



LHCb Use of CernVM and CernVM-FS


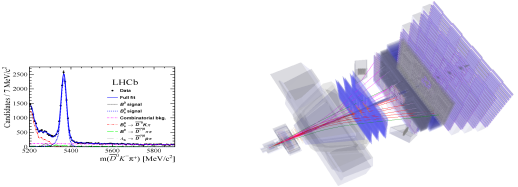


CernVM User Workshop

CERN March 5th-6th 2015

*Ben Couturier
for the LHCb Computing Team*

CernVM(FS) in LHCb

- CernVM and CernVM-FS Use cases

	 <i>Production</i>	 <i>User</i>
 <i>CernVM-FS</i>	<ul style="list-style-type: none"> • Productions jobs on physical hosts • Release distributions • Conditions distribution 	<ul style="list-style-type: none"> • Software development • Nightly builds deployment • Local analysis <i>On local clusters</i>
 <i>CernVM</i>	<ul style="list-style-type: none"> • VAC / Vcycle jobs • Volunteer computing (BOINC) 	<ul style="list-style-type: none"> • Software development • Local analysis • OpenData / Masterclass release

This feedback



To cover all the Use Cases:

- The librarian's opinion
- Feedback from the production team
- Tried to Survey the usage in LHCb...
 - ~700 physicists and engineers from 69 institutes...

