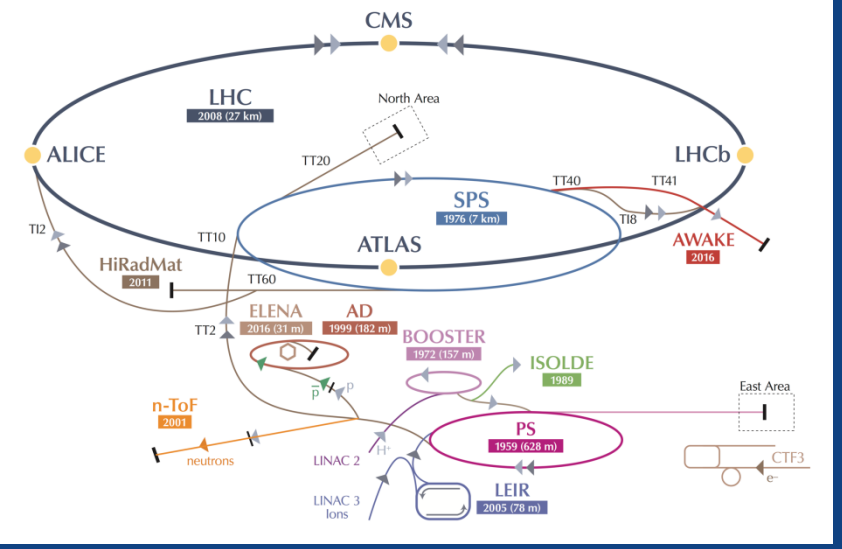

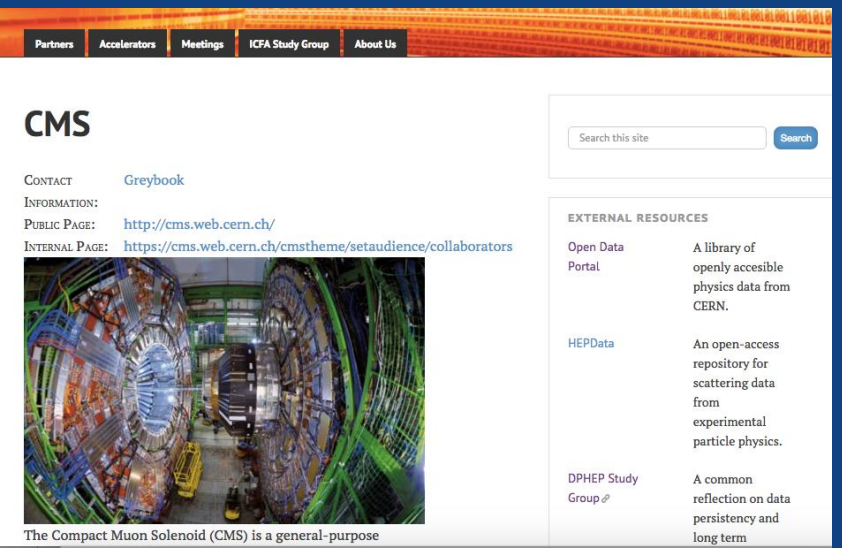

CERN

Jamie.Shiers@cern.ch

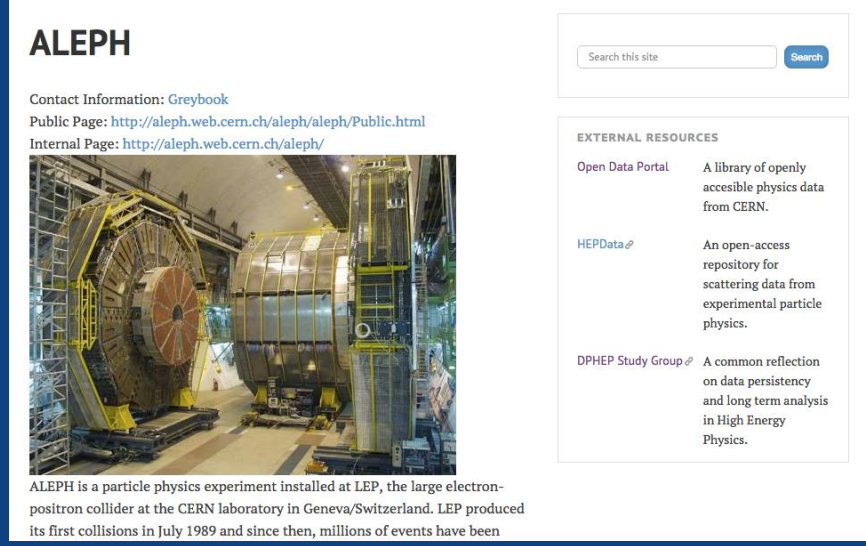

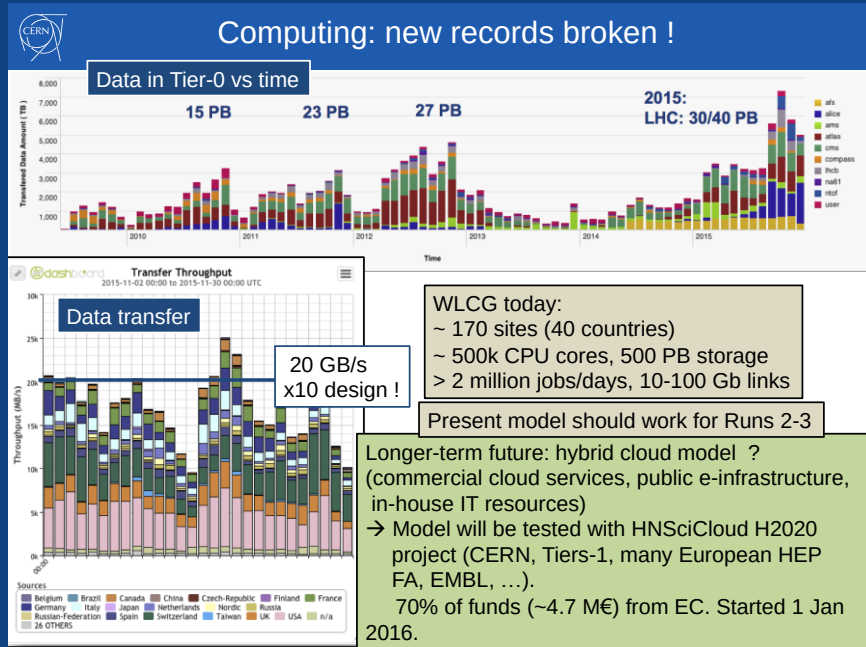

International Collaboration for Data Preservation and
Long Term Analysis in High Energy Physics

