Long-Term Data Preservation

Jamie.Shiers@cern.ch

DPHEP Project Manager

DPHEP

International Collaboration for Data Preservation and Long Term Analysis in High Energy Physics

Overview

- Data Preservation activities in HEP (DPHEP);
- Significant events in past year relating to Open Access to Data and Long-Term Data Preservation;
- Implications and wish-list for HEPiX community;
- Outlook: EU H2020, NSF and other funding

Nota Bene

- People often confuse “data preservation” with “bit preservation” (we need a better name…)

- This is **part** of the problem but by no means all (and even this is far from trivial and can be (very) costly over time

➢ What is required is the ability to **fully interpret** and **use** the data in the long-term (includes s/w, build systems, meta-data, DCIs…)

- What is long? Different timescales / discipline and per Use Case

- Rule of thumb: “long” means adapting to **inevitable**, and possibly unforeseeable, changes: **nothing** can be safely assumed to be constant!

Reliability

- File loss is unavoidable and needs to be factored in at all stages
- Good news: it has been getting better for both disk and tape
- Disk storage reliability greatly increased by EOS over CASTOR disk
  - RAID-1 does not protect against controller or machine problems, file system corruptions and finger trouble
- Tape reliability still ~O(1) higher than EOS disk
  - Note: single tape copy vs. 2 copies on disk

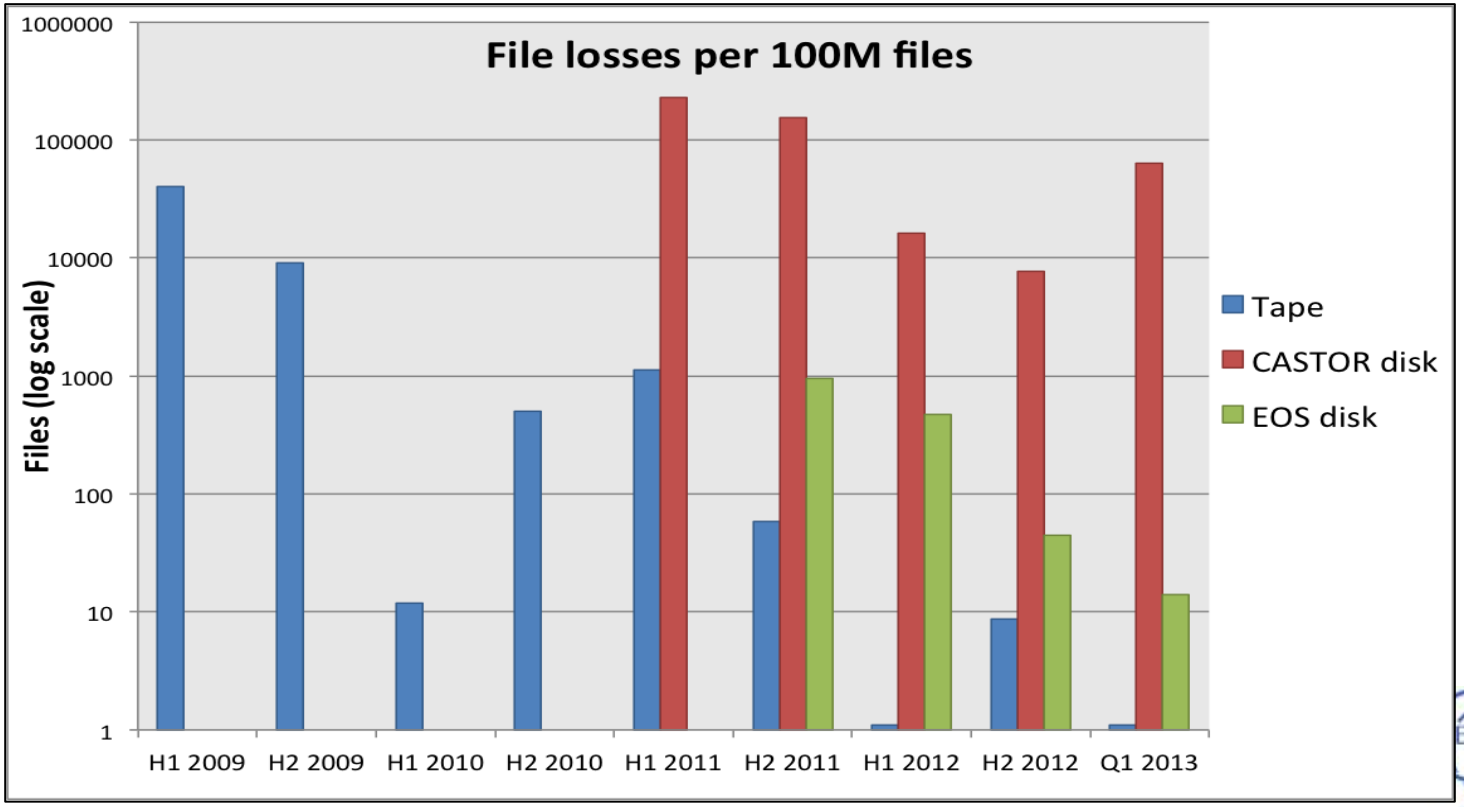

[CERN Tape Reliability Figures]

- On 17th March at 17:30:45 we have cleared the complete tape verification backlog and therefore scanned all full tapes.
- It took 2 years, 7 months and 4 days.

