Data Preservation Activities at DESY

David South (DESY)

on behalf of the DESY Data Preservation Group

Fourth Workshop on Data Preservation and Long Term Analysis

Tsukuba, Japan, 8 - 10 July 2010

Contents

- DESY Data Preservation Group
- Preservation models, virtualisation and validation
- Data formats for preservation
- Hardware
- Documentation
- Governance
- Towards requests for funding
- Summary

The DESY Data Preservation Group

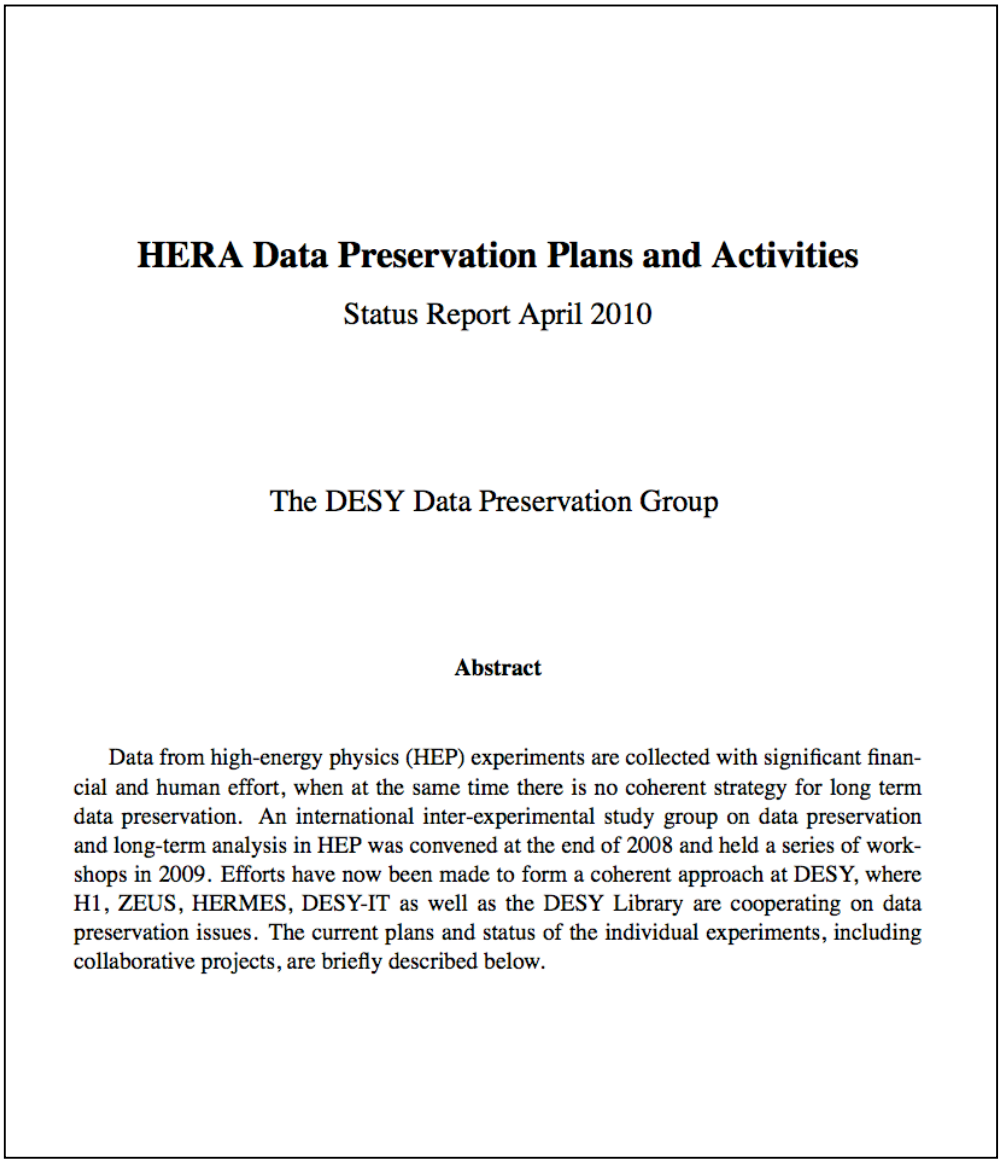

HERA Data Preservation Plans and Activities

Status Report April 2010

The DESY Data Preservation Group

Abstract

Data from high-energy physics (HEP) experiments are collected with significant financial and human effort, when at the same time there is no coherent strategy for long term data preservation. An international inter-experimental study group on data preservation and long-term analysis in HEP was convened at the end of 2008 and held a series of workshops in 2009. Efforts have now been made to form a coherent approach at DESY, where H1, ZEUS, HERMES, DESY-IT as well as the DESY Library are cooperating on data preservation issues. The current plans and status of the individual experiments, including collaborative projects, are briefly described below.

- HERA data preservation document prepared for the PRC in April, well received and efforts congratulated
  - Significant contributions from H1, ZEUS, DESY-IT, DESY-Library
  - Data preservation now seems to be "part of" the PRC agenda

Data Preservation Models Identified by DPHEP

Preservation Model	Use case
1. Provide additional documentation	Publication-related information search
2. Preserve the data in a simplified format	Outreach, simple training analyses
3. Preserve the analysis level software and data format	Full scientific analysis based on existing reconstruction
4. Preserve the reconstruction and simulation software and basic level data	Full potential of the experimental data

Cost, complexity, benefits

- Only with the full flexibility does the full potential of the data remain
  - Level 4 type programme was required by JADE and ALEPH re-analyses
- H1 aims for level 4, ZEUS somewhere between 3 and 4
  - Different approaches, can benefit from each other's experiences
  - A level 2 scheme for outreach using HERA data, collaborating via DPHEP, could (should, will?) also be pursued

The Virtualisation Route

- Virtualisation of the ZEUS software, running environment and MC production
- The prospect to simulate new MC after the end of the current analysis model led us to develop the concept of virtualising MC production
  - Based on the current GRID production scheme
  - Include the whole chain from generation to simulation to common ntuple production (see later)
- Prototype environment created using VirtualBox
  - All of ZEUS software included with additional environmental setup for recompilation and running test analysis programs
  - Removed present dependencies (AFS, oracle, storage)
  - Calibrations, conditions etc. put inside (3-4 GB)
  - Still a prototype – not fully usable for a regular ZEUS member

More like a Rolling Preservation Model

- At least for the analysis level H1 plans a rolling model of preservation, with a timescale of say 3 months interval
  - Regular recompilation of analysis level software
  - Full data production of μODS/HAT (analysis level) files, MC as well..

