



Site Report



Stephan Wiesand
DESY -DV -

October 19th , 2004

Outline



- Platforms
 - Windows
 - Linux
 - others
- Grid
- Mass Storage
 - dCache related projects
- APE
- Spam
- other

Windows



- new domain win.desy.de in production since January
 - [XP](#) (SP1 + Multi Language Pack German), [2003](#) (Servers)
 - [SP2](#) not yet rolled out due to problems with some applications
 - installation: [RIS](#) (servers & clients)
 - patches: [SUS](#)
 - virus scan: [McAfee VirusScan](#) Enterprise V7.1, AutoUpdate Architect
 - user file recovery: [Shadow Copy Client](#)
 - backup: [TSM](#), [Networker](#)
 - now using [Samba 3.0](#) for file and print services
 - applications made available via [NetInstall 5.7](#)

Windows



- **server hardware:** HP ProLiant DL 360, DL380, BL20p
- **storage hardware:** HP MSA 1000, Storageworks
- **migration** from old domain in **progress**
 - 1900 accounts in 60 OUs, 1500 PCs done
 - goal is 80% of PCs migrated by December
- see talk by Reinhard Baltrusch tomorrow afternoon
- **next projects:**
 - MS **Exchange** 2003 (hardware has arrived, system installed)
 - **terminal server** clusters
 - introduction of **ePolicy Orchestrator**

