
CERN

# Challenges in data preservation in High Energy Physics (HEP)

# Challenges in data preservation in High Energy Physics (HEP)

- **Chapter I**: Data preservation: what it is
- **Chapter II**: Aspects of data preservation
- **Chapter III**: Lessons learned from LEP
- **Chapter IV**: Summary and conclusions

# Challenges in data preservation in High Energy Physics (HEP)

- **Chapter I**: Data preservation: what it is
- **Chapter II**: Aspects of data preservation
- **Chapter III**: Lessons learned from LEP
- **Chapter IV**: Summary and conclusions

# Let’s start with an example !

The following slides are based on a

Talk by Siggi Bethke at KEK 2010

with his kind permission.

Have a look!

# The JADE experiment @ PETRA (DESY)

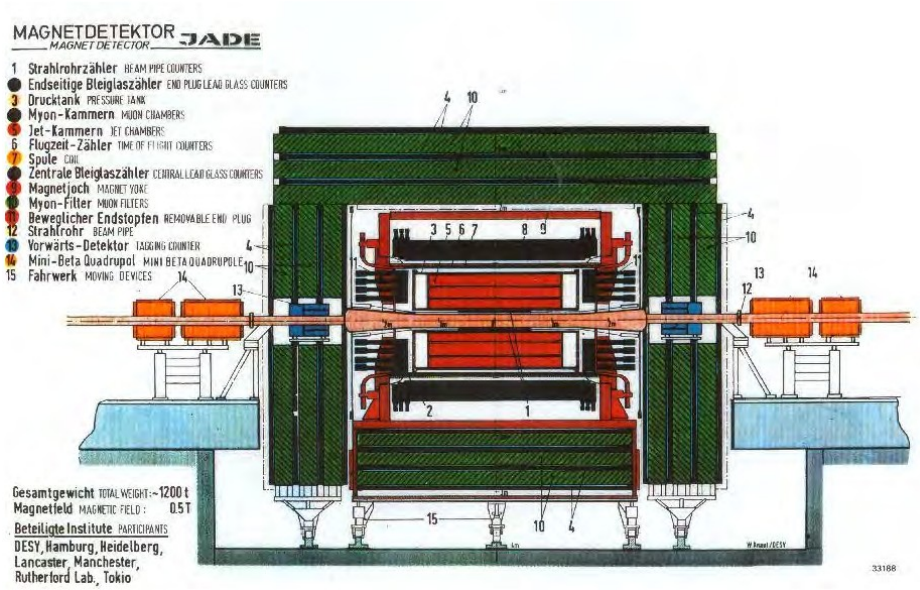


*Courtesy: Siggi Bethke*

- Data taking **1979-1986**
- $e^+e^-$ collision data at energies of **12-46 GeV**
- Data taking stopped 3 years before LEP was finished in 1989

# The JADE experiment @ PETRA (DESY)

**Idea:** (around ~1995)

- Re-analyse this old data
- Use new tools and exploit experiences from LEP
- Combine this with LEP data

  (90 GeV-209 GeV)

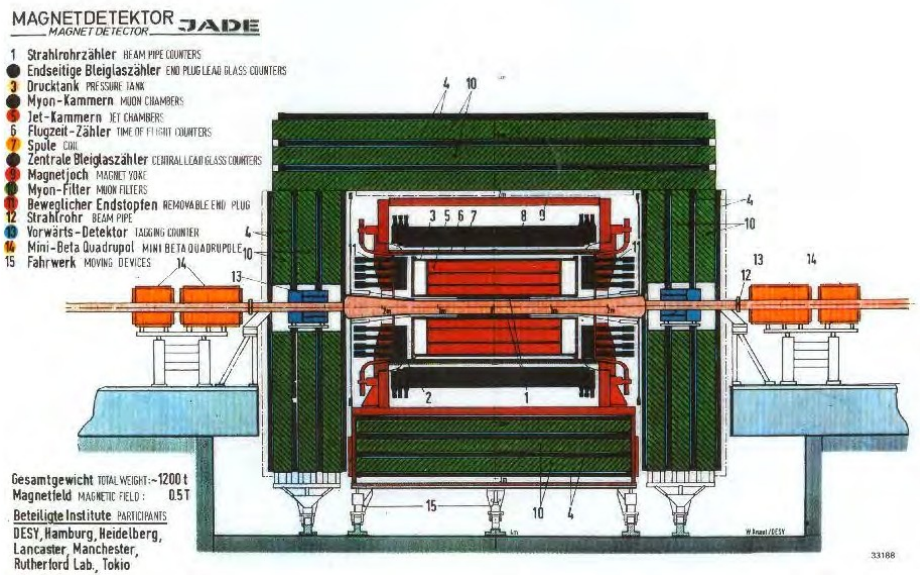


*Courtesy: Siggi Bethke*

Siggi Bethke et. al.

