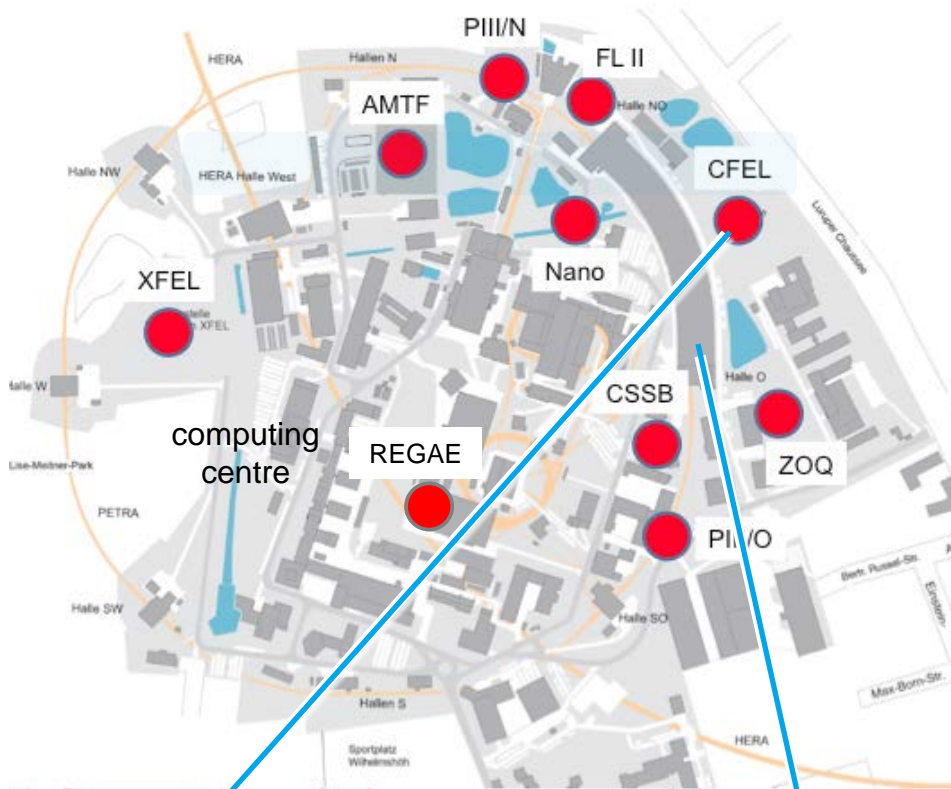


DESY Site-Report

HEPiX fall 2011 – Vancouver

Dirk Jahnke-Zumbusch
HEPiX fall meeting
Vancouver, 2011-10-25

News from the Lab: ongoing construction activities



- > PIII = PETRA III, north (N), east (O)
 - after DORIS III shutdown PETRA III will be substitute for hard x-rays
- > FLII = **FLASH II** extension
Free-Electron Laser in Hamburg
- > **Centre for Free-Electron Laser Science CFEL**
 - DESY / Max-Planck-Society / University of Hamburg
 - coherent imaging, coherent molecule imaging, ultrafast x-rays
- > **Centre for Structural Systems Biology CSSB**
research on infections
- > **Centre for Optical Quantum technologies ZOQ**
- > **Relativistic Electron Gun for Atomic Exploration REGAE**
- > Nano-Spectroscopy cooperations
- > last but not least: **XFEL**


