# Data Preservation and Long Term Accessibility in High Energy Physics

# Why is the Preservation of Large Databases an Issue for Particle Physics

- Most major High Energy Physics projects involve either making precision measurements or detection of tiny signals from **huge** samples of data.
- The experiments (accelerators and detectors) are difficult to build and the number of analyses are numerous. Furthermore, the analyses usually benefit from the **full data sample** and **complex**, **computing intensive** data reduction techniques.
- Result: Analysis of the huge amount of data continues many years after data collection.

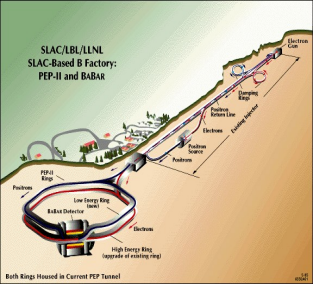


# An Elephant among Databases

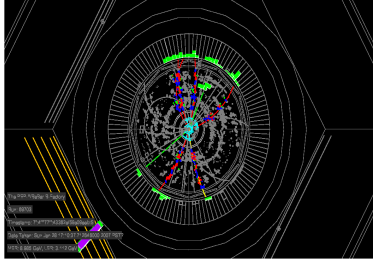

## Press Release: World's Largest Database reaches 500,000 Gigabytes

http://home.slac.stanford.edu/pressreleases/2002/20020412.htm

Date Issued: April 12, 2002

Contact:Neil Calder, Stanford Linear Accelerator Center

Last week at the Stanford Linear Accelerator Center (SLAC), the BABAR experiment's database stored its 500,000th Gigabyte - a milestone that makes it the largest known database in the world. The BABAR experiment - a collaboration of 600 physicists from nine nations - observes collisions between subatomic particles to understand how the behavior of matter and antimatter shaped our universe. BABAR, also known as the "B Factory," mass-produces huge quantities of scientific data with industrial efficiency. Up to 500 Gigabytes of data is sent relentlessly to the experiment's database daily.

The half million Gigabytes of data in the BABAR database, printed out, would fill one billion books. That's nearly 60 times the number of books in the Library of Congress, the largest library in the world. "The need to store the avalanche of information coming from the experiment and then efficiently search and retrieve specific data samples has driven physicists and computer experts to create innovative technology," said SLAC Director Jonathan Dorfan. "Governments, commercial corporations and institutes will face similar needs in the near future and the knowledge and experience we have gained will be passed on."

http://www.slac.stanford.edu/BFROOT/www/Public/Computing/Databases/proceedings/CIDR05.pdf

**Lessons Learned from Managing a Petabyte**

Jacek Becla*
Stanford Linear Accelerator Center
2575 Sand Hill Road, M/S 97
Menlo Park, CA 94025, USA
becla@slac.stanford.edu

Daniel L. Wang*
Stanford Linear Accelerator Center
2575 Sand Hill Road, M/S 97
Menlo Park, CA 94025, USA
danielw@slac.stanford.edu

BaBar's vibrant research effort continues to demand more data, more quickly from the detector, as well as higher levels of data service. The first generation eventstore was undoubtedly a great success, providing storage and service throughput well beyond its original design goals. Data rates, for example, were several times higher than originally designed. Hundreds of users analyzed data in BaBar. Its complexity and size has put it beyond today's scalability frontier: in 2003, it was larger than the largest 200 relational databases combined, earning the grand prize in Winter Corporation's TopTen Program (a survey of world's largest databases).

Further details on this ODBMS-implementation can be found in [1], [2], and [3].

# More than Elephants Roaming at SLAC these days

- At SLAC the management and preservation of large database is no longer just a BaBar issues.
- We (SLAC/NASA/...) have a particle detector in space. **FERMI** is investigating building an archival system.
- **LSST** will produce a gargantuan amount of data and no means of doing the archiving has been conceived yet.
- SLAC is part of **ATLAS** which is has now collected … Pbytes of data
- SLAC is providing machinery for the very high luminosity **SuperB** project next door in Tor Vergata

# Other Reasons for Preservation

- We may want to re-do previous measurements
  - Increased precision, reduced systematics
  - New and improved theoretical calculations / MC models
  - Newly developed analysis techniques
- We may want to perform new measurements
  - At energies and processes where no other data are available (or will become available in the future)
  - Particularly relevant to HERA $e^{\pm}p$ data (and also Tevatron)
- Investigate if new phenomena found today
  - Go back and check in the old data

# The JADE Example

End of Data Taking,
3. Nov. 1986, 05:45

**Most of the HEP world didn't even know that the JADE data had been preserved yet > 20 years after data collection a new significant result was published using resurrected JADE data and analysis framework.**

## Successful resurrection of JADE data analysis

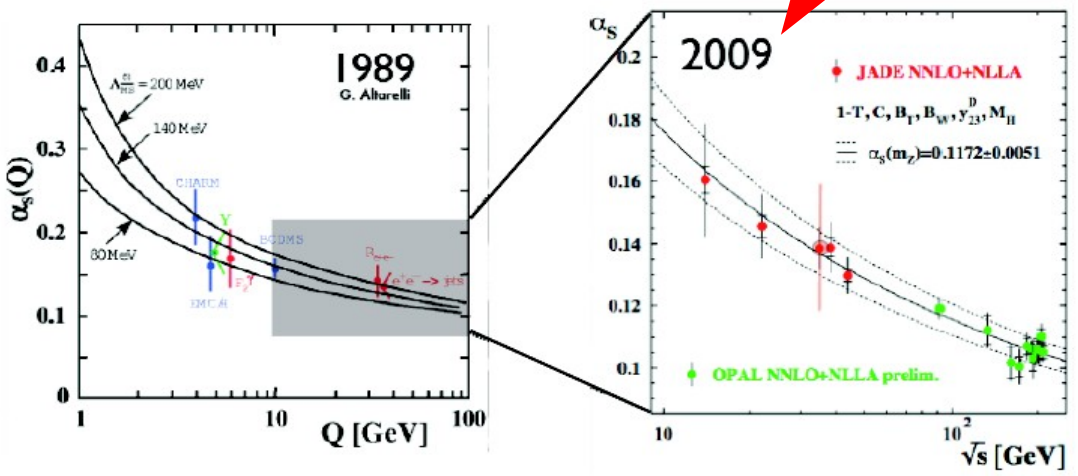


- New analyses required full raw data preservation, software revitalisation, needed many individual initiatives…
- 6 JADE publications since 1997

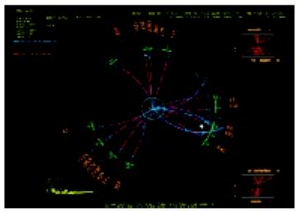

*arXiv:0810.1389*
*arXiv:0707.0392*
*hep-ex/0106066*
*hep-ex/0001055*
*hep-ex/9903009*
*hep-ex/9708034*

## JADE data and software

- 1995: „private" (*neither* collaboration *nor* lab) initiatives to :
  - rescue data from original archive tapes and copy them onto more modern media (IBM cartridges & Exabyte) (J. Olsson @ DESY)
  - reanalyse data using modern (LEP-like) methods and observables plus improved theoretical calculations (S. B. and P. Movilla-Fernandez @ RWTH Aachen
  - revitalise JADE software on modern computer platforms to enable generation of new MC data files (P. Movilla Fernandez, J. Olsson)
- so far, the only example of reviving and still using 25-30 year old data & software in HEP
- since 1996, O(10) publications, O(10) conf. contributions; no competition in $e^+e^-$ data analysis at $E_{cm}$ ~14 ... 200 GeV

Re-Analysis of PETRA Data | Data Persistency Workshop, DESY | January 26-28, 2009 | S. Bethke MPP Munich | 8

# CMS' Data Curation Plan

From Elizabeth Sexton-Kennedy

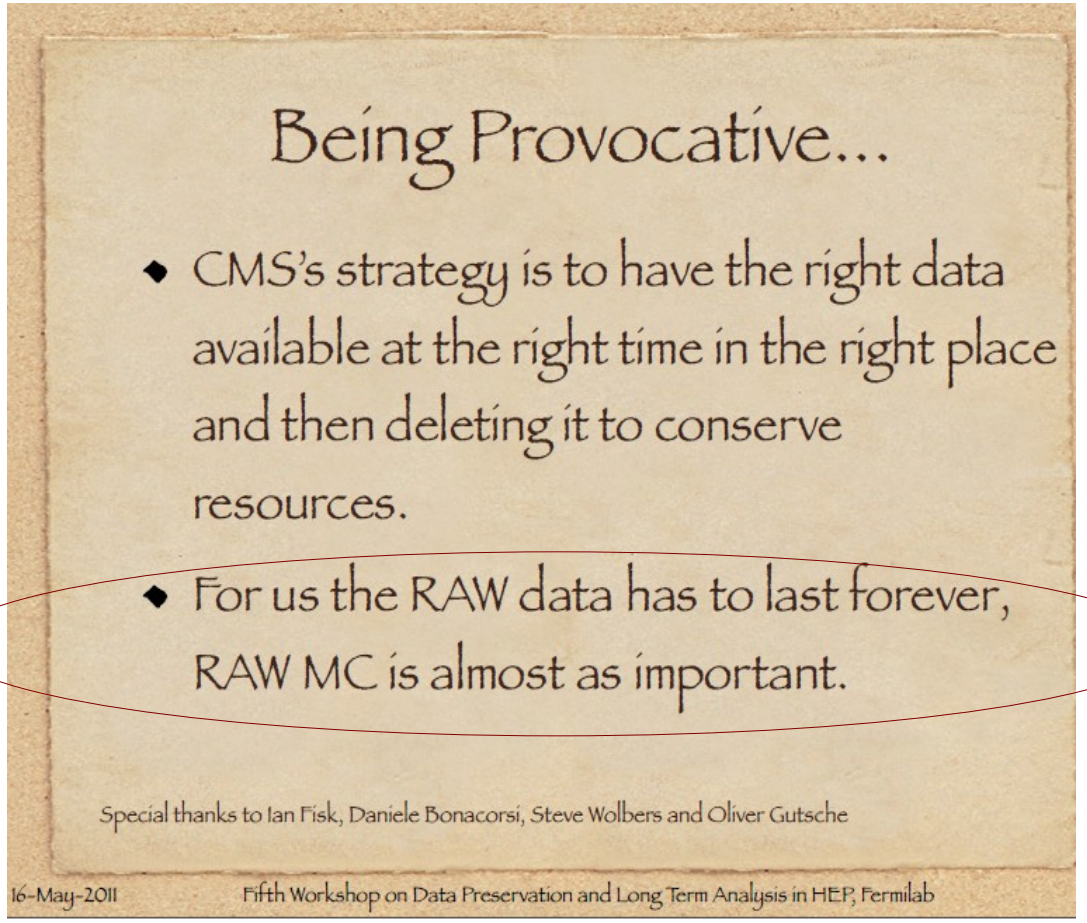

# Ill Prepared Projects Face a tough Challenge

- Pressure to **reduce resource usage** on experiments that have completed data taking so as to free up sources for new projects
- The **computing infrastructure (Operating systems, hardware, remote resources, etc...) change** while the computing expertise is drained
  - Access to expert support reduces
- Meanwhile the analyses are finally in the stage of having the maximum data and are frequently doing more complicated fits than ever before!

# The resources begin to disappear just when you need them

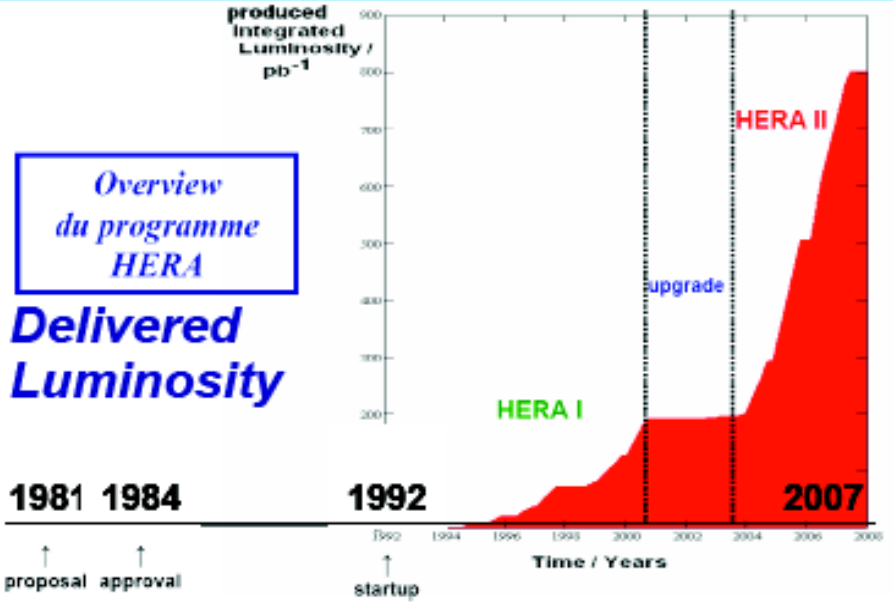


- Good data taking period is towards the end of running
- The existing resources (funding and expertise) then decrease when the data taking stops
- Dedicated resources for preservation need to be planned early!

