Long-Term Data Preservation in the EOSC according to FAIR principles

2nd ASTERICS-OBELICS Workshop

16-19 October 2017, Barcelona, Spain.

H2020-Astronomy ESFRI and Research Infrastructure Cluster
(Grant Agreement number: 653477).

Outline

- **Long-Term Data Preservation for sharing & re-use: multi-decade, multi-hundreds of PB (to tens of EB)**
- **A Science Demonstrator in the European Open Science Cloud (Pilot) – experience and extension to ESFRIs**
- How might this apply to **INFRAEOSC-04-2018**?
  Connecting ESFRI infrastructures through cluster projects
  - **FAIR principles & Open Science**
    - Including *domain-specific **(& cross-domain)** issues on acquisition, deposit, **preservation, curation, access, sharing & re-use***
  - **Skills needs & development**
    - Best practices, open standards and inter-operability

HEP LTDP SD

- **Motivation: Funding agencies today require (FAIR) Data Management Plans, explaining how data acquired or produced will be preserved for re-use, sharing and verification of results.**

- **Goal: use generic services to build a system capable of storing and preserving Open Data at a scale of 100TB+**

1. **Trustworthy digital repositories (TDR) + PIDs;**
2. **Scalable "digital library" services with DOIs;**
3. **A versioning file system to capture and preserve the associated software and needed environment;**
4. **A virtualised environment that allows the above to run in Cloud, Grid and many other environments.**

Targets & Stretch Targets

- **Given that:**
  - Numerous TDRs exist; **[100 TB scale – not 200 PB!]**
  - **CernVM / CVMFS** offered in **production** to **multiple disciplines**;
    - EGI InSPIRE, EGI Engage, WLCG, other HEP labs;
  - **Invenio-based** services:
    - CDS, INSPIRE-HEP, Zenodo, B2SHARE, … – almost ubiquitous

➢ ***What were we going to do after coffee?***

- **Stretch targets:**

1. Understand / implement F.A.I.R. in multi-disciplinary environment
   - Benefit from FAIR expertise within project
2. Understand (and potentially prototype) use of generic services by other disciplines

FAIR DMPs & TDRs

- *If we want to be able to **share data**, we need to store them in a **Trustworthy Digital Repository** (TDR).*
  - ***Data created and used by scientists should be managed, curated, and archived in such a way to preserve the initial investment in collecting them.***
  - *Researchers must be certain that data held in archives remain **useful** and **meaningful** into the future.*
  - *Funding authorities increasingly require continued access to data produced by the projects they fund, and have made this an important element in **Data Management Plans** (DMPs – H2020 Guidelines).*
  - *Indeed, some funders now stipulate that the data they fund must be deposited in a **trustworthy repository**.*

What is a data repository?

- **For us (HEP), it is much more than just a "bit repository"**
  - And even that probably has several components
    - **Long-term archive (“tape”)**; **cache(s) for production & analysis (“disk”)**; "**Open Access**" area (not necessarily "**immediate Open Access**")
    - What data is accessed when, by whom, access patterns
- **It includes also documentation, software + environment in which it runs, "knowledge"**
- These are probably supported by different services - some of which may already be "remote" - that evolve on different timescales
- ***Something*** is changing all the time!
  - **If you believe in transparent and seamless migrations you probably don't have a sustainable sustainability plan (or have never done a migration)**
- **Sustainable: financially** + **technically** + "**logically**" (holistically?)

Certification for TDRs

- Several certification procedures exist but only one developed by a Scientific Community **– ISO 16363**
  - All based on OAIS model: ISO 14721

- Follows OAIS structure:
  3. **Organisational Infrastructure;**
  4. **Digital Object Management;**
  5. **Infrastructure and Security Risk Management.**

- EOSC Pilot experience: even at “modest” scale (100TB), HEP data formats and long-term needs mean that a “generic” TDR is unlikely to work
  - “Domain” or institutional repositories?
  - Can still share knowledge, experience and even tools in building / operating these!

FAIR Data Principles

Expert Group on turning FAIR into reality

TO BE FINDABLE:

- F1. (meta)data are assigned a **globally unique and eternally persistent identifier.**
- F2. data are described with rich metadata.
- F3. (meta)data are registered or indexed in a searchable resource.
- F4. metadata specify the data identifier.