    Some numbers from current times:

    Read and copy 13.5 Tb of HERA II DST format data to Grid working nodes

    900 Grid jobs each running on average 20 hours

    Produce 1.3 Tb of HERA II μODS/HAT format data

    In ideal conditions: 1 day to produce data, 1 day to download from Grid

- Defining a strategy for such a rolling preservation model
  - Always use newest versions or freeze external software?
  - Would aim to at least incorporate ROOT updates
  - Continue using the database / have a snapshot of it?
  - More extreme (level 4) version: Adopt change in OS, include Fortran..

    This requires guaranteed manpower

Isn't it obvious, *Virtualisation* will solve everything?

My first and very naïve ansatz

- OK, why don't we just put everything an a virtual machine?
  - Data archival is done elsewhere, just need "to plug that into the VM"
  - Your VM contains everything you need to develop and run code and analysis
- The problem would then be reduced to maintain virtual images, and maintain their ability to run. In the Cloud era, seems like a trivial task
- Problems: Everything in IT is a moving target:
  - Will your network always be the same?
  - Will your access protocol always be the same?
  - Are you sure you do not need new software (e.g. MC generators) that require a new OS?
  - Are you sure your i386/SL4 VM will produce the same results when emulated on a quantum computer in NN years?
  - What about service you need, like CondDB,…
- Naïve virtualization will not work… but still, virtualization can help

Yves Kemp | Long Term Data Preservation and Virtualization | 22.6.2010 | Page 5

Yves Kemp (DESY-IT)

Freezing vs Rolling (or "Test-driven migration")

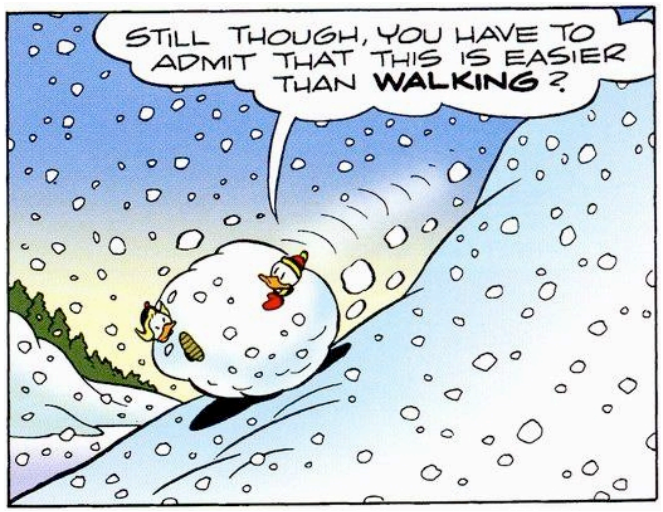

- **Pro Freezing**
  - One-time effort, very small maintenance outside of analysis phase
  - Also allows software w/o code (but might fail with DRM / licensing issues)
- **Pro Test-driven migration**
  - Usability and correctness of code is guaranteed at every moment
  - Data accessibility and integrity can be checked as well
  - Fast reaction to standard/protocol changes
  - General code quality can improve, as designed for portability and migration
- **Cons Freezing**
  - Rely on certain standards and protocols that may evolve
  - Potential performance problems
- **Cons Test-driven migration**
  - Needs long-time intervention, more man-power and resources needed
  - Some knowledge of the frameworks must be passed to maintainers

Yves Kemp (DESY-IT)

H1: Validation Tools

- If we want to have anything like a dynamic "rolling preservation model", then we will need good validation tools
  - Such a scheme already exists to validate file content of analysis level software between releases
  - In addition we have nightly builds
  - We will expand this to include full analysis selections
  - Validation tools now being developed for the Fortran (simulation/reconstruction) part of the H1 software

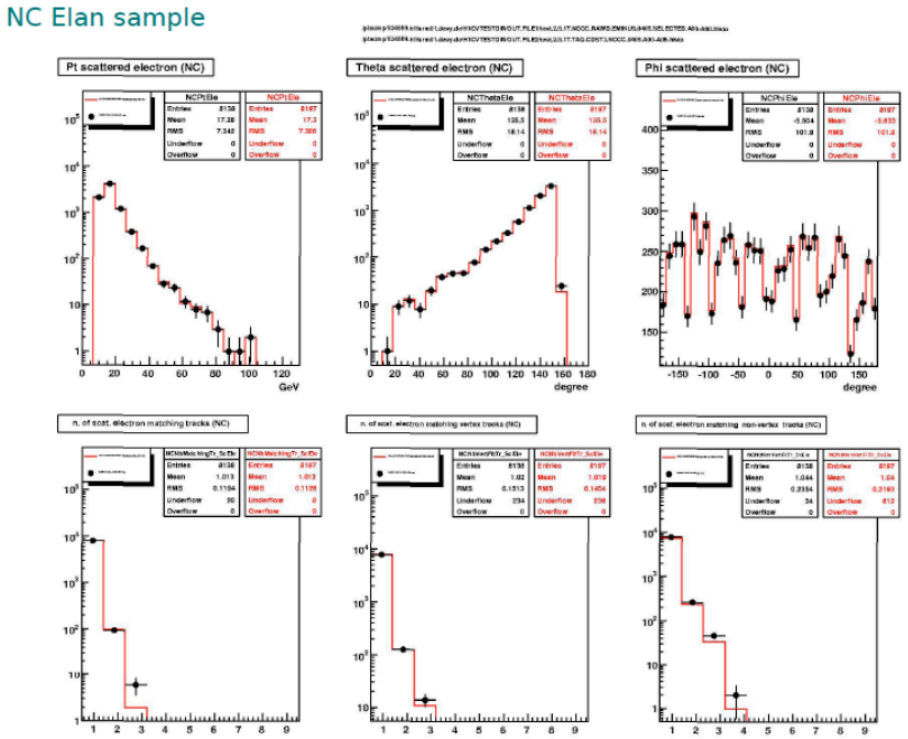

- What about a unified validation suite, which can compare different DSTs, releases and even analyses running under different operating systems… *for H1, ZEUS, Experiment X..?*

