ASPIRE

# Middleware and Managing Data and Knowledge in a Data-rich World

SEPTEMBER 2012

# Introduction

ASPIRE - **A** **S**tudy on the **P**rospects of the **I**nternet for **R**esearch and **E**ducation

The ASPIRE foresight study has been exploring the implications of potential developments of the Internet up until 2020 and assessing their impact for the Research and Education networking community.

In May 2011, a consultative workshop was held to ascertain what the community considers to be the four topics that are most likely to have a significant impact on the sector.

The topics chosen as a result of the workshop were:

- Middleware and Managing Data and Knowledge in a Data-rich World
- Cloud Services
- Adoption of Mobile Services
- The Future Roles of NRENs

Four panels of experts were convened during the latter part of 2011, and worked until the spring of 2012, gathering material and reaching a consensus on the major issues.

This document is the work ASPIRE panel on:

## Middleware and Managing Data and Knowledge in a Data-rich World

The conclusions and recommendations from each of the panels will be discussed in a second ASPIRE workshop in September 2012. The workshop will validate the work of the panels and determine a community strategy for the future.

The ASPIRE study team at TERENA wish to express their sincere thanks and appreciation for the work undertaken by the panel members and leaders.

John Dyer Magda Haver

The ASPIRE foresight Study was funded from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 238875, relating to the project 'Multi-Gigabit European Research and Education Network and Associated Services (GN3)'.
TERENA is solely responsible for this publication, which does not represent the opinion of the European Community; nor is the European Community responsible for any use that may be made of this report



back to contents page

# Contents



back to contents page

back to contents page

back to contents page

back to contents page

back to contents page

# 1 EXECUTIVE SUMMARY

This report explores the important aspects of data handling and storage in the context of future research networks and the associated services. The study encompasses networking requirements, storage, middleware, data policies, and data origin, each of which is considered from the standpoint of five disciplines: Genomics, High Energy Physics, Digital Cultural Heritage, Radio Astronomy, and Distributed Music Performance.

From the specific requirements of each discipline, it has been possible to identify some common requirements for National Research Network (NRENs), GÉANT, and others involved in providing network connectivity:

- to collaborate with user communities to ensure the networking requirements associated with the data deluge are well understood. Adequate network services need to be put into place in a timely manner and in an economically viable fashion. Aspects of the speed of provision, throughput, privacy, persistence of connection, and other important parameters need to be addressed;

- to define standardised datasets in order to profit from economies of scale, which would follow from the availability of generic, cross-discipline middleware. Standardised datasets, metadata, and middleware and applications for easier data accessibility need to be defined. A common metadata standard that takes into account multi-disciplinary use of data needs to be adopted;

- to adopt a globally recognised Authentication and Authorisation Infrastructure (AAI) based on recognised standards for the exchange of assertions and security tokens that can be utilised by all user communities, e-Infrastructure providers, and ICT service providers. Examples include eduroam®, a system dedicated to authenticating network access, and eduGAIN, which is increasingly being used to connect interfederations;

- to create common mechanisms and procedures to enable disciplines to certify and authenticate data. This could include automated signing of data before they leave instrument, whether they are large devices such as the LHC or SKA, or smaller devices such as cameras or scanners;

- to facilitate collaboration between disciplines in order to create common policies, procedures, and tools to assist in the curation of data and its selection for preservation.

It has been an interesting exercise to gather experts from a diverse range of disciplines. The panel members are people who work day-to-day with ever-increasing datasets. They already see the new problems that are arising as the world's store of data grows at an ever-increasing rate. Those involved have found it a revelation that so many requirements have common threads, despite the diversity of the disciplines studied.
The challenge is to ensure collaboration between originators of data, users, and those who preserve and provide access to the data, and to address these requirements in an efficient manner.

back to contents page

# 2. INTRODUCTION

Data are being created and collected at an ever-increasing rate. This is partly because there are an increasing number of devices in the world that are capturing data. These devices range from small sensors that continuously generate a few kilobytes of data per day, to digital cameras that create images of a few megabytes, to large experimental devices, such as LOFAR and LHC that create fifteen petabytes per year, and to SKA, which plans to generate up to one exabyte per day.

Life sciences and humanities have made their needs and requirements clear for the handling of their rapidly increasing data, which currently stands at fifteen petabytes. Add to this, the other data-sources, such as smaller experiments, surveys, e-books and the creation of documents in general, digital art, and twitter and blog information, and the scale of the problem becomes apparent.

The "data challenge" or "data deluge" is being tackled at national, international, and global level. The European Commission has given the green light for a large number of data-related projects. The recent EUDAT [18] has raised expectations for the creation of a European Data Infrastructure. The United States government has launched the "Big Data Initiative", which will expand the arsenal of tools for handling massive data, including collection, storage, preservation, management, analysis and sharing of data, and will fund projects to manage data in meaningful ways.

In 2010, the European High Level Expert Group on Scientific Data published its report, "Riding the wave – How Europe can gain from the rising tide of scientific data" [17]. This report discusses the enormous growth in the amount of data, the changing ways in which data are used, the trust needed between all stakeholders, and the funding requirements. It provides a vision and recommendations for a global framework for data handling.

back to contents page

# 3. ASPECTS OF DATA AND DATA HANDLING

The panel looked at what aspects of data are important for the data they create and/or handle in their disciplines and identified four main categories: networking, middleware, data policies and data origin.

## 3.1 Networking

### 3.1.1 Bandwidth requirements

Three types of data transfer have been observed: SKA/HEP transfers, Human Genome Project/Digital Cultural Heritage transfers, and music performance transfers.

**SKA/HEP Network Model**

For astronomy, (SKA) and HEP, the organisation of the data is similar. Raw data are "inside" the instrument, meaning that the raw data are initially in a central repository at the place where it was collected. Data prepared for further handling is organised in a very structured, hierarchical way (Tier-0, Tier-1, Tier-2). Hence, the data streams are rather controlled, in the sense that the transfer of the data is defined between the different parts of the tier structure. In astronomy, the bandwidth requirements between the data archive and regional user centres are estimated to increase to 100 Gbps, once SKA data becomes available.

**HG-DCH Network Model**

Personalised medicine will create an enormous amount of data in the near future. Many individuals, both volunteers and patients, will have their genomes sequenced. It is likely that, for ethical and legal reasons, these data will remain at the hospital or other institution where they were generated. However, to be useful, these data will need to be compared with the reference datasets associated with the Human Genome Project, which are located at a small number of institutions around Europe, particularly the EMBL-EBI in Cambridge, United Kingdom. At the moment, these data need about fifteen petabytes of disk to serve them, but they are doubling in size approximately every eighteen months. Whether these data will be moved to the compute or if the compute will be moved to the data will depend on the ethical and legal constraints as well as on the relative cost of storage and bandwidth.

There will also be other modes of network utilisation, including the international exchange of reference datasets in order to share the data-collection task and to provide disaster recovery. There will also be periodic peaks of utilisation when the datasets associated with projects such as the Thousand Genome Project and the Cancer Genome Project are published. Finally, these reference datasets are used by a vast community of academic scientists, clinicians, policy makers, and commercial scientists across Europe and the world. Although EMBL-EBI does not authenticate these users, analysis of IP addresses suggests that this community may already number as many as three million and it is growing fast.

By 2020, we estimate that the collections of reference data will be in the exabyte range and that a substantial proportion of this will need to be transferred regularly, at least between reference sites, if not more widely.

back to contents page

The Digital Cultural Heritage (DCH) world may be a bit different in size and the type of interaction between the depositories of data, but their data are distributed around the world, and rather large data-access and transfer capability between DCH institutions is needed. Researchers will need access to data of an intermediate size, and millions of people (the public at large) will need access to data for many small transfers. The amount of data that will become available online is not yet known or even estimated, but it grows rapidly as digitalisation projects emerge and start to deliver. The high growth-rate of the amount of data from DCH is also due to the emergence of new, high-resolution scanning technologies, the use of 3D scanning, and the digitisation of video (e.g., films, documentaries, and the virtual-reality representation of archaeological sites).

**Musical Performance Network Model**

The needs for network bandwidth for music performances are rather different. In this case, there is not a huge transfer of data, but the latency of the network is very important. DVTS and Conference XP software, H.323 standard Video Conferencing and, to a lesser extent, Skype are currently used for musical education at a distance. Recently, LOLA middleware proved to be a solution to the latency problem that was encountered with all of these applications. When full LOLA functionalities are used, the bandwidth requirement can increase to 7 Gbps. Another requirement is that the latency does not go beyond 1 ms per 100 km, that there is zero packet loss, and that the network jitter is very stable (< 3ms at 30 fps).

## A general concern shared by all disciplines

Currently, for all of the observed disciplines, the network link capacity/usage is often below the required bandwidth for a variety of reasons: the network links are not available where they are needed, the network links are too expensive, and not all institutions are allowed to connect to the national research network and hence, to the international research network.

### 3.1.2 Storage, mirrors, preservation, disaster recovery

For the Human Genome Project, many sites will need to store exabyte-quantities of data. Not all datasets will need long-term preservation and a social infrastructure is needed to decide about the preservations issue for the different types of data. Mirrors will be used when adequate bandwidth is not available. A split topology of the infrastructure of the reference data of the HG project is set up to guarantee disaster recovery.

A study for a global preservation scheme is on-going for DCH. In addition, a social infrastructure will deal with preservation issues. Mirrors will be used for performance reasons and disaster recovery.

In HEP, data are currently stored, at least for the LHC experiments, in a tier-structure, in which each organisation takes care of its own storage needs. Tier-0 and Tier-1 centres have to archive data. All data are replicated often and at several places. At present, long-term preservation of data (with all the related aspects of software and knowledge) is not currently considered.

Astronomy data is already enormous and it is very expensive to store all of the raw data or to secure long-term preservation of all stored data. Depending on the project (e.g., eVLBI, ALMA, LOFAR, or SKA) not all raw data are stored but final products are derived from the raw data and stored in an archive.

### 3.1.3 Content

The content of the data is, of course, very dependent upon the discipline.

The Elixir Core Data Collections cover the molecules of life, including DNA, RNA, proteins, small molecules of biological interest, and also include scientific literature in electronic form and the links between the latter and the data collections.

In the DCH sector, the majority of the content is composed of digital images but also includes text, audio, and 3D models. Different formats are used. All data are accompanied by metadata that usually follow a standard.

back to contents page

In radio astronomy, the raw data are the RF/IF data from radio receivers and are transformed by the correlator in a set of visibility data or a "measurement set", which is then processed into a data-cube, image, or similar data product. The RF/IF data has associated metadata that includes synchronisation, timing, and essential ID information. Data models for radio astronomy have been standardised.

In HEP, the raw data are delivered by the LHC particle detectors of the four experiments. The detectors register the passage of these particles with a vast number of sensors and, finally, a digitised summary of this is recorded, and is referred to as an "event".

### 3.1.4 Cost aspects

All disciplines mention cost considerations for obtaining the necessary bandwidth and for storing data in the short or medium term and for long-term preservation.

Costs considerations for bandwidth are diverse:

- bandwidth that is now available for free may incur tariffs in the future;
- bandwidth requirements are sometimes very high and will lead to higher costs as these requirements are not included in the standard offerings of the research networks;
- sometimes dedicated connections (lightpaths) are necessary and have to be paid for;
- some parts of a country or region have more expensive connections than others. This can influence the place where data are stored;
- the last mile of the connection can bring considerable costs.

Cost considerations for storage are common among most of the disciplines. Disciplines that traditionally had a small amount of data and are now changing at a rapid pace, and also have the new burden of defining the organisation of their data.

## 3.2 Middleware

Middleware is a vast domain and all of these disciplines already use middleware for access to their data. A lot of middleware will be needed in the area of authentication of persons and sources, including federated access. For the ELIXIR Data Collections in the genomics discipline, data are usually made available by complete downloads or through interactive webpages and realised with locally developed middleware solutions. In the future, it is expected that a number of generic solutions, probably via the cloud mechanism, will be developed.
In the DCH sector, the most-used middleware is that of the harvesters of metadata, thanks to the existing protocol, OAI-MPH, for metadata harvesting. There are also special middleware tools for the indexing and retrieval of audio-visual sequences. In the future middleware will be developed for the automatic or semi-automatic enrichment of cultural metadata belonging to distributed (multimedia) repositories.

HEP has developed very important middleware for handling data within the framework of WLCG (the Worldwide LHC Computing Grid), but also still uses many experiment-specific products.

The models for radio astronomy data have been standardised and the middleware currently in use operate using those standards.

back to contents page

For distributed music performances, middleware was almost non-existent. However, the LOLA audio-visual streaming system has been recently developed opening the way for high-quality distributed music performances. Although LOLA is primarily application-level software, because it interacts directly with the data-transport stream, it could be regarded middleware that controls the network.

### 3.3 Metadata

Metadata describes what your data are, how others can use them, and how they are physically put together. It allows other researchers to determine how, when, where, why and by whom the data were originally produced or collected. Metadata is generally broken down into three different kinds: descriptive metadata, administrative metadata and structural metadata. For data to be successfully stored, managed, shared, and re-used throughout their life cycle, documentation and description of the data is essential. Documentation and metadata requirements will vary according to the discipline and the nature of the research that is being carried out. Metadata and other necessary documentation to make data usable by and discoverable by others should be identified during data-management planning.

In the disciplines involved in this study, radio astronomy, high energy physics, and the digital cultural heritage have metadata attached to their data. The cultural heritage sector has been working on the definition of highly structured metadata and controlled vocabularies. The simplest type of metadata used by cultural institutions to associate to the data is a textual description of the digital representation of the cultural object. Two major formats to represent this text in a structured way are XML and RDF. Dublin COR (DC), ISO standard 15836:2009 is the most widely-used metadata standard.

