

# Report from DPHEP Collaboration W/S

[Jamie.Shiers@cern.ch](mailto:Jamie.Shiers@cern.ch)



DPHEP Collaboration Workshop  
CERN, 8/9 June 2015



# Goals

1. *Establish the motivation for long-term data preservation in HEP in terms of succinct Use Cases*
  - *Are there a common set of Use Cases, such as those that were recently agreed for the 4 main LHC experiments but in a more global scope?*
2. *Review the existing areas of "Common Projects"*
  - *Can these be extended (similarly) from their current scope - often LHC - to become more global?*
3. *Perform a site-experiment round-table to capture the current situation HEP-wide*
  - *A summary paper / presentation will be made to PV 2015 (accepted as Oral presentation)*
  - *(We may also wish to submit a paper / talk to iPRES 2016, October 2016 in Bern)*
  - **A longer document – closer to a Blueprint update – is foreseen ~by next workshop – see agenda for outline**

# Themes

- **Collaboration**
  - This is not the first DPHEP workshop, but the first one since the Collaboration Agreement has been signed (7 sites; more coming)
- The move (transition?) to **Open Science**
  - This is happening “externally” – but also matches quite well what is happening inside the project and the experiments

# Use Cases – January GDB

- Preserve data, software, and know-how in the collaborations
  - Foundation for long-term DP strategy
  - Analysis reproducibility: Data preservation alongside software evolution
- Share data and associated software with larger scientific community
  - Additional requirements:
    - Storage, distributed computing
    - Accessibility issues, intellectual property
    - Formalising and simplifying data format and analysis procedure
    - Documentation
- Open access to reduced data set to general public
  - Education and outreach
  - Continuous effort to provide meaningful examples and demonstrations
- Bit preservation
  - Data taken by the experiments should be preserved
- Strategy and scope in approved policy documents for all (LHC) collaborations
  - <http://opendata.cern.ch/collection/data-policies>

# Use Cases – “all HEP”

1. Bit preservation
    - Data taken by the experiments should be preserved
  2. Preserve data, software, and **know-how** in the collaborations
    - Foundation for long-term DP strategy
    - Analysis reproducibility: Data preservation alongside software evolution
  3. Share data and associated software with **(larger)** scientific community
    - Additional requirements:
      - Storage, distributed computing
      - Accessibility issues, intellectual property
      - Formalising and simplifying data format and analysis procedure
      - Documentation
- Open access to reduced data set to general public
    - Education and outreach
    - Continuous effort to provide meaningful examples and demonstrations
  - Strategy and scope in approved policy documents for all **(LHC+LEP)** collaborations
    - <http://opendata.cern.ch/collection/data-policies>
- LEP (and other?) access policies exist (L3?) – need to be uploaded & given DOI**

# Use Cases – “all HEP”

1. Bit preservation – basically ok (at CERN) but not a formal policy
    - Data taken by the experiments should be preserved
  2. Preserve data, software, and know-how in the collaborations
    - Foundation for long-term DP strategy
    - Analysis reproducibility: Data preservation alongside software evolution
  3. Share data and associated software with (larger) scientific community
    - Additional requirements:
      - Storage, distributed computing
      - Accessibility issues, intellectual property
      - Formalising and simplifying data format and analysis procedure
      - Documentation
- Open access to reduced data set to general public
    - Education and outreach
    - Continuous effort to provide meaningful examples and demonstrations
  - Strategy and scope in approved policy documents for all (LHC+LEP) collaborations
    - <http://opendata.cern.ch/collection/data-policies>
- LEP (and other?) access policies exist (L3?) – need to be uploaded & given DOI

# CAP Use Cases (I) (=know-how?)

1. The person having done (part of) an analysis is leaving the collaboration and has to hand over the know-how to other collaboration members.
2. A newcomer would like join a group working on some physics subject
3. In a large collaboration, it may occur that two (groups of) people work independently on the same subject
4. There is a conflict between results of two collaborations on the same subject

# CAP Use Cases (II)

5. A previous analysis has to be repeated
6. Data from several experiments, on the same physics subject, have to be statistically combined
7. A working group or management member within a collaboration wishes to know who else has worked on a particular dataset, software piece or MC
8. Presentation or publication is submitted for internal/collaboration review and approval: lack of comprehensive metadata
9. Preparing for Open Data Sharing



# Knowledge capture – beyond the grave

- No-one(?) believes that this is possible today
- But LHC (&FCC) experiments have to solve “the succession problem”
- The above may be enough whilst the collaborations exist (+ a bit longer)
- This goes way beyond the ambition of many past experiments (targetting ~2020)


# Joint Projects

- Bit Preservation
- Virtualisation (CernVM + CernVMFS)
- Analysis Capture & Preservation
- Open Access, including Discoverability etc
- Open Data → Open Science
- **Concern: how to “drive” such projects between DPHEP workshops – particularly involving the broader DPHEP Collaboration (community)**