Overall statistics:

- total tapes analyzed : 51545
- total tapes with data: 51519 ; empty tapes: 26
- tapes with [ERROR] : 69
- tapes with [WARN] : 1693
- tapes with [OK] : 49780

- total verified data : 106958.45 TB
- average performance per drive and tape : 138.697 MB/s
- aggregate performance for all drives : 1.26 GB/s

WLCG Service Incident Reports & Repack experience e.g. from KEK also relevant

DPHEP: From Birth to Blueprint

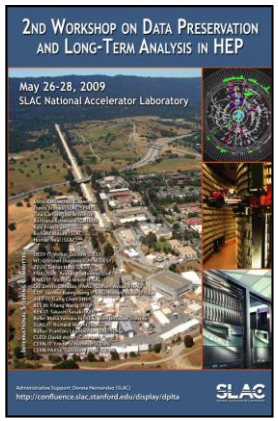

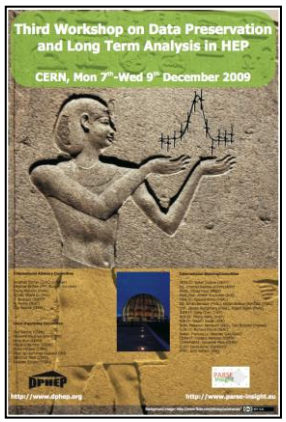

Jan 2009: DESY | May 2009: SLAC | Dec 2009: CERN | Jul 2010: KEK | May 2011: Fermilab | Nov. 2012: Marseille

- Study Group started in January 2009
  - Workshop in DESY
- 2 more workshops in SLAC (May) and CERN (Dec)
  - First recommendations released in December 2009
- Tour of labs continues
  - Workshops in KEK (July 2010) and FNAL (May 2011)
  - Release of the Blueprint in May 2012;
- Priorities:
  - Secure data in all experiments
  - Consolidate the on-going international cooperation: DPHEP “organisation”
  - Promote common multi-experiment projects and/within interdisciplinary cooperation

Slide: Cristinel Diaconu, DPHEP chair

The DPHEP Organisation (aka “Collaboration”)

- Retain the basic structure of the Study Group, with links to the host experiments, labs, funding agencies, ICFA
- **Installation of a full time DPHEP Project Manager, who acts as the main operational coordinator**
- The DPHEP Chair (appointed by ICFA) coordinates the steering committee and represents DPHEP in relations with other bodies

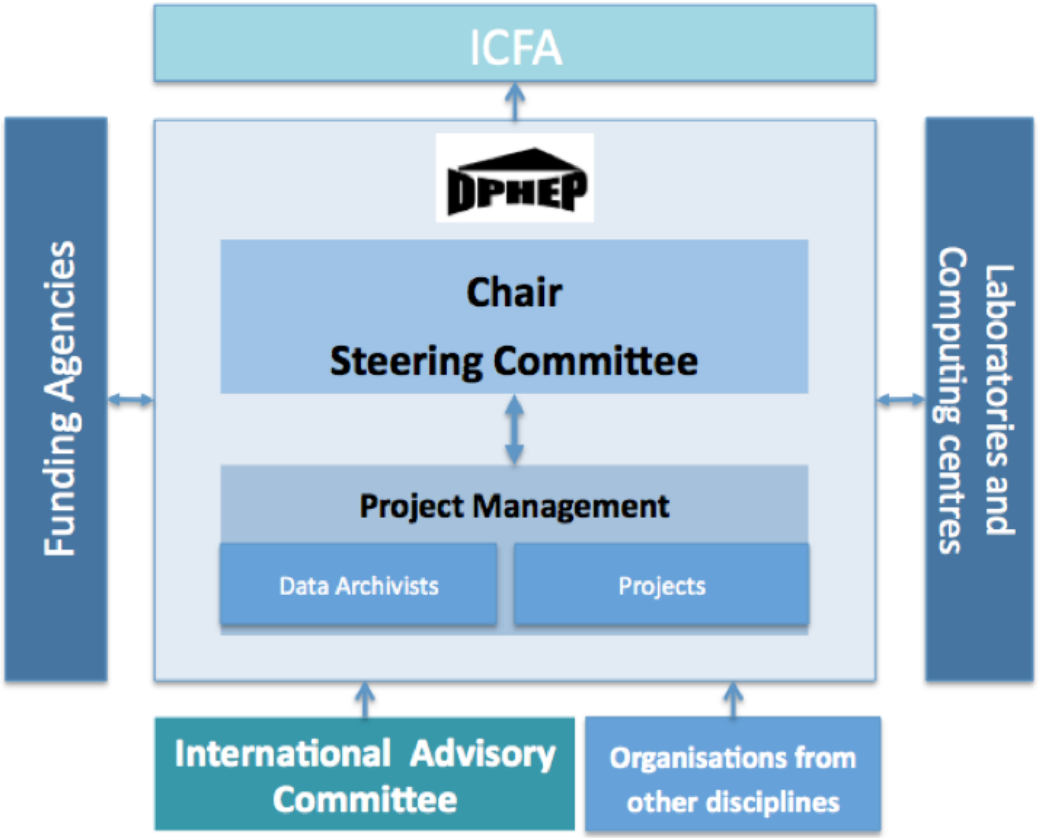

The DPHEP Organisation/Collaboration

- Retain the basic structure of the Study Group, with links to the host experiments, labs, funding agencies, ICFA

