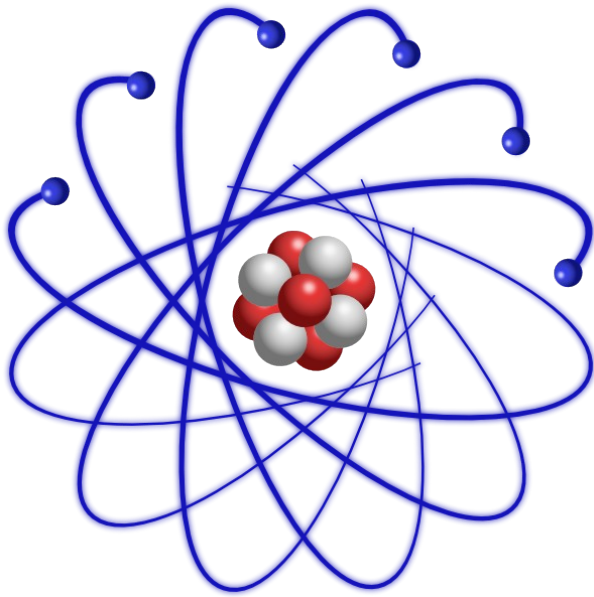


# Scientific Linux 6.

at DESY, location Zeuthen



Stephan Wiesand  
Technical Seminar  
Zeuthen, 2012-01-10

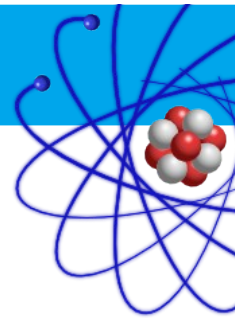


## Introduction

- > 4.5 years since the SL5 introductory presentation
- > Linux in Zeuthen is moving on
- > SL4 EOL is three weeks from now
  - We'll switch off the few remaining systems
- > SL5 approaching end of “full support phase”
  - The next hardware generation coming soon
  - It will *probably* still work with SL5 (at least 5.8, at least on servers)
    - > but we don't know yet
- > SL5 EOL in about 2 years
- > SL6 is available and considered ready for general use



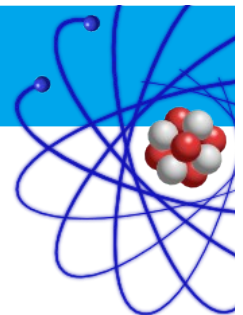
# Outline



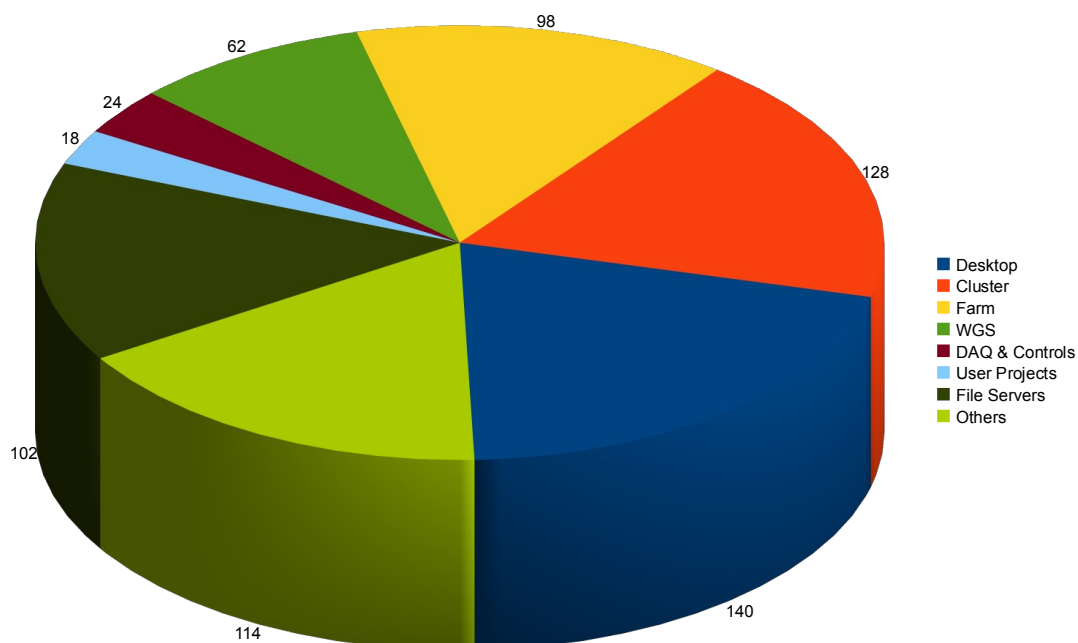
- > What's this about?
  - Linux in Zeuthen
  - RHEL/SL
- > What's new with SL6?
  - And what isn't?
- > What's next?
  - New options
  - New requirements?
  - New ideas?