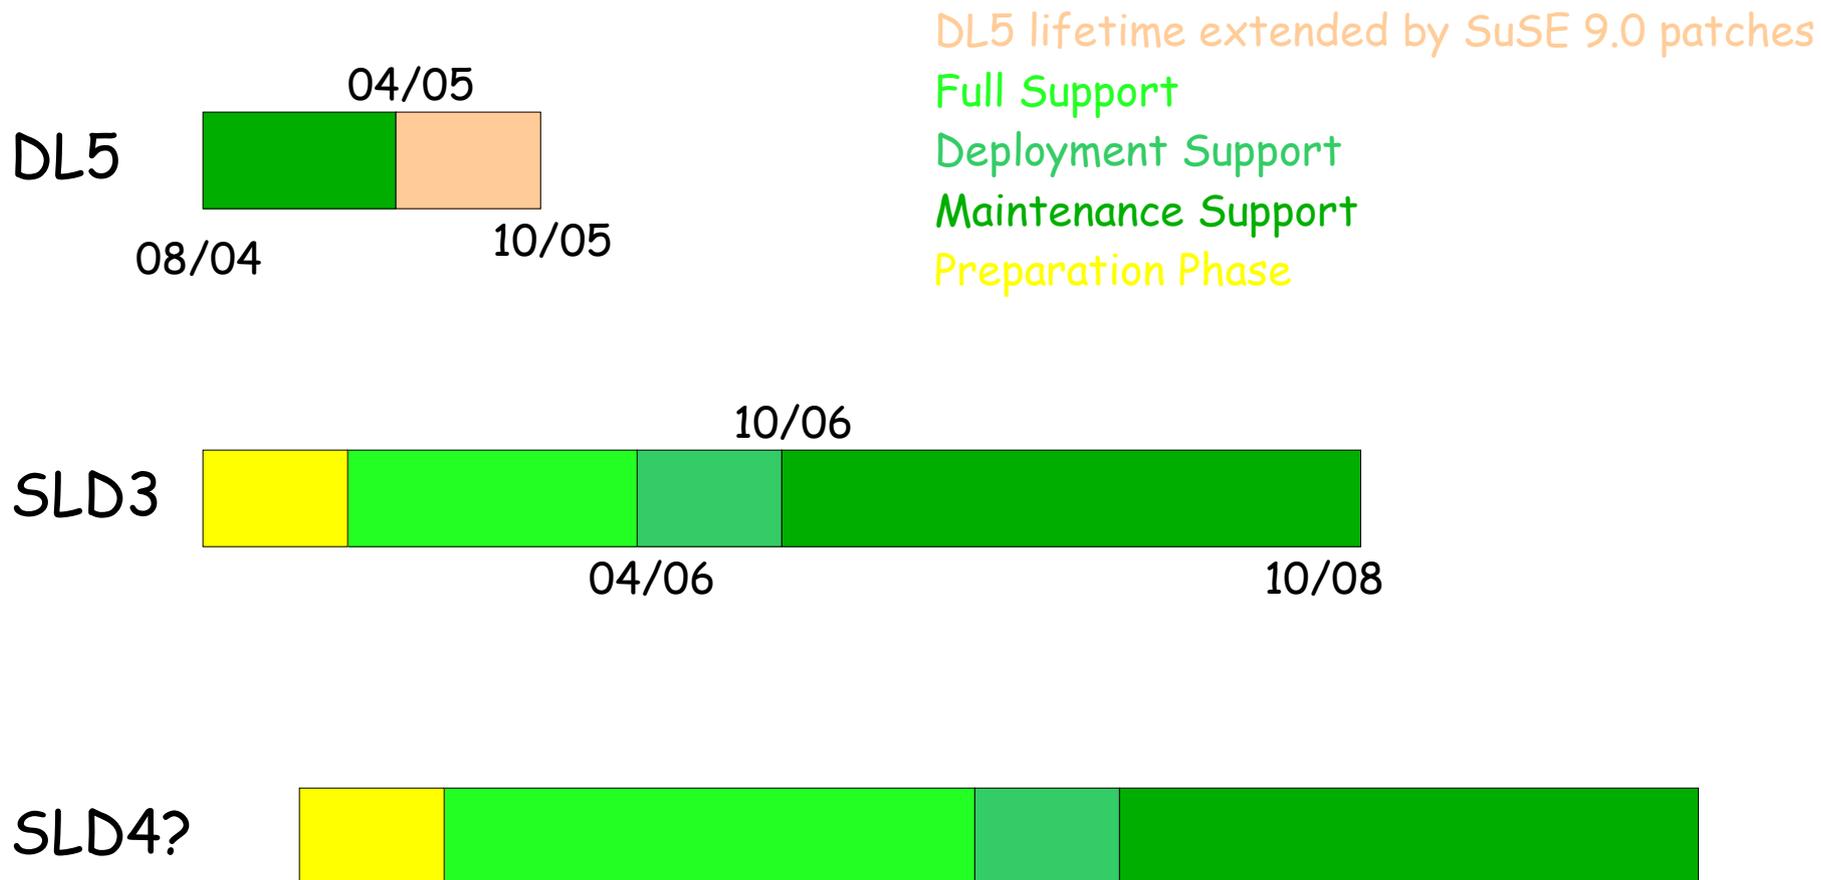


Linux: DL5 -> SL3

- DL5 (SuSE 8.2) rollout in progress (45%)
 - we clearly need longer living base distributions
- preparing introduction of SL3
 - user access to public preview system since end of July
 - detecting 1st unexpected incompatibilities
 - glibc locale functions
 - incompatible with SuSE and Red Hat <=8
 - affects older libraries & 3rd party products
 - photon injector group is porting DOOCs to NPTL
 - a few other groups have tested their code
 - migration should be smooth



Linux: Time Plans





Linux: AMD64

- first production systems running
 - still SuSE 9.0 professional
- more are ordered
- Opteron systems ordered for Infiniband test cluster
- SL3 **problems**:
 - AFS/krb5 - we now have workarounds
 - 32bit compatibility - improved with SL 3.0.3
 - see thursday talk
- new installations will be SL3



Other Platforms

- growing interest in **Solaris 10**
 - SPARC only
 - 1st installation server was set up
- recently had to reanimate one retired **HP-UX** host
 - one user group has code that won't compile on anything else
 - no plans to reestablish HP-UX as a user platform
- no plans for central OS X support
- remaining AIX and IRIX servers will go eventually
- still have a large number of debian linux hosts on site (w/o central support)

GRID



- participation in EGEE, D-GRID, ILDG
- "DESY Production Grid" up since August
 - successor of "Testbed" (EDG-1.4) and "Testbed2" (LCG-2)
 - LCG-2_2_0, official LCG-2 site in LCG Test Zone
 - RB, BDII, PXY, LRC, MDC, VOMS, UI, CEs, WNs, dCache-based SE
- HERA experiments H1, ZEUS are adapting their MC production
- support for University of Hamburg (CMS)
- see talk by Patrick Fuhrmann thursday afternoon



