DESY Linux 5 at Zeuthen

Stephan Wiesand

DESY Zeuthen – DV -

November 4th, 2003

Agenda

- What's DL5 anyway?
  - base distribution
  - DESY/HEP add-ons
  - automatic installation and maintenance
- Changes w.r.t. DL4
  - what's new for users (and admins)
- Status and Roadmap
- Questions & Discussions (anytime)

DL5 Base Distribution

- SuSE 8.2 Professional
  - frequent releases
    - could pick an up to date release
    - 9.0 is too young
      - SuSE releases mature after a few months
      - we need a few months for adaptation as well
  - supported by SuSE for 2 years
    - compatible bug/security fixes
  - negligible monetary cost
    - use on private/external computers feasible

Competing Distributions

- SuSE/RedHat Enterprise Distributions
  - could not yet negotiate reasonable conditions
- RedHat Linux
  - had 10-12 months of support, now discontinued
- Fedora (RedHat Community Linux)
  - 6-9 months of support
- debian
  - no defined release cycle ("when it is time")
  - "stable" out of date when released
  - not for notebooks, little to no ISV support

Add-ons: Kernel

- having own kernel proven useful when ptrace bug hit
  - we had a secure kernel out much faster than SuSE
- stock kernel from kernel.org w/ few modifications
  - max symlink depth (stock: 5, SuSE: 8, DESY: 15)
  - preferrably few other patches
    - possibly additional drivers
    - in future, possibly security enhancements
      - exec shield, st. jude, ...
- smart packaging for non-volatile updates
  - old version not removed before new one works
  - includes kernel specific part of AFS, ...

Add-ons: AFS

- OpenAFS from openafs.org
  - built together with kernel
  - some components go into kernel package
    - kernel module
    - cache manager
    - kernel specific debugging tools (kdump-`uname -r`)
  - oldfashioned transarc path layout
    - used in too many scripts...
    - but some compatibility links for new paths
  - many additional executables and other files
    - mainly debugging tools, but also additional file server

Add-ons: HEPix environment

- set of login scripts
  - login shells (zsh, tcsh)
  - X session
- in use at most HEP sites
- consistent customization scheme
  - per user, per group
  - compliant customizations forward compatible
- stable window manager configuration (fvwm2)
  - identical on all platforms
  - may be scrubbed eventually

Add-ons: Security/Privacy

- security update packages
- many configuration changes to help prevent
  - remote root exploits
  - local root exploits (from cracked/sniffed user account)
- device permissions
  - should be fully accessible locally, and only locally:
    - floppy, cd-rom, cd writer, usb memory stick
    - audio devices
    - video devices

Add-ons: Convenience

- automounter
  - consistent with other platforms
  - SuSE comes with autofs only (linux specific)
- integrated login
  - kerberos
  - AFS
- customized ssh
  - kerberos ticket & AFS token forwarding
  - secure "single sign on" for all Unix hosts

Add-ons: Application Software

- compilers and libraries
  - several gcc, intel, portland group, ...
- scientific software
  - ROOT, cernlib, maple, mathematica, ...
- customized and/or updated versions of
  - browsers (cache in /tmp, certificates, ...)
  - mail readers (kerberized pine)
  - document viewers (security, features)
  - TeX (logos, additional styles)

Availability

- most of all this is available by NFS
  - /net/z/SuSE-8.2/ = z.ifh.de:/net1/z/SuSE-8.2/
  - iso/ CD iso images
  - orig/ CD content
  - update/ update packages, mirrored regularly
  - install/ updated CD content + addons
  - nbinstall/ like install, for notebooks
    - kernel & AFS from SuSE not replaced
- exported to registered hosts only

Automatic Installation

- now based on AutoYaST2
  - completely different from YaST1
  - from updated repository (/net/z/SuSE-8.2/install)
- partitioning, selection etc. specified in ".ch5" file
- CH5.pl translates .ch5 to .xml for YaST2
  - and sets symlink to host IP in hex notation...
  - find script in /project/linux/DL5/CH5/
  - no input from VAMOS DB (host database) yet
    - may be added later
- YaST installs host into well-defined, simple state

Automatic Maintenance

- after YaST is done, sue/cfengine kicks in
- desired state defined in VAMOS
- agent (sue.update) runs cfengine "features"
  - triggered by
    - postinstall script executed by YaST
    - cron
    - reboot
- additional mechanisms by init scripts
  - during boot & shutdown
  - kernel & glibc updates, ...

Maintenance: Remarks

- mechanism is complex and fragile (by law of nature)
- it's also essential
  - for security & usability in our environment
    - of hundreds of hosts, within manpower constraints
- consequences:
  - desktops, farm hosts handled like commodity items
  - even true for many servers, workgroup servers
  - handing out root passwords to users is not feasible
    - too easy to break something
      - unless configuration and mechanisms fully understood
  - nonstandard configurations are very expensive

Remark: DL5 on Notebooks ?