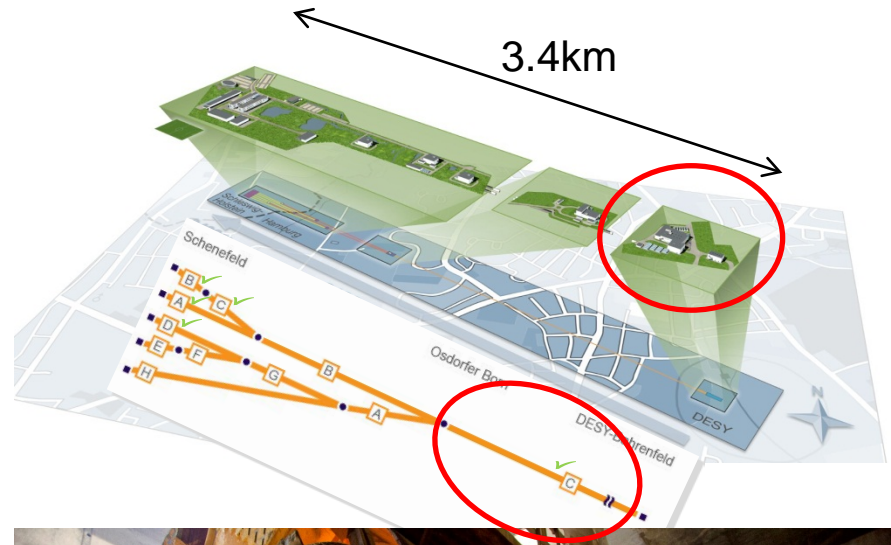


Courtesy: hammeskrause architekten



News from the Lab: tunnel construction for XFEL

- > tunnel boring in the backyard of the computing center ended
- > construction of buildings are ongoing
- > as is the „fan-out“ drilling (until summer 2012)
- > despite of vibrations no increase of defect hardware was noticed



> DESY-Zeuthen

as of Oct 1st Christian Stegmann is Zeuthen's representative in DESY's directorate; he holds an Astroparticle Physics professorship at the University of Potsdam

> DESY-Hamburg

- Knut Woller has left DESY as of Sep 30th
- Yves Kemp is team lead of the Systems Group
- Kars Ohrenberg is team lead of the Operations Group



Astroparticle Physics – in general

- > DESY will become a national Astroparticle Physics center inside the Helmholtz Association (HGF)
- > Grid IceCube–Tier1 continued, i.e. hosting of all 2nd level data (filtered raw), >50% of simulation data, 300~400 cores in batch farm, Grid support
- > use of ~10 GPUs for simulation of photon propagation in ice planned
- > simulations for theoretical Astroparticle Physics on HPC cluster started
- > participation in all gamma ray experiments – Magic, Veritas, HESS, CTA
➔ strong involvement of the computing center especially into the Cherenkov Telescope Array (CTA – see next slide)



- > Alma Common Software ACS from European Southern Observatory ESO
 - Containers – Component based framework
 - Employs several standard CORBA services like:
Notification service, Naming service, C++/Java/Python Orbs, Interface Repository – IR Browser GUI
 - Generic GUIs and tools for:
ACS Command Center, Logs Displayer, Object Explorer
Data base Explorer, Alarm Display System, ...
 - Code generation frameworks are available for ACS – Java (UML based) and C++ are available
- > The Control & DAQ software for the Medium Size Telescope prototype (MST) to be installed spring 2012 Berlin Adlershof will be based on ACS
- > Prototype Array Control system comprises Drive control, Weather station, Alignment System (CCD cameras), Advanced Mirror Control and a Camera emulation
- > The CTA remote Array Control Center approach will be tested between DESY at Zeuthen and Berlin Adlershof

Storage – dCache instances

- > DESY's dCache Tier-2 instance now around 4.7 PByte
- > DESY- & Grid instances – dCache used for Grid & NAF
 - 3 instances internally (HERA, H1, Photon-Science w/NFS 4.1)
 - 3 Grid instances (ATLAS, CMS, DESY (ILC, CALICE, CFEL)) in Hamburg
 - more Grid instances (ATLAS, LHCb, Astroparticle Physics, IceCube, CTA) in Zeuthen
- > 1 GridLab for **your** tests incl. assistance
 - **everyone who is interested in tests, e.g. NFS 4.1**
→ go for it, ask the GridLab-Team, Dima is here (Dmitry Ozerov) !
 - also xrootd, dcap possible
 - 100 TB storage
 - 32 clients with 256 slots possible
- > capacities
 - 3.5 PB in Hamburg
 - 1.2 PB in Zeuthen, ATLAS (750 TB), LHCb (180 TB), Astroparticle Physics (300TB)
- > versions
 - production: 1.9.12
 - lab instance: 2.1



Storage – continued

> AFS

- EAKC2011 – European AFS and Kerberos Conference at DESY
→ see indico.desy.de for slides
- also concerning HEPiX-Storage-Working-Group (see Andrei Maslennikov's talk)

> Lustre

- ◆ v 1.8.5 on servers, v 1.8.6-wc1 on SL6 clients
- ◆ systems are very stable
- ◆ quotas enabled for groups with daily e-mail reports
- ◆ capacities ~500TB locally, ~400TB NAF Hamburg and Zeuthen combined
- ◆ connected by DDR InfiniBand, bonded 1GE or single 10GE

> NetApp 6280

- virtualization
 - ◆ multi tenancy
 - ◆ virtual controller
 - ◆ privacy & isolation
 - ◆ delegated administration
- management software suite allows for configuration, deployment, performance analysis
- installed capacity of 1.2 PB (up to 4.3 PB)


