INFN OPEN SCIENCE

Archives and Data Management Systems in the Big Data Era
Feb. 26-28, 2025

Open Data
Long-Term Data Preservation

Luca dell'Agnello,
Stefano Dal Pra
INFN-CNAF

INFN Istituto Nazionale di Fisica Nucleare

FAIR Principles

- **Findable**
  - **Persistent Identifiers**: Assign globally unique and persistent identifiers to data
  - **Rich Metadata**: Describe data with detailed metadata to facilitate discovery
- **Accessible**
  - **Standardized Protocols**: Use open, free, and universally implementable protocols for data retrieval
  - **Metadata Accessibility**: Ensure metadata remains accessible even in case of restricted data access
  - **No economic barriers** to access.
- **Interoperable**
  - **Common Languages and formats**: Use open data formats, and "popular" languages for data representation
  - **Vocabularies**: Employ vocabularies that adhere to FAIR principles
- **Reusable**
  - **Detailed Documentation**: Provide adequate descriptions and provenance information to support data reuse
  - **Community Standards**: Ensure data meets domain-relevant community standards

DOI

WebDAV

(Digital) Archive

- Collection of data with the following key characteristics:
  - **Preservation –** Ensuring long-term protection of data
  - **Organization –** Records are systematically arranged and catalogued for easy retrieval
  - **Authenticity:** The integrity and authenticity of the records is maintained
  - **Access –** Providing (standard) access protocols; access can be restricted (*open as possible, closed as necessary*)
- In an archive, DOIs can be assigned to data
- While it is not mandatory to adhere to FAIR principles, open access to data in an archive allows sharing it with other communities and/or future generations of scientists

Why Long-Term Data Preservation (LTDP)

- Scientific research, economic value, legal compliance, historical record-keeping are strong drivers for LTDP
  - cost to collect data much greater than cost to keep them.
- **Ensures Accessibility**: LTPD ensures that original datasets remain accessible for subsequent validations
- **Enables (late) Reusability**: Access to preserved datasets allows researchers to conduct secondary analyses, uncovering new insights
- **Irreplaceable Data**: experiments might not be easily reproducible
  - E.g., data from CDF ($p\bar{p}$), LEP experiments (e^+e^-), BABAR

Long-Term Data Preservation: how

Establishing a comprehensive data management framework is essential:

- standardized formats
- Data Management Plans
- stakeholder engagement.
- Preserving data usability over time involves multiple levels
  - Storage – Physical protection of data
  - Integrity – Data curation
  - Control – Access Control to data
  - Metadata – Curation of inventory of content
  - Content – Documentation about data

} *Bit preservation* (Storage, Integrity)

} *Analysis framework preservation* (Metadata, Content)

Storage level (1/2)

- Keep at least 2 complete copies in different locations (possibly with different technologies)
  - For experiments at LHC, data are duplicated among CERN and Tier1 centers using automatic procedures and “standard” tools (e.g., FTS and Rucio)
- Have a plan and execute action to address obsolescence of storage hw, sw and media
  - At CNAF we periodically migrate data from a generation of tapes to a newer one
  - Operating system of servers is regularly upgraded (upon agreement with the collaborations)
- For no more active collaborations, data access applications become obsolete after a few years, and they are instantiated (as Virtual Machines) when needed

Storage level (2/2)

- Tape storage usually offers the best compromise between technical and economic requirements
- Special devices to meet legal requirements for data immutability (e.g., WORM devices) are also used to ensure the integrity and security of critical information

Integrity level

- Ensure that the data remains accurate and unaltered over time against both deliberate tampering and accidental damage
- Performing regular (automatic) checks to verify data integrity (and restore from a secondary archive if data is compromised)
  - Common practice in HEP is generating a checksum of the data unit (file) and verify when transferring and reading from tape
  - In WLCG a corrupted copy of a file is invalidated and then retransferred from another site (o regenerated in some cases)