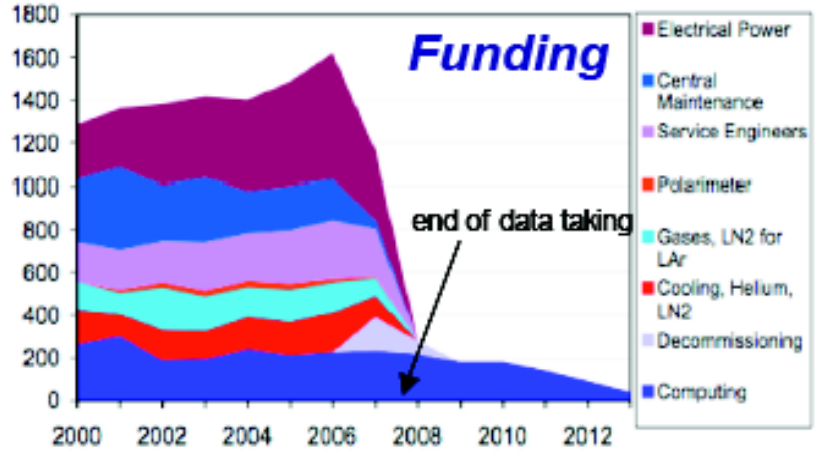


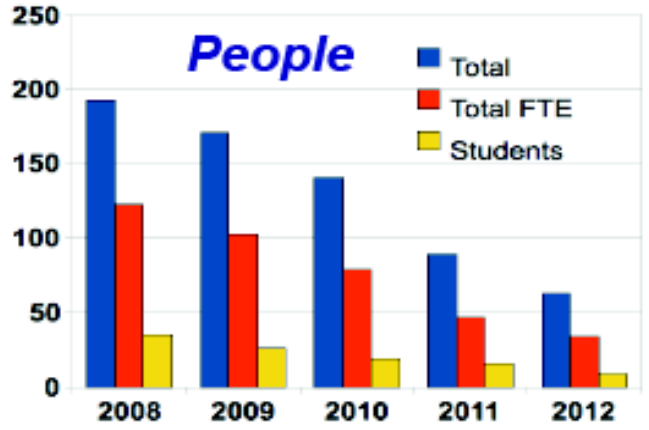

# The importance of preservation

## Future independent checks

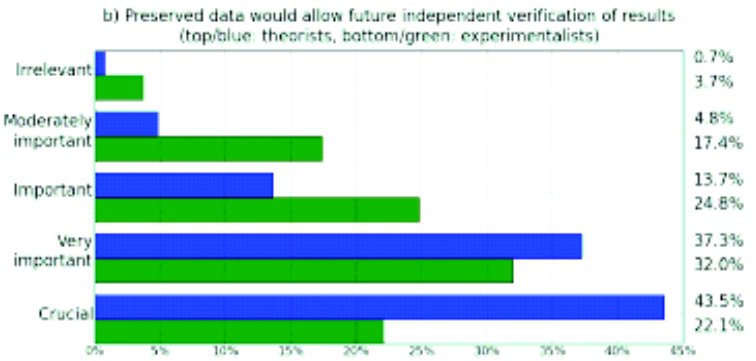


## Combine with future data

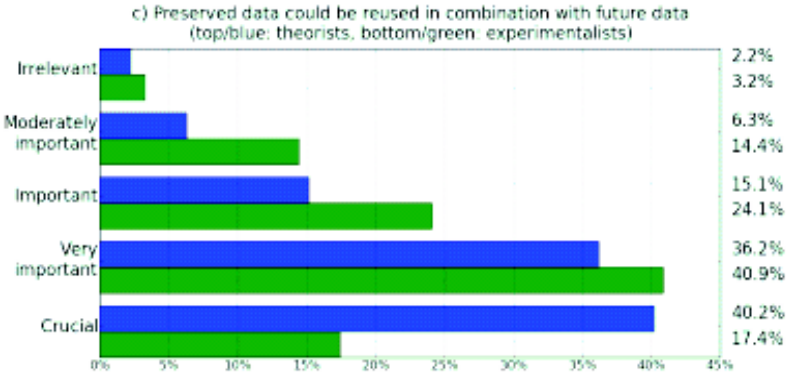


## Re-analyse for future theories

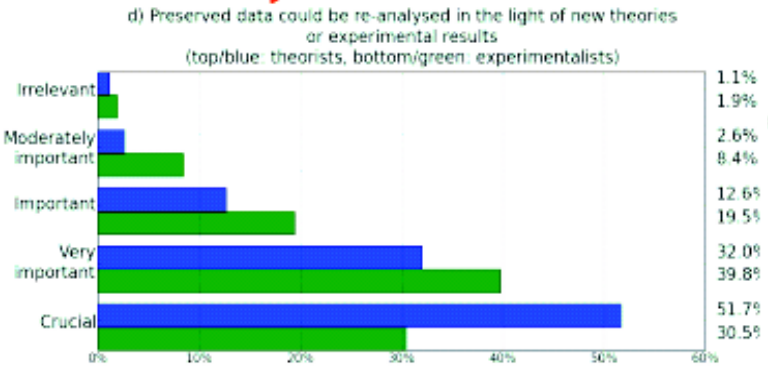


## Teaching and outreach

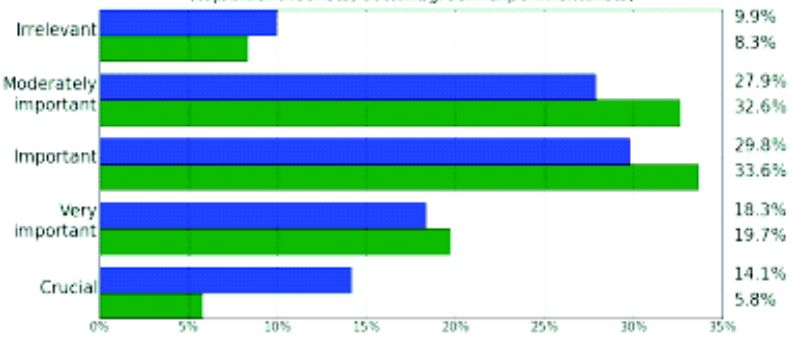

# A Problem of Magnitudes

- BaBar raw data output rate 530 KB/s - 2Pbytes to preserve
- ATLAS raw data output rate 300 **MB/s**

  - 15Pbytes per year!!!

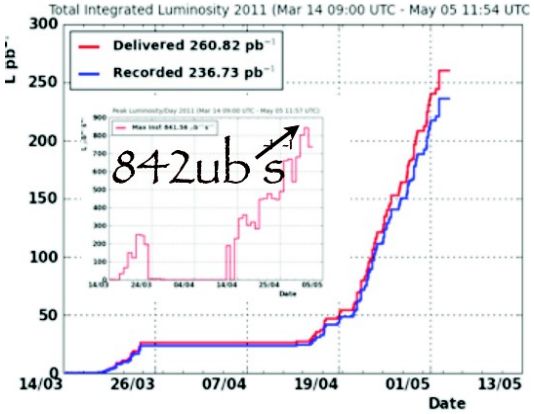


From Elizabeth Sexton-Kennedy

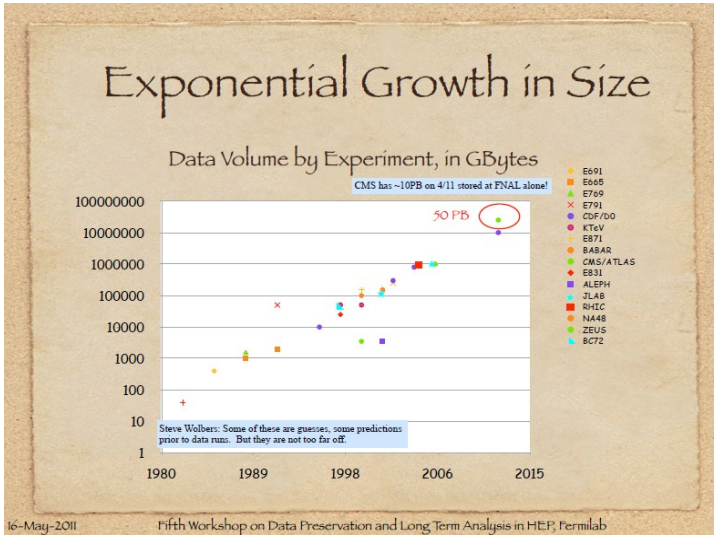


- LSST calibrated data output rate 60000 MB/s every 40 seconds
- LHC and LSST data will be unique and contain rich new physics that will take decades fully exploit
- The steps beyond LHC and SuperB are too far in the future for anyone to know when they might occur; their data will be analyzed for a very very long time.

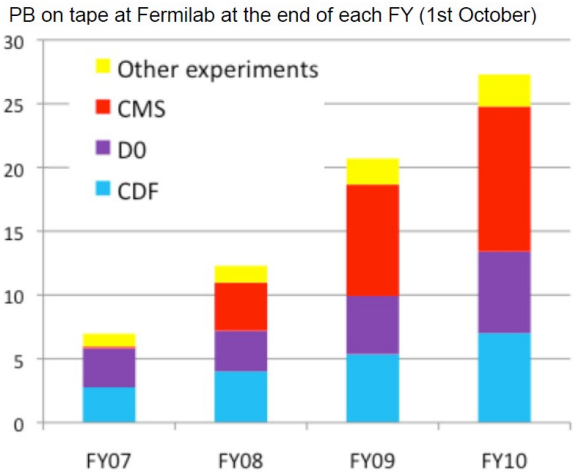

# Existing and Future HEP $e^+e^-$ Datasets

| | BaBar | Belle | BES-III | CLEO |
|---|---|---|---|---|
| End of Data Taking | 07/04/08 | ~2010 | ~2017 | 01/04/08 |
| Collaboration end date | end of 2012 | end of 2012 | 2017-2022 | |
| Type of data to be preserved | raw + sim/recon (ROOT) | raw + MDST | raw+DST (ROOT) | OBJY/PDS (too dificult) preserve analysis data |
| quantity | 2 Pbytes | ~4 Pbytes | ~6 Pbytes | |
| desired longevity of long term analysis | unlimited | 5 years (until super KEKB) | 15 years | superseded by B-Factories and BES-III |
| Simulation | Geant4 | Geant 3 | Geant4 | Geant3 |
| Platform | SL3,4,5 | | SLC4 | |
| code | C++ | Fortran | C++ | Fortran, C, C++ |

**Super-B**

An «e+e– collider that produces a nearly **two-order-of-magnitude increase** in luminosity over the current generation of asymmetric B Factories »


**CERN Courier**

Jul 19, 2011

## SuperB Factory set to be built at the University of Rome 'Tor Vergata'

Roberto Petronzio, president of INFN has announced that the SuperB Factory, will be built at the University of Rome 'Tor Vergata'. The facility tops the list of 14 flagship projects of the National Research Plan of the Italian Ministry for Education, Universities and Research ( CERN Courier January/February 2011 p8).

The SuperB project involves the construction underground of a new asymmetric high-luminosity electron–positron collider. It will occupy approximately 30 hectares on the campus of the University of Rome 'Tor Vergata' and be closely linked to the INFN Frascati National Laboratories, located nearby.

