



HEPiX report

Helge Meinhard, Jan Iven,
Alan Silverman / CERN-IT
Computing Seminar / After-C5

4 June 2004

Outline

- Overview, misc. topics (H. Meinhard)
- Linux releases and RedHat policy (J. Iven)
- Mass Storage workshop (A. Silverman)

HEPiX

- Global organisation of service managers and support staff providing computing facilities for HEP
- Covering all platforms of interest (Unix/Linux, Windows, Grid, ...)
- Aim: Present recent work and future plans, share experience
- Meetings ~ 2 / y (spring in Europe, autumn in North America)

HEPiX Spring 2004 (1)

- Held 24 – 28 May 2004 at NeSC, Edinburgh
- Format:
 - Mon – Wed Site reports, HEPiX talks
 - Thu – Fri Large Cluster SIG on mass storage and WAN data movement
- Very good (remote!) organisation by David Kelsey / RAL, with very active help by NeSC staff
 - Even the weather was good
- Full details: <http://www.nesc.ac.uk/esi/events/291>

HEPiX Spring 2004 (2)

- 100 participants, of which 17 from CERN-IT
 - Barring, Baud, Cancio, Dimou, Draper, Foster, Guijarro, Iven, Keeble, Kunszt, Meinhard, Panzer, Polok, Robertson, Siket, Silverman, T Smith
 - Other sites: FNAL, DESY, SLAC, PSI, NeSC, RAL, NIKHEF, Glasgow, IN2P3, BNL, Braunschweig, Jefferson Lab, U Edinburgh, Oxford, QMU, DAPNIA, GSI, Triumpf, CNAF, ICL, Barcelona, Caspur, Bristol, FZK, Stanford U, Manchester, UCSD, LBNL
 - Vendors: IBM, RedHat, Storage Technologies
- 60 talks, of which 19 from CERN

Next meetings

- Brookhaven 18 – 22 October 2004
- Karlsruhe, date to be confirmed (propably 9 – 13 or 23 – 27 May 2005)

NeSC: National e-Science Centre

- Founded in context of UK e-Science program (213 MGBP for personnel in 2001...2006)
 - which in turn is linked with GridPP, EGEE, LCG, ...
- Roles: help coordinate and lead UK e-Science, run e-Science Institute, undertake R&D projects
 - Projects focused on data management
- Industrial partners

Highlights: Storage

- SAN vs. NAS (disk arrays or storage in a box)
 - NFS: scaling problems (BNL)
- Better MTBF for disk arrays than for 3Ware-based systems (DESY)
- Horizontal disk mounting better than vertical one (RAL)
- Disk arrays
 - 4 Infortrend, 1 HW problem in 1.5 years (CASPUR)
 - Bad experience with another vendor (Glasgow)
- RAID, file systems
 - Moving to HW RAID 5 plus SW striping, using XFS (CERN)
- Lustre being looked at (LBNL, CASPUR, ...)

Highlights: Hardware, CPUs

- 6 HW failures/month on 1360 dual-CPU machines (1U rackmount) (BNL)
- 1 HW failure/day on 504 dual-Xeon blades (UBC)
- Fermi: Additional machine room in refurbished building 1.5 miles away
 - Lights-out operations
 - Cyclade AlterPath for console access, KVM, remote power on-off (authentication with krb5)
 - All cables overhead, false floor for cooling only
- Growing interest in Opterons
 - DESY and LAL mentioning 20 dual systems each for testing

Highlights: Operating systems

- Linux (see J. Iven's talk)
 - LBNL: CHOS, Linux Virtual Server
- Windows
 - Everyone moving to XP (client), 2003 (server)
 - Windows Terminal Server, Citrix getting ever more popular
- Solaris
 - Everything between 2.6 and 9 (10 in beta test)
 - CERN looking at N1
- Everything else (on its way) out... except
- MacOS X
 - Support requested by desktop users
 - Interesting features for file servers

Highlights: Security

- Bob Cowles (SLAC)
 - Local exploits as bad as remote ones
 - XP has shortened patch time for Windows, but very confusing for users
 - Attacks coming faster, getting smarter
- Linux: St. Michael (finds changes caused by root kits), Patchfinder (compares kernel addresses with known good status) (LBNL)
- Password sniffing getting a real problem, hence considering one-time passwords (LBNL)
- Non-centrally managed machines are a problem everywhere
 - Some sites have scanned imported laptops during Sasser
- CVS exploit during the meeting...

Highlights: AFS, Kerberos

- Transarc on its way out, OpenAFS (soon) used everywhere for server and client
- 'AFS best practices' workshop at SLAC
 - 100 attendees (mostly users including industry, some developers)
 - 70 signed up for tutorial
 - Commercial support for OpenAFS available
 - Much progress recently on Windows
- Kerberos
 - Many sites mentioned Kerberos 5 (Heimdal)
 - DESY wrote k5cron (replaces acron) and arcx (replaces arc – supports krb4 and krb5, portable across wide range of platforms, easily extensible via plugins)

Highlights: Miscellanea (1)

■ Batch schedulers

- LSF (BNL, SLAC, CERN, DESY, LBNL), Condor (BNL), Sun Grid Engine (DESY, LBNL), OpenPBS, Torque (FNAL), BQS (IN2P3), ...
- No clear tendency... except there is a clear interest to test Sun Grid Engine and Condor

■ Farm management

- Quattor and Lemon (displays) at CERN
- Some sites mentioned interest in Quattor and Lemon

Highlights: Miscellanea (2)

■ Grid

- LCG talks (Status, user registration and VO management, testing suites)
- GridPP overview and status

■ InDiCo

- Grown out of CDS Agenda
- Interest to use it for future HEPiX meetings