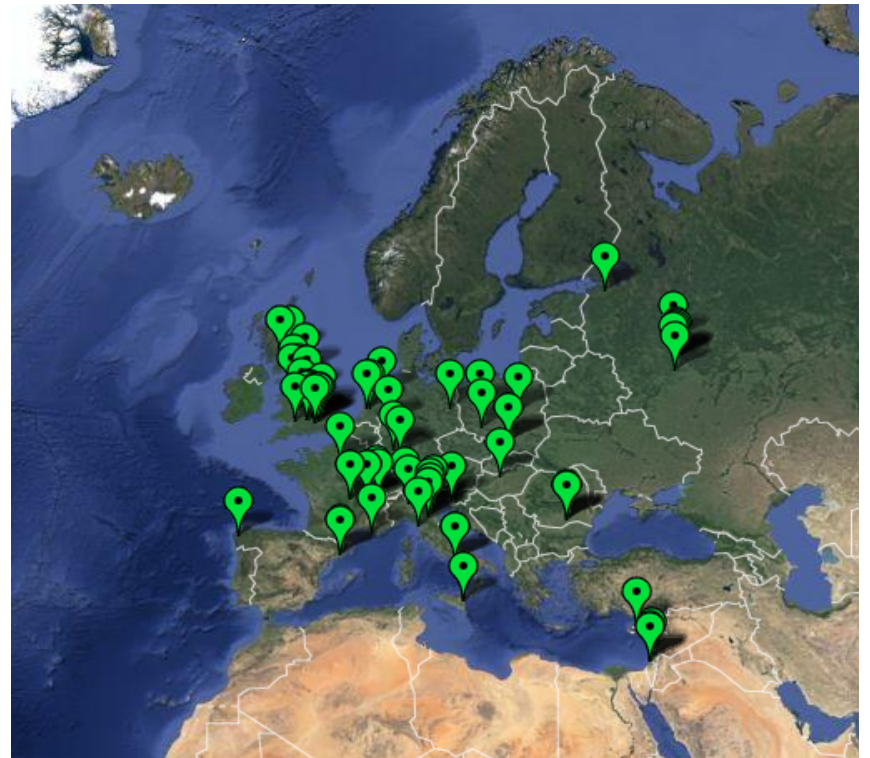


CernVM-FS in LHCb



CernVM-FS@LHCb

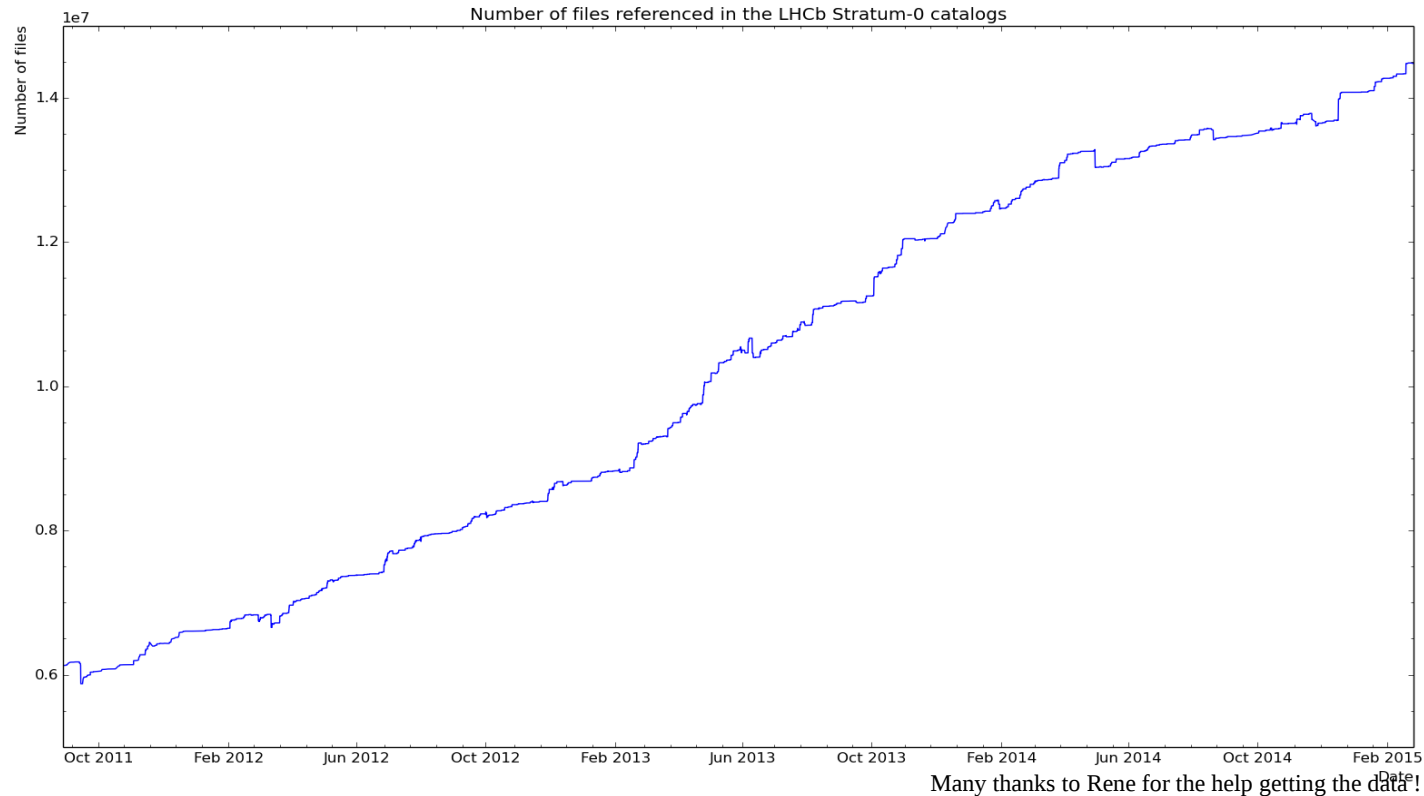
- In production since 2011
- THE deployment mechanism since spring 2013
 - Required for all distributed computing sites used by LHCb (~80 sites)
- Used for the release of software
 - 200K files per month in average (frequent releases)
 - Detector Conditions updated every hour (needed for offline reconstruction).





Production Setup

- /cvmfs/lhcb.cern.ch - LHCb Production volume
 - 614 GiB
 - 14.5 million file in catalog
 - 4.8 million physical files after de-duplication
- Tests have shown that we can scale to 10 times that...





Cern VM-FS Experience

- Many benefits to our switch to CVMFS only in 2013
 - Abandoned our custom deployment code
 - Including for Dirac (our grid software)
 - Significant time saving on the deployment side
- Not without errors of course
 - Transient errors do happen...
 - ~80 GGUS tickets with CVMFS failure in the last year
 - But the file system is sometimes blamed when it shouldn't
 - And we ran ~ 35M jobs during that time...
- Propagation delay to the grid can be difficult to handle
 - Any reduction in delay would be helpful
 - Guaranteed maximum delays would be useful