Control level

- Determine who has rights to access (read, write, delete, etc..)
- Control and log accesses
- For scientific data this is implemented at a basic level
  - E.g., production managers of a collaboration can write and delete data of that collaboration, other users can only read
- For administrative staff (e.g. Alfresco) stronger security is needed

Metadata and content levels

- Maintain detailed information about content of the data and location of the storage
- In HEP it is done normally by the collaborations using catalogues
- For no more active collaborations, ensuring accessibility of data requires that metadata is also preserved at the data centers

Data Preservation in HEP

- Problem of LTDP addressed since a few years
- **DPHEP collaboration**: Data Preservation in High Energy Physics
  Past experiments have already successful DP projects in place (e.g., Babar, CDF, Aleph)
- All **LHC experiments** are devoting efforts to data preservation
- A data preservation project in HEP can be better described in a functional way:
  - *Bit preservation: how preserve data*
  - *Analysis framework preservation: code preservation, virtualization*

DPHEP-2023-01-PUB
April 2023

First community report 2012 - Updated with experience last ten years

Data Preservation in High Energy Physics: Decennial Global Report

DPHEP Collaboration*

Abstract

Data preservation is a mandatory specification for any present and future experimental facility and it is a cost-effective way of doing fundamental research by exploiting unique data sets in the light of the continuously increasing theoretical understanding. This document summarizes the status of data preservation in high energy physics. The paradigms and the methodological advances are discussed from a perspective of more than ten years of experience with a structured effort at international level. The status and the scientific return related to the preservation of data accumulated at large collider experiments are presented, together with an account of ongoing efforts to ensure long-term analysis capabilities for ongoing and future experiments. Transverse projects aimed at generic solutions, most of which are specifically inspired by open science and FAIR principles, are presented as well. A prospective and an action plan are also indicated.

Submitted to Eur. Phys. J. C

A The DPHEP Collaboration

DPHEP Collaboration

T. Basaglia[1], M. Bellis[I,2], J. Blomer[1], J. Boyd[1], C. Bozzi[3], D. Britzger[4], S. Campana[1], C. Cartaro[5], G. Chen[6], B. Couturier[1], G. David[II,7], C. Diaconu[8], A. Dobrin[9], D. Duellmann[1], M. Ebert[10], P. Elmer[11], J. Fernandes[1], L. Fields[21], P. Fokianos[1], G. Ganis[1], A. Geiser[12], M. Gheata[9], J. B. Gonzalez Lopez[1], T. Hara[13], L. Heinrich[1], M. Hildreth[21], K. Herner[14], B. Jayatilaka[14], M. Kado[1], O. Keeble[1], A. Kohls[1], K. Naim[1], C. Lange[20], K. Lassila-Perini[15], S. Levonian[12], M. Maggi[16], Z. Marshall[18], P. Mato Vila[1], A. Mečionis[1], A. Morris[17], S. Piano[16], M. Potekhin[7] M. Schröder[1], U. Schwickerath[1], E. Sexton-Kennedy[14], T. Šimko[1], T. Smith[1], D. South[12], A. Verbytskyi[4], M. Vidal[1], A. Vivace[1], L. Wang[6], G. Watt[19], T. Wenaus[7]