# An Issue for Many Fields

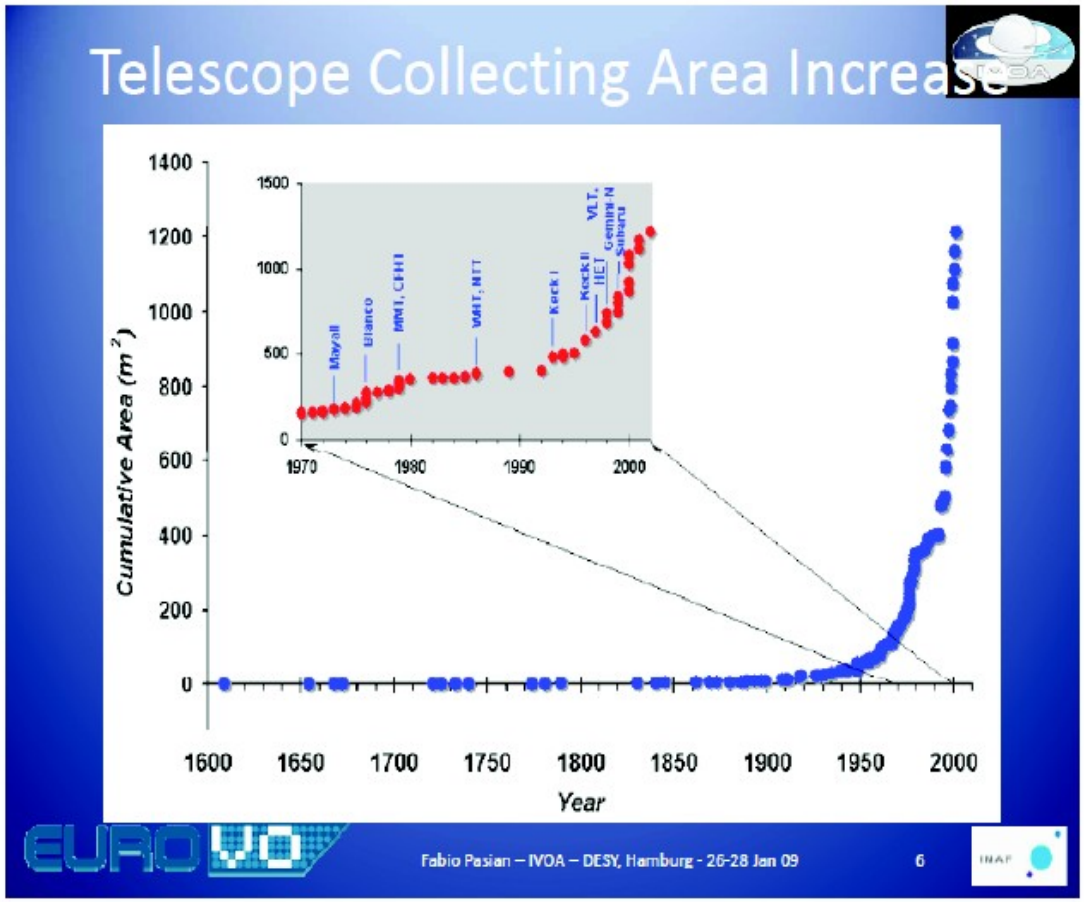


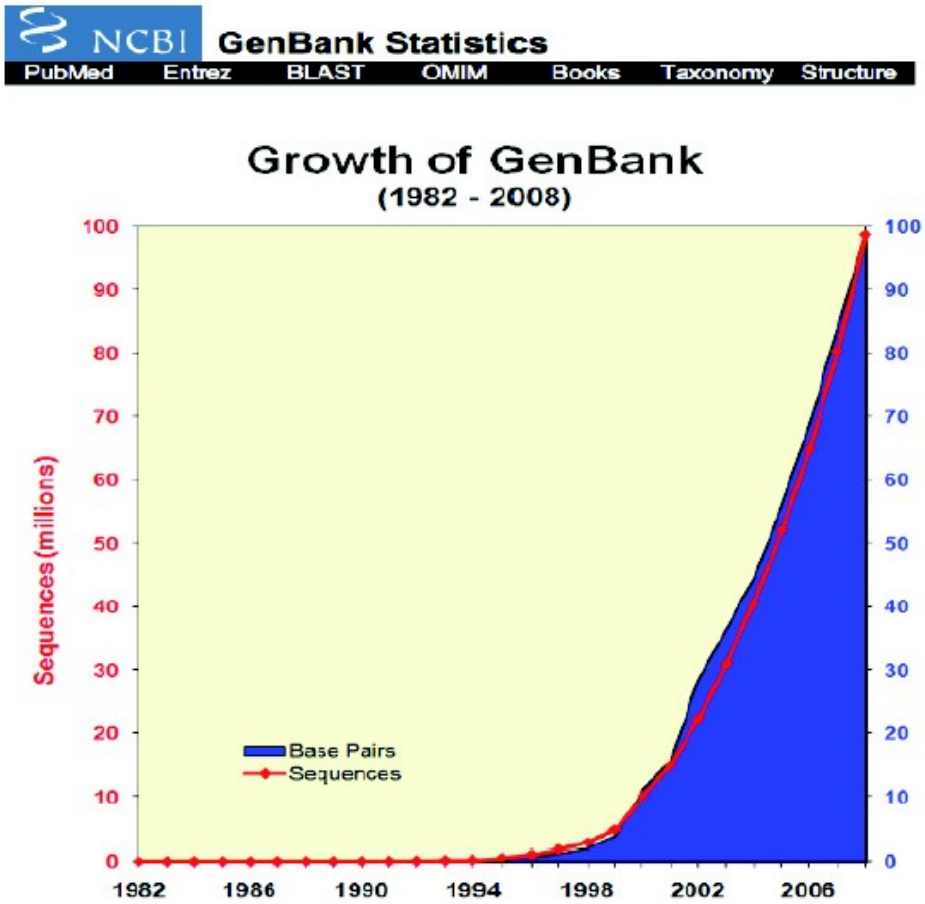


Other scientific fields are also experiencing a dramatic increase in data and are questioning the long term future of this data

> Task forces already in place to address this issue in a generic way, defining standards, e.g. Blue Ribbon, APA, DPC, eSciDIR, …

# Data Preservation: What does it mean?

**But a few communiqués suggest a common problem…**

*To Whom it may concern,*

*In the tape storage area we still have 4132 tapes of type 3840 containing HERA data.*

*We do not have a functioning reading device anymore and the storage area was polluted recently, so it is likely that the tapes are damaged.*

*Would you like us to send you these tapes or should we destroy them directly?*

*Yours Sincerely,*

*Tape admin. service [a large computing centre]*

> Some other choice quotes:

*"We cannot ensure data is stored in file formats appropriate for long term preservation.*

*"We cannot ensure those data are still usable. The software for exploiting those data is under the control of the experiments.*

*"We are sure most of the data are (not easily) accessible!"*



**This Data:**

> **Digital information:** Data event files, database
> **Software:** Simulation, reconstruction, analysis, user
> **Publications:** Journals, arXiv, Spires/INSPIRE, HEPDATA
> **Documentation:** Publications, notes, manuals, slides
> **Meta information:** Hyper-news, messages, wikis, forums
> **Expertise (people):** Often the hardest to secure

← A sign of the problem

## Looking for an answer →

> Only with the full flexibility does the full potential of the data remain
- Level 4 type programme was required by the JADE and ALEPH re-analyses

> BaBar, H1, HERMES aim for DPHEP level 4, ZEUS between levels 3 and 4
- Still some different approaches, can benefit from each other's experiences

> Levels 1 and 2 still require some work!

| 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 |

benefits and cost

# Why has data preservation become so important now

**BEFORE**:

- Experiments didn't run so long, the datasets weren't too big and improved experiments covering similar physics were always on the horizon

**NOW:**

- Enormous datasets with far ranging physics potential beyond the the data collection period now exist
- Renewed appreciation of how archived data can provide crucial input to new models
- Lengthy projects with no quick followup project that would completely replace the data

# Data Taking Ends but thePhysics Extraction Continues Long After

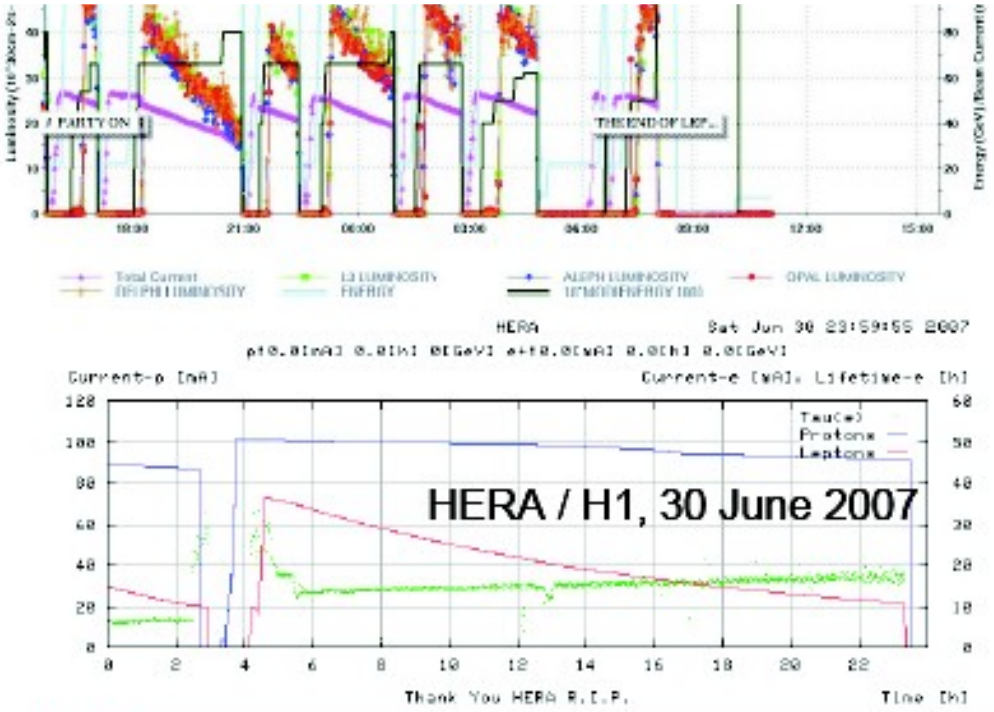


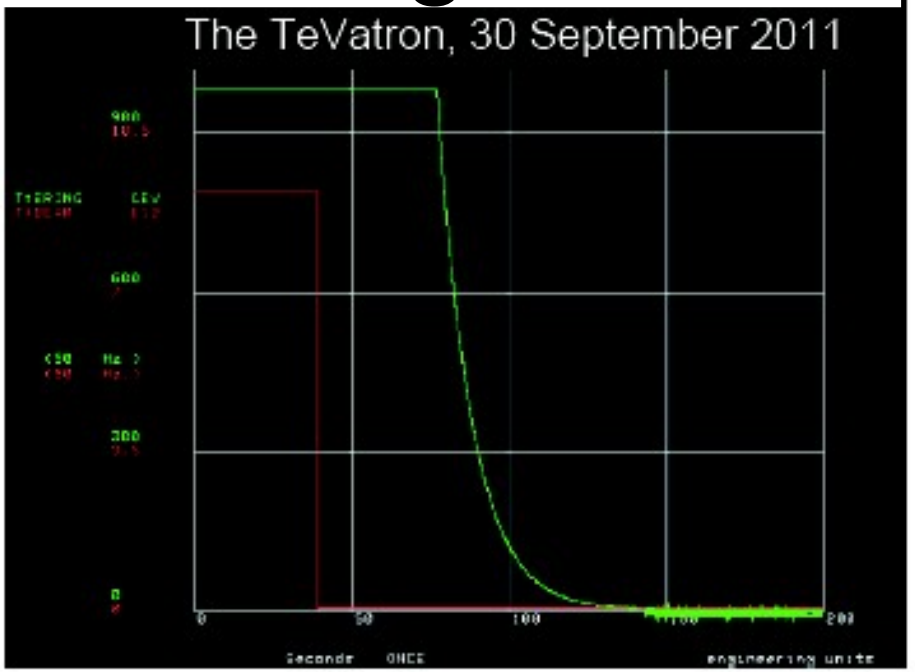


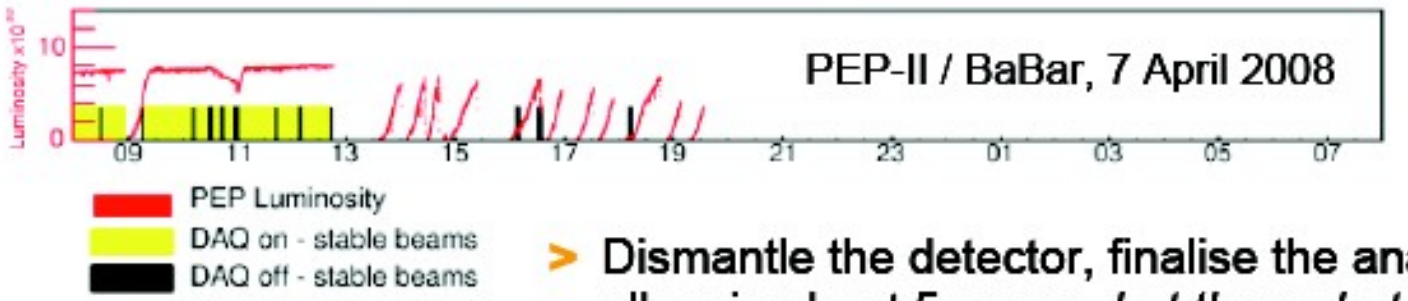


> Dismantle the detector, finalise the analyses and wrap it all up in about 5 years.. *but then what to do with the data?*

# Addressing the Concerns

nature International weekly journal of science

nature news home | news archive | specials | opinion | features | news blog | nature jour

comments on this story

Stories by subject
- Lab life
- Physics
- Policy

Stories by keywords
- Tevatron
- Data preservation
- Cristinel Diaconu

This article elsewhere
- Blogs linking to this article
- Add to Connotea
- Add to Digg

Published online 27 May 2011 | *Nature* **474**, 16-17 (2011) | doi:10.1038/474016a

News

## Tevatron's legacy set to disappear

**Lack of long-term preservation plan threatens to leave key information inaccessible for future analysis.**

Eugenie Samuel Reich

Four months before the Tevatron shuts down for good, physicists at Fermilab's giant particle collider near Batavia, Illinois, are pulling out all the stops to collect every last bit of data that they can. But some worry about what will eventually happen to the trove of data — approaching 20 petabytes ($20 \times 10^{15}$ bytes) — amassed over the machine's 26-year life.

Data from Fermilab, such as this reconstruction of two top quarks, could be lost to physicists forever. *Fermilab*

At BaBar, an experiment that produced B mesons at the Stanford Linear Accelerator Center in Palo Alto, California, until 2008, physicists have begun building a US$500,000 archival system that will save the raw data and software, and are also setting up virtual interfaces to run the older software on modern machines. "We have decided to save everything at least to 2018," says Tina Cartaro, BaBar's computing coordinator.

Currently running and future experiments cannot wait to face the issue:

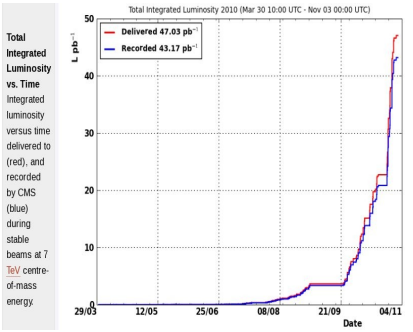


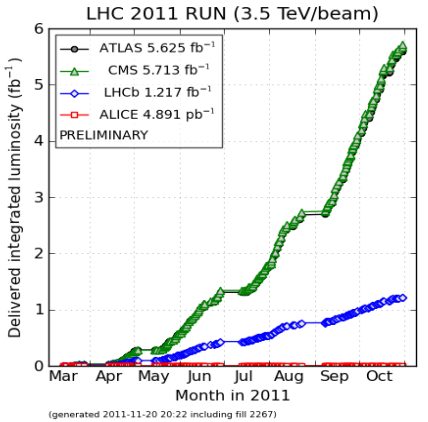


***Note:*** FNAL just made two hires to work on data preservation and we just received a very positive report on the D0 efforts last week,

Study Group for Data Preservation and Long Term Analysis in High Energy Physics

Started by C. Diaconu


- **First contacts established in September 2008**
  - Group has since grown to over 100 contact persons
  - Endorsed as an ICFA panel summer 2009
  - LHC experiments joined in 2011

- **The first task of the group was to establish the working directions**
  - Confront data models, clarify the concepts, set a common language, investigate technical aspects, compare with other fields such as astrophysics handling large data