Data Preservation for Re-Use

From Tens of TB to Tens of EB … for (several) Tens of Years

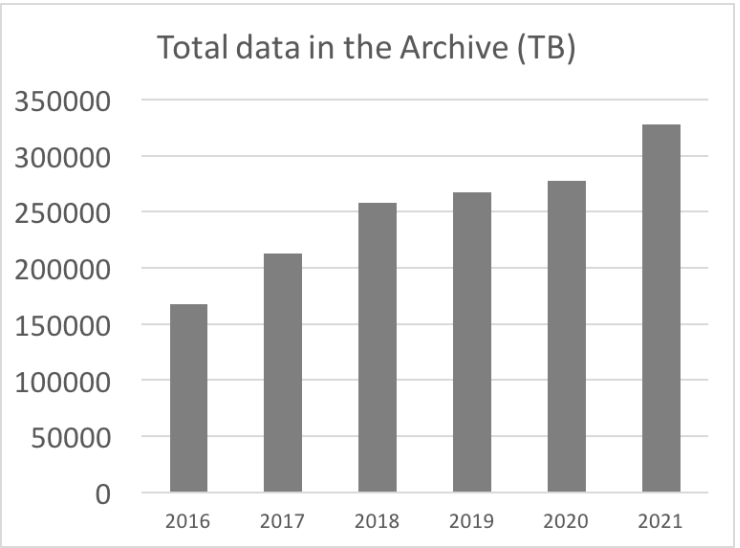

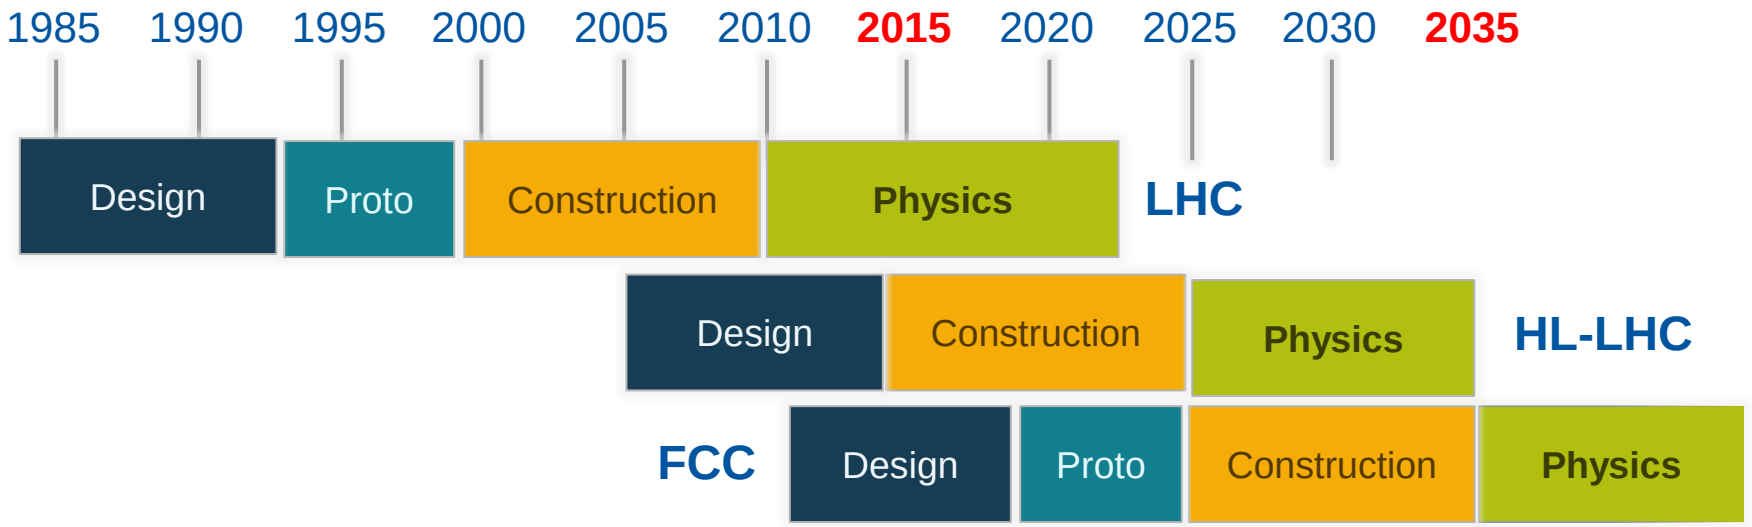

Overview

- This talk will (loosely) follow the H2020 Guidelines for DMPs
- These are complemented by some additional requirements from other FAs (NSF, DoE, …)
- **And – most importantly – by some of our own**
  1. A Common Way of presenting DP Status across the main HEP labs & experiments worldwide
  2. Common Metrics (LHC) to ensure that the data can be shared, re-used and (some) analyses reproduced

- N.B. HEP data are **measurements** – not **observations** – and are only recorded after very de-selective triggers

Some Specific Questions…

- **How do we scale data preservation and data access to address petabyte volumes of data & allow access by users across the world?**

- **How can we preserve material data infrastructures: storage media, software, and the machines to run it?**

- **What are the challenges in data description and curation, and how are different academic disciplines managing this change?**

What is Data?

- H2020 talks about data / datasets
- We (DPHEP) talk about **data**, **software**[1] and **documentation**[2]
  1. **Including also the environment in which to run it**
  2. **Sometimes extended to "knowledge" (later)**
- ➢ **These are the 3 pillars of our DP strategy: distinct services run by different teams**
- **(Very) loosely coupled (which I think is right)**

Documentation – CERNLIB example

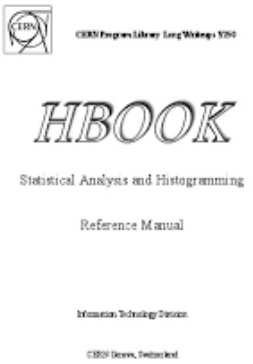

- **Short & Long writeups – last revised in 1995**
  - CERNLIB itself dates back to ~1964 – 1995
- Postscript and HTML (not consistent) still exist but no longer accessible through Web
- Re-formatted in 2015 to consistent PDF/A & HTML, stored in Digital Library (CERN Document Server – Invenio-based)
- **Lifetime of repository: ~decades + forward migration (not always transparent or complete)**