Affiliation Notes

[I] Also at: Siena College
[II] Also at: Stony Brook University

Collaboration Institutes

[1] CERN, Geneva, Switzerland
[2] Cornell University, USA
[3] INFN Ferrara, Italy
[4] Max-Planck-Institut für Physik, München, Germany
[5] SLAC National Accelerator Laboratory, USA
[6] Institute of High Energy Physics, IHEP, CAS, Bejing, China
[7] Brookhaven National Laboratory, BNL, USA
[8] Aix Marseille Univ, CNRS/IN2P3, CPPM, Marseille, France
[9] Institute of Space Science, ISS, Bucharest, Magurele, Romania
[10] HEP Research Computing, University of Victoria, BC, Canada
[11] Princeton University, USA
[12] Deutsches Elektronen Synchrotron, DESY, Hamburg, Germany
[13] High Energy Accelerator Research Organization, KEK, Tsukuba, Japan
[14] Fermi National Accelerator Laboratory, Batavia, USA
[15] Helsinki Institute of Physics, Finland
[16] INFN Trieste, Italy
[17] University of Bonn, Germany
[18] Lawrence Berkeley National Laboratory, Berkeley, USA
[19] IPPP, Durham University, Durham, UK
[20] Paul Scherrer Institut, Villigen, Switzerland
[21] University of Notre Dame, Notre Dame, USA

The CDF use case

- Experiment running at Tevatron (FNAL) in the period 1985-2011
  - Main achievement: discovery of top quark (1995)
- End of 2011: Start of discussions for implementing LTDP of CDF at CNAF
- Goal: preserve a complete copy of CDF data and MC samples and services (access, data analysis capabilities)...
- In September 2012, CSN1 agreed to fund LTDP of CDF data

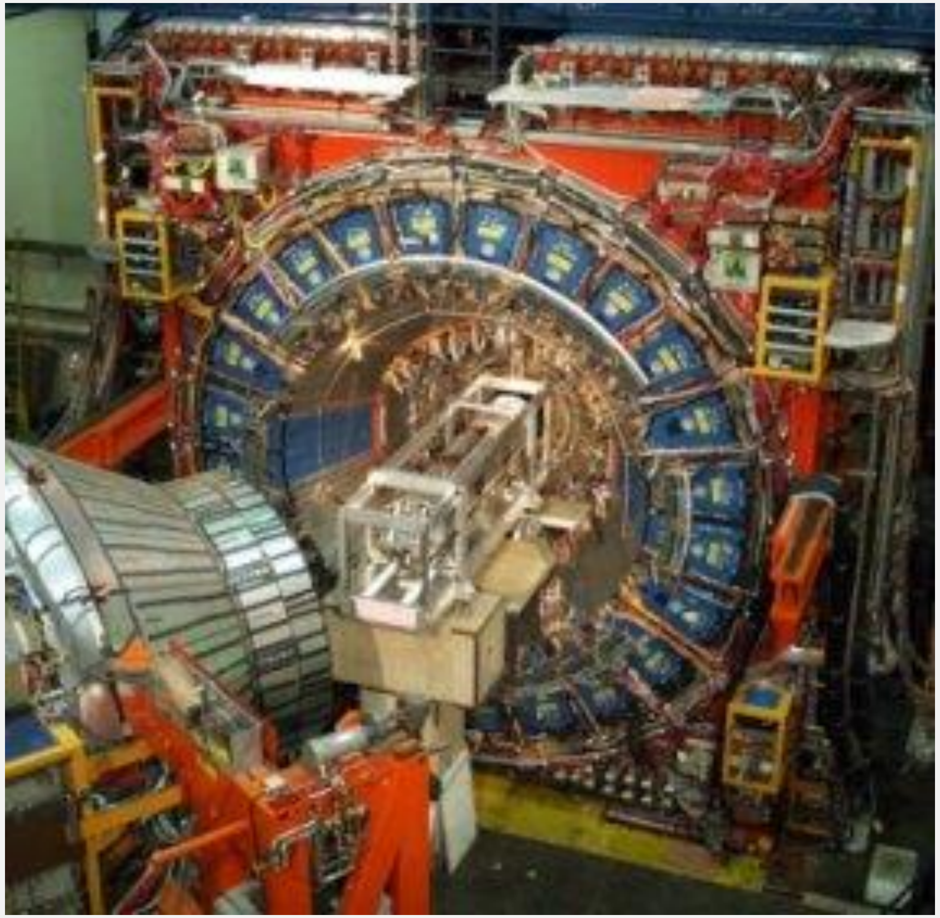

Why preserve CDF data?