TO BE ACCESSIBLE:

- A1 (meta)data are retrievable by their identifier using a standardized communications protocol.
- A1.1 the protocol is open, free, and universally implementable.
- A1.2 the protocol allows for an authentication and authorization procedure, where necessary.
- A2 metadata are accessible, even when the data are no longer available.

TO BE INTEROPERABLE:

- I1. (meta)data use a formal, accessible, shared, and broadly applicable language for knowledge representation.
- I2. (meta)data use vocabularies that follow FAIR principles.
- I3. (meta)data include qualified references to other (meta)data.

TO BE RE-USABLE:

- R1. meta(data) have a plurality of accurate and relevant attributes.
- R1.1. (meta)data are released with a clear and accessible data usage license.
- R1.2. (meta)data are associated with their provenance.
- R1.3. (meta)data meet domain-relevant community standards.

Long Term Data Preservation

- e-IRG definition:

Long-term *is defined as a period of time long enough for there to be concern about the loss of integrity of digital information held in repositories, including deterioration of storage media, changing technologies, support for old and new media and data formats (including standards), and a changing user community.*

- DPHEP definition:
  - **"Disruptive change";**
  - **Target periods: 25 – 30 – 50 years.**

➢ **Let's look back: 25 / 30 / 50(!) years…**

In the year…

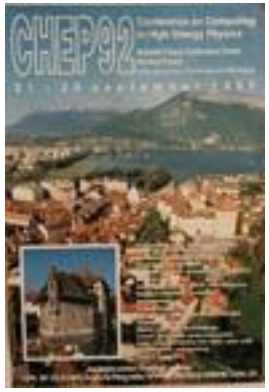

- ***T – 25:***
  - www was ***just*** emerging with the first X-based browsers…
  - CERN (HEP) had just begun migration to "distributed computing" and Unix…
  - Fortran, VMS, VM/CMS still dominant…

- ***T – 30:*** mainframe era; open reel 6250 bpi tapes (no-one in HEP even dreamed of LTDP)
- ***T – 50:***

(moon landing: T – 48)

How Has FAIR evolved in 2017?

- Increasingly, FAIR has been taken to include not just data + meta-data but also **software**
- What started as "source code" preservation has now evolved to **"running s/w and its environment"**
  - **Much better IMHO**
- But there is still a lot to define / do
  - **How is the data Findable?**
    - **Navigation?** Search? Is there an API? …
  - **How to implement this in a scalable & sustainable way**
    - E.g. how many PID / DOI lookups per unit time, for how long is the service "guaranteed", … **"eternally?"**
  - **How to implement cross project / discipline searches?**
- **I have heard claims that people have been doing this for 20 – 100(!!!) years**
  - **(These people clearly don't need any more project money)**

What is (HEP) data? (And its not just “the bits”)

Digital information

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

Other digital sources such as databases to also be considered

Software

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

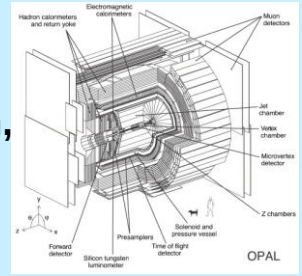

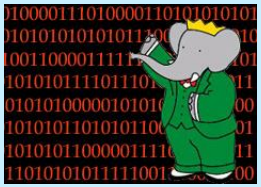

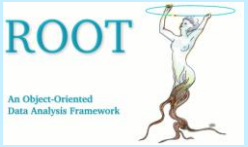

CERNLIB Access

- Access to the CERN Program Library is free of charge to all HEP users worldwide.
- Non-HEP academic and not-for-profit organizations: 1KSF/year

Meta information

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

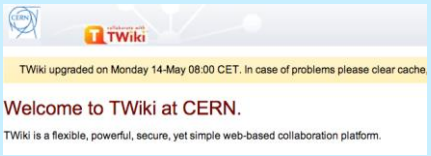

Publications

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

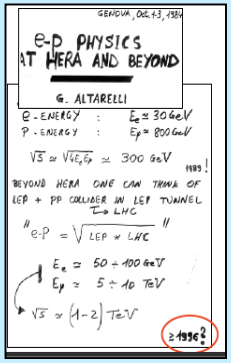

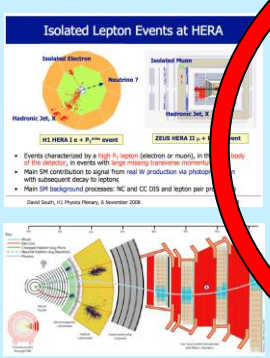

Expertise and people

Typical EU H2020 Call Text

- *Research Infrastructures, such as the ones on **the ESFRI roadmap** and others, are characterised by the **very significant data volumes** they generate and handle.*
- *These data are of interest to **thousands** of researchers across scientific disciplines and to other potential users via **Open Access** policies.*

➢ ***Effective data preservation and open access for immediate and future sharing and re-use are a fundamental component of today's research infrastructures.***

How Are We Going To Solve This?

