



PaN-data ODI

Deliverable D5.1

Specific requirements for the virtual laboratories

Grant Agreement Number	RI-283556
Project Title	PaN-data Open Data Infrastructure
Title of Deliverable	Specific requirements for the virtual laboratories
Deliverable Number	D5.1
Lead Beneficiary	STFC
Deliverable Dissemination Level	Public
Deliverable Nature	Report
Contractual Delivery Date	01 Apr 2012 (Month 6)
Actual Delivery Date	23 Apr 2012

The PaN-data ODI project is partly funded by the European Commission under the 7th Framework Programme, Information Society Technologies, Research Infrastructures.

Abstract

This document discusses the requirements of the virtual laboratories for the PaN-data ODI service and research activities.

Keyword list

PaN-data ODI, Virtual laboratories

Document approval

Approved for submission to EC by all partners on 21.04.2012

Revision history

Issue	Author(s)	Date	Description
1.0	Thorsten Kracht	21 Apr 2012	Complete version for approval

Table of contents

	Page
1 INTRODUCTION.....	4
2 A COMMON AUTHENTICATION SYSTEM.....	4
3 STANDARD DATA FORMAT	5
3.1 THE NEXUS LIBRARY.....	5
3.2 THE APPLICATION DEFINITIONS.....	6
3.3 CONVERTING NEXUS FILES TO OTHER FORMATS	6
3.4 CDMA – AN ADDITIONAL ABSTRACTION LAYER INTERFACING FILES AND APPLICATIONS.....	7
3.5 DATA COLLECTION	7
3.6 SUMMARY.....	7
4 DATA CATALOGUE.....	7
5 CENTRAL SOFTWARE DATABASE	8
6 WORK FLOW MANAGEMENT SYSTEM.....	9
7 PROVENANCE.....	10
8 PRESERVATION.....	10
9 SCALABILITY	10
10 SUMMARY	10

1 Introduction

The PaN-data ODI project is a common initiative of eleven European neutron and photon sources which has the goal to create an Open Data Infrastructure among the participating partners. It covers various services and research activities that will have a significant impact on the work of the users of the facilities. The services are a common authentication system, a common data catalogue and a public software repository. In addition PaN-data ODI coordinates joint research activities in the following areas: recording the data continuum (provenance), establishing a system for long-term data storage (preservation) and developing a parallelized data processing framework (scalability). A general deployment of the standard data format is essential for the success of the whole project.

Since PaN-data ODI comprises so many activities, it may be worthwhile to illustrate the expected outcome of the whole project with the help of the following scenario: Scientists who wish to carry out investigations at a European neutron or photon source will be able to submit their proposals to any of the Digital User Offices using the same credentials. During their beam time they will create files having a standardized format. The data processing speed and the I/O performance of the file systems at the facilities will cope with the output rates of all relevant detectors. Self-explaining metadata will be created that fully describe the experiment and the measurement. The facilities will archive the data and take care of long term storage. Catalogues will be created to store metadata. Keyword-based searches will be possible across facilities using a single sign on system. Data portals will support search-runs and data downloads. In addition they will be the interface to workflow systems for data processing. The data continuum will be recorded to allow for a validation of scientific results.

Work package 5 (virtual laboratories) will setup three test cases to demonstrate the outcome of the PaN-data ODI activities:

- VL1: Structural 'joint refinement' against x-ray and neutron powder diffraction data
- VL2: Simultaneous analysis of SAXS and SANS data for large scale structures
- VL3: Access to tomography data exemplified through paleontological samples

WP5 will continuously interact with other work packages: software will be installed once a new version is released. It will be tested within the virtual laboratories and feedback will be returned to the developers.

The first task of WP5 is to gather the requirements for the virtual laboratories as an input for the other work packages. The following sections contain the requirements for each service or research activity.

2 A common authentication system

For scientists it is common practice to perform measurements at different facilities. They would benefit from a system which allows them to access the resources of the individual institutes in a unified way. A single sign on system as proposed by the Umbrella initiative should be implemented for the following services:

- Digital user office
- Data portal, data catalogue
- Software database. Credentials are needed, if users want to upload information.
- In the future, beamline control and data acquisition could also use the same credentials.