- Historical value
  - Unique proton-antiproton collisions
  - Discovery of the top quark
- Educational value
  - Training of students
- Scientific value
  - Physics results are still being extracted from the data
  - <u>2022:</u> <<Using data collected by the Collider Detector at Fermilab, or CDF, scientists have now determined the particle’s mass with a precision of 0.01% – twice as precise as the previous best measurement.>>

CNAF-CDF project

- Some “easy” steps
  - Copy all data from FNAL to CNAF
    - Implementing bit preservation
  - Copy metadata and related documentation
  - Set up VMs with all the needed services and computing platforms
    - Implementing analysis framework
- The first issue was getting the data
  - Data from Run1 (1992-1996) 4125 Exabyte 8mm tapes, 2.5 or 5 GB
    - Valuable for educational purposes (discovery of top quark)
  - Data from Run2 (2001-2011) available on Oracle tape libraries at FNAL
    - Valuable physics content

CNAF-CDF project: copy of Run2 data (1/2)

- ~4 PB of data from Run2
  - All data and MC user level n-tuples (2.1 PB)
  - All raw data (1.9 PB)
  - Databases (O(100 GB))
- Decision to transfer all the data over the network using “standard” (i.e. Grid) tools
  - Not feasible to use GPN (too slow) or LHCOPN/ONE (used for LHC)
  - Dedicated 5 Gbps link FNAL-CNAF provided by GARR
    - Link provisioning and set up took ~10 months (April 2012 – February 2023)
    - Not only GARR involved (also GEANT, ESNET and FNAL…)
  - Estimated (very optimistic) time for the copy: ~3-6 months