# DPHEP Organization

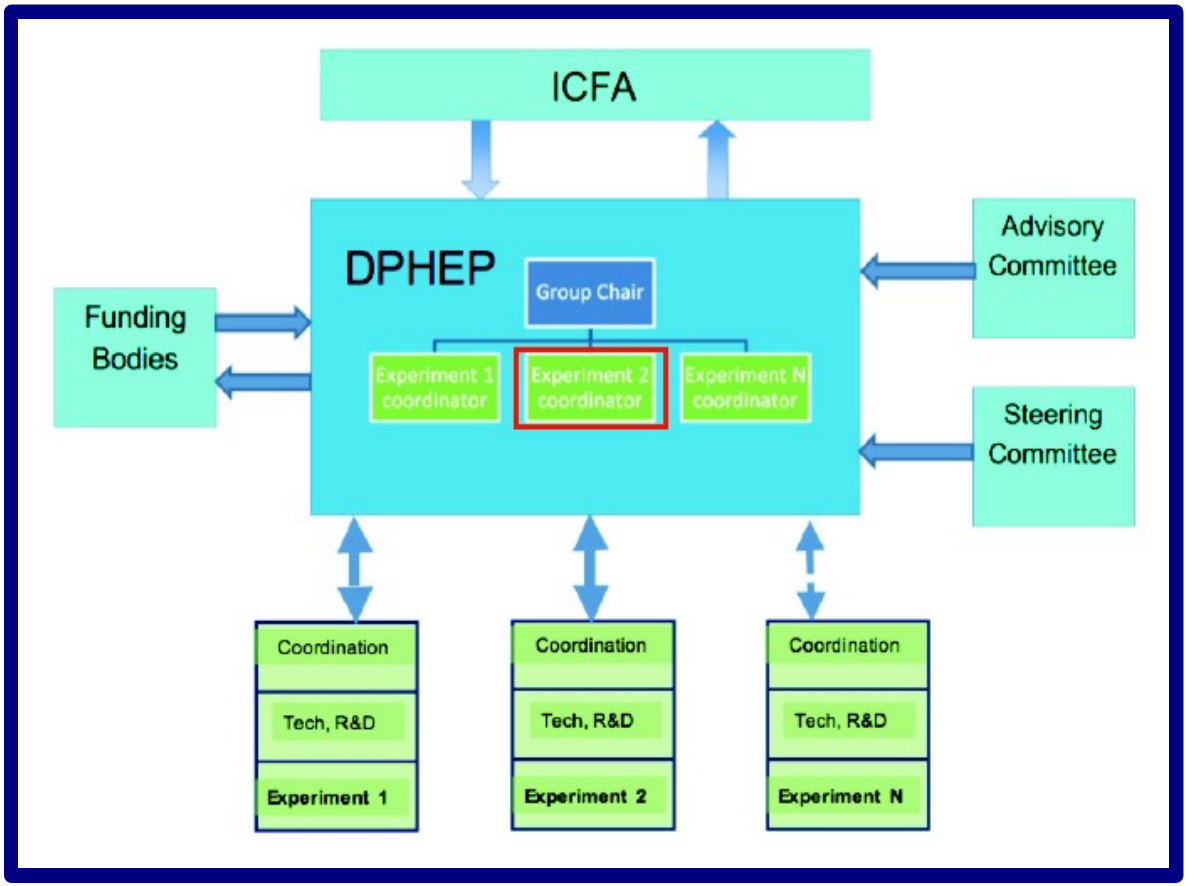


Support for DPHEP from CERN, DESY, Fermilab, IHEP, SLAC, as well as a variety of HEP committees: ICFA, HEPAP (DOE and NSF), FALC

> Contacts within DPHEP with astrophysics and other general initiatives

# Global Data Preservation Efforts


PARSE insight

Permanent Access to the Records of Science in Europe
www.parse-insight.eu

Common issues with many fields on:

- Data integration
- Technologies for data curation
- Work on standards
- (Open) access
- Policy issues
- Funding models

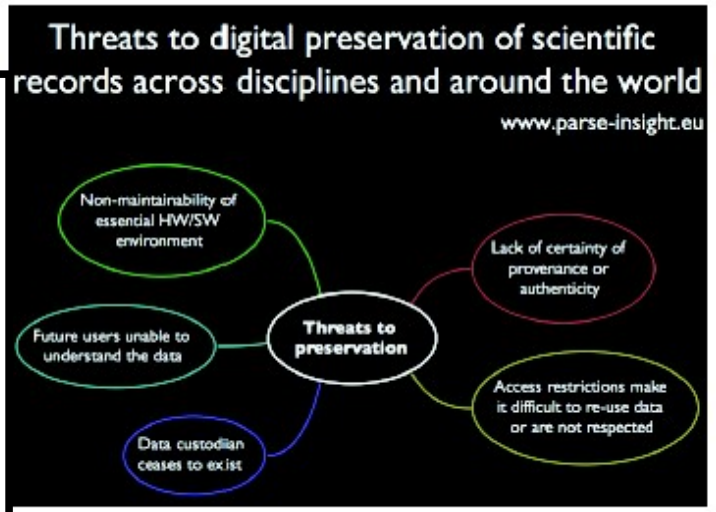




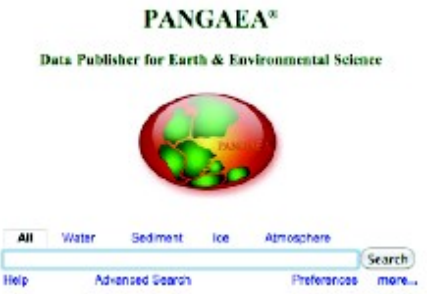


# DPHEP Members

Study Group for Data Preservation and Long Term Analysis in High Energy Physics

- Group has grown since 2008 to over 100 contact persons
- Endorsed by ICFA summer 2009
- **LHC** experiments joined in 2011

- Chair: Cristinel Diaconu (DESY/CPPM)
- Working Groups
  - Physics Cases: François Le Diberder (SLAC/LAL)
  - Preservation Models: D. South (DESY), Homer Neal (SLAC)
  - Technologies: Stephen Wolbers (FNAL), Yves Kemp (DESY)
  - Governance: Salvatore Mele (CERN)
- International Steering Committee
  - Participants from ee, ep and pp collider experiments
  - Associated computing centres at the labs
  - Some funding agencies
- International Advisory Committee
  - Chairs: Jonathan Dorfan (SLAC), Siegfried Bethke (MPIM)
  - Advisers: Gigi Rolandi (CERN), Michael Peskin (SLAC), Dominique Boutigny (IN2P3), Young-Kee Kim (FNAL), Hiroaki Aihara (IPMU/Tokyo), Alex Szalay (JHU)

# The Data Preservation in High Energy Physics ICFA Panel

Series of DPHEP workshops held since 2009

- Jan 2009: DESY May 2009: SLAC Dec 2009: CERN Jul 2010: KEK May 2011: Fermilab

Initial findings published in an interim report at the end of 2009

- Physics Case for Data Preservation
- Preservation models,
- Technologies,
- Governance

Blueprint of Data Preservation in preparation, providing recommendations and cost estimates for past, present and future HEP experiments

# NSF DataNet

Support for data preservation at the national funding level

http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503141&org=OCI&from=home

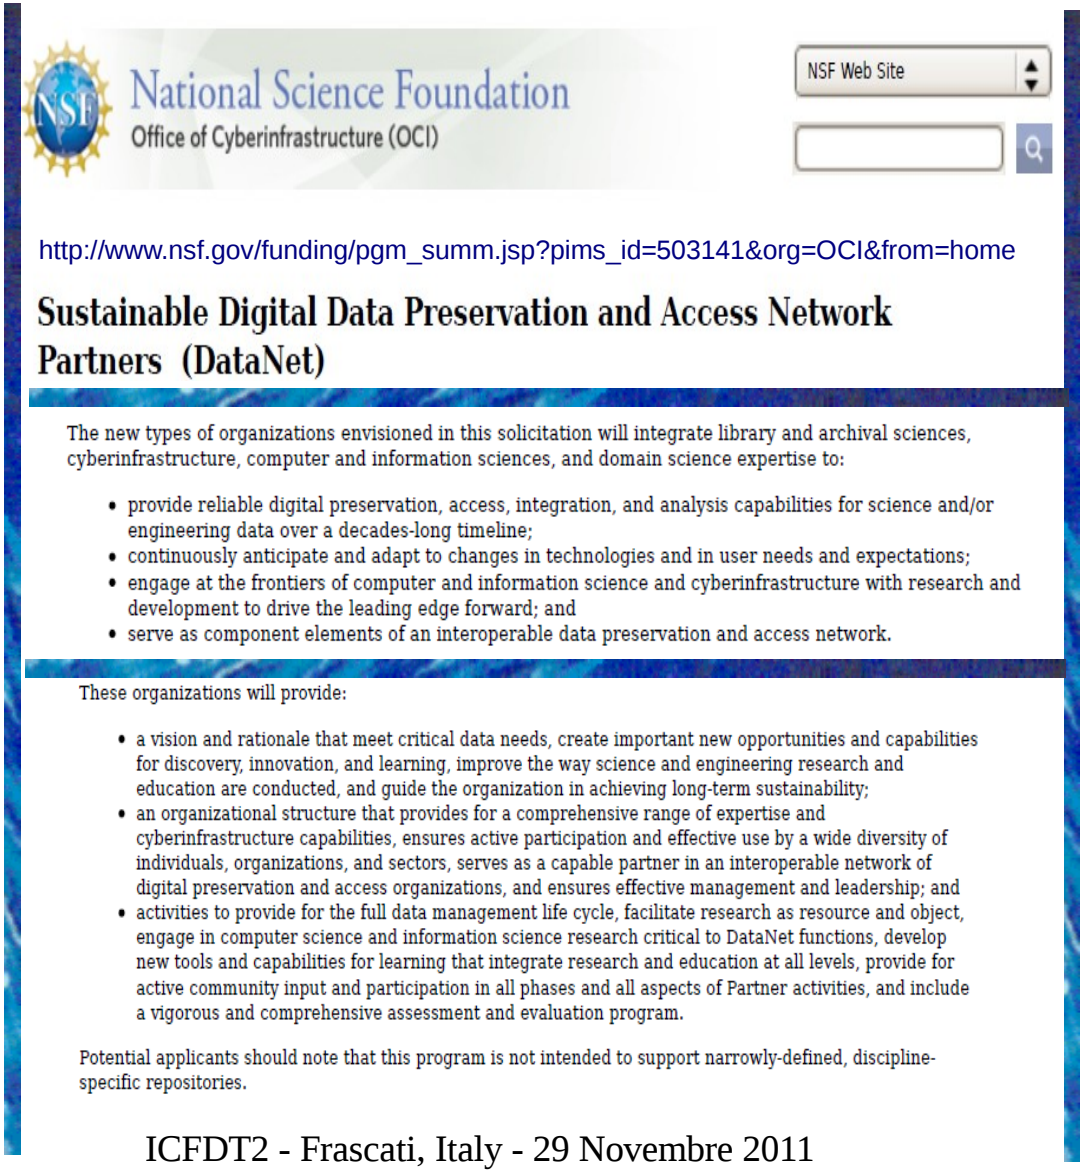

National Science Foundation
Office of Cyberinfrastructure (OCI)

NSF Web Site

http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503141&org=OCI&from=home

**Sustainable Digital Data Preservation and Access Network Partners (DataNet)**

The new types of organizations envisioned in this solicitation will integrate library and archival sciences, cyberinfrastructure, computer and information sciences, and domain science expertise to:

- provide reliable digital preservation, access, integration, and analysis capabilities for science and/or engineering data over a decades-long timeline;
- continuously anticipate and adapt to changes in technologies and in user needs and expectations;
- engage at the frontiers of computer and information science and cyberinfrastructure with research and development to drive the leading edge forward; and
- serve as component elements of an interoperable data preservation and access network.

These organizations will provide:

- a vision and rationale that meet critical data needs, create important new opportunities and capabilities for discovery, innovation, and learning, improve the way science and engineering research and education are conducted, and guide the organization in achieving long-term sustainability;
- an organizational structure that provides for a comprehensive range of expertise and cyberinfrastructure capabilities, ensures active participation and effective use by a wide diversity of individuals, organizations, and sectors, serves as a capable partner in an interoperable network of digital preservation and access organizations, and ensures effective management and leadership; and
- activities to provide for the full data management life cycle, facilitate research as resource and object, engage in computer science and information science research critical to DataNet functions, develop new tools and capabilities for learning that integrate research and education at all levels, provide for active community input and participation in all phases and all aspects of Partner activities, and include a vigorous and comprehensive assessment and evaluation program.

Potential applicants should note that this program is not intended to support narrowly-defined, discipline-specific repositories.

Year 2011 awards

- DataNet Full Proposal: Sustainable Environment through Actionable Data (SEAD)
- DataNet Full Proposal: Terra Populus: A Global Population/Environment Data Network
- DataNet Full Proposal: DataNet Federation Consortium
- MRI: Instrumentation for Enabling Data Analysis, Sharing, Storage, and Preservation
- CRCNS Data Sharing: An open data repository for cognitive neuroscience: The OpenfMRI Project
- CRCNS Data Sharing: An open data repository for cognitive neuroscience: The OpenfMRI Project
- CRCNS Data Sharing: An open data repository for cognitive neuroscience: The Open fMRI Project
- CRCNS Data Sharing: An open data repository for cognitive neuroscience: The OpenfMRI Project
- CRCNS Data Sharing: An open data repository for cognitive neuroscience: The OpenfMRI Project
- Future Career Opportunities and Educational Requirements for Digital Curation
- EAGER: Knowledge and Data Transfer: the Formation of a New Workforce
- EAGER: Field Computational Ecology Course
- Undergraduate Training at NSF Teragrid XD RDAV Center
- EAGER: Advanced Erasure Coding Technology for Storage Networks

# Required Data Preservation Plans

**Press Release 10-077**

## Scientists Seeking NSF Funding Will Soon Be Required to Submit Data Management Plans

**Government-wide emphasis on community access to data supports substantive push toward more open sharing of research data**

**May 10, 2010**

During the May 5th meeting of the National Science Board, National Science Foundation (NSF) officials announced a change in the implementation of the existing policy on sharing research data. In particular, on or around October, 2010, NSF is planning to require that all proposals include a data management plan in the form of a two-page supplementary document. The research community will be informed of the specifics of the anticipated changes and the agency's expectations for the data management plans.