- **October, 2012: CERN endorses the blueprint and appoints the DPHEP Project**

- The collaboration agreements are being defined and will become effective in 2013

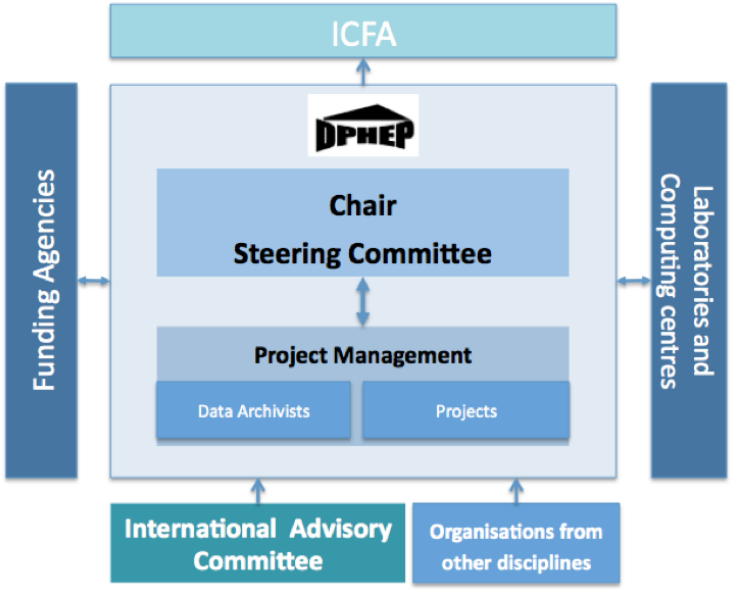

Dear Dr. Diaconu,

Following the delivery of the final DPHEP blueprint, various inputs received into the European Strategy for Particle Physics symposium earlier this week and after consultation with my colleagues, I would like to inform you that CERN offers to provide the role of the initial DPHEP project manager.

We would propose to appoint Jamie Shiers in this role for an initial period of 3 years starting 1 January 2013, after which the role may be assumed by another laboratory, as suggested in the blueprint.

We would anticipate that during this period the DPHEP organization will be launched (year 1) and that the initial deliverables defined in the blueprint would be achieved.

CERN would also foresee participation in the other activities described in the document in areas such as R&D into the use of virtual machine technology for data preservation purposes (PH-SFT input to ESPP) and into the management of very large data stores.

Yours sincerely,

Sergio Bertolucci
Director for Research and Computing

Slide: Cristinel Diaconu, DPHEP chair

From Blueprint to Now

- **A summary of DPHEP Blueprint was submitted to ESPP (September) – data preservation retained in draft update**
- Research Data Alliance launched (US, EU, AUS) – output **will** influence EU Horizon 2020 and other funding lines
- EU Commission Recommendation on "*access to and preservation of scientific information*" (July)
- EU Coordination Workshop on above (October)
- e-IRG workshop on Data (December)
- Collaboration with numerous other disciplines strengthened significantly (e-IRG, EUDAT, RDA, APA, …)
- CERN Director for Research and Scientific Computing proposes to provide DPHEP Project Leader 2013 – 2015
- RDA kickoff + EU / EIROforum workshop (March)

➢ **DP is one of the areas targetted in H2020 (also by NSF) (see backup)**

2020 Vision for LT DP in HEP

- ***N.B. Long-term – e.g. LC timescales****: disruptive change(s)*
  - By 2020, all archived data – e.g. that described in Blueprint, including LHC data – easily findable, fully usable by designated communities with clear (Open) access policies and possibilities to annotate further
  - Best practices, tools and services well run-in, fully documented and sustainable; built in common with other disciplines, based on standards
- Many of the tools / services involved are those with which HEPiX is concerned
- Opportunity for improving coordination across sites and / or considering the impact of an important Use Case

DPHEP – S.W.O.T.

