

LHCb Software and Virtualisation : A Summary of User Experiences

C.Jones

Introduction

- ✦ Collect and report a selection of user experiences of CERNVM (and other solutions) running the LHCb Software
- ✦ Disclaimers :-
 - ✦ Distillation of the (many) responses I received.
 - ✦ Many thanks to all those who sent me feedback.
 - ✦ Apologies if anything gets 'lost in translation'
 - ✦ Please jump in and correct me.
 - ✦ I am not the biggest LHCb VM user... Just the messenger.
 - ✦ A mixture of different CERNVM versions. Apologies in advance for reports of problems in 1.2.0 already fixed in DEV releases.
 - ✦ Some issues not specific to CERNVM, but in general to VMs.
 - ✦ In general I think issues should be report via CERNVM mailing list or savannah.

Outline

- ✦ Various themes came up several times in the reports I received.
 - ✦ Initial installation of the VM
 - ✦ Speed
 - ✦ Expected Software
 - ✦ Both LHCb and general linux applications.
 - ✦ Differences in usage models
 - ✦ Simple X-less server
 - ✦ Full Gnome/KDE X session

Installation

- ✦ Wide range of host OS'es and VM applications
 - ✦ Windows (XP, Vista), Mac and linux
 - ✦ VMware (Fusion and player) and Virtual Box VM applications.
- ✦ Wide range of experiences
 - ✦ Ranging from extremely smooth to various strange issues (some anecdotal reports, cannot really comment further myself).
 - ✦ Better documentation was a theme. Is improving through.
 - ✦ Many resorted to Atlas specific documentation, converting on the fly to LHCb...

General Issues

- ✦ Slow boot up, particularly when outside CERN
 - ✦ Seems related to significant network activity at a certain point in the boot sequence.
 - ✦ Happens **every** startup (even if I reboot immediately), so not just initial file system caching.
- ✦ Some reports of applications running too slowly
 - ✦ Some, I suspect related to initial caching of application binaries.
 - ✦ Need to be careful what you are comparing. CERNVM SLC4-32bit versus native SLC5-64bit is not entirely fair.
 - ✦ Still, some suggestion that CERNVM is slower than other similar VMs persist...
- ✦ More problems with 64-bit CERNVM than 32-bit version.
- ✦ Multi core problems with parallel processing

System Software

- ✦ Some things missing in current PRO release (1.2.0)
 - ✦ subversion (needed by 'getpack')
 - ✦ procmail (???) provides lockfile for CMT
 - ✦ LHCb environment setup scripts too old
- ✦ Should all be fixed in more recent DEV releases

Experiment Software

- ✦ LHCb software available in CERNVM often lags behind that available on AFS or via `install_project.py`
 - ✦ Seems to be getting better.
- ✦ Requests for inclusion of ganga
 - ✦ Currently uses own installation script. Plans to move to using LHCb “`install_project.py`” like other projects, which should allow this.
 - ✦ Maybe also requires a local config file for CERNVM in order for ganga to work out the box.
- ✦ LHCb emacs environment package.

Other Software Requirements

- Data Access
 - In general users want access to data ...
 - E.g. castor. Does not work out-the-box with the LHCb software, like on lxplus.
 - LHCb 'castorfs' fuse filesystem works well in one case.
- AFS. Not strictly needed by CERNVM, but wanted by several people
 - Easily Installable, but does not work well. Probably network issues (NAT).
 - Issues not really CERNVM specific
- Shared Folders.
 - Useful to share files between host and various VMs. But problems..
 - Does not work always out-the-box - Requires fiddling with some configuration files (e.g. /etc/fstab)
 - Building LHCb software uses some filesystem features (sym/hard links) not supported on shared filesystems (hgfs, etc.)
 - Again, not a CERNVM specific issue ...

Use Case 1 : Minimum Server

- ✦ Simple X-less server.
 - ✦ Login (via ssh for instance) from host and export windows back to host X11 server.
 - ✦ Network communication purely internal, so very fast
- ✦ In general those using this approach found it to work very well.
- ✦ Those generating their own data seemed most happy, since no issues with external data access.

Use Case 2 : Full X Session

- ✦ Much high requirements placed on virtual machine
 - ✦ Requires decent VM 'Hardware' drivers etc.
 - ✦ Depends on what the user expects.
 - ✦ Resizable X, mouse driver, "Unity" mode, Full screen etc. etc.
- ✦ Some success, but in general more problems encountered.
 - ✦ Mostly I think due to SLC4 being somewhat ancient in the linux world. SLC5 better but still not exactly recent.

Use Case 3 : Other Approaches

- ✦ Ubuntu 9.04 VM.
 - ✦ Fully binary compatible with SLC5(64bit) binary releases. Installation trivial via *install_project.py*
 - ✦ Some users prefer the predictability of installing software themselves. Others prefer the exact opposite.
 - ✦ More modern Linux Distribution
 - ✦ Better support by VM applications. Particularly for full X sessions.
 - ✦ Larger choice of more modern applications.
 - ✦ Larger disk usage than CERNVM. But with disks now O(1TB) ...
 - ✦ Maybe faster as a VM than CERNVM ...
 - ✦ AFS available.
- ✦ Standard SLC5 installation in a VM also used successfully.
- ✦ SLC(4/5) in a chroot environment on top of another Linux installation.

Conclusions

- ✦ Users appreciate CERNVM providing a full LHCb software environment
- ✦ CERNVM Wish List (From a desktop user perspective) :
 - ✦ Faster availability of new LHCb software releases
 - ✦ More LHCb software working out-the-box
 - ✦ ganga, LHCb Emacs etc.
 - ✦ Wider selection of general linux applications (tkdiff, crossvc etc.)
 - ✦ Access to data (Castor, Grid ?)
 - ✦ Better AFS ...
 - ✦ Causal work OK. Heavy usage (LHCb nightlies) poor.
 - ✦ SLC5 64-bit (coming with 1.4 ?)
- ✦ CERNVM not the only solution to running LHCb software on non-SLCX systems.
 - ✦ Ubuntu 9.04 as a VM does not seem to suffer from many of the problems CERNVM does.