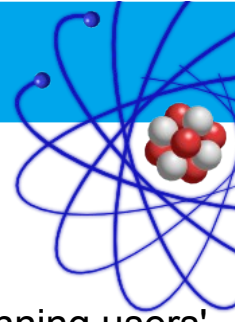


## Linux in Zeuthen



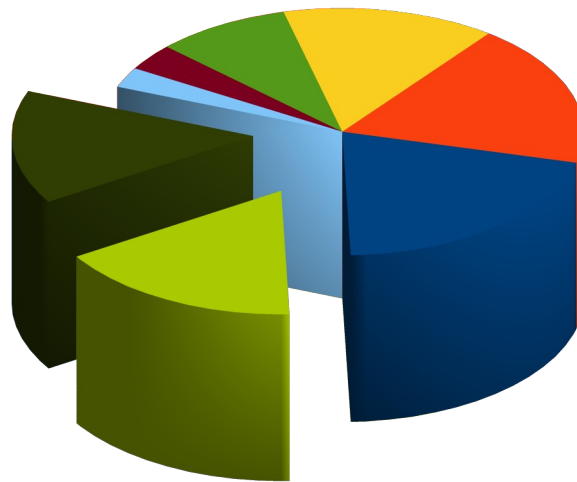
- > 687 live systems (counted 2012-01-05)
  - omitting ~ 200 Grid/Tier2 and ~ 70 NAF systems





serving users'  
data

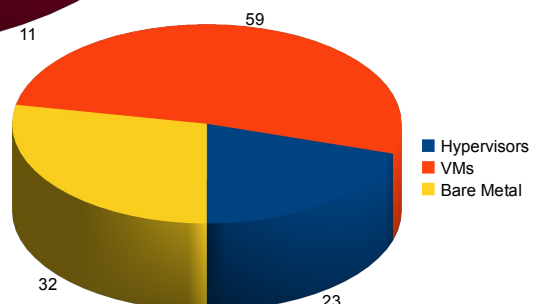
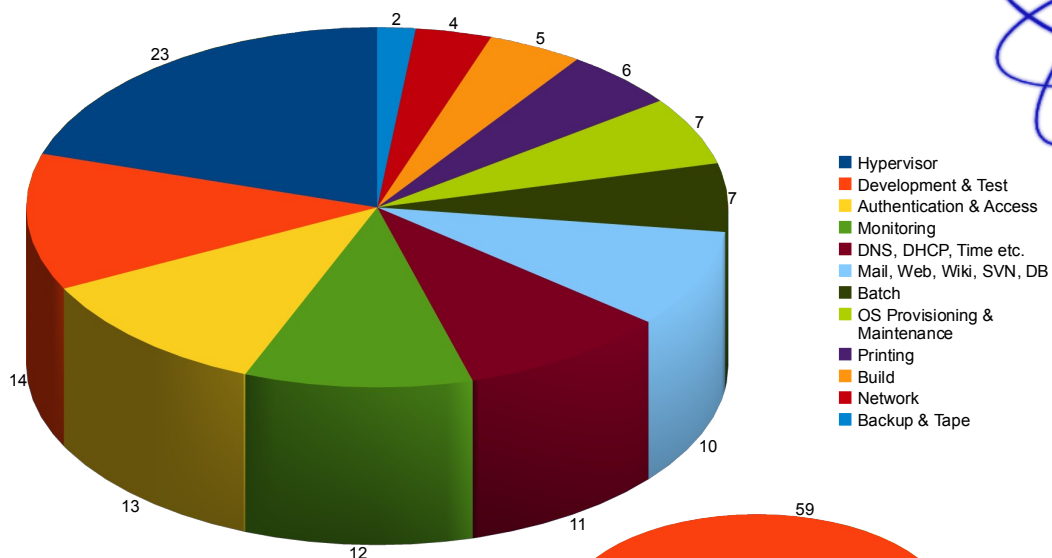
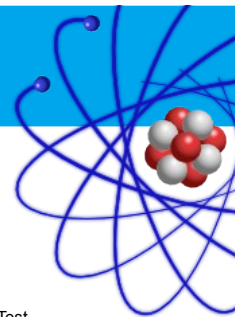
running users'  
software