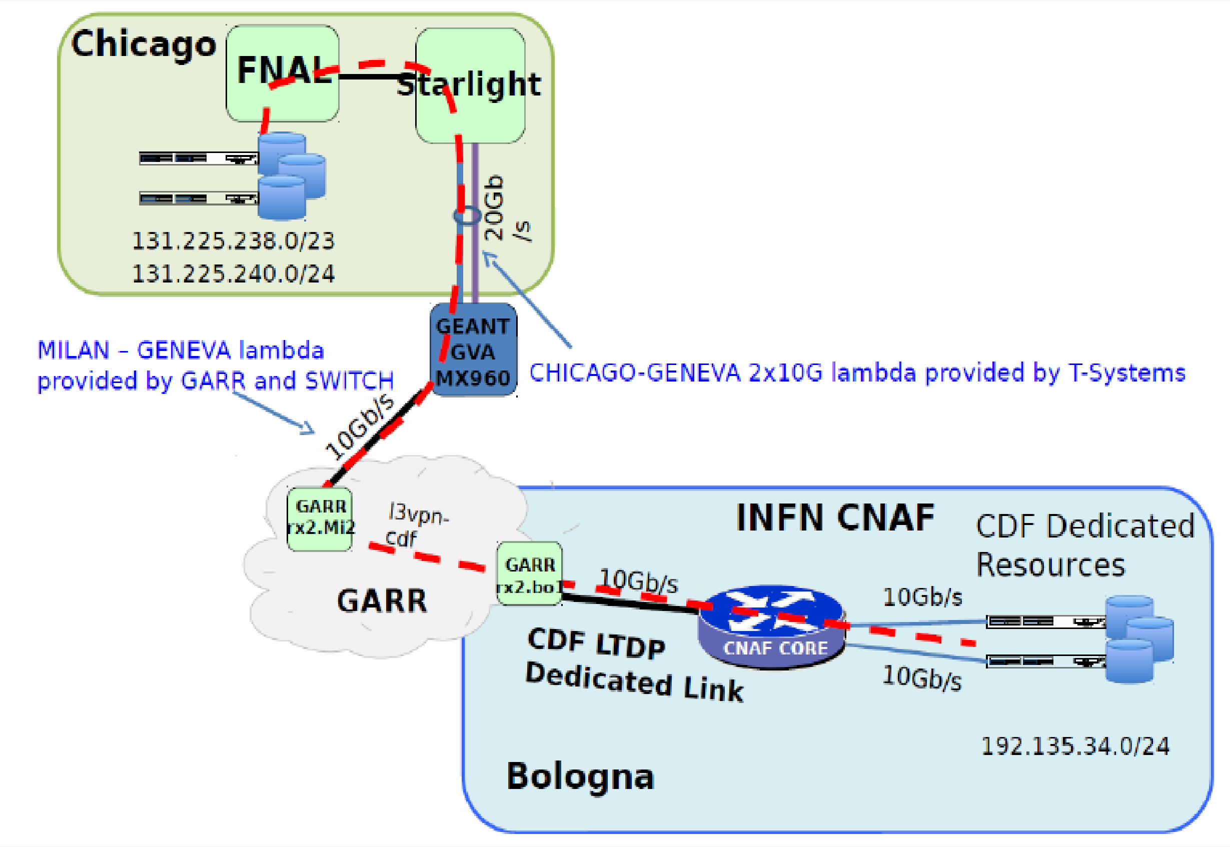

CNAF-CDF project: copy of Run2 data (2/2)

- Network was not the only issue
- Tapes available by the beginning of 2014
  - Funding for tapes at CNAF delayed until June 2013
  - Contingent upon the availability of the project manager (on CDF side)
- Some modifications of CDF tools needed to use third party gridftp transfers and interface with CNAF Mass Storage System
- Unable to use the network steadily at full capacity due to a bottleneck in the transfer to the library
- After the start-up minimal FTE overhead
- Copy process took nearly 12 months

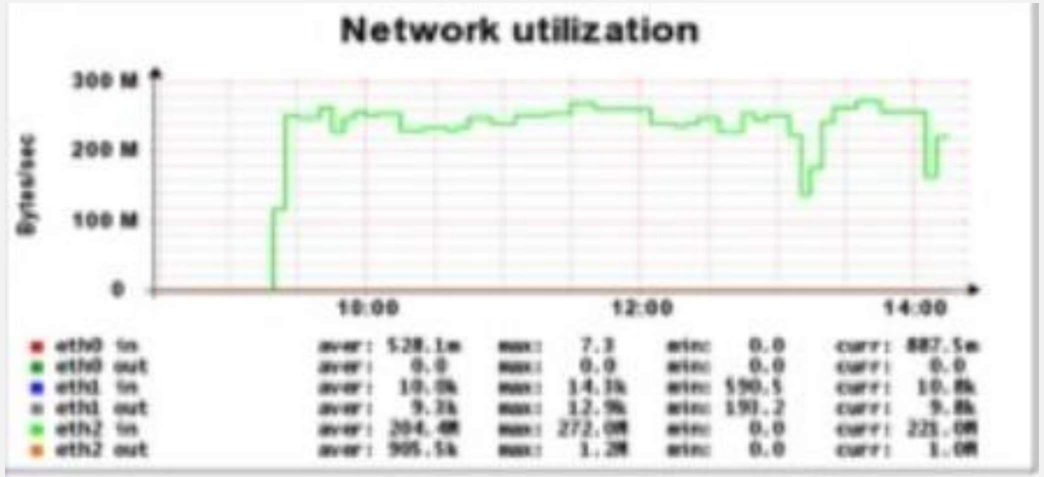

Timeline of the project (1/2)

- Project in September 2012
- Funds available in June 2013
- First tapes available in January 2014
- Copy started in February 2014
- Copy completed in February 2015

Costs

- 2013: 89 k€
  - Tape drives, servers, tapes
- 2014: 99 k€
  - Tapes, servers

+ 0.2 FTE for 2 years
link FNAL-CNAF provided by GARR pro bono

Timeline of the project (2/2)

- Other activities from February 2015 to 2019
  - Recovery of missing files (and then periodic checking of copied files)
  - Installation of a new data access system and new CDF code
  - Data analysis with the new *jobsub* system at CNAF (2017)
  - Production of a web page (documentation)
- **To complete the above tasks, we needed to travel to FNAL in search of an expert (a nearly retired guy)**

What about data from Run1?

