CVMFS

- Rod Walker, LMU Munich 3rd Dec 2010

- What is it?
- Why good for prod sites?
- Concerns
- Current deployment
- Conclusion and proposal

What is it?

- Read-only web FS developed for CERNVM
  - to provide exp SW but keep small VM image
  - Web server has all SW. WN caches files used
    - caching web proxy per site

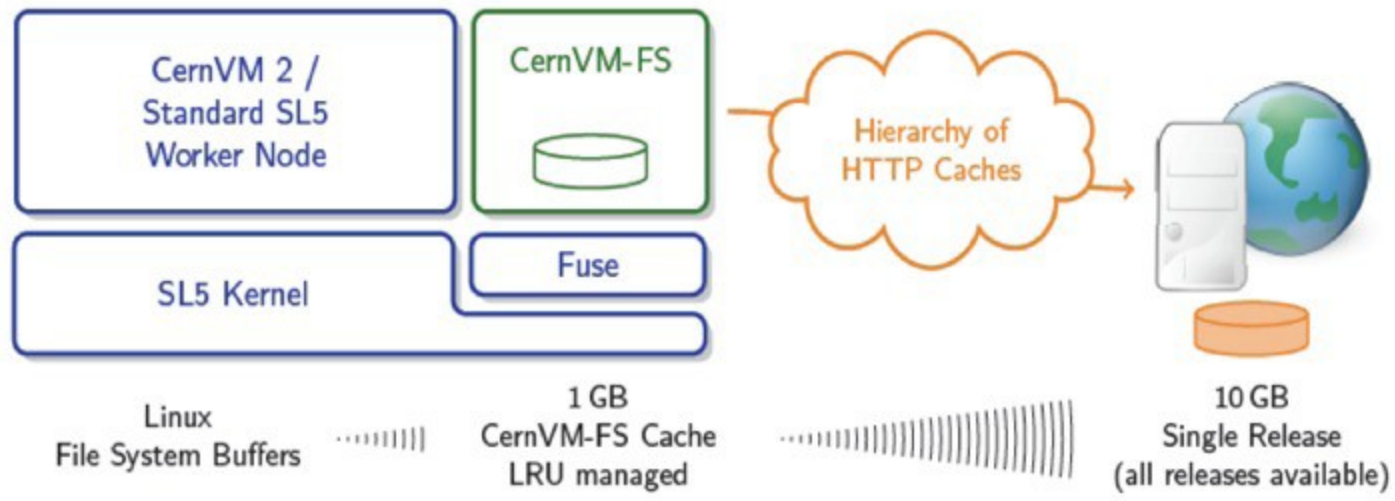

Why consider for prod sites?

- Negatives of current distbn scheme
  - NFS,AFS and esp. Lustre, GPFS do not scale well
    - cmt causes many metadata operations
    - setup varies with load between 10s and 10mins
  - space required in high performance disk limited
  - some understandable delays in install via WMS
    - 180 caches on 100+ sites
- SW install would instead just validate
  - checks native libs are there

CVMFS advantage

- All sites see all releases instantly
- Scales and has local disk performance
  - important for build jobs
- All DBReleases available, unpacked
- Conditions data flat pool files available
  - no need for HOTDISK/PFC
- Lowers the requirements to run prod

Concerns

- Single point of failure
  - Depend on web server, connection to CERN
    - if WN has the files, the should be independent
      - but ttl consistency checks would fail (I think)
  - Mistake would break all sites at once
    - all fixed/patched at once too
  - Investigate mirror with transparent failover
- /opt/atlas path burned into rel setups
  - must be free on all sites, and not an nfs mount
  - let's make release relocatable with $ENV base

Current Deployment

- Many CERNVM's using it
- QMUL,RAL,Wuppertal
  - for production(not RAL) and analysis
- 3 lxplus nodes
  - and prepared to put on all and lxbatch
    - not requested yet. Reduce afs load and speed-up builds.
- LHCb uses it on some prod sites
  - with requests to problematic sites

Installation

- Reported to be easy to install
  - installed as root
  - 2 rpms in yum repository
  - requires /opt/atlas path
    - non-relocatable ATLAS SW
- Potential to install non-root – but best not
  - still needs Fuse permission
  - leave cache on WN disk /tmp – tricky
- Other VOs use it
  - probably WLCG can help with deployment

Conclusions

- CVMFS validated for production and analysis
  - validated by existing SW install mechanism
    - can fasttrack because do not install
- Proposal:
  - gain more experience from early adopters
  - invite sites to install it, but not compulsory
  - more encouragement to problematic sites
    - and you know who you are?
- Need to provide instructions and migration strategy