

Scientific Linux 5

at



Standort Zeuthen



Outline



- **SL5: available now**
 - what's new
 - what's not
- **migration**
 - how long / to which extent keep support for SL3/4?
 - use occasion for changes to **WGS model**?
- **SL5 will be status quo for ~ 2 years**
 - > 1 year before work on next Linux starts
 - => this is the time for thoughts about the future
 - is our current **Linux desktop support model** still right?



Is SL still the right choice?



- most of **HEP (&GRID)** uses **RHEL** or a rebuild
 - **CentOS**: exact clone, closed development, many users
 - **SL(C)**: more add-ons, open development, fewer users
 - but user/developer community active & strong, low noise mailing lists
- **SLES, RHEL**: significant cost, no SLES rebuilds (possible?)
- **OpenSuSE, Fedora, ...**: insufficient life time (2 years or less)
- **debian**: completely undefined release cycle & support life
 - yet, gaining support by hardware vendors (HP, Dell)
- **Ubuntu**: paid developers make up for debian's shortcomings
 - free, incl. enterprise edition (3/5 years support)
 - paid support available; alas: quite different from RHEL



How to try/test/start using SL5



- Details: http://dvinfo.ifh.de/SL5_User_Information
- public login systems s15.ifh.de and s15-64.ifh.de
- public PCs [smeagol](#) and [deagol](#) in terminal room 2L01
 - using guest account avoids picking up any KDE/GNOME settings
 - switching forth and back with KDE may cause minor problems
- SL5 [batch](#) queues
 - currently 16 of the fastest cores in the farm, 2 GB/core
 - addressed with `-l os=s15` (not used by default yet)
 - `-l os='s13|s15'` should work; will become default eventually (July ?)
 - more systems will be migrated depending on # of jobs



Desktop Upgrades



- good reasons:
 - much **better support for USB storage devices**
 - more **recent versions of applications** (TeX, ...)
 - much better **interactive response during I/O** than SL3
 - more fun to work with, especially on old hardware
 - supporting SL3 on desktops probably hard to impossible soon
 - may even be true for SL4
- **SL5 does not require new hardware**
 - if it runs SL3, it runs SL5
 - but: 6 GB / required (8 GB for 64- bit)
- mail request to uco (see dvinfo for details)



SL 4 and on: Changes



- SL3 was our first Linux ever with many years of support
- => SL4 was the occasion to make a few major changes
- general directions:
 - make systems more independent
 - install software, fonts, ... locally if possible
 - reduce dependency on central services (font, LDAP, NIS, ...)
 - as few modifications as possible to upstream system
 - don't replace profiles, just extend
 - end of HEPiX profiles, HEPiX11 (and fwm2)
 - use unmodified software from distribution wherever possible



What hasn't changed



- **scientific software** equipment
 - cernlib, root, maple, mathematica, ...
 - latest version only yet, older ones on request where possible
- **browsers, mail readers, ...**
 - **firefox** is the recommended browser
 - flash & java plugins, ...
 - **alpine** is the recommended mail reader
 - direct successor of pine
 - much better UTF-8 support, open source license
 - thunderbird, evolution, ... available as is



Legacy Applications



- many still available
- many users still prefer these over "modern" equivalents
- sometimes they're simply better
- more suitable for old PCs (4 years or older, 256 MB or less RAM)
- no intention to provide **xemacs**

The screenshot displays a classic X11 desktop environment with several overlapping windows:

- plan**: A window showing a calendar for June 2007, with the date Saturday, 2.6.07, and time 17:59. The calendar grid shows dates 2, 9, and 16.
- gv: guide.dvi**: A window displaying a document titled "A L^AT_EX survival guide for Unix systems". The window includes a menu bar (File, State, Page, Portrait, 1.000, BBox) and a toolbar with buttons for File Size (112 x 495), Open, Print All, Print Marked, Save All, and Save Marked. A vertical toolbar on the left contains navigation and editing tools.
- mgdiff**: A diff utility window comparing two files: `/tmp/osi_groups.c` (line 623) and `roups.c.afs_pag_destroy-hold-glock` (line 615). The code snippets show differences in the `osi_keyring_init` function.
- xv 3.10a-20050410 <unregistered>**: A window displaying a colorful desktop background featuring a blue sky with white clouds, several red fish, and large, stylized orange letters.