# Re-analysis of JADE data

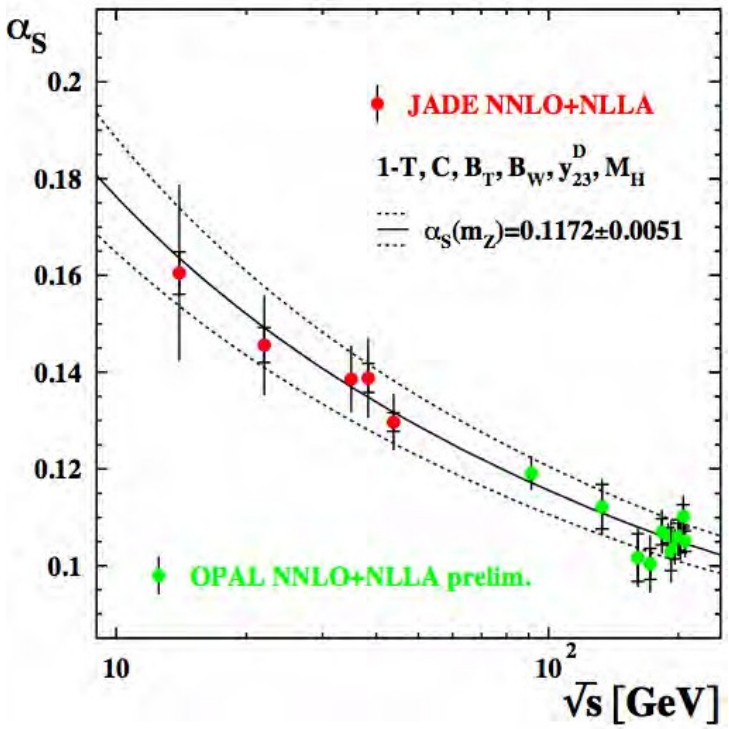


*Courtesy: Siggi Bethke*

- Re-vitalisation of **JADE software**
  - 1995 - 2003
  - Software partly rewritten
  - Full stack recovered
  - Simulation, reconstruction, event display
- Revitalisation of **JADE data**
- Result: **proof of Asymptotic Freedom**
- **See: Eur.Phys.J.C64:351-360,2009**
  - Siggi Bethke et. al.

# Some anecdotes along the line ....

- one important „calibration“ file, containing the recorded luminosities of each run and fill, was stored on a private account and therefore lost when DESY archive was cleaned up.

  *Jan Olsson, when cleaning up his office in ~1997, found an old ASCII-printout of the JADE luminosity file. Unfortunately, it was printed on green recycling paper - not suitable for scanning and OCR-ing. A secretary at Aachen re-typed it within 4 weeks. A checksum routine found (and recovered) only 4 typos.*

- an old version of the original BOSlib 1979 version was found, on our request, at the Univ. of Tokyo computer centre.

- Peter Bock, when cleaning out an old lab at the Physics Institute at Heidelberg University, found a few 9-track tapes containing original JADE MC files which were very valuable for validating results of our first re-analyses in ~1997

*Taken from a talk given by Siggi Bethke at KEK 2010, With kind permission from S. Bethke*

# What is data preservation about ?

(Large) experiments collect a **huge amount of scientific data** and software to analyse it.

- What happens with this **after the end** of an experiment ?
  - Throw it way ?
  - Or keep it ?
- And then ?
  - Who owns that data ?
  - Who can access it ?
  - Who maintains it ?
  - For what ?

Dismantling of the DELPHI detector at LEP (soruce: CERN/CDS)

**“(...)The main goal of data preservation is to protect data from being lost or destroyed and to contribute to the reuse and progression of the data.”**

Source: https://en.wikipedia.org/wiki/Data_preservation

# Lessons learned and consequences

Data preservation in high energy physics is important for

- Continued **improvements in theory, experiment and simulation** which may require a re-analysis of old data
- **New ideas** coming up
- Unexpected **discoveries**

DPHEP group since 2009

DATA PRESERVATION

## Study group considers how to preserve data

For experimentalists in high-energy physics, the data are like treasure, but how can they be saved for the future? A study group is investigating data-preservation options.

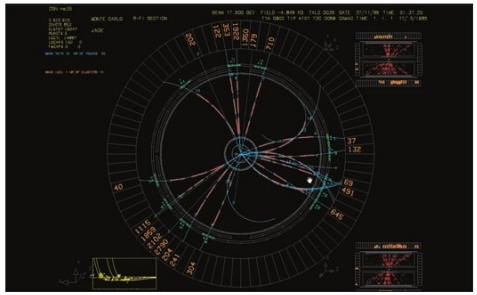

*A simulated event in the JADE detector, generated using a refined Monte Carlo program and reconstructed using revitalized software more than 10 years after the end of the experiment. (Courtesy Siggi Bethke.)*

High-energy-physics experiments collect data over long time periods, while the associated collaborations of experimentalists exploit these data to produce their physics publications. The scientific potential of an experiment is in principle defined and exhausted within the lifetime of such collaborations. However, the continuous improvement in areas of theory, experiment and simulation – as well as the advent of new ideas or unexpected discoveries – may reveal the need to re-analyse old data. Examples of such analyses already exist and they are likely to become more frequent in the future. As experimental complexity and the associated costs continue to increase, many present-day experiments, especially those based at colliders, will provide unique data sets that are unlikely to be improved upon in the short term. The close of the current decade

the complexity of the hardware and a more dynamic part closer to the analysis level. Data analysis is in most cases done in C++ using the ROOT analysis environment and is mainly performed on local

CERNCOURIER 29 April 2009

# Challenges in data preservation in High Energy Physics (HEP)

- **Chapter I**: Data preservation: what it is
- **Chapter II**: Aspects of data preservation
- **Chapter III**: Lessons learned from LEP
- **Chapter IV**: Summary and conclusions

# Aspects of data preservation

- Access policies
- Bit preservation
  - e.g. raw data, simulations
  - calibration databases, ...
- Software preservation
- Document Archival
- Analysis preservation


