

# HEPiX Spring 2008 Meeting May 5<sup>th</sup> to 9<sup>th</sup>, CERN

## After C5 Report

Alan Silverman, Tim Bell, Tony Cass, Alex Iribarren, Andreas Peters





- Held at CERN for the first time since 1992
- Attracted 93 registered attendees
- Over 65 submitted contributions
- Started each day with a plenary on LHC, LCG, openlab, etc
- Split agenda into topics of which 4 will be summarised today
- All overheads online at <http://indico.cern.ch/conferenceDisplay.py?confId=27391>
- Full report available in CDS as 2008-05-19



- Some random highlights
  - Best attended talk was from Google on their impressive (to say the least) cluster sizes and the volumes of data they deal with
  - Most entertaining talk was from Charles Curran on how to make better use of magnetic tapes
  - Most scary talks were from Romain Wartel and Lionel Cons on (lack of) security, especially on the web
  - Talk from Sweden on the good news that IPv6 is (still) coming. [Oh yeah? But when?]
- Next meeting is in Taiwan on Oct 20 to 24.



## • Excerpts from Site Reports

- FZK has merged with Uni Karlsruhe and they now run a combined computer centre
- LAL fighting harmonics in their power network
- RAL had a fire alarm which closed down access control and prevented support staff re-entering to fix the problems
- RAL moving to CASTOR (from dCache) while CNAF is moving away (to GPFS and TSM)
- TRIUMF is one of only a handful of sites taking Windows VISTA seriously
- BNL has installed a SUN/STK for their HPSS service
- SLAC dominated by staff layoffs due to DoE budget cuts (as is FNAL which sent only one person)
- SLAC have installed a second SUN Black Box container

- Vista Deployment Policy
  - Do not force people to run Vista if they do not want
  - Encourage people who could run Vista to try and to run it
    - allows smoothly increase number of Vista computers before XP will have to be phased out
- Vista SP1 brings performance and reliability improvements, however requires significant effort to deploy
- Vista x64 will soon be available at CERN



# CD player story – episode 2



**DON'T STAY  
WITH OLD TECHNOLOGY  
TOO LONG!!!**

Subject: [REDACTED]

Dear [REDACTED],

(...)

About cassette player, I asked to factory. Unfortunately, this equipment is no longer available but it's still possible to play music on cassette through AUX input. You will just need a male-male cable (7.- Sfr at Mediamarkt store) and a walkman which could be plus on it. Easiest and cheapest solution.

Have a sunny day,

Oliver

- **Scientific Linux**

- Talk from Troy Dawson (Fermi, SL maintenance team)
- Maintain older versions (SL3 and SL4) until October 2010
  - Grid middleware availability on SL5 ? Autumn 2008
  - Some ongoing work to update OpenOffice and firefox
- Roll-out SL5
  - Redhat Enterprise 5.2 due soon
  - Base for SLC5
- Enhance XFS integration since widely used ?
  - Include XFS kernel module in SL5 as standard
- SL6 outlook
  - Synergies with CentOS ?
  - Why not drop SL all together ?



# CPU Technology Track

- CPU-Level Performance Monitoring with Perfmon  
- Andreas Hirstius
- Status Report from the Benchmarking Working Group - Helge Meinhard
- The CERN Benchmarking Cluster - Alex Iribarren
- Benchmarking ATLAS Applications - Franco Brasolin
- Benchmarking CMS Applications – Gabriele Benelli
- Benchmarking ALICE Applications – Peter Hirstov
- Benchmarking LHCb Applications – Hubert Degaudenzi
- Benchmarking in Production Environment – Iwona Sakrejda (LBNL/NERSC)





# CPU-Level Performance Monitoring with Perfmon



- By Andreas Hirstius (CERN)
- Perfmon: interface to CPU's performance counters
- Tools improved to cope with HEP/CERN software
- No measurable performance impact

Test	SPECint2006	SPECfp2006
std. SLC4.6 w/ std. Kernel	34.26	25.60
std. SLC4.6 w/ 2.6.24.2 Kernel & perfmon patches & pfmon <i>inactive</i>	38.16 (111%)	27.77 (108%)
std. SLC4.6 w/ 2.6.24.2 Kernel & perfmon patches & pfmon <i>active</i>	38.23 (112%)	27.78 (108%)