Strengths	DPHEP is **well established** within the community and recent contacts to other disciplines are **very encouraging**
Weaknesses	Effort is very **scarce** within the project at a time when manpower is already stretched to the limit elsewhere
Opportunities	Through a convergence of events there are clear **possibilities** for significant funding and collaboration in the EU's Horizon 2020 programme and most likely corresponding programmes in other areas of the world, e.g. NSF-funded projects
Threats	**Failure** to invest now would jeopardise attempts to "rescue" LEP data as well as to take other preservation events (BaBar, Tevatron, Hera etc.) to a stable and sustainable state. It could also limit our ability to prepare for – and hence participate in – future projects

Recommendations to ICFA

1. **Adopt** the OAIS model across HEP Long-Term Data Preservation Projects;
2. Actively **participate** in the Research Data Alliance and its Working Groups with the intent to influence EU, US (NSF) and other funding agencies;
3. **Build** on existing contacts with other disciplines to share best practises and where possible also tools and services;
4. Ensure that there are **sufficient resources** (2013 / 2014) to allow the vision of Long-Term, Sustainable Archives to be realized;
5. Address true **long-term** preservation by R&D into handling **change** in all areas of the archive environment and in particular that of the software and offline infrastructure. **Long-term commitment** also required.

ICFA Statement on LTDP

- *The International Committee for Future Accelerators (ICFA) supports the efforts of the Data Preservation in High Energy Physics (DPHEP) study group on long-term data preservation and welcomes its transition to an active international collaboration with a full-time project manager. It encourages laboratories, institutes and experiments to review the draft DPHEP Collaboration Agreement with a view to joining by mid- to late-2013.*

- *ICFA notes the lack of effort available to pursue these activities in the short-term and the possible consequences on data preservation in the medium to long-term. We further note the opportunities in this area for international collaboration with other disciplines and encourage the DPHEP Collaboration to vigorously pursue its activities. In particular, the effort required to prepare project proposals must be prioritized, in addition to supporting on-going data preservation activities.*

- *ICFA notes the important benefits of long-term data preservation to exploit the full scientific potential of the, often unique, datasets. This potential includes not only future scientific publications but also educational outreach purposes, and the Open Access policies emerging from the funding agencies.*

- 15 March 2013

After the collisions have stopped

- Finish the analyses! But then what do you do with the data?
  - Until recently, there was no clear policy on this in the HEP community
  - It's possible that older HEP experiments have in fact simply lost the data
- Data preservation, including long term access, is generally not part of the planning, software design or budget of an experiment
  - So far, HEP data preservation initiatives have been in the main not planned by the original collaborations, but rather the effort a few knowledgeable people

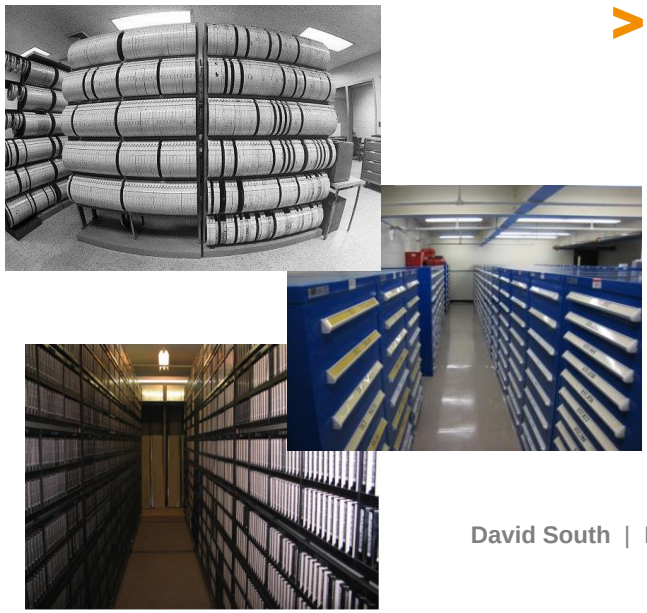

- The conservation of tapes is not equivalent to data preservation!
  - *"We cannot ensure data is stored in file formats appropriate for long term preservation"*
  - *"The software for exploiting the data is under the control of the experiments"*
  - *"We are sure most of the data are not easily accessible!"*

The difficulties of data preservation in HEP

- Handling HEP data involves large scale traffic, storage and migration
  - The increasing scale of the distribution of HEP data can complicate the task
- Who is responsible? The experiments? The computing centres?
  - Problem of older, unreliable hardware: unreadable tapes after 2-3 years
  - The software for accessing the data is usually under the control of the experiments
- Key resources, both funding and person-power expertise, tend to decrease once the data taking stops
- And a rather key ingredient to all this is: *why do it?*
  - Can the relevant physics cases be made?
  - Who says we want to do this anyway?
  - Is the benefit of all this really worth the cost and effort?

DESY

What is HEP data?

Digital information

The data themselves, volume estimates for preservation data of the order of a few to 10 PB

Other digital sources such as databases to also be considered

Software

Simulation, reconstruction, analysis, user, in addition to any external dependencies

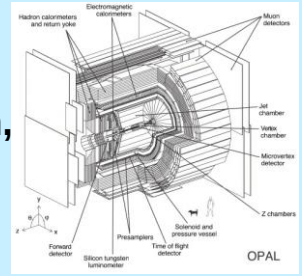

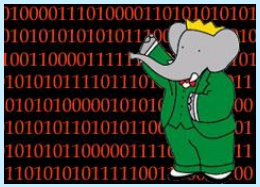

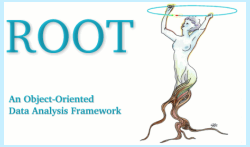

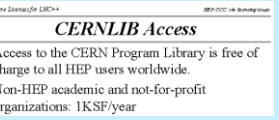

Meta information

Hyper-news, messages, wikis, user forums..