*cernvm screenshot OPAL event loop, curtesy: Frank Berghaus*

# Access policies

- Who can **access the data** and under which conditions ?
- Usually experiment specific
  - Experimental data in HEP is **very complex**
  - Access conditions for real data often **restrictive**
- Several initiatives addressing this

# Bit preservation

- Long term storage of data, including
  - Raw data
  - Reconstructed data
  - Simulations
  - Databases, …
- Data format and representation
  - Typically **compressed binary data**
  - E.g. BOS, ZEBRA, ROOT
  - Human readable form would take too much space
- Technically considered to be a **solved problem**

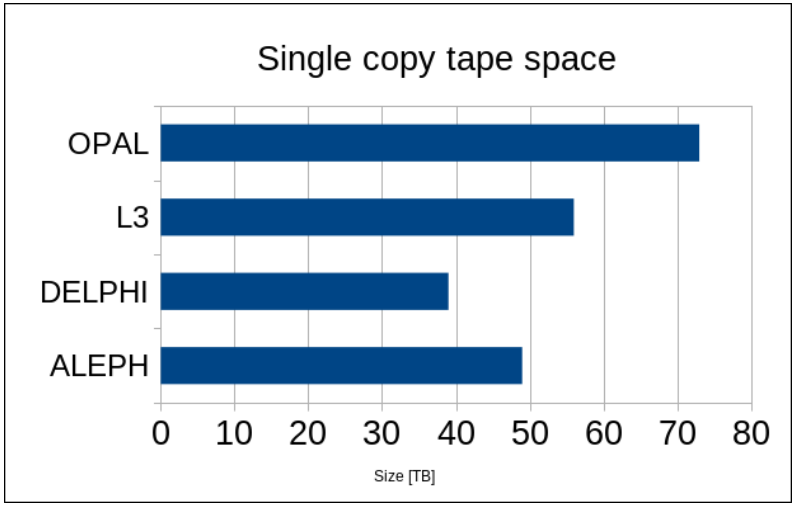


Tape storage in use for LEP data

# Bit preservation

- Storage media options
  - **CD, DVD**: typically not enough capacity
  - Spinning **disks, SSD**: expensive, needs power
  - **Tape**: best for archiving of little used data
    - Can be stocked in a tape silo
    - High storage capacity
- Currently mostly on Tapes archives
  - Offline archival, needs time to get access
  - Multiple copies needed
  - One or more off-site copies

Tape silo at CERN

# Tape storage

- **Different technologies** used by different sites
  - HPSS
  - dCache
  - CERN Advanced STORage manager (CASTOR)
  - CERN Tape Archive, CTA (replacing CASTOR at CERN)
  - ...
- Storage media is **periodically replaced**
  - Typically increasing capacity per tape
  - **Reduces storage costs** over time (Kryder's law)
  - Caveat: **vulnerability in case of failures**
- Tape vendors become an issue. They are getting rare.
  - (Automatic) migration to new media and/or technologies
  - Cost for tape and new media ?