The changes are designed to address trends and needs in the modern era of data-driven science.

"Science is becoming data-intensive and collaborative," noted Ed Seidel, acting assistant director for NSF's Mathematical and Physical Sciences directorate. "Researchers from numerous disciplines need to work together to attack complex problems; openly sharing data will pave the way for researchers to communicate and collaborate more effectively."

"This is the first step in what will be a more comprehensive approach to data policy," added Cora Marrett, NSF acting deputy director. "It will address the need for data from publicly-funded research to be made public."

Seidel acknowledged that each discipline has its own culture about data-sharing, and said that NSF wants to avoid a one-size-fits-all approach to the issue. But for all disciplines, the data management plans will be subject to peer review, and the new approach will allow flexibility at the directorate and division levels to tailor implementation as appropriate.

This is a change in the implementation of NSF's long-standing policy that requires grantees to share their data within a reasonable length of time, so long as the cost is modest.

# US Department Of Energy Survey

- The request

http://science.energy.gov/~/media/hep/hepap/pdf/HEPAP_Charge.pdf

**Department of Energy**
Office of Science
Washington, DC 20585

**Office of the Director**

February 25, 2011

As a first step in assessing the policies for researchers funded by the Office of Science, I am requesting your assistance. Please submit to me, no later than July 1, 2011, a report describing current policies and practices for disseminating research results in the fields relevant to the High Energy Physics program. For the purposes of this report, "dissemination" refers to the circulation of research results outside of the originating institutions or scientific collaborations; "research results" refers to both written research findings (scholarly papers, presentations, reports, etc.) and digital data; and "practices" refers to accepted practices within a scientific discipline. Policies from DOE and other federal and non-federal agencies, including foreign institutions and international scientific collaborations, should be considered within the scope of this report provided that these policies have notable impact on the dissemination of research results in your fields. Examples of relevant government policies include provisions in grants and contracts as well as overarching guidance as set forth in federal regulations and DOE orders[1].

W. F. Brinkman
Director, Office of Science

- The criteria for dissemination and who makes this determination.
- How access is provided and controlled.
- Whether access is limited in any way.
- Whether the access comes with any additionalfunctionality.
- The version of the written material or data provided.
- Whether peer review is a condition ofdissemination.
- The institution, DOE user facility, or other body by which the policy is currently upheld.
- Whether, in addition to dissemination, long-term stewardship is accounted for by the existing policy or practice.

# DOE Survey Results

http://science.energy.gov/~/media/hep/hepap/pdf/Dissemination_Report.pdf

## Report of the HEPAP Sub-Committee on the Dissemination of Research Results

June 3, 2011

### Dissemination of Experimental Research Results

The committee contacted experiments across the spectrum of OHEP funded efforts. For the dissemination of research results in the form of papers and presentations, all of the experiments (ATLAS, BaBar, CDF, CDMS, CLEAN, CMS, D0, LHC-b) employ a very similar set of criteria and follow the same dissemination practices. The same is largely true regarding the dissemination of digital data. Since all OHEP funded experimental efforts are international in nature, their dissemination policies are affected by multiple funding agencies and contributing institutions. The policies described below are not formally upheld by any institution outside of the collaborations themselves and instead reflect long established practices within HEP.

### Digital Data

Historically, the digital data collected by HEP experiments are not disseminated for analysis by people outside the originating collaboration. There are several reasons for this. The raw data set acquired for most of the collaborations is large, measuring from hundreds of tera-bytes to tens of peta-bytes in size. In order for the data to be useful a good deal of high level processing must be performed and an understanding of the details of that processing is important in extracting physics results. In addition, corrections for resolution and acceptance effects as well as background contributions must be accounted for using dedicated Monte Carlo simulation samples. Thus, there are a number of technical challenges that need to be overcome in order to make dissemination of the digital data a useful thing to do. Besides addressing the data-handling and database challenges associated with these large data sets, a great deal of additional functionality would also need to be made available. Implicit in all of this is the availability of clear and thorough documentation. Moreover, any potential user must also have access to a large scale computing facility in order to generate and process the necessary Monte Carlo samples and in order to perform the analysis itself. While there is a general consensus that these issues probably all could be addressed, significant additional personnel and capital resources would be required to do so.

The code is disseminated in an open access manner via the internet on dedicated web sites stewarded by the collaborations themselves. Much of the available code is also available via the HepForge repository. The dissemination of the digital data collected by HEP experiments has so far been restricted to limited releases for outreach and education purposes. These releases are typically open access via the internet and stewarded by the collaborations and their host laboratories. They come with some limited additional functionality in the form of software that enables some simple visualization or manipulation of the data. To date no HEP experiment has provided large-scale open access to its raw form digital data, although limited access to processed data has sometimes been granted upon request. The size and complexity of these datasets present significant technological, governance, and support challenges. The IUPAP and ICFA sanctioned DPHEP Study Group is an international effort working to develop solutions to these challenges and to provide common guidelines for use by future collaborations. The preservation of HEP data and its dissemination requires organized action from the experimental collaborations, the participating laboratories, and the funding agencies.

# Public Use of CMS Data

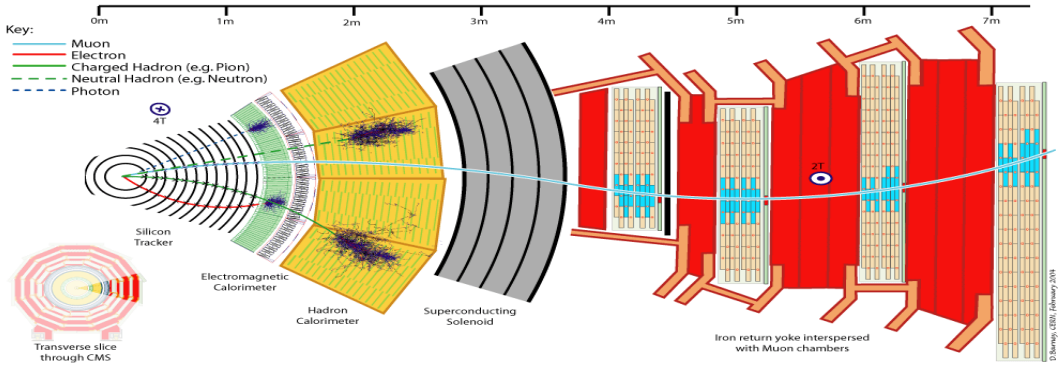


http://www.mattbellis.com/index.php?title=LHC_Data_Hack


https://twiki.cern.ch/twiki/bin/view/Main/CMSPublicData

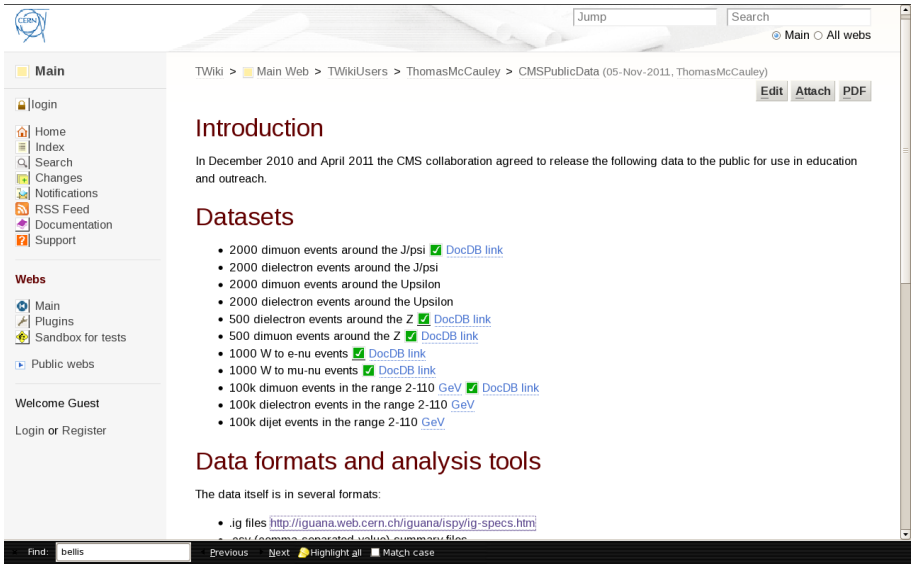


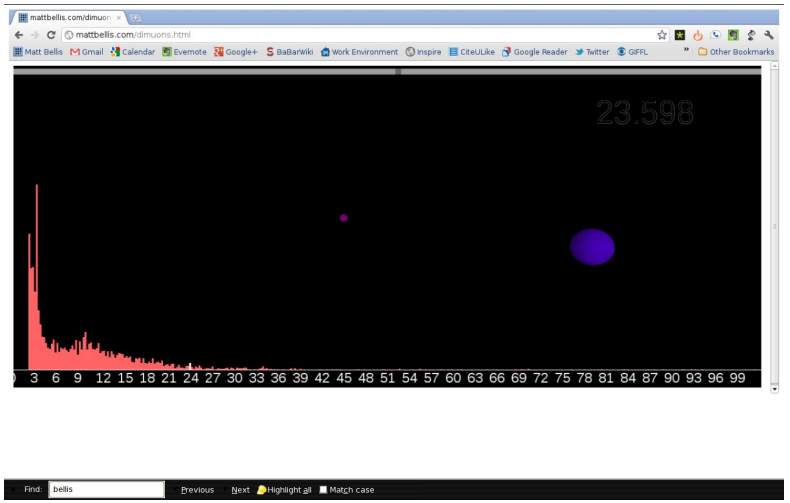

# Outreach Use of BaBar/HERA Data

## The Particle Physics Windchime

http://www.mattbellis.com/windchime/science-hack-day/

Home | How to Use It | What it is | Download | Science Hack Day | Related media | Search

In the middle of Nov, 2010, I participated in Science Hack Day SF. The idea was to get some scientists, coders, web designers and artists together in a room for a weekend and see what they could hack up. David Harris, who I know from his time with SLAC Communication and Symmetry Magazine was one of the organizers and invited me to bring an idea. I'd been thinking about sonfication of data in general and BaBar data (Monte Carlo) in particular, and he had been considering much the same.

At the Hack Day, we presented our concept, got some volunteers to contribute code, ideas and design and by the end of the weekend we had a working demo! The idea was to allow users to take properties of the particles that we observe in our detector (energy, distance from the interaction region, type of detector it is interacting with, etc.) and map that onto sonic characteristics (volume, timbre (instrument), pitch, duration, etc.). In this way, the user can explore the data themselves and find mappings which either make sense to him or her, or are simply more aesthetically pleasing.

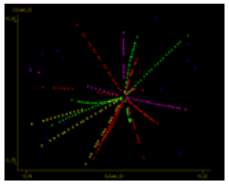

Viewpoints imaging of the final state particles used in the sonfication.

**Release of some data for public use (education) is currently under discussion.**

In our presentation, we explained what we had done and then played samples of Monte Carlo samples of $e^+e^- \to \tau^+\tau^-$ events and $e^+e^- \to B\bar{B}$ events. We mapped the energy of each particle onto the volume for that particular "note" and used the distance the particle flew from the interaction region (center of the detector) to represent the pitch. For those of you who don't work on $e^+e^-$ colliders, we run the machine so that each collision has the same total energy as any other collision. So the total energy of the final state particles is the same for all events. We may not measure every particle that comes out, but to first order they have the same *total* energy. But the $B\bar{B}$ events result in more particles in the detector when they decay compared to the particles produced

Previous | Next | Highlight all | Match case

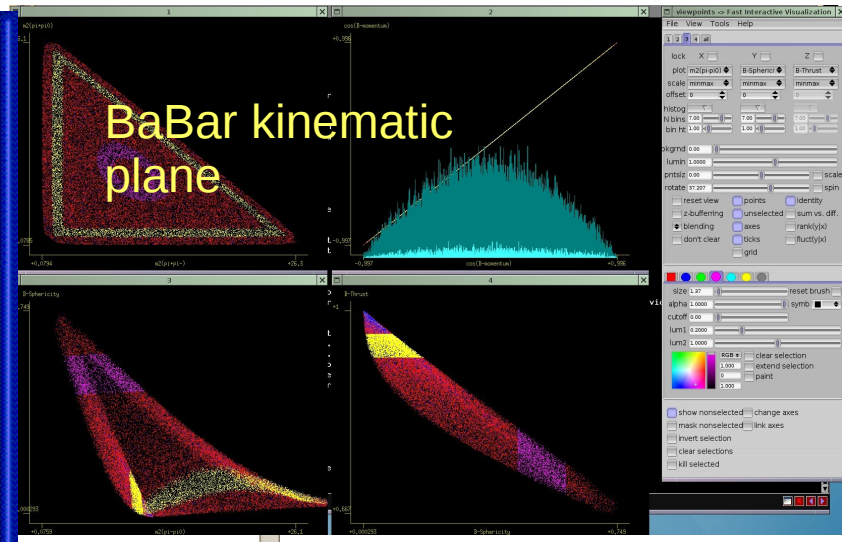


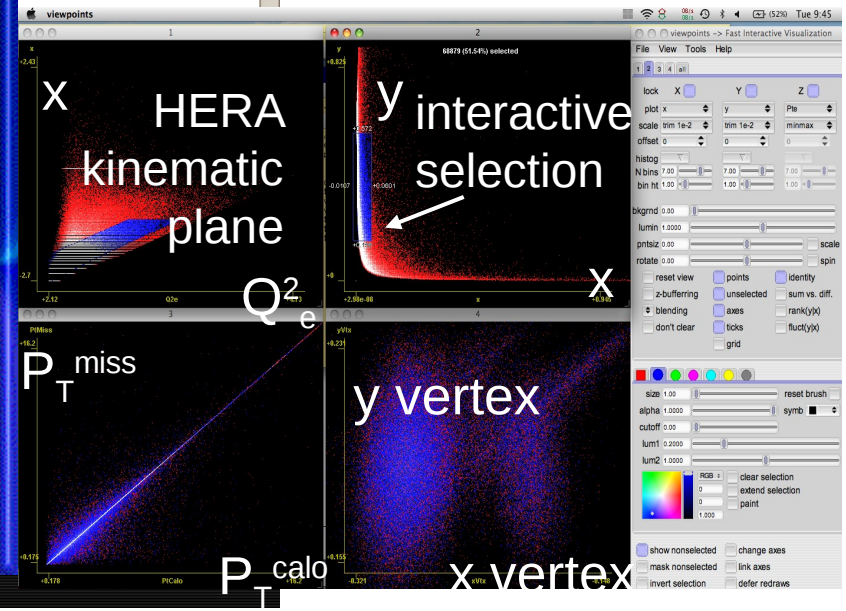


•*Viewpoints* from NASA

# ALEPH Data Preservation Policy

Statement on the use of Aleph data for long-term analyses.