Publications

arXiv.org

INSPIRE

Durham University — Science & Technology Facilities Council

HEPDATA: REACTION DATA Database

...containing numerical values of HEP scattering data such as total and differential cross sections, fragmentation functions, structure functions, and polarisation measurements, from a wide range of experiments. It is compiled by the **Durham Database Group (UK)** with help from the COMPAS group (Russia,) and is updated at regular intervals.

Journal of High Energy Physics

A refereed journal, written, run and distributed by electronic means

Documentation

Internal publications, notes, manuals, slides

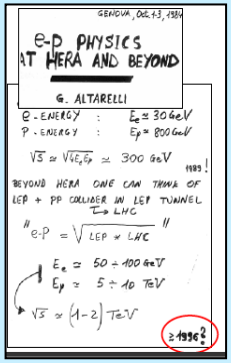

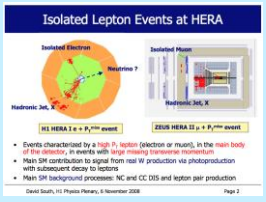

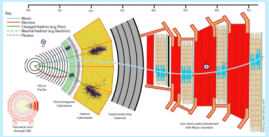

Expertise and people

DESY

DPHEP models of HEP data preservation

Preservation Model		Use Case	
1	Provide additional documentation	Publication related info search	**Documentation**
2	Preserve the data in a simplified format	Outreach, simple training analyses	**Outreach**
3	Preserve the analysis level software and data format	Full scientific analysis, based on the existing reconstruction	**Technical Preservation Projects**
4	Preserve the reconstruction and simulation software as well as the basic level data	Retain the full potential of the experimental data	

- These are the original definitions of DPHEP preservation levels from the 2009 publication
  - Still valid now, although interaction between the levels now better understood
- Originally idea was a progression, an inclusive level structure, but now seen as complementary initiatives
- Three levels representing three areas:
  - Documentation, Outreach and Technical Preservation Projects

DESY

Open Access Policies

- Simplified example from LHCb. (CMS similar)
- Can we harmonize policies:
  - Across experiments? Across labs?

Level-1 data: Published results	**All scientific results are public. …**
Level-2 data: Outreach and education	[Samples made public.] The data are for educational purpose only, not suitable for publication
Level-3 data: Reconstructed data	LHCb will make reconstructed data (DST) available to open public; 50% 5 years after data is taken, 100% after 10 years. Associated software will be available as open source, together with existing documentation. Publications must include disclaimer.
Level 4 data: Raw data	[Not directly accessible to collaboration] But must still be preserved!

Suggested Topics for DPHEP7

- “Ingest Issues” (10’)
  - How did you (the experiment) decide what data to save, how to make it discoverable / available, how is it documented, where is the data / meta-data etc. What are the access policies and target communities?
  - What tools do you use?

- “Archive issues”: (10’)
  - How is the archive managed? How are errors detected and handled? What is the experience?
  - What storage system / services are used?

- “Offline environment issues”: (20’)
  - What have been the key challenges in keeping the offline environment alive? What are the key lessons learned / pitfalls to be avoided? What would you have done differently if long-term preservation had been a goal from the early days of the experiment?

“Conclusions” from DPHEP7

- “Ingest Issues”
  - *How did you (the experiment) decide what data to save?*
  - **The raw data, the latest reprocessing & MC samples and associated calibrations etc.**
  - *How to make it discoverable / available, how is it documented, where is the data / meta-data etc. What are the access policies and target communities?*
  - *What tools do you use?*
  - **Strong overlap in terms of tools -> more coherence & coordination could help both now and in the future.**
  - **IMHO a “DPHEP portal” is called for to provide a single entry point to all of the above.**
  - **This could be achieved by a TECH or FELL working between IT & the experiments**
- “Archive issues”
  - *How is the archive managed? How are errors detected and handled? What is the experience?*
  - *What storage system / services are used?*
  - **Site specific – coordination (& standard interfaces?) could help a lot here: HEPiX?**
- “Offline environment issues”:
  - *What have been the key challenges in keeping the offline environment alive? What are the key lessons learned / pitfalls to be avoided? What would you have done differently if long-term preservation had been a goal from the early days of the experiment?*
  - **An automated test system is likely to be installed for / by the LHC experiments**

Where to Invest – Summary (WLCG OB)

Tools and Services, e.g. Invenio: could be solved. (2-3 years?)