- The only copy was on ~4000 Exabyte tapes (~20 TB in total)
  - Robots to read this type of tapes no more available at FNAL
- 2016: tapes shipped to CNAF
- 3 Autoloaders procured (total cost: 5 k€)
- Not negligible effort foreseen (load tapes) for ~1 year (to recover 20 TB....)
- We started a pilot project to recover the most from these old media.

BUT.....

On Nov 9, 2017, Tier1 data center experienced a flooding, resulting in severe damage to those tapes....

Next step was Disaster Recovery...

Lesson learned (1/3)

- Comply with the guidelines described for bit preservation and analysis framework preservation
  - e.g., multiple copies in different locations using different technologies, save metadata, document all the procedures,...
- International collaborations usually follow these policies (at least until they are active)
- For smaller collaborations, this may not always be the case: we should ensure these guidelines are enforced
  - During the CNAF flooding, some tapes from another experiment were also damaged
  - They had a second copy.... in the same library at CNAF (no data loss luckily)

Lesson learned (2/3)

- Whenever possible, utilize standard data and metadata formats, along with standard access protocols and frameworks
  - The convergence towards de facto standards (WLCG) is a step in the right direction
- The INFN Datacloud project, which aims to create a data lake by merging Tier1 and Tier2 resources into a national cloud, could address some of these requirements
  - Copies of data, standard infrastructure and, possibly, protocols

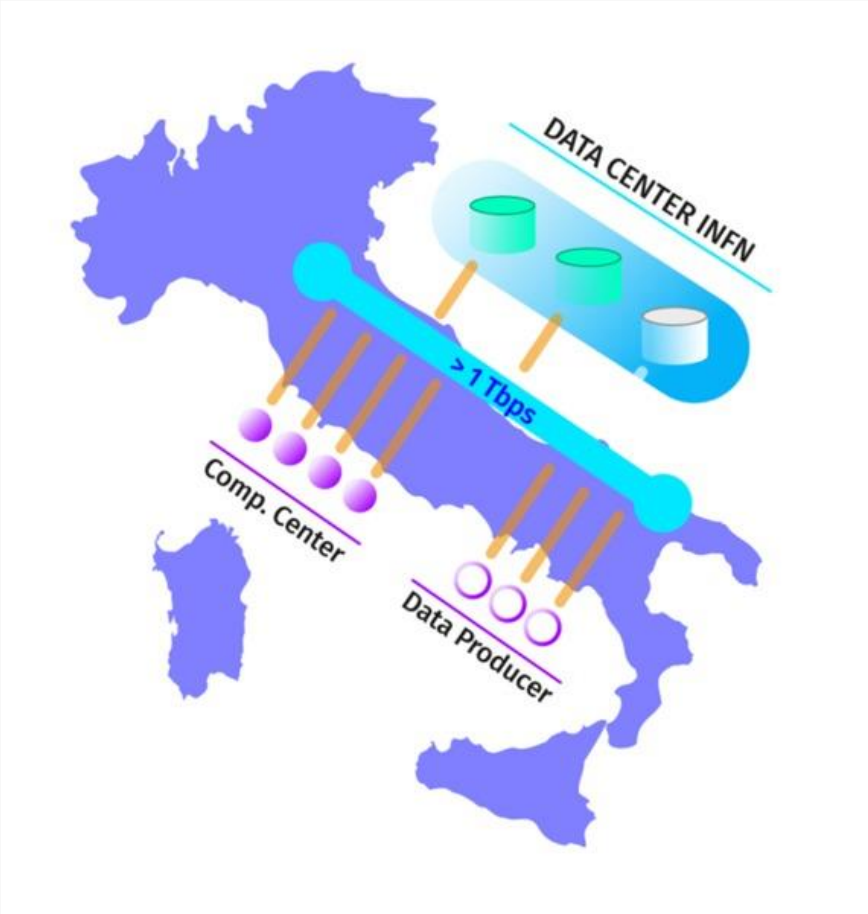

Lesson learned (3/3)

