DESY site report

Peter van der Reest et al.
HEPiX site report
Oxford, 2015-03-23

HELMHOLTZ ASSOCIATION

DESY

Zimbra Migration

- Migration is running smoothly
- All services are running virtualized on VMware
- 60% of mailboxes now migrated
  - migration is planned to be finished by early May

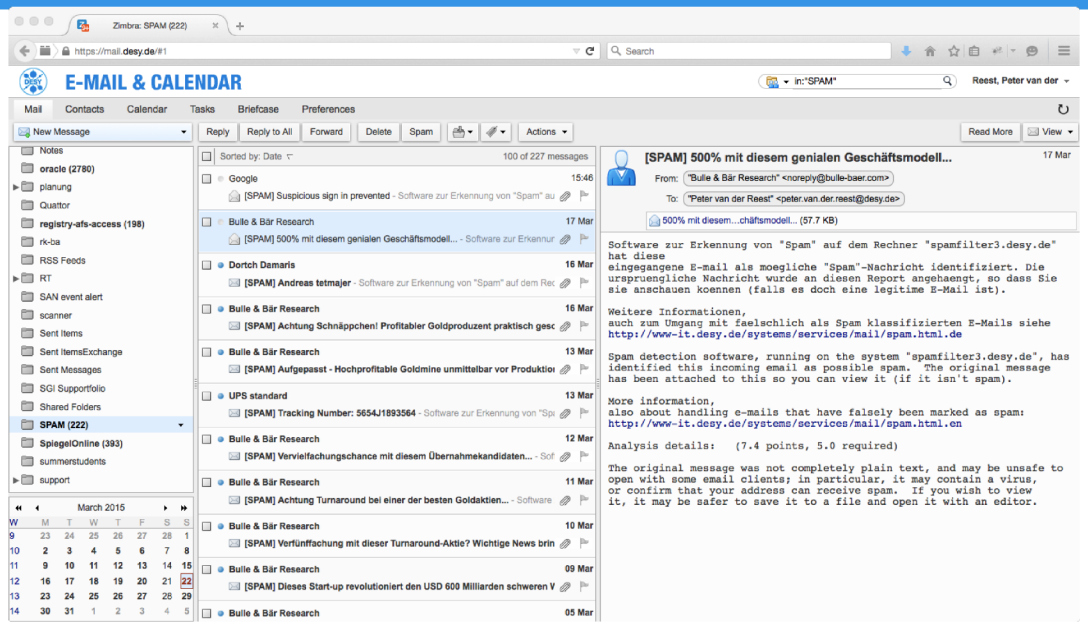

- Response from users is mostly positive
  - like new features
  - like higher performance
  - some are switching from Outlook, Thunderbird etc to webGUI

- one downside for Android users of ActiveSync interface is the high requirements security-wise: “server may wipe device”
  - alternatively use IMAPS, CalDAV, CardDAV

Batch Systems – Zeuthen site

- deploying Univa and started to use UGE 8.2
  - main motivation is cgroup support
  - read only qmaster threads help to let qstat not disturb scheduling
  - share tree calculation bug fixed

Batch Systems – Zeuthen site (cont'd)

- cgroup support
  - gives us core binding and limitation of resident RAM
  - problem when using resident memory limit (m_mem_free): job does not get killed but memory gets paged out

    When using h_rss instead then m_mem_free gets set higher than h_rss and the job is killed when over h_rss

  - bug in RHEL6 Kernel 2.6.32-504 lets nodes crash sometimes

    happens when removing the job-cgroup (race condition?) (also reported by other site)
    https://bugzilla.redhat.com/show_bug.cgi?id=1197612

    therefore production still relying on UGE 8.0

Batch Systems – Hamburg site

- Grid
  - currently running Torque, with a home-build scheduler
  - scheduler now able to handle multi-core jobs, and these are well used by experiments
- local batch farm
  - running SoGE, seeing performance and stability issues
  - setup a testbed to investigate HTCondor and SLURM in our environment
  - Interested in knowing how other people are running HTCondor and/or SLURM with $HOME=/afs/…
- HPC cluster
  - calender based reservation tool - working well, but will not scale for future
  - SLURM might be of interest, HTCondor most likely won't

DESY

Linux status at DESY

DESY central IT groups will provide a RHEL 7 derivative:

- we will call it EL7
- we will offer EL7, whether SL or CentOS still under discussion ...
  - but in the end: makes no difference for users
  - and little difference for admins
- observing the 7.0->7.1 transition ... SL or CentOS decision soon
- we will not offer EL7 on desktops, only on (work group) servers
- will try to minimize dependencies to either SL or CentOS to be flexible (possibly offer a choice of both)