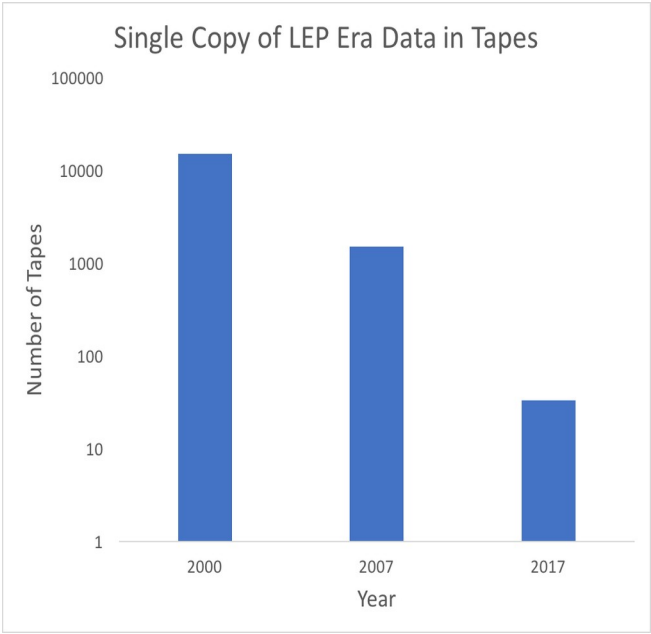


*Curtesy: Frank Berghaus*

# Tape storage

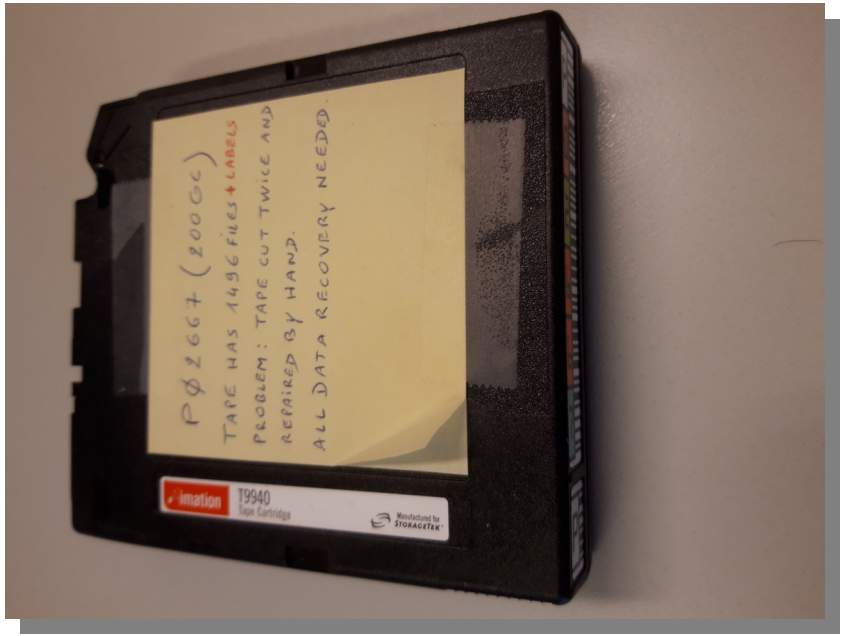


- Damaged tape from CASTOR containing DELPHI **data** files
- Damaged files could be restored from an external copy in Spain
- **No data loss**, but we've been lucky

# Software preservation

**The bare data is useless without the software to read it.**

- Shared software between the experiments
  - HEP libraries, e.g. CERNLIB
- Experiment frameworks
  - Reconstruction code
  - Simulation
  - Analysis frameworks

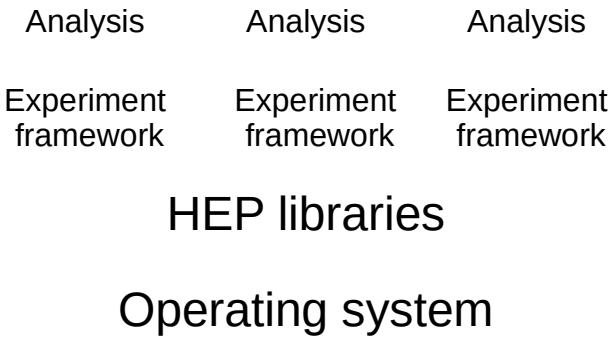

# Software preservation: options

Keep **software sources** alive

- Requires access to the full sources
- Keep adapting source code to changing computing environment
  - In version control system
  - As software CD with build scripts
- **Frequently rebuild and validate** the results
- Most flexible but also most **labor intensive** approach

# Software preservation: options

## Encapsulation and emulation

- Install binaries on **physical hardware** and archive that
  - Likely to break physically with the years
  - Does not scale
- Install binaries in **VM** or in a **container**
  - Or keep only the environment on the VM/container
  - Install the binaries on shared file system like AFS or cvmfs
  - No guarantee that the images will still work after many years, as technology changes
- External dependencies and interfaces, e.g. with storage systems can be an issue
  - E.g. changing access protocols (like rfio which was used largely at LEP)

# Knowledge preservation

Data and software are useless without the documentation on how to use them:

- Software without documentation is just another piece of data
- Knowledge preservation has to start early on in the process, before the experts leave, change, forget (or die ...)

# Knowledge preservation

- Manuals
  - Detector specifications
  - Software, interfaces, …
- Internal and technical notes
- Theses
- Conference contributions
- Publications
- Web pages and documentation

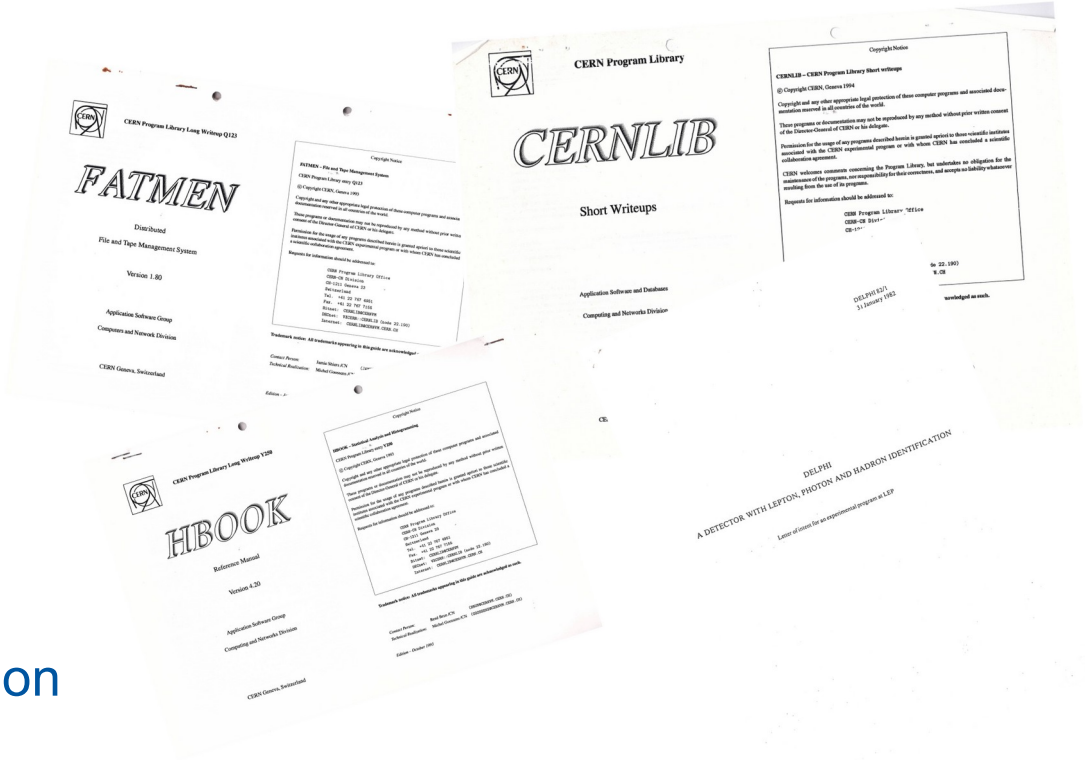

# Analysis preservation

- Archiving of Analysis
  - Requires **action during the run time** of the experiment already
  - A **common format** simplifies this a lot
  - Both software (sources) and building
  - Rapid re-run and reproduction of e.g. published results
  - **Re-use of existing and tested analyses** and processes
- Relevant for current experiments, e.g. LHC
  - Too late for LEP ...