- DL = base distro + add-ons + maintenance mechanism
  - add-ons and maintenance mechanism still depend on permanent, unlimited network connection to servers
    - AFS
    - kerberos
    - VAMOS
- there is no DL for notebooks
  - and won't be for quite a while
- but possible to install a DL5-compatible SuSE-8.2
  - most application software now rpm-packaged
  - allows developing & running DL5-compatible apps

DL5 Changes: AFS Sysname

- now i586_linux24
- DL4: i386_linux24
- watch out for `~/bin -> ~/.@sys/bin`
- right commands to find out the current sysname:
  - fs sysname
  - livesys (OpenAFS >= 1.2.2)
- wrong command:
  - sys (gives compile time, not run time sysname)

Changes: KDE

- KDE is now the default, recommended session type
  - many enhancements & new features
  - fvwm2 may eventually vanish
- version 3.1.1, many improvements over version 2
- switching back and forth may not work perfectly
- before running it the 1st time, keep a copy of
  - ~/.kde and ~/.kde2
  - ~/.kderc and ~/Desktop
- terminal app: xterm with Vt fonts unbeaten quality
  - prefer over konsole

Changes: Application Software

- provided in a completely different way
- location: `/opt/products`
- users should forget about /products
  - admins, generally, should forget it as well
    - and some of their former habits
- few compatibility links maintained in /usr/local
  - kept to minimum
  - additional ones can be added, if justified
- everything in /opt/products is from an rpm package
  - nothing is maintained manually there

Reminder: Application Software

- new scheme kicked off in spring, objectives:
  - well-defined, reproducible state on all clients
    - no need for client backup, easy rollbacks
  - suitable for
    - notebooks, critical systems (local software installation)
    - systems with small disks (remote installation, symlinks)
    - local installation of some (vital) software on all systems
  - flexible, classable, hierachic, delegatable configuration
    - QA environments, group specifics, (workgroup) servers
  - self documenting builds (preserve lessons learned)
  - easy user access to information
  - dependency management
  - long term man power savings

Application Software ctd.

- solution chosen: RPM packages
  - separate DB: `rpm --dbpath /opt/products/RPMDB`
  - packages generally split in two parts
    - base package (gcc-3.3.1-4)
      - lives in /opt/products/gcc/3.3.1
      - which may be a link to /afs/<cell>/@sys/products/gcc/3.3.1
        - link to "reference installation", from same packages
      - several versions can be installed in parallel, each locally or not
    - default links package (gcc-default-3.3.1-4)
      - contains /opt/products/bin/gcc -> /opt/products/gcc/3.3.1/...
      - /opt/products/bin in user PATH
      - only one version can be the default (have its -default installed)
      - is always installed locally (but cheap in terms of storage)

Application Software: Tools

- installation/query tool: ppm
  - queries: `ppm -q [/perl regular expression/[ix]]`
    - `ppm -q`
      - gives full list of what is installed which way

```
type    name        version         release    links
=============================================
link    acroread    5.0.8           1          default
link    cernlib     2002            1
link    cernlib     2003            0.030830   default
link    db          3.3.11          1
local   db          4.1.25          3          default
...
```

    - `ppm -q /xml/i`
      - gives list of all packges with name containing xml or XML or...
  - installation mode accepts input similar to -q output

Detailed queries

- shortcut for rpm --dbpath /opt/products/RPMDB:
  - `prpm -qi gcc`
    - information about all installed gcc versions
  - `prpm -q --changelog gcc-3.3.1-4`
  - `prpm -ql gcc-3.3.1`
    - lists all files from package
  - `prpm -qd gcc-3.3.1`
    - lists all documentation files
  - `prpm -yv gcc`
    - verifies that package installation is intact, unchanged
    - user, group, timestamp may differ for link installs
      - but MD5 sum will not
- eventually: graphical/web interface

Software Repository

- /afs/ifh.de/packages/
  - SOURCES/
    - tarballs
  - SRPMS/
    - (no)src-rpms (specs, additional sources)
  - RPMS/i586_linux24/
  - RPMS/noarch/
  - RPMDB/
    - full DB of all available packages
    - RPMDB/locatedb is a locatedb of all available packages
      - including noarch