The physics-data of LCH is also accompanied by metadata and the same holds true for radio astronomy.

An exploration of these and other disciplines quickly reveals that the definition and usage of metadata is generally undertaken on a per-discipline basis, with little consideration for cross-disciplinary application or usage. A universal but extensible approach to middleware that is capable of conveying research and semantic attributes is required.

### 3.4 AAI

There is consensus on the principle of an AAI, but there is much disagreement about the definition and realisation of such an AAI.

In the entire ecosystem of the infrastructure (networking, computing, and data), the topic of authentication and authorisation, and the availability of an AAI (Authentication and Authorisation Infrastructure) has been the source of much debate and many requirements. Authentication and authorisation is needed to make infrastructures as secure as possible. AAI systems should ensure that only authorised access can be made to protected resources, and that even authorised access can be traceable using audit trails.

As the HEP part of this report describes, the LHC community adopted an AAI based on X.509 certificates issued by a globally distributed network of certification authorities to provide the trust relationship and authenticate members of its Virtual Observatory (VO). The Human Genome Project does not make use of authentication and authorisation schemes, but has other mechanisms for quality assuring data depositions. In the future, when dealing with collections that have to be protected for legal or ethical reasons, a lot of attention will be given to an adequate AAI.

DCH does not use an AAI at this time, but will need such an infrastructure to allow users to make annotations or to provide access to data that is protected, for example, to preserve authors' rights.

back to contents page

In general, everyone agrees that there is a need for a globally accepted AAI system that is usable across disciplines and e-infrastructures, and that is efficient and easy to use. AAI systems must not create a barrier for access to e-Infrastructure resources. The e-Infrastructure Reflection Group (e-IRG) has made a recommendation for such an AAI system in its White Paper 2011 [18].

The federations for authentication and its global organisation, eduGAIN [19] is the GÉANT service that interconnects identity federations around the world and is an excellent move in the direction of a globally accepted AAI system.

## 3.5 Data Policies

### 3.5.1 Availability of data (public, private, and cost)

Current policies on data access tend to be discipline-specific. However, there is now a tendency to move to "open data". In astronomy, American politicians demand as much data as possible be made public. In Europe, there is a belief that "publicly funded data should be made public". "Open data" is in the public domain, meaning that it can be studied by a much larger population of users than if it were private. The consequence is that there is an increased chance that someone may make an insightful discovery. However, open access to data may contradict the funding policy of the project that generated the data or the legal, ethical, or societal implications of the data.

For example, radio astronomy has traditionally operated using the principle of 'Open Skies'. This principle allows access to radio telescopes, irrespective of the country of origin of the request, and only subject to peer review of the scientific value of a proposal. This model has proved difficult to maintain in the era of very large and expensive radio telescopes with large output datasets. The bodies that fund the instruments expect privileged access to the instrument that they are funding for their own scientific community. In new telescopes, the data are accessed via regional data centres in the countries or regions of funding. How this change in operating mode will manifest itself remains to be seen.

The human genome was decoded using charitable funds and belongs to us all. It should therefore be made freely available, without the need for authenticated access. This right of open access may apply to many other, similar datasets. It should be noted that this does not preclude creating value-added services based on these datasets, which could be commercialised. The pharmaceutical industry, in particular, is very interested in this approach as a means of pooling resources and sharing costs in non-competitive areas. Such systems will feed back into the public domain and benefit the community at large. Data in the DCH sector are generally subject to Intellectual Property Rights (IPR) rules and constraints. High-resolution images are accessible only to authorised users while open access is provided for low-resolution images.

Data-confidentiality is not usually a problem in HEP. However, there is a "long way to go" before data are actually made public.

Storage and curation of data is expensive, so the concept of charging for access to certain types of data may be required in order to provide a sustainable infrastructure.

### 3.5.2 Data-preservation and replication

Data preservation and replication, particularly for the long term, can pose huge problems for creators of data. They have to provide reliable physical storage itself, cope with changes of storage media as technology develops, preserve the related metadata, undertake curation of the data, and ensure the continued availability software that is capable of accessing the data in a meaningful way.

back to contents page

Since our ability to create data exceeds our ability to store it, it will not be possible to preserve all of the data that is created. Some approach must be found to decide what data to preserve and what data to discard. The selection of what to preserve is currently undertaken by human appraisal, but innovative technologies to support the selection process would be beneficial in increasing the efficiency of the task. The same is true for the active and on-going management of data through the process of curation. Somebody or something needs to decide what to curate and what not to curate.

## 3.6 Data Origin

### 3.6.1 Authentication of source

For any data, its **integrity** and **source authentication** are important.

At present, there is little or no generalised mechanism of data-source authentication. For the data of large instruments, there is the metadata that provides information concerning the data, and this can be used to create the necessary authentication. In many cases, it should be possible to ensure that the data are authenticated by a unique digital signature during the creation process, unequivocally linking the data to its creation source or instrument. There are segments of the community that believe that even if such a process only adds a tiny amount to the cost of the data-creation, this would result in an additional cost that is too large for the research community to bear.

For many of the major biological data collections, arrangements are in place for allocating globally unique identifiers to coincide with publication. However, datasets are increasingly being generated for research use and are not destined for peer-reviewed publication. Hence, they do not receive these unique identifiers, and an alternative authentication process is required.

There has been little in the way of control over the authenticity of the historic data of astronomy that is used for virtual astronomy processes, such as research on time-series variations. The saving factor is that such data are rather impenetrable to the non-astronomers and consequently, it has remained quite reliable. Similarly, it is considered that the data in the human genome datasets is not usable for people who are not working in this scientific field, and consequently, can be regarded as "safe".

There is some concern about data produced by non-professional projects, because of the possibility that data could be faked and injected into the results, alongside the legitimate data.

### 3.6.2 Data integrity

For some data, it would be quicker and less expensive to recreate the data than to store it. However, this is very dependent upon the discipline and the way the raw data are obtained. HEP cannot lose its raw data because the circumstances of creating a particular dataset can never be reproduced exactly. However, data integrity is extremely important for them and checksums are used extensively throughout their systems to ensure the integrity of the data.

back to contents page

# 4. FUTURE DATA AND DATA-HANDLING

The ever-increasing creation of data and the growing interest in undertaking collaborative, multidisciplinary exploration of this data leads to new discoveries, new knowledge, and new possibilities. This multidisciplinary approach to data exploration is likely to accelerate in the future. The consequence is that access will be required to larger datasets by more distributed teams and this will increase the challenge for the network infrastructure and services, which are used for sharing and distribution.

Chapter 3 has already shown a trend in the growth of the amount of data and a larger demand for the services that will be fundamental to the success of the involved research. Chapter 5 will provide more details of these future requirements for each discipline. An extract of some of the future requirements is presented below.

**Drastically increasing amount of data**
Astronomers have currently access to around one petabyte of archived data and it is anticipated that this will grow to 60 petabytes of archived data by 2020.

The reference datasets associated with the Human Genome Project currently amount to fifteen petabytes and it is estimated that by 2020, the reference data collections will be in the exabyte range.

Other disciplines are also expecting some large increases.

**Increasing bandwidth requirements**
The bandwidth of data links needed to handle astronomical amounts of data will need to increase from the current maximum, 30 Gbps, to at least 100 Gbps within the coming decade.

Digital Cultural Heritage data is growing fast and this will lead to a large number of varying bandwidth requirements.

Distributed music performance is now in its infancy, but recent experiments show that bandwidth requirements are going up to 7 Gbps, with very stringent requirements on latency, network jitter and packet loss.

For the Human Genome Project, a substantial proportion of the data collections (in the exabyte range) will need to be transferred regularly, at least between reference sites if not more widely.

**Increased use of middleware**
All disciplines expect an increase in the use of middleware in order for them to handle the vast quantities of data they are amassing. It is clear that standardisation is a need to create generic middleware that is usable for the whole user community in a given discipline. Easy access to the middleware via cloud services is envisaged.

back to contents page

**Harmonisation of metadata**

To be able to make efficient use of the very large amount of data that is available on an increasing number of topics, metadata attached to all data is necessary. However, the disparity in the definition of metadata content, the diversity of metadata standards, and the non-existence of metadata for multi-disciplinary use means that the harmonisation of metadata is a significant problem that needs to be addressed.

**Global AAI**

For the future, there is a strong need for a globally accepted AAI that goes across disciplines and e-infrastructures, that is efficient and easy to use, and that does not create an extra barrier for access to the resources of the e-infrastructure.

**Availability of data**

In the future, the principle of "open data" can be used as a starting point for availability of research data. However, adequate policies have to be put into place to protect ethically or socially sensible data, and to cope with the legal issues of ownership and intellectual property rights.

In light of the rising costs associated with the growing amount of data, it might become necessary to charge for added-value services associated with certain data.

**Preservation and replication of data**

With the growing amount of data, it will not be possible to preserve it all, and decisions about what to preserve have to be made. Currently this selection is mostly done by human appraisal, but innovative technologies to support the selection process would be beneficial for increasing the efficiency of that task. The same holds true for the curation of data.

**Authentication of source**

It is agreed that authentication of the source of data will become much more important than it has been until now. In multi-disciplinary research, the creation of new data from existing data will be much less controlled than it was with the original data. Availability of data to a large public may also lead to the creation of non-controlled new data. Hence, mechanisms and procedures have to be established to ensure that the source of the data is known.

back to contents page

# 5. DESCRIPTION OF DATA-RELATED ISSUES FOR THE SELECTED DISCIPLINES

## 5.1 Human Genome Project

### 5.1.1 ELIXIR: A sustainable infrastructure for biological information in Europe

ELIXIR is a European Strategy Forum - Biological and Medical Sciences Research Infrastructure (ESFRI BMS RI) with a remit that covers the provision of substantial data resources to all of the life sciences and medical research communities, including industrial research - in particular, the pharmaceutical and biotechnology sectors.

### 5.1.2 Networking

#### 5.1.2.1 Bandwidth requirements

ELIXIR is a distributed infrastructure constructed from a hub and nodes. The hub will coordinate the infrastructure as well as host the ELIXIR Core Data Collections and Services (ECDD&S). Currently, these take up fifteen petabytes and are doubling in size every year.

Bandwidth requirements will depend on the modality of the interaction. It will be necessary for ELIXIR to support, at least, the following modalities.

- Very large data deposition at the ELIXIR Hub by a few collaborators
  EMBL-EBI works with a small number of partners who generate very large datasets, which they deposit. A good example of this would be the Sanger Institute, which was one of the centres that took part in the Human Genome Project. These datasets are currently of the order of petabytes per year but will move to exabytes per year over the course of the coming decade.

- Large data exchanges between the ELIXIR Hub and International Partners
  The bio-molecular data collections, for which EMBL-EBI is responsible, and which will form the core data collections for ELIXIR, are managed by long-standing international collaborations. Typically, the data-acquisition and curation tasks are split between EMBL-EBI and other members of the collaborations, who are located in mainland Europe, the United States, and Japan. Exchanges of data take place daily in order to ensure that all of the sites are kept up to date and to provide resilience against regional disasters.

- Periodic publication of large datasets
  ELIXIR will be responsible for maintaining the data collections for international flagship data-generation projects such as The Thousand Genome Project and The Cancer Genome Project. Periodically, a new version of these data collections will be released to all members of the project at the same time. Currently, this

back to contents page

would correspond to hundreds of gigabytes, downloaded by a few dozen sites. This is likely to increase as the projects become more numerous and larger.

- Small depositions from incidental data-generators
  All over the world, scientists generate data in the course of their research that is likely to be of wider interest. Many of the main scientific journals encourage or even insist that these data be deposited in the international data collections as a prerequisite of publication. This will probably increase substantially as well.

- Diverse modes of service utilisation by the user base
  It is now impossible to carry out experimental biology without regular recourse to the data-collections that are currently managed by EMBL-EBI and will form the core of ELIXIR. The availability of new experimental techniques and the projected rapid increase in the size of the datasets will undoubtedly increase this practice dramatically.

The nodes will also need these modalities. A few of the larger nodes may need all of them, while others will need a subset of them.

#### 5.1.2.2 Storage

The amount of storage at the Hub will continue to increase very fast. Currently it stands at fifteen petabytes, with a doubling time of slightly over a year. In addition to the Hub, many nodes will construct very large storage infrastructures. By 2020, many sites will be handling exabyte quantities of data.

#### 5.1.2.3 Mirrors

Mirrors will be used when bandwidth is not sufficient to provide acceptable performance.

#### 5.1.2.4 Preservation – lifetime

Not all genomics data will need to be preserved for the same length of time. Some of it, such as the data used to create the reference copy of The Human Genome Project is part of the record of science and will need to be kept for a long time. Other datasets will have a much shorter life. Sequencing of the whole genome is already being used as a tool for clinical diagnosis, and once the diagnosis has been made, the data are destroyed. These two cases represent the extremes. In addition to these two, there will be datasets covering many other eventualities. Part of the role of ELIXIR is to put into place the necessary social infrastructure to allow rapid and appropriate decision-making on these issues.

#### 5.1.2.5 Preservation – disaster recovery

ELIXIR data will be held in multiple data centres that are geographically distributed in order to provide resilience against local disasters. Protection against regional disasters will be provided for by daily data exchanges with international partners.

#### 5.1.2.6 Content

The ELIXIR Core Data Collections cover the molecules of life, including DNA, RNA, proteins and small molecules of biological interest, as well as many other kinds of data. Of utmost importance is the scientific literature in electronic form. The links between this literature and the data collections is one of the most important tools in modern biological research. Typically, each collection is managed under the auspices of global collaborations that involve Europe, the United States, and Japan. Many of them have been running for decades.

