

DESY Linux Status & Strategy

Yves Kemp et al., DESY IT & DESY DV
HEPiX Fall 2021
25.10.2021

DESY: Compute infrastructure strategy and communities

- Minimize number of central compute infrastructures
- Maximize number of communities and use cases
- Pool resources whenever possible, enable cross-usage

- Communities (random order)
 - “WLCG” community (LHC and Belle II): Grid as well as local/regional/non-Grid
 - Other HEP communities, as well as theory
 - Astroparticle community, different experiments
 - Photon science community, strong support for on-site facilities, as well as support for remote facility for in-house scientists
 - Accelerator R&D and operation
 - Helmholtz, EOSC, ... science communities

DESY: Compute infrastructure technologies

Compute infrastructure have different types of components:

- Commercial components (e.g. IBM GPFS, Auristor)
- Community components (e.g. HTCondor)
- Community components with strong own involvement (e.g. dCache)
- Own developments (e.g. ASAP::O (Photon science DAQ))

→ Maximize Scientific output, minimize manpower and resource costs

Operating system equilibrium at DESY

- Windows: Mainly administration, but also a fair amount of usage in accelerator operations and data taking
- Linux: (Some) Criteria for choice of distribution:
 - User: software availability, stable operation, right balance between stable interface and modern developments
 - Admin: stable operation incl. updates, ease of operation, reliable timeline, adequate lifetime
- Current situation (centrally supported by IT & DV):
 - Ubuntu LTS: Server & Desktops (~1.500 installations)
 - 100% RHEL Clone (CentOS 7 & Scientific Linux 7): Server incl. Compute & Storage (~5.000 installations)
 - Special purpose (original RHEL, Debian, Raspbian,...) with restricted support (~500 installations)

Our take on RHEL 8 and a potential RedHat agreement

- We acknowledge that RedHat is doing an excellent job and that we benefit from them!
- However, commercial relations restricted to few selected products
 - Enterprise commercial software requiring RHEL for support reasons (E.g. Siemens, IBM BPM, ...)
- We have evaluated and discussed with RedHat the possibility of a much larger agreement, did not follow up:
 - Scaling to all compute & storage systems would be beyond our budget
 - Licensing model (per box) induces management overhead
 - Some technology choices would taint the Kernel, thus reducing RedHat support to userland only

Should there be an agreement elsewhere that we could jump on:

- We have ~5.000 clones running, which is probably well above any radar, would need to negotiate
- Our science portfolio is rather broad, a potential deal would need to cover all disciplines
- What about science related infrastructure? ("LDAP-server")

Our take on CentOS Stream

- Not enough confidence in timely security updates
 - e.g. 8 days delay w.r.t RHEL and clones
- Not enough confidence in application upgrade policy
 - e.g. broken Podman upgrade and its versions going forth and back
 - RHEL does use CentOS Stream for testing purpose
- No support for certain commercial products
 - e.g. GPFS client availability rather unlikely (whereas some commitment for clones)
- Some sites and community codes go with clones
 - first DESY partners in Germany (Photon science only) go Rocky Linux

- Some points can only partially alleviated e.g. by internal deployment rings, would push much load on users and admins for testing

Our take on the Clones

- We consider Rocky Linux and AlmaLinux to
 - be projects that have well settled and have some prospects for long future
 - deliver rebuilds in a timely manner
 - produce stable rebuilds, based on long time experience
 - be quite similar, only minor differences
- We have a (very small) preference for AlmaLinux
 - A bit faster than Rocky
 - Availability for mirroring the repositories
- ... should there be a very strong momentum for Rocky Linux (or something bad happening to AlmaLinux), we are probably ok to change here
- no, we will not consider Oracle Linux

Our take on something completely new

- Need to further leverage use of containers to separated host OS from user land
- otherwise: No new OS considered (Debian, Arch, Gentoo, Plan 9, BSD...)

Status

- Disclaimer: We only have small number of CentOS 8 systems deployed
 - Basically some infrastructure nodes, no user machines
- AlmaLinux 8:
 - Repositories mirrored
 - Ready for installation
 - Observation:
 - Base server works
 - Some Puppet modules require “CentOS” ... modules need to be adapted (also 3rd party modules)
 - Testing general server with users will start
 - Testing Maxwell HPC compute nodes with users will start
 - Migration plan CentOS 8 → AlmaLinux 8 still to be defined (probably reinstallation)
 - Migration plan WLCG & HEP: Need to coordinate with users and VOs once future is clear

Wishlist to WLCG

- What we do not want:
 - to be forced to purchase RHEL licenses for all of our cluster to comply with only one customer (WLCG)
 - to be forced to roll out another distro, that we seem not fit
- What we want:
 - RHEL 8 *clones* to be accepted as a supported platform for future WLCG Linux environment
 - on the same level than CS 8 or RHEL 8
 - We can offer to run containers build e.g. by the application developers (middleware services, VO compute jobs)
- We understand that WLCG cannot support a plethora of distros.
 - Clones should be close enough to RHEL and CS (in contrast to e.g. Debian or SUSE)

... and two words about Scientific Linux

- Scientific Linux will not resurrect. We understand the reasons.
- Scientific Linux is not dead. SL 7 will stay with us for some more years. But the end is in sight.
- Some things we will miss from the overall excellent work of the SL team:
 - faster errata
 - separation of security updates and others
 - continuation of security updates for older minor releases
- One thing we will surely miss is the community, seen every day via its mailing list.
 - Maybe HEPiX can provide for a (e.g.) redhat-family-users@hepik.org mailing list. Maybe open to the public, but with a certain professional level, low noise and with subjects of interest to our community.
 - ... or is there a wish to have a closer cooperation within HEPiX in the finding and running of \$FUTURELINUXDISTRO and offering shelter for the community?