Linux at DESY
- The RHEL & clones situation
- Other distros

Yves Kemp et al., DESY IT
Hepix Fall 2023

HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

DESY.

The clusters and distributions

- Central IT managed
  - Currently mixture of CentOS Linux 7, Ubuntu LTS and some Debian
  - Grid & NAF clusters & services; Maxwell HPC cluster (RHEL clones)
  - General purpose machines (mix)
- Accelerator controls (FLASH, XFEL) (Petra-III: Windows)
  - Currently Ubuntu LTS (every other LTS, 16.04 → 20.04 (→ 24.04 ?)
  - Planning to move to Debian
- Experiment controls (Petra-III, FLASH, XFEL)
  - Lots of Debian: Control software (Tango…), support also for older machines (crates…)

- Side note: Accelerators and experiment controls integrate into their worldwide communities. They are often grounded in the Debian/Ubuntu world. ESRF is using mostly Debian (AFAIK).

A closer look at the IT managed RHEL&clones machines _ 1

- past/present: basically, mostly running CentOS Linux 7, and initial plan migration to AlmaLinux 9 ... but:
- future: general concerns about AlmaLinux 9 long-term stability

- NAF & Grid cluster:
  - applications rely on RHEL or clones. We offer containers in batch (Apptainer), but unclear whether Ubuntu or Debian + Container would do the job
  - prefer staying in the RHEL/Clones world for some time

- NAF & Grid services:
  - e.g. HTCondor backends: Also available for Debian, Ubuntu, as well as Docker images
  - Grid services: Only RPM available. Maybe containerized deployment on another distro?

A closer look at the IT managed RHEL&clones machines _ 2

- Maxwell HPC Cluster (mostly used by photon science and accelerator R&D)
  - need a GPFS client. Officially: RHEL, Ubuntu, SLES.
  - So far, 1:1 clones worked technically – though not officially supported by IBM
  - ABI compliance of AlmaLinux might break GPFS client.

  → Need either RHEL subscriptions, or Ubuntu (no Debian since no GPFS supp.) (looking at the known DESY distros)
  - Pro RHEL: We heavily rely on NFS kernel client, excellent work from RedHat for this component in the past
  - Pro Ubuntu: More software available for photon & accelerator R&D

- General purpose machines: e.g. Group web server, Experiment-own database, ….
  - Need a case-by-case decision
  - Sometimes also hard requirements w.r.t RHEL/Clones dependency

  → In the past, the 10-years support was happily accepted, especially for VMs.

  → Users should prepare to a much shorter lifecycle, be it with AlmaLinux 9/10/11… , Ubuntu LTS or Debian+LTS

What happens to AlmaLinux? The Clones? My personal view:

- We opted for AlmaLinux as a clone rebuild from RHEL
- So far were very happy with this decision, no regrets so far!
- But: Alma chose to create an enterprise distro from *CentOS Stream*
  - Valid choice, given RedHats restriction to access RHEL source code // Rocky & Oracle (&OpenELA)
  - Sooner or later, GPFS modules will no longer compile
  - Sooner or later, there might be a situation where issues cannot be fixed easily in AlmaLinux, since distance is too far from CentOS Stream. (At the moment, Alma does a good job, and sometimes is even faster with some fixes compared to RedHat)
  - Alma N relies mostly on CentOS Stream N, which has ~only half the lifetime of RHEL N. What will Alma N look like when CentOS Stream N+1 is out? … we might learn this for Alma 8 after 2024-05-31 (end of CentOS Stream 8)
- Other clones take from RHEL source (however they achieve to do so)
  - GPFS; fix situation & lifetime are probably OK
  - There is danger of license violation, and legal actions against other clones
- OpenELA
  - currently no alternative, since only ABI compatible … but might become an interesting alternative in the future
- **Usage of clones in production must be carefully followed**

Containers to the rescue? … my personal view

- Maybe Yes … but who builds them? And how?
- Naïve idea:
  - A couple of central RHEL machines compile software (experiment, WLCG middleware, …)
  - Sites run $PREFERRED_DISTRO, and execute these containers (or users do so via the batch system)
- Yes, but have a look at the details on: https://developers.redhat.com/articles/ubi-faq
- My keywords:
  - Red Hat Universal Base Images (UBI) only has a very limited set of RPMs that allow running without subscription
  - If other RPMs from RedHat are used, then *running* UBI requires subscriptions!
  - Special care must be taken by the builders of UBI.
  - Sites must inspect each container and check compliance.
- Centrally building RHEL-based UBI containers is feasible. What about user-compiled code?
- Sites executing these containers take a risk of licence violation when containers are not correctly built.
- **RHEL UBI containers can only be part of a short-term migration strategy to a non-RHEL world.**

What the DESY Linux distro future could look like:

- short-term: Need to upgrade the CentOS Linux 7 machines!!!
  - In some cases, no alternative to RHEL or clones at the moment.
  - RHEL? AlmaLinux? ... ?
- mid/long-term:
  - Migrate everything that has no hard RHEL dependencies to Ubuntu or Debian
    - The photon science, accelerator R&D and theory community probably are OK, and might even welcome this step
  - Where there are dependencies on RHEL:
    - Can they be solved in a secure and sustainable way using containers, without subscriptions?
    - Can we minimize the use of subscriptions to a minimum, e.g. just some portal machines
    - ... probably the LHC analysis code is the hardest part here
- Moving away from the RedHat world will be a change, not only technical:
  - Long-term-supported, enterprise distros were/are much appreciated