The following list details the major data-collections that are currently held by EMBL-EBI. Functionally equivalent copies of most of these are also held at similar organisations in Japan and the United States. Currently, it takes

back to contents page

approximately fifteen petabytes of magnetic disk to deploy these collections. The doubling time is about eighteen months and has been so for at least a decade.

1. **Ensembl:** a joint project with Sanger Institute - high-quality annotation of vertebrate genomes
2. **Ensembl Genomes:** an environment for genome data from other taxons
3. **1000 Genomes:** a catalogue of human variation from major world populations
4. **EGA (European Genotype Archive):** genotype, phenotype and sequences from individual subjects and controls
5. **ENA (European Nucleotide Archive):** all DNA & RNA, nextgen reads and traces
6. **ArrayExpress:** an archive of transcriptomics and other functional genomics data
7. **Expression Atlas:** differentially expressed genes in tissues, cells, disease states and treatments
8. **UniProt:** an archive of protein sequences and functional annotation
9. **InterPro:** an integrated resource for protein families, motifs and domains
10. **PRIDE:** a public data repository for proteomics data
11. **PDBe:** protein and other macromolecular structure and function
12. **ChEBI:** chemical entities of biological interest
13. **ChEMBL:** bioactive compounds, drugs and drug-like molecules, properties and activities
14. **IntAct:** a public repository for molecular interaction data
15. **Reactome:** biochemical pathways and reactions in human biology
16. **Biomodels:** mathematical models of cellular processes
17. **GO:** Gene Ontology, consistent descriptions of gene products
18. **CiteXplor:** bibliographic query system

#### 5.1.2.7 Cost aspects

Traditionally, biology experiments only produced relatively modest amounts of data and the cost of the Internet Technology necessary to deal with them was only a small part of the total cost of the experiment. In Europe, the core data collections are managed by the European Bioinformatics Institute (EBI), which is an outstation of the European Molecular Biology Laboratory (EMBL). One of the main purposes of ELIXIR is to develop a new organisation and new funding streams that can take on the increasing costs, and to ensure that the data collections remain available to European scientists in the future.

### 5.1.3 Middleware

#### 5.1.3.1 What do you have now?

Currently, most of the ELIXIR Core Data Collections and Services are deployed, without access control, from the webservers of EMBL-EBI and other organisations. These are generally made available as complete downloads and through interactive webpages using web services. As far as the user is concerned, the latter approach is indistinguishable from cloud deployment using the 'Software-as-a-Service' deployment model (SaaS), although it is done with locally-developed, bespoke middleware solutions rather than with cloud software per se.

#### 5.1.3.2 What do you want?

It seems inevitable that, in due course, all of the ELIXIR Core Data Collections and Services will be moved to a generic middleware solution of some kind. Indeed, EMBL-EBI is a member of the Helix Nebular Science Cloud Project, which is investigating cloud solutions for scientific computing in the commercial sector. In addition, EMBL-EBI is also investigating a number of generic solutions using experimental deployments of various services and production processes.

back to contents page

#### 5.1.3.3 Generic middleware

Whether the solution that is chosen for this is commercial or in the public domain, and which particular one is chosen will depend on the development of the marketplace. It seems likely that the Helix Nebular Science Cloud Project will catalyse the development of solutions in this area.

#### 5.1.3.4 AAI for access

At the moment, most of the data are available without authentication through a web interface. This is entirely appropriate for data, such as the human genome data, that are generated with public money and, in some way, "belong to us all". For data-collections where there are legal or ethical reasons for protecting the data, AAI will be needed (see below).

#### 5.1.3.5 AAI for creation

The large data-collections listed above are managed by international consortia that have been in existence for many decades. These consortia have developed data quality assurance mechanisms that are well known and trusted in their user communities. Typically, these involve the cooperation of publishers, although the details of the submission process are different for the different collections. In essence, before an investigator can publish findings that are supported by a large dataset, he or she has to obtain a globally unique identifier from the database operator. While this has proved satisfactory thus far, the increasing prevalence of projects where data-generation is the objective rather than a step on the way to a scientific publication means that other mechanisms will need to be (and are being) developed. Bearing in mind that the depositors receive little direct benefit from sharing their data, it is important that the process of sharing does not impose an onerous burden on the provider of the data.

### 5.1.4 Data policies

#### 5.1.4 .1 Availability of data

ELIXIR has an open-access remit so all of the data collections will be available without restriction, wherever possible. The one exception to this is when access to the data has ethical, legal, or societal implications (ELSI). In this case, there will be a process by which investigators can be certified as appropriate people to use the data. Typically, this will be performed by a data-access committee. ELIXIR never owns the data collections for which it is responsible; it acts as a custodian on behalf of third parties. These third parties may stipulate particular requirements for the authorisation and authentication of those who use the data.

#### 5.1.4 .2 What to preserve

Conceptually, an analogy exists between the process of scientific publication and the processes by which data are deposited in the bio-molecular data collections. Scientific results are published at the point at which the researchers feel that it is appropriate to present them to the scientific community. This is the same for data depositions in international bio-molecular data collections. At least in the eyes of the depositors, deposited datasets will have reached a level of completeness where it makes sense to "publish" them. It is these deposited datasets that are preserved. Any other data associated with the development of the deposited dataset, which is not deposited, remains the responsibility of the depositor. The deposited data are the responsibility of the custodian and is entered into the long-term archive. It is then distributed as part of the standard service.

### 5.1.5 Curation and quality control of data

The custodians of the international bio-molecular data collections, such as EMBL-EBI, expend substantial effort on curation and quality control of the data that they accept as depositions. Large numbers of domain-specific experts are employed to oversee automatic and semi-automatic curation and quality control of deposited data, and these experts also carry out extensive manual curation.



back to contents page

### 5.1.6 Origin of data

The data collections for which ELIXIR will be responsible are mainly managed by international collaborations between institutions in many different countries. These consortia, many of which have been in existence for several decades, are responsible for managing and developing the procedures by which data are deposited in the collections.

#### 5.1.6.1 Authentication of source

For many of the data collections for which ELIXIR will be responsible, special arrangements have been made for data deposition. So far, for public-domain data, it has not been found necessary to use formal authentication mechanisms for data deposition. In the case of data collections where there are ethical, legal, or societal implications (ELSI), these will have been collected under the auspices of an ethics committee and this committee will be responsible for ensuring that appropriate authentication is in place.

#### 5.1.6. 2 Charging for data

Data are deposited freely by the owners and are held by the various institutes that act as custodians. They contain data deposited by millions of scientists from most of the countries in the world. The whole infrastructure is in the public domain and built on the premise that the data remain freely available to all who want them. While this does preclude charging for the data, it does not preclude charging for value-added services constructed using the data.

## 5.2 Digital Cultural Heritage

### 5.2.1 Introduction

The digital cultural heritage (DCH) is the result of the process of creating digital data from the digitisation of physical cultural heritage and of the creation of born-digital cultural data. Examples of born-digital cultural data are digital texts, artistic creations, and virtual reconstructions of archaeological sites. The digital cultural data are used by researchers in the humanities. These researchers can remotely access information about digital cultural objects at a scientific level.

### 5.2.2 A description of data in the DCH sector

#### 5.2.2.1 Data flow

The data in DCH are mainly generated via a digitisation process. The physical cultural object is subject to 2D or 3D scanning, which reproduces the physical object in a digital format, using a particular file format. For 2D scanning, the most widely used file format is JPEG. Other file formats, such as the TIFF format, are still in use. 2D vector images of a scanned image can also be used to produce 2D animations. The formats are numerous; Scalable Vector Graphics (SVG) and Adobe Flash are the most well known. For 3D scanning, there are still several formats in use, depending on the actual result of the digitisation. The digitisation can generate either a real vector image or a still, 2D raster image with visualisation of some 3D characteristics, such as depth of the objects. The formats used to represent real vector 3D objects derived from the 3D scanning process are those already established in the Computer Aided Design (CAD) industry. Another formatting issue is related to the transmission of 3D videos, including the transmission to 3D devices. There are still different formats in use by the 'entertainment' industry. Several working groups have been created to work on standardisation, but the results are not yet fully agreed.

back to contents page

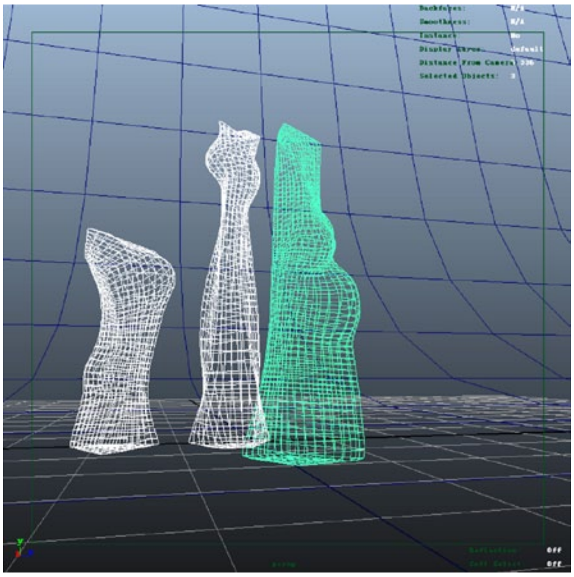

*Xiaochun Situ digital for sculpture (source: Promoter Digital Collection)*

Finally, a lot of attention is devoted to the geo-referencing of cultural data. This information is used both in the creation of specific applications for cultural tourism and for the management of cultural objects (for example, to trace the movement of a certain object during shipment for a temporary exhibition).

The case of born-digital data [14] is different. This is cultural material that is digital by origin. It can refer to textual information, such as the e-books. The formats for the representation of the e-books are many and there is not yet a 'winning' standard. In the case of digital documents, the most widely used format is PDF. Born-digital can also refer to digital photography, harvested web archives, digital manuscripts, electronic records, static datasets and dynamic data, digital art, and digital media publications.

In general, DCH data are produced with the intervention of a human being, making the cost of the generation of DCH data very high.

### 5.2.3 Networking

The networking requirements of DCH vary according to the application and the location of the cultural object. In general, DCH data are stored in the local repositories of each individual cultural institution. Those local repositories are hosted either in the data centres of the institutions or in the data centres of their service providers. Service providers can be either private or public, although the cultural sector prefers to use public resources whenever possible.

The actual capability, in terms of connectivity of the data centre, becomes a central issue for the provision of access services to the final users and to researchers. In fact, if the data centre is well connected and if the data are accessed by a 'well-connected' researcher, then the data transmission will appear to be of high performance. On the other hand, if the data centre is not well connected (e.g., the cultural institution is not on the GÉANT network), then the transmission speed will be limited, reducing actual accessibility to the data.

An infrastructure for the digital cultural heritage is still under study and the implementation phase has not yet started. There are metadata registries implemented at both European level (e.g., Europeana) and at national/ thematic levels (national cultural portals and thematic portals). The actual repositories of data are implemented by each individual cultural institution. However, there are still many medium-small cultural institutions that have not yet implemented repositories that are accessible online. The protection of high-resolution resources is an issue for the content owners. The future DCH e-Infrastructure will offer facilities to these institutions to deposit

back to contents page

their data in safe and accessible repositories. The future e-Infrastructure for digital cultural heritage will need to take the existing systems (portals and local repositories) in a federated infrastructure into account. Some European projects have begun to study the feasibility of infrastructures dedicated to the DCH sector, including DC-NET [21] and DARIAH [22].

All of the data of all of the cultural repositories are then expected to be connected in a linked data space. Linking data is a very important area of investigation in the domain of digital cultural heritage. All of the data are expected to be openly accessible to the users from all over the world, for research, education, and possibly commercial re-use by the creative industry. This means that there will be a very large number of users with different levels of access.

In this light, the following types of interaction, and the consequent bandwidth requirements will need to be taken into account for the implementation of the future e-infrastructure for digital cultural heritage. Access will be required for:

- content providers (cultural institutions) who want to use the infrastructure to store their content (data and metadata). There will be thousands of cultural institutions accessing the infrastructure, each with volumes of data in the range of tens of megabytes per image;
- researchers, students, and teachers who want to search in the metadata and download images and other data files (e.g., sound, 3D representations, and audio-visual material);
- researchers, students and teachers to add annotations;
- creative companies, in a similar way as for the researchers, but with different authorisation mechanisms;
- portals and other software that is re-directed to the infrastructure to find data and other information.

### 5.2.4 Storage: mirrors, preservation

A global preservation scheme for the digital cultural heritage is under study in a new project named DCH-RP (Digital Cultural Heritage Roadmap for Preservation). This will include curation, certification, and persistent identification mechanisms. The infrastructure is expected to work on the basis of characterisation of the federated repositories. Characterisation tools have been developed in research projects funded under the FP7 ICT programme.

In addition to the storage resources, the management of the preservation facilities provided by the infrastructure will need computing resources to run periodic quality checks to assess the quality and integrity of the data as well as to identify obsolete information.

There is a need for interaction with the owners of the data, who need to agree with the selection processes, with the preservation cycles to be applied to the data, and with the rules for erasing obsolete data. Mirrors will be used to guarantee a good response time, good performance when traffic increases, and as an additional resource for disaster recovery.

### 5.2.5 Content

Content for the DCH sector is composed of the digital collections of the cultural institutions, museums, libraries, archives, 3D models of artefacts, and virtual reconstructions of cultural places (e.g., archaeological sites).

