

# Experiences, Opinions and Plans

## General Thoughts

Talk from April GDB on “Future Linux Distribution Scenarios update”:

<https://indico.cern.ch/event/876788/contributions/4309932/attachments/2226066/3770705/gdb-linuxfuture.pdf>

“Containers are the future” - or at least a mitigator.  
CEPH deployment model is picked up.

What do the VOs want to run their jobs on?

CMS - 100% Singularity. Same with ATLAS.  
WN problem “solved” with can it run Singularity (and probably cvmfs).  
How to limit impact on Local users? - Don’t want to run different OS.

## Requirements

Singularity  
Stability  
Security Patched  
Hardware support

## QMUL

Stream not a “bug-for-bug” equivalent of RHEL8.  
Does it matter if we’re running a bug/binary compatible version the exact flavour of linux?  
But not running the official supported OS is making a rod for your own back.  
(QM run OracleLinux8)

## Rocky Linux

Oxford: condor pkgs installs fine once power tool repo enabled and processes starts.  
Having trouble with cobbler - but imagine it would be the same for any CentOS-clone.  
NFS server and Apache works.

## AlmaLinux

Lots of experience at Edinburgh (see slides).

## CentOS8 Stream

Not a bug for bug equivalent of RHEL.

Short lifetime

Probably not as stable

## “Paid for” RedHat

16 “developer” licences.

## Ubuntu/Debian

Brunel: I am testing Ubuntu for new storage and worker nodes

## Round Up

Come back to this at GridPP, collate efforts on it in the Tuesday Ops Meeting.

Want a bug-for-bug RHEL equivalent.

Rejecting Stream (as this isn't it)

Alma and Rocky are pretty much a toss up at the moment.

Attempt to be in a position to make a decision at GridPP.

Action on PMB to invite a relevant expert to talk at GridPP

## Chat Window

11:02:56 From Brew, Chris (STFC,RAL,PPD) To Everyone : I do like the fact that CERN now has TWO Single sign on systems.

11:11:03 From West, Matthew To Everyone : There need to be some efforts to get best-practices identified and shared amongst the community. Whether that's container orchestration or something else.

11:13:53 From Luke Kreczko To Everyone : Aren't workernodes relatively easy (singularity/Docker)? What about services? I've only seen containers used for DEV for grid services, but not really in production

11:14:19 From West, Matthew To Everyone : The more HPC minded folks have very intense opinions about compiler flags and do not trust containers.

11:15:16 From Luke Kreczko To Everyone : but trust can be built over time.

11:16:16 From West, Matthew To Everyone : @Luke, I agree. But that takes people and folks are very busy. I try in my own little corner but it's a non-trivial social task.

11:16:41 From Luke Kreczko To Everyone : yeah, completely agree.

11:17:37 From Luke Kreczko To Everyone : So any OS we pick needs to support:

1. containers (Docker/singularity)

## 2. CVMFS

### 3. A firewall that is easy to configure

Anything else?

11:18:03 From Luke Kreczko To Everyone : yes, a united solution --> easier to debug as a community

11:18:14 From Robert Currie To Everyone : Do pilots require python2 or 3 then, and what external modules and such or does the pilot launch singularity almost straight away?

11:18:16 From Adam Boutcher To Everyone : Does anyone know what commercial/enterprise is moving to? Reduced overall security efforts...

11:19:46 From West, Matthew To Everyone : Heard HPC sites sticking with RHEL

11:20:07 From Luke Kreczko To Everyone : user jobs need "behind the scenes containers"

11:20:13 From Peter Clarke To Everyone : But maybe they can afford the licence

11:20:16 From Peter Clarke To Everyone : we cannot

11:21:00 From West, Matthew To Everyone : @Peter, absolutely.

11:21:11 From Alessandra Forti To Everyone : ATLAS runs everything in singularity on the grid

11:21:28 From Alessandra Forti To Everyone : On the T3 there is also a push to run in containers

11:21:49 From Alessandra Forti To Everyone : Exactly to allow users to run the same stuff everywhere

11:22:31 From Luke Kreczko To Everyone : containers also need vulnerability tracking

11:23:03 From Mark Slater To Everyone : Just to mention: I don't \*think\* LHCb are really using containers at all. Maybe someone else knows but more but I believe that's the case

11:23:39 From Luke Kreczko To Everyone : But they have the option, Ganga supports container jobs

11:23:51 From Alessandra Forti To Everyone : LHCb can easily run containers too

11:23:55 From Crooks, David (STFC,RAL,SC) To Everyone : @Luke oh very much so

11:24:15 From Mark Slater To Everyone : I stand corrected :)

11:24:16 From Crooks, David (STFC,RAL,SC) To Everyone : I am intensely interested in this

11:24:20 From Luke Kreczko To Everyone : even LZ picked up Ganga in places for this

11:24:35 From Alessandra Forti To Everyone : All the covid stuff I run on gridpp using Dirac was containerized (not to talk about last and ska containerised jobs)

11:24:44 From Alessandra Forti To Everyone : lsst

11:25:46 From Brew, Chris (STFC,RAL,PPD) To Everyone : CentOS8Stream is not bug for bug compatible!