# Large scale media migration

- Challenge:

- ~85 PB of data
- 2013: ~51 000 tapes
- 2015: ~17 000 tapes
- Verify all data after write
  - 3x (255PB!) pumped through the infrastructure (read->write->read)
- Liberate library slots for new cartridges
  - Decommission ~35 000 obsolete tape cartridges



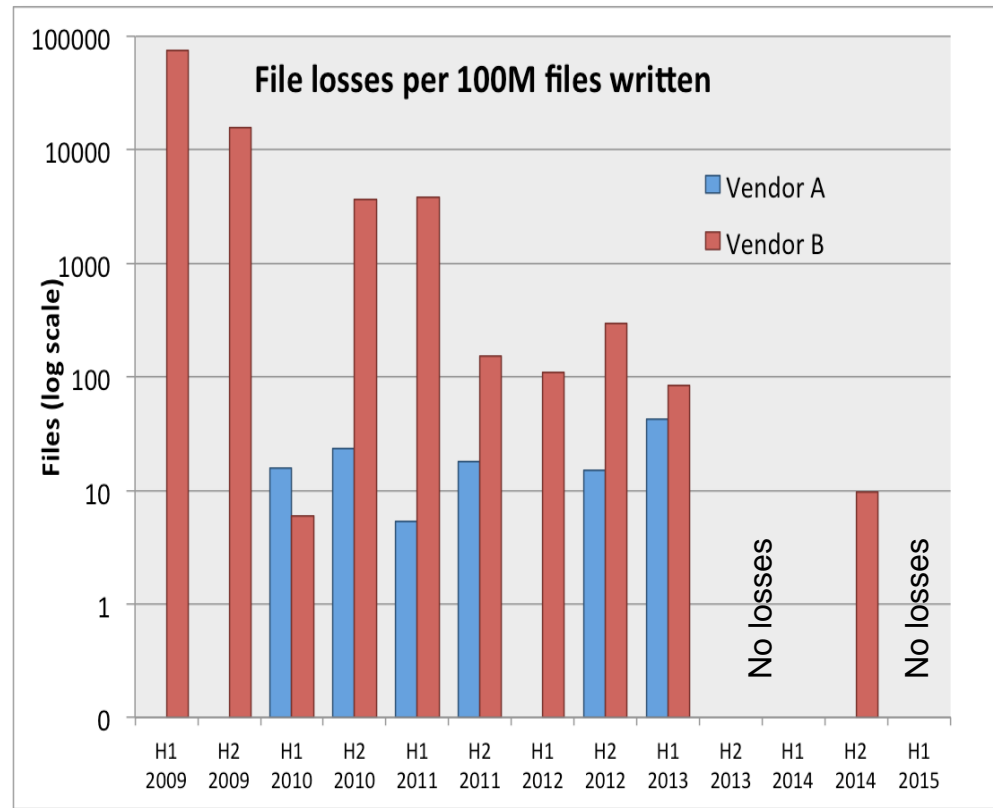
- Constraints:

- Be transparent for user/experiment activities
- Preserve temporal collocation
- Finish before LHC run 2 start


# CERN Archive Reliability

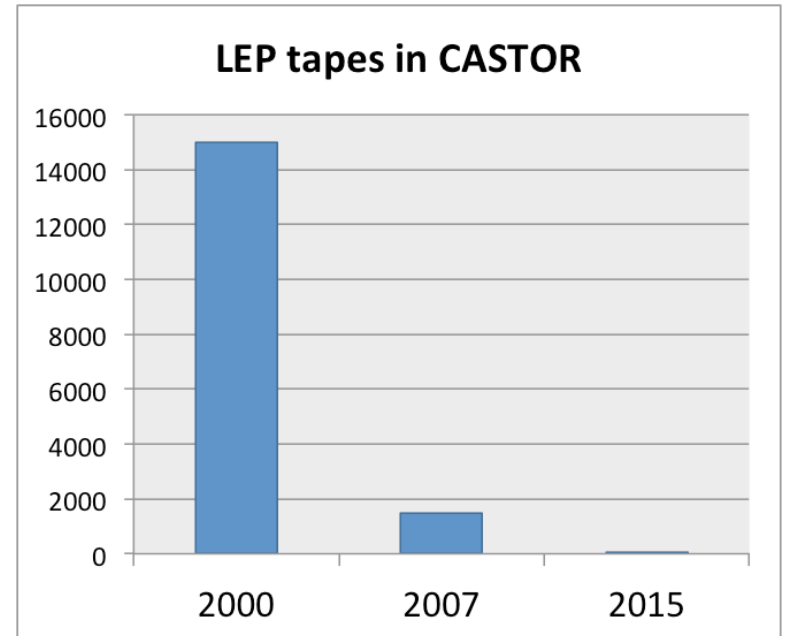
## Ongoing activity to improve archive reliability