Digital collections are repositories of digitised cultural objects.



back to contents page

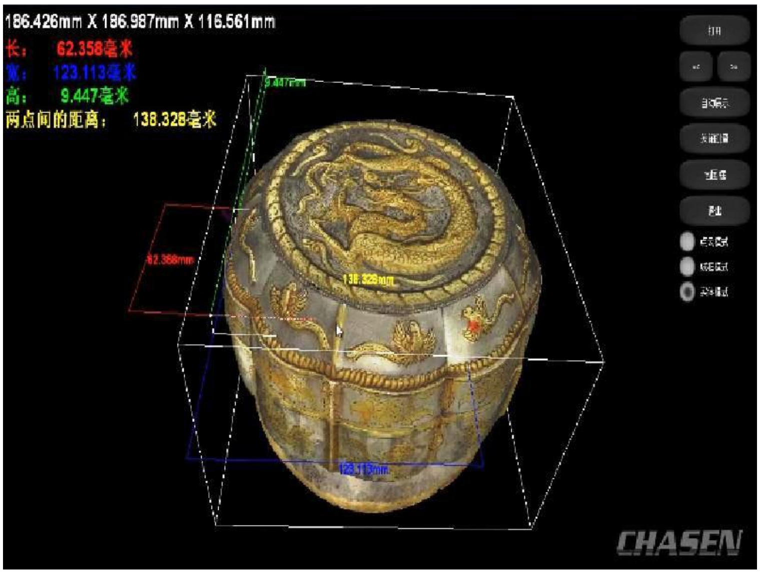


*Chasen 3D reconstruction (source: Promoter Digital Collection)*

So far, the majority of the digitisation is made of JPEG images. However, the volume of audio and audio-visual content is growing very rapidly. Audio is normally used to represent voice recording for linguistic research and intangible heritage records. Of course, audio is also used for music objects. Audio-visual content comes from performing arts recordings, films, television programmes, and documentaries. 3D models are another type of cultural content. 3D models come both from 3D digitisation of cultural objects and 3D reconstruction of cultural landscapes. Another important type of cultural content is text, which can be represented as both images (JPEG and TIFF) and as text files. All of these different types of content have associated metadata. Metadata generally follow standards, with Dublin Core (DC) being the most widely-used, cross-domain standard. In addition, each sector – museums, libraries and archives – have their own domain-specific standards. The new CIDOC-CRM standard is now being diffused, an important step ahead towards the possibility of establishing semantic links among data, because CIDOC-CRM introduces the semantic layer to DC, which is a totally flat representation of the information.

There are so many DCH repositories that it is impossible to provide a full list of those that are expected to be accessible through the future DCH infrastructure. The European digital library, Europeana[1], provides a good starting point to explore the cultural digital content available in Europe.

### 5.2.6 Cost aspects

Each institution currently pays costs in the DCH sector independently. The new DCH infrastructure will include a cost model to consider the sharing of costs among the DCH users and the exploitation of the facilities of the e-Infrastructure through ad-hoc agreements between the national authorities of research and education on one side, and the national authorities of cultural heritage on the other side.

NRENs and NGI s are not always willing to offer their services to the cultural sector. This is often because the national authorities that fund the e-Infrastructures are different from those in charge of the cultural heritage.

### 5.2.7 Metadata and middleware

The creation and maintenance of highly-structured metadata and controlled vocabularies has been the core business of cultural heritage institutions since the end of the nineteenth century [15]. However, because there are many standards to represent these structures, it is necessary to build systems capable of supporting the

1 www.europeana.eu

back to contents page

interoperability of repositories coming from different sources. The simplest type of metadata used by cultural institutions to associate with the data is a textual description of the digital representation of the cultural object. In order to represent this description in a structural way, two major formats are used: XML and RDF.

The Dublin Core (DC)[2] is the most widely used metadata standard in the DCH sector. A DC resource can be represented using RDF or XML. It can be seen as a namespace for a description of resources and it can be used to describe a single resource. DC is the ISO Standard 15836:2009.

To represent complex domains and special cultural resources, the concept of classes is elaborated and each resource has a certain range of properties. These are the ontologies that are represented as RDF Schema (RDFS) and by the Ontology Web Language (OWL). CIDOC-CRM[3] is the most relevant format for DCH ontologies. The CIDOC Conceptual Reference Model (CRM) provides definitions and a formal structure for describing the implicit and explicit concepts and relationships used in cultural heritage documentation. CIDOC-CRM is the ISO Standard 21127:2006.

The most-used middleware in DCH is represented by the harvesters of metadata through the protocol defined by the Open Archives Initiative (OAI) known as OAI-PMH (Protocol for Metadata Harvesting).

Today, metadata registers are used throughout in the DCH sector. One well-known example is Europeana, the European cultural portal managed by the Europeana Foundation, hosted by the Royal Library in The Hague. Europeana currently has fifteen million records[4] containing metadata related to digital representations of paintings, sculptures, historical buildings, books, maps, videos, sound recordings. Europeana is based on the implementation of an OAI-PMH harvester.

Middleware for the authorisation and authentication mechanisms is a very valuable resource for the DCH sector, especially considering the high value of high-resolution images.

Another element that needs to be processed at the level of middleware is the creation of semantic links among the different cultural objects, using reasoning and intelligence that is embedded in the content itself.
The tools for the indexing and retrieval of audio-visual sequences are another special case of middleware that is in use in the DCH sector. Several implementations have been realised in the last years, but a standard protocol has not yet been established.

Middleware for automatic or semi-automatic semantic enrichment of cultural metadata belonging to distributed (multimedia) repositories represents an area of rapid development and intense research. This middleware is used, for example, to add geo-references and multilingual features.

### 5.2.8 Preservation and curation of data

The selection of what to preserve is usually made by human appraisal. Innovative technologies to support the selection process would be beneficial to increase the efficiency of this task. Data are curated at the level of the individual cultural institution.

### 5.2.9 Data policies and data origin

Data in the DCH sector are generally subject to IPR rules and constraints. High resolution images are accessible only to authorised users (those images are often not available online because of lack of trust from the content providers about security issues on the network), while open access is provided for low resolution images.

2 http://www.dublincore.org/
3 www.cidoc-crm.org /
4 Towards the 30M objects 2015 target



back to contents page

Another important issue is the permanent identification of the digital objects. This is a pre-requisite for preservation, interoperability, and linked data services. Cultural institutions are working on feasibility studies on this matter, but a final, shared solution among all of the participants, at European level, is far from being established.

The localisation of the storage of DCH data is another concern for the sector, because the cultural institutions are very careful about tracking where the data are actually taken. In this light, the cultural institutions look on the emerging cloud storage services suspiciously.

The problem of the 'Orphan' works is being addressed by the DCH sector. Orphan works are the works whose authors and/or rights holders are not known. In this case, the digital object suffers from the same problem. This adds complexity because digital copies can easily be produced and it becomes very difficult to trace who is using what. A set of recommendations have been prepared by high-level experts groups appointed by the European Commission, by online registries and by other online solutions. These solutions have been encouraged and supported by the EC to provide better information to users on rights, and to facilitate clearance of these rights.

#### 5.2.9.1 Authentication of source

Authentication of source is very important in the DCH sector, both for the data (the digitised cultural object) and for the metadata (the description of the digital object that should be provided by professionals with qualifications in the domain).

To improve the knowledge about a cultural object, annotations coming from generic users can also be useful. Wikipedia is an experiment in the DCH sector that involves the final users in the creation of the encyclopaedia. There are still doubts about the value of the annotations coming from non-professional users.

AAI services would be very useful in this context because they automatically authenticate the person and decide if the annotation is authorised or not, and at which level.

#### 5.2.9.2 Data-integrity

The quality of data is selected at the time of the digitisation, usually based on the requirements imposed by the public tender that finances the digitisation project. The procedures for the assessment of the preservation of the integrity of the data are still widely used in the DCH sector. Some tests of the use of automatic checksum over large datasets are on going, but they are still at a very experimental stage.

## 5.3 High Energy Physics (HEP)

### 5.3.1 Introduction

High Energy Physics (HEP) has been pushing the boundaries in scientific data-management and computing for a long time. The Large Hadron Collider (LHC) went into production in 2009. The LHC accelerator creates high-energy particle collisions, which, in turn, create new particles that decay in very complex ways as they move through the detectors. The detectors register the passage of these particles with a vast number of sensors and a digitised summary of this is recorded as what is called an "event". The raw data per event is around a million bytes (1 Mb), produced at a rate of about 40 million events per second. To reduce the amount of data, and to keep only events that are potentially worth studying, an intelligent selection filter is applied to reduce the number of events recorded to around 100,000 per second. In a second selection stage, this is reduced again to around 100 or 200 events per second by processing the data using more specialised algorithms. These raw data are then recorded at a rate around 1.5 CDs per second (~1050 Mb/sec). Hence, the Large Hadron Collider produces roughly fifteen petabytes (fifteen million gigabytes) of data annually.

back to contents page

Thousands of scientists around the world want to access and analyse these data, so CERN is collaborating with institutions in 34 different countries to operate a distributed computing and data storage infrastructure: the Worldwide LHC Computing Grid (WLCG). It became clear very early that the size of the LHC and its associated data volumes would require a new e-Infrastructure and the WLCG has successfully met these challenges.

### 5.3.2 Networking

#### 5.3.2.1 Bandwidth requirements

Data from the LHC experiments are distributed around the globe, with a primary backup recorded on tape at CERN. After initial processing, these data are distributed to eleven large computer centres – in Canada, France, Germany, Italy, the Netherlands, the Nordic countries, Spain, Taipei, the United Kingdom, and two sites in the USA – with sufficient storage capacity for a large fraction of the data, and with round-the-clock support for the computing grid. These so-called "Tier-1" centres make the data available to over 160 "Tier-2" centres for specific analysis tasks. Individual scientists can then access the LHC data from their home country, using local computer clusters or even individual PCs (Tier-3).
Optical fibre links working at 10 gigabits per second connect CERN to each of the eleven major Tier-1 centres around the world. This dedicated high-bandwidth network is called the LHCOPN (LHC Optical Private Network). Tier-2s are linked to a Tier-1 or to other Tier-2s via the corresponding NREN and/or dedicated GÉANT links.

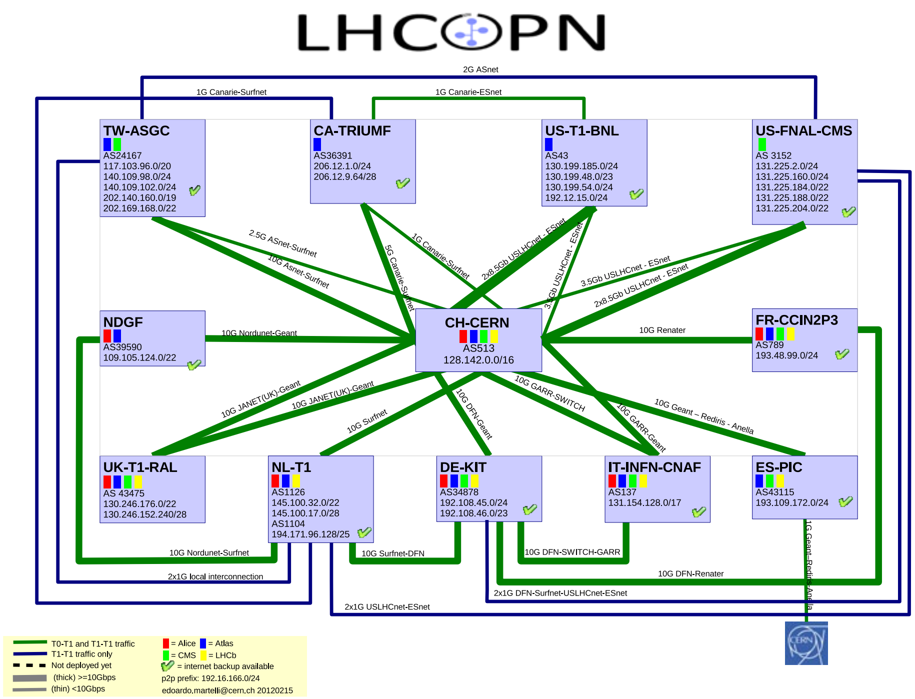


*Author: Edoardo Martelli, CERN*

Thousands of physicists around the world count on the continuous availability of data to do their research and thus on the continuous availability of a network with a good performance.

Transferring the huge amounts of data and large files associated with the LHC can have a noticeable effect for other users of the network if is not adequately managed. Transfers can be managed by a file transfer service (FTS) that only attempts to start or re-try a transfer when there is sufficient bandwidth available on the link. GridFTP uses the approach of using concurrent multiple streams to improve throughput and provide additional

back to contents page

functionality, such as third-party transfers, allowing a user to transfer between two remote sites without transferring it through the initiating site.

Recognising the need for a higher-level interface to storage systems, the Storage Resource Manager (SRM) protocol was developed by some of the participating HEP institutions. This is a web-services protocol, which provides information about the underlying storage system. A major advantage of SRM is that it can optimise transfers by separating the control of transfers from the transmission of the data.

Researchers throughout the Worldwide LHC Grid (WLCG) community have spent much time optimising data transfers over networks: tuning TCP windows, buffer sizes, parallel streams, congestion backoff policies, etc., but these approaches require human intervention and monitoring. Dedicated point-to-point optical links provide some solutions for routes where large data-flows are constantly transferred.

However, the most important issue is that the NRENs and the user communities' work together to ensure that adequate network services are available to support the science.

#### 5.3.2.2 Storage, mirrors, preservation, disaster recovery

The WLCG data model is hierarchical: imagine concentric circles with the Tier-0 at the centre, a small number of Tier-1 centres distributed around, a larger number of T2 centres around those, and T3s around those again. As a rule, it is a strict hierarchy in the sense that a T2 "belongs" to a single T1, and the T1s themselves "belong" to the T0. The Tier-0 centre (CERN for the LHC) is where the data are generated. From here, data are copied to the Tier-1 centres – in WLCG, there are eleven of these centres, distributed across the world, but not all centres support all experiments. The European centres are in Germany, France, Italy, the United Kingdom, the Netherlands, and Spain. There is also one in the Nordic countries (where the T1 itself is "distributed" but with the centre at NDGF in Copenhagen), so seven of the eleven are European.