Towards a Generic Solution

- Validation and standards using virtualisation project by DESY-IT

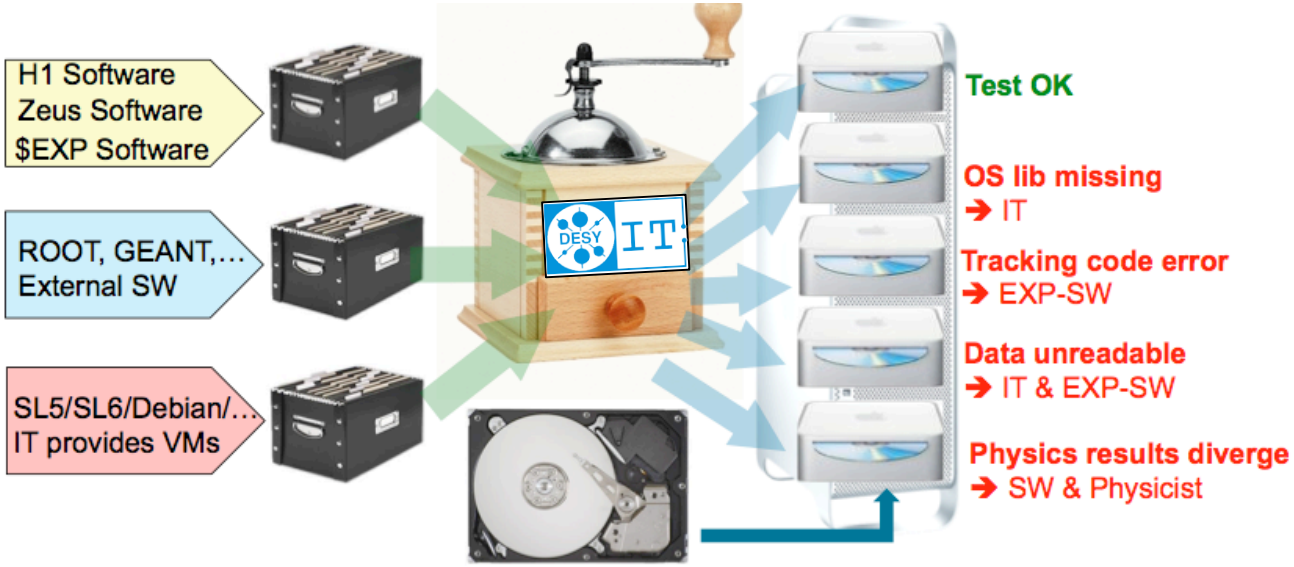

Clear separation between providers of input.

Automated VM image generator provided centrally.

Tests defined by $EXP.
Test data store provided by IT.

Different VMs run SW and tests. Depending on results, different action needed.

  - Detect incoherence in absence of intensive human survey
  - Useful collaboration for future OS transitions and preservation
  - Person power needs are being evaluated: Test with 5% pilot project

Yves Kemp, Dima Ozerov (DESY-IT)

Workflow: One test in detail (5% mock-up)

Experiment XYZ:
`application.sh`
`test.sh`
`files.tgz`
Put into VM image

application.sh:
-Can be precompiled executable
-Better: Compile source code
test.sh
-Do something with the binaries obtained in application.sh

Experiment XYZ:
Check extracted logs and analyse resulting files

Prepare test VM

Run VM and perform tasks within it

Extract results from VM and analyze them

IT:
Raw VM image
system_rpm.sh
Preparation for automatic execution

IT:
Run VM
Monitor execution

IT: extract
`system.log`
`app.log`
`test.log`
`output.root`
Put on webserver
Analyze errors

Example result: July 6th 2010 runs:

	SL4	SL5	Fedora 13	
ROOT V5.26	-no F77 compiler gfortran found -libX11 MUST be installed	Estimated ROOTMARKS: 1534.29	Estimated ROOTMARKS: 1512.76	**Compilation**
H1Data analysis	Processed 47243 events with J/Psi candidates Histogram written to jpsi_mods.root	Processed 47243 events with J/Psi candidates Histogram written to jpsi_mods.root	Processed 47243 events with J/Psi candidates Histogram written to jpsi_mods.root	**Run pre-compiled tgz using compat libs**
ZEUS MC prod	> ls -lh ZEUSMC.HFSZ627.E8954.GRAPE.Z01 4.2 MByte	> ls -lh ZEUSMC.HFSZ627.E8954.GRAPE.Z01 4.2 MByte	> ls -lh ZEUSMC.HFSZ627.E8954.GRAPE.Z01 4.2 MByte	**Run pre-compiled tgz using compat libs**
HERA-B	Compilation OK DB connect fails	Compilation OK DB connect fails	Compilation failed – needs code change	**Compilation**

+ Hermes to come soon

Migration to Latest* Operating System

- Both collaborations plan to move to SLD5
  - Progress made in all parts of H1 software, but perhaps surprisingly the Fortran world causes some headaches, especially with the compiler
  - SLD4 → SLD5 means gcc-3 → gcc-4, which means g77 → gFortran
  - And SLD5 is gcc-4.1 whereas there is now gcc-4.4..
  - Care needed with options (allowed line length, variable initialisation..)
  - Remaining problem is GKS, which does not compile under gFortran
- Further external dependencies at H1OO (analysis) level:

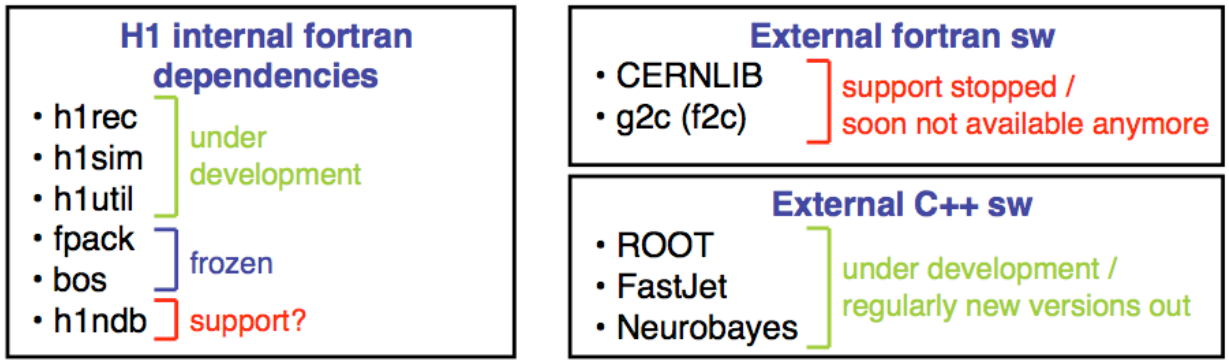

- Widespread use of ORACLE could also cause problems

**SLD6 comes in 2011*

H1 Virtualisation on the CERN VM

H1OO Virtualization Using CernVM Software Appliance.