- Continued systematic verification of freshly written + “cold” tapes
- Less physical strain on tapes (HSM access, buffered tape marks)
- With new hardware/media, differences between vendors getting small
- For smaller experiments, created dual copies on separated libraries / buildings



# ... and the past

- LEP-era data: ~370TB
- 2000:
  - ~ 15'000 tapes
- 2007:
  - ~ 1500 tapes
- 2015:
  - 30 tapes... x 2 (replicated in separate buildings)
  - Cost: 
- LEP **data**: 2 (3?) copies at CERN + outside



# CERN Archive current numbers

## Data:

- ~105 PB physics data (CASTOR)
- ~7 PB backup (TSM)

## Tape libraries:

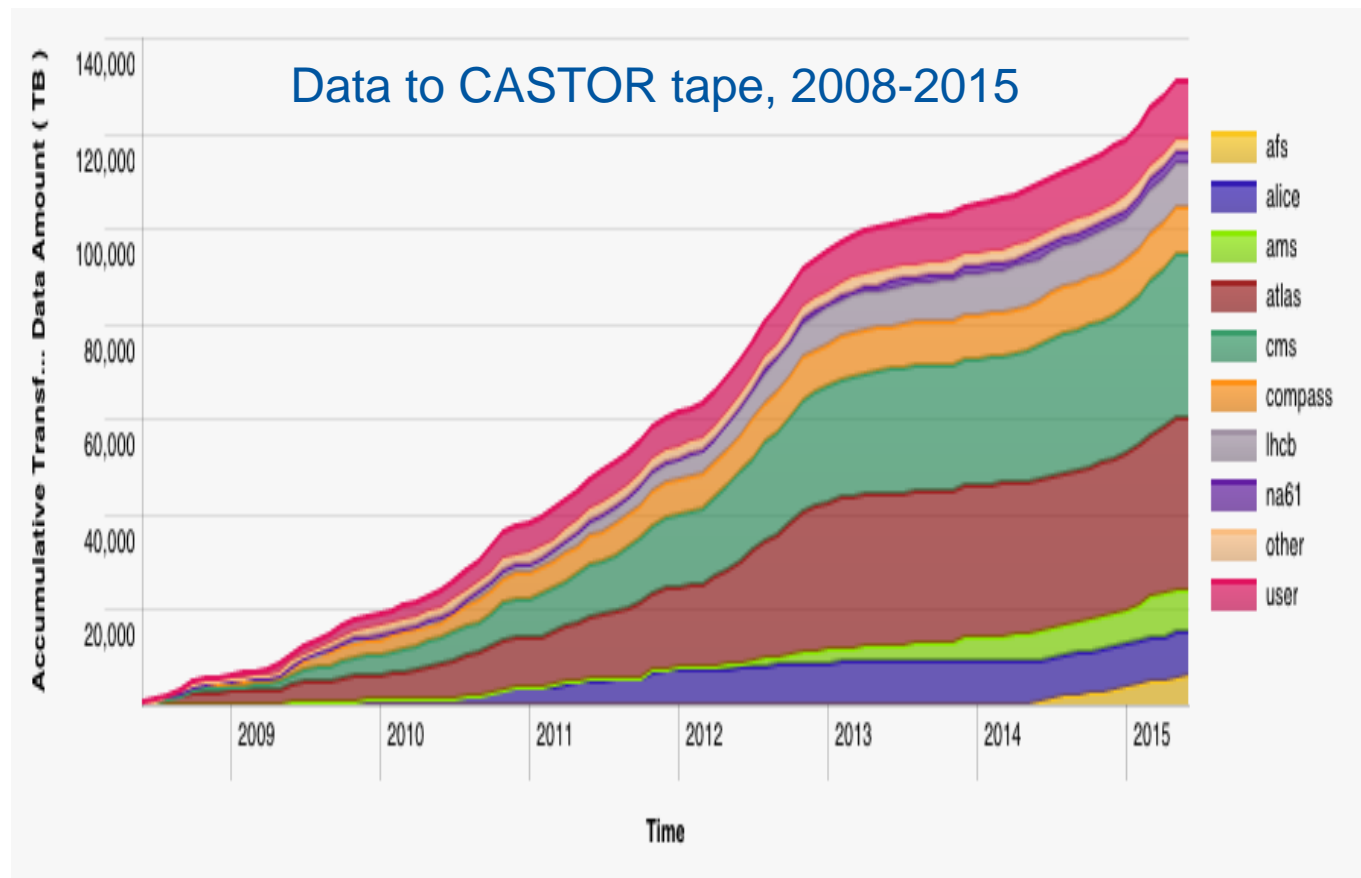
- IBM TS3500 (3+2)
- Oracle SL8500 (4)