While the original design was strictly hierarchical, the boundaries are now looser, and a T2 may copy data to or from any T1, not just "its own T1." The T1s themselves may copy data between each other. The high-level architecture resembles the concentric circles in the diagram.

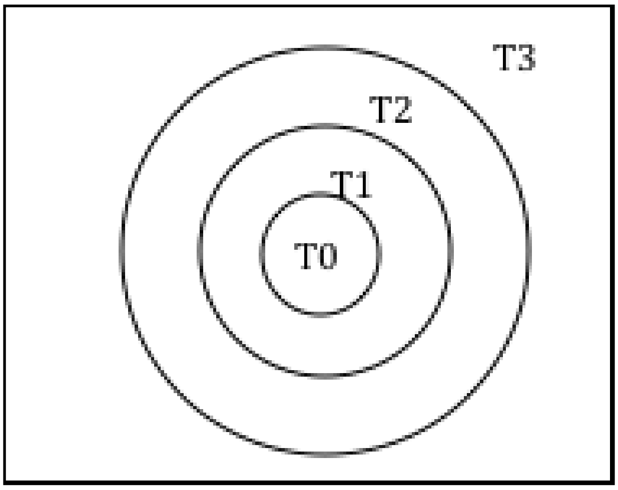


*The WLCG data model*

Each experiment has its own usage model. For example, the ATLAS experiment has a policy that the data are replicated across the T1s, but no T1 needs to hold all the data. In the case of ATLAS, data (Event Summary Data,



back to contents page

ESD) is held at T0, at one primary T1, and at one backup T1. Thus, if data are lost for some reason from a T1, or there are network problems, the data can be fetched from another.

Tier-0 and Tier-1 sites generally provide tape-backed storage, as some of the data they store is considered "precious." The data model considers three types of data "retention policies".

- Custodial: This originally meant that the data cannot be recreated, such as raw data from a beamline, and the site storing the data "shall make reasonable efforts to prevent its loss". As custodial data are available at other T1 sites (and most likely at the T0, as well), the meaning of the attribute quickly changed to "data shall be stored on tape," because tape is considered more reliable for long-term storage than disk.

- Output: This term describes data that is the output of an analysis job, which may have run for a week before creating its output. Thus, this data could be recreated from its original input data, but at some cost. At this time, this attribute is not being used directly by WLCG.

- Replica: The "replica" data are data that are genuinely replicated from elsewhere, and could be recovered at the cost of the data transfer. This is generally interpreted in WLCG as "data are not backed up to tape."

Almost all data are replicated, but in a targeted way. People in designated data-management roles within the experiments – rather than the end users - decide which sites should receive which datasets, according to pre-defined data models of the experiment. As the data are copied, the system keeps track of replicas (and checks their checksums) in replica catalogues. Interestingly, the LHC experiments have all opted to implement and use their own catalogues, along with their own data-tracking tools. In contrast, the non-LHC VOs do not have resources to develop additional tools, and generally use the stock infrastructure - European Middleware Initiative (EMI) (in Europe) or Open Science Grid (OSG) (in the United States).

Once the data are replicated, the job – the program that analyses the data – is then sent to the location where the data are stored. Because the data are replicated in advance, and the locations of the replicas are known, this requires planning ahead. There are also ways of creating local replicas of "hot" data files that are expected to be - or discovered to be - "popular" and need replicating to several disk servers. Some of the storage systems do this local replication automatically.

Experiments also have introduced monitoring systems that collect information about the "popularity" of files, so replicas that are no longer required can be deleted (or scheduled for deletion).

Users of the infrastructure can be forgiven for thinking that the grid is brittle, but in any distributed system that has to cope with a similar workload, errors are inevitable. Moreover, designing and implementing software that copes gracefully and recovers in a meaningful way in a distributed infrastructure is not an easy task. Finally, the error-recovery is not always obvious; one proposal currently being discussed is "federated storage" in which data are recovered over the networks from another replica at another site when needed. The protocol is seen as a key part of this recovery; the most likely implementation is based on a transfer protocol from the Stanford Linear Accelerator Center (SLAC) called xroot, and on its capability for redirection, but xroot currently lacks the security needed for WAN use. Other "federation" protocols are being considered, such as WebDAV.

If an "error" is permanently lost data, or disaster recovery, it makes sense to first update the catalogues and then restart the controlled data migration from other sites. It is not feasible to query the storage systems about individual files; rather, a dump is made in a format known as "syncat" (catalogue synchronisation), which is then compared against the contents of the catalogues, and discrepancies can be investigated and repaired.

back to contents page

The HEP community have spent a lot of time and effort on developing methodologies and systems for the long-term preservation of their data.

The most mature curation (data preservation) activity in HEP – looking beyond bit-level storage, appears to be the group led by Cristinel Diaconu from IN2P3 (the National Institute of Nuclear and Particle Physics in France). The work of the ICFA study group on data preservation and long-term analysis in high-energy physics is reported on the website.[5]

Currently, preservation is not currently considered a problem for WLCG, because the data are very much "live", but it can be a problem for some of the older physics experiments whose funding has closed. Preservation of data implies preservation also of the software that was used to access and analyse the data, or even the knowledge of the people who worked on the data and the software, as they move on to other things.

### 5.3.3 Middleware

Interoperation is important for building a shared infrastructure. The most important grid middleware stacks in WLCG are: EMI (European Middleware Initiative), combining the key middleware providers of ARC, gLite, UNICORE and dCache, Globus Toolkit developed by the Globus Alliance , OMII from the Open Middleware Infrastructure Institute and the Virtual Data Toolkit, albeit complemented with many experiment-specific products (such as higher level data catalogues and data movers) added by the experiments, with relatively little sharing between the experiments.

### 5.3.4 Metadata

Physics metadata (as opposed to the metadata needed for the infrastructure) is held in a grid-enabled database known as 3D, Distributed Database Deployment. Its implementation relies heavily on Oracle databases.

### 5.3.5 AAI

Users access the grid infrastructure through X.509 certificates obtained by a globally distributed network of certification authorities (CAs), operating to a common baseline defined as IGTF (International Grid Trust Federation[6],). The trust model enables users from one part of the world to get a certificate from their (usually national) CA, and access resources (assuming they have permission) in any other part of the world.

Authorisation is normally granted by virtue of membership of a Virtual Organisation (VO). The infrastructure provides a set of VO Membership Services (VOMS), which provide membership and role, attribute in the form of attribute certificates. (These mainly follow RFC3281 but can also use SAML format.) In addition to membership, VOs can grant special roles to members, which the members can choose to optionally exert. Typically, there are very few roles, but an LHC VO would have a "production" role, assigned to the few individuals who have permission to write data to a T1. Individual members will generally do their analysis work at T2s (or T3s.)

Establishing and maintaining trust between the resource providers and the users/VOs is a complex matter. To summarise it in a few words, it is based on the acceptance of basic policies compiled by the Joint Security Policy Group, JSPG[7] (it has changed its name in EGI but the principle remains the same.) In brief, VOs request resources from a site describing their intended work, and users sign an agreement with the VO (using their digital certificate.) Trust is also helped by regular "security challenges," where a designated user will upload a file or run a job, or delete a file, and the site will be challenged to respond to the "incident."

5 http://www.dphep.org/.

6 www.igtf.net

7 www.jspg.org

### 5.3.6 Data policies

#### 5.3.6.1 Public availability of data

In general, data confidentiality is not a problem for HEP: as a rule, a VO understands its own data, and no one else does. Moreover, if data were encrypted in-flight or at-rest, data transfers would slow down a lot, incurring also a heavy CPU load on sender and recipient systems (or at the time of use, for at-rest encryption.) Indeed, world-readable data is being considered primarily to disable access control and speed up the access to the data, i.e., for performance reasons. However, there is a long way to go in actually making the data public in a way that can be used by the general public. The closest the public will get – in the short to medium term – to making sense of the data is LHC@Home.

#### 5.3.6.2 Curation of data

The T0 and T1 sites are expected to archive the data as well as making at available in as working repositories. While in the past, tape archives have been used for both archives and working repositories, WLCG are proposing to change this to use mainly tape for the archival copies of the data and to rely on disk copies – and network transfers – for the working copies of the data. As with other data stores, there is an on-going debate about cost/effectiveness/reliability/"greenness" of tape versus disk for long-term storage: currently tape is still being relied on as a cost effective archiving solution, and while technological developments are tracked, this is not expected to change in the near future.

### 5.3.7 Data origin

#### 5.3.7.1 Data integrity

Data integrity is extremely important, and checksums are used extensively to ensure the integrity. As data volumes are large, these are generally computationally lightweight checksums with a good chance of detecting common errors, rather than cryptographically secure and computationally expensive types. ADLER32 is a popular choice.

### 5.3.8 Likely future developments in the topic area

WLCG does not rest on its laurels, but continuously looks to improve systems and performance, particularly as users or VOs push resources to the limit, and defects are found in the supporting software. When the grid was small, there was a danger of a small group getting together to decide how to "improve" things, but WLCG, to its credit, has learnt to consult more widely. Of course, the problem is to consult so widely that nothing gets done, but, the Grid Deployment Board (GDB), the steering group for WLCG, generally follows up on proposed improvements.

In 2010, a collaboration "Jamboree" was held in Amsterdam, throwing the floor open to brainstorming improvements. Among the improvements suggested was using bittorrent for data distribution (not for the first time, nor for the last), using Apache Cassandra, and other potential ideas. All ideas were challenged to demonstrate their impact within six months; some of them did, others did not.

In late 2011, WLCG came around to visiting the performance and efficiency problem again, looking for optimisations. This time, four Technical Evolution Groups were set up, covering storage, data management, databases, and operations[8]. Their work is currently on going (as of Feb.2012); the storage and data management TEGs have quizzed the experiments and are currently working on re-evaluating the baseline requirements (e.g., file access, directory functions and data transfers) from the experiments, with the aim of identifying the software sites that may still be running three to five years from now. WLCG is once again consulting widely, bringing

8 As of March 2012, there are now two more working groups: workload management and security



back to contents page

representatives of experiments, resource providers, middleware experts, and support staff together. Draft reports from these working groups were used as input to this report.

## 5.4 Radio Astronomy

### 5.4.1 Introduction

The study of astronomy and cosmology requires high-sensitivity, high-resolution instruments in order to image the heavens. New instruments, designed for the advanced study of astronomy and cosmology, such as the origins of galaxies, stars and planets, require ever-higher sensitivity and resolution.

Radio telescopes use aperture synthesis[9] techniques in order to simulate a telescope with a very large diameter. While the resolution of an observation is proportional to the diameter of the array, the sensitivity is a function of the effective area of the array and the square root of the bandwidth.

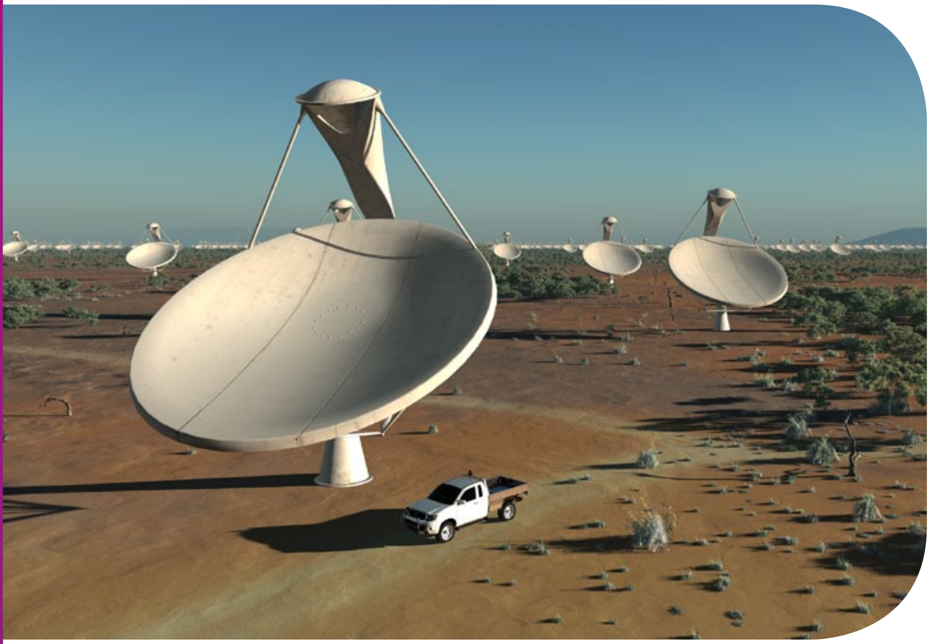

*SKA - Dishes*

The consequence is that modern radio interferometers achieve the characteristics of a very large instrument by utilising an array of many smaller antenna or receiver stations.

Already astronomers have access to around one petabyte of archived data. It is anticipated this will grow to more than 60 petabytes of archived data by 2020 [1]

### 5.4.2 A Description of the data in radio telescopes

#### 5.4.2.1 Data flow

The data in a radio telescope flows from an individual antenna to a central processing facility where it is correlated. These data are the digitised Radio Frequency/Intermediate Frequency data from radio receivers. In general, the digitised RF/IF data flows through private networks associated only with the telescope. A notable exception to this rule is eVLBI [2], which uses NREN links to deliver data from telescopes across Europe to a correlator at JIVE [3] in the Netherlands.

The correlator generates a set of visibility data or a 'measurement set' from the RF/IF input data. [3] The measurement set is then processed into a data cube, image or similar data product. In the case of an image, this process involves the calibration and flagging of data followed by a gridding and Fourier transform process. These

9 Aperture synthesis involves the transfer of signals from separate antennas in an array, to a central location, where the radio signals are correlated and a cosmic source signal is extracted



back to contents page

data are often held in FITS files or in more recent telescopes, in HDF5 format. The processed data has metadata associated with it.
In the past, end users have had access to the correlated visibility data in the form of a measurement set. In modern radio telescopes, the measurement sets are large enough that this becomes untenable. In this case, visibility data may be archived at a central node and data products and images distributed through regional centres.

