

Stephan Wiesand
Technical Seminar
Zeuthen, 2017-05-09




Agenda

- > Linux @ Zeuthen
- > Recent or upcoming changes
 - SL5 EOL
 - SL6 → EL7
 - > except on desktops
 - SL6 → Ubuntu 16.04 LTS
 - > only on desktops
 - New desktop concept
- > Singularity




Linux @ Zeuthen

> ~ 1100 systems




Linux @ Zeuthen

> ~ 1100 systems



Linux @ Zeuthen


> Infrastructure, dissected



Linux @ Zeuthen

> Operating Systems

■ all x86-64



Changes: SL5 End of Life

- > RHEL5 and thus SL5 went EOL end of March
 - No special announcement
 - > there had been no more desktops, WGS, compute nodes for a while
- > Remaining systems must be replaced or retired a.s.a.p.
 - DV, PITZ



Stephan Wiesand | Linux Update | 2017-05-09 | Page 7

Changes: SL6 → EL7

- > “EL” ?
 - “Enterprise Linux”, generic term for Enterprise distributions derived from RHEL (public sources)
 - > Scientific Linux or CentOS or ...
 - > which one doesn't matter for users
- > Public preview system available: el7.zeuthen.desy.de
 - Please try & test
- > Groups are welcome to request WGS updates
- > Farm nodes will be upgraded step by step
 - in agreement with stakeholders
 - `qsub -l os=sl6`
- > The new pax nodes (in testing) are all EL7 already
 - and managed by SLURM rather than UGE



Stephan Wiesand | Linux Update | 2017-05-09 | Page 8

EL7: What's new?

- > For users: not that much
 - newer versions of much software
 - we expect to be able to support the AFS client for the EL7 life time
 - > sysname list starts with `amd64_rhel70` now
 - relying on AFS sysnames is deprecated
- > Known issues:
 - we just about finished EL7 support, and it may need some polish
 - > some software may still be missing
 - no more implicit linking of shared libraries
 - > suppose `libX.so` was linked against (depends on) `libY.so`
 - used to be sufficient: `gcc mycode.c -lX`
 - now required: `gcc mycode.c -lX -lY`
- > As ruled by the desktop commission, EL7 will not be available as a desktop OS.



Stephan Wiesand | Linux Update | 2017-05-09 | Page 9

Changes: SL6 → Ubuntu

- > Now available for desktops: Ubuntu 16.04 LTS
 - The same caveat applies as for EL7: just finished, may need polish
 - New desktops: ask for Ubuntu on the request form
 - Existing desktops: send upgrade request to uco
 - > for the time being – this will become unnecessary soon
 - > Precision 390 (“oreade”) and older systems are not supported
- > Lots of changes w.r.t. SL6 of course
 - Default (and supported) desktop environment: Unity
 - Home directories are local and not backed up centrally
 - > Recommendation: user should arrange a dejadup backup
 - Late news: We do install an AFS client after all
 - > Support by the OpenAFS project has improved
 - > But no guarantee that it will be available and work forever
 - > No special AFS sysname (using the default `amd64_linux26`)



Stephan Wiesand | Linux Update | 2017-05-09 | Page 10

More on Ubuntu 16.04

- > On SL6, we permit installing additional software packages
 - Install only, not remove – and there's a blacklist
- > On Ubuntu, we also permit removing packages
 - and there's no blacklist
- > => More ways for users to hurt themselves
 - supported remedy: reinstallation from scratch
- > No more automatic updates during the boot process
 - automatic updates still happen at night
 - systems shut down when new updates become available will prompt the user to install them
 - > or install them the next time they run over night
 - users will also be prompted to reboot the desktop after certain updates (kernel, libc, ...)



Stephan Wiesand | Linux Update | 2017-05-09 | Page 11

New Desktop Concept

- > In the past, desktop (re-)installation required action by DV
 - change in the configuration database
 - preparation of installation process
 - > disabling all automatic maintenance
- > “New style” desktops can be (re-)installed by any local user
 - configuration is now generic: “it’s a desktop”
 - PXE (=network) boot is unlocked
 - upon PXE boot, our servers will offer a menu with the available OSs
 - > currently: SL6 and U16.04
- > We plan to support a desktop OS for one year after the availability of a successor.
 - > SL6: one more year
 - > U16.04, U18.04, ...: ~3 years



Stephan Wiesand | Linux Update | 2017-05-09 | Page 12

And now for something completely different...