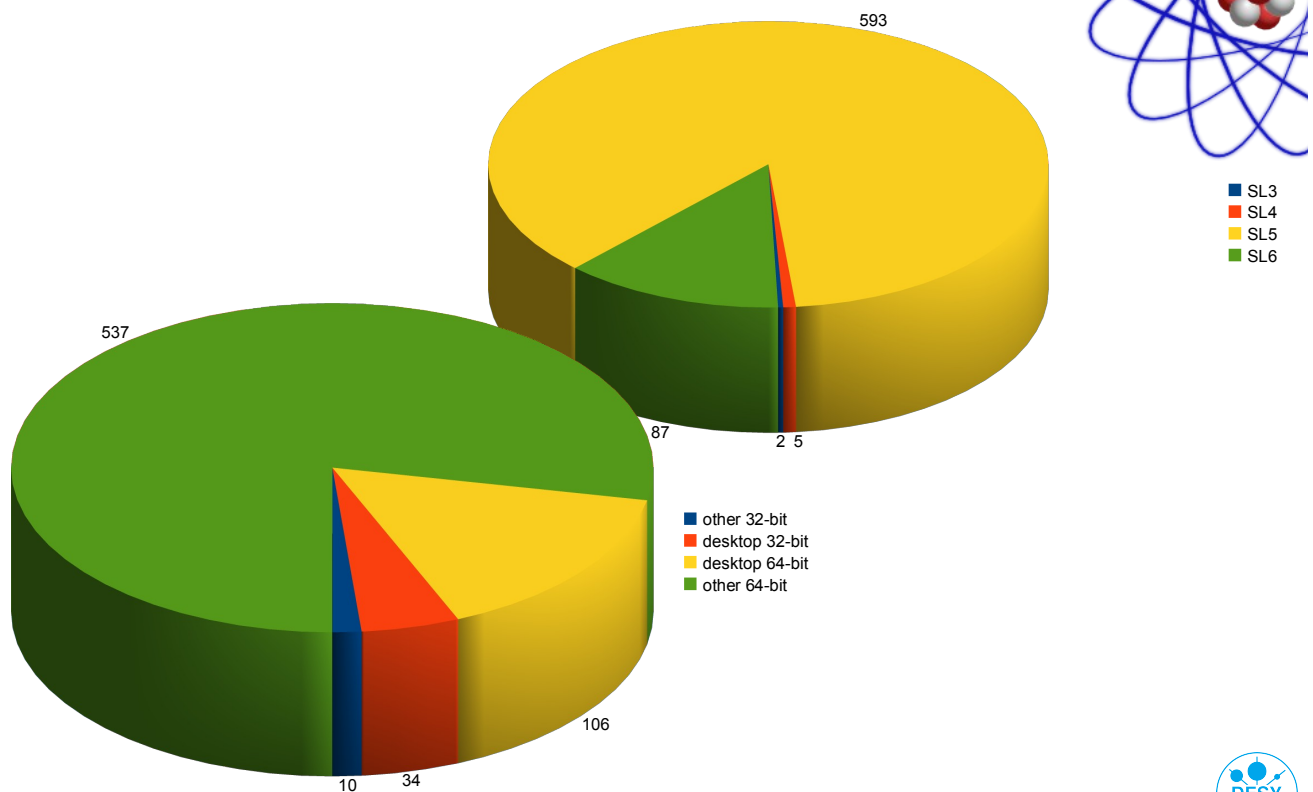
providing services  
&  
running the CC



## Others, dissected



# Operating system flavours



Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 7



# Systems management

- > All systems are centrally maintained in a **uniform** way
  - Configuration database (VAMOS)
  - Agents (cfengine, scripts, packages) work on the systems
  - Admins work on the configuration and the agents, not the systems
- > This keeps systems **consistent**
  - In a known, well defined, reproducible state
  - It often alleviates the need for backup, even on servers
- > No **root access** for users (with exceptions where required)
  - It compromises the privacy and security of other users
  - It makes it easy to break our maintenance mechanisms by accident
    - > unless you're very familiar with them
  - It tends to bring systems to a state not known and defined
    - > Harder to debug
    - > Much harder to replace or reinstall

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 8



## Reminder: Scientific Linux

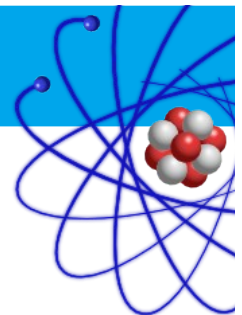


- > **Rebuild** of the free sources of Red Hat Enterprise Linux
  - As far as they're available in a public place
    - > ftp.redhat.com
    - > This is not true for several RHEL add-ons
      - Extended Update Support
        - > Security patches for previous minor releases (now called “service packs”)
          - backported
        - > SL adds such support for all minor releases
          - but security fixes always come from the latest minor release
            - > no backports
      - Extended Life cycle Support
  - > SL is supposed to be **binary compatible with RHEL**
  - > Creating such a clone is not trivial, and a lot of work
    - **FNAL** doing a very good job – Thanks!
  - > Other choices for RHEL clones:
    - CentOS, SLC, PUIAS (all free); OEL (fee)

