHAPDF
- CRAB clients
  - tools for submission of analysis jobs
- SITECONF
  - local configuration needed for each site
- CMS@Home

# Automated deployment

- Executed via cron every hour
- CMSSW, CRAB client, and gridpack installs are fully automatic
  - Check for new CMSSW releases in “CMS tag collector” and/or “release.map”
  - Check for new CRAB versions in certain repository directories
  - Gridpacks are installed via rsync from an agreed EOS directory and new Gridpacks are pushed to CVMFS automatically
- CMSSW git repository is also mirrored daily
  - this does not go well sometimes due to a high load

# Automated deployment

- LHAPDFs are downloaded from the web and installed in CVMFS per GEN Group request
  - Some steering files to be manually edited before installation
- SITECONF is synchronized from [gitlab.cern.ch](https://gitlab.cern.ch)
- Feedback
  - Deployment machine was sometimes very slow
    - was upgraded last year, now much more responsive
  - CVMFS publication tools are updated by CERN IT per CMS request either by [cvmfs-talk](mailto:cvmfs-talk) mailing list or SNOW ticket



# Integration builds

- Periodic automatic builds of CMSSW are called integration builds
  - Several combinations of platform, architecture, GCC version
  - 7 base versions of CMSSW plus several flavour variations
  - 19 releases produced twice a day
- The latest 2 weeks of IBs are always available for developers

# CVMFS & HPC resources

- Lots of activity, mainly in the US
- Different classes
  - “Easy”
    - CVMFS available available on worker nodes
  - “Challenging”
    - not possible to install CVMFS, no outbound network connectivity and/or non-x86 architectures
- Have used Parrot to provide access to CVMFS in user space, e.g. Carver (NERSC)
- Currently no plans to continue using Parrot to provide CVMFS
  - Instead, using Docker images to provide access to CMS software
  - Originally a CVMFS snapshot, now just a specific CMSSW version


# CernVM

# Virtual machines in CMS

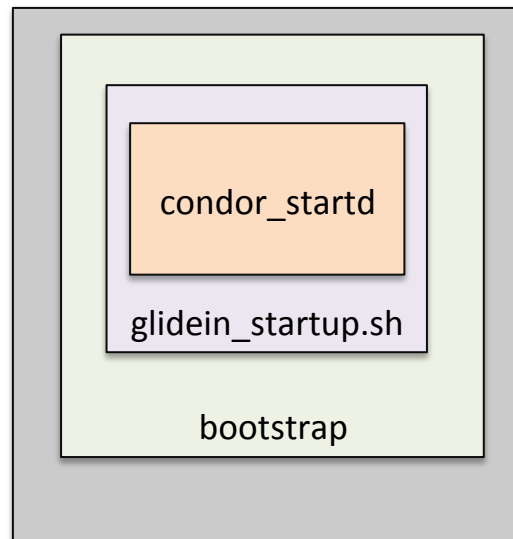
- Virtual machine images
  - custom VMs for Tier-0 & other activities
  - custom (& different) VM for HLT cloud
- glideinWMS
  - pilot framework used by CMS
  - glidein factories able to both submit jobs to grid sites & instantiate VMs using EC2
  - generates user data in a specific form for a specific VM image
  - bootstrap in the VM extracts the required information to run the standard pilot script

# Virtual machines in CMS

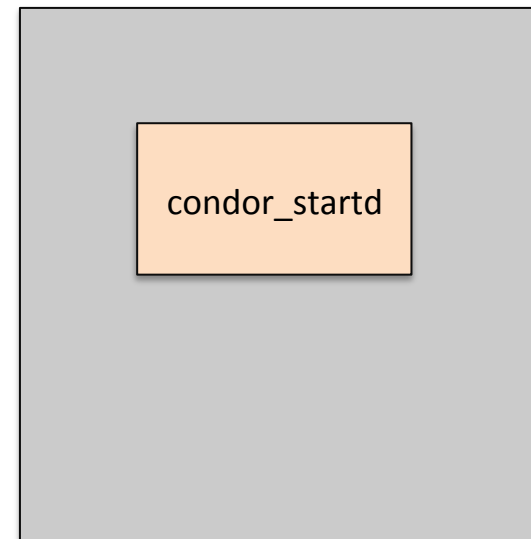
- Comparison between grid jobs & the two different types of VMs



**Grid job on WN**



**glideinWMS VM**



**HLT VM**

# Virtual machines in CMS


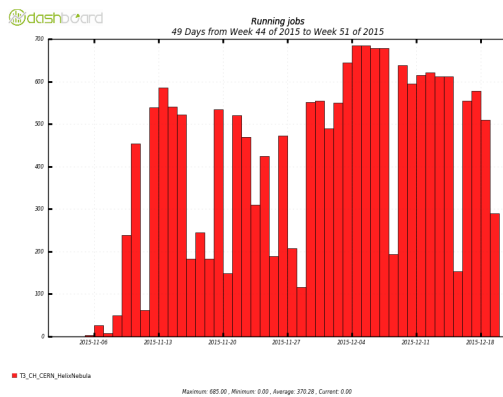
- As of 3.5, CernVM detects glideinWMS user data
  - can be instantiated by existing glideinWMS factories
  - automatically runs the CMS VM bootstrap
- Squid proxies & CMS\_LOCAL\_SITE setup automatically
  - uses information in <http://frontier.cern.ch/geolist.txt>
  - only works at existing grid sites
  - CMS\_LOCAL\_SITE will be the same as the grid site
    - some sites want something different or are more complex (e.g. CERN)
- Can be changed via user data, but there's a glideinWMS factory limitation
  - designed for a specific custom VM image, user data can't be changed
  - not enough effort available to change this any time soon

# Virtual machines in CMS

- CernVM 3.6 can extract the proxy & site name from glideinWMS user data
  - CVMFS\_HTTP\_PROXY, CVMFS\_PAC\_URLS, CMS\_LOCAL\_SITE
- Have run some successful tests on CERN OpenStack, but testing is ongoing
  - Also want to test CernVM on another sites with cloud resources
- Plan is to switch CMS Tier-0 from the glideinWMS image to CernVM
- Ideally CernVM should be available on CERN OpenStack by default

# CernVM usage by CMS

- CMS@Home (*see Ivan's talk*)
- HelixNebula
  - Part of ongoing series of tests by CERN IT with commercial cloud procurement
  - 4000 cores for all LHC experiments
  - CernVMs provisioned using vacuum model (Vcycle)
  - WMAgent MC production run successfully (Nov-Dec 2015)





# Summary

- CernVM-FS
  - Successfully being used throughout CMS computing
- CernVM
  - Already have used CernVM for some activities
  - Aim is to have CernVM as the recommended image for all cloud resources in CMS