11:25:52 From Daniela Bauer To Everyone : @Luke: Yes, but LZ containers didn't exactly work out of the box on the grid as you know ;-)

11:26:18 From Robert Currie To Everyone : Can Panda/Dirac devs be asked to 'by-default' drop a user into a container so they don't worry/notice?

11:26:38 From Luke Kreczko To Everyone : yeah, there are always teething issues :). But if containers are the main way to run jobs on the grid --> these issues should become very rare

11:27:41 From Luke Kreczko To Everyone : @robert: This can also be done at site level. Andrew Lahiff used to put every job into a Docker container and VOs would start singularity within it

11:28:23 From Robert Currie To Everyone : @Luke, yes but it would be nice if there is a shared solution I'd argue not requiring all sites to re-architect if possible

11:28:39 From Luke Kreczko To Everyone : one day, one day ;)

11:31:38 From Adam Boutcher To Everyone : Which has the best community? Why couldn't we just become a SIG on one of the distros? Keeps man-power on a mainline effort and we're not playing catch-up with our own distro.

11:31:50 From Alessandra Forti To Everyone : We were a sig in centos

11:32:05 From Dewhurst, Alastair (STFC,RAL,SC) To Everyone : sig?

11:32:17 From Adam Boutcher To Everyone : Special Interest Group

11:42:21 From Alessandra Forti To Everyone : Dropping old hardware is normal though depends how old

11:46:21 From Adam Boutcher To Everyone : FYI Secureboot is not yet supported with Alma and Rocky as they've not got a signed shim but that is coming soon.

11:48:19 From Alessandra Forti To Everyone : Is also the same guy that commercialised singularity

11:48:45 From West, Matthew To Everyone : The group behind Rocky are also currently fighting Syslabs over Singularity if I remember correctly.

11:51:16 From West, Matthew To Everyone : We are all talking about Gregory Kurtzer for those unaware of the name.

11:51:23 From Alessandra Forti To Everyone : At the very least he might be naive

11:53:00 From West, Matthew To Everyone : Is the variety of products from Fedora not worth looking at? Not stable enough?

11:55:23 From Mark Slater To Everyone : I use Fedora for the desktops and it's good (though I don't have a lot of experience with other things!). Problem is the short life cycle. There are releases every 6 months with each release only having a ~1 year life time

11:55:35 From West, Matthew To Everyone : That's what I thought

11:55:39 From Mark Slater To Everyone : Saying that though, upgrades are generally easy

11:57:45 From West, Matthew To Everyone : My question is: for those who are setting up a new cluster, pool, grid now, what should people be investing time in?

12:02:38 From Luke Kreczko To Everyone : and moving from Alma-->Rocky is very easy, apparently ;)

12:06:04 From Alessandra Forti To Everyone : They'll have to take a decision because the experiments might not go back to support multiple OSs

12:06:30 From Alessandra Forti To Everyone : They are already struggling with diversification of hardware

12:12:22 From Raul Lopes To Everyone : Is it clear that all kickstart and ansible will work with Alma or Rocky?

12:13:03 From Matthew Steven Doidge To Everyone : I don't see why it wouldn't, but I haven't tried it yet (just interactive installs).

12:13:35 From Raul Lopes To Everyone : I don't see why it wouldn't is different from it has been tested

12:14:08 From Adam Boutcher To Everyone : If you require secure boot you can't currently use either Alma or Rocky

12:14:23 From Raul Lopes To Everyone : Did Edinburgh tried completely automated installs?

12:14:58 From Robert Currie To Everyone : @Raul unfortunately no I haven't tried this yet

12:15:43 From Raul Lopes To Everyone : Actually it's concern for Ubuntu. They've just changed automated management

12:16:10 From Robert Currie To Everyone : @Adam can't you just add their keys to the trusted keystore for secure boot? (Never played with this in the server space but have built and signed my own kernels on a laptop)

12:17:01 From Robert Currie To Everyone : @Raul, I can speak to the dev who are deploying Ubuntu at scale in the school to see how they're managing it

12:20:31 From Raul Lopes To Everyone : Yes, rob, it would be very helpful

12:27:12 From Daniel Traynor To Everyone : also only redhat are doing centos stream. Where there are several options for the clones of 8

12:28:14 From Alessandra Forti To Everyone : lunch

12:31:00 From Brew, Chris (STFC,RAL,PPD) To Everyone : We can always finish it with a ceremonial coin toss!

12:31:29 From Mitchell Perry To Everyone : Best of 5, 4Tails 1Heads. Tails has it

12:32:02 From Daniel Traynor To Everyone :  
[https://www.youtube.com/watch?v=4\\_n-HiDe5hQ](https://www.youtube.com/watch?v=4_n-HiDe5hQ)

12:34:01 From Alessandra Forti To Everyone : They are arguing behind the scenes