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 9



## RHEL Life Cycle



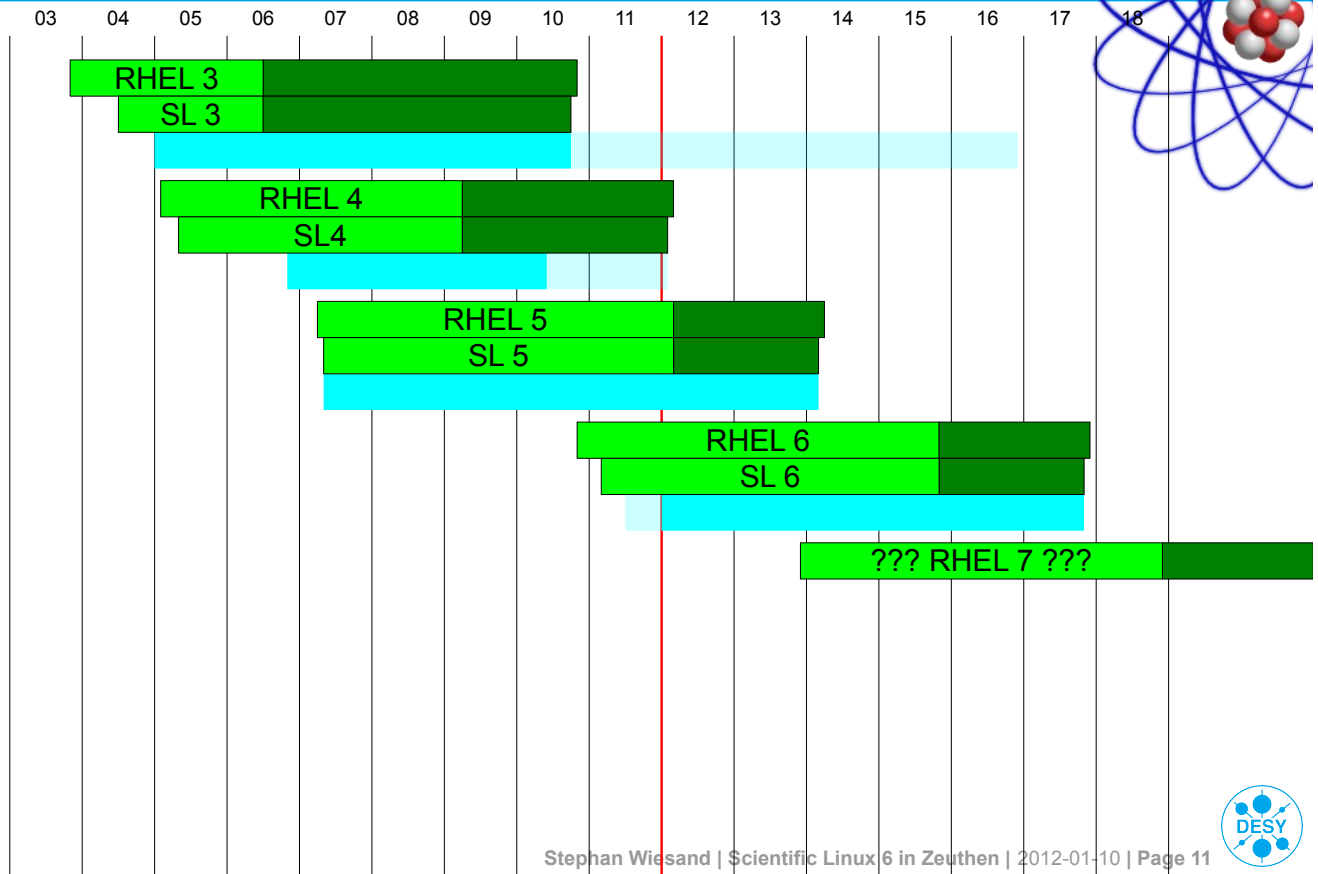
1	2	3	4	5	6	7	8	9	10
Production 1				2	3		ELS		

- > “Production” phases
  - 1: Minor releases (every 6-8 months)
    - > Bug fixes, enhancements (new software, new versions)
    - > Support for new hardware
  - 2: Minor releases
    - > Bug fixes
    - > Limited support for new hardware
  - 3: Security and other critical bug fixes
    - > Support for running as VM under current major release
- > Extended Life Cycle phase
  - Like production 3, for an additional fee

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 10



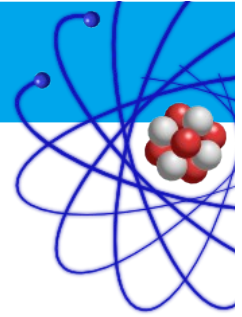
# Timeline



## Why are we late?

- > ~4 years of upstream development from EL4 to EL5
- > => Many changes under the hood – examples:
  - different output from `uname -r`
  - dumbed down gdm (really hard to display the hostname w/o ifh.de)
    - > and before you ask: no, we can't enlarge the font...
  - retrieval of per-host kickstart file via http broken
  - hypervisor changed from Xen to KVM
- > We wanted it to mature a bit (waited for 6.1)
- > We wanted to use OpenAFS 1.6
- > We're using a new (much better!) way to package kernel modules
- > SL itself took a while
- > The missing Lustre client was a showstopper for a while
- > There was no pressing need, and limited manpower

## What's new for users



- > (Few to) **No spectacular changes**
- > GNOME looks & feels very much the same as on SL5
- > Firefox *is* the same version as on SL5
- > But **improvements** all over the place
  - evince works much better
    - > This helps avoiding Adobe Reader
      - Critical issues, being exploited in the wild, unfixed for weeks/months on Linux
        - > This is the current state (advisory December 6<sup>th</sup>, fix announced for today)
    - > By the way, we recommend giving **mupdf** a try
      - lean & fast, displays very well, even problematic PDFs
      - lacks printing, forms support, and a GUI though
  - **AFS client** should be much faster for cached data
  - Much improved **power saving**, especially for idle systems
  - Expect **performance** to improve by a few % on average

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 13



## New default locale



- > Past default: LANG=C
- > SL6+ default: LANG=**en\_US.UTF-8**
  - This can cause all kinds of interesting effects
    - > especially when the terminal is running under a different locale than the application
      - Remote logins
  - UTF-8 can slow down software
    - > and make it use more (virtual) memory
  - **Recommendation: Run batch jobs under LANG=C**
- > Users can set their personal default in **~/.i18n**
- > We set LC\_PAPER to de\_DE.UTF-8 by default
  - Makes many applications use A4 format

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 14



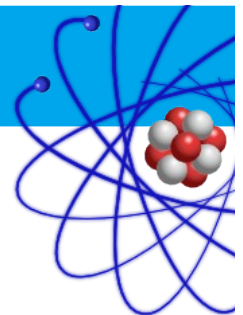
## New features (1)



