

33rd Linux User Meeting

2006-09-20

SemR 2a/2

A large, faint watermark of the DESY logo is visible in the background, featuring the word "DESY" in a large, light blue sans-serif font and a network diagram above it.

DESY

Agenda



- Administrivia
- SLD3 Migration Status
- SLD4 Status
- Notebook Support
- Server and Storage Hardware
- Asset and Lifecycle Management
- User Space - Discussion of Plans and Requirements

DESY

A Word in Advance



- I am speaking on behalf of a growing and very busy Linux team.
- We have racked and installed about 100 new servers since the last meeting, often with completely new hardware
- Work presented here has been done by colleagues from several IT subgroups at Zeuthen and Hamburg.
- I hope to do them justice with what I wrote up here.
- Some things written here are ideas, not services.
 - I encourage discussion and feedback
- Any mistakes are mine

DESY

Knut Woller
IT Systems and Operations

- Next meeting proposed for December 6, 2006
 - Agenda items: SLD4
 - DESY Software
 - AFS?
 - Oracle?

- Linux on the move
 - Linux team members are moving from 2b/2xx to 2b/0xx
 - Starting early October
 - Consult phonebook before visiting us in the South Base of 2b

SLD3 Migration

Current Statistics

Migration Campaign

AFS System Tzpe Names

SLD3 News and Facts

Linux Wiki

A large, faint, light blue watermark of the DESY logo is visible in the background, consisting of a stylized particle detector structure and the word "DESY" in a large, bold, sans-serif font.

Migration Statistics (Site Hamburg)



OS	Jul → Sep	
DL3	69 → 67	-2
DL4	324 → 308	-16
DL5	914 → 872	-42
SLD3	234 → 340	+106
SLD4	51 → 70	+19

**We're moving in the right direction.
Thank you for your support!
We'll continue and try to increase pace.**

Announcements:

PAL Migration, Pool PCs



- pal cluster nodes will migrate to SL3 by mid October
 - after public announcement
 - requires migration of k5cron jobs
 - may require changes to scripts (path or command names)

 - Linux pool PCs in bldgs. 1 and 2 will be replaced
 - by Igel Thin Clients with secured X11 tunnels
 - after PAL migration
 - Use of local media (CD, Floppy) will no longer be possible
 - Is that needed in pool PCs?
-

AFS System Type Name



- The AFS system type name serves two purposes
 - (1) to designate the combination of hardware architecture and operating system on which a particular set of AFS binaries are supported
 - (2) as values for the magic '@sys' variable, which is used to select files based on the client system architecture.
- The `sys` command returns the first. It is defined at compile time
- The `fs sys` command returns the second. It can be set and modified at any time.
- The two are often confused for each other
- Sites which don't care will have them identically and set to defaults
 - Which, in case of Linux, reflect kernel instead of library versions and does not suit software distribution needs
- Canonical sysnames are listed on <http://grand.central.org/numbers/index.html>
- New AFS releases support a list of compatible system names for (2)

AFS System Type Defaults at DESY



- DL1 i386_linux20
- DL2 i386_linux22
- DL3 i386_suse63
- DL4 i386_linux24
- DL5 i586_linux24
- SLD3 (1) i386_linux24 (2) i586_rhel30, i386_linux24
- SLD4 (1) i386_linux26 (2) i586_rhel40, i386_linux26

You need to know these when you use @sys for software distribution or compatibility links in AFS

Local Administrator Feature



- Enable users to perform root privileged commands
 - by allowing sudo for the owner(s) of a machine
- SALAD flag 'ladm' creates sudo right for registered owner

```
$ sudo -l
```

User woller may run the following commands on this host:

```
(root) NOPASSWD: /usr/local/etc/post_install,  
/usr/local/bin/salad, /usr/local/salad/ALL/etc/jammi,  
/bin/kill, /bin/mount, /bin/umount, /sbin/fdisk,  
/usr/bin/yumex, /usr/bin/yum, /usr/bin/system-config-  
display, /usr/bin/system-config-soundcard, /usr/bin/apt-  
get, /usr/bin/apt-cache
```