Tape Storage

- Hamburg: > $\frac{1}{2}$ PByte, > 4M files on tape
 - 4 STK powderhorn robots
 - 20 x 9940 STK drives (OSM)
 - 17 x 9840 STK drives (OSM) + 9 x 8940 STK (TSM backup)
 - migration d2 ampex adic -> STK 9940B: finished
 - migration 9840 -> 9940B: 50% finished (ETA: end of 2004)
- Zeuthen: ~ 60 TB online (compressed)
 - ADIC AML-J mixed media robot
 - 6 x DLT 4000 drives (now r/o)
 - 4 x LTO1 drives (OSM) + 4 x LTO1 (backup)
 - migration DLT -> LTO: ETA end of 2004



Tape Storage: Access via dCache

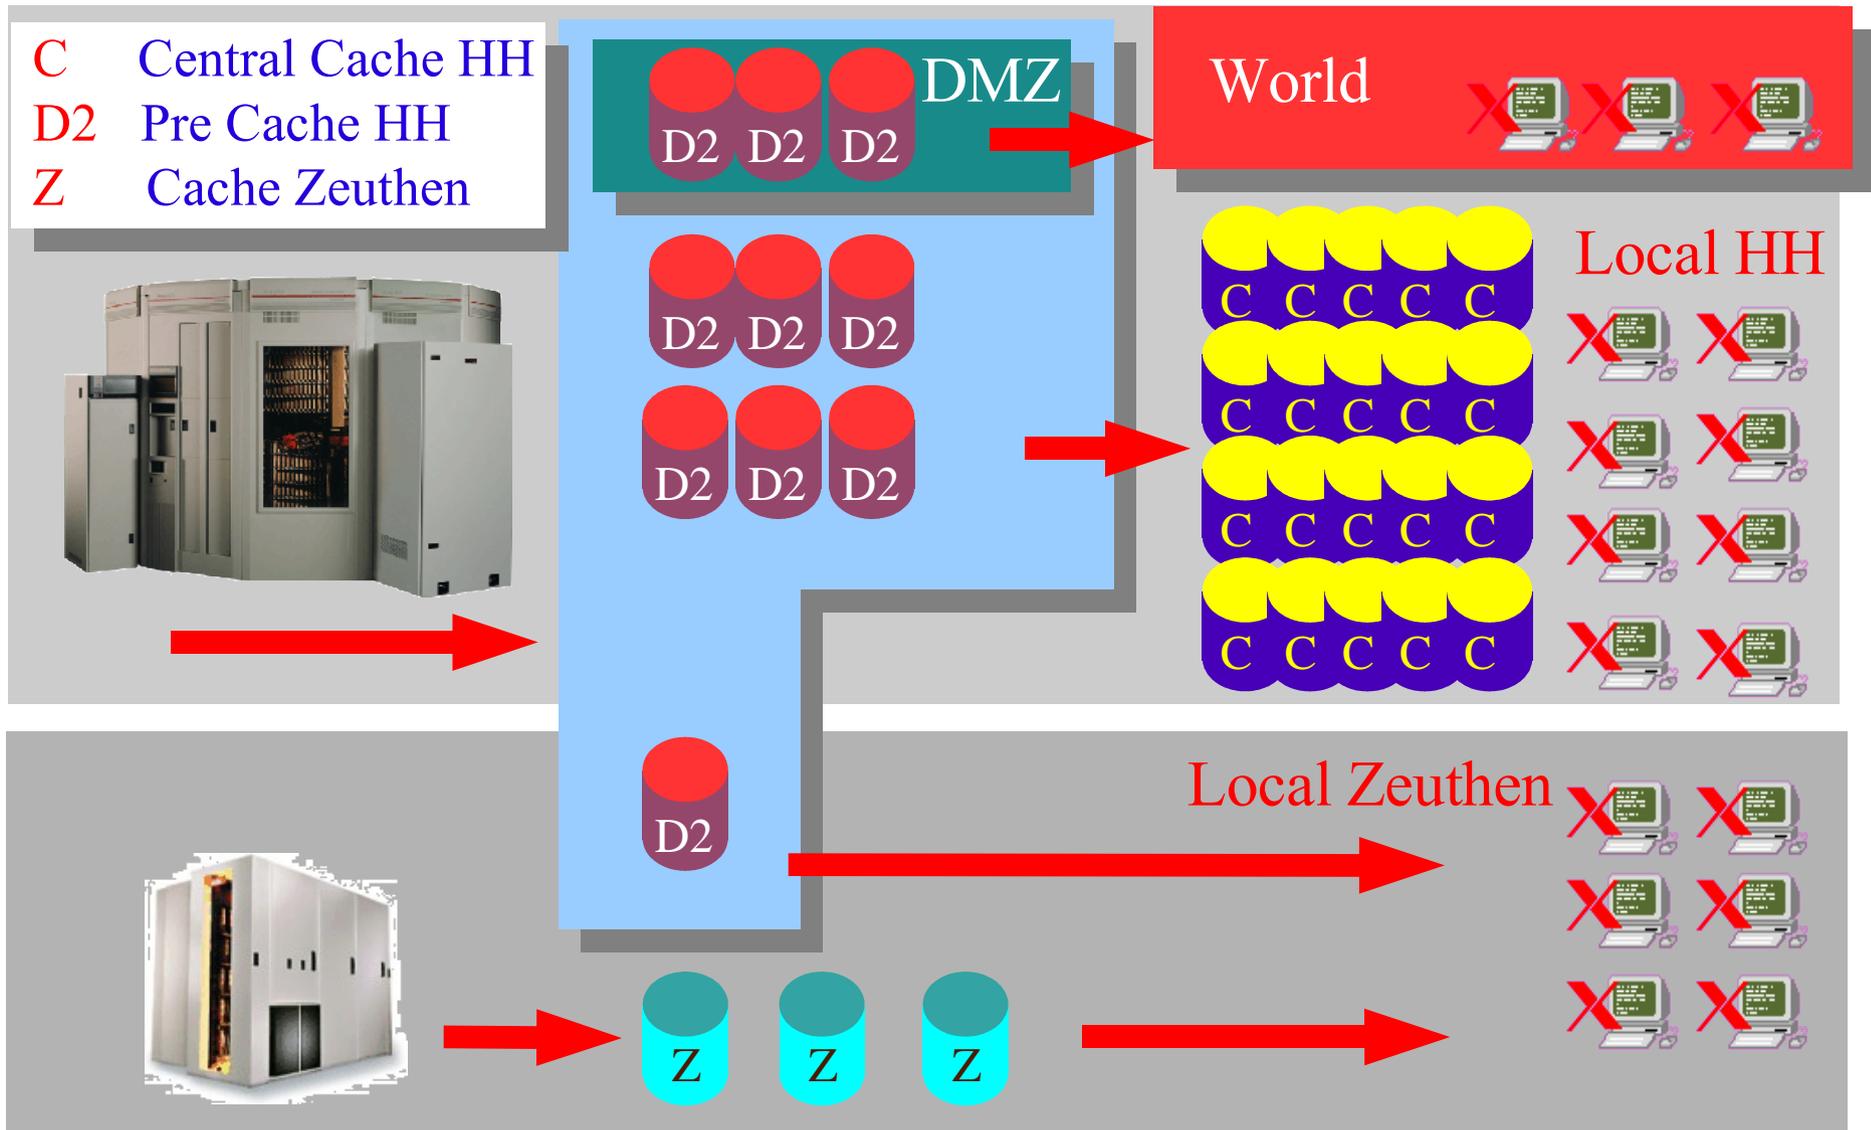
- full tape repository is accessible via dCache
 - local access by **dCap** protocol (POSIX like I/O functions)
 - remote access by **gridFtp / SRM** protocol
- DESY mass storage system is **LCG Storage Element**
- dCache disk space
 - **60 TB** (experiment owned) + **30 TB** (IT owned) **read cache**
 - **5 TB** (IT owned) **write cache**
- access profile:
 - **10-15 TB/day** (~ 150000 files) delivered to clients
 - **2 TB/day** restored from tape to dCache