In extreme cases, it is possible that visibility data are not archived and that automatic processing generates the data products. This approach may be necessary when data archives reach beyond the petabyte scale in short timeframes. It is possible to take this approach in astronomy because the majority of objects under observation change very slowly (often on the scale of 100's or 1,000's of years). Parochially, this is known as "the great database in the sky".

### 5.4.3 Networking

The networking requirements of radio astronomy vary according to the application and the location of the network element within the data flow of the telescope.

Data flowing from dishes to a correlator over public networks, as is the case in eVLBI, require 1-10 Gbps links, and are dependent upon the available bandwidth in public networks. Bandwidths as large as 30 Gbps are required in the VLBI National Radio Astronomy Facility, e-MERLIN, and are served by private networks. The bandwidth is determined by the observing frequency of the array and the bandwidth of the antenna receivers. Data in this mode of transmission do not require high reliability, but do need high-speed throughput and, ideally, deterministic arrival times.

Data flowing from archive facilities will depend on the scope of the operations of the telescope. Large telescopes, such as ALMA and the SKA have regional centres serving data to users within a region. In the SKA, the archive itself is expanding at petabytes per day. Serving regional centres requires 100Gbps input links and output links, determined by the needs of the regional users. The data flow in this mode needs high reliability, and their throughput depends on the backup strategies determined by operational plans, as yet undefined.

The future of data transfer in radio astronomy in the next decade will be driven by the new radio telescopes that are being designed and developed. These new radio telescopes are characterised by high sensitivity and a high field of view. In data terms, this means high data rates within the telescope infrastructure and very large data archives. The large capital costs of these new observatories mean that there are likely to be fewer telescopes and only a few central archives. This era of large telescopes and large archives means that radio astronomers are going to place pressure on research networks when they access large data sets remotely. The precise nature of the distribution of the data archives has yet to be defined. The economics and limitations of network access will play an important part in the development of the strategies for the distribution of this data.

### 5.4.4 Metadata and middleware

RF/IF data has synchronisation and timing metadata associated with it (usually as a header) as well as essential ID information. This metadata allows the central processor to associate a field system log to the data. The field system log has important instrument data associated with it that are essential for the calibration and flagging of datasets in the generation of data products.

Visibility data, in the form of the measurement set has calibration and flag tables embedded in it, for use in the processing stages. Data models for radio astronomy have been standardised in recent years, with the advent of the International Virtual Observatory Alliance (IVOA) initiative [5]. To enable the use of searchable catalogues,

back to contents page

IVOA has developed standard model formats for describing astronomical datasets. The middleware currently in use uses these standard models.

Metadata associated with data products in radio astronomy include many types of data associated with the observation including:

- proposals and papers;
- scheduling information;
- environmental data;
- array instrument data;
- calibration data;
- time and synchronisation data;
- coordinate data;
- data product information (such as, total power measurements);
- processing history.

The subject of ontology as applied to astronomy is starting to be explored, in order to expand on the capability of searchable catalogues. [6]

### 5.4.5 Data policies and data origin

Traditionally, radio astronomy has operated using the principle of 'Open Skies'. This principle allows access to radio telescopes, irrespective of the country of origin, but based upon peer review of the scientific value of a proposed observation.

This model has proved difficult to maintain in the era of very large and expensive radio telescopes with large output data sets. Funding bodies want privileged access for their scientific community, to the instrument that they are funding. In these new telescopes, the data are accessed via regional data centres in the countries or regions of the funding. How this change in operating mode will manifest itself in practise remains to be tested. There is no moral imperative or sensitivity associated with radio astronomy data, as demonstrated by many years of successful operation of the 'Open Skies' policy. Access to the data will be based on policy decisions at the management layer of these large telescopes. It will be determined, in part, by the attitude to data access and the investment return of the funding bodies and their associated governments.

### 5.4.6 Modern radio telescopes – usage model

#### 5.4.6.1 eVLBI

Very Long Baseline Interferometry (VLBI) is carried out by a collection of large telescopes located across the globe, which collaborate to produce the world's highest resolution instrument for radio astronomy. In the past, VLBI observations were restricted because the data was recorded onto tape and then shipped to a central processing facility for analysis. Consequently, radio astronomers were unable to judge the success of their endeavours until many weeks after the observations were made. Using NREN connections and a standardised set of digital recording equipment, real time observations are now standard between telescopes across Europe and a correlator at JIVE.

back to contents page

The processed data are available in FITs files. The figure below shows the increase in the size of the archive at JIVE over time.

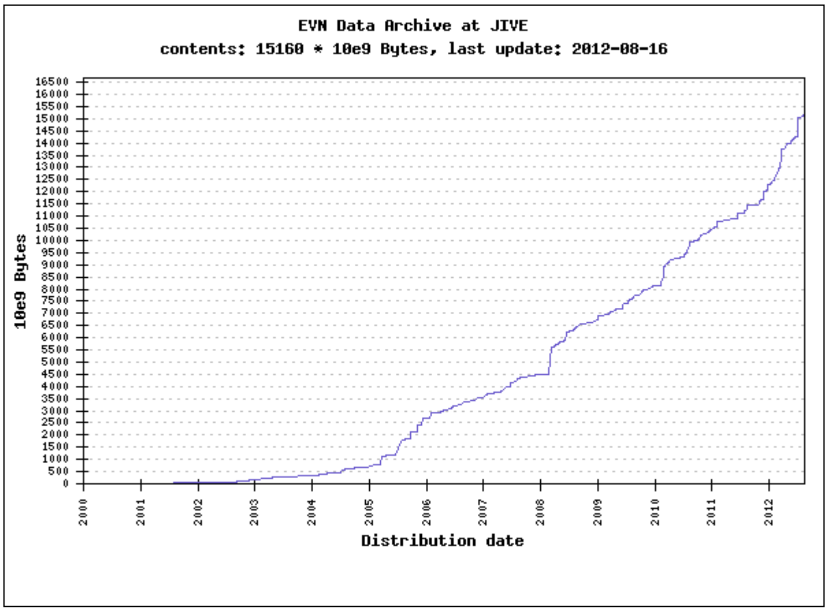


### 5.4.6.2 ALMA

The Atacama Millimetre Array is a new telescope being commissioned in the Atacama Desert in Chile. It specialises in observations at millimetre wavelengths. ALMA will operate five data archives in total: one at the Operations Support Facility (OSF), one in Santiago and one at each of the three ALMA Regional Centres (ARCs).

The OSF archive receives visibility data from the correlator, together with additional monitor and weather data. The average data rates are 6.6 MB/s, with peak data rates of 66 MB/s. The data are stored for periods of up to a year. The architecture is based on the New Generation Archive System (NGAS) with Oracle technology for the metadata.

The data in the OSF Archive are copied via the network to the Main Archive, located in Santiago. The Pipeline is operated in Santiago and its data products are ingested into the Main Archive. The data in the Santiago archive are replicated to the archives at the three ALMA Regional Centres:

- ESO – Europe;
- NRAO – USA;
- NOJ – Asia/Pacific.

The data are held in the ALMA Science Data Model (ASDM) format, which is common to the EVLA and ALMA telescopes. [6]

### 5.4.6.3 LOFAR

LOFAR is a new telescope based mainly in the Netherlands with outlying stations across northern Europe. The main data archive is based, along with the software correlator, at ASTRON.

LOFAR produces very large data streams. One month of observing results in petabytes of data. Thus, systematic long-term storage for such data volumes becomes extremely expensive.

back to contents page

Therefore, LOFAR has adopted a processing model in which final data products for routine observations are formed in a highly-automated mode, with user interaction at well-defined moments. Only final data products are exported. The file format adopted for distribution of the data products is HDF5. [8]

#### 5.4.6.4 SKA

The Square Kilometre Array (SKA) will be an ultra-sensitive radio telescope, built to further the understanding of the most important phenomena in the universe, including some pertaining to the birth and eventual death of the universe itself. Upon construction, the SKA will be the largest radio telescope in the world with an equivalent aperture of one million square metres made up of thousands of antennas spread across thousands of kilometres. It will be located in Southern Africa or Western Australia. The SKA is anticipated to begin science observations with a 10% array in 2017 and observations with the full array by 2022.

The data flow in the SKA is of unprecedented scale. The transport of the RF/IF data within the telescope will require a network capacity that exceeds the global Internet traffic predicted for years to come. Fortunately, this network will be, for the most part, a private network and will be a part of the infrastructure of the instrument.

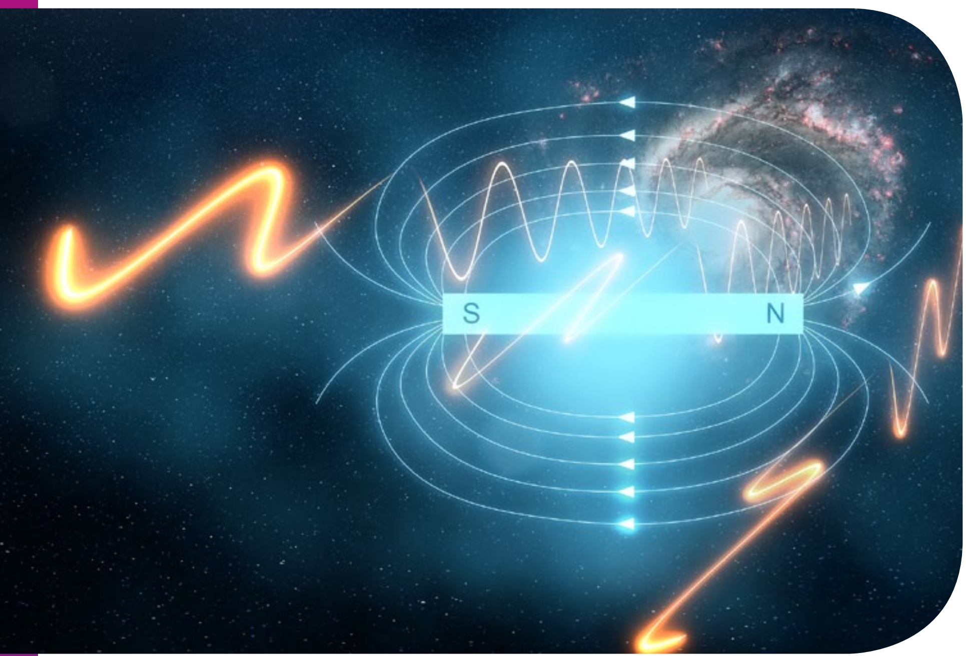


SKAScience - cosmicmagnetism

The measurement-set data will be too large for distribution or even for archiving, and it is anticipated that automatic processing of this data will deliver data products and images to the astronomers. Even with efficient processing of the data and compression of the image files, the final SKA is predicted to generate many petabytes of data per day into archive. The SKA is truly in the next generation of scientific instrument, taking information management and processing into the exascale domain. A 100 Gbps link will stream image data to regional centres around the world.

On the way to the final SKA, precursor projects are providing milestones, testing the technologies and processes in advance of the full system. ASKAP, the Australian SKA Precursor, is 1% of the final SKA and will dump 1 Gbyte per second into archive [8]. In a 10% step, SKA1 will have an exabyte archive storage facility and a 10 Gbps link to stream images to regional centres around the world [10]. Astronomers would like more, but the upper limits have been capped by the anticipated costs of such a system.

Such large data management challenges are encouraging institutions to consider the need for collaboration tools and visualisation techniques. In addition to the IVOA, projects such as CyberSKA, funded by NRC in Canada [11], and the ICRAR science archive facility in Australia are looking at infrastructure and collaborative tools to support the projects with the high volume of data associated with radio astronomy, with the aim of supporting the needs of the final SKA.

back to contents page

## 5.5 Distributed Music Performance

Music institutions are beginning to make use of NRENs to deliver traditional practices such as rehearsals, auditions, master classes, composition, and performance workshops via Video Conferencing technology. This activity fosters links between music institutions in Europe and worldwide and creates opportunities for cross-cultural music composition and collaboration.

Faced with funding cuts, innovating with technology has become increasingly important, offering new opportunities to students both nationally and internationally.

### 5.5.1 Networking

#### 5.5.1.1 Bandwidth requirements

H.323/SIP standard Video Conferencing requires bandwidth of no less than 384 Kbps at the minimal configuration. The highest available video resolution at present is HD 1920x1080 pixels with progressive scan at 30 fps. Polycom equipment transmits data over 1024kbps. DVTS requires bandwidth of around 30 Mbps and has latency between 200 – 400 ms. Conference XP requires bandwidth from 2 Mbps up to 30 Mbps (or higher for multiple stream connections) and has latency between 200 – 400 ms. Jack Trip requires bandwidth of 1 – 2 Mbps per audio channel and has latency of less than 10 ms. LOLA, at its minimal configuration (640x480 pixels, black and white, 30 fps, 2x44100 24 bit audio) requires 96 Mbps of end-to-end bandwidth. A colour camera at the same configuration requires 500 Mbps of end-to-end bandwidth. An HD colour camera and multiple audio channels will further increase bandwidth requirements up to 7 Gbps. The general rule of thumb for latency when using LOLA on high-speed NRENs and backbones such as GÉANT and Internet2 is 1 ms per 100 km.

Research and education networks with a reasonable amount of capacity should be able to support LOLA performances at its minimal configuration but, where available, end-to-end layer 2 or layer 1 optical circuits (lambdas) should be used and are required for higher configurations.

LOLA requires the underlying network to provide a very reliable transmission system of 1Gbps with zero packet loss and very low jitter. This is because the LOLA software does not include any audio/video buffering or include any mechanism for recovering from data loss.

