

Linux Future(s) at DESY

WLCG/HEP, other Sciences, etc.

Yves Kemp, Thomas Hartmann and many more
GDB 2024.Feb.14



Support by central IT


Situation so far

- Supported RedHat (mostly derivatives ~ CentOS, some limited RHEL)
 - used for most of science computing, incl. Grid, dCache, NAF, Maxwell-HPC
 - used for most services around science (Wikis ,...)
 - used for some enterprise computing
 - -> 10 years of support



- Supported Ubuntu (LTS releases)
 - used for some some services around science
 - used for some enterprise computing
 - used for desktops
 - -> 5 years of support



- Limited support for Debian 
 - limited to groups with own expertise
 - -> effectively ~5 years of support

The Future

(future horizon ~5a)

- Support for RedHat RHEL



- made possible via 5y campus subscription
- aiming to channel users to RHEL 9, leaving out RHEL 8

- Support for AlmaLinux



- hope to channel to AL 9
- CentOS Stream 8 being close to EOL: following development of Alma8 closely as benchmark for Alma9

- Ubuntu LTS



- Debian: ... keeping the limited support



General Conclusions

Our message to our users

- Expect 3-5a life cycle of an installation
- Uncertainties >5a
 - RHEL: potentially expensive if no follow-up attractive campus agreement
 - Alma: reliability development/support process
- Foremost lessons learnt (admin perspective):
 - Trying to avoid lock-ins
 - Have to be more proactive in the future
i.e., switching base OS more readily if conditions change
- **Plan and prepare for more frequent changes**

Users

User Recommendations

Users/non-central IT service admins

- Users/Admins have to evaluate, which distribution fits best their use case
 - Can they decouple from central IT base OS (containers?)
 - Have to ensure ongoing ops
 - When couple to base OS with IT support
 - If no specific requirements: Ubuntu
 - Ubuntu features:
 - Support available for ~5a
 - not better/worse than *EL flavour uncertainties
 - Reasonable software support
 - (Debian also ~<5a support; hardware support occasional difficult)



Clusters

Compute Clusters: RHEL9

HTC and HPC



- Traditionally Redhat-flavoured oriented
- HTC: no special HW
 - Grid HTC: all workloads (👉) containerized independent of base OS
 - NAF HTC:
 - WLCG VO to provide user environments/containers (users ready...?)
 - Smaller user groups still running native base OS
- HPC: Hardware requiring enterprise OS (infiniband & GPFS) – excludes Alma/Debian
 - Maxwell HPC: somewhat similar to NAF wrt. users → even more broader communities from Photon etc. fields
 - No dependency on user code → keep eyes also on Ubuntu to gain experiences wrt. NFS, InfiniBand, GPFS
- All heavily relying on NFS: classical Redhat field, doing a better job than Ubuntu, e.g., backporting into LTS

Storage Clusters

dCache & GPFS Instances

- dCache w. Java not as OS sensitive - but
 - RHEL 9
 - Support can be utilized
 - NFS
 - A number of pool nodes on Alma8/9 from the transition period
- Recently migrated ~500 nodes EuXFEL dCache instance
 - O(1week)
 - Massive migration is possible with reasonable pain
- GPFS/SpectrumScale
 - Hard client dependencies on enterprise OS
 - → RHEL9 (user facing) or Ubuntu (backend)




(Current) Transition

Transition Pains & Middleware Readiness

Ongoing EL7 → EL9

- **Only 4.5 month left until EL7 EOL**
 - **Challenging to debug and prepare the whole system parallel to daily operations**
- All the small stuff (python 2.7 remnants, SELinux,...)



- **Post-EL7 readiness of critical HEP middleware unclear**
 - Needs to be integrated and tested
 - 4.5 month left sufficient?
 - **Lesson? separate middleware requirements from our base OS**

Conclusion

After the migration is before the migration

- we/central IT have to be able to adapt to changing conditions
 - Have to learn from the current pains...
- Users have to realize the same
 - In step with central IT base OS? → bound to our schedule
 - Decouple from central IT base OS? → cultural shift
 - Same for middleware → decouple application requirements from HW requirements