d2Cache Project



- the problem:
 - STK 9940B characteristics:
 - 200 GB/tape => very large number of files per tape
 - (dis)mount counters already approaching tape spec limits
 - tape is nonsharable resource => long wait queues
 - mount/dismount and positioning operations are very slow
 - => long wait queues
 - we will not survive using the 9940B the same way as the 9840 before
- the idea (under evaluation):
 - introduce additional disk storage layer, logically located between HSM and central dCache

d2Cache Project



d2Cache Project



- contemplated solution:
 - d2Cache should **store** (nearly) the **full tape repository**
 - operates in "**virtual HSM**" mode:
 - **virtual tapes**
 - only very **limited number of concurrent read streams** allowed
 - **unused disks** will need to be **powered down**
 - **HSM** becomes a **backup system only**
- we need:
 - modifications to the dCache software (**easy**)
 - **inexpensive, very high density disk storage**
 - **partners with similar problems or goals**



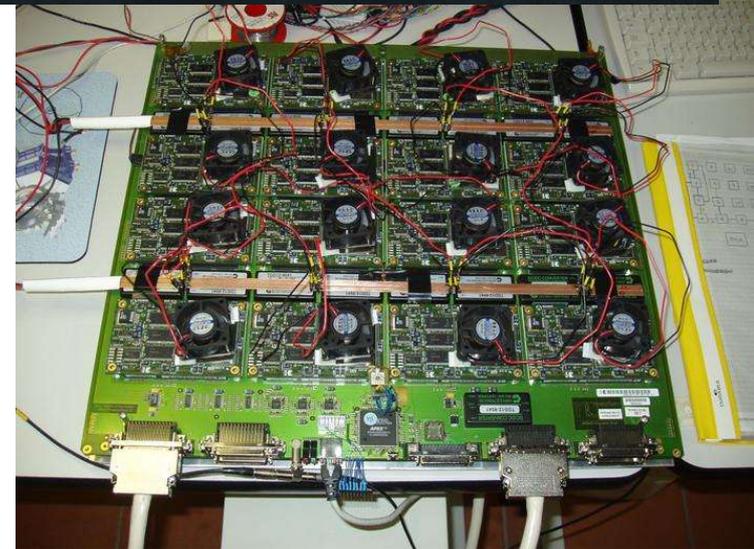
Project: Resilient dCache for H1 Experiment

- Goal:
 - make use of unused disk space on farm nodes
 - provide reliable, highly performant scratch space for jobs
- Solution: dCache in resilient mode
 - dCache server running on all farm nodes
 - using unused disk space
 - providing a distributed file space with a single virtual root
 - data sets are replicated to different farm nodes
 - compensates for node failures and scheduled downtime

apeNext



- new generation LQCD machine
 - INFN, DESY, Paris Sud
- 3 small prototype systems now
 - 64/64/8 CPUs
 - Ferrara, Rome, Zeuthen
- 1st physics codes have been ported
- two larger prototypes in Q1/05
 - 512 CPUs, 0.8 TFlops each



SPAM Fighting



- reported by Wolfgang Friebel for Zeuthen mail server:
 - tagging efficiency of SpamAssassin 2.x had dropped
 - from 99% to 95% (at a Score of 5.0)
 - SpamAssassin 3.x released in September
 - installed with features Razor2, DCC, SPF, Bayes, AutoLearning
 - efficiency is back to 99.5 %, close to no false positives
 - Mail exceeding a score of 10 is now rejected
- Hamburg plans introduction of commercial product
 - from Crocodile
 - SpamAssassin + rules maintenance



Other Activities & Plans

- considering use of **quattor** for Linux/Solaris management
 - servers, compute nodes, desktops, notebooks
- new **user registry** going into production now
 - step by step
 - now handles password changes
- **Web Office**
 - Zope/ZMS system **rollout in progress**
 - hosts **16+ virtual web sites now**
 - migration of **www.desy.de** main pages **in progress**
- use of **IT AMS** for accelerator magnets (tests)