- To exploit that, make sure that end users are registered
 - It's in the user/responsible field of the wboom or AMS records

wboom record



```
$ wboom -view linuxtest3
```

Variable	List
hostname	linuxtest3
alias	
ip	131.169.137.108
subnet	137
mac	001320B8D5B5
serial_number	
desy_number	
equipment_class	pc
boot_version	Linux
hardware_type	i686
standard_memory	1GB
special_memory	
netgroups	zdv-desktop-hosts 0001
user-responsible	flemming
phone	4667
location	2b/213
group-owner	IT
comment	
config	zdv
subconfig	afs auto ppm desktop pub p da-d ladm
bootp_configtag	
status	ok-flemming-06/09/19

Release Cycles



- Professional Linux Releases (DL): 2 years
 - Create repository on release
 - Maintain this repository with irregular updates
 - Freeze at end of life
- Enterprise Linux Releases (SL): up to 7 years
 - Create initial repository on release, maintain fixes
 - Create new repository for quarterly updates, maintain fixes
 - 28 repositories per major release after 7 years? No way!

Enterprise Linux requires different methods of repo maintenance and client synchronization, owing to its lifetime

SL Jumbo Patch



- The SL Jumbo Patch updates a previously installed system to the latest repository
 - `cd /usr/local/salad/system`
 - `/usr/local/bin/salad -p Linux.JUMBO`
- Implemented as yum wrapper due to yum's incorrect handling of package exclusion lists
 - Preview of effects possible by running 'jammi' script directly
- Implemented for SLD3 and SLD4
- Executed by root user or (on request) in nightly update
- Our goal is to have no more than two repos in the field
 - This would require quarterly updates (maintenance downtimes)

- A Wiki service for the Linux communities at DESY
 - has been suggested and discussed long ago
- Web Office Wiki services are under development
 - but will take more time before general availability
- We would like to start a linux developer platform before
 - based on a simple MoinMoin Wiki engine
 - with low service level (single server, may have downtimes)
- Need a kick-off with community members in October
 - to collect expectations and prepare a basic layout
- Who's interested?

SLD4 Status

Operating System

DESY Software Packages

Open Issues

Public Preview

A large, faint watermark of the DESY logo is visible in the background. It consists of a network diagram with four nodes and connecting lines, and the word "DESY" in a large, light blue font below it.

SLD4 OS Status



- SLD4 32bit base installation built in wboom/SALAD
 - current maintained repositories are 4.3 / 4.4
 - AFS and Kerberos integration is ready
 - OS level is production grade
 - Is already used in user-less production servers
- Very similar to SLD3
 - in terms of package selection and system configuration
- Maintenance and patch mechanisms have been set up
 - including Jumbo patch for version updates
 - most of this work is done for SLD3/SLD4 in parallel
 - SLD4 directly benefits from work done for SLD3

SLD4 Software Status



- 600 additional software packages have been provided
 - significant effort in Zeuthen and Hamburg
 - almost the same amount as for SLD3

SLD4 (32-bit) DESY Addons

- * Indexed Packages: 597
- * Available Architectures: i386, i586

Latest packages:

- * 08-Sep-2006: f95-default-5.1-1
- * 08-Sep-2006: f95-5.1-1
- * 01-Sep-2006: oracle-client_10.g-config-2.0-1
- * 28-Aug-2006: gsl-default-1.8-2
- * 28-Aug-2006: gsl-1.8-2
- * 25-Aug-2006: glue-3.SL-75
- * 09-Aug-2006: pgi-default-6.1-6
- * 09-Aug-2006: pgi-alias-6.1-6
- * 09-Aug-2006: pgi-6.1-6
- * 08-Aug-2006: valgrind-default-3.2.0-1
- * 08-Aug-2006: valgrind-3.2.0-1
- * 07-Aug-2006: rfio-default-1.0-4
- * 07-Aug-2006: rfio-1.0-4

Open Issues



- We now need application and stability testing
 - before we can roll out on large scales (infrastructure protection)
 - test interoperability with network services (NIS, KRB, ...)
- Probably need fine tuning of services in real life
 - like mail configurations
- Need a working solution for application printing
 - Apps assume CUPS, which does not scale well
 - Active work in progress in Hamburg and Zeuthen

DESY IT

Public Preview



- We are putting the final pieces together now
 - installation tests starting on IT computers
- We will provide public preview systems in early October
 - will be announced in IT News and to DESY Linux Admins
- Will be able to provide test systems soon afterwards
- We're currently on our proposed release schedule
 - I still hope to have usable systems in January
- We encourage use of and feedback on test systems

DESY

SL4 Notebooks

Installation

DESY Add-Ons

User Responsibility

A large, faint watermark of the DESY IT logo is visible in the background, consisting of a network diagram and the text "DESY IT".