How to get H1 Collaboration OO analysis framework and the grid User Interface on your laptop.

Mihajlo Mudrinic (H1, Belgrade)

Step-by-step Instruction

- Install VMware Player or VirtualBox
- Download Latest CernVM
- Untar the file and open it with your VM Software.
- Play your CernVM image and wait until the end of boot process.
- Read out your IP Address.Fig1
- Open a web browser on your computer, and point to the IP address.Fig2
- Type user: admin password: password.
- Change the admin password.Fig3
- Setup an local user (Group Must be hone!!).Fig4
- Click on preference, "VO set to hone", open advance option and choose "enable grid user interface".Fig5
- Wait until CernVM rebootsFig6
- Login and type: source /opt/hone/etc/login.Fig7
- Good Luck!Fig8

Special Note for VirtualBox Software

Our suggestion is to use "Bridged Networking". The guest will obtain its IP address in the same way that the host does.

You can find on CERN wiki page special instructions for VirtualBox Users.

Introduction

Virtual machine software (VMware, VitrualBox, ...) is software allowing one to run two operating systems simultaniously on a single machine (you laptop). To promote the idea of using HEP data in scientific training, education and outreach we build "virtual" H1OO Linux image using CernVM Software Applience. CernVM is a CERN based R&D project which delivers a thin Virtual Software Appliance already used by LHC experiments (ATLAS,ALICE,CMS,LHCb). We would also like to thank the members of CernVM R&D project for dedicating one VM on the CERN domain to The H1 Collaboration on which we can build and publish "virtual" Linux images with the preinstal H1OO analysis framework (release 3.4.14).

System minimum requirements

- **Windows users** : desktop or laptop PC running Windows with the VMware Player/VirtualBox software installed (free software)
- **Macintosh users** : desktop or laptop Macintosh running Mac OS with the WMware Fusion/VirtualBox software installed (shareware/free software)
- at least a **1 GHz** processor
- at least **1 Gb** of RAM for the PC or Macintosh
- **2 Gb** of free disk space available

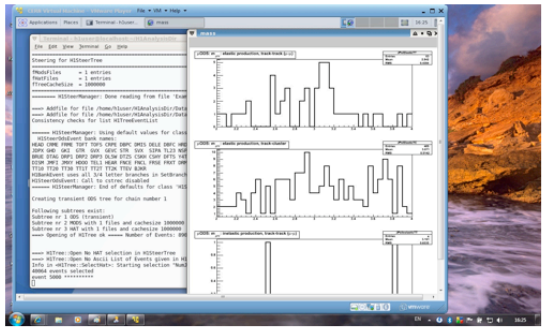

CernVM
SoFTware Appliance

The H1 Collaboration CernVM

Last updated 12.06.2010

- Studies of virtualisation ideas using H1OO within the CERN VM
  - Nice example of running analysis without need for a network connection, using a virtual image of the H1 environment
  - New form of simple laptop installation of H1 software
  - Access to the [*large scale*] data remains an outstanding problem

ZEUS Data Format: Common Ntuple Project

- The archival system deployed by ZEUS for long term preservation is based on Common Ntuple project
  - Simple flat ROOT format
  - Wide content to allow full physics analysis
  - Iterative production incorporating new content and many improvements / additions in physics analysis tools – driven by the usage
- 4th (full) iteration of HERA II data and corresponding MC samples already available last year
  - Several preliminary results presented already at DIS 2010
- 5th (partial) iteration for content revision – end of 2009
- 6th (full data, partial MC) iteration – used already by several analysis for *ICHEP10* preparation
  - Due to internal constraints of CERNLIB – restricted content of PAW ntuples
  - Access enhancements tests with ROOT 5.26.00 are ongoing

H1 Data Formats for Preservation

- Data formats to be preserved
  - RAW data of good and medium runs: 75 TB
  - At least one full set of DSTs, total for HERA I+II: 18 TB
  - A version of common analysis level format, μODS and HAT (< 3 TB)
  - In addition to calibration and cosmic runs, total data about 100 TB
  - Amount of MC to be decided, but will at least be of the same order

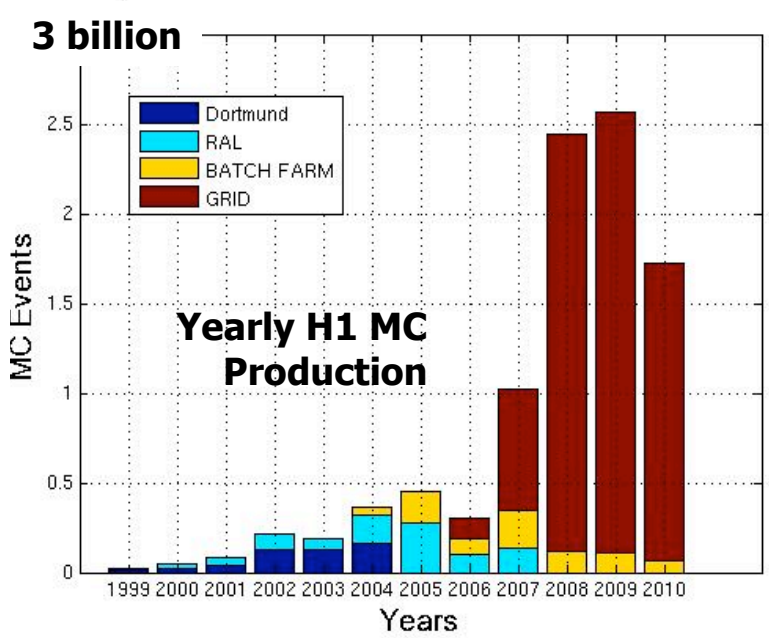

ELAN High Q^2 MCs	NC						CC
	$Q^2>60$	$Q^2>100$	$Q^2>1000$	$Q^2>10000$	$Q^2>15$, $y>0.3$	$Q^2>40$	$Q^2>100$
	2M	*2M*	*500k*	*100k*	*2M*	*2M*	*200k*
0304LH	7007	7004	7005	7006	7030	7031	7008
0304RH	7013	7010	7011	7012	7032	7033	7009
0405LH	6995	6992	6993	6994	7024	7025	6996
0405RH	7001	7059	6999	7000	7026	7027	6997
2006LH	6991	6988	6989	6990	7028	7029	7002
2006RH	6987	6982	6983	6984	6985	6986	7003
0607LH	7017	7014	7015	7016	7034	7035	7018
0607RH	7023	7020	7021	7022	7036	7037	7019
460GeV	-	7038	-	-	7039	7040	-
575GeV	-	7041	-	-	7042	7043	-

To be updated

MC sets for long term now being defined

- Conservatively (x2) estimate total amount to preserve at 500 TB

HERA Outreach Format, Global HEP Outreach Project?