- Together with code & references to papers: probably good for several / many decades (longer than the code itself can be compiled…)

- **LEP offline documentation also being reformatted and captured**
- ➢ **For LHC, documentation constantly used to train new collaboration members and for exploitation of Open Data releases**

(Big) Data (Today – Peanuts Tomorrow)

DPHEP Data Preservation Status	H2020 DMP Guidelines (Annex 1)
• Bit Preservation • Data (volume, storage strategy including replication etc.) • Documentation • Software • Use Cases • Target Communities • Value • Uniqueness • Resources (how funded – sustainability) • Status • Issues • Outlook	Data set reference and name Data set description Standards and Metadata Data Sharing Archiving and Preservation (including storage and backup)
	H2020 Annex 2
	1. Discoverable 2. Accessible 3. Assessable and intelligible 4. Useable beyond the original purpose 5. Interoperable to specific quality standards

New IBM drives and media:

Just In Time for LEP!

The last 200 MB tape – now ~10 TB
in the same form factor

- Tape reliability took **a big jump forward** with the self-enclosed and lightweight IBM 3480 type of cartridge. There were still many drive builders (now just **Oracle** and **IBM**) and media suppliers (now just 1…)

Peta- to Exa-byte Storage

- Multiple copies worldwide – distributed and processed via WLCG
- Primary copy is on tape with active data cached to disk
- ➢ **Regular scrubbing of volumes, migration to new media every ~3-5 years**
- **Cost** per unit time ***goes down*** with time – reliability ***goes up***
  - But not the integrated cost…
- Tapes still have a future (~2030) – beyond that hard to tell

- WLCG "cost model" widely shared – e.g. through EU 4C project

- **How (when) will this evolve to Cloud Storage?**

Case B) increasing archive growth

Start with 10PB, then +50PB/year, then +50% every 3y (or +15% / year)