Cern VM-FS Client Experience

- Appropriate caching is crucial to good operation
 - Peak of file-system errors do appear when starting new productions
 - *Production managers are aware of the fact*
- Improved diagnostic tools would help
 - Tools to debug more easily would help sorting the problems
 - and of course reduce the number of tickets wrongly attributed
- Sometimes CVMFS is mounted via NFS
 - In which case central monitoring difficult to do (NFSv4 xattrs could help)

**Of course there are always problem but overall
A VERY POSITIVE EXPERIENCE !**



Cern VM-FS Server Experience

- Using Stratum-0/1 Service provided by CERN-IT
 - Many thanks!
- Version 2.1 is a great improvement
 - Not a big issue if a publish is interrupted...
 - Considering it took ~48 hours for a filesystem check
- Publish is a very CPU intensive operation!

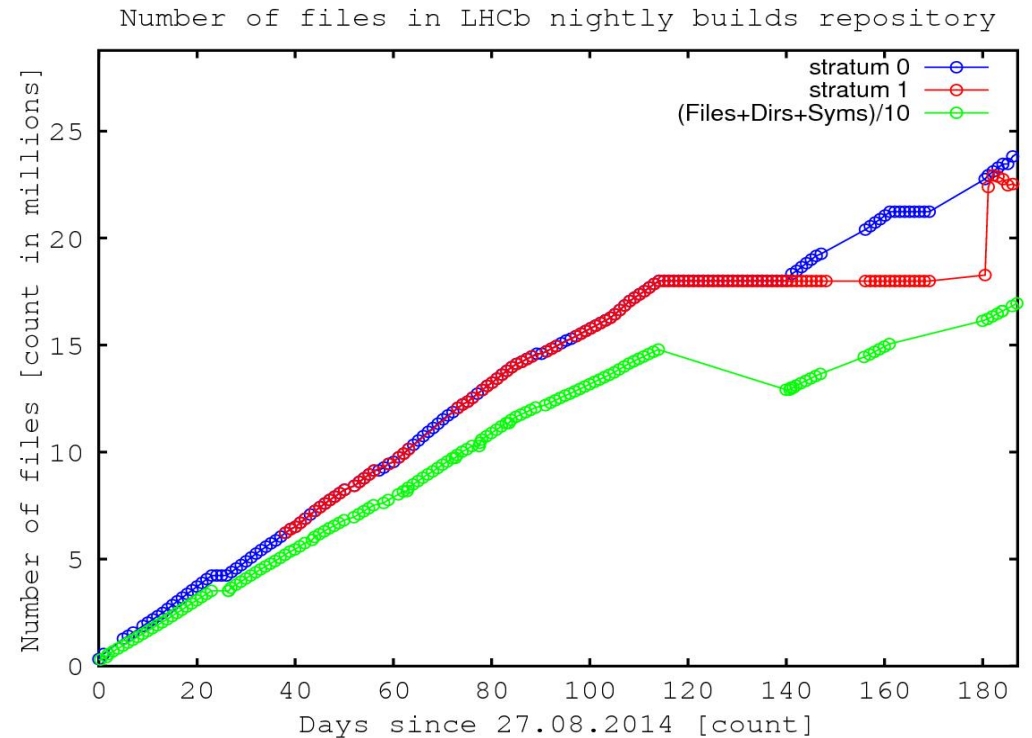
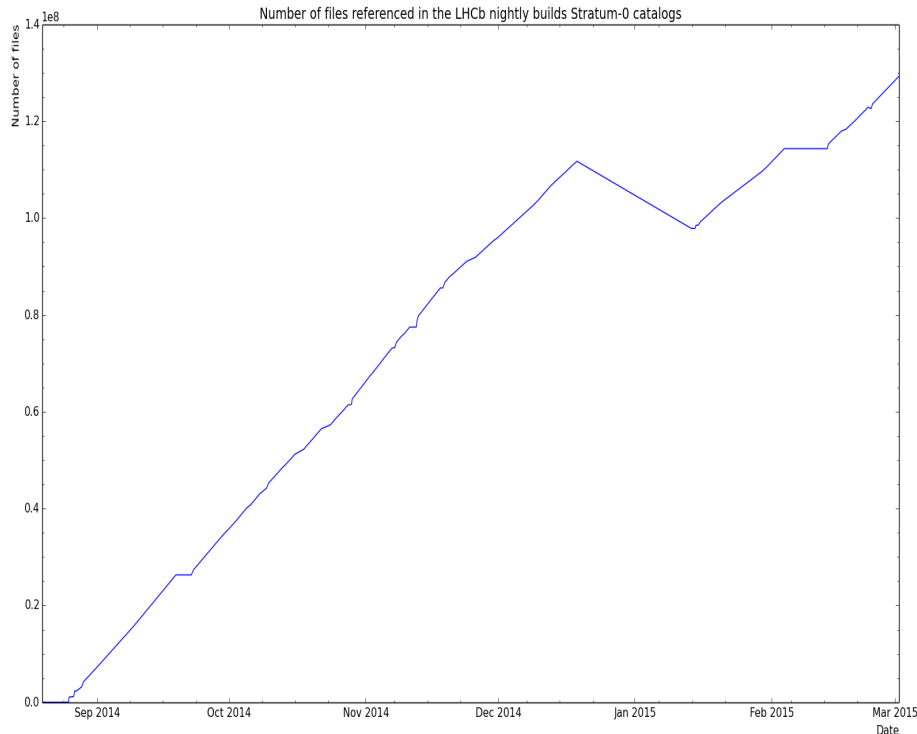