The above mentioned services are operational only, if the authentication is integrated into an authorization system. Authorization is not standardized among the PaN-data partners because it depends very much on the specific local IT infrastructure. Therefore it is important that services are presented to the users through interfaces that hide implementation specific features.

Umbrella works with a central database containing unique user IDs. It does not replace the user databases at the facilities. This procedure generates inconsistent records of user information, if scientists change their affiliation. The solution is a central affiliation database service which has procedures to maintain the data and to distribute them to all participating institutes; for this the users must provide their agreement to avoid data protection issues.

Performance can also be an issue. Currently the PaN-data institutes have > 30000 users. The common authentication system has to cope with this number.

3 Standard data format

A standard data format has to serve several purposes:

- It defines an interface that facilitates the interchange of applications and data across facilities and disciplines,
- It supports efficient disk or tape storage by gathering lots of data pieces into a single container file.
- It supports data management by storing data and metadata in a self-describing, well-defined structure.
- The neutron and photon science community can take advantage of the developments of work package 8 (scalability) only, if the standard data format is generally deployed.

PaN-data Europe selected HDF5 as the underlying format and considered NeXus to be the starting point for a specification of file structures and application definitions. Implementing NeXus/HDF5 raises two different kinds of questions: on the one hand there are technical issues concerning the library and the bindings and on the other hand suitable application definitions have to be found.

3.1 The NeXus library

This is the list of requirements for the NeXus library:

- The NeXus library must contain at least a C++ API and Python bindings because these languages are widely used in the neutron and photon community.
- The interface should be strictly object oriented and type-safe. The latter means that data have to be fully qualified when they are transferred between an application and the library. The aim of this restriction is to avoid misinterpretations of the data structures and thereby make the usage of the NeXus library more reliable.
- The NeXus library should not cause substantial performance losses compared to HDF5 as far as I/O speed is concerned. A few percent overhead seems to be acceptable.
- The library code has to be maintained in a public repository to allow for a cooperative code development.
- Detector vendors should be encouraged to write NeXus/HDF5 files.
- Data compression is of increasing importance. The library should make the default HDF5 filters available and allow for additional, discipline-specific filters which are provided by application programmers.

- Detector vendors may install FPGA- or ASIC-based data compression in the near future. Future releases of NeXus/HDF5 have to implement a scheme for treating external filters in a way that they are transparent from the users point of view.
- A complete documentation of the NeXus library, including a user guide, is mandatory.
- All NeXus developments have to be coordinated with the NeXus International Advisory Committee (NIAC).
- Future NeXus/HDF5 versions should implement the multiple writer/multiple reader functionality.

3.2 The application definitions

Application definitions (ADs) allow analysis programs to interpret the contents of NeXus files. They consist of metadata like the name of the investigator, the name of the instrument, the description of the experiment, etc. and of information that describes how the measured quantities relate to variables in physical space. In general, application definitions are specific for experimental techniques. However, since the virtual laboratories VL1 and VL2 are both scattering experiments they can use the same AD: "Application Definition of a Scattering Experiment"¹. This AD has been developed within the High Data Rate Initiative (HDRI) of the PNI institutes of the Helmholtz Society. It comprises a list of metadata and quantities that describe the beam, the sample and the detectors including definitions for the corresponding coordinate system.

The selection of the variables that are represented in an AD is usually a subject of intense discussions. For the scattering AD the choice of the coordinate systems is of special importance because it determines how experiment coordinates (beam, sample, detector) are mapped to reciprocal space. The scattering AD has been discussed in several workshops. It has reached a mature state now.

At least some of the metadata that are part of the AD will appear as keywords for the data catalogue, the software catalogue, the long term preservation and the process that records the data continuum. The integration of these services is a major objective of the PaN-data ODI project. Therefore a dictionary consisting of standardized the keywords needs to be developed. It will facilitate keyword-based searches.

During the setup of the test cases VL1 and VL2 NeXus will be deployed at two scattering experiments. The analysis applications will be customized accordingly and files will be ingested into a data catalogue. This process will show whether the scattering ADs are sufficient to describe the experiments. If necessary the AD will be extended. Eventually the scattering AD will be presented to the NIAC and once it is approved it will be distributed to the international community of neutron and photon science.