DESY IT

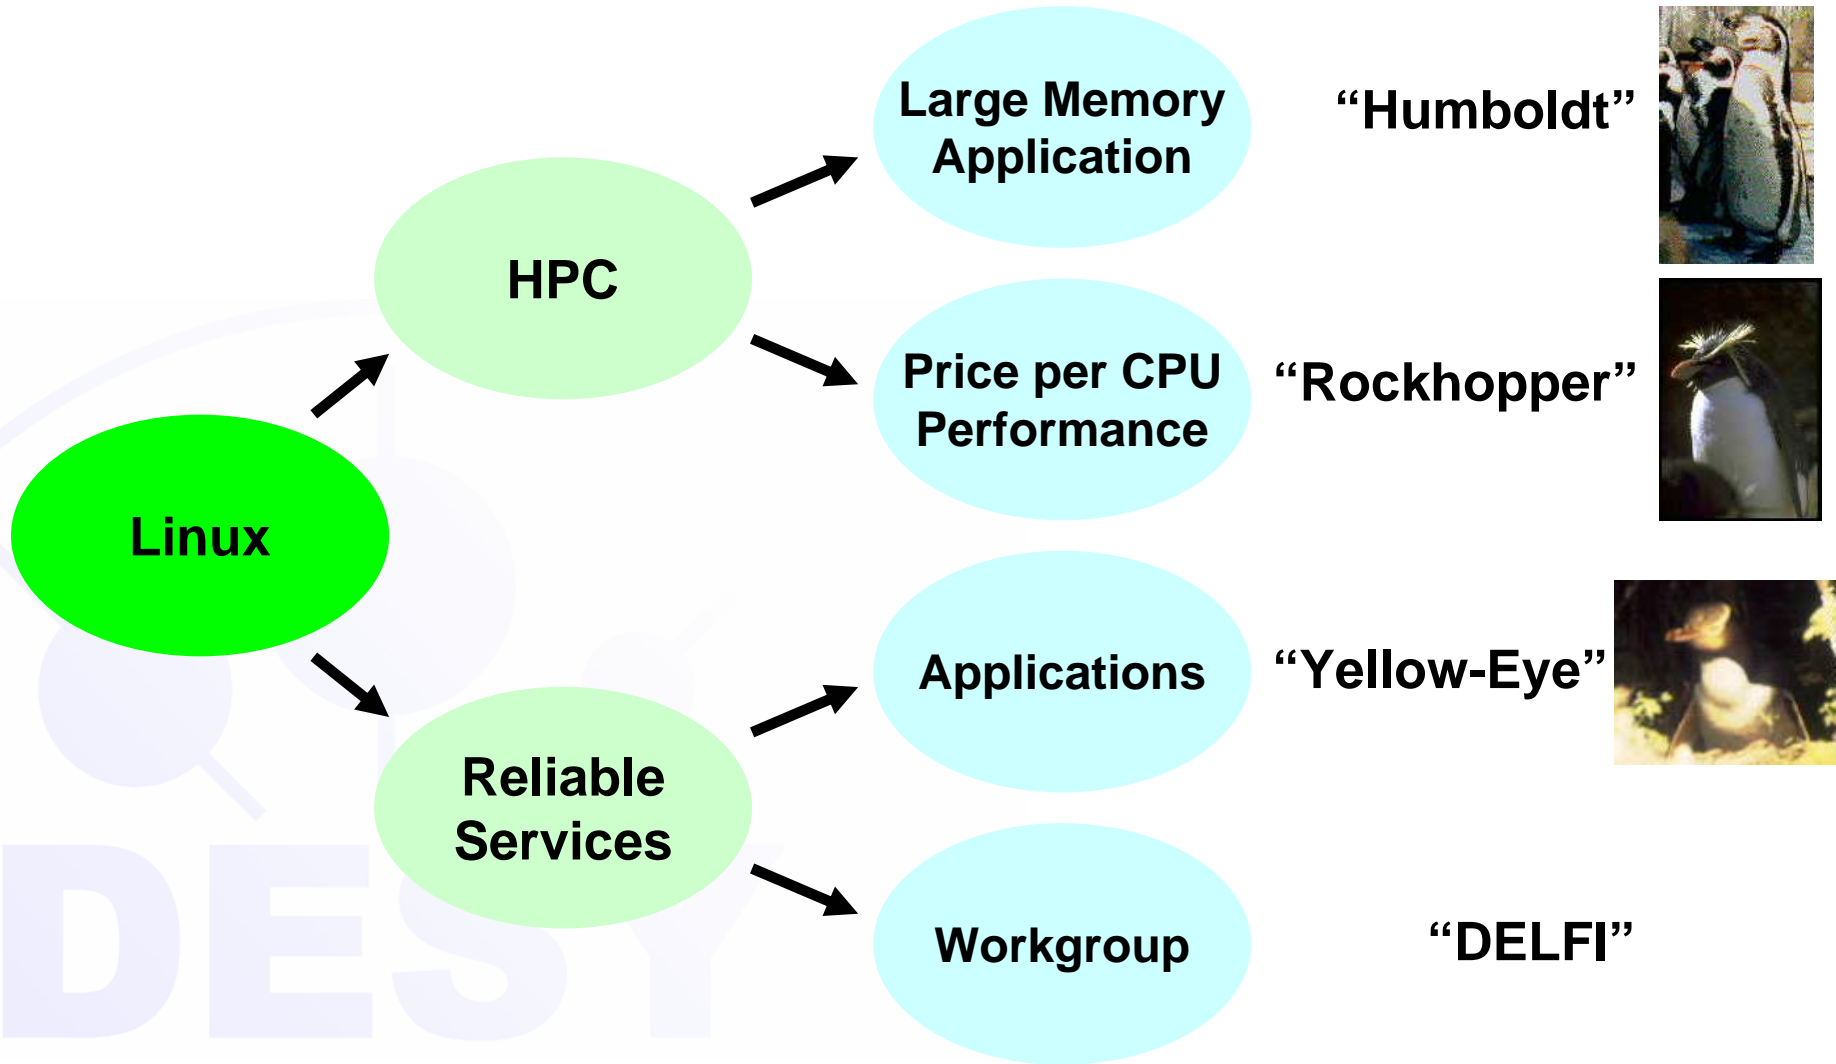
- When you order a new notebook via AMS webshop, you can request a system installation along with it
 - Fill in the additional fields and and your OS of choice
 - Both Windows and Linux (SL4) are available
- SL4 standalone installation comes with
 - DESY network settings
 - Preconfigured user login (AFS user id, dummy password)
 - AFS client
 - Preconfigured yum updates
- Note that this is not a 'managed' installation
 - wboom/SALAD or Quattor are not updating it
- This installation is taylored to DESY Standard Notebooks only