Nightly builds on CVMFS

- A long standing requirements
 - To reduce dependency on AFS
- Not easy to implement on CVMFS:
 - Deployment propagation time can still be problematic
 - Too costly be done on standard production Hardware
- Tried to use a VO box with ~ 2TB of storage
 - Managed to fill it up in a month...
 - Could reset every month but that is a bit painful
 - *Garbage collection* is the solution for this use case

Nightly builds on CVMFS

- Trial of the Openlab S3 storage
 - Already 1 order of magnitude more files than production!
 - c.f. presentation by Seppo



* Plot from S.Heikkilä

Nightly builds on CVMFS



- The system scales to our production requirements
 - Stressed the system:
 - 1e6 files per day (Around 2:00 for copy/publication)
- Great collaboration between LHCb, CernVM-FS developers and CERN-IT/Openlab
- S3: Interesting solution for scalable stratum 0 backend !



Conclusion on CernVM-FS

- CernVM-FS is a key component for LHCb
 - Offline Production team very happy with the move to CVMFS only!
- Software distribution is critical to running the experiment
 - The whole CernVM-FS deployment chain is therefore critical
 - CVMFS will not be used on the LHCb Online Farm trigger nodes
 - As this would require 24/7 support of the LHCb stratum-0
- LHCb interested in the new features!
 - e.g. Garbage collection
- Many thanks to the teams involved !
 - Great collaboration, their competence and attitude was key in the success



CernVM in LHCb




Use of CernVM

- Use in production:
 - Sites using VAC/Vcycle
 - Volunteer computing using BOINC
- Also used for:
 - Software development
 - Data analysis



Experience with VAC/VCycle

- As presented by Andrew McNab
 - Small percentage of LHCb jobs (~ a T2 site)
- Uses CernVM 3
 - It just works!!!!