**The Aleph Collaboration**

The data collected by the Aleph experiment in the years 1990-2000 have been archived to allow their use for physics analyses after the closure of the Collaboration. The archiving includes the last set of simulated events and the most updated version of the analysis software.

**Limitations.**

The available information is not sufficient to repeat all analyses, particularly when systematic effects play an important role as, for instance, for precision measurements in the electroweak sector. Examples of physics analyses that cannot be repeated on archived data are

- The measurement of the Z lineshape
- The measurement of the W mass
- The measurement of the tau polarization
- The measurement of lepton and quark forward-backward asymmetries
- Most heavy flavour measurements, such as the measurement of $R_b$, of the CKM matrix elements, of $B_d$ and $B_s$ oscillations
- The searches for the Higgs boson
- Many searches in the Susy sector

**Authorized Users.**

The use of archived Aleph data is authorized to former members of the Aleph Collaboration and their collaborators. The use of a subset of data for teaching and pedagogical purposes, under the guidance of former members of the Collaboration, is allowed.

**Authorship.**

The publication of results based on archived Aleph data is not allowed until 1 year after the official termination of the Collaboration, foreseen for the end of 2004. The authors of the analysis take full responsibility for the publication. Any figure, plot or table using Aleph data should contain the label "ALEPH Archived Data". A reference to the present document "Statement on the use of Aleph data for long-term analyses" must be present in the publication.

Approved by the Aleph Steering Committee
CERN
4 December 2003

# BaBar

**Preserving the essence of an experiment that was conceived 16 years ago (LOI 1994) and collected frontier HEP data over almost a decade through the efforts of 10 countries and ~600 collaborators**

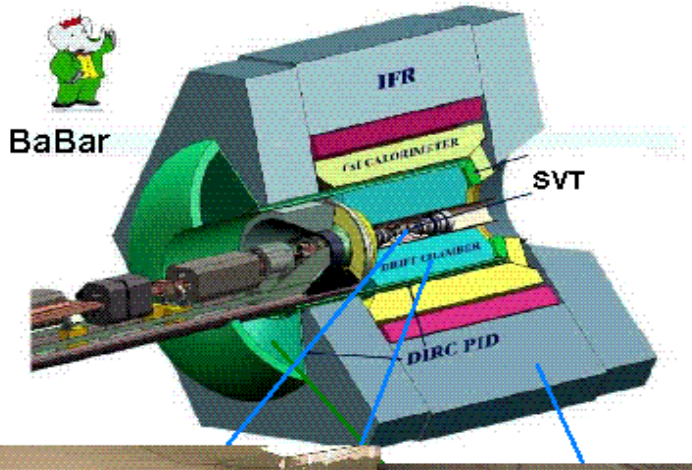


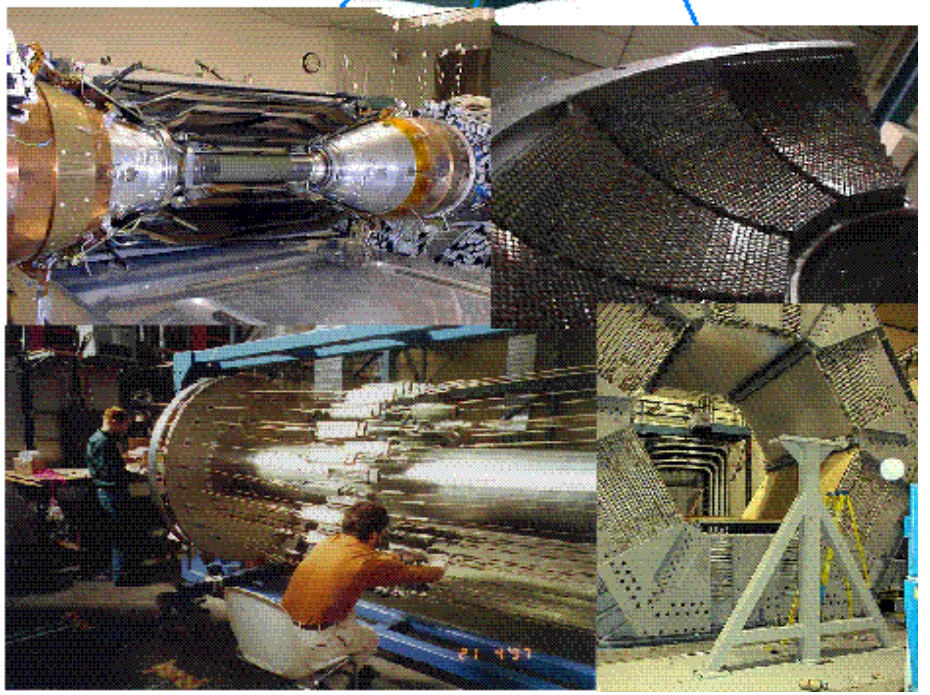

- BaBar has collected good data from Oct 22nd1999 to Apr 7th2008
  - 800TB of raw data, 1.2 PB from R22 and R24 processings
- BaBar (and Belle) data will not be superseded by LHC data
  - Belle II and SuperB will do it in 5-10 years
- **The BaBar Long Term Data Access project aims to preserve both the data and the ability to do analysis until at least 2018 and will provide support for >50 publications foreseen beyond 2012**
  - Need to account for the dwindling resources (both manpower and money

# The BaBar Archival System

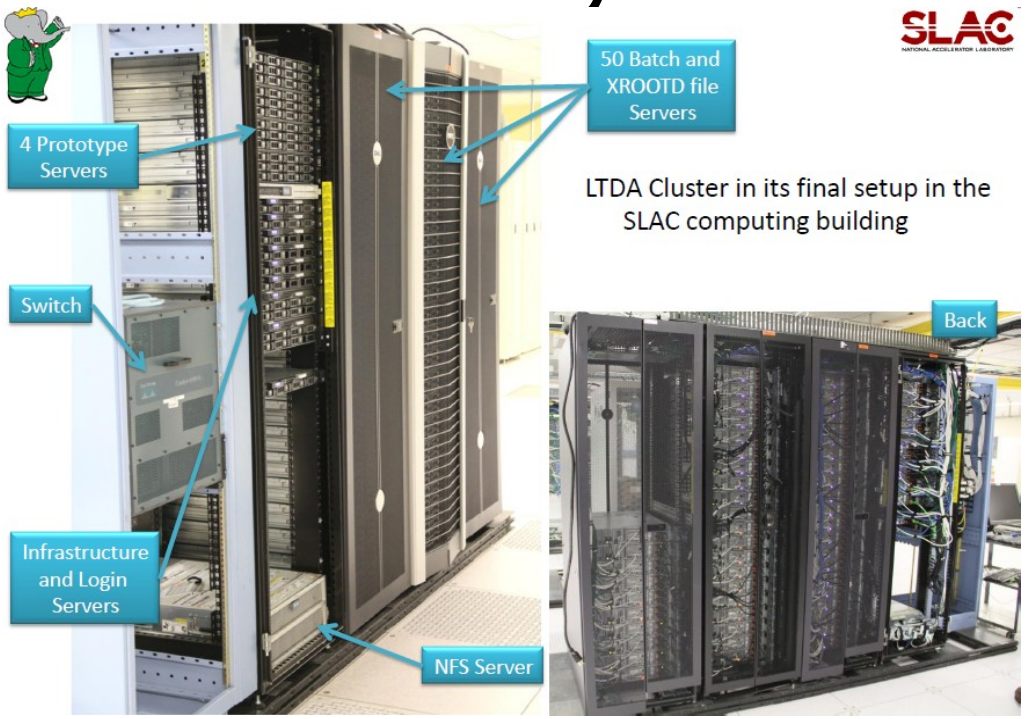


LTDA Cluster in its final setup in the SLAC computing building

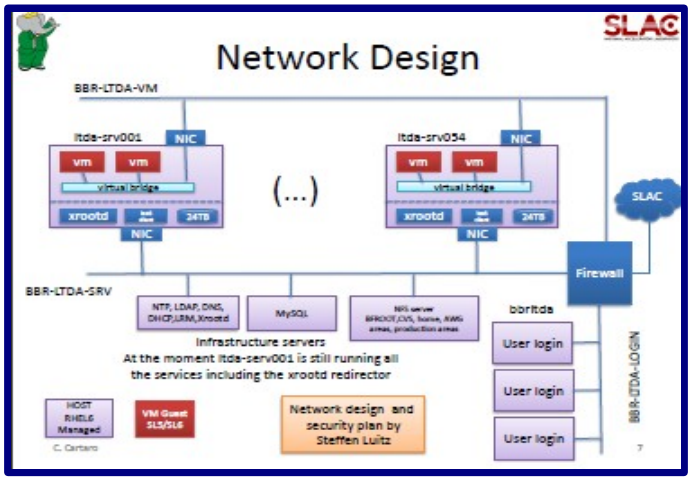


A secure yet functional network design was essential.

- 9 infrastructure servers
  - 3 front end machines, 1 cronserver, 1 test server, 2 service servers (1 hot, 1 cold), 2 database servers (mirrored)
  - Ad hoc hardwareconfigurations
- 50 batch servers
  - Same as prototype servers

## Main Milestones 2010-2011

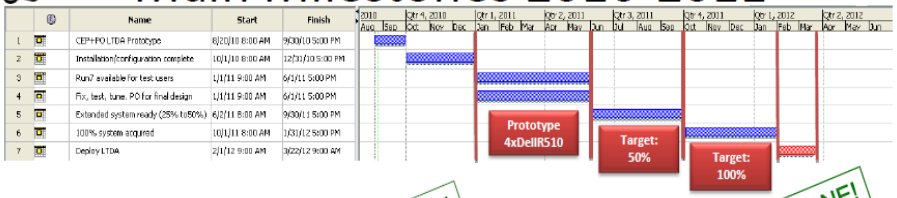


- Sep 2010 — DONE!
  - PO for the prototype — DONE!
  - Prototype on site by October — DONE!!
- Dec 2010
  - Installation/configuration complete — DONE!
- Jan 1st 2011
  - System available for test users
- Jun 1st 2011 — DONE!
  - Test phase ends, LTDA final design ready
  - PO for first 50%
- Jul 1st 2011 — DONE!
  - First 50% of LTDA available
- Oct 1st 2011 — DONE! 100% on site
  - PO for 100% of LTDA
- March 21st, 2012
  - Deployment of the LTDA
  - All new analyses will use LTDA

# Bringing LTDA to the Users

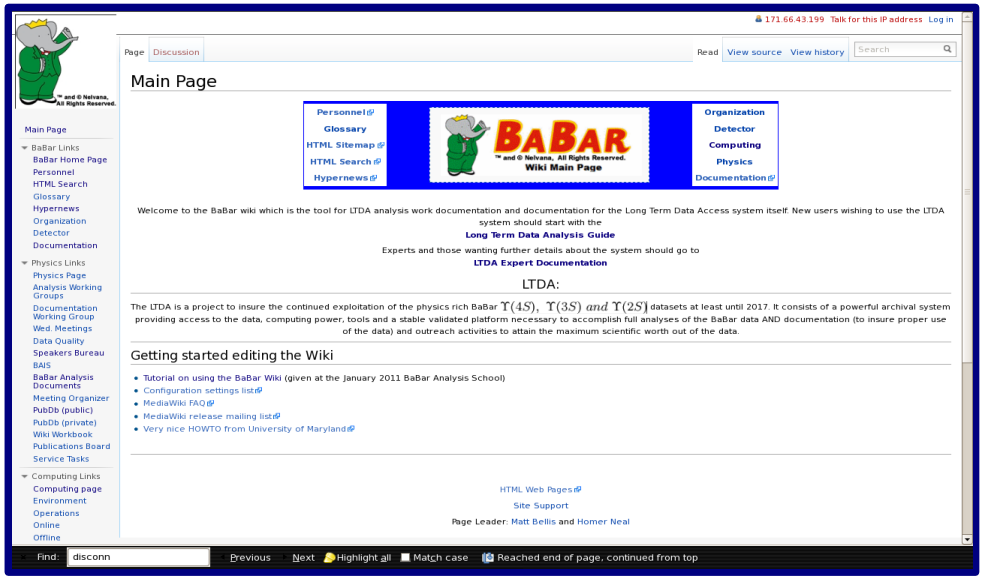


Production phase will start
March 2012
but already > 20 users trying out analysis and production now

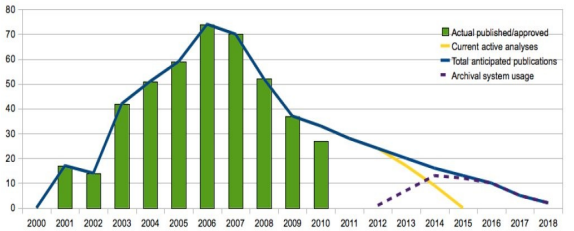


**50 Batch/Storage servers**

Each with:

**12 x 3.1 GHz cores** (simulates 24 with hyperthreading – 60% gain for full load)

**24 Tbytes** of storage on each server

2 Tbytes for scratch and the rest for XROOTD storage

Batch system using Torque/MAUI

SL6 and SL5 virtual machines

# Lessons learned … so far ...

- Easy to get trapped in unexpected details (security!)
- Beware of unexpected costs (licenses!) 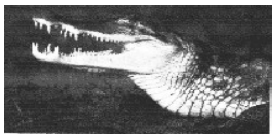
- Complete independence from the existing computing infrastructure is difficult to achieve and not necessarily desirable.
- Storage costs are not negligeable
- No twisting of arms needed to get experts to voluntarily contribute

We were very fortunate in the support we received from the DOE, SLAC, BaBar Management and collaboration experts who although were heavily involved in new projects voluntarily jumped in when needed.