- A HERA outreach format is technically within reach and several ideas are there
  - See H1 talk from last DPHEP workshop
- Can we start off running by opening up the BaBar outreach wiki effort to be HEP wide?
  - By making a true HEP data portal for outreach from the beginning
  - Support for mediaWiki from DESY-IT
  - But no progress here due to lack of manpower

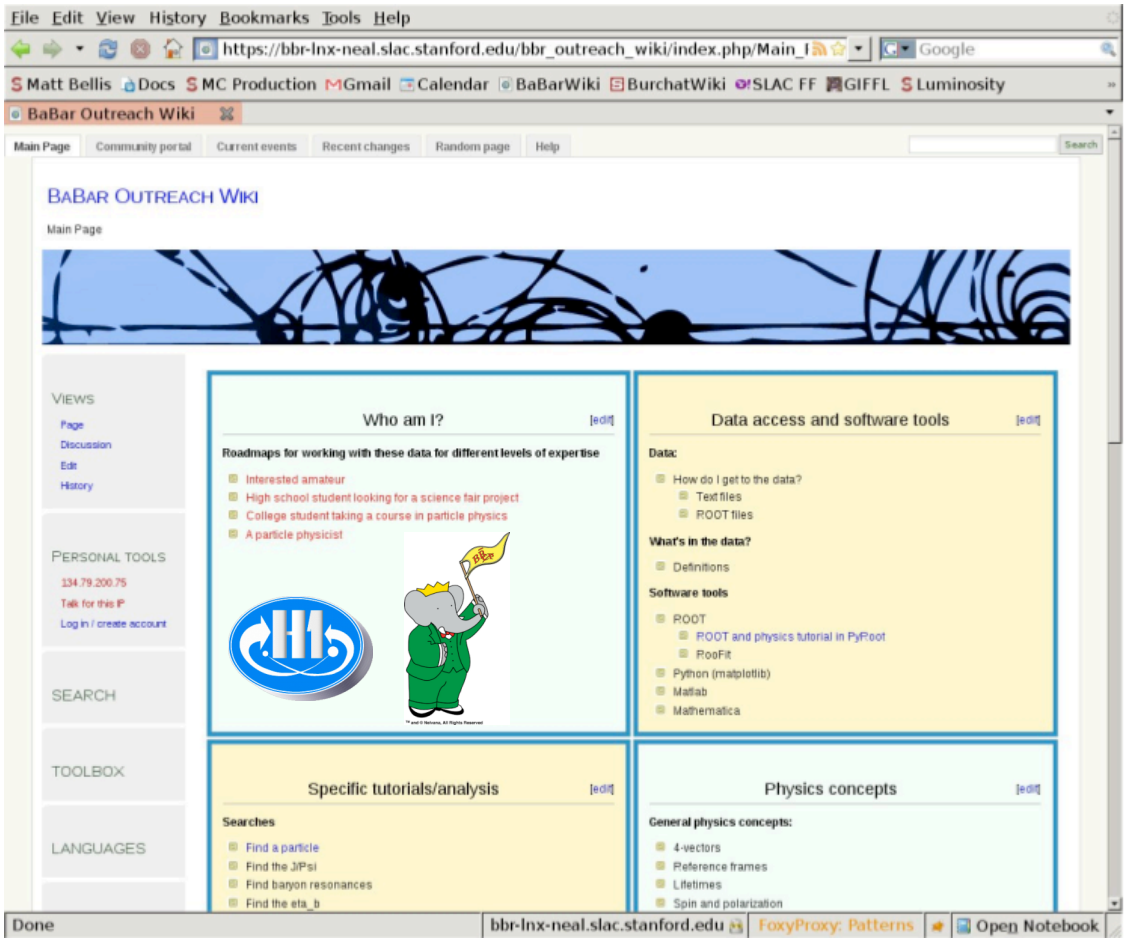

Matt Bellis (SLAC)

Full Survey of H1 Hardware Performed

- Try to retire oldest machines (least efficient, slowest, lowest spec..)
- Oldest farm machines from 2006 removed from batch and re-used
- Move the H1 web server to DESY-IT virtual SLD 5 environment?
- Current age of the H1 hardware
  - H1 batch farm (104) : oldest machines from 2007
  - H1 contributions to dCache (39) : oldest machines from 2006
  - h1wgs, general purpose (85) : oldest machines from Nov/Dec 2005 (except 2*)
- RZ totals for 2009/2010
  - New machines: 63
  - Retirements: 141
  - Total H1 machines: 228

  46% of older machines retired
  or 25% reduction in total
- The desktops are trickier
  - 140 in total, certainly not all needed and all cannot use SLD5

New H1 Hardware

- Newest H1 CPUs: DELL M610 blades
  - One 16 blade enclosure gives 256 cores, with 2*4*2=16 cores, and each with 48 GB RAM (3 GB per farm slot) and 600 GB disk per blade
  - Current farm capacity 868 cores across 104 machines, all now 3 years old or less
  - Integral part of analysis at H1, most users run parallel analysis jobs and does some MC production

- Newest H1 storage: DELL PowerEdge R710 servers
  - Capacity: 24TB (18TB usable with RAID6)
  - Just to have a number: 40 kB /event (on DST) x 2 billion/year = 75 TB

Total H1 disk storage on dCache at DESY: 250 TB + 150/2 TB

H1 Documentation Effort

Call for (old) papers

```
*****************************************
*
*      H1 non-digital  DOCUMENTATION
*
*****************************************

Dear colleagues.

As you probably already know, an effort is made to collect
and organize the older H1 Documentation, i.e. those paper
documents which are not available already in digital form.

This effort is part of the H1 Data Preservation Initiative,
and Eberhard Wuensch and myself are volunteering the work
connected with this special part of the H1 DP Initiative.

The H1 Documentation Room (Room 338, Lab. 1c) is the main
storage room for such paper documentation. It is already
quite full, and Eberhard and I are finding also other
places for this storage.

The final goal is a permanent H1 Documentation Room for the
longtime saving of the material, maybe located within the
new DESY Library. Of course, the scanning of old and important
documentation is foreseen, so that eventually it will also
be available for everyone in electronic form.

Presently quite a lot of office cleaning is going on in
Lab.building 1, due to the moving of people from their
old offices to new and temporary offices during the ongoing
renovation of the building. In this relocation process a lot
of old documentation is usually uncovered from the tops and
bottoms of the shelfs and cupboards, sometimes material which
has not been touched for a decennium or more.

*****************************************************

     Please, DO NOT THROW AWAY such material!

     Talk first with Eberhard and me, so that
     we can judge whether this material should
     be moved to rooom 338 (or any other place
     which we would recommend).

*****************************************************
```