- > On desktops, users can **install additional software** packages
  - from the distribution, and our add-on repositories (only)
  - Caveats:
    - > There's a blacklist
      - Not everything can be installed, even if visible
    - > The blacklist is probably incomplete – please be careful
      - Refrain from adding low level packages that could do harm
        - > If you spot those, please let us know
    - > Installation is a one way road: Removing packages is not possible
- > Display configuration relies on card/driver/monitor autonegotiation
  - This should always set the right mode
    - > except when it doesn't
      - Expect problems with very old or slightly broken hardware



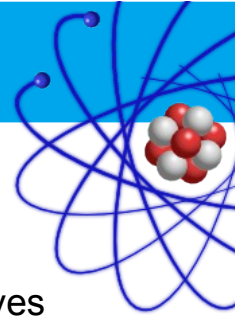
## New features (2)



- > **sshfs** - Mount a remote filesystem through ssh
  - `mkdir /mount/point`
  - `sshfs my-wgs:/some/fs /mount/point`
  - `ls /mount/point`
  - `fusermount -u /mount/point`
- > **CUDA** is installed on systems that support it
  - That's the latest desktop model (T3500) only
    - > And it's a low end card, not supporting all features
  - It may still help getting started with GPU computing
  - Serious GPGPU systems for general use will become available soon



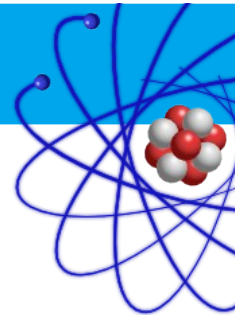
## SL5/6 software compatibility



- > **Binary**: Software built for SL5 *should* work on SL6
  - All available compat-\* packages are installed
  - If some are missing, we try to create and add them ourselves
    - > This was recently done for mysql
- > **Source**: GCC
  - SL5: 4.1 default, also available:
    - > gcc44, g++44, gfortran44
    - > gcc34, g++34, g77
  - SL6: 4.4 default, also available:
    - > gcc34, g++34, g77
    - > no gcc41
  - 4.4 is meant to be the link



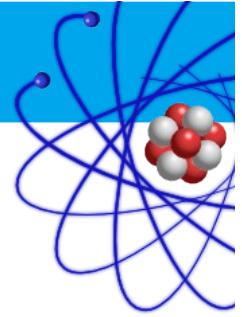
## Application software



- > General policy: Use what comes with the **distribution**
  - even if it's not the latest and greatest
  - Next, try **3<sup>rd</sup> party repositories**, preferably EPEL
  - If no luck with those, and effort is justified, next try:
    - > **Rebuilding** a package from another distribution (Fedora, SuSE)
    - > **Building** and packaging ourselves
- > Preferred: Distribution as an ordinary package, installed locally
- > Few exceptions: Very **bulky software, typically commercial**
  - Maple, Mathematica, Matlab, Intel & PGI compilers
  - Now located in **/opt** instead of /opt/products (which is basically gone)
    - > Usually a link into AFS, except on special systems
  - No default versions – run maple15, math8.0, icc2011, ...
  - Use of **ini** is being replaced by **environment modules**



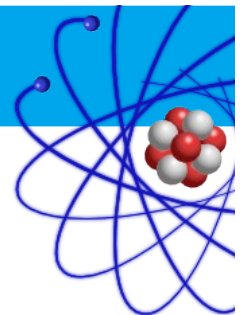
## Software: Other changes



- > Firefox is now the 64-bit version
- > So is the local **Java** (and still JDK 6)
  - oops: **JDK 6 EOL is in July** – yes, this year
  - Poll:
    - > Who needs JDK 7 now, or will need JDK 6 after July?
    - > Any thoughts on OpenJDK vs. Oracle?
- > freemind → **vym**
  - Poll: Who is using mindmapping software?
- > No **scribus** yet
  - Poll: who's using this?
- > nvu → **kompozer**
  - Alas, this seems dead now as well
  - Poll: Who knows what nvu/kompozer are?



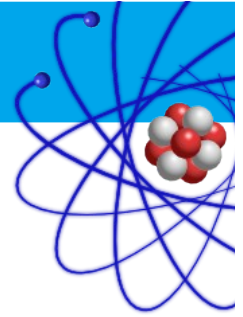
## Software no longer installed



- > Some software is probably still missing
- > But some is deliberately not installed
  - **KDE**
    - > Users can install it on their desktop
      - Some components are not available though
        - > no kdepim – the akonadi backend doesn't work with \$HOME in AFS
  - **root**
    - > This also affects SL5 already
      - No way to satisfy all users at the same time
      - Per-group installation rather simple
      - (almost?) All groups are doing this anyway
  - **cernlib**
    - > Has been declared dead for a decade now
    - > SL6 build has problems
    - > **Workaround**: Use the SL5 build in AFS



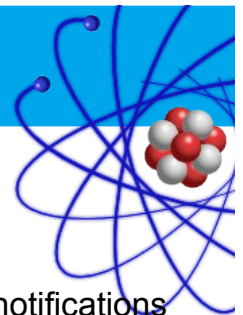
## What hasn't changed



- > Basic user environment
  - Shell profile customizations
  - bash is still not available as login shell – zsh and tcsh only
- > Fonts added
  - Free MS truetype core fonts
  - Bitmap fonts and aliases from the early Unix days
- > Legacy software added
  - gv, plan, xv, xcalc, xdvi, ...