> high performance bulk file store

- evaluation system with IBM SONAS installed with possible native GPFS
- UNIX-Windows-Integration issues, e.g. Kerberos cross-realm trust



NAF – National Analysis Facility

- > recent purchase of DELL C-blades
 - 2U with 4 systems, each 2 x hexa-core with 4 GB/core
- > DELL AMD-C blades might be interesting in future
 - 2U with 2 systems, each 4 x 12-core – purchased one testsystem
- > Lustre
 - replaced IB with 10 GE on Lustre server, good results
 - clients up to now do not saturate 1 GE – no need for 10 GE this will change with 48 core machines
 - Lustre does not fit our access patterns: in search of a replacement as “bulk file store”
- > users express interest in interactive system with graphic tools for collaborative work (across institutes)
- > fair share batch processing for all VOs



batch and GRID activities

> batch processing (fair share for all VOs)

▪ Hamburg


- ◆ problems with Torque & Maui when 5'000 jobs are scheduled
- ◆ 4'784 cores installed by 38kHS
- ◆ NAF and BIRD batch farms still on S(un)GE
- ◆ U(niva)GE w/ some startup problems
- ◆ S(on of)GE as first test soon

▪ Zeuthen

- ◆ Univa Grid Engine 8.0.0p1 in production, with support contract and good experiences with Univa
- ◆ 1'024 cores with tight MPI integration into GE, QDR-InfiniBand
- ◆ 960 cores for other jobs

> Grid

- gLite 3.2 in production
- Zeuthen: UMD-1 tests successful, planned migration during next months
- Hamburg: some systems running EMI-1 release
- EMI: European Middleware Initiative
- UMD: Unified Middleware Distribution
- some kernel problems with virtual worker nodes have been observed



- > DISH – DESY Image Sharing service is coming up
 - images for KVM provided
 - could also be other hypervisors
 - designed for our science community
 - the service will start next month
 - see also Owen's talk this week

- > cloud – plans for 2011
 - Open Nebula testbed now being converted
 - new switches ordered to enable multi subnet access from hypervisors
 - 1'000 IPs on dedicated public and private network
 - opening the new testbed to some users in 2011 planned
 - ask Thomas, he's around...



> Windows

- Office-XP → Office-2010 migration has been completed
- Windows-XP → Windows 7 / 64bit is ongoing, ~80% migration rate as target until the end of 2012; not many problems, except some older applications

> SL6

- starting deployment in Zeuthen
- around 50 machines running SL6.1, including 20 desktops
- hitting bug of Nehalem CPUs (system freezes for several minutes)
- workarounds (disabling C-States in the BIOS) lead to higher power consumption for idle machines
- see http://bugzilla.redhat.com/show_bug.cgi?id=710265
 - beta-tests in Hamburg
 - also ~50 machines
 - SL5 desktop / graphic-card issues vanished with SL6
 - general availability planned for December



> LHCone

- HEPPI → national network infrastructure between the Tier-1 and Tier-2s
- currently involved KIT (Karlsruhe), GSI (Darmstadt), RWTH (Aachen) and DESY
- uplink into LHCone: first connections with CERN and IN2P3

> IPv6 – HEPiX working group

- participation in testbed for IPv6-grid

> Ethernet news

- FCoE / Fibre-Channel-over-Ethernet
 - tests within data center
 - reducing cabling inside cold aisles racks
- raising port numbers for 10G-Base-T infrastructure (~400 ports) with Arista hardware

> XFEL call for tenders (as tunnel will be equipped soon)

- three distinct networks: office, data acquisition, machine controls

GO – Optimization of business processes

- > objective is the efficient and transparent processing of business processes
 - establish organizational and technical foundations
 - determine business processes
 - analyze, optimize and implement specific business processes
 - introduce a central identity and access management system (IAM)
 - connect selected information systems to the IAM
 - install a central workflow system
 - launch a central web-portal

- > some challenges are
 - acceptance & comprehension of different “cultures”:
e.g. business administration needs and researchers traditions
 - many identities of different origin
 - manifold roles
 - privacy of personal data in conjunction w/ non-DESY persons

- > helpful dialogue with friendly labs & institutions about methodology and experiences



HEPiX – other DESY contributions this week

> please note the following talks apart from DESY's Site Report

- Owen Syngde “Distributing images with image lists”
(Wednesday 10h00)
- Patrick Fuhrmann “EMI, the second year”
(Thursday 09h30)
- Thomas Finnern “Hepi-X-Perience”
(Friday 10h00)

> questions anyone?

> and thank you for your attention !