- Ubuntu Server: supported at Hamburg site only

Linux status at DESY (cont'd)

- Ubuntu Desktop:
  - Ubuntu LTS chosen as the only supported Desktop Linux after SL6
  - SL/EL: 4 year release cycle and 10 year life-time too long for rapidly changing desktop software environment (e.g. Chrome, Mendeley, …)
- One details of future Linux Desktop layout: $HOME
  - in the past, $HOME=/afs/... : transparent for users, bad for e.g. KDE
  - Proposal to users $HOME=/home, with some directories linking to /afs for data safety and global accessibility
  - Zeuthen users opted for continuation of AFS-only scheme, Hamburg users still digesting the proposal
- new model is a major cultural shift for our users ... Input from other labs?

DESYcloud

- Service has been reported on before
  - Based on ownCloud and dCache
  - Currently with about 800TB of dedicated backend storage
  - Long term plan is transparent melt with DESY dCache instances
- We mentioned going into full production around Christmas 2014
- In December we actually found some interesting shortcomings in the Linux NFS v4.1 client (kernel module)
  - which we fixed and submitted
  - and we used the opportunity to improve the dCache NFS server
- Currently we are in end-user pilot phase
  - Adding about 15 users a week on word of mouth alone
- Full production scheduled for mid-April.

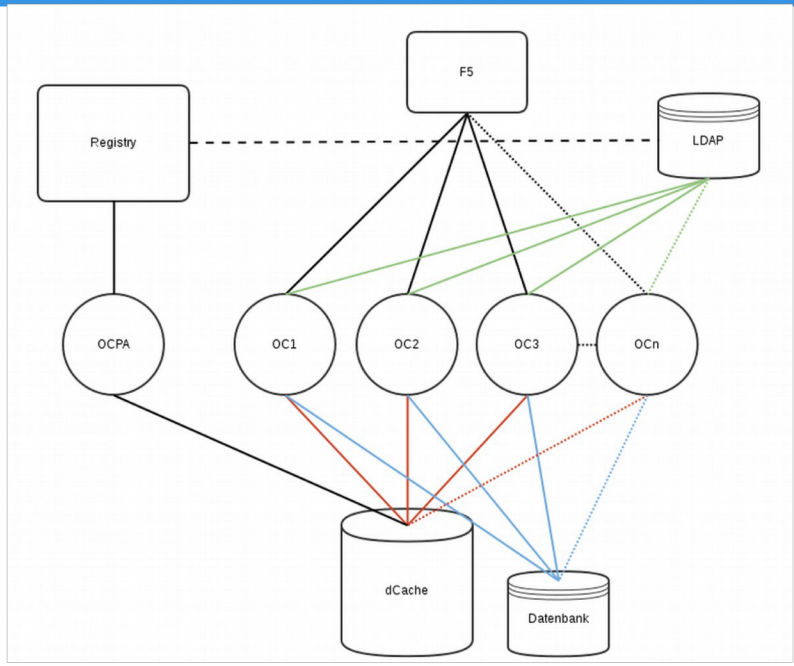

Computer center upgrade in progress

- Adding water cooling
  - in order to use floor space more efficiently
- Adding more electrical power
  - ramping up into single digit MW-Range for Photon Science Computing
- Adding new installation infrastructure for network and power connectivity
  - installed 'bathing tub' for water installation
  - electrical and networking installation from above
- Upgrade is happening while running the computing facility
  - 'Christo' tape library
  - large (heavy) plastic drapings
- Also upgrading core network link bandwidth to 40 Gb/s for PETRA III data taking ops