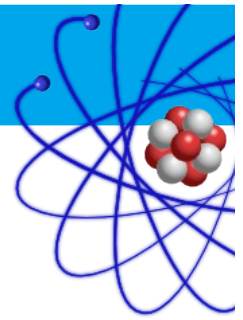
## More information



- > [https://dvinfo.ifh.de/SL6\\_User\\_Information](https://dvinfo.ifh.de/SL6_User_Information)
  - By the way, it's a Wiki
    - > Users can create an account and then subscribe to change notifications
    - > Participation is also welcome
      - Ask us to add your account to EditorsGroup
- > Technical notes on what we did to make it work, open problems, ...
  - [https://dvinfo.ifh.de/SL6\\_Development](https://dvinfo.ifh.de/SL6_Development)



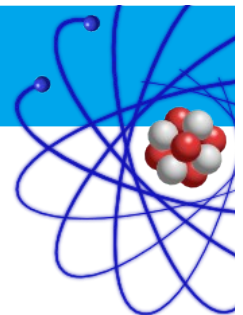
## Public preview systems



- > sl6.ifh.de
  - WGS setup
  - VM with 2 cores and 6 GB of RAM
  - Meant for testing, not running production jobs...
- > nomos23
  - One of the public **desktop** systems in 2L01
- > **farm** nodes
  - `qsub -l os=sl6`
  - `qssh -l os=sl6`
  - If your jobs work on both SL5 and SL6: `qsub -l os='sl5|sl6'`



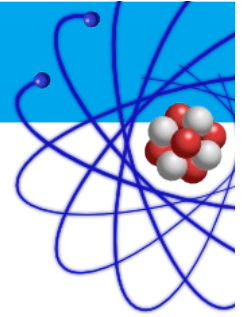
## Migration Plan



- > Desktops
  - Install SL6 by default on new systems
    - > Effective immediately
    - > SL5 is still available – Please note on the request form
  - Users / group admins may request upgrade (= remote reinstallation)
    - > Mail to uco-zn (host name, when, who will reboot the system)
- > Farm
  - Available today: 8 x 12 Cores 3 GHz Westmere, 4GB/Core
    - > + 2 older systems, often used for tests
  - Migrate more nodes when
    - > it's requested by major stakeholders
    - > it's required to provide assigned shares for jobs requesting SL6 nodes
    - > SL5 nodes are idle & SL6 nodes busy for a while
- > Anything else: Case by case



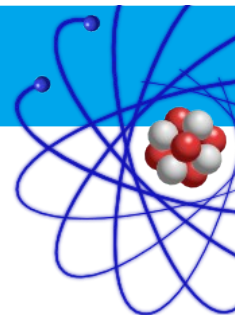
## Milestones reducing SL5 support effort



- > The last “dryade” desktop (5)
- > The last 32-bit SL5 desktop (+ 29)
- > The last 32-bit SL5 system (+ 10)
- > The last SL5 desktop
- > The last SL5 WGS & farm node



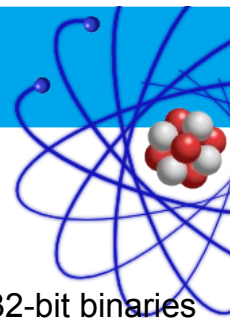
## Caveats



- > Early users *will* find bugs and omissions
  - Fixing those will have high priority now
    - > much higher than next year
- > Lustre client
  - Changes in the RHEL 6.2 kernel broke it – this was not anticipated
  - No visible activity to fix it
    - > and we failed (maybe it's doable with high effort)
  - Currently, it's not critical to install this kernel on lustre client systems
  - With the next security update, this may change
    - > In the worst case, we could be forced to uninstall lustre clients
      - or hope for the best and reinstall all systems later...
- > Breaking news: kernel crashes when uptime approaches 200 days?
- > Policy changes



# Policy changes



## > Only 64-bit systems

- The full set of 32-bit compatibility packages is installed
  - > including the development packages, compilers can create 32-bit binaries
- Desktops older than the Precision 370 (from 2004) must be replaced
  - > The Precision 370 itself has not been tested (lately)
    - and will have limited graphics support

## > The local firewall is enabled on all systems

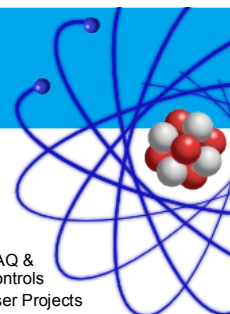
- All ports are closed by default
  - > This breaks qrsh on desktops
    - A large range of ports is required to be open for farm nodes – this is only configured on the WGS
  - > It requires tunneling VNC through ssh
    - so far, this has just been a very good idea - that's a clear text protocol
  - > Running services may require requesting further port openings

## > The automounter will only mount what's required, not what's available.

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 27



# Desktops



## > The single largest class of systems

## > Actually, not homogeneous:

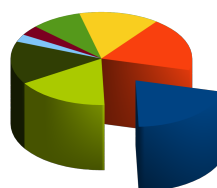
### ■ Group desktops:

#### > ssh login (WGS-- + local display)

- Must be limited to owning group
  - > Or you'll find other groups' diploma students abusing your CPU and disk
  - > Assignment to group is sometimes fuzzy
  - > And sometimes, it changes frequently
- Potential security issues
  - > Access to audio & video devices may be achievable remotely
    - We have not yet figured out how to restrict access to /dev/nvidia\* on SL6

#### > Persistent local storage

- This actually makes it an individual, stateful system – like a server:
- Hard to replace, urgent to repair
  - > Yet, cheap disk, no redundancy
  - > Some users understand what /usr1/scratch is suggesting
    - Others don't
- Requires ssh login to access data – and a running system

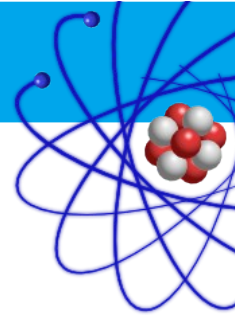


■ Desktop ■ DAQ & Controls  
■ Cluster ■ User Projects  
■ Farm ■ File Servers  
■ WGS ■ Others

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 28



## SL6+ desktop policy



- > Only the “nogroup” setup is supported
  - Login is allowed for **any user**
  - But only on the **local** console – **no ssh** logins
  - No permanent local storage – **only /tmp**
    - > Files in /tmp are retained for two weeks after last access
      - Data for which this is unacceptable should not live on an SATA drive in a PC
      - Reliable storage is provided as AFS, Lustre, dCache & accessible from WGS
        - > even after some colleague unplugged your “PC”
- > This turns a “personal computer” into a piece of office equipment
  - Stateless, with identical setup for all devices
  - Could be regarded as a “thin client”, but is more useful
- > It saves effort, and improves security
- > A growing share of SL5 PCs has been set up like this lately
  - Thanks to all who accepted this setup in the past already!

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 29



## Open question: How to handle minor releases?

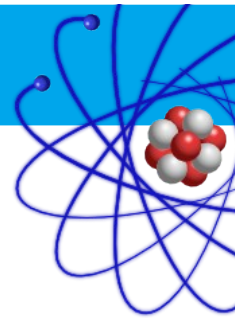


- > Every 6-8 months:
  - Fixes for bugs that aren't critical
  - Limited version upgrades for some software
- > SL allows “sitting on a release”
  - Developers filter out security updates + dependencies
  - Nice feature, but:
    - > These updates are from the latest minor release
      - And that's all Red Hat QA covers
    - > The feature update one wanted to avoid may now roll in as a security one
      - And now it's urgent
- > Policy with SL3/4/5 was: roll out minor updates after a test phase
- > IceCube (Madison) would like us to stay with 6.1
  - “At least for a while”

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 30



## What next?

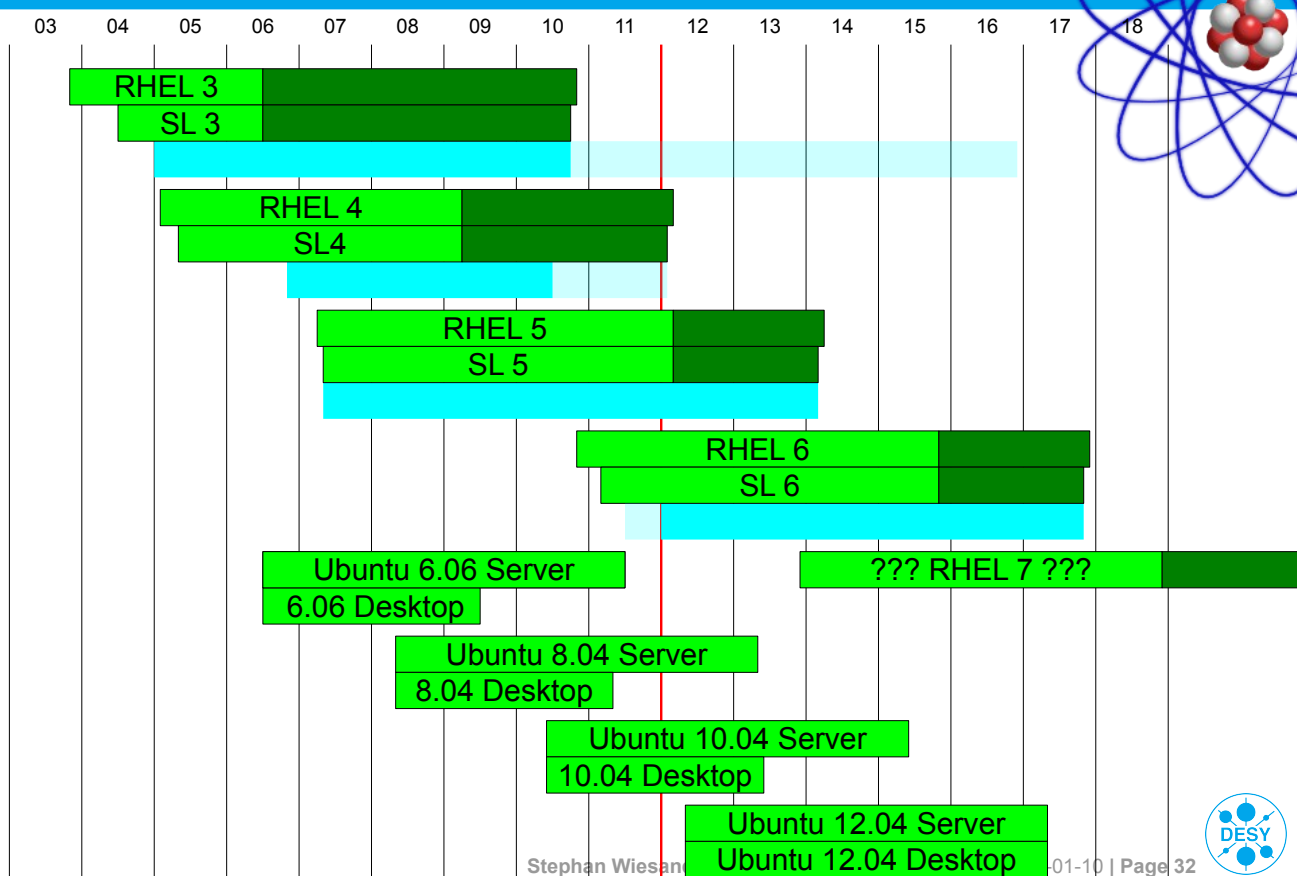
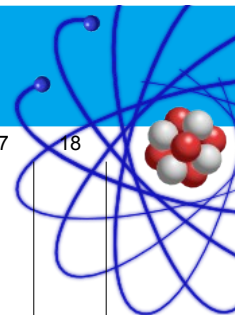