- **General agreement on potential (short-medium term) BASIC components**

➢ **But that is not the same as offering scalable, sustainable, long-term production services!**

- **The scientific communities must be directly involved in defining the parameters and additional services!**

❑ **And we must not forget that long-term implies CHANGE!**

▪ **(No-one else is going to do this for us)**

Services are (just) services

- No matter how fantastic our { TDRs, PID services, Digital Library, Software repository } etc is, they are there to support **the users**

- **Who have to do the <u>really</u> hard work!**
  - **E.g. write the software, documentation, acquire and analyse the data, write the scientific papers**

- Getting the degree of public recognition as at the Higgs discovery day was a **target KPI**!

Open Science: A 5-Star Scale?

- **We have a 5-star scale for Open Data**
  - **Sir Timothy Berners-Lee**

- **We have a proposed 5-star scale for FAIR data management (+TDRs)**
  - **Peter Doorn and Ingrid Dillo**

➢ **How about a 5-star scale for "Open Science: Open to the World"?**
  - **The EOSC & Friends**

What are the right metrics?

- **As easy to use as Amazon?**
- **Cheaper (and better) than doing it in-house?**
- **A majority of ESFRIs use it as their baseline?**

➢ ***"To find dark matter, you need the EOSC"?***

LTDP: How do we measure progress / success?

- **Practice:** through Open Data releases
  - Can the data really be (re-)used by the Designated Community(ies)?
  - What are the support costs?
  - Is this sustainable?
- **Theory:** by applying state of the art "preservation principles"
  - Measured through ISO 16363 (self-) certification and associated policies and strategies
  - Participation in relevant working & interest groups

One, without the other, is probably not enough. The two together should provide a pretty robust measurement...

(Some) Basic Needs

- **TDRs** capable of handling formats, data types, scale and duration of **ESFRI**(-like) projects
- Understanding of what **FAIR** means to our communities – and how to implement / tailor corresponding services
- **Synergies between ESFRI projects – e.g. using FAIR DMPs to identify them**
- Identification of needed skills → training; direct support to users, including for **MIGRATIONS!**
- **All this can & should be done in cooperation with e-i service providers and other users / communities…**

PANORAMICS…

- **P**reservation
- **AN**d
- **O**pen Data for
- **R**esearchers **A**nd
- **M**embers of the public
- **I**ncorporating
- **C**loud
- **S**ervices


CERN

Acknowledgement

- H2020-Astronomy ESFRI and Research Infrastructure Cluster (Grant Agreement number: 653477).

LEP / (HL-)LHC Timeline

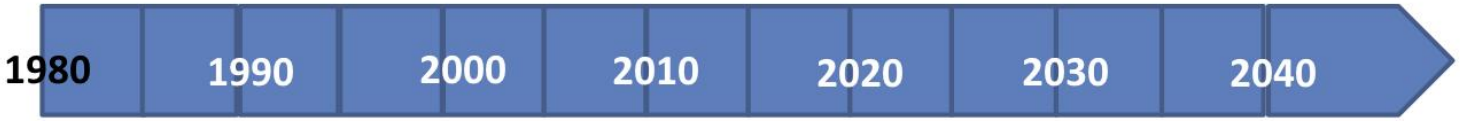

LEP: Constr. | Physics | Upgr.

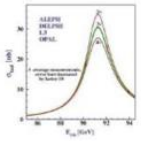

Database / data management support,
CERN Program Library, Distributed Computing

LHC: Design, R&D | Proto. | Constr. | Physics

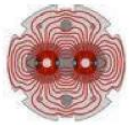

DM R&D, DBs, WLCG, EGI
Major Data Migrations(!)