# Other BaBar Archival Actions

- Migration to new media completely done
- Old tapes stored away safely … but old drives would have to be purchased to read them
- Backing up all legacy data to ccin2p3 (~1 year to transfer)
- Tape library for accessing raw data copy at Padova restored
- Backup of analysis storage areas
- History of analyses being preserved in the joint BaBar/BELLE physics of B-factories book (PBF)
- Looking into INSPIRES for long term archiving of analysis, appropriate support documents and data taking logs

# DESY data preservation

## Reminder: HERA experiments aim for level 4 preservation

| 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 and HERMES aim for DPHEP level 4, ZEUS between levels 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 also be pursued..



## Validation Project at DESY in 2011

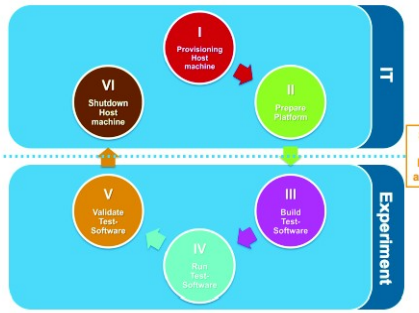


- Latest developments due to influx of Masters student manpower in the last months and full project to begin now
  - More details: Yves' talk 16.05.11
- Now we have to examine what the experimental part of this looks like
- Clear need for synergic input from both parties
- H1, ZEUS and HERMES to participate - I use H1 as an example



## Summary

- HERA Data Preservation Group now in place at DESY
  - Successful manpower requests almost now in place
- Joint validation project underway, expect much progress in coming months
  - Will require some intense effort from all parties in the next period
- Other work with INSPIRE also looks promising
- Streamlining of the hardware, the documentation and the collaborations themselves is also in motion

# RECAST

## RECAST

K.C., Itay Yavin [hep-ex/1010.2506], JHEP.

Center for Cosmology and Particle Physics

The proposal now is to build RECAST, a framework for recasting analyses

- cuts don't change, no need to look back at the data or estimate backgrounds and systematics,
- "the signal is the easy part" -- Conservative Physicist
- data is kept private, still goes through necessary approval process as determined by collaboration, original paper receives citation & recognition

RECAST front-end is just an information broker

- standardize format, API allows for process to be automated,
- back-end implementation is up to collaboration (perhaps supported by national labs)
- ... and it's technically easier than archiving data

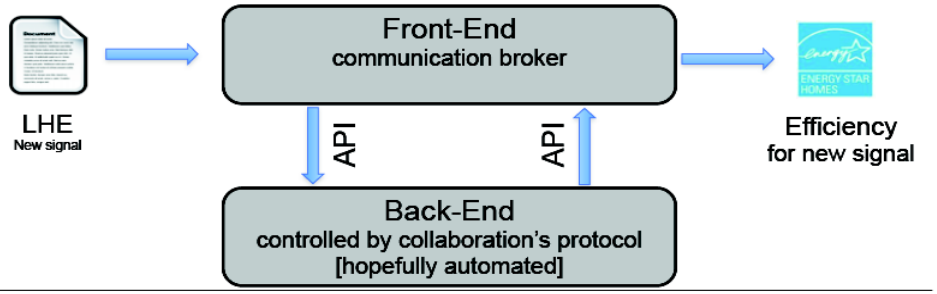




## Conclusions

Center for Cosmology and Particle Physics

DPHEP's preservation models are are very valuable right now

- not just about long-term, "what-if" scenarios
  - strategy: extend value proposition to areas where impact is immediate
- efficient techniques to address inverse-problem need:
  - likelihood functions of our existing measurements
  - ability to recast existing searches

Analysis archival is complementary to Data archival

- requires Level-4 infrastructure, but it is not inclusive of L1-3
- Analyses archival eliminates many costs/risks/...
  - a potential answer to data management plans for LHC experiments
- it also poses new challenges:
  - decay time is weeks, not years! Must act while analyses are still alive.

# Must Prepare The Path for managing ever Growing Quantities of Data

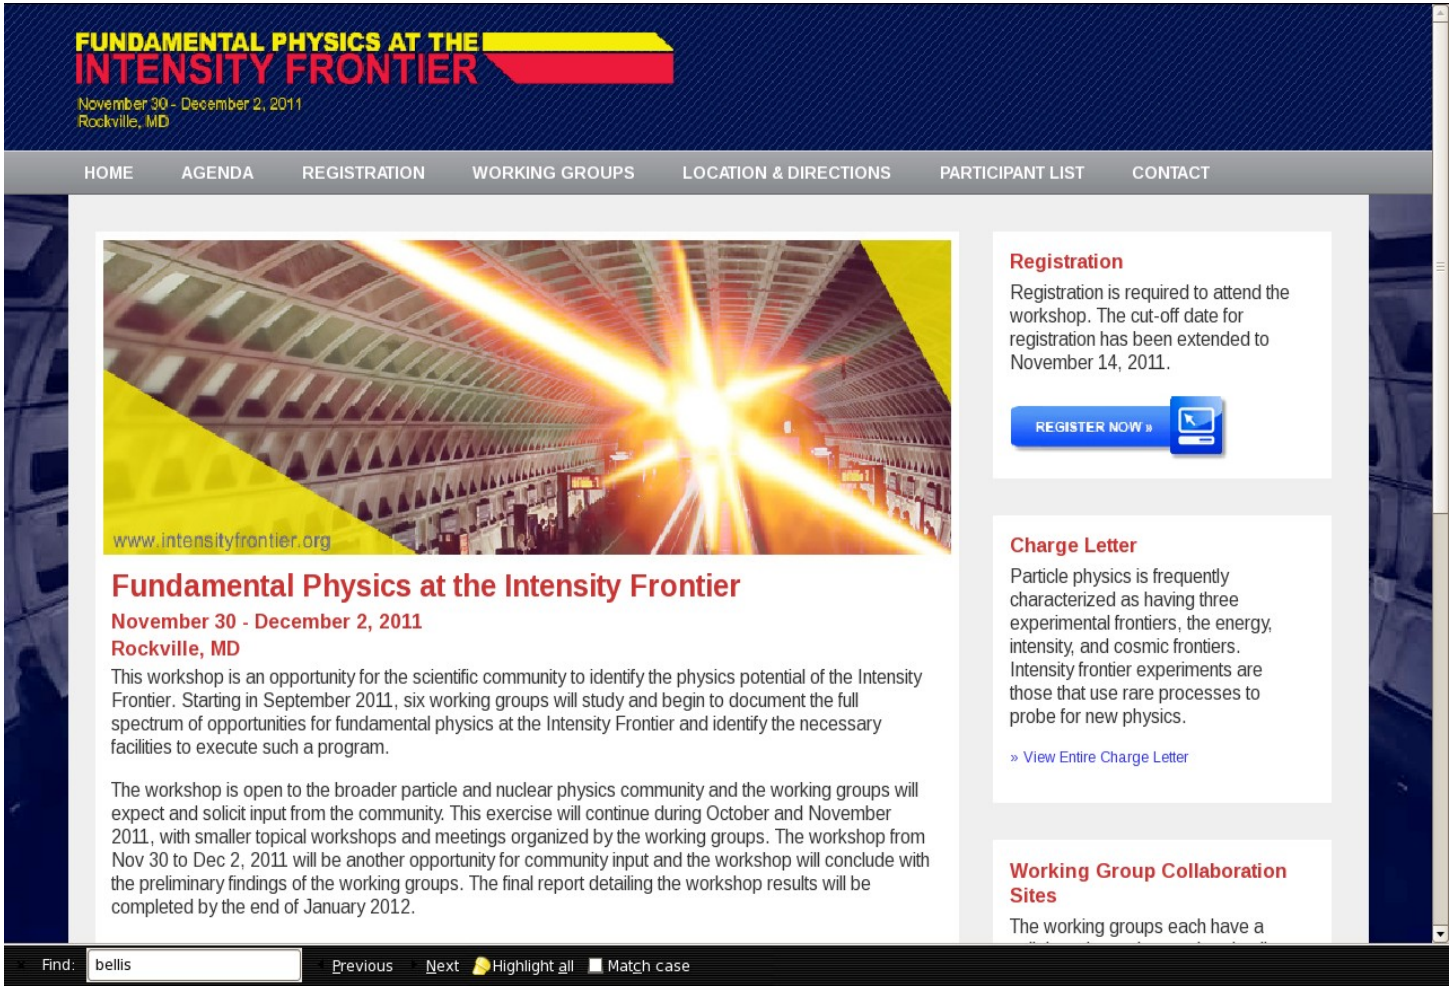

# A Multi-Laboratory European Effort

- CERN+DESY+CCIN2P3 are forming a proposal to support data preservation in HEP
- They've requested and received US support for the concept
- Archivist needed to implement data preservation need a source of funding, the appropriate tools (guidance in particular) and backing for the mission.
- This looks like a solid step in that direction.

# Where Guidance is Needed

**Implementing a feasable plan from the start**

- Personnel
- Data format
- Migration
- Testing
- Enduring platform
- Management
- Analysis approval
- Openness of the data
- Code and documentation
- Authorship
- Joint experiment combination
- Data reinterpretation

# Guidance from DPHEP

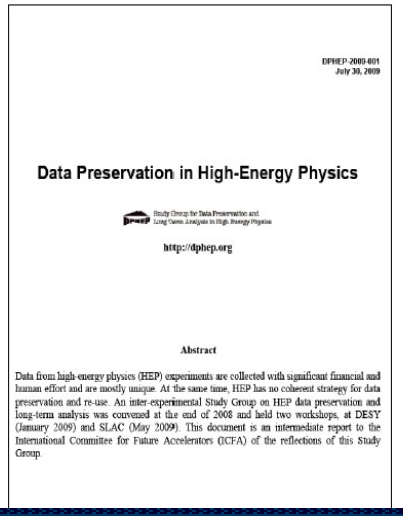
DPHEP-2009-001
July 30, 2009

Data Preservation in High-Energy Physics

Study Group for Data Preservation and Long Term Analysis in High Energy Physics

http://dphep.org

Abstract

Data from high-energy physics (HEP) experiments are collected with significant financial and human effort and are mostly unique. At the same time, HEP has no coherent strategy for data preservation and re-use. An inter-experimental Study Group on HEP data preservation and long-term analysis was convened at the end of 2008 and held two workshops, at DESY (January 2009) and SLAC (May 2009). This document is an intermediate report to the International Committee for Future Accelerators (ICFA) of the reflections of this Study Group.

> First recommendations of the group published November 2009 arXiv: 0912.0255

> The report covers the four key areas

- Physics Case for Data Preservation
- Preservation Models
- Technologies
- Governance

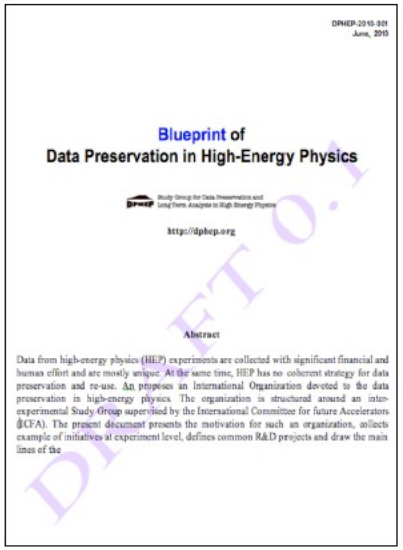
DPHEP-2010-001
June, 2010

Blueprint of
Data Preservation in High-Energy Physics

Study Group for Data Preservation and Long Term Analysis in High Energy Physics

http://dphep.org

DRAFT 0.1

Abstract

Data from high-energy physics (HEP) experiments are collected with significant financial and human effort and are mostly unique. At the same time, HEP has no coherent strategy for data preservation and re-use. An proposes an International Organization devoted to the data preservation in high-energy physics. The organization is structured around an inter-experimental Study Group supervised by the International Committee for future Accelerators (ICFA). The present document presents the motivation for such an organization, collects example of initiatives at experiment level, defines common R&D projects and draw the main lines of the

> More details and specifics, focussing on:

- An expanded physics case for data preservation
- A description of the existing experiment-level and lab-level strategies for data analysis and preservation
- Global transverse data preservation projects, such as outreach, validation, documentation
- Future international coordination of activities

> Cost estimates also included

> Should provide a skeleton for local, national and international proposals, for past an future experiments

- LEP!
- Other branches of HEP (HI, neutrinos etc.)

# DPHEP Executive Summary

The DPHEP Study Group has identified the following priorities, in order of urgency:

> Priority 1: **Experiment Level Projects in Data Preservation.** Large laboratories should define and install data preservation projects in order to avoid catastrophic loss of data once major collaborations come to an end. The recent expertise gained during the last 18 months indicate that an extension of the computing effort within experiments with a person power of the order of 2-3 FTEs leads to a significant improvement in the ability to move into a long-term data preservation phase. Such initiatives exist already or are being defined in the participating laboratories and are followed attentively by the Study Group.

> Priority 2: **International Organisation DPHEP.** The efforts are best exploited by a common organisation at the international level. The installation of this body, already prefigured by the ICFA Study Group, requires a Project Manager (1 FTE) to be employed as soon as possible. The effort is a joint request of the Study Group and could be assumed by rotation among the participating laboratories.