# Challenges in data preservation in High Energy Physics (HEP)

- **Chapter I**: Data preservation: what it is
- **Chapter II**: Aspects of data preservation
- **Chapter III**: Lessons learned from LEP
- **Chapter IV**: Summary and conclusions

# What was LEP again ?

- $e^+e^-$ collider at CERN
  - 1989-2000
  - 90-209 GeV
- 4 experiments:
  - ALEPH
  - DELPHI
  - L3
  - OPAL
- **Largest circular lepton collider so far**

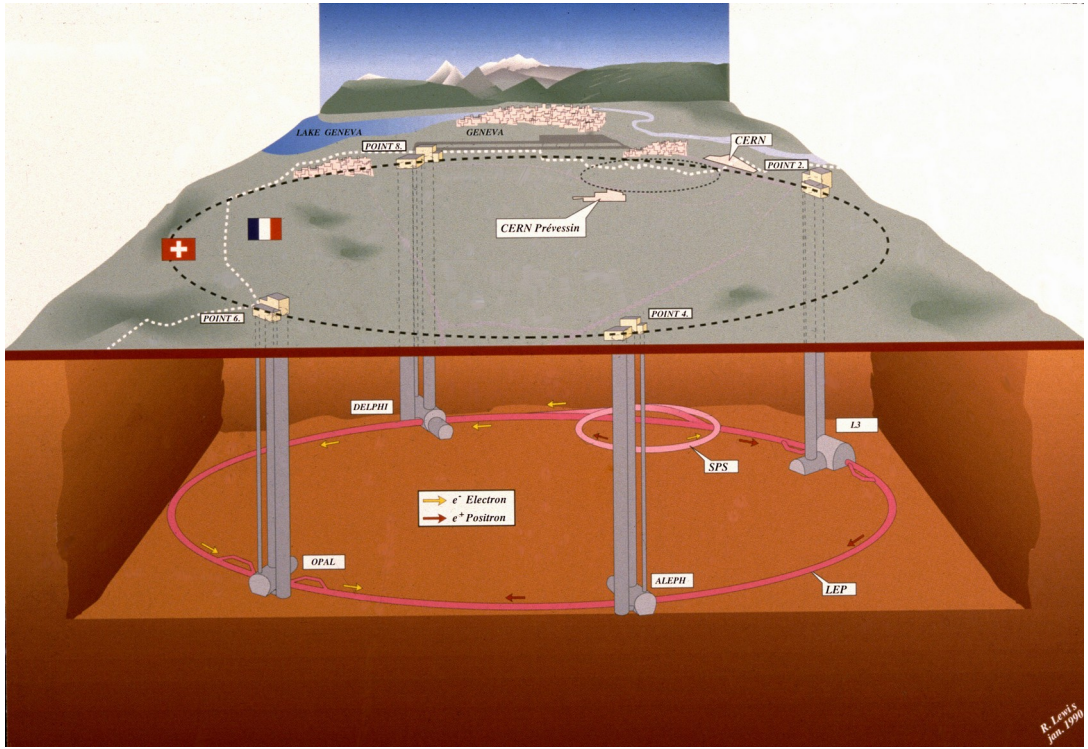


*Courtesy:CERN*

# LEP: some impressions

First Z0 event,
OPAL 1989

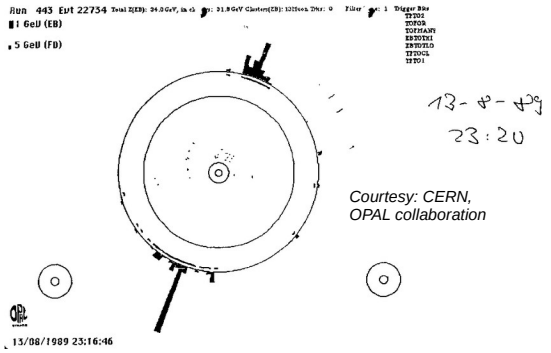


*Courtesy: CERN, OPAL collaboration*

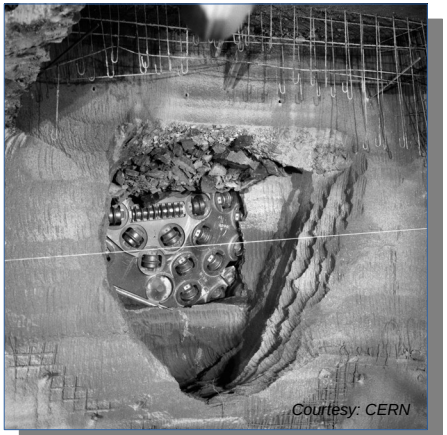

*Courtesy: CERN*

LEP Tunnel
break through,
1986

*Courtesy: W. Liebig*

LEP RF galleries

*Courtesy: W. Liebig*

LEP accelerator

*Courtesy: CERN*

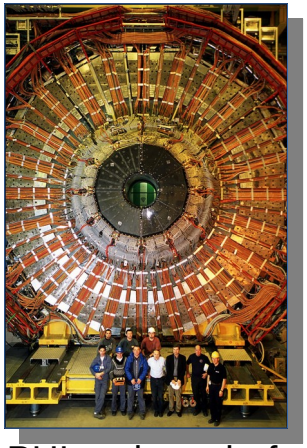

DELPHI end-cap before
dismantling

# LEP: experiment publications and status

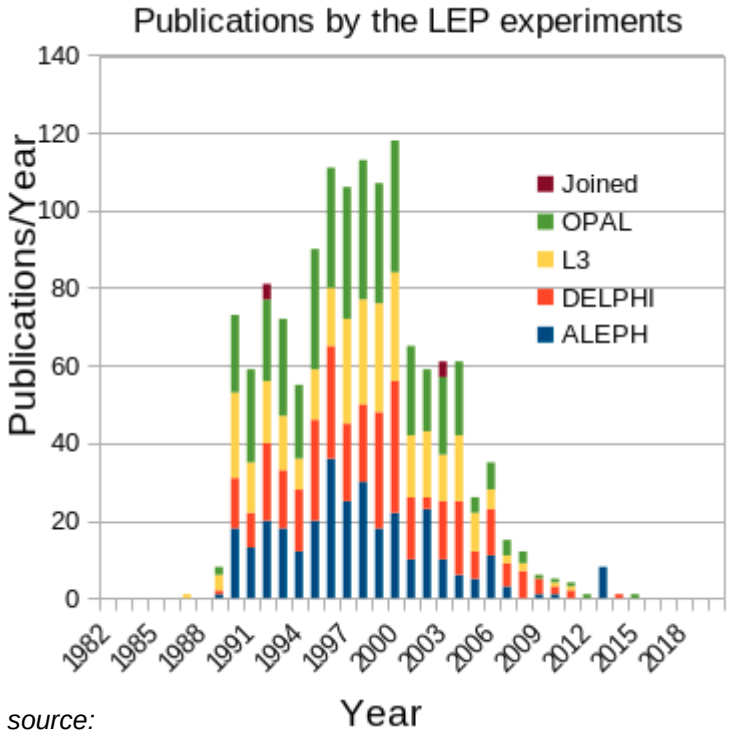


*Data source:*
*https://inspirehep.net*

- Lots of publications after the end of data taking
  - Shown are publications only
  - Many more notes and 3rd party articles up to recently
- 3 still active
  - Several papers in internal review
  - Partly works by external people
  - E.g. Phys. Rev. Lett. 123, 212002 by ALEPH

*Courtesy: CERN ALEPH, DELPHI, OPAL and L3 collaborations*

# Interlude: Glossary of terms

- **C**ERN **A**dvance **STOR**age Manager
  - CASTOR home page
  - Hierarchical storage (disk and tape)
- **C**ERN **T**ape **A**rchive
  - CTA home page
  - Tape based, replaces CASTOR
- **E**OS **O**pen **S**torage
  - EOS home page
  - Disk based, Kerberos secured mass storage system
- **CernVM**