- > Transition periods are painful
 - SL5 gone
 - mixed farm
 - mixed WGS
 - very different desktop



Stephan Wiesand | Linux Update | 2017-05-09 | Page 13

And now for something completely different...

- > Transition periods are painful
 - SL5 gone
 - mixed farm
 - mixed WGS
 - very different desktop



Stephan Wiesand | Linux Update | 2017-05-09 | Page 14

Containers

- > A lightweight way to provide a runtime environment for software
 - almost independent of the host operating system
- > Typically, aiming for maximum isolation



Stephan Wiesand | Linux Update | 2017-05-09 | Page 15

Containers

- > A lightweight way to provide a runtime environment for software
 - almost independent of the host operating system
- > Typically, aiming for maximum isolation
- > Popular example: docker



CC BY-SA 3.0, <https://commons.wikimedia.org/w/index.php?curid=1923693>



Stephan Wiesand | Linux Update | 2017-05-09 | Page 16

Containers

- > A lightweight way to provide a runtime environment for software
 - almost independent of the host operating system
- > Typically, aiming for maximum isolation
- > Popular example: docker
 - alas:
 - > geared towards “micro services”
 - > not suitable for HPC
 - containers spawned by daemon
 - > resource management
 - > AFS PAGES & tokens



CC BY-SA 3.0, <https://commons.wikimedia.org/w/index.php?curid=1923693>



Stephan Wiesand | Linux Update | 2017-05-09 | Page 17

Singularity

- > developed at LBL, goals
 - “mobility of compute”
 - “reproducible science”
 - allow users to bring their own execution environment
- > geared towards HPC/batch
- > aims for minimum isolation
 - processes in the container have the same permissions as outside
 - > including AFS tokens
 - > no privilege escalation in containers allowed (suid)
 - many things just work (X11, ...)
 - resource management unchanged



CC BY-SA 3.0, <https://commons.wikimedia.org/w/index.php?curid=463020>



Stephan Wiesand | Linux Update | 2017-05-09 | Page 18

Singularity – how to use

- > A preliminary set of “standard” images is available:

```
% ls /afs/zeuthen.desy.de/project/singularity/images
SL-5.11-2017-05-05.img
SL-6.9-2017-05-05.img
SL-7.3-2017-05-05.img
U-16.04-2017-05-06.img
```

- > Images are executable, and by default run the user’s shell

```
% lsb_release -d
Description: Red Hat Enterprise Linux Server release 6.9 (Santiago)
% /project/singularity/images/U-16.04-2017-05-06.img
% lsb_release -d
Description: Ubuntu 16.04.2 LTS
```

- > One can directly invoke arbitrary applications too:

```
% singularity exec U-16.04-2017-05-06.img gcc -v |& tail -1
gcc version 5.4.0 20160609 (Ubuntu 5.4.0-6ubuntu1~16.04.4)
```

- use U16’s evince to view a brand new PDF on SL6
- run atom on SL6



Stephan Wiesand | Linux Update | 2017-05-09 | Page 19

Singularity: a few details

- > Containers are immutable
- > Filesystems shared with the host:
 - default:
 - > /tmp \$PWD \$HOME
 - > /proc /sys /dev (incl. /dev/shm)
 - we add:
 - > /afs /lustre /cvmfs /acs /batch (where appropriate)
 - user can specify additional binds
- > No limitation to a “runtime environment”
 - build
 - develop
 - debug



Stephan Wiesand | Linux Update | 2017-05-09 | Page 20

Singularity: Caveats

- > While it's available on desktops, compute nodes and WGSs now...
- > The "standard" images are certainly not final
 - need user input, more tailored images may be required for batch
 - some oddities to be sorted out yet (shell environment)
 - NB don't like our images? - just bring your own
- > Issue with automounting filesystems
 - mount must be triggered outside the container on SL6 (works on EL7)
- > dbus woes (the newer the software the worse)
- > some applications would require compatible software in- & outside
 - CUDA, MPI across nodes
- > while AFS works fine, the sysname is always that of the host



Stephan Wiesand | Linux Update | 2017-05-09 | Page 21

Summary

- > EL7 available now
 - for WGSs and compute nodes only
- > Ubuntu 16 available now
 - for desktops only
 - "New" desktops grant more control to the user
 - > Including installation of a choice of OSs, anytime
- > Singularity available now
 - on desktops, WGSs and compute nodes
- > Users are encouraged to give all those a try
- > All are brand new (in our environment)
 - Early adopters will find additional problems



Stephan Wiesand | Linux Update | 2017-05-09 | Page 22