- H1 paper documentation effort started, more and more appearing in the documentation room
  - Now to the cataloguing and organisation
  - Jan Olsson (our expert from JADE) contributing
  - Lots of things moving due to building 1 renovation
  - Digitisation, including old theses, being considered
- H1 digital documentation also examined
  - Old online shift tools may be particularly vulnerable to losses, mostly not updated since July 2007
  - Electronic logbooks: H1, trigger, components, detailed run information
  - Calibration files on old hardware, was it all rescued from the North Hall / can it be rescued now?
  - Web-pages: Streamline the content, try to rescue the dead links, increase performance
- This is no small work…

BaBar task force could help us?

ZEUS Documentation Efforts

- Non-digital documentation: notes, transparencies, technical drawings
  - Storage of collaboration talks from pre-web days already in good shape
  - Consolidation, creation of electronic catalogue, handing over custody to DESY library
- Digital documentation: mostly resides on the main ZEUS web server
  - Specific technical documentation (detectors, trigger), electronic log book distributed over several machines
  - Consolidation ongoing

Project Between H1 and INSPIRE

- Start test project with INPSIRE to host H1 paper histories
  - INSPIRE beta launched: http://inspirebeta.net/
  - In discussions with Zaven Akopov, DESY/INSPIRE, nice collaboration
  - We try extreme H1 example (Isolated leptons: 12 preliminaries!)

INSPIRE structure for publication history

- The preliminary reports will have each their own record, since they have information associated with them (varying figures, varying abstract, mostly varying presentation)
- Another record will be dedicated to the T0 stage (pre-T0, T0 and possibly addendum)
- Each drafting stage has it's own record (1st, 2nd, …) with corresponding figures and answers to draft
- Referee's report – a presentation of paper with summary of changes done to reflect the comments made by the collaboration
- Final version – the one that will be directly linked to the published paper (and is probably identical with it, unless there have been revisions submitted afterwards). In case of the mentioned revisions, they should be listed in this record as well.

  - Some other test ideas H1 notes, CB or other meetings..?