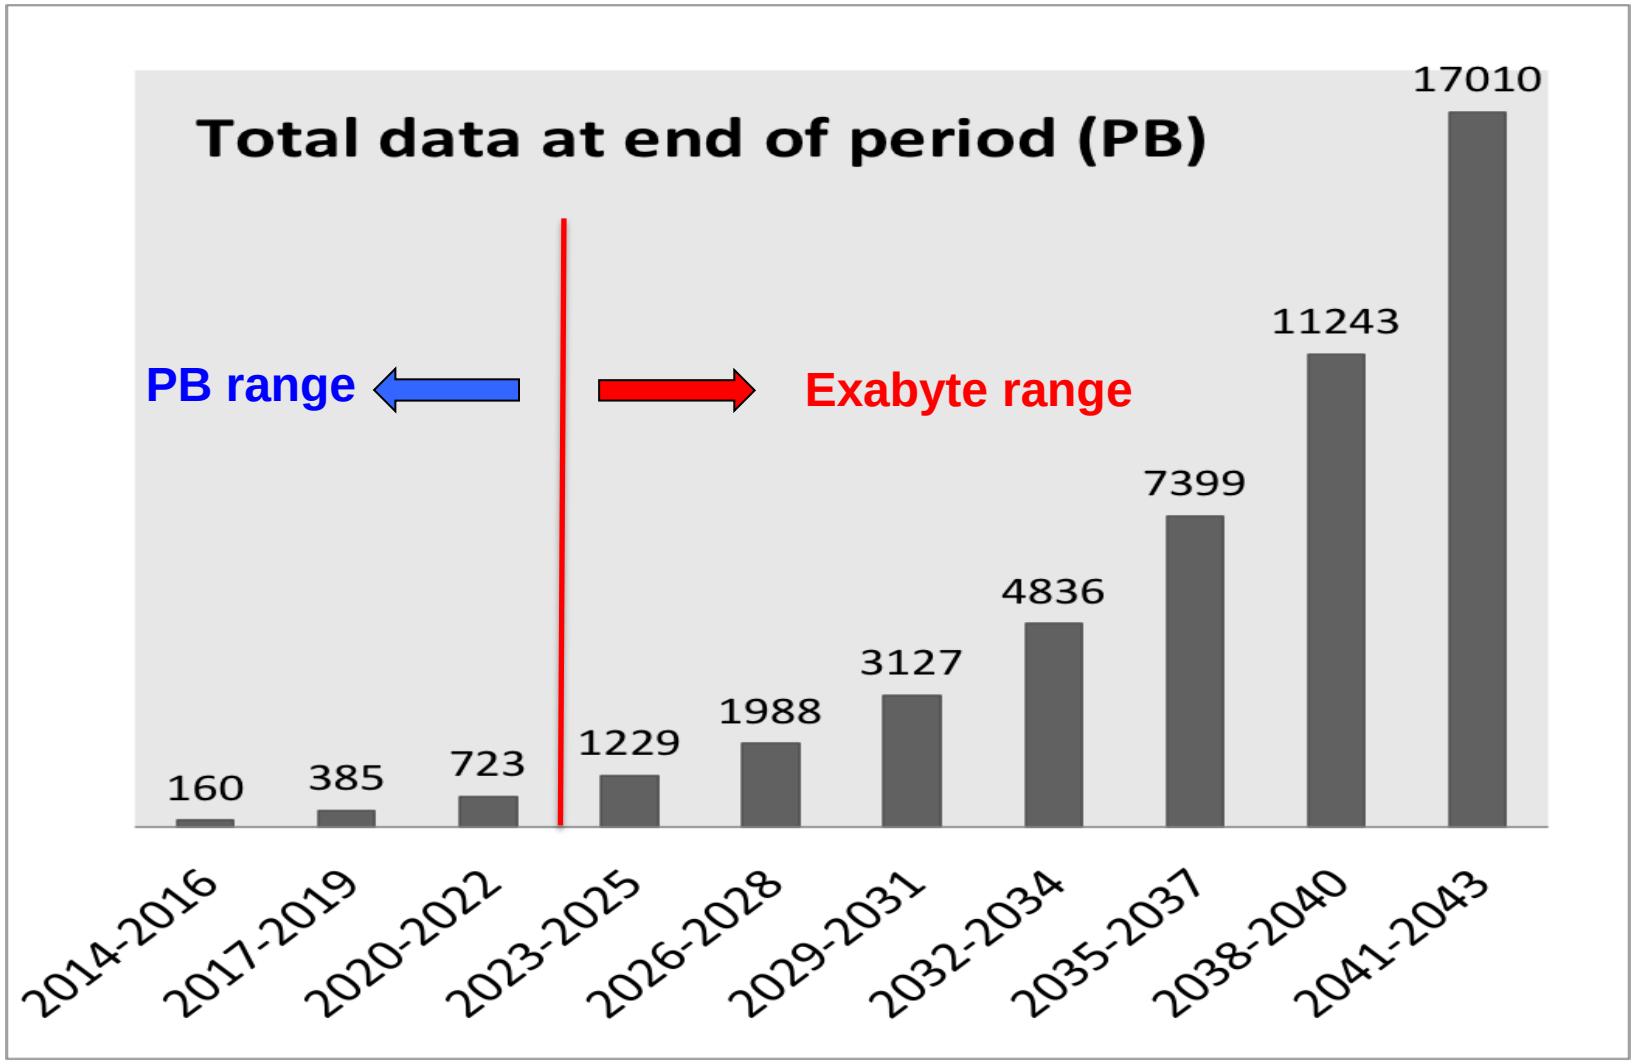

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Case B) increasing archive growth