## Tape drives:

- ~100 archive

## Capacity:

- ~70 000 slots
- ~25 000 tapes



# CernVM[FS] – Summary


- CernVM ecosystem: natural potential for software environment preservation
  - CernVM-FS embedded versioning of SW and OS
- The *bootloader* technology has simplified the way we can regenerate the target VM
  - Demonstrated with ALEPH and CMS Open Data Pilot
- Investigating the use of container technology
  - Potential to add smoothness to the whole process
  - Promising tests with ALEPH/SLC4
- On the way for a full provenance system

# CernVMFS – “Request”

- Store past experiment s/w (including non-CERN, e.g. JADE) in CernVMFS
- IMHO & CERNLIB too!
- (Plus “archiving” of CERNLIB documentation in a Digital Library)
- Web copy of CERNLIB doc “unavailable” (again & again & again...)



# Projects and PP estimates In the released document Table 8



	Project	Goals and deliverables	Resources and timelines	Location, possible funding source, DPHEP allocation
<b>Experiment and laboratory</b> Priority: 1	Experimental Data Preservation Task Force	Install an experiment data preservation task force to define and implement data preservation goals.	1 FTE installed as soon as possible, and included in upgrade projects	Located within each computing team. Experiment funding agencies or host laboratories. DPHEP contact ensured, not necessarily as a displayed FTE.
	Facility or Laboratory Data Preservation Projects	Data archivist for facility, part of the R&D team or in charge with the running preservation system and designed as contact person for DPHEP.	1-2 FTE per laboratory, installed as a common resource.	Experiment common person-power, support by the host labs or by the funding agencies as a part of the on going experimental programme. A fraction 0.2 FTE allocated to DPHEP for technical support and overall organisation.
<b>Multi-experiment</b> Priority: 3	General validation framework	Provide a common framework for HEP software validation, leading to a common repository for experiments software. Deployment on grid and contingency with LHC computing also part of the goals.	1 FTE	Installed in DESY, as present host of the corresponding initiative. Funding from common projects. Cooperation with upgrades at LHC can be envisaged. Part of DPHEP.
	Archival systems	Install secured data storage units able to maintain complex data in a functional form over long period of time without intensive usage.	0.5 FTE	Multi-lab project, cooperation with industry possible. Included in DPHEP person-power.
	Virtual dedicated analysis farms	Provide a design for exporting regular analysis on farms to closed virtual farm able to ingest frozen analysis systems for a 5-10 years lifetime.	1 FTE	The host of this working group should be SLAC. Funding could come from central projects and can be considered as part of DPHEP.
	RECAST contact	Ensure contact with projects aiming at defining interfaces between high-level data and theory.	0.5 FTE	Installed with proximity to the LHC, the main consumer of this initiative, with strong connections to the data preservation initiatives that may adopt the paradigms.
	High level objects and INSPIRE	Extend INSPIRE service to documentation and high-level data object.	0.5-1.5 FTE	Installed at one of the INSPIRE partner laboratories.
	Outreach	Install a multi-experiment project on outreach using preserved data, define common formats for outreach and connect to the existing events.	1 FTE central + 0.2 FTE per experiment	A coordinating role can be played by DPHEP in connection with a large outreach project existing at CERN, DESY or FNAL. The outreach contributions from experiments and laboratories can be partially allocated to the common HEP data outreach project and steered by DPHEP.
	<b>Global</b> Priority: 2	DPHEP Organisation	DPHEP Project Manager	1 FTE

Table 8: Resources required by projects of the DPHEP study group.

# A note on the physics case

- Many hypothesis and concrete examples discussed in the past workshops: re-analysis, re-cast, combinations etc.
- Did all this continue to happen?
  - Do we have now continued evidence for the physics case of preserved data?
- Did data preservation initiatives within experiments played any role in enhancing the physics output flow for the ending experiments?
- Does it play any role in the running of the present experiments?
- ... and in the planning for new experiments?

# Site / Experiment Reports

- Belle I & II
- BES III & others
- HERA (DESY + MPP)
- LEP
- LHC
- CDF + FNAL (Tevatron)
- BaBar
- TRIUMF, SNOLAB etc
  
- **Many similarities in approaches**
  - But not so much that can really be called “common projects”
- **Numerous issues with manpower**
  - Would more coordination / collaboration help?
  - **Experiment man-power is often the issue...**
- **“Its never too early to consider DP”**
  - But this advice is often not followed...
- **Bit preservation is not (should not be) an issue...**
  - Nor preservation of documents that actually exist...
  - Nascent “definitions” of what “knowledge capture” might mean

# Progress(?)

- Its just over 3 years since the DPHEP Blueprint was published and ...
- The DPHEP session @ CHEP 2012 and ...
- The DPHEP input to ESPP update
- **The progress has simply been remarkable**