HL-LHC: Design, R&D | Constr. | Physics

ESFRI roadmap as "landmark project"

- Robust, stable services over **several decades**
- Data preservation and re-use over **similar periods**
- "Transparent" and supported **migrations**

~30 years of LEP – what does it tell us?

- Major migrations are **unavoidable** but hard to **foresee**!
- **Data** is not just "**bits**", but also **documentation, software + environment + "knowledge"**
  - "**Collective knowledge**" particularly hard to capture (remember)
    - Documentation "refreshed" after 20 years (1995) – now in Digital Library in PDF & PDF/A formats (was Postscript)
- Today's "**Big Data**" may become tomorrow's "**peanuts**"
  - **100TB** per LEP experiment: **immensely challenging** at the time; now "trivial" for both CPU and storage
  - With time, **hardware costs** tend to zero
    - O(CHF 1000) per experiment per year for archive storage
  - **Personnel costs** tend to O(1FTE) **>> CHF 1000!**
    - Perhaps as little now as 0.1 – 0.2 FTE per LEP experiment to keep data + s/w alive – no new analyses included

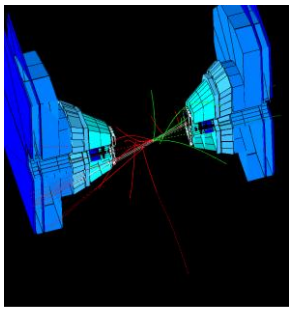

See DPHEP Workshop on "Full Costs of Curation", January 2014: https://indico.cern.ch/event/276820/

ODBMS migration – overview (300TB)

- **A triple migration!**
  - Data format and software conversion from Objectivity/DB to Oracle
  - Physical media migration from StorageTek 9940A to 9940B tapes
- **Took ~1 year to prepare; ~1 year to execute**
- **Could never have been achieved without extensive system, database and application support!**
- Two experiments – many software packages and data sets
  - **COMPASS** raw event data (300 TB)
    - Data taking continued after the migration, using the new Oracle software
  - **HARP** raw event data (30 TB), event collections and conditions data
    - Data taking stopped in 2002, no need to port event writing infrastructure
  - In both cases, the migration was during the "lifetime" of the experiment
  - System integration tests validating read-back from the new storage