  - CernVM home page
  - Virtual software appliance for LHC
- **CernVM-FS**
  - CVMFS home page
  - Scalable, reliable, low-maintenance software distribution system
- **A**ndrew **F**ile **S**ystem
  - OpenAFS home Page
  - Global file system with strong authentication
  - Kerberos based

*Courtesy: CERN ALEPH, DEPHI, OPAL and L3 collaborations*

# LEP: data access policies

Different for **each experiment**

- Full data access is still handled relatively **restrictive**
- Requires access to CERN and approval from the experiments
- In some cases requires an experienced senior person from the experiment to assist in the analysis

Still **possible to access the data** for external people

https://dphep.web.cern.ch/accelerators/large-electron-positron-collider-lep

# LEP: bit preservation status

- LEP data on CASTOR at CERN
  - Migrated recently to the new system, CTA
  - **2 copies on tape**
  - Same for all 4 experiments
- Additional copies on EOS at CERN
  - **Direct file based access**, e.g. from batch or interactive services at CERN
  - Full data set access usually protected
  - Some simplified data sets available on EOS, 4 vectors (e.g. ALEPH)
- **External copies** exist in some cases
  - E.g. DELPHI in Santander (Spain), for Opal in Munich
- **Open data initiative** for educational purpose,
  - Initiatives in ALEPH and DELPHI (early days there)

# LEP: software preservation status

- Source code
  - Gitlab at CERN (DELPHI)
  - CVMFS (ALEPH, DELPHI, OPAL)
  - CERN VM (DELPHI, OPAL)
- Binaries
  - CVMFS (ALEPH, DELPHI, OPAL)
- Computing environment
  - Originally on AFS
  - CERNVM mounting CVMFS (DELPHI, OPAL)
  - VM images (ALEPH, ...)
  - Archived PCs (inofficial, DELPHI)