Archival Storage Functionality: should be solved. (i.e. "now")

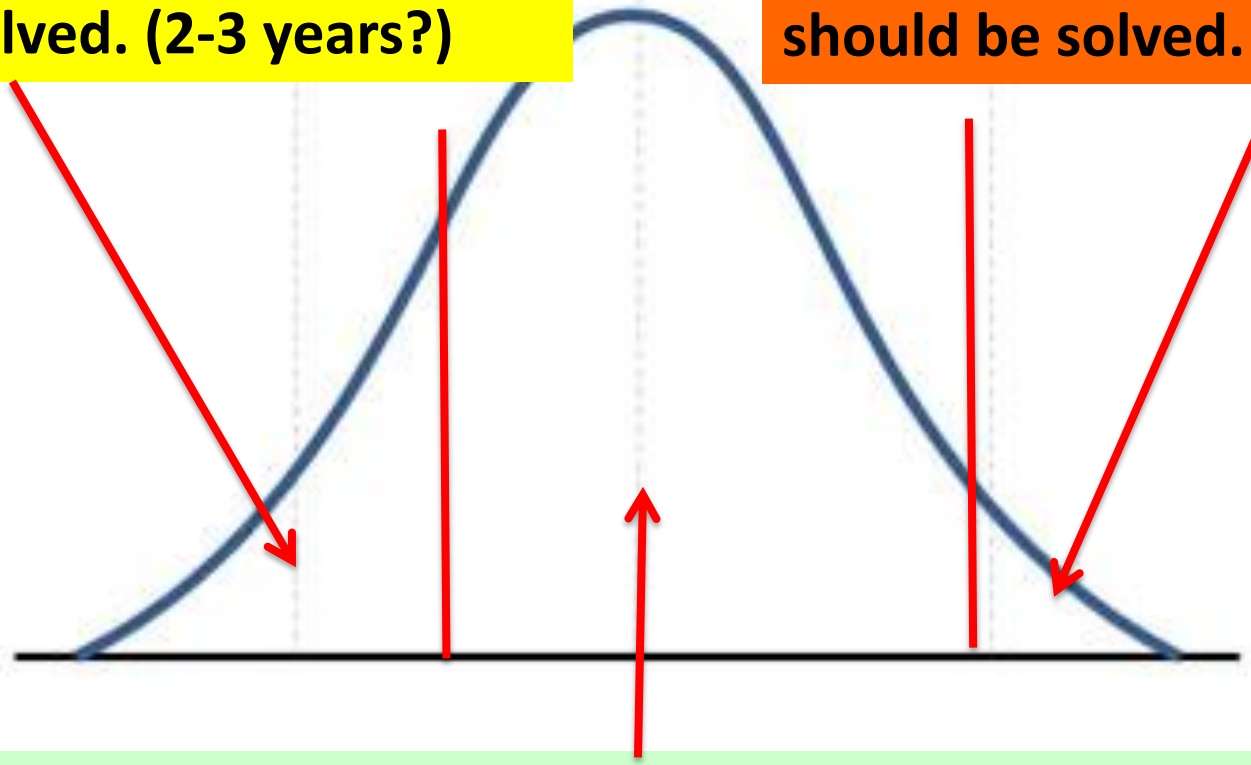

Support to the Experiments for DPHEP Level 4: *must be solved – but how?*

Digital Curation Centre Guidelines

- *Collect as much information as possible about your data at the time of creation and processing, when rich information is available and might be automatically captured*
- *Appraise your data and select what is really worthwhile preserving*
- *Ingest, secure and maintain both the physical data as well as its content, syntax and semantics*
- *Data and tools will need to evolve to keep pace with both IT technological developments, but also scientific demands including data transformation, new analytics, changed descriptions and so on*
- *Plans for data preservation need to be regularly reviewed and updated*

Technical Projects: DPHEP preservation levels 3 and 4

- This is really the main focus of the data preservation effort
  - Level 3: Access to analysis level data, MC and the analysis level software
  - Level 4: Access to reconstruction and simulation software, retain the full capability
- Deciding on level 3 or 4 depends on the scope of your project
  - What do you want to be able to do in *N* years time?
  - Only level 4 gives full flexibility, but this also means not relying on frozen executables and binaries but rather retaining the ability to recompile: more work

The majority of DPHEP experiments aim for DPHEP level 4 preservation

- Remember: it's not about the data, but about still being able analyse it
  - Either keep your current environment alive as long as possible
  - Or adapt and validate your code to future changes as they happen
- Two complimentary approaches taken at SLAC and DESY
  - Both employing virtualisation techniques, but in rather different ways

DESY

The BaBar Long Term Data Access archival system

- New BaBar system installed for analysis until at least 2018
- Isolated from SLAC, and uses virtualisation techniques to preserve an existing, stable and validated platform
- Complete data storage and user environment in one system