> Priority 3: **Common R&D projects.** Common requirements on data preservation are likely to evolve into inter-experimental R&D projects (three concrete examples are given above, each involving 1-2 dedicated FTE, across several laboratories). The projects will optimise the development effort and have the potential to improve the degree of standardisation in HEP computing in the longer term. Concrete requests will be formulated and the activity of these projects will be steered by the DPHEP organisation.

These priorities could be enacted with a funding model implying contributions from the three regions (Europe, America, Asia) and strong connections with laboratories hosting the data samples.

# Summary

- Many thanks for inviting me back to ICFDT
- Very honored to be presenting this topic near the future site of the ***SuperB*** experiment for which data preservation is a very important issue.
- Much progress in data preservation has occured since the 1st ICFDT
- Data preservation in HEP is important because:
  - Relevant physics cases for future use can be made
  - It is timely, given the current experimental situation and plans
  - Enhances the return on the initial investment in the experimental facilities
  - It provides additional research at particularly low cost
- It requires a strategy and well-identified resources
  - International cooperation is the best way to proceed
- Unique opportunity to build a coherent structure for the future: DPHEP
- DPHEP Blueprint on Data Preservation is on the way

# Thanks


- Thanks to the Department of Energy for their backing and funding of data preservation activities at SLAC
- Thanks to the DPHEP group for their collaboration and contribution of material for this presentation
- Thanks to the LTDA international Advisory Committee for their guidance
- Thanks to CCIN2P3 for their continuing great support of BaBar and many other HEP projects

# Best wishes

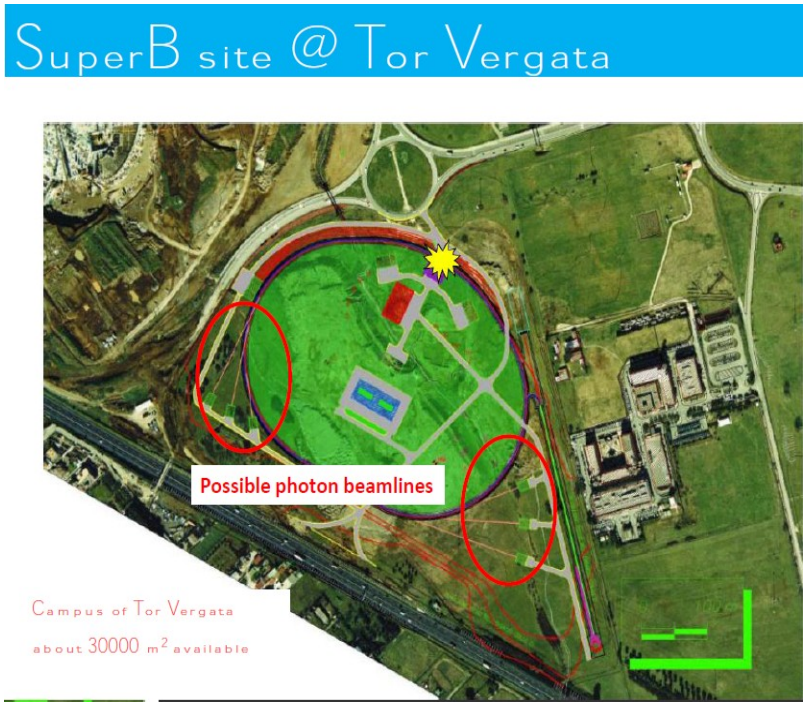
SuperB site @ Tor Vergata
Possible photon beamlines
Campus of Tor Vergata
about 30000 m² available

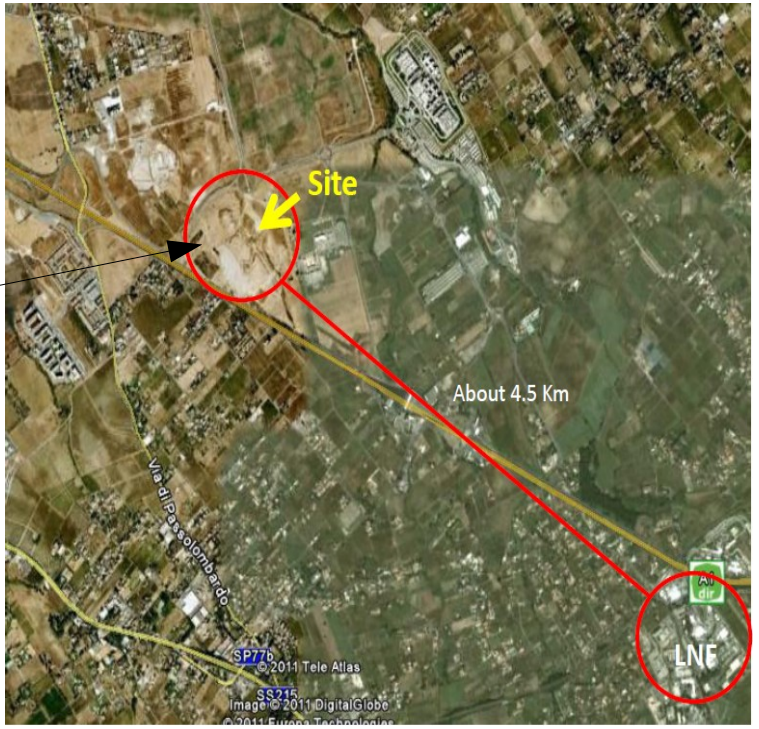
Site
About 4.5 Km
LNF
© 2011 Tele Atlas
Image © 2011 DigitalGlobe

# Extra

# The LTDA Developers

–BaBar Computing Coordinator: Tina Cartaro
–BaBar software expert and beta tester: Homer Neal and Marcus Ebert
–Network design: Steffen Luitz
–Virtualization expert: Kyle Fransham
–System performance expert: Igor Gaponenko
–Documentation: Matt Bellis
–Databases, tools and production: Douglas Smith and Tim Adye
–Computing Division experts
•System setup and adminstration: Booker Benseand Lance Nakata
•Xrootdexperts: Wilko Kroeger
•Network setup: Antonio Ceseracciu
•BaBar-SLAC liason: Len Moss

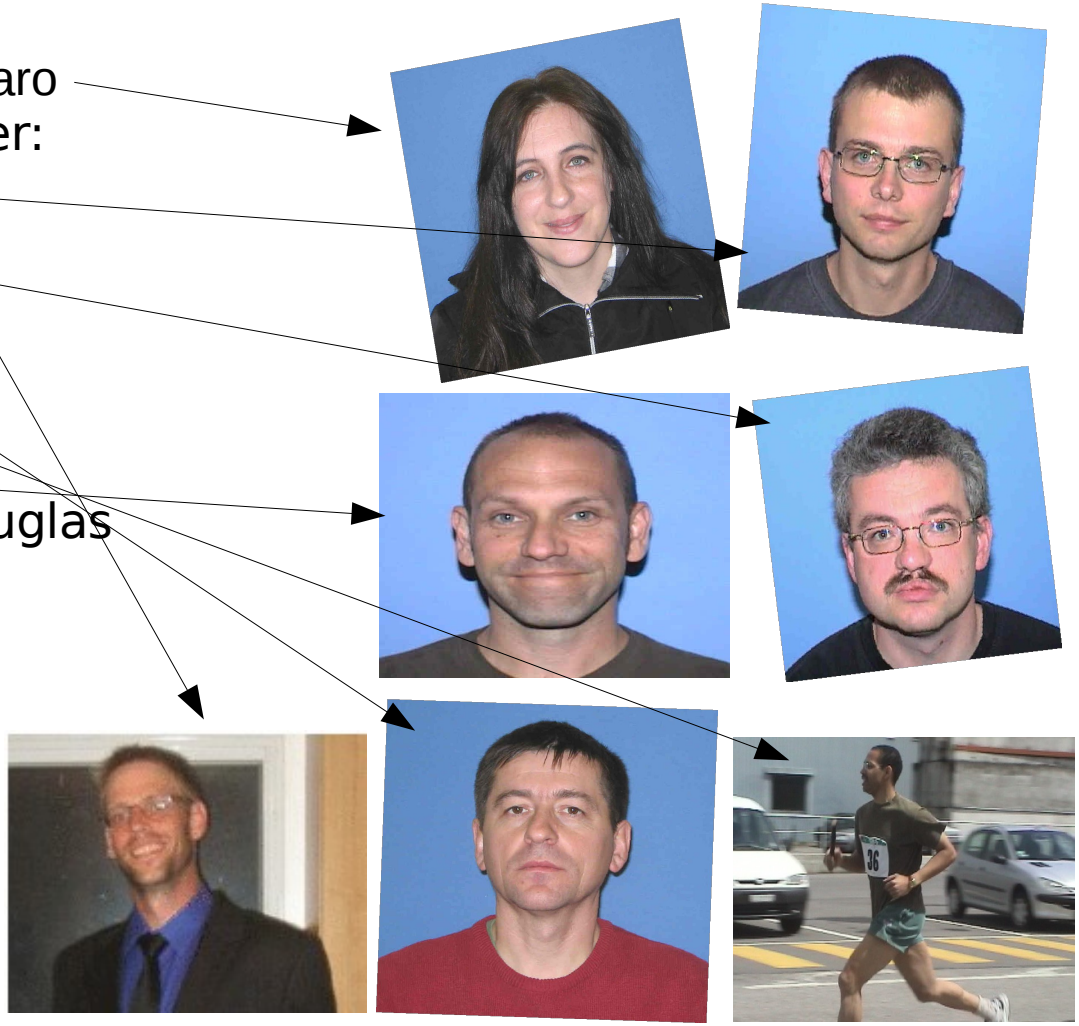

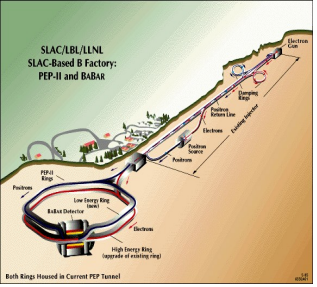


# An Elephant among Databases

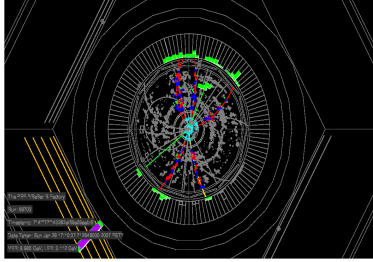

## Press Release: World's Largest Database reaches 500,000 Gigabytes

http://home.slac.stanford.edu/pressreleases/2002/20020412.htm

Date Issued: April 12, 2002

Contact:Neil Calder, Stanford Linear Accelerator Center

Last week at the Stanford Linear Accelerator Center (SLAC), the BABAR experiment's database stored its 500,000th Gigabyte - a milestone that makes it the largest known database in the world. The BABAR experiment - a collaboration of 600 physicists from nine nations - observes collisions between subatomic particles to understand how the behavior of matter and antimatter shaped our universe. BABAR, also known as the "B Factory," mass-produces huge quantities of scientific data with industrial efficiency. Up to 500 Gigabytes of data is sent relentlessly to the experiment's database daily.

The half million Gigabytes of data in the BABAR database, printed out, would fill one billion books. That's nearly 60 times the number of books in the Library of Congress, the largest library in the world. "The need to store the avalanche of information coming from the experiment and then efficiently search and retrieve specific data samples has driven physicists and computer experts to create innovative technology," said SLAC Director Jonathan Dorfan. "Governments, commercial corporations and institutes will face similar needs in the near future and the knowledge and experience we have gained will be passed on."

http://www.slac.stanford.edu/BFROOT/www/Public/Computing/Databases/proceedings/CIDR05.pdf

### Lessons Learned from Managing a Petabyte

Jacek Becla*
Stanford Linear Accelerator Center
2575 Sand Hill Road, M/S 97
Menlo Park, CA 94025, USA
becla@slac.stanford.edu

Daniel L. Wang*
Stanford Linear Accelerator Center
2575 Sand Hill Road, M/S 97
Menlo Park, CA 94025, USA
danielw@slac.stanford.edu

BaBar's vibrant research effort continues to demand more data, more quickly from the detector, as well as higher levels of data service. The first generation eventstore was undoubtedly a great success, providing storage and service throughput well beyond its original design goals. Data rates, for example, were several times higher than originally designed. Hundreds of users analyzed data in BaBar. Its complexity and size has put it beyond today's scalability frontier: in 2003, it was larger than the largest 200 relational databases combined, earning the grand prize in Winter Corporation's TopTen Program (a survey of world's largest databases).

Further details on this ODBMS-implementation can be found in [1], [2], and [3].

As the only commercial software in a sea of 5 million lines of home-grown C++, Objectivity/DB was the only non-source-accessible component. Though support was excellent, their priorities and release cycle did not always align with BaBar's. Dependence on their libraries, for example, locked an otherwise open system in certain compilers and operating systems. Monetary cost was also an important factor.

BaBar wanted to eliminate dependence on commercial software, and so chose ROOT I/O [13], an open-source, almost BSD-licensed persistence technology that had wide acceptance in the HEP community. ROOT I/O lacked many standard ODBMS features, never claiming to be a "database," but early tests indicated that it could be adapted to meet BaBar's demands.

Improving data administration was another goal. Users found the array of database files representing their data unnecessarily difficult to manage. They eschewed the organization of database files as clustering constructs for their data, whose logical structure depended on collections storing references to events with references to data objects. The new computing model simplified the mapping, defining the one-collection-one-set-of-files model. This model also simplified exporting data and facilitated unconnected laptop analysis.

#### 5.2 Scaling the system

When dealing with petascale data sets, scalability issues come up everywhere: number of servers, files, persistent objects, connections to servers, crashes, lines of code, and so on. Managing such data sets is a challenge easily underestimated. Very large systems are complicated by nature, therefore the simpler the solution, the better. At this scale, managing the metadata of the data set itself becomes a major problem that requires special consideration in the early stages of design. Being a