3.3 Converters

We cannot expect that all important reconstruction and analysis programs from the neutron and photon science are instantly adapted to the standard data format. For this reason converters should be developed that extract data from NeXus files and store it in a convenient format. The opposite direction, i.e. the conversion of legacy data to NeXus files is also a valuable option pro-

¹ www.pni-hdri.de

vided that sufficient metadata is available. Such a conversion has the advantage that all data can be treated uniformly. It will be the task of the facilities to provide users with suitable software.

3.4 CDMA – An additional abstraction layer interfacing files and applications

Experience shows that NeXus files which are created at different beamlines or at different facilities do not necessarily have the identical structure although they come from the same experimental technique. The CDMA² initiative has been started by SOLEIL and ANSTO to overcome this problem. DESY has joined this initiative. CDMA uses a dictionary mechanism to relate application specific notions to file entities. New file formats are included by providing a suitable plugin. APIs have been developed for Java and C++. Python bindings will follow soon. CDMA can become an important component of WP5 when applications need to be interfaced to data files

3.5 Data collection

Measurements are carried out in a network environment with one PC, the experiment PC, controlling the data acquisition. This PC opens the output files and stores scalar and most likely also 1D data. However, area detectors are typically delivered with extra PCs that write the images to separate file systems. For performance reasons the data stream is not transferred through the experiment PCs. After the measuring sequence has finished, the output file has to be extended to store all data of the sequence, including the images. This can be done by moving the camera frames into the output file and thereby converting them into a device independent format. The other option is to use external links. The images stay where they are and the output files store their path names only. The first option facilitates the data management because a whole measuring sequence is contained in a single file and it has the additional advantage that analysis programs do not have to convert the data. The second option, external links, is less time consuming because no data transfer or conversion is involved. The virtual laboratories will show which procedure is preferable.

3.6 Summary

To summarize the issues concerning the standard data format

- A standard data format allows applications to process data across platforms and facilities.
- NeXus/HDF5 is officially accepted by the PaN-data institutes. Still it has to be deployed everywhere.
- The PaN-data community has to agree on application definitions.
- The development of a common keyword dictionary is essential, not only for the implementation of NeXus/HDF5 but also for the integration of the other PaN-data ODI services.

4 Data catalogue

A data catalogue is a service which includes a database that stores file names and related metadata. It provides an API and the end-user interaction can be through a web portal. For the virtual laboratories the following points are essential:

- An authentication system has to be implemented which allows for operations across facility borders.
- Keyword-based searches have to be supported.

² code.google.com/p/cdma

- The metadata model must be extendable in order to support discipline specific annotations.
- Unique identifiers must be implemented to allow for citations.
- The catalogue should be implemented on a DB without utilizing vendor specific technologies.
- A data portal should include the following features
 - Download of files
 - Upload of files
 - Workflow support
 - A basic data management system must be implemented.
- The functionality as described for the portal should be available from an API.
- The API should support mounting the catalogue as a file share or via nfs.
- The performance of the catalogue should be sufficient to allow for interactive work.
- The API and the portal should be documented. A user guide should be prepared.
- The interfaces to the local authorization and file systems should be prepared in a way that the catalogue can easily be deployed at a facility.

PaN-data Europe denotes ICAT³ to be a strong candidate for a common data catalogue. It comprises many of the required features. WP4 will perform a detailed evaluation and comparison with other products.

5 Central software database

A central software database gives an overview of the applications that exist in the field of neutron and photon science. Scientists will search this database when they encounter a problem that can be solved with software. In the best case they find a program that meets their needs. If not, there is still a chance that the database returns some information about a project with a similar scope. It may be possible to adapt the code to the new problem. Hence, a software database is valuable for those who search software. From the perspective of a developer, the software database is also beneficial because the code improves, if other users run it. Missing features will be added and bugs will be removed.

PaN-data Europe D6.2 proposes that the metadata should comprise the following items:

- Name, version
- URL of the home page
- Author, maintainer, incl. affiliation
- Language
- Hardware/software requirements, supported platforms
- Supported data formats
- Description of the software, including supported scientific categories and instruments, using a 'common vocabulary'
- License type
- Usage guide

A discussion forum, if it is established, would be the place to report bugs or missing features and to ask for help. It is questionable whether a forum should explicitly display some kind of ranking be-

³ www.icatproject.org

cause this might prevent developers from publishing their software in the database. Scientists will find other means to evaluate the software anyhow.