# CPU-Level Performance Monitoring with Perfmon, II



- SPECint2000:
  - ~1% floating point
  - Slightly faster in 64bit mode
- SPECint2006:
  - ~0.1% floating point
  - Slightly faster in 64bit mode
- SPEC2006 C++:
  - 10-14% floating point
  - Slower in 64bit mode
- LXBATCH Public sample:
  - ~10% floating point (~70% 32bit, ~30% 64bit)

# Status Report from the Benchmarking WG



- By Helge Meinhard (CERN)
- Introduction, Motivation
- Agreed environment:
  - Default SL4 64bit, 32bit benchmarks
  - GCC3 compiler
  - LCG-SPI flags (`-O2 -fPIC -pthread -m32`)
- Agreed methodology, benchmarks:
  - Multiple parallel runs
  - SPECint2000, SPECint2006, SPECfp2006
- 15 machines (DESY, RAL, INFN, CERN)
- Summary of preliminary results

# The CERN Benchmarking Cluster

- By Alex Iribarren (CERN)
- Description of the Benchmarking cluster
  - 11 machines
  - Representative sample of our CC
- 9 benchmarks
  - SPEC2000, SPEC2006, 32bit, 64bit, rate, parallel
- Conclusions:
  - SPEC2006 scales well with SPEC2000
  - 64bit is better than 32bit
  - Parallel is better than rate



- By Franco Brasolin, Alessandro De Salvo (INFN)
- ATLAS Benchmark Suite
  - KitValidation: software testing framework
  - Global KitValidation Portal
- Preliminary results based on ATLAS release 12.0.6
  - Currently rerunning with release 14.1.0



- By Gabriele Benelli (CERN)
- Benchmarks with CMSSW\_2\_0\_0\_pre5 release
  - 7 different physics processes
  - GEN, SIM, DIGI, RECO
  - Future work: CMS benchmark suite
- Some results:
  - CMS scales well with SPEC
  - Some benchmarks are much better with AMDs
    - Especially RECO for simpler events
  - CMS scales very well with the number of cores



# Benchmarking ALICE Applications

- By Peter Hristov (CERN)
- Benchmarks:
  - 2 different types of events
  - GEN+SIM, DIGI, RECO
- Results:
  - ALICE scales well with SPEC
- Raised interesting questions about error estimates
  - Currently being investigated by the WG



# Benchmarking LHCb Applications

- By Hubert Degaudenzi (CERN)
- Benchmarks:
  - Gauss: generation and simulation
  - Brunel: reconstruction
- Results:
  - Scales well with SPEC
  - Scales better with int than fp
- Need more tests





# Benchmarking in Production Environment




- By Iwona Sakrejda and Jay Srinivasan (Lawrence Berkeley National Laboratory NERSC)
- Hostile environment
  - No control of network or storage load
  - Different storage pools
- Observations:
  - AMD architecture well-suited for their workloads
    - But you need to have enough memory
  - Multi-core systems need more bandwidth
    - Need to upgrade network infrastructure



# DATA CENTRE MANAGEMENT

# Power and Cooling...





# Infrastructure Trends...



- Power Density: Up
- Machine Room Temperature: Up
  - Machine Room Noise: Up
    - Why on earth do you want to be in there?
- Cold/Hot aisle enclosure
- Rack standardisation
- Increasing rack depth
  - Sun and IBM going in the opposite direction...
- Increased environmental monitoring
- “Go” to “In Service”: ~2 years
  - BNL managing 18 months, through keeping size and budget down; 30 months seen as more typical.
  - Do we have a ~2 year model now?

# A worrying issue?





- IPMI is good...
  - ... and CERN is exploiting it well
    - Why on earth do you need to be in the Machine Room?
- Monitoring is becoming (serious) automation
  - And, again, CERN is well placed
    - Who needs people in the machine room?
    - But, as usual (?) we undersell ourselves.