Installing on any SuSE-8.2 host

- start AFS client, make sure sysname is i586_linux24
- `mkdir -p /opt/products/RPMDB /opt/products/perl`
- ```
  rpm --dbpath /opt/products/RPMDB -ivh --nodeps \
  /afs/ifh.de/packages/RPMS/@sys/perl/perl-5.8.0-19.i586.rpm \
  /afs/ifh.de/packages/RPMS/noarch/ppm/ppm-0.9-14.i586.rpm
  ```
- install -default packages accordingly, set PATH
- create ppm input file my.cf
- ppm -v my.cf
  - dry run
- ppm -vx my.cf
  - once satisfied with result my.cf
  ```

Remark: Copyrights

- most packages are licensed under the GNU GPL
  - or free in some other sense
- some are absolutely not
  - installing these on non-DESY systems is illegal, even if it's possible
    - the same for copying installed files!
- copyrights are clearly declared in package description and copyright tag (output of prpm -qi)
- we try to get rid of license manager dependencies
  - only possible if not abused

Changes: GCC Compiler

- default on DL5 is gcc 3.3.1
  - default on DL4 was gcc 2.95.3
  - C++ ABIs are incompatible
    - all C++ code in program must be compiled with same compiler
  - 2.95.3 is also provided
    - and a root version built with it
  - SuSE 8.2 comes with 3.3 prerelease
    - updated runtime libs are needed for running C++ software compiled with our DL5 compilers on plain SuSE 8.2
    - /afs/ifh.de/packages/RPMS/i586_linux24/System/
  - 2.91.66 and earlier no longer work

GCC Compiler continued

- gcc 3.3 no longer accepts K&R C
- varargs.h is gone
  - has been deprecated for a while
    - but still code around using it
  - code must be changed to use stdarg.h instead
- errno.h must be included explicitly
  - symbols no longer available without
- C++ code should run 5-7% faster

Changes: HEP Software

- cernlib
  - only versions 2002 and 2003 are provided yet
  - older versions hard to build on current systems
- root
  - only version 3.05.07 is provided yet
  - for both gcc 3.3.1 and 2.95.3

Remark: nonstandard libs

- if shared libraries are used in nonstandard versions
  - it is not sufficient to link against them
  - make sure you are using the right headers
  - make sure you pick up the right runtime libraries

```
ROOTVERS:=3.05.08

ROOTCFG:=/opt/products/root/$(ROOTVERS)/bin/root-config
ROOTINCDIR:=$(shell $(ROOTCFG) --incdir)
ROOTLIBDIR:=$(shell $(ROOTCFG) --libdir)

my_program:
    $(CXX) -o $@ $(OBJECTS) -L$(ROOTLIBDIR) \
    -Wl,-rpath,$(ROOTLIBDIR) -Wl,-rpath-link,$(ROOTLIBDIR) $(LIBS)

%.cc: %.o
    $(CXX) -c -I$(ROOTINCDIR) -o $@ $<
```

Changes: Browsers

- customized version of mozilla 1.5 now available
  - disk cache in /tmp (preserve AFS space, backup)
  - mail, certificates preconfigured (strong ciphers only)
  - only effective if no ~/.mozilla or ~/.netscape
- netscape4 is available
  - legacy version 4.80
  - deprecated
- netscape7 is not available
  - discontinued by AOL
  - use mozilla instead

Changes: TeX

- resources for maintaining TeXLive installation in /products no longer available
- using SuSE's TeTeX for the time being
- DESY add-ons provided by rpm package
  - styles (foiltex, ...)
  - logos, DESY letter, ...

DL5 Status: Upcoming

- not quite finished
  - some software still missing
  - some desktop specific features not finalized
    - sound
      - ALSA or default kernel drivers ?
    - CD writing
      - use native mode for IDE recorders ?
      - handle all CD/DVD drives by ide-scsi ?
    - USB/hotplug
- current status documented on www.ifh.de/linux5/
  - kept up to date - check regularly

DL5 Roadmap

- public preview available: dl5.ifh.de
  - please check: what's missing for you ?
- running on 2 other servers, 5 desktops (CC)
- no major problems, very stable
- next week:
  - available to early adopters (volunteers/guinea pigs)
  - farm host, dedicated queue
- end of November: begin rollout
  - new installs (also after hardware repairs) will be DL5

Upgrades

- can be prepared remotely
  - no physical access to desktops needed
- after negotiating with user/group admin:
  - upgrade prepared by DV (new default boot entry)
    - may be scheduled in advance
  - triggered by reboot (by user or DV)
  - maintenance mechanisms must be disabled
    - reboot same day by user, or next morning by DV
- problem: user/group of many systems unknown
  - users take "commodity item" too literally

Upgrades: Preserving /usr1

- /usr1 can usually be preserved
  - no guarantees
    - OS installation is a major procedure
  - but haven't lost any data during upgrades to DL4
- but only on hosts with root filesystem >= 3 GB
  - otherwise
    - disk must be repartitioned, or
    - only a subset of applications can be installed
      - no TeX
      - no OpenOffice (yet), ...
- /usr1/data will be renamed to /usr1/scratch

DL4 End of Life

- will be established once DL5 fully available
- as soon as possible
- maybe as soon as Christmas
  - 7.2 no longer supported by SuSE
  - keeping DL4 systems secure increasing burden
  - would rather spend time on DL5 improvement
- even after EOL, users not forced to upgrade
  - for about 3 months
  - but insecure functionality may be disabled
- DL5 will have absolute priority once rollout started

Summary

- DL5 is going to happen
  - RSN
- please, check out the preview
  - what's missing ?
  - what needs to be changed ?
  - what could be improved ?
  - anything that could/should be dropped ?
    - are we wasting our time on something you don't need ?