INSPIRE Record for an H1 Paper

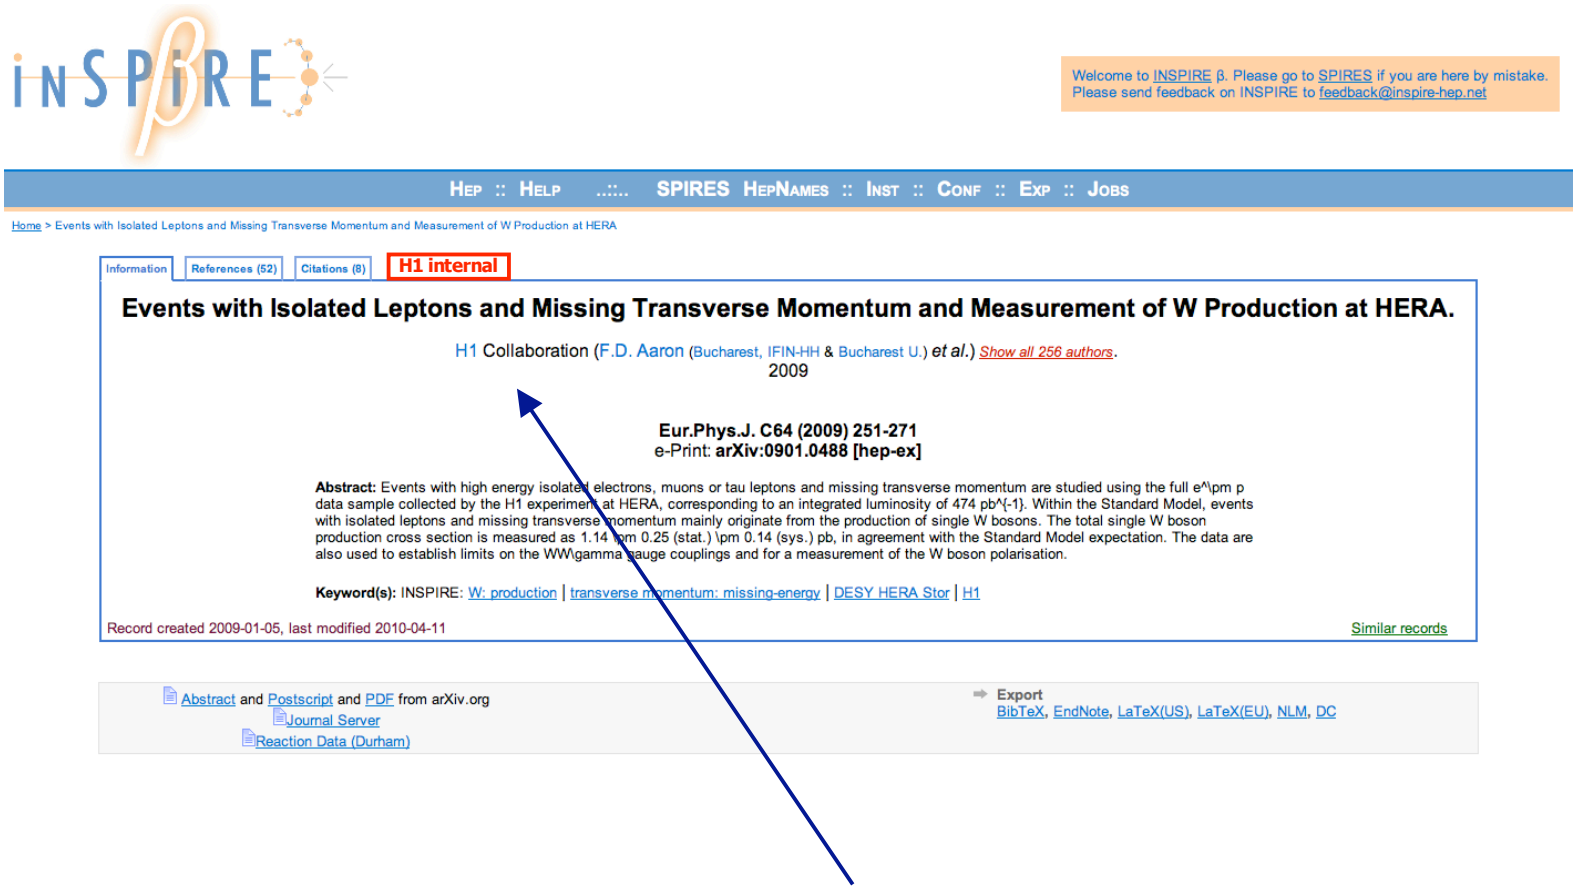

Envisage an additional link for H1 members

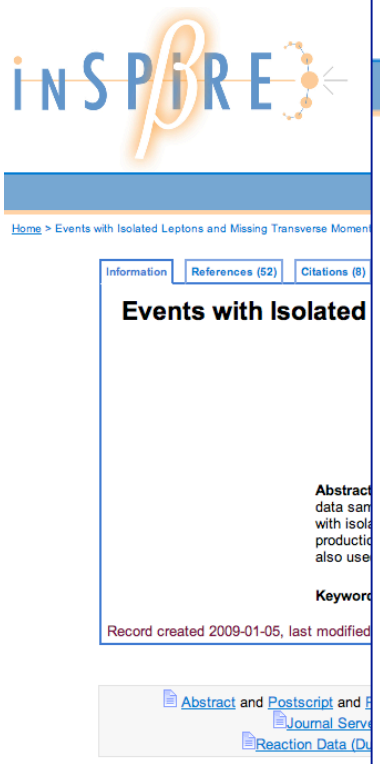

Welcome to INSPIRE ?. Please go to SPIRES if you are here by mistake.
Please send feedback on INSPIRE to feedback@inspire-hep.net

HEP :: HELP ..... SPIRES HEPNAMES :: INST :: CONF :: EXP :: JOBS

Home > > Search Results

Events with Isolated Leptons and Missing Transverse Momentum and Measurement of W Production at HERA

PUBLICATION HISTORY

Preliminary Results

HEP-EPS 2007 conference paper. | July 2007
Prepared for Deep Inelastic Scattering 2007 | April 2007
Prepared for 42nd Rencontres de Moriond (Electroweak) | January 2007
Prepared for the 62nd DESY PRC | October 2006
ICHEP 2006 conference paper |July 2006
Prepared for the 60th DESY PRC | November 2005
HEP-EPS 2005 conference paper | July 2005
Lepton Photon 2005 conference paper | June 2005
Prepared for Deep Inelastic Scattering 2005 | April 2005
Prepared for the 58th DESY PRC |October 2004
Analysis of High Pt HERA II Data | ICHEP 2004 conference paper | August 2004
High Pt Analysis of the HERA II Data |Prepared for Deep Inelastic Scattering 2004 |April 2004

T0 talks

Pre-T0 Talk | 08.02.2008
T0 Talk | 24.07.2008
T0 Addendum | 14.08.2008

Paper Drafts

First Draft | Answers to Draft |15.08.2008
Second Draft | Answers to Draft | 19.11.2008
Referee Report | 20.11.2008
Final Version | 06.01.2009

Enhanced Presentation of H1 Results

New Publication of the H1 Collaboration

DESY-09-135 **Prompt Photons in Photoproduction at HERA**

arxiv:0910.5631
H1-187

Reference	H1 Collab., F.D. Aaron et al., Submitted to EPJC (10/09) , 10/09
Figures	(1a) (1b) (1c) (1d) (2) (3) (4) (5) (6) (7) (8) (9) (10)
Links	back to overview Abstract from hep-ex Spires pdf version
Comments	