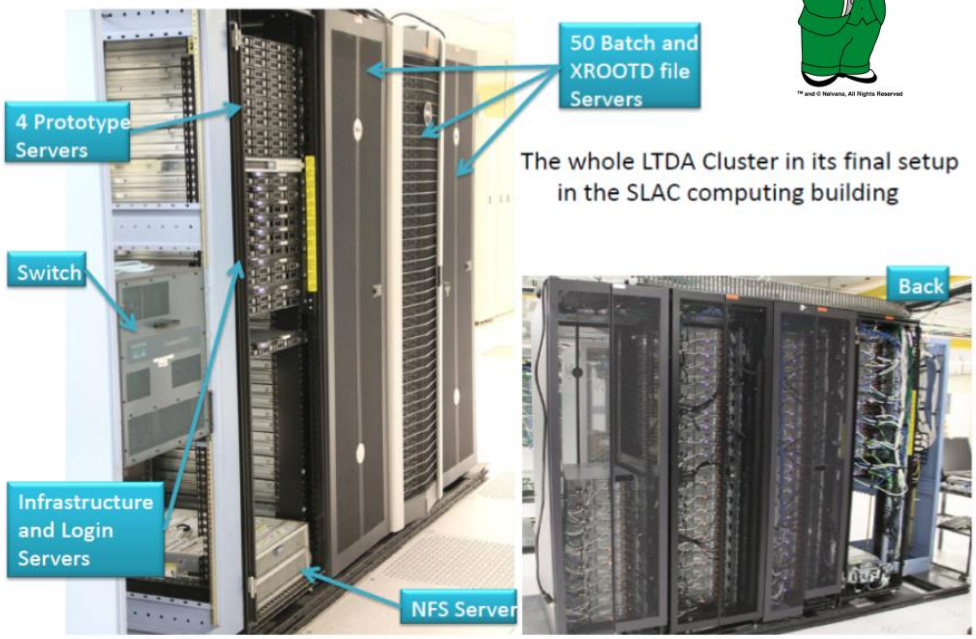

- Required large scale investment: 54 R510 machines, primarily for data storage, as well as 18 other dedicated servers
  - Resources taken into account in experiment's funding model during analysis phase!
- From the user's perspective, very similar to existing BaBar infrastructure

The BaBar Long Term Data Access archival system

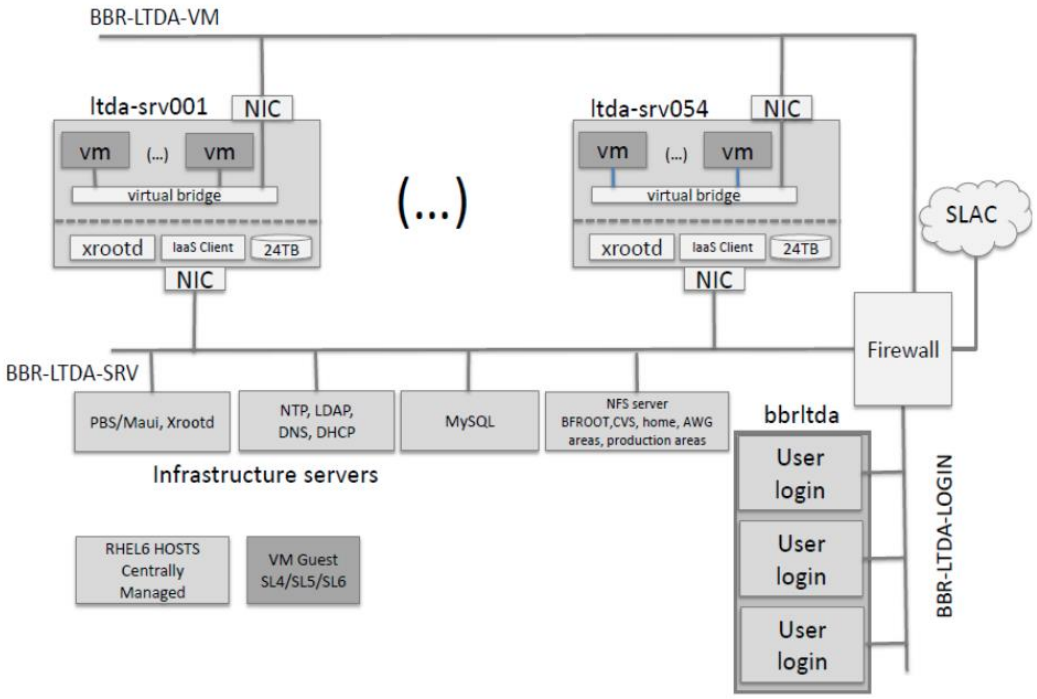

- Crucial part of design is to allow frozen, older platforms to run in a secure computing environment
- *Naïve* virtualisation strategy, not enough
  - Cannot support an OS *forever*
  - Security of system under threat using old versions

- Achieved by clear network separation via firewalls of part storing the data (more modern OS) and part running analysis (the desired older OS)
- Other BaBar infrastructure not included in VMs is taken from common NFS
- More than 20 analyses now using the LTDA system as well as simulation