Bandwidth issues exist at several layers:

- circuits and lightpaths may be available at large research universities, but they are not within reach of smaller institutions;
- security features including firewalls and NAT introduce latency and should be avoided on end-to-end paths;
- the 'last mile' problem is inherent in many music institutions and within campus networks, where classrooms, studios, and performance spaces are not adequately connected to the LAN. Also, the LAN itself may not be adequately connected to the NREN;
- previous tests have shown that on the GÉANT backbone, there can be conflicts between different applications for bandwidth allocation, for example, with data transfer from the LHC.

back to contents page

## Content

**Now**
Music institutions worldwide that have adopted Video Conferencing technology to enhance music education include the New World Symphony, the Australian National University's School of Music, the Manhattan School of Music and the Royal Danish Academy of Music. Their students are given access to some of the world's leading composers and musicians via videoconference. In some cases, master classes and outreach programmes to schools provide a source of income for institutions.

The LOLA software takes innovating with technology one step further, by allowing musicians to play together over the national academic networks in real time. Music institutions in the United States, including the New World Symphony and the Northern Illinois School of Music are currently using LOLA for education activities and performances.

**The future**
An increase in the number of music institutions using LOLA will lead to multi-site performances that increase demands on the network.

At present, LOLA uses a single camera to capture performances. As presentation becomes more refined, multi-camera capture will be introduced. HD, Super HD and 3D cameras will require an exponential increase in network bandwidth.

As commodity network bandwidth increases, Video Conferencing technology will be more widely used by commercial music organisations. Orchestras will have the option of rehearsing with soloists and conductors at a distance in advance of performances, thus saving time and expense.

#### 5.5.1.2 Cost aspects

The LOLA system makes heavy demands of up to 1Gbps on the underlying network. The consequence is that whilst usage is slight, it can be accommodated within the existing networks capacity. The UK JANET lightpath policy, for instance states that such a service will be provided without direct cost to the associated project. LOLA has proved an extremely capable tool within the performing arts sector. However if its adoption becomes more widespread there will be a real impact for network capacity requirements, resulting in increased costs which will have to be met.

### 5.5.2 Middleware

**What do you have now?**
DVTS and Conference XP software, H.323 standard Video Conferencing and, to a lesser extent, Skype are currently used for music education at a distance.

**H.323/SIP**
Video Conferencing codecs using this standard were designed for voice, not music. Audio and video is transmitted in a compressed format. The Manhattan School of Music has worked with Polycom to develop the 'Music Mode' version, which disables the acoustic echo cancellation and automatic noise suppression features found in traditional models and, therefore, enhances the quality of music.

back to contents page

**DVTS (Digital Video Transport System)**

This is available on Mac - and Windows-based operating systems. It transmits uncompressed audio at 48kHz/16 bit and standard definition video using DV25 compression. An HD version is also available. The latency is approximately 200 – 400ms.

**Conference XP**

This is Windows-based software, which transmits uncompressed audio and video and can handle audio and video transmission independently. It has multicast capability and the latency is approximately 350 ms.

**Skype**

Skype is peer-to-peer and client-server system software. Network administrators in many educational establishments, particularly schools, block the use of Skype because of security concerns. The latency is unpredictable.

**JackTrip**

This is Linux- and Mac-based software, which transmits uncompressed audio at up to 192Khz. It has multichannel capability and latency of less than 10 ms.

**Echo Damp**

This is a Mac- and Windows-based multi-channel audio mixer and echo controller, designed for use in high-bandwidth, music videoconferences. Unlike other echo cancellers, it maintains the full frequency spectrum of the audio hardware used to capture the sound.

**What do you want?**

Musicians have not been able to play together with any of the technologies above, due to a number of reasons, the most significant of which is latency. Acceptable latency for the human ear is <60s (RTT).

Since 2005, musicians and developers from Conservatorio Tartini in Trieste and network specialists from GARR, the Italian National Research and Education Network, have been working together to overcome this problem with the development of the LOLA (Low Latency) project.

The LOLA audiovisual streaming system uses audio and video transmission software to reduce latency and jitter, sending pictures and sound in real time over high-speed NRENS, giving the impression that performers are in the same venue.

The first public demonstration took place during TERENA's Network Performing Arts Production Workshop at IRCAM, Paris in November 2010. Two pianists, one in IRCAM, the other in Trieste, successfully played together in real time. The two sites were connected using an end-to-end link that used the Conservatorio Tartini LAN, the Trieste Lightnet Metropolitan Optical Network, the GARR backbone, GÉANT, the French RENATER research network backbone, the IRCAM 1Gbps last mile and the IRCAM LAN.

The second public demonstration took place during TERENA's Network Performing Arts Production Workshop at Gran Theatre del Liceu in Barcelona in June 2011. Two violinists, one in Barcelona, the other in Trieste, performed together in real time. Both sites were connected using an end-to-end link that used the Conservatorio Tartini LAN, the Trieste Lightnet Metropolitan Optical Network, the GARR backbone, GÉANT, the Spanish RedIRIS research network backbone, the Anella Científica network and the Gran Teatre del Liceu link to Anella Científica.

During the Internet2 and Fall Member Meeting 2011, a violinist from the University of North Carolina in Raleigh and a cellist from the Northern Illinois University School of Music in Chicago played together in real time using

back to contents page

LOLA over Internet2 using an end-to-end link that used the North Carolina Research and Education Network and the Illinois Century Network.

The LOLA project team hosted a 'hands-on' workshop at Conservatorio Tartini, Trieste, in April 2012. This gave music practitioners and technical staff from conservatoires and universities across Europe the opportunity to experience the system first hand. Performers from a wide range of backgrounds, including classical music students, music teachers and students and professional musicians demonstrated playing at a distance in real time at an average of <50 ms (RTT).

LOLA software is freely available for non-commercial use.

### 5.5.3 A view of the future

Through the use of digital technologies, music institutions will be able to overcome the constraints imposed by physical location in the future. Using LOLA, they will have the ability to create new opportunities for students through cross-cultural collaboration. In turn, new audiences will be reached and new business models will be developed, adding new sources of economic and cultural value.

back to contents page

# 6 COMMON REQUIREMENTS FOR THE DISCIPLINES STUDIED

The ASPIRE study of the data and data-handling requirements between now and 2020 has indicated that these diverse disciplines have a number of common requirements.

The economics of networking will play a part in the development of strategies adopted in of all disciplines - for access to their large data sets, for the fast transmission of those datasets, and for real-time transmission of specific data. The need to access large amounts of data is common across a number of high-performance applications and middleware, discussed in this study. The ability to share services, infrastructure, and therefore cost, would be a highly desirable outcome for all these services.

Storage, preservation, and curation of data are common to all. For most of these disciplines, it is not possible or will not be possible to store and preserve all of the data. Clear rules and procedures concerning what to store and preserve and how to do curation are required. All of the disciplines studied welcome new techniques and automatic procedures to assist in data curation and the choice of data that has to be preserved.

Datasets in the various disciplines are very different, but often, they are also very often different within a discipline, depending upon where the data have been generated. Standardisation of datasets in some generic way will facilitate the development of middleware and/or applications.

Metadata is a vital element of storing, handling, and using the vast qualities of data that are being amassed. There are many standards for metadata and they are not always compatible. There is a common requirement for the adoption of one standard that takes the inter-disciplinary use of data into account.

The users of e-Infrastructures need an easy-to-use, globally recognised authentication and authorisation infrastructure, usable across disciplines.

There is a need for the authentication of the source of data, with more data and data being produced by less-controlled or uncontrolled sources. Everyone agrees that authentication procedures that are now in use will not scale and will not be adequate for authentication of source of data in the future. A common solution is welcomed.

There is support for the concept of an "open data" policy, declaring all data available for everyone. All disciplines understand this tendency, and all agree that the policy opens doors to new research and that it makes inter-disciplinary research easier. However, data can be ethical or socially sensitive and there can be legal issues of the ownership of data and intellectual property rights. There can also be financial issues related to the creation and preservation of data. There is consensus on the need for clear, global rules for the availability of data.

back to contents page

# 7 RECOMMENDATIONS

**Issue:** The amount of data collected and stored is increasing exponentially and along with it, the need for the bandwidth to transport the data in order to make it available to researchers and users. Because the capacity, throughput, jitter, and delay requirements of the network can be stringent, commercial network providers cannot make these connections available quickly and at an affordable price.

**Recommendation 1**

National Research and Education Networks (NRENs); GÉANT and others involved in providing network connectivity need to collaborate with the user communities to ensure that the networking requirements associated with the deluge of data are well understood. Adequate network services need to be put into place in a timely manner and in an economically viable fashion. Aspects of speed of provision, throughput, privacy, persistence of connection and other important parameters need to be addressed.

**Issue:** Many disciplines create their own standards for data storage, and consequently, create their own middleware and applications. Due to the incompatibility of the datasets, scientific disciplines cannot benefit from the common generic developments that would be possible if cross-discipline standards were adopted.

**Issue:** The increased move towards "open-data" and the improvement of e-infrastructures bring with them new possibilities for cross-disciplinary research. This approach will be hampered by the parochial nature of datasets, middleware, and applications.

**Issue:** Many originators of data add metadata at the moment of the creation. However, since there are so many public and private standards available, there is a considerable chance of incompatibility across disciplines and even within disciplines.

**Recommendation 2**

Define standardised datasets in order to profit from the economies of scale that would follow from the availability of generic, cross-discipline middleware. Define standardised datasets, metadata, middleware and applications for easier accessibility of data. Adopt a common metadata standard that takes the multi-disciplinary use of data into account.

**Issue:** Even though there is a strong desire to hold data in the public domain, it is necessary to safeguard some resources for privacy, commercial, or legal reasons. A pervasive Authentication and Authorisation system is required to provide a scalable system of protection that will establish authorisations at the appropriate administrative level.

**Recommendation 3**

Adopt a globally recognised Authentication and Authorisation Infrastructure (AAI) based on recognised standards for the exchange of assertions and security tokens that can be utilised by all user communities, e-Infrastructure providers, and ICT service providers.

back to contents page

Examples include eduroam®, a system dedicated to the single purpose of authenticating network access, and eduGAIN, which is increasingly being used to connect interfederations.

**Issue:** With the massive increase in the volume of data and increasingly open accessibility, it will be necessary to be able to understand the provenance of the data along with any manipulation or processing that may have taken place. It would be helpful to researchers to be able to ascertain if data have come from a particular instrument, if they have been processed, or if they have originated from other sources.

### Recommendation 4

Create common mechanisms and procedures for all disciplines to certify and authenticate data. This could include automated signing of data before they leave the instruments where they are created, whether the instrument is a large device, such as the LHC or SKA, or a smaller device, such as a camera or scanner.

**Issue:** All disciplines have a need for the preservation and curation of data. Currently, much of the decision-making process in selecting what should be preserved and/or curated is undertaken by human appraisal. Human intervention in this way is not scalable in the longer term.

**Issue:** There is an increasing tendency to make results of publicly funded work "open data" that is accessible by all. Whilst this is a laudable development, there maybe issues of ethical or social sensitivity, legality, industrial property rights, or other reasons for keeping some data protected from public access.

**Issue:** If datasets are to be maintained for long periods or indefinitely, continued funding will be required,

### Recommendation 5

Facilitate collaboration between disciplines to create common policies, procedures, and tools to assist in the curation of data and the selection of data for preservation.

back to contents page

# 8 REFERENCES


[1] How will astronomy archives survive the data Tsunami?, Berriman G.B, Groom S.L, Cornell University Library, Instrumentation and Methods for Astrophysics, http://arxiv.org/abs/1111.0075.

[2] eVLBI with the EVN. http://www.evlbi.org/evlbi/evlbi.html

[3] JIVE http://www.jive.nl

[4] MS2 Description for LOFAR CS1. http://www.lofar.org/wiki/lib/exe/fetch.php?media=public:documents:ms2_description_for_lofar_v05.pdf

[5] About the IVOA. http://www.ivoa.net/pub/info/

[6] Ontology of Astronomical Objects, Cambrésy L.et al, IVOA Technical Note 17, http://www.ivoa.net/Documents/Notes/AstrObjectOntology/

[7] The Alma Science Data Model, Viallefond F, Astronomical Data Analysis Software and Systems XV ASP Conference Series, Vol. 351. http://adsabs.harvard.edu/full/2006ASPC.351.627V

[8] LOFAR and HDF5 Toward a New Radio Data Standard, Anderson K, Cornell University Library, Instrumentation and Methods for Astrophysics, http://arxiv.org/abs/1012.2266

[9] SKA Exascale Software Challenges, Cornwell T, Humphreys B, SKA Memo 128, http://www.skatelescope.org/publications/

[10] SKA Phase 1, Preliminary System Description, Dewdney P. et al, SKA Memo 130, http://www.skatelescope.org/publications/

[11] The Canadian CyberSKA project, Willis T, Workshop - The growing demands on information processing and connectivity in radio astronomy from VLBI to the SKA, IT, PT, 24th-25th May 2011, http://www.av.it.pt/workshops/program/tony_willis_vlbi2ska.pdf

[12] DPHEP interim report http://arxiv.org/ftp/arxiv/papers/0912/0912.0255.pdf

[13] LHC@Home: http://lhcathome.web.cern.ch/LHCathome/

[14] WLCG TEG interim presentations: http://indico.cern.ch/conferenceDisplay.py?confId=158775


back to contents page

[15] "Born Digital" An Essay by Ricky Erway, OCLC Research.
http://www.oclc.org/research/activities/hiddencollections/bornditgital.pdf November 2010 ©2010 OCLC Online Computer Library Center, Inc.