Data Management / Access Policies

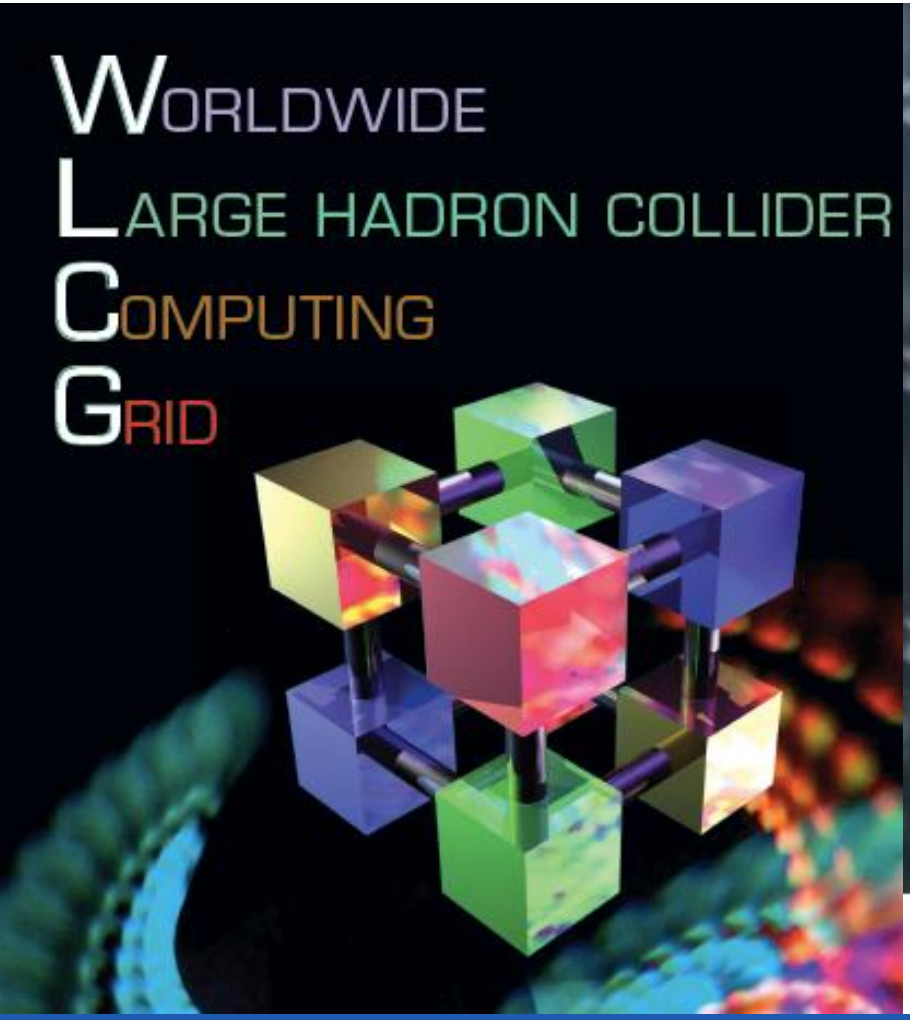

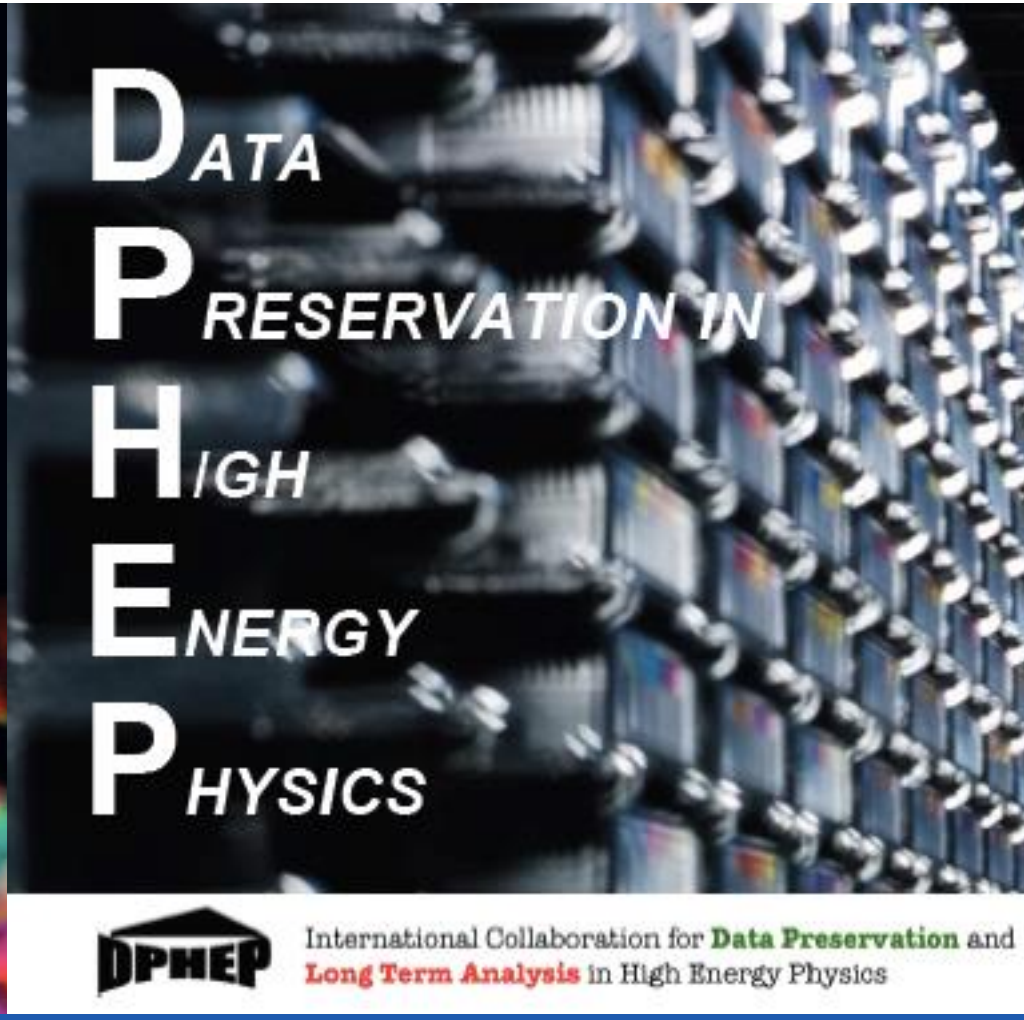

You can't share data, nor re-use it, unless you have preserved it!