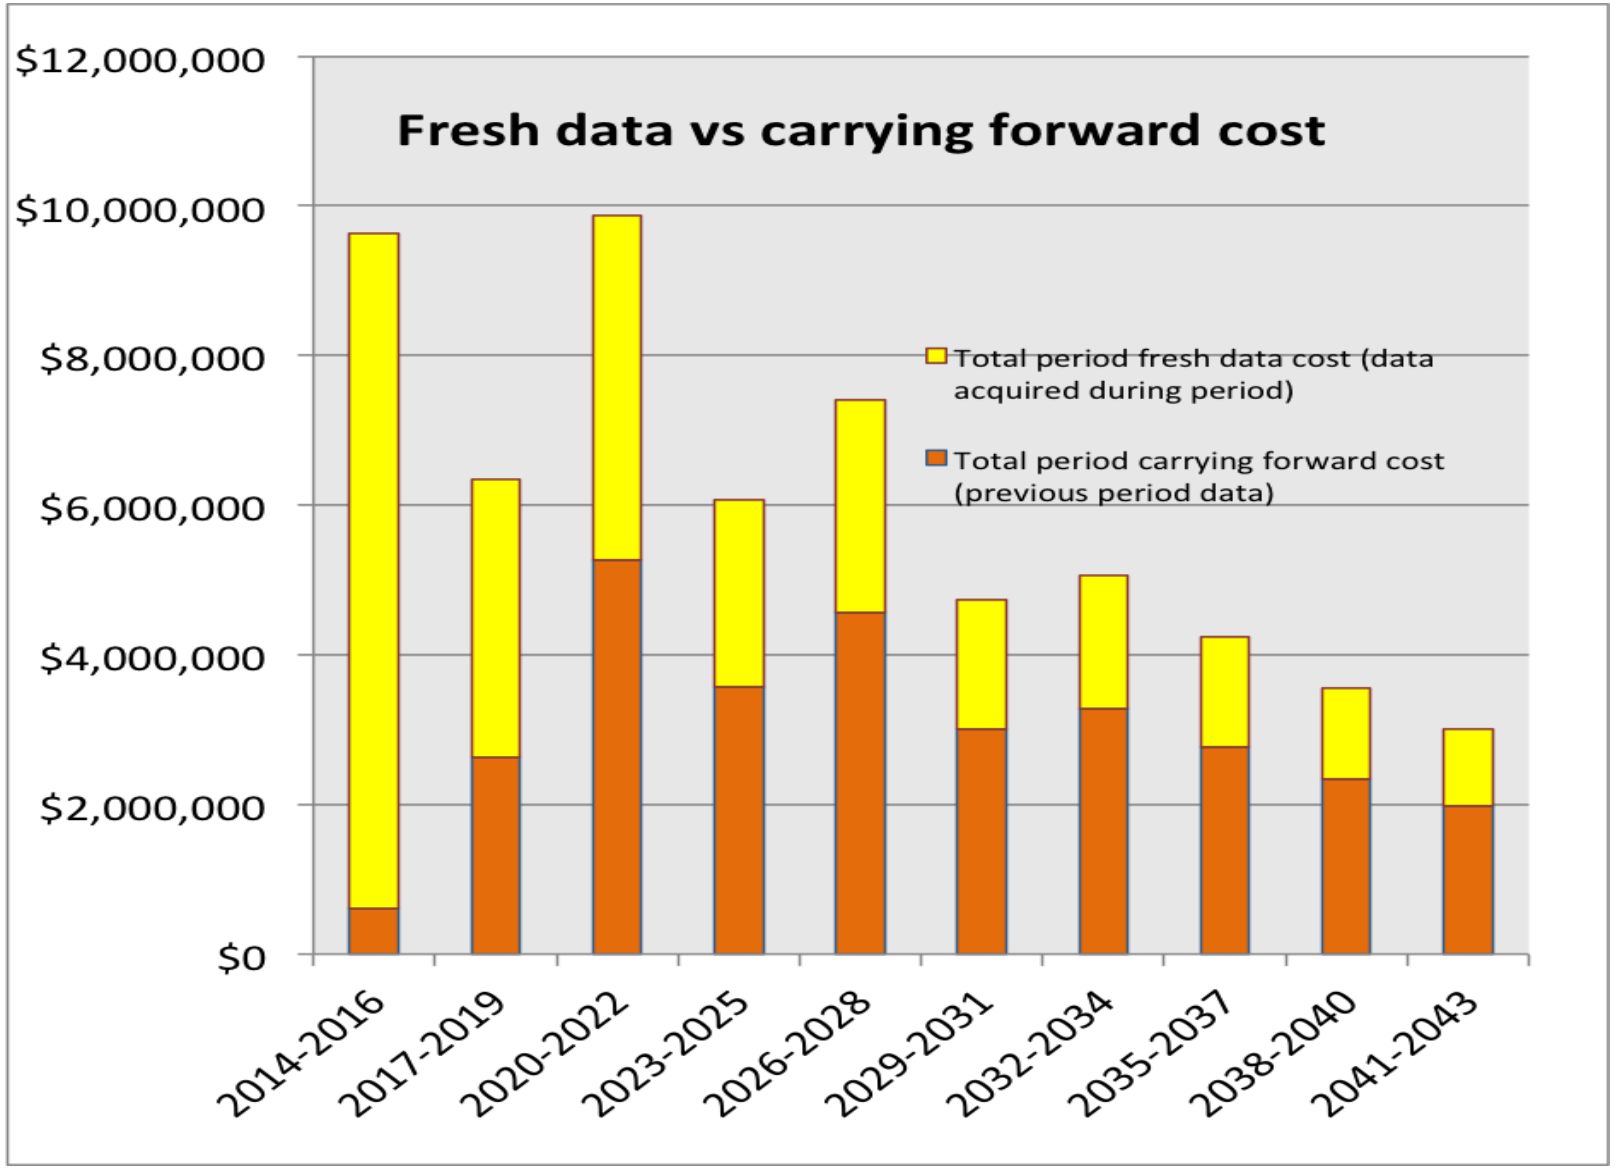

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

LEP Era Data (aka “peanuts”)

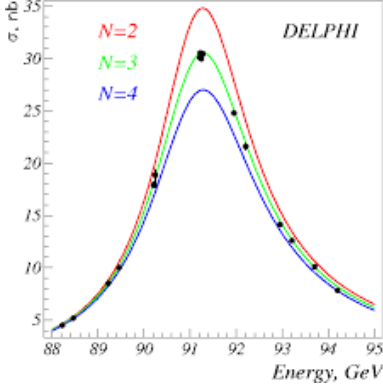

- A **full copy** of the data from the 4 LEP experiments is now on disk (EOS) at CERN
- Plus **2 further copies** on tape at CERN
  - And **multiple copies** at outside institutes
  - **ALEPH: originally a physical machine per collaborating institute – later a VM**

- Will LHC data also be “peanuts” one day?
- (The “bits” are not the difficult bit…)

Some Specific Questions…

- **How do we scale data preservation and data access to address petabyte volumes of data** & **allow access by users across the world?**

- **How can we preserve material data infrastructures: storage media, software, and the machines to run it?**

- **What are the challenges in data description and curation, and how are different academic disciplines managing this change?**

“Software Preservation”

CernVM and CVMFS

CernVM (-FS) (*caveat emptor)*

- Main ingredients
  - CernVM-FS embedded versioning (time machine)
  - Well defined and versioned recipes to regenerate VMs
- **Being actively used for past & present experiments – also for "open data releases"**
- ➢ Complementary to – not a replacement of – forward porting and validation (the latter always required)
- **Big progress in recent years – similar approaches also in other disciplines**

Demonstrators for Software Environment Preservation

- ALEPH
  - Software was last validated on Scientific Linux 4
  - Dedicated VM and containers
- CMS Open Data Pilot
  - Released data were taken in 2011
  - Requires Scientific Linux 5 from same period