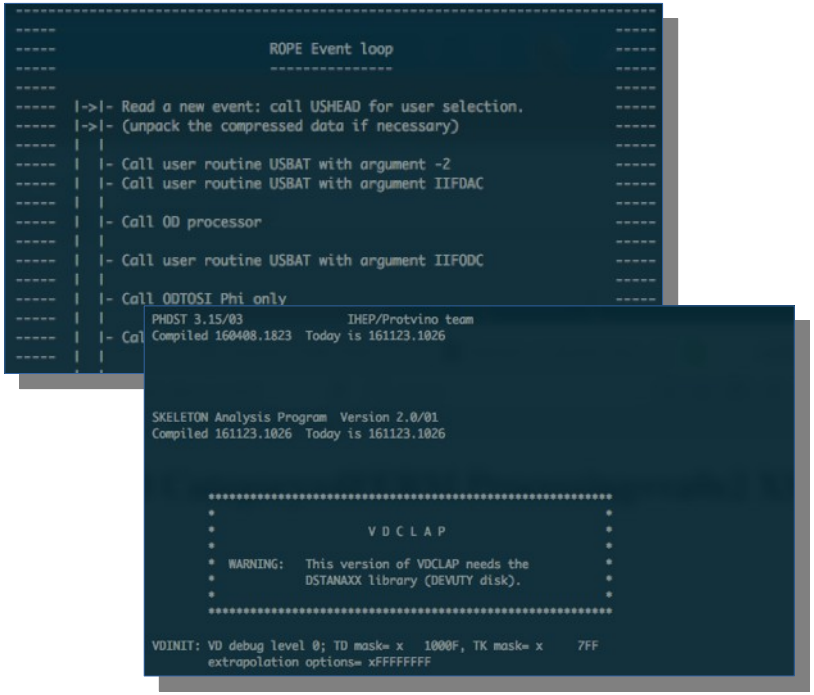

# LEP: knowledge preservation status

- Documentation - Manuals
  - CERN Document server (https://cds.cern.ch)
  - CVMFS (DELPHI, OPAL)
- Sample code, examples
  - CVMFS (ALEPH, DELPHI, OPAL)
- Papers, internal notes, theses, photos…
  - CDS (https://cds.cern.ch)
  - Inspire (https://inspirehep.cern.ch)

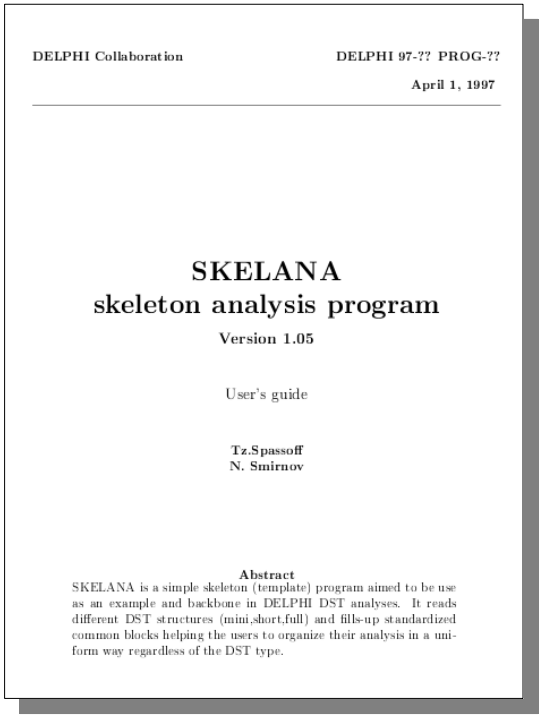

DELPHI Collaboration DELPHI 97-?? PROG-??

April 1, 1997

**SKELANA**
**skeleton analysis program**

Version 1.05

User's guide

Tz.Spassoff
N. Smirnov

**Abstract**

SKELANA is a simple skeleton (template) program aimed to be use as an example and backbone in DELPHI DST analyses. It reads different DST structures (mini,short,full) and fills-up standardized common blocks helping the users to organize their analysis in a uniform way regardless of the DST type.

# LEP: knowledge preservation status

- Early internal notes and technical documentation only on paper
  - No electronic copies available at the time
  - Archived e.g. in the CERN library
- Recovery attempt started in 2020 in DELPHI
  - About 2000 documents processed and checked
  - Roughly 30000 pages in 1175 documents were recovered and imported into CDS at CERN

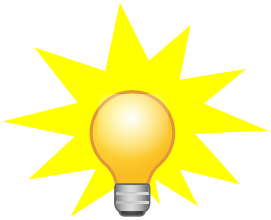

# LEP: Analysis preservation

**No** general analysis preservation was done for LEP

- Some code preserved, along with their output (ALEPH, DELPHI, OPAL)
- **Missing documentation** and meta data on how to interpret these sets

# LEP: Upcoming challenges

- **32bit support** being abandoned
  - At LEP time, Linux based computers were running in 32bit mode only (and usually single-core)
  - Nowadays, everything is 64bit (apart from some older gaming applications)
  - Example: **Motif 32bit libraries gone** from Ubuntu 20.04 and newer
- CERNLIB status
  - Several 64bit ports and forks of CERNLIB exist
  - Missing **validated 64bit** port
  - Most LEP experiments rely on this
  - A working group being established to address this
- Then the experiment code needs to be
  - **ported to 64bit**
  - **and validated**