Gallery

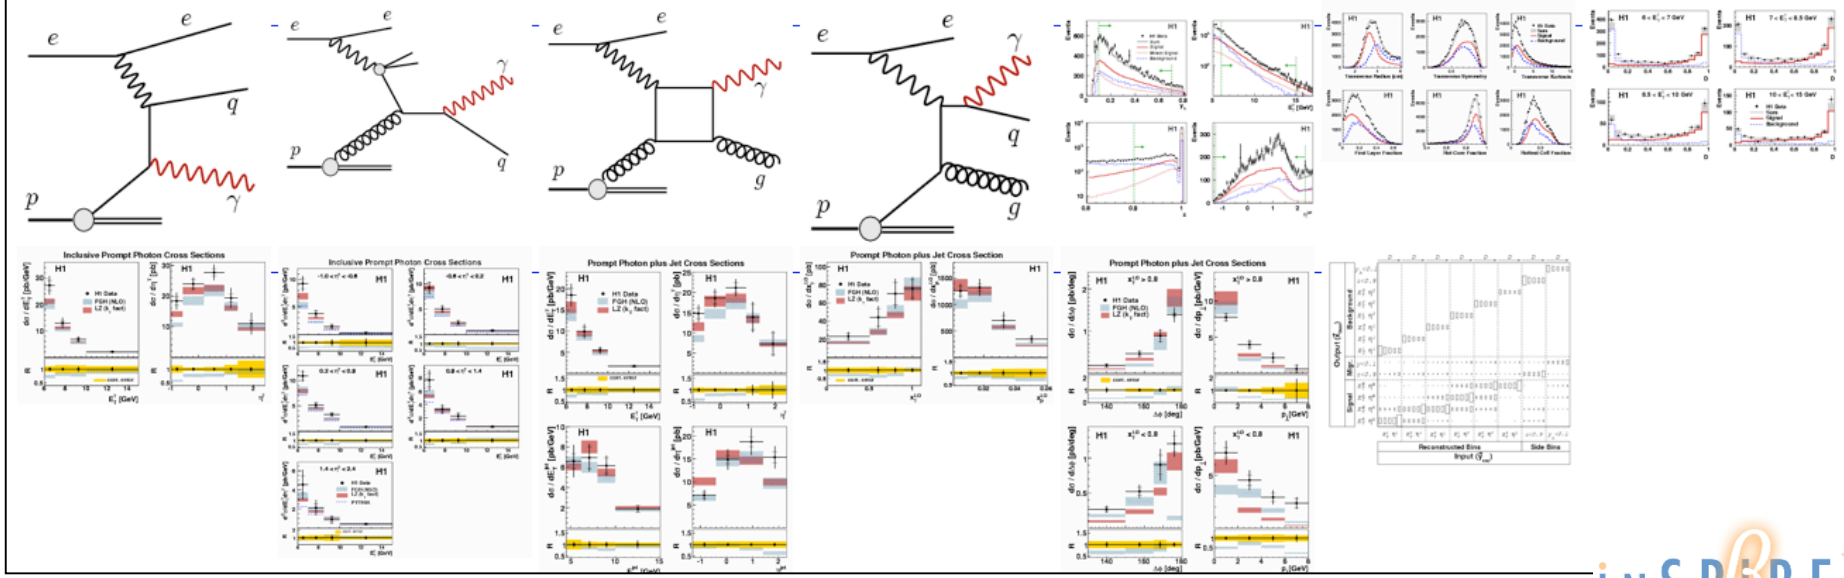

Inspired by / inspiring INSPIRE

Future Governance of the Collaborations

1. **Define the end of the H1 Collaboration in the present form**
   - 06/2013 defined by the common funding
2. **Adopt long term organization scheme:**
   - H1 Physics Committee overview the operations: nominations and appointment by September 2010
   - Infrastructure: DESY responsibilities to be defined
   - H1 Emeritus Collaboration is linked to the usage of data
3. **Data access should remain possible in the next period**
   - Consolidate the analysis environment and hardware in the present configuration
   - One more technological step (SL6+new h/w) near end 2012
   - A technical proposal is in progress (target 2013-2023, to be defined)
4. **H1 plays a major role in DPHEP**
   - Extra resources for the extension of the data access to be identified outside the internal computing effort
   - Close connection with DESY/IT, ZEUS, HERMES

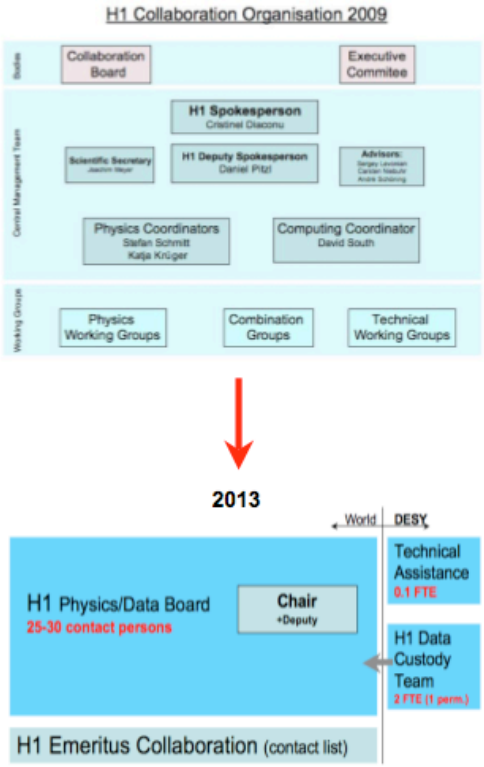

- The discussion on the various models of the organisation of ZEUS after 2013 has started in Collaboration Meeting in Kiev in October 2009
  - Due to change of management of the ZEUS collaboration postponed until the coming collaboration meeting in autumn 2010
- The model should not be much different, if not the same, as H1

H1 Data Preservation: Areas of Interest

- Try now to package the H1 project into manpower requirements
  - To be written up formally now, from existing documents: PRC, DPHEP, (KEK),..
  - This is just an initial collection of different areas we've examined into groups
  - The formal proposal may differ, and will certainly go into more detail

Documentation Master (0.5 FTE) *1 year*
Non-digital documentation
Web pages and other digital
Includes some database dependencies

Virtualisation / Validation Projects (1.0 FTE) *1-2 years*
Project with DESY-IT for validation suite
Stand alone effort with CERN VM
SLD5(6) migrations
External dependencies
Web server migrations and virtual environments

Outreach *(parallel project..)*
Would be nice, but needs some thought about how to proceed

Project with INSPIRE *(parallel project..)*
Highly attractive, and important even, promising ideas

H1 Hardware Supervision
GRID infrastructure and interfaces
Batch farm maintenance
H1 desktops, h1wgs (user level)
Digital preservation, hardware archaeology

(1.0-2.0 FTE)
2-3 years

H1 Software Supervision
H1OO maintenance + validation
Fortran (more validation)
GRID, MC production
Data and MC sets for preservation

Summary

- The $e^{\pm}p$ collisions collected at HERA are a unique data set!
  - Physics motivation detailed, full flexibility desirable, we know the arguments by now for data preservation
- Data preservation effort at DESY now more unified between the different contributors , project starting to emerge
  - Task force set up, survey of the relevant preservation issues
  - Data and data formats; technologies and resources; documentation
  - Clean up, validation, software migrations, hardware: safeguarding the data
- Isolate projects from the survey and start to attribute cost (FTEs)
  - Individual projects, those with DESY-IT (virtualisation/validation), the DESY Library (INSPIRE), as well as potential global initiatives (outreach) via DPHEP
- Written proposals to follow this workshop

DPHEP Seminars

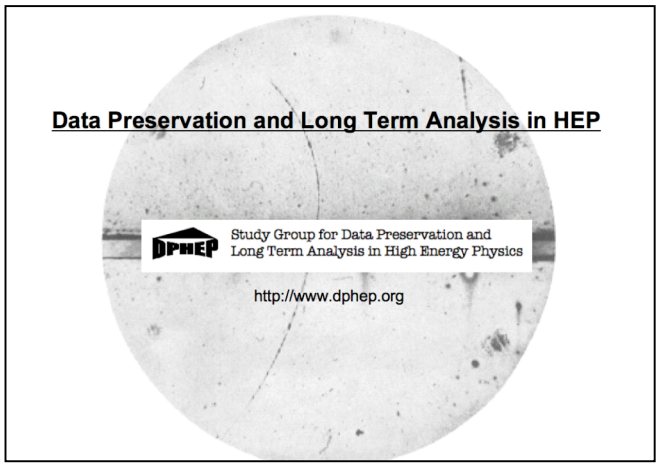

DESY

technische universität dortmund

universitätbonn

IN2P3

INSTITUT NATIONAL DE PHYSIQUE NUCLÉAIRE ET DE PHYSIQUE DES PARTICULES

BROOKHAVEN NATIONAL LABORATORY

"Very interesting, good to know someone is thinking about this"

"I'm not sure you will be able to do level 4, it seems like quite a task" [people mostly won over a little with arguments about validation procedures]

"No hope already for ATLAS software to be aligned and unified as you suggest. Oh dear"

"You should rather turn the argument round and say *BECAUSE* it costs so little in FTE with respect to initial outlay, that it would be wrong *NOT* to do it"