Language Support, I18N



- default is still **LANG=C**
 - UTF8 slows things down, still bugs in some applications
- users can set **LANG=en_US.UTF-8** for themselves
 - in `~/ .i18n` (->shell) and/or `~/ .dircolors` (->>window manager)
- [Win] modifier no longer works for äöüßÄÖÜ
 - instead, R-Alt is "Compose Character" key
 - [R-Alt],["],[a] yields ä, also works for ç ø æ Å ñ ô é è €...
 - "european" languages, UTF-8 independent (ISO-8859-15)
- **SCIM** ("smart common input method") for others, UTF-8 only
 - こんにちは , Дубна Ресторан (getippt: Dubna Restoran)
- only English supported for UI



Desktop Environments



- recommended and default: **GNOME**
- alternative for low memory desktops: **IceWM**
- also available: **KDE, WindowMaker**
- most GNOME/KDE apps work well under IceWM
 - including tray support
- example for combining environments:

run alpine in UTF-8 mode under IceWM w/ LANG=C:

```
[pc] % LANG=en_US.UTF-8 gnome-terminal -e alpine &
```



Compilers



- default: **GCC4**
 - C++ code may need to be adapted
 - FORTRAN frontend is **gfortran**, not **g77**
 - Fortran95 (mostly complete?)
 - may not accept all **g77** code
 - FORTRAN runtime library is **libgfortran**, not **libg2c**
- **GCC 3.4.3 (SL4 default)** is available
 - **g77** command is from this one
- also available: **PGI**, **Intel** compilers
 - **SUN studio 12** released yesterday -> provide?



Backward Compatibility



- RHEL provides runtime compatibility with 1 previous version
- => if it's built for SL4, it should work on SL5
 - that's different from "works on SL4"
- in reality, much software built for SL3 works
 - SL5 ships a C++ compatibility library for g++-3.2.3 (SL3)
 - we add missing shared libraries if required and possible
- some software does not work and never will
 - no more kernel support for **LinuxThreads**: NPTL only
- some software requires some extra attention due to SELinux



SELinux



- additional, fine grained permission system
 - consulted after traditional UNIX permissions grant access
- processes and files (and soon: network packets) carry a "security label" (aka "type", "security context")
 - visible with new -Z option for many commands (ls, ps, id, ...)
- (targeted) policy defines permitted access vectors
 - process of type httpd_t can read files of type etc_httpd_t, but not files of type etc_t
- denials logged with prefix "avc:denied", visible through dmesg
- introduced with SL4, does not affect user applications there



Potential SELinux problems



- **on SL5, the default policy can affect user processes**
 - general aim: make memory either executable or writable
 - protects against buffer overflow attacks
 - other measures in place, some since SL3
 - execshield, position independent executables, stack protector, ...
 - may prevent execution of binaries that are buggy or incorrectly built, but work on SL3/4
 - **several fixes/workarounds available (see dvinfo Wiki page)**
 - special label for executables to allow violations
 - build everything to go into a shared object with -fPIC
 - generally right, and required on 64-bit anyway
- **similar problem: new malloc checks in glibc**
 - workaround: `MALLOC_CHECK_=1` (see Wiki)



New features under the hood



- **XEN Virtualization**, coming in two flavours:
 - **full** virtualization, for unmodified guest systems
 - requires hardware support (latest servers & desktops only)
 - not completely stable yet
 - **paravirtualization**
 - guest system knows it's virtual and plays along
 - good performance, stability seems production grade
 - no hardware support required (host must support PAE, though)
 - SL4 (4.5+) or SL5 guests, Solaris in progress (!)
 - host+guest must have same memory model (32/64-bit)
- **ionice** (like nice, but for I/O, not CPU), **I/O schedulers** per queue and runtime configurable, **CPU sets**, improved **power management**,...



SL5: Summary



- major progress across the board
- good balance of new features and maturity
 - 2nd RHEL distribution with kernel 2.6
 - 2nd RHEL distribution with SELinux
 - 2nd SL with OpenAFS 1.4.x
 - 2nd SL in Zeuthen with major changes in management
- => although we're providing it very early (few weeks after RHEL5 GA), should be more mature than SL4 after a year
- SL4 was a good candidate for skipping
- SL5 is not. Let's use it.

03 04 05 06 07 08 09 10 11 12 13 14

RHEL3

SL3 under discussion

SLC3 64

RHEL4

SL4 under discussion

SLC4

certification
deadline

LHC

RHEL5

SL5 to be discussed...

SLC5 "release in 2007"

Server Hardware Generations

DL3 DL4 DL5 SL3 SL4 SL5 SL6 SL7 ...

full support
maintenance
uncertain

How desirable is this ?!