- CernVM 3 is an enabling technology for the Open Data Project
  - http://opendata.cern.ch

ALEPH: regenerating SLC4

```
/cvmfs/cernvm-slc4.cern.ch          (OS template)
/cvmfs/aleph.cern.ch                (ALEPH software)
/cvmfs/sft.cern.ch/lcg/external/cernlib   (CERNlib)
```

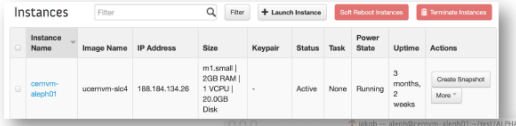

cernvm-aleph01
VM machine on
CERN openstack

```
pb-d-128-141-134-74:~ jakob$ ssh -X aleph@cernvm-aleph01
aleph@cernvm-aleph01's password:
[aleph@cernvm-aleph01 ~]$ source setaleph.sh
[aleph@cernvm-aleph01 ~]$ cd test/ALPHA/
[aleph@cernvm-aleph01 ALPHA]$ sh alpha.sh
     ***********************************************
     **********          ALPHA RUN          **** 11.6 ***
     ***********************************************
     ***********************************************

Wed Mar 19 16:10:27 CET 2014

   ************************************************************
   ***    Compilation  and creation of the makefile 6lep.mk
   ************************************************************
gmake -f /home/aleph/test/ALPHA/6lep.mk
gmake: `6lep' is up to date.
```


CMS Open Data Pilot: deployment

```
/cvmfs/cernvm-slc5.cern.ch          (OS template)
/cvmfs/cms.cern.ch                  (CMS software)
```

- Deployed as OVF/OVA bundle[1]
  - Open specification for bundling VMs, stable since 2009
  - OVA: tarball with a hard disk image and an XML specification

Easy auto-installation
in VirtualBox

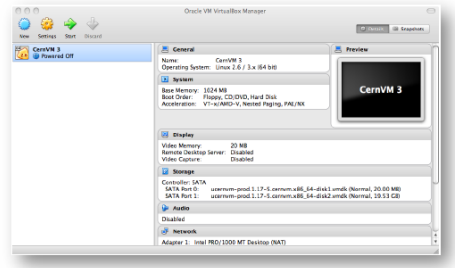

Same solution for
LHCb @ OpenData
http://opendata.cern.ch

[1] Open Virtualization Format / Open Virtual Appliance: http://www.dmtf.org/standards/ovf