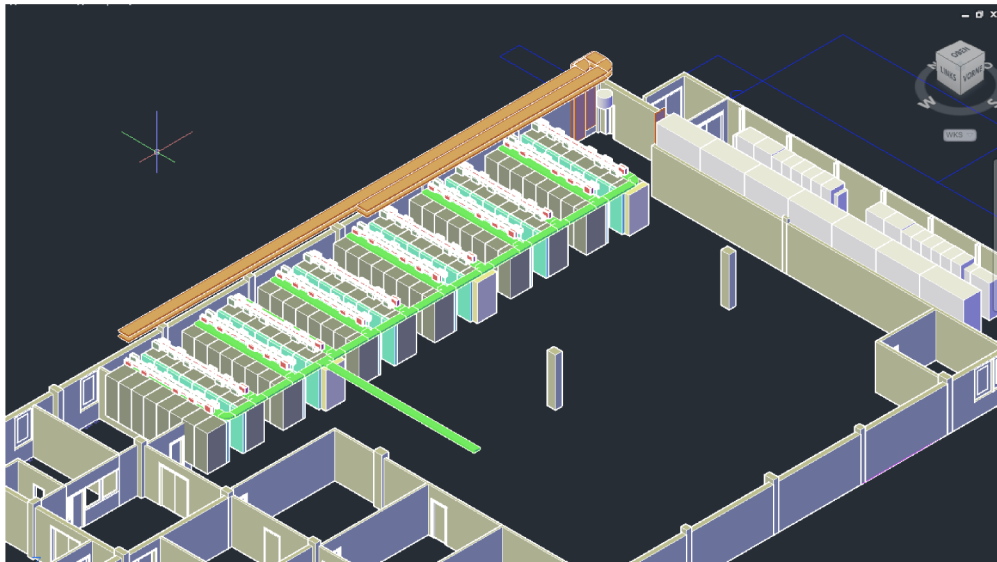

Photos courtesy of Stefan Bujack, DESY - IT

Photos courtesy of Stefan Bujack, DESY - IT

Photos courtesy of Stefan Bujack, DESY

More news in talks from

- Stefan Dietrich
  - ASAP3: New data taking and analysis infrastructure for PETRA III
  - Deployment and usage of MCollective in production
- Yves Kemp
  - ~~Storage on a disk ->~~
    Building large storage systems with small units: How to make use of disks with integrated network and CPU
  - BeeGFS at DESY
- and maybe Dirk Jahnke-Zumbusch
  - Migration of UNIX mail & Microsoft Exchange to Zimbra Collaboration Suite

Q & A

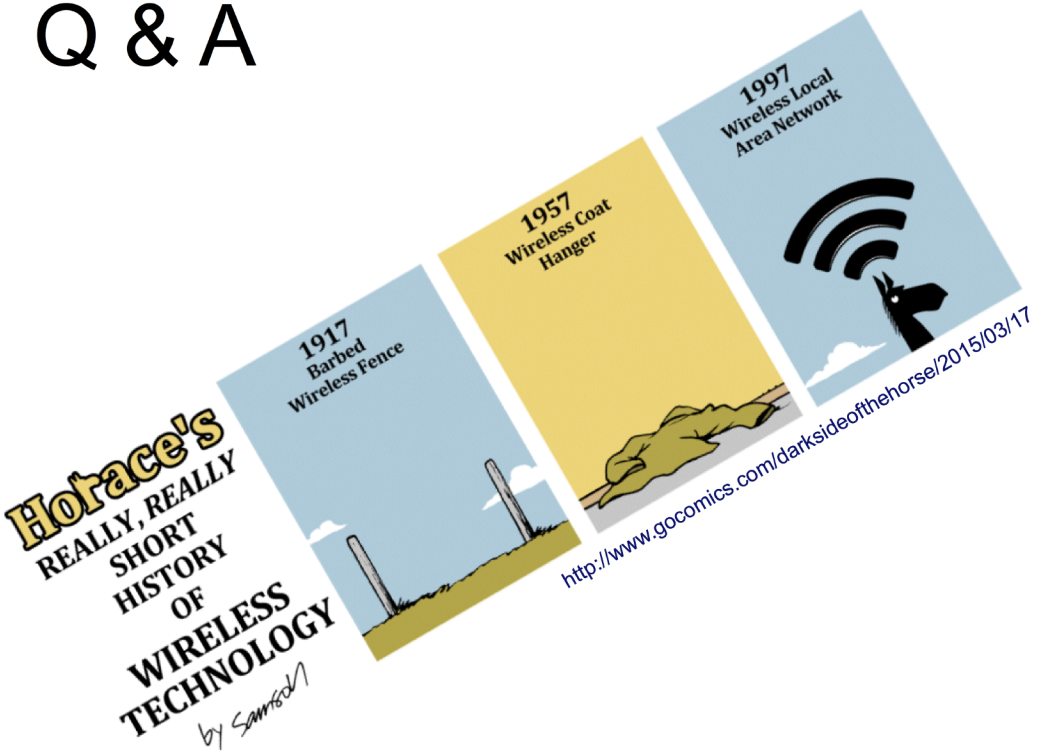
Horace's
REALLY, REALLY
SHORT
HISTORY
OF
WIRELESS
TECHNOLOGY
by Samson
1917
Barbed
Wireless Fence
1957
Wireless Coat
Hanger
1997
Wireless Local
Area Network
http://www.gocomics.com/darksideofthehorse/2015/03/17