Generated on 2015-02-26 16:19:00 UTC



Volunteer computing



<http://lhcbathome.cern.ch/Beauty/>

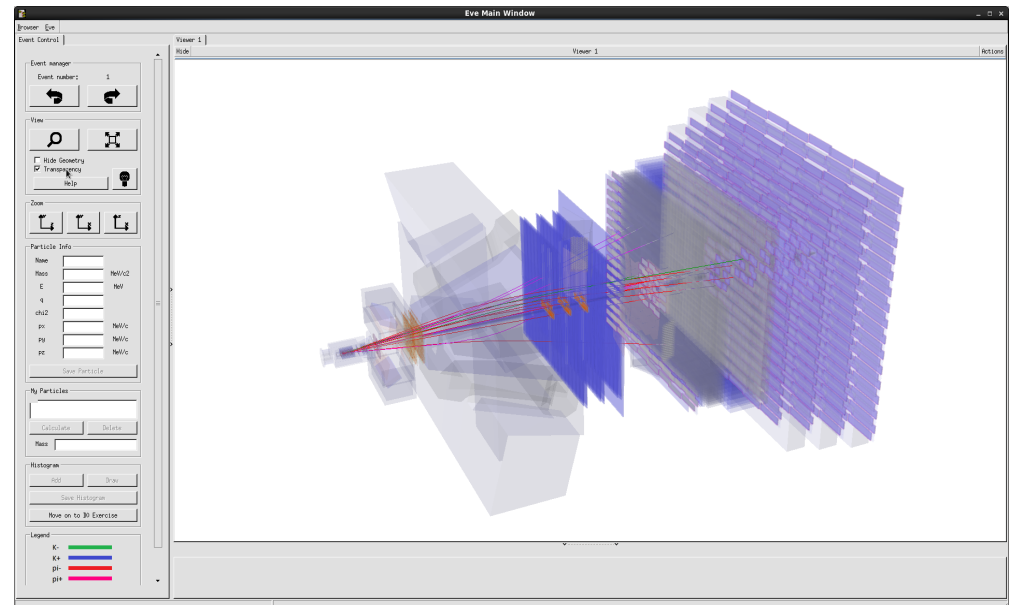
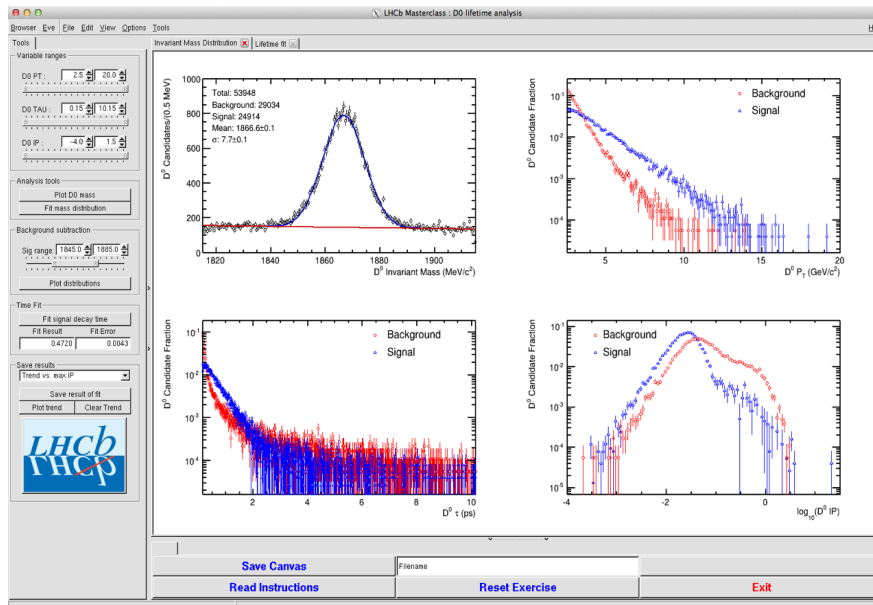
- CernVM 3 with BOINC
- Same contextualization as production



opendata
CERN

Open Data release

- LHCb Masterclass exercise now distributed in CernVM as part of OpenData
- Used in the 2015 Masterclasses





Data Preservation

- Legacy software to stay on CernVM-FS
- Need to keep SLC5 compatibility
 - As used in Run I
 - CernVM is the ideal candidate



User feedback



Cern VM-FS

- Everybody uses CVMFS
 - at least for user Jobs
- Overall positive feedback on the file system
 - Sometimes cache problems...
- Interest in running on a machine without ROOT privileges
 - Parrot plugin
- Contrasted feedback on documentation
 - Good for people who know their OS well
 - Not necessarily appropriate for new users
- Would be nice if directly in the SLC6/CC7 installation



CernVM

- Varying opinion on CernVM:
 - Very practical way to deploy the LHCb stack
 - Praises for the small image, updates, etc...

but

- Some people prefer lightweight containers:
 - Chroot/docker vs running a VM
 - With CernVM-FS already configured
- CernVM 3 appliance model felt more appropriate for Data Centers
 - Some use cases were difficult to fulfil:
 - e.g. Getting latex to work
 - Running OpenGL Software (Event Display) is not necessarily easy



Outlook



CernVM in LHCb plans

- CernVM-FS is THE deployment mechanism for LHCb production
- CernVM is only a small proportion of the Jobs
 - But we feel it's crucial to be ready for cloud sites
- Great tools for Data Preservation
 - Even if our strategy is not fully defined, we know that CernVM-FS and CernVM are at its core



Conclusion

- LHCb Committed to CernVM(-FS)
- Happy with the direction of the tool
 - More diagnostics would be helpful though !
- *Many thanks to the development and operation teams!*