2020 Vision for LT DP in HEP

- ***Long-term – e.g. FCC timescales**: disruptive change*
  - By 2020, all **archived data** – e.g. that described in DPHEP 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
  - **DPHEP portal**, through which data / tools accessed
    - **"HEP FAIRport": Findable, Accessible, Interoperable, Re-usable**

➢ **Agree with Funding Agencies clear targets & metrics (via DMPs?)**

Portals & Websites

Official Database of CERN Experiments	"The Grey Book" – Experiments, Institutes and Scientists
DPHEP Portal	Access to Data Preservation Status of HEP institutes worldwide (and, where applicable, other portals)
CERN Open Data Portal	Release (eventually large) subsets (copies) of the data with documentation, software and environment to run it. Access (release) policies of the expts
CERN Analysis Portal	(Currently only within .cern.ch) Detailed Use Cases (= knowledge capture?)
HEPData	Repository of data from publications

A Look Ahead… ("open science")

- Following a pilot, half-day event in Nov 2015, we foresee a larger, multi-disciplinary workshop in **~April 2017**
- Target 200 – 250 people, focussing on practical experience of **data sharing**, **re-use**, **reproducibility** of results, **linking** publications to data (and other objects with DOIs) etc.

➢ **Programme committee & dates ~ April 2016**

Summary

- We believe we have "**bit preservation**" under control – technically and financially – with a **~2+ decade** outlook
- Good solutions exist for **documentation** but there are often tensions between short-term needs (aka Wikis) and LTDP (also **~2+ decade** outlook)
- **Virtualisation** looks increasingly promising but probably can't protect against "disruptive change"
  - e.g. proprietary O/S + h/w to commodity + Linux
  - Possibly **~1 decade retrospective + 1 prospective**

➢ **"The Proof of the Pudding" – test, validate, repeat constantly (and Collaborate!)**


CERN

2020 Vision for LT DP in HEP

- Long-term – e.g. FCC / mescales: disrup/ve change
  - By 2020, all archived data – e.g. that described in DPHEP 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
  - DPHEP portal, through which data / tools accessed ("HEP FAIRport": Findable, Accessible, Interoperable, Reusable)
- Agree with Funding Agencies clear targets & metrics

http://science.energy.gov/funding-opportunities/digital-data-management/

- *"The focus of this statement is sharing and preservation of digital research data"*
- **All proposals submitted to the Office of Science (after 1 October 2014) for research funding must include a Data Management Plan (DMP) that addresses the following requirements:**
1. **DMPs should describe whether and how data generated in the course of the proposed research will be shared and preserved.**

   If the plan is not to share and/or preserve certain data, then the plan must explain the basis of the decision (for example, cost/benefit considerations, other parameters of feasibility, scientific appropriateness, or limitations discussed in #4).

   At a minimum, DMPs must describe how data sharing and preservation will enable validation of results, or how results could be validated if data are not shared or preserved.

U.S. DEPARTMENT OF ENERGY | Office of Science

1 – 'Long Tail' of Papers

2 – 'New' Theore+cal 'Insights'

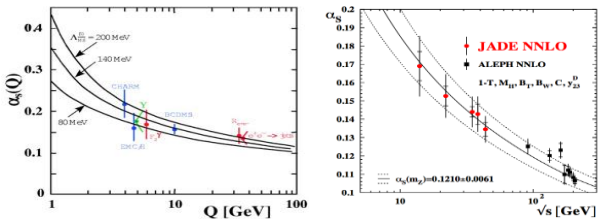

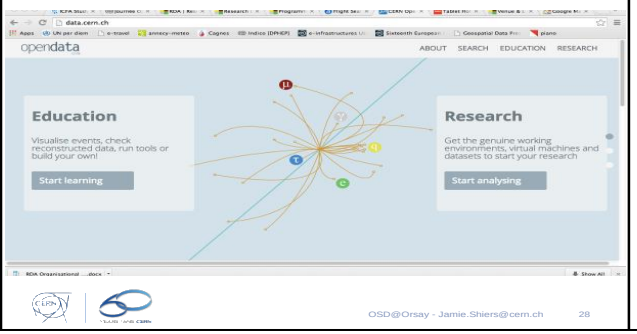

DSS Repack

CERN IT Department

http://indico.cern.ch/event/CERN-ITTF-2014-09-26

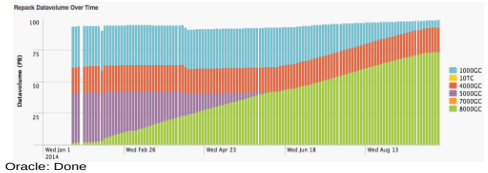

- Oracle: Done
  - 39PB self-repacked (5->8TB), 27PB 1TB emptied
- IBM: Dec'14-Mar'15
  - 20PB of IBM 4TB to self-repack and 5.6PB 1TB tapes to empty
- All repacked media has been verified
- All problem source tapes identified and being handled (cf next slides)
- Cleanup of tape pools and (properly) establishing double copies
  - across buildings
  - complete second copies where missing (ie OPAL)