2020 Vision for LTDP 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**

What is?

- Preservation
  - **Data preservation** refers to the series of managed activities necessary to ensure continued access to digital materials for as long as necessary.
- Curation:
  - **Digital curation** involves maintaining, preserving and **adding value** to digital research data throughout its lifecycle.
- Stewardship:
  - **Even more** – including decisions on what data to preserve, what is the necessary meta-data (and perhaps also data management during active life of the data).
  - (From cradle to grave, according to EU HLEG report claiming a missing 500,000 data scientists)
  - 5% “total project” tax proposed (and disputed by some)

ISO 16363 certification of CERN

- ISO 16363 follows OAIS breakdown:
  3. **Organisational Infrastructure;**
  4. **Digital Object Management;**
  5. **Infrastructure and Security Risk Management.**

- Many of the elements in 3) and 5) covered by existing (and documented) CERN practices
  - **Some "weak" areas – being addressed – include disaster preparedness / recovery (together with EIROForum)**
  - **And we haven't really started to address 4) yet…**

➢ **Next step is "stage 1" external audit to high-light those areas requiring attention**
  - **May just be a question of documentation, e.g. CERN is not going to change its financial practices (MTP etc) as a result of ISO 16363!**

Bit Preservation: Steps Include

- Controlled media **lifecycle**
  - **Media kept for 2 max. 2 drive generations**
- Regular media **verification**
  - When tape written, filled, every 2 years…
- **Reducing** tape mounts
  - Reduces media wear-out & increases efficiency
- Data **Redundancy**
  - For "smaller" communities, a 2nd copy can be created: separate library in a different building (e.g. LEP – **3 copies at CERN**!)
- **Protecting** the physical link
  - Between disk caches and tape servers
- Protecting the **environment**
  - Dust sensors! (Don't let users touch tapes)

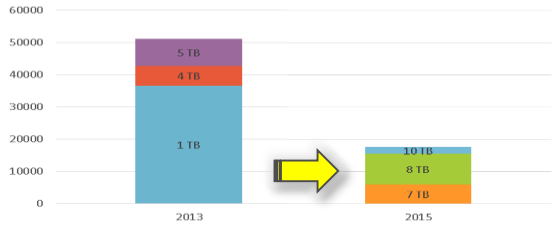

Constant improvement: reduction in bit-loss rate: 5 x 10^{-16}

Organisational Infrastructure

3.1	**Governance & Organisational Viability**	**Mission Statement, Preservation Policy, Implementation plan(s) etc.** **Operational Circular, DPHEP Reports**
3.2	**Organisational Structure & Staffing**	**Duties, staffing, professional development etc.**
3.3	**Procedural accountability & preservation policy framework**	**Designated communities, knowledge bases, policies & reviews, change management, transparency & accountability etc.** **Generic descriptions refined by project DMPs**
3.4	**Financial sustainability**	**Business planning processes, financial practices and procedures etc.**
3.5	**Contracts, licenses & liabilities**	**For the digital materials preserved…**

Infrastructure & Security Risk Management

5.1	Technical Infrastructure Risk Management	Technology watches, h/w & s/w changes, detection of bit corruption or loss, reporting, security updates, storage media refreshing, change management, critical processes, handling of multiple data copies etc
5.2	Security Risk Management	Security risks (data, systems, personnel, physical plant), disaster preparedness and recovery plans …

Digital Object Management

4.1	Ingest: acquisition of content	
4.2	Ingest: creation of the AIP	Archival Information Package
4.3	Preservation planning	
4.4	AIP Preservation	
4.5	Information management	“FAIR” etc
4.6	Access management	

The plan is to address these after metrics 3 & 5…

Need to agree on scope: only “Open Data”?

Open (Linked) Data

★ Available on the web (whatever format) but with an open license, to be Open Data

★★ Available as machine-readable structured data (e.g. excel instead of image scan of a table)

★★★ as (2) plus non-proprietary format (e.g. CSV instead of excel)

★★★★ All the above plus, Use open standards from W3C (RDF and SPARQL) to identify things, so that people can point at your stuff

★★★★★ All the above, plus: Link your data to other people's data to provide context