- > SL6 can be used for almost 6 years from now
- > SL7 will not arrive before late 2013 (and rather 2014)
- > Time to discuss our Linux **strategy**
  - In particular, which **distribution(s?)** to use:
    - > **RHEL** & clones
      - Defined, sufficient life time (~ 5 + 2 years after release)
      - Completely **undefined release cycle**
        - > Indicator: Public beta ~ 8 months before
    - > **Fedora, OpenSUSE**: Insufficient **life time** (~ 18 months)
      - That's also true for ordinary **Ubuntu** releases
    - > **SLES**: Probably not bad, but no clones, not cheap, not common in HEP
    - > **debian**: Completely **undefined release cycle and lifetime**
      - EOL: 12 months after next release – whenever that may happen
    - > **Ubuntu LTS**
      - Time based releases, **every 24 months**
      - **5 years of support**, now including the desktop (was: 3 years)

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 31



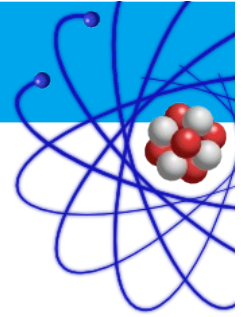
## Timeline - Ubuntu



Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 32



## Why we'd want to use Ubuntu Linux

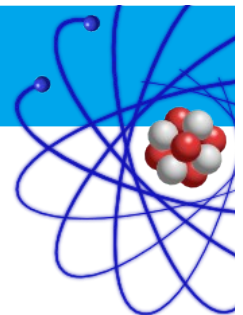


- > Well defined **release cycle & life time**
  - 5 years are probably sufficient if we can plan ahead
- > No need to clone it – the **original** is free
  - Commercial **support** available if needed (+ IP infringement protection)
- > Allows using a **younger** base distribution where required
- > Growing **user share**
- > **Popular** with certain communities
  - M, Photon science
  - Example: Recent request regarding “Tango” control software
    - > Developed and used by ESRF, DESY and others
    - > Packages available for Ubuntu 10.04 LTS
      - but not for SL
- > Next release finally has proper 32-bit support on 64-bit systems

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 33



## Why we'd not want to use Ubuntu

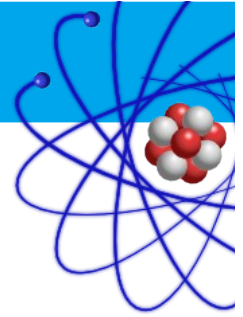


- > Very **different** from RHEL
- > Different **package format**
  - We're using RPM a lot, significant know-how
- > We won't be able to replace SL completely
  - LHC computing, Grid, IceCube
  - CTA controls (ACS)
  - Certified hardware & software? (TSM?)
- > What's the right choice for file servers?
- > Future prospects?
  - Red Hat is profitable
    - > Competition by Oracle may force them to make cloning harder though
  - Canonical probably isn't
    - > yet?

Stephan Wiesand | Scientific Linux 6 in Zeuthen | 2012-01-10 | Page 34



## Summary & Discussion Starter



- > SL5 → SL6 is hopefully **evolution, not revolution**
- > There's ample time for migration
  - There's no chance to completely skip SL6, either
- > There's also time to reevaluate our Linux **strategy**
  - RHEL + clones has worked reasonably well in the past 7 years
    - > SL is still our clone of choice, but alternatives are available, just in case
  - Ubuntu LTS has become an interesting option
    - > Who wants it, needs it, or rejects it?
    - > Replace SL with Ubuntu on desktops?
  - Do Linux **desktops** still make sense?
    - > Leave administration of end devices to users?
  - Is the SL6 desktop model completely unacceptable?
    - > We *have* to cut corners *somewhere*
- > Other **thoughts?**