Boundary Condition: Hardware Support



- RHEL receives **drivers for new hardware for 2.5 years**
 - “at the discretion of Red Hat”
 - **RHEL3 exceeded expectations** in this respect - **on servers**
 - SL3 works well on server hardware generation released last summer, even on models released this year
 - many of their customers wanted to skip EL4
 - on the other hand, no (decent) support for common **desktop** hardware released during RHEL3 “full support phase”
 - X works on our nvidia, boards, but only with the VESA driver
 - BIOS only supports 60Hz => CRT/dual head requires proprietary driver
 - significant extra work with every kernel update
 - Intel HDA onboard audio not supported at all, requires ALSA
 - inordinate amounts of work with every kernel update
- **Sound problems even with SL4 on current desktops**



Conclusion from hardware support



- **no way around SL5 on desktops**, sooner or later
 - rather sooner: new Intel desktop chipsets launched today
 - a few months from now, no more choice
 - very unlikely to work well with SL3, may not work at all
 - unlikely to work well with SL4, even may not work at all
 - few hardware certifications for EL4 (different for SL5!)
- **SL3 desktops already very expensive**, except for old models
 - => we should at least:
 - drop sound support where it does not work by default
 - get rid of SL3 systems with dual head display or CRT
- **SL4 situation is similar** w.r.t. sound
 - 4.5 may help, but for how long?



Boundary Condition: SL



- announced SL3 end of life date is October 2007
- SL3 "legacy support" under discussion
 - final release 3.0.9 corresponding to RHEL3 Update 9
 - currently in beta, no new drivers
 - patches as long as RHEL3
- if this doesn't happen, we can use CentOS instead
- in both cases, some current add-ons will be missing
 - firefox, thunderbird



SL3/4/5 Roadmap: Proposal



- freeze SL3 and SL4 now
 - no new features, applications, software versions
 - may need to drop Software (acroread) or fall back to SL (openoffice) eventually
- prefer SL5 on new systems, especially desktops
 - skip SL4 where possible; ATLAS may need it - really required on desktops?
- aim for upgrading all desktops to SL5 a.s.a.p.
 - drop expensive part of SL3 desktop support with 3.0.9 (October)
 - drop ALSA, nvidia driver
 - drop software if required
- migrate farm, pubs, ... step by step
- keep remaining SL3/4 systems (servers) alive as required
 - up to 2010



The Future of Workgroup Servers?



- pub1-6 are aging, and to be migrated soon
 - simply upgrade/replace them, or rethink the model?
 - don't panic, these are **ideas** - not even proposals
- current **problems**:
 - pubs are often **abused**, may become unusable for others then
 - some **important use cases are not supported well**:
 - building large software projects (parallel make)
 - requires several cores (4 typically optimal), as fast as possible, fast RAM, ...
 - interactive analysis, possibly with threaded application
 - others?
- simply replacing the pubs with leading edge systems will make the abuse problem much worse



Idea for Future Workgroup Servers



- **replace pubs by 2 minimal systems** with a single purpose: login from outside to connect to internal systems
 - little CPU/RAM/disk, limited network speed
 - no access to bulk data, maybe no AFS at all
- **provide all groups with dedicated workgroup servers**
 - several groups already have them anyway
 - coordinated use much easier to achieve on group level
 - small set of people, usually working together anyway
- HW example: 2xWoodcrest 3GHz, 8GB RAM, 146G RAID-0
 - replace systems after 2 years, then use in farm until retired
- **WGS must not be file servers; maintenance slots required**



The Future of Linux Desktops?



- current support model is a straight descendant of the "central server + X-Terminal" one (+local disk/CPU)
 - fully centrally managed, no root access for user
- advantages:
 - **uniformity** across all desktops, WGS, farms
 - good **security**
 - physicists can simply use their PC, **no administration required**
- problems:
 - no **flexibility**, no individual **customization**
 - **results** often **less than optimal** (least common denominator)
 - actual **use (as intended) declining** steeply



Desktop problems continued



- possible reasons for declining use:
 - give a student a 5y old PC with a 10y old monitor, a dirty keyboard, and a mouse without a scroll wheel
 - if this student can afford a notebook, what will happen?
 - what if the PC is up to date, but some wanted feature is not supported? (Firewire, or Midi playback, or ...)
 - what if just some favourite software is not installed? if the user prefers a different keyboard layout? wants a different mouse? a different screen resolution? a different linux distribution? a cyrillic user interface? ...
- people are no longer forced to use our centrally managed PCs, and that's good, and they stop doing it (and that's bad)



Linux Desktops: Questions



- are **centrally maintained desktops still required ?** (at all)
- if yes:
 - **how many?**
 - current estimates: 20-90
 - what's the right **model?**
 - true **thin client ?**
 - only local access
 - no permanent storage
 - **user managed ? with opt-in to limited central management ?**
 - remote access for "owner" only
 - "owner" is responsible, DV only provides working base installation
 - **centrally managed** general purpose PC, like today?
 - **hybrid models are not feasible; multiple models are**



Final Summary



- SL5 should become the platform for user computing asap
- SL3/4 still available (for years)
 - but limit use to cases where it's really required
 - special case: ATLAS depends on CERN schedule
 - others?
- change WGS model ?
 - could be done on the occasion of migrating to SL5
- change desktop model ?
 - would require design, and take much more time
 - if we want this, discussion should start now