The `sp-system` at DESY

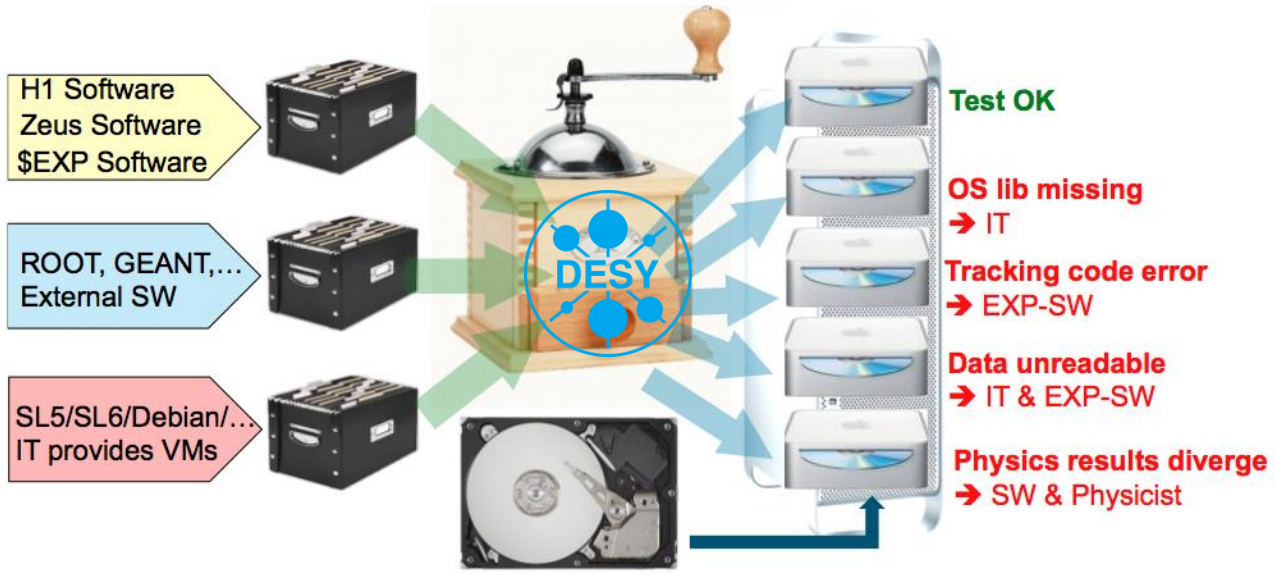

- Automated validation system to facilitate future software and OS transitions
  - Utilisation of virtual machines offers flexibility: OS and software configuration is chosen by experiment controlled parameter file
  - Successfully validated recipe to be deployed on future resource, e.g. Grid or IT cluster
  - Pilot project at CHEP 2010, full implementation now installed at DESY
- Essential to have a robust definition of a complete set of experimental tests
  - Nature and number dependent on desired preservation level
- **Such a system will likely be established for the LHC expts**

Take-Home Messages for HEPiX

- There is significant, if not complete, overlap with the core IT services required for DPHEP & those coordinated by HEPiX
- Can HEPiX expand its activities to include the LTDP Use Case?
  - e.g. coordination of management of long-term archives; inter-site data recovery; long-term commitment to LTDP requirements
  - IMHO, experiments should not be talking about media migration at DPHEP workshops!
- N.B. there are tensions between long-term & short-term needs but these will need to be balanced, in particular for the LHC experiments

DPHEP Collaboration Agreement

- Intent is for main accelerator labs to sign the CA
- This could also be relevant for those sites providing primary / secondary copies of data
  - If not done at a higher level, e.g. that of funding agency

- I have not talked about collaboration with **other** disciplines, nor about H2020 / NSF funding, but these are important issues to be resolved <(<) to the Fall HEPiX
  - E.g. prior to RDA-2 in September

Conclusions

- My top priority for this meeting is to seek agreement / confirmation that management of the data is a long-term site responsibility
  - e.g. checking for and recovering from faulty media
  - Migrating data to new media (repack)
  - (Other actions described in OAIS reference model / ISO std)
- And that the techniques and policies will be coordinated across sites
- This is something valuable that we can bring to future H2020 (and other) projects

➢ **It is not too late to prepare – but it is certainly not too early either!**

Long-Term Data Preservation: It's About Time!

Jamie.Shiers@cern.ch

DPHEP Project Manager

DPHEP

International Collaboration for Data Preservation and Long Term Analysis in High Energy Physics

storing, managing and preserving research data

- Scenario: Diversity is likely to remain a dominant feature of research data – diversity of formats, types, vocabularies, and computational requirements – but also of the people and communities that generate and use the data.
- Europe needs to develop an integrated and service-driven approach to e-infrastructures for the data of a wide range of research communities.
- Europe should therefore step-up the available capacity to cope with extremely large, heterogeneous and complex datasets incorporating advanced computing and software.
- Furthermore, costs of storing and preserving data can be significant if one needs to keep it well managed for long periods to be used by others.
- Different institutions archive their research data in different ways - making access difficult from outside the institution - whereas storage and computing media evolve and become quickly obsolete.
- How will we preserve and maintain future access to priceless research data?
- Data management plans are intended to help researchers, funders and data repositories to get maximum value from research data at minimum cost.

Source: Research Data e-Infrastructures: Framework for Action in H2020