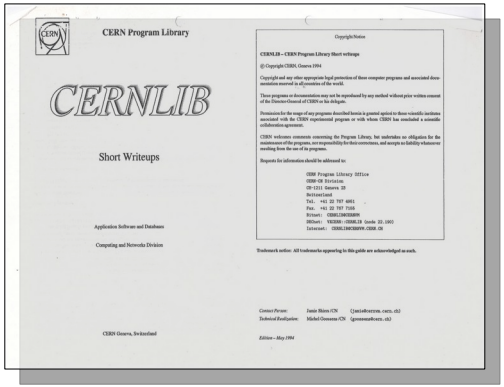
CERN Program Library

CERNLIB

Short Writeups

# Lessons to be learned from LEP


- General status
  - ALEPH, DELPHI and OPAL still in a fairly good shape, situation for L3 is a bit unclear
  - No strict analysis preservation has been done
- Avoid commercial and closed source software
  - DELPHI and OPAL used a commercial 3d visualisation tool called GPHIGS
  - New versions to adapt to new glibc versions cost money
  - Licenses are expensive and need long term funding. Need a dedicated infrastructure to allow users to run it.
- All three still active experiments try to keep the sources alive
  - Gives most flexibility albeit labor intensive

# DPHEP collaboration

## DPHEP collaboration

- 1st workshop **Jan 2009**,
- 140 events in the past since then
- https://indico.cern.ch/category/4458/

https://dphep.org

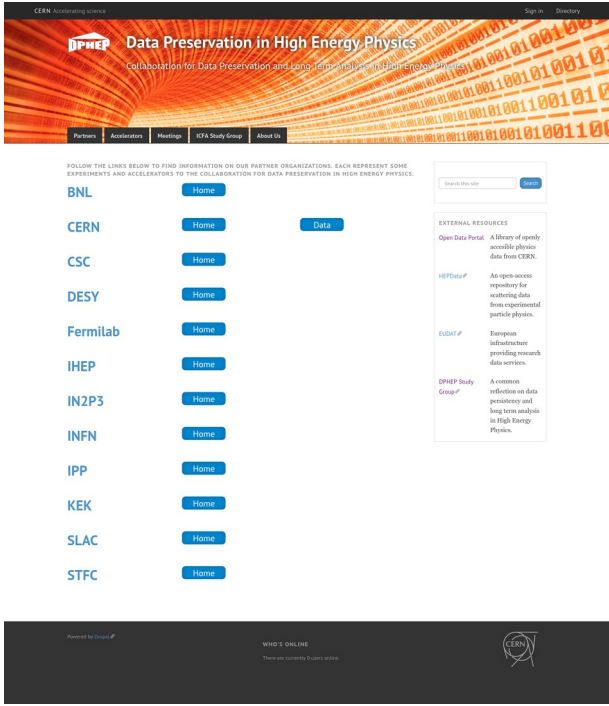

# Lessons learned

## Developments around **access policies:**

- Common **access policy document**, aiming at LHC experiments:
  https://cds.cern.ch/record/2745133/files/CERN-OPEN-2020-013.pdf
  - Defines 4 levels of policies, covering all aspects including **education** and **collaborations EOL**
  - Eventually interesting for other experiments

# Lessons learned

## Developments around **bit preservation**:

- A subset of **data publicly available**,
  - e.g. https://opendata.cern.ch/, https://hepdata.net
  - Produces about 90 PB of data per year
  - Existing data is of EB scale

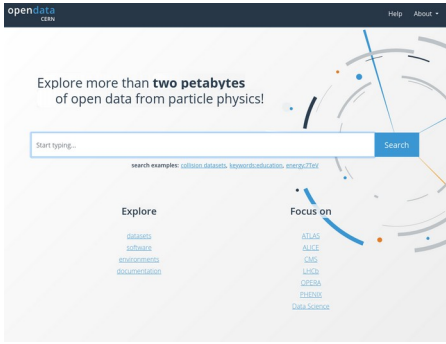


- (Shared) **common data format for archiving** (under discussion)
  - Easier to port things to new technologies
  - Not as flexible as full long term preservation (e.g. MonteCarlo)

# Lessons learned

Tools for **software preservation** and distribution:

- Gitlab, Github
  - Software versioning and change tracing
  - CI for functional tests
- CVMFS (https://cernvm.cern.ch/fs/)
  - Software distribution
- ...

# Lessons learned

## Portals for **knowledge preservation**:

- CERN document server,(https://cds.cern.ch)
  - Manuals, photos, papers, (internal) notes, ...
- INSPIRE, (https://inspirehep.net)
  - “...share and find accurate scholarly information in high energy physics.”
- ZENODO, https://zenodo.org/
  - Not specific to HEP

# Lessons learned


## **Analysis preservation** tools

- Virtual machine images, containers, ...
- **Notebooks**, e.g. Jupyter
- **Reana** initiative https://www.reana.io/
  - Being adopted by several LHC experiments, e.g. CMS
- **CAP**: https://analysispreservation.cern.ch
- Portals like **https://hepdata.net**
  - Published (peer-reviewed) papers
  - High level data and meta data

reana

Reproducible research data analysis platform

CERN
Analysis Preservation

HEPData

# Challenges in data preservation in High Energy Physics (HEP)

- **Chapter I**: Data preservation: what it is
- **Chapter II**: Aspects of data preservation
- **Chapter III**: Lessons learned from LEP
- **Chapter IV**: Summary and conclusions

# Summary and conclusions

- There are many reasons why to preserve data in HEP
- This activity needs to **start while the experiments are still alive** and taking data
  - Early on at the design phase of the software
- Lots of lessons can be learned from past attempts to recover data.
  - **DPHEP** collaboration for exchange of experience
  - Lots of **new tools** available nowadays to help preserving experimental results
  - **Efforts continue after the end of the experiment** to keep data access alive
  - Long term data preservation **is not for free** and needs funding
- Long term data preservation is **not restricted to HEP**.

# Questions ?


www.cern.ch