6 Workflow management system

Facilities have to provide computing resources for data processing, simulation and modeling to help scientists who do not have the sufficient IT infrastructures at their home institutes or the necessary expertise. The computational tasks are best organized with a modular workflow system. Workflows may consist of a single step. But in general workflows are a sequence of data processing tasks. Each step has a specific purpose and is carried out by a module with well-defined input/output interfaces. This way, modules can be combined to form an arbitrary order and they can be reused by other workflows. The workflows are prepared by the facility experts according to the discipline-specific needs.

Workflow systems have several advantages:

- Sequences which have been installed for external users can also be used for in-house research and the near real-time analysis. The latter is essential for monitoring the data quality for ongoing measurements.
- Recording the data analysis process is a major goal of PaN-data ODI which can only be achieved if the logging process is implemented in an automated way. It should not depend on the willingness of the user to provide all relevant information. Each module that is part of a workflow should automatically create logging information which fully describes the analysis process.
- A workflow system helps to improve the quality of the program code.
- Modules can run on dedicated hardware.

The following requirements apply to a workflow system:

- Ideally the workflow system is integrated into the data portal because this is the most natural access point for an external user. In addition, portals comprise services that are needed by workflow systems: remote access, authentication, authorization, selection of input data, download of derived data, etc.
- The portal should also be used to start the process, to monitor the progress, to inspect the results and to download derived data.
- It has already been mentioned that workflow modules have to have well-defined interfaces. This is also true for the first step, the input of the data. Hence, the standard data format is very beneficial for workflow systems.
- The workflow system should be easy to maintain. There are several ways how such a system can be implemented. Two possible solutions have been presented recently: DAWN⁴ and DPDAK⁵. The first is based on the Eclipse⁶ IDE, the second is a plugin framework for Python scripts. The future will show which approach suits our needs best.

⁴ DAWN (<https://sites.google.com/site/dataanalysisworkbench/home/latest-announcements>)

⁵ <http://dpdak.desy.de>

⁶ <http://www.eclipse.org>

7 Provenance

The aim of this work package is to install a conceptual framework which allows the validation of scientific results. It will be based on an automated logging mechanism that records all data processing steps. It has been pointed out in a preceding section that a workflow system could be a solution for this problem. It will be necessary to unambiguously identify raw data and software versions. The deployment of the logging mechanism requires a discipline-specific vocabulary for beamline, instrument, equipment, technique, etc. It has to be coordinated with the other PaN-data ODI work packages.

8 Preservation

A framework which guaranties that data can be re-analyzed in the future depends on the following components:

- Long term data storage has to be provided.
- A scheme for metadata management following the OAIS standard has to be implemented.
- The integrity of the data has to be assured.
- A system that handles persistent identifiers for data and related software has to be established.

9 Scalability

Image sensors that are currently available produce data rates up to 700 MB/s. The next generation of area detectors will produce even higher output rates posing increasing demands for processor speed and disk I/O performance. WP-8 addresses this challenge by investigating multicore environments, parallel file systems and pHDF5 which is the parallelized version of HDF5. Ideally these components should be presented to a user as an integrated framework.

- A NeXus interface has to be developed for pHDF5.
- Python bindings should be provided.
- The framework should be configurable in terms of chunk sizes and cache layout.
- Optimal configurations should be provided for the relevant detectors and file systems.
- EPICS or Tango servers have to be created.
- Windows based detectors should be supported.
- The whole system should be scalable to cope with future needs.

10 Summary

When collecting the requirements for the virtual laboratories the following points turned out to be of special importance.

- Several PaN-data ODI activities use keywords. They appear in application definitions, the documentation of the analysis process, the long term data storage and the software database. All activities should agree on the same vocabulary.
- Python appears to be the standard scripting language which is referred to by various work packages.

- Data management and data processing will be done with specialized tools. Users need well defined interfaces, easy to use graphical interfaces and good documentation
- A generally deployed standard data format is the basis of the PaN-data ODI project.