Server and Storage Hardware

Server Classes and Models

Intel vs. AMD Server CPUs

Linux Server Classes



Rockhopper Class



- High compute power, low redundancy
- Two CPU chips, 1-2 cores each
 - Dual Opteron or dual Xeon
- Moderate amount of memory
 - 1-2 GB per core
- Single SATA system disk
 - optional single data disk
- Gigabit network
- Non-redundant power
- SL3/SL4 with 32 or 64 bit
- e.g. SM6015V-T, Tyan GX28, Tyan GT24, DELL PE1950, ...



e.g. SM 6015V-T

1u enclosure
max 2 disks
max 2CPU/4cores

Yellow Eye Class



- **Moderate compute power, max. redundancy**
- 1-2 CPU chips, 1-2 cores each
 - Xeon or Opteron
- Moderate amount of memory
 - 1-2 GB
- Mirrored SCSI system disk
 - optional mirrored data disk
- Gigabit network
- Redundant 2-phase power
- SL3/SL4 with 32 or 64 bit
- e.g. SM 6015P-8R, Sun Fire X4100, HP DL360, DELL PE 1950, ...



e.g. SM 6015P-8R

1u enclosure
max 4 disks
max 2CPU/4cores

Humboldt Class



- **Four or more CPU chips**
 - Currently 4-8 chips with 1-2 cores each
- Very large memory
 - typically 32 GB upward
- Redundant local disk space
 - system mirror, data mirror or data RAID
 - typically on 3ware controller
- Gigabit network
- Redundant Power
- SL3/SL4 64 bit



e.g. Tyan VX50

5u enclosure
max 10 disks
max 8CPU/16cores
max 64 GB

DELFI Class



- Low cost, high redundancy
- Two CPU chips, 1-2 cores each
- Moderate memory
 - 1-2 GB per core
- Redundant local disk space
 - system mirror, data RAID
 - built around 3ware controller
- Gigabit network
- 2-phase redundant power
- DL5; SL3/SL4 32 or 64 bit
- DESY IT design and specs
- Designs ranging from 1u/4disks to 4u/16disks



e.g. DELFI-1-12-DO

3u enclosure
max 12 disks
max 2CPU/4cores

The StorBox



- low cost, performant dCache Server
- 2 CPU chips, 1-2 cores each
- Moderate memory
 - 2-4 GB
- Redundant local disk space
 - system mirror, data RAID
 - built around Areca controller
- (multiple) Gigabit network
- 2-phase redundant power
- SL4
- DESY IT design and specs
- Not for general file server usage due to lack of standard drivers



StorBox 4

4u enclosure
2+16 SATA disks
5+ TB

About the DESY Server Designs



- Both DELFI and StorBox ...
 - ... have been and still are very economic
 - ... match operating and installation systems well
 - ... fit perfectly into our computer center environment
 - ... allow early adoption of next generation hardware
 - ... require careful selection of components
 - ... require careful matching of firmware and kernel versions
- I think their bell will toll soon because
 - ... increasing disk capacities reduce need for multi-controller
 - ... onboard SATA/SAS RAIDs start working with Linux (ZCR, CERC)
 - ... industry starts delivering competitive servers or disk boxes
 - like SUN E4500, DELL MD1000, ...
- We are actively evaluating alternatives


- Tyan recently had large delays in deliveries
 - combination of RoHS laws and company relocation
 - now recovering slowly, but still lacking motherboards and IPMI's
- DELL offers impressive range of new servers
 - Xeon Dual Core available now (PE1950, PE2950, PE6850)
 - Opteron 2xxx/8xxx by end of year
 - With very aggressive pricing
- SUN X4500 targets low cost storage
 - 48 SATA disks in 4u, 20TB net storage
 - Prices in DELFI/StorBox range



Opteron vs Xeon



- Xeon “Woodcrest” dual core systems available now
- Impressive benchmarks!
 - high memory bandwidth
 - best price/performance
 - low power consumption
 - beat all Opterons (incl. Socket F)
- run with SL 3.0.8 / SL 4.4
- have more potential yet
- Most vendors have both
 - Supermicro, Tyan, DELL, ...
 - except SUN: AMD only



Asset and Lifecycle Management

Operations Workflow

A large, faint watermark of the DESY logo is visible in the background, featuring a network diagram and the word "DESY" in a large, light blue font.

DESY

IT Asset Lifecycle



1. Decision (Group)
2. Design (Group, IT)
3. Acquisition (Group, IT)
4. Configuration (Group, IT)
5. Delivery (Vendor to IT)
6. Installation (IT)
7. Maintenance (IT, Group)
8. Phase Out (IT, Group)

**The Asset Management System (AMS)
and IT Request Tracker (RT) support this
by providing information and workflows**

Computer Center Servers



- We are monitoring orders of IT equipment
- Ordering a server for the CC triggers a workflow
 - Collection of information about the system (Host name, IP)
 - Pre-configuration of DNS and installation mechanisms
 - Allocating rack space
 - Preparation of cables and connections
- This information can be available weeks before delivery
- If so, the server can be operational within a week or less

**If you do not order through DESY,
please contact rz@desy.de in advance!**

Increasing AMS Usage



- AMS has good reporting capabilities
 - Provided it contains valid data
- We are working on the AMS data quality
- We intend to use it (and its delegation model) as front end to the installation mechanisms
 - There is still significant work needed
- We intend to implement automated notifications
 - E.g. end of warranty or support contracts

DESY IT

User Space

Your room for requests and discussions

A large, faint watermark of the DESY IT logo is visible in the background, consisting of a network diagram and the text "DESY IT".

DESY IT