3 – "Discovery" to "Precision"

possible long-term time line

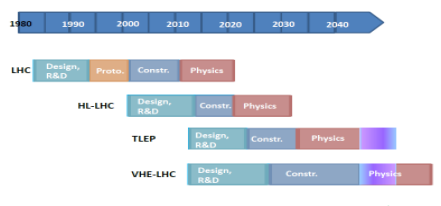

Alain Blondel TLEP design study r-ECFA 2013-07-20 Zimmermann

Use Case Summary

1. Keep data usable for ~1 decade
2. Keep data usable for ~2 decades
3. Keep data usable for ~3 decades

Volume: ~100PB + ~50PB/year (+500PB/year from ~2025)

4C Roadmap Messages

A Collaboration to Clarify the Costs of Curation

1. Identify the **value** of digital assets and make **choices**
2. Demand and choose more **efficient** systems
3. Develop **scalable** services and infrastructure
4. Design digital curation as a **sustainable** service
5. Make funding **dependent** on costing digital assets across the whole lifecycle
6. Be **collaborative** and **transparent** to drive down costs

Balance sheet – Tevatron@FNAL

20 year investment in Tevatron	~$4B	
Students	$4B	~$50B total
Magnets and MRI	$5-10B	
Compu>ng	$40B	

Very rough calculation confirms our gut feeling that investment in fundamental science pays off

I think there is an opportunity for someone to repeat this exercise more rigorously

cf. STFC study of SRS Impact

http://www.stfc.ac.uk/2428.aspx

Science & Technology Facilities Council

Approximation of (HL-)LHC Growth

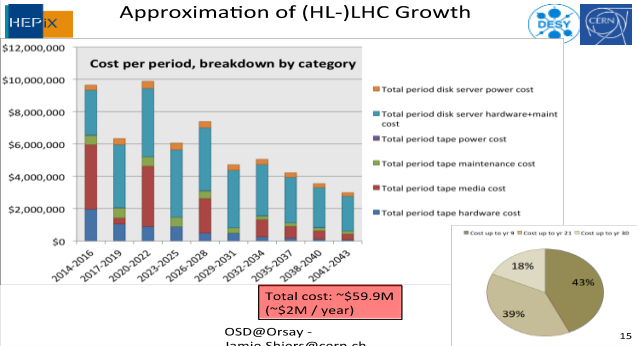

Sustainability – Funding

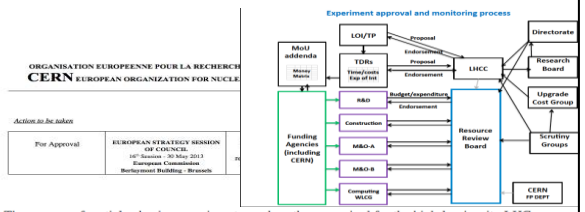

i) The success of particle physics experiments, such as those required for the high-luminosity LHC, relies on innovative instrumentation, state-of-the-art infrastructures and large-scale data-intensive computing. *Detector R&D programmes should be supported strongly at CERN, national institutes, laboratories and universities. Infrastructure and engineering capabilities for the R&D programme and construction of large detectors, as well as infrastructures for data analysis, data preservation and distributed data-intensive computing should be maintained and further developed.*

What Next?

- **Training on, and certification of, sites as "Trusted Digital Repositories"**
- **Expanding "DPHEP Portal" to other (non-LHC) experiments and external sites**
- **Supporting key experiment Use Cases / Funding Agency Requirements**
  - **Reproducibility, Open Access for Outreach, DMPs**

➢ **Ensuring everything is sustainable, documented, "standards-based" and complete**

CMS Open Data release

- Data
  - CMS collision data in format used in analysis by CMS physicists (AOD)
  - For the next release, a partial set of simulated MC included (for the first release no corresponding MC available)
  - For future releases, include "miniAOD" (less complete, but more compact and cleaner)
- Tools
  - VM image of the computing environment
  - Access to the corresponding software and condition data
  - Access to data through xrootd or direct download
- Instructions
  - Basic instructions to get started ($\approx$ 15 mins to setup)
  - Basic description of the physics objects
- Examples of derived datasets to be used in different education and outreach contexts
  - Event display, online histogramming
  - Code to produce the derived datasets

Examples of open data usage

- Ongoing analysis at MIT on jet substructure
  - a small group with a theorist, a post-doc and undergraduate
  - got started with the instructions on portal, and got help on volunteering basis from MIT and US CMS colleagues
  - aiming for a publication
  - willing to contribute to the documentation to help other users
- Research into cloud computing security
  - testing data deletions and operations by the local file system
  - the nature of the data itself is not relevant, but LHC data ideal.
- Pilot project on teaching applicatios for high-schools
  - ideas from physics teachers on further education course at CERN
  - based on the existing tools online tools (event display...)
- External resources have been generated
  - IFCA provides computing resources https://cmsopendata.ifca.es/

Use Cases – “all HEP”

1. Bit preservation – basically OK (at CERN) but *not a formal policy*
   - **On the path to Certification of WLCG “digital repositories” (Tier0/Tier1)**
2. **Preserve data, software, and know-how** in the collaborations **(see CAP Use Cases – backup)**
   - Foundation for long-term DP strategy
   - **Analysis reproducibility:** Data preservation alongside software evolution
3. **Share data and associated software** with (wider) 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 policy documents for LHC collaborations
  - http://opendata.cern.ch/collection/data-policies

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