- Preserve the environment for data analysis and document all procedures to analyze data too
- Besides the risk of losing data, there is a high probability of losing know-how (e.g., retirement of relevant people)
- All these steps could be defined or planned with assistance from a *Data Steward*

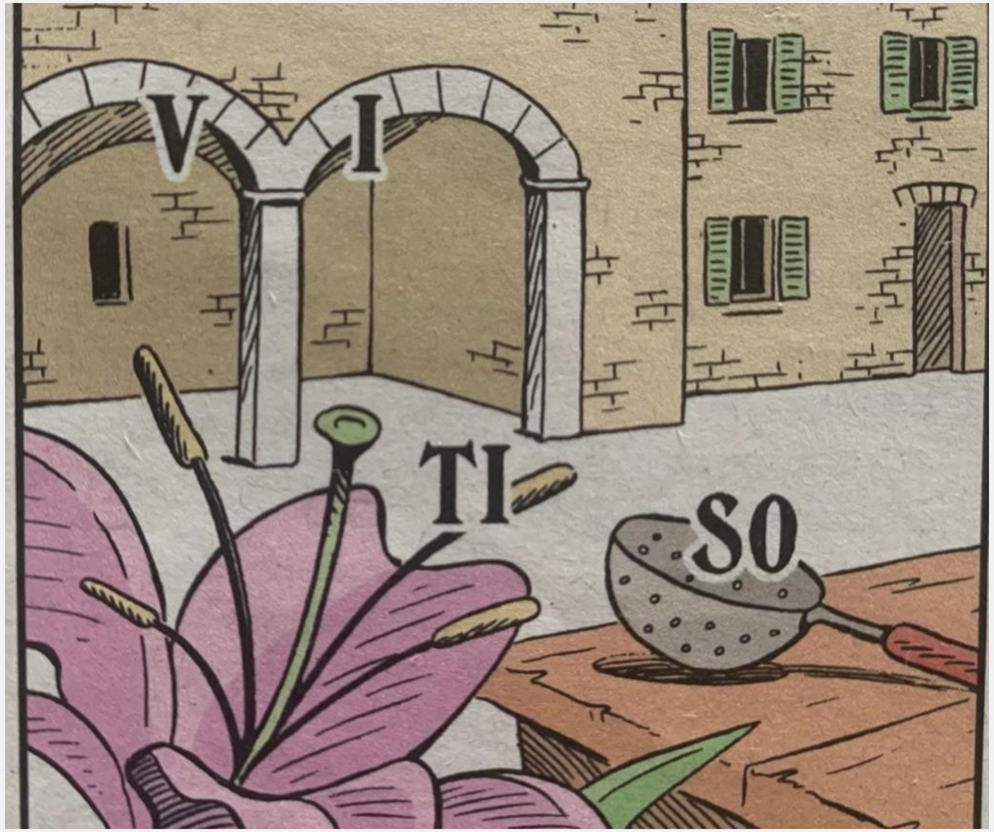

10,10

Some final remarks

- A Data Management Plan (DMP) *including LTDP* should be adopted by all collaborations
  - DMP is already mandatory for all EU grants
  - INFN is working to standardize DMP for the National Laboratories
- Adhering to FAIR principles helps: the usability of open data (including both integrity and accessibility of the data) is also verified by the end user
- It is never too early to start the planning of data preservation :-)

INFN
Istituto Nazionale di Fisica Nucleare

Thanks!

Contacts

- https://www.cnaf.infn.it
- Gruppo di lavoro Open Science dell’INFN: openscience@lists.infn.it