[16] Seth van Hooland, Ruben Verborgh, Max De Wilde, Johannes Hercher, Erik Mannens, and Rik Van de Walle (2011). Free your metadata: Integrating cultural heritage collections through Google Refine reconciliation. Pre-submission paper available on http://freeyourmetadata.org/publications/freeyourmetadata.pdf

[17] Riding the wave – How Europe can gain from the rising tide of scientific data"
http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/hlg-sdi-report.pdf

[18] e-IRG White Paper 2011: http://www.e-irg.eu/publications/white-papers.html

[19] EUDAT website: http://www.eudat.eu/

[20] EDUGAIN website: http://www.geant.net/service/edugain/pages/home.aspx

[21] DC-NET website: http://www.dc-net.org/

[22] DARIAH website: http://www.dariah.eu/



back to contents page

# 9 GLOSSARY

| | |
|---|---|
| **3G** | 3rd Generation (mobile telecommunications technology) |
| **3GPP** | 3rd Generation Partnership Project |
| **AAI** | Authentication and Authorisation Infrastructure |
| **AKA** | Authentication and Key Agreement |
| **ALMA** | Atacama Millimetre Array |
| **API** | Application Programming Interface |
| **APN** | Access Point Network |
| **ARC** | ALMA Regional Centre |
| **ASDM** | ALMA Science Data Model |
| **ASKAP** | Australian SKA Precursor |
| **ASPIRE** | A Study on the Prospects of the Internet for Research and Education |
| **ATLAS** | A particle physics experiment at the Large Hadron Collider at CERN |
| **AUP** | Acceptable Use Policy |
| **AWS** | Amazon Web Service |
| **BYOD** | Bring Your Own Device |
| **CA** | Certification Authority |
| **CAD** | Computer Aided Design |
| **CAI** | Community Anchor Institutions |
| **CAPEX** | Capital Expenditure |
| **CEF** | Connecting Europe Facility |
| **CEF/DSI** | Connecting Europe Facility/Digital Service Infrastructure |
| **CERN** | European Organisation for Nuclear Research |
| **CERT** | Computer Emergency Response Teams |
| **CIDOC-CRM** | International Committee for Documentation - Conceptual Reference Model |
| **CP** | Connection Policy |
| **CPU** | Central Processing Unit |
| **DANTE** | Delivery of Advanced Network Technology to Europe |
| **DARIAH** | Digital Research Architecture for the Arts and Humanities |
| **DC** | Dublin Core |
| **DCH** | Digital Cultural Heritage |
| **DCH-RP** | Digital Cultural Heritage Roadmap for Preservation |
| **DC-NET** | Digital Cultural heritage NETwork |

back to contents page

| | |
|---|---|
| **DEAS** | Delegate eduroam® Authentication System |
| **DL** | Distance Learning |
| **DNA** | Deoxyribonucleic acid |
| **DRDB** | Distributed Replicated Block Device (software) |
| **DSI** | Digital Service Infrastructure |
| **DVTS** | Digital Video Transport System |
| **EAP** | Extensible Authentication Protocol |
| **EC2** | Elastic Compute Cloud (Amazon) |
| **ECDD&S** | ELIXIR Core Data Collections and Services |
| **eduGAIN** | Education GÉANT Authorisation Infrastructure |
| **eduroam** | Education Roaming |
| **EEA** | European Economic Area |
| **EGI** | European Grid Infrastructure |
| **EIRO** | European Industrial Relations Observatory |
| **ELIXIR** | A sustainable infrastructure for biological information in Europe |
| **ELSI** | Ethical, Legal and Social Implications |
| **EMBL-EBI** | European Molecular Biology Laboratory - European Bioinformatics Institute |
| **e-MERLIN** | VLBI National Radio Astronomy Facility |
| **EMI** | European Middleware Initiative |
| **ESD** | Event Summary Data |
| **ESFRI BMS RI** | European Strategy Forum - Biological and Medical Sciences Research Infrastructure |
| **EU** | European Union |
| **EUDAT** | European Data Infrastructure |
| **FITS** | Flexible Image Transport System |
| **FTP** | File Transfer Protocol |
| **FTS** | File Transfer Service |
| **GA** | General Assembly |
| **GB** | Gigabyte |
| **Gbps** | Gigabits per second |
| **GÉANT** | Gigabit European Academic Network Technology |
| **GN3** | Multi-Gigabit European Academic Network |
| **GPRS** | General Packet Radio Service |
| **GPS** | Global Positioning System |
| **GUI** | Graphical User Interface |
| **HDF5** | Hierarchical Data Format |
| **HEP** | High Energy Physics |
| **HG** | Human Genome Project |
| **HPC** | High Performance Computing |
| **HPC/Grid** | High Performance Computing and Grid |
| **HTTPS** | HyperText Transfer Protocol Secure |
| **IaaS** | Infrastructure as a Service |

back to contents page

| | |
|---|---|
| **ICFA** | Study Group on Data Preservation and Long Term Analysis in High Energy Physics |
| **ICRAR** | a science archive facility in Australia |
| **ICT** | Information and Communication Technologies |
| **IEEE 802.1X** | e Institute of Electrical and Electronics Engineers – standard for port-based Network Access Control |
| **IETF** | Internet Engineering Task Force |
| **IGTF** | International Grid Trust Federation |
| **IN2P3** | the National institute of nuclear and particle physics in France |
| **IOS** | iPhone Operating System |
| **IP** | Internet Protocol |
| **IP** | Intellectual Property |
| **IPR** | Intellectual Property Right |
| **IRCAM** | Institut de Recherche et Coordination Acoustique/Musique |
| **IRG** | e-Infrastructure Reflection Group |
| **IRU** | Indefeasible Right of Use |
| **ISO** | International Organization for Standardization |
| **ISP** | Internet Service Provider |
| **IVOA** | International Virtual Observatory Alliance |
| **JIVE** | Joint Institute for VLBI in Europe |
| **JSPG** | Joint Security Policy Group |
| **K-12 schools** | primary and secondary schools |
| **km** | kilometre |
| **KVM** | Kernel-based Virtual Machine |
| **LAN** | Local Area Network |
| **LHC** | Large Hadron Collider |
| **LHCOPN** | LHC Optical Private Network |
| **LIPA** | Local IP Access |
| **LMS** | Learning Management Systems |
| **LOFAR** | Low Frequency Array |
| **LOLA** | LOw LAtency audio visual streaming system |
| **LTE** | Long Term Evolution - a standard for wireless communication of high-speed data |
| **MAN** | Metropolitan Area Network |
| **mID** | Unique Identification of person per device |
| **MiFi** | Mobile Broadband Wi-Fi |
| **MMS** | Multimedia Messaging Service |
| **ms** | millisecond |
| **NDGF** | Nordic DataGrid Facility |
| **NFC** | Near Field Communication |
| **NGAS** | New Generation Archive System |
| **NGI** | National Grid Initiatives |
| **NIST** | (US) National Institute of Standards and Technology |
| **NOC** | Network Operations Centre |

back to contents page

| | |
|---|---|
| **NRC** | National Research Council |
| **NREN** | National Research and Education Network (can also refer to the operator of such a network) |
| **NREN-PC** | National Research and Education Network Programme Committee |
| **NSF** | National Science Foundation |
| **OAI-MPH** | Open Archives Initiative Protocol for Metadata Harvesting |
| **OECD** | Organisation for Economic Co-operation and Development |
| **OMII** | Open Middleware Infrastructure Institute |
| **OPEX** | Operating Expenditure |
| **OSF** | Operations Support Facility |
| **OSG** | Open Science Grid |
| **OTP** | One Time Passwords |
| **OWL** | Ontology Web Language |
| **PaaS** | Platform as a Service |
| **PII** | Personally Identifiable Information |
| **PKI** | Public Key Infrastructure |
| **PMH** | Protocol for Metadata Harvesting |
| **PoP** | Point of Presence |
| **R&E** | Research and Education |
| **RADIUS** | Remote Authentication Dial In User Service |
| **RAM** | Random Access Memory |
| **RDF** | Resource Description Framework |
| **REST** | Representational State Transfer |
| **RF/IF** | Radio Frequency/Intermediate Frequency |
| **RNA** | Ribonucleic acid |
| **RTT** | Round-Trip Time |
| **S3** | Simple Storage Services (Amazon) |
| **SaaS** | Software-as-a-Service |
| **SAML** | Security Assertion Markup Language |
| **SIM** | Subscriber Identification Module |
| **SIP** | Session Initiation Protocol |
| **SIPTO** | Selective IP Traffic Offload |
| **SKA** | Square Kilometre Array |
| **SLA** | Service Level Agreement |
| **SLAC** | Stanford Linear Accelerator Center |
| **SMIL** | Synchronized Multimedia Integration Language |
| **SRM** | Storage Resource Manager |
| **SSID** | Service Set Identifier |
| **SVG** | Scalable Vector Graphics |
| **SWOT** | Strengths, Weaknesses, Opportunities, Threats |
| **TERENA** | Trans European Research and Education Networking Association |
| **TLS** | Transport Layer Security |

back to contents page

| | |
|---|---|
| **U.S. UCAN** | United States Unified Community Anchor Network |
| **UMF** | University Modernisation Fund (Greece) |
| **UMTS** | Universal Mobile Telecommunications System |
| **VLAN** | Virtual Local Area Network |
| **VLBI** | Very Long Baseline Interferometry |
| **VLE** | Virtual Learning Environment |
| **VM** | Virtual Machine |
| **VO** | Virtual Observatory |
| **VoIP** | Voice over Internet Protocol |
| **VOMS** | VO Membership Services |
| **WAN** | Wide Area Network |
| **WAP** | Wireless Application Protocol |
| **WebDAV** | Web Distributed Authoring and Versioning |
| **Wi-Fi** | Wireless exchange of data |
| **WiMAX** | Worldwide Interoperability for Microwave Access |
| **WLAN** | Wireless Local Area Network |
| **WLCG** | Worldwide LHC Computing Grid |
| **XML** | Extensible Markup Language |



back to contents page

# 10. CONTRIBUTORS

**Rosette Vandenbroucke,** Vrije Universiteit Brussel*, Belgium*

**ASPIRE DATA Study Leader**

ROSETTE VANDENBROUCKE is the ASPIRE Data Panel leader. She is a computer scientist at the Vrije Universiteit Brussel and represents the University in the Flemish Supercomputer Centre and in the Belgian Computing Grid for Research. During the past several years, she has been active in the e-Infrastructures environment at the international level (EGEE, EGI, e-IRG and several European projects related to e-Infrastructures).

**Gill Davies,** Royal Conservatoire*, Scotland*

GILL DAVIES is a multimedia producer and consultant and runs her own company, Gill Davies Production Ltd. She has twenty years' experience with the BBC, producing content for radio, web and TV. As Senior Interactive Development Producer with the BBC Scottish Symphony Orchestra, she was responsible for the orchestra's multimedia strategy. Gill is currently studying part-time for a Masters Research degree in Collaborative Music Performance over High Speed Research Networks at Edinburgh Napier University.

**David Foster,** LHC/CERN*, Switzerland*

DAVID FOSTER has both a BSc and PHD from Durham University, Department of Applied Physics and Electronics and an MBA from Durham University Business School. He has spent more than thirty years at CERN and has held many technical and managerial roles, He currently chairs a number of international committees on networking and supercomputing as well as representing CERN on a number of projects with the European Commission. David is currently Deputy IT Department Head at CERN and has particular responsibility for international network strategy as well as infrastructure strategy in the areas of information management. He is a Fellow of the Institute of Physics and a member of the Chartered Management Institute and the Association of MBAs.

back to contents page

**Antonella Fresa,** DC-NET*, Italy*

ANTONELLA FRESA is an ICT expert and was a researcher at Olivetti Pisa, Ivrea, and at the Olivetti Advanced Technology Centre in Cupertino (California). She was product manager at Tower Tech. and project officer at the European Commission between 1999 and 2002. Since the early 90s, she has been working on European cooperation projects and policy development frameworks. She is technical coordinator of several EC projects: FP7-ICT, FP7-e-Infrastructures, CIP Pilot B, CIP BPN. Her personal interest and current research focus is related to e-Infrastructures for digital cultural heritage.

**Jens Jensen,** STFC*, United Kingdom*

JENS JENSEN received his PhD in maths from the University of Aarhus, Denmark. He currently works for the Science and Technology Facilities Council, one of seven UK research councils, where he is head of the data services group in e-Science. He leads the grid storage and data management group in GridPP, the UK grid for particle physics. He is Area Director for security in the Open Grid Forum (OGF), and a member of the OGF standards council. He leads the security work package in the EU-funded Contrail project, a project building federated cloud services for, among others, life sciences and high performance scientific data processing. He is the CA manager for the UK e-Science certification authority, and occasionally does other things for the national grid service. Most of his work is in data security, data management, and identity management infrastructures. In the past he has also done work in maths, quantum information theory, quantum cryptography, and various aspects of e-Science.

**Andrew Lyall,** EBI*, United Kingdom*

ANDREW LYALL is the ELIXIR Project Manager. The purpose of ELIXER is to build a sustainable infrastructure for biological information in Europe. Before this, Andrew worked in the Pharmaceutical and Biotechnology Industries. He has a PhD in Bioinformatics and a long-standing interest in handling and analysing large datasets.

**Roshene McCool,** SKA*, United Kingdom*

ROSHENE MCCOOL is the Element System Engineer for Signal Transport and Networks for the SKA (Square Kilometre Array). She has eighteen years' experience in the field of optical transmission. Roshene joined the SKA Organisation in 2009. As well as her role as a specialist in the field of signal transport she supports the system engineering for the SKA and has a management role in the coordination of work packages for the development of solutions in the domain of signal transport and networks for the SKA. Roshene is a chartered member of the Institution of Engineering and Technology (IET).

back to contents page