

Towards a Scalable Long-term Preservation Repository for Scientific Research Datasets

Arif Shaon, Simon Lambert, Erica Yang, Catherine Jones, Brian Matthews, Tom Griffin
STFC Rutherford Appleton Laboratory, UK

Introduction

Despite the considerable progress made in the area of digital preservation over the past few years, the task of preserving digital information over the long-term remains markedly challenging. This is mainly due to the varying levels of complexity of the structural and semantic dependencies (i.e. representation information) of digital data that also need to be preserved to enable reusability. Moreover, for substantially large volumes of data, there is also the need to devise scalable preservation solutions to ensure effectiveness - this makes the task of long-term digital preservation incredibly daunting.

In essence, scientific datasets inherit the preservation challenges inherent in all digital information - e.g. capturing adequate representation information, facilitating discovery and performing preservation operations, such as migration or emulation to address technological obsolescence. Additionally, effective long-term preservation of scientific data would also need to address the unique requirements to **preserve the original context** (such as processed data and final publications) of the experiment or research which generated the data in the first place. Preserving context could also involve preserving software or algorithms used in experiments or analyses to provide the complete record of provenance associated with a scientific dataset, and thereby facilitating reproducibility of the analyses or experiments as well as effective re-use of the data [1].

The Science and Technology Facilities Council (STFC) has the responsibility to preserve and curate the large amounts of scientific data generated from its various scientific facilities, such as the ISIS suite of neutron and muon instruments. Recognising the aforementioned need for scalable preservation solutions for large volumes of scientific datasets, STFC is currently participating in a number of preservation-related EU Projects, notably the EU SCAPE¹ project that is developing cloud-based preservation solutions using Apache Hadoop². The outcomes of STFC's work in these projects are intended to improve the existing STFC data archives towards a scalable long-term preservation infrastructure. In this paper, we report the work done to date and outline our vision for building a scalable long-term repository for STFC's scientific research datasets.

Current STFC Data Management Infrastructure

STFC supports a comprehensive range of research services and infrastructure for scientific endeavours, with particular emphasis on the science conducted in its scientific facilities, such as ISIS. The "**Existing System**" segment of Figure 1 gives a general overview of the systems involved in the existing data management infrastructure³. The workflow embedded in the infrastructure begins

¹ <http://www.scape-project.eu/>

² <http://hadoop.apache.org/>

³ It should be noted that there are many other systems and services that STFC supports which are not illustrated in this diagram. We primarily focus on the systems that are directly involved in the dealing of experimental data.

current SDB-based system. In future, we might consider integrating the SCAPE system - once the prototype services are fine-tuned and transformed into production-level services - with the SDB-based system or vice versa (whichever proves more feasible) in order to provide a highly scalable and efficient long-term preservation system for the STFC scientific data.

The new preservation-related services being developed are intended to facilitate the preservation workflows for the STFC scientific data assets. This involves the development of the following *new* preservation workflow services:

The **Staging Service** acts as a 'staging dock' to gather all the datasets and the information related to an investigation to be ready. This is a key service introduced to start the preservation workflow. The information includes metadata from the ICAT Ingest process, experimental data from the distributed file system, information extracted from facility website (e.g. ISIS) or supplied by the facility and software description and copies retrieved from software repositories related to data reduction and analysis processes of the current experiments. How much collection-level metadata is needed is defined by the STFC preservation metadata policy from the policy store accessible to the staging service. Once the staging service performs all the steps, it hands it over to the next phase – the **Pre-Ingest** service.

The aims of the **Pre-Ingest** service are to validate the integrity of the investigation-level and collection-level metadata, both specified in XMLs, generate checksum, message digest, OR digital signature and generate the SIP for an experimental dataset (comprising of all the data related to an experimental investigation) and hands it over to the SCAPE Ingest service. This service has access to the workflow store - a *new* persistent storage for workflows that is being developed as part of the SCAPE project. The workflows are not generated by this service at runtime but rather the workflow instances relevant to the datasets are invoked. We envisage running these workflow instances on the **STFC Hadoop-based Workflow Execution Platform** that is also being developed by the SCAPE project.

The **Ingest** service receives the Submission Information Package (SIP) for an experimental dataset from the Pre-ingest service and performs fixity validation on it. This would involve calculating checksum/message digest/digital signature and matching it against the one recorded at the pre-ingest stage. A successfully authenticated SIP is passed onto the SCAPE Planning/Action service. In common with the pre-ingest service, we also envisage executing the workflows for the ingest service on the local STFC Workflow Execution platform.

The primary objective of the **Action** service is to perform a set of pre-defined preservation-related actions on an SIP received from the Ingest service. These actions are defined based on the analysis of the dependencies and preservation objectives articulated in the Preservation Network Model (PNM - described below) associated with an SIP. The preservation actions performed by this service will be orchestrated in a pre-created workflow, which will also run on the local STFC workflow execution platform to ensure scalability. The execution of a workflow and/or the actions embedded therein may need to take into account any preservation or archival rules or policies (from the Policy Store) applicable to the associated SIP. The final output of this service is an OAIS-compliant Archival Information Package (AIP), which is stored in the STFC e-Science Preservation archive.

As noted before, STFC aims to use PNM records to underpin **preservation planning** and other related activities needed for the effective preservation of its scientific datasets. Notably, the concept of PNM was developed by the CASPAR project in order to represent the output of a preservation analysis conducted for an OAIS -compliant Archival Information Package (AIP) (i.e. a digital object to be preserved) in a preservation archive or repository [2]. In effect, a PNM record can be used to articulate the result of preservation analysis as a network of related objects along with the preservation decisions associated with the relationships between the objects.

The PNMs are created manually based on understanding of the collections, i.e. the experimental investigation and the instruments. The tasks of creating and managing PNM records, the corresponding workflows, and possibly preservation policies and rules will be facilitated by the **STFC Preservation Planning Service**. The underlying architecture of this service will be decided later in the project but we envisage it being a Web Service or portal, possibly integrated with existing related tools, such as Plato⁵, and supplemented by other new manual or semi-automated standalone tools and/or client applications –e.g. a desktop graphical tool for creating and managing PNM records.

As noted above, the new STFC preservation system is designed to leverage the **STFC Hadoop-based Workflow Execution Platform** (being developed by the SCAPE project) to execute the workflows in a scalable fashion. The platform will be hosted locally at STFC. Finally, the system also interacts with a Preservation Watch service, which is a service developed by the SCAPE project to monitor content and the external environment for changes bearing on the preservation of the datasets.

Conclusions and Future Directions

In our participation in the SCAPE project, we have adopted a two-phase approach to developing the proposed scalable preservation repository for the STFC scientific datasets. In the current phase, we are developing various preservation operations, such as characterisation and format migration, needed for the STFC scientific datasets as well as exploring the scalability aspects of these operations with respect to the Hadoop-based platform on which they will be executed. The future phases of this work are expected to focus on capturing rich context and linking of datasets and associated representation information among other activities. This will be driven by analysis of the Preservation Network Models, using semi-automated tools and services. Finally, we are also developing a small scale testbed to allow us to develop an understanding of the resource implications of the proposed preservation infrastructure and its suitability to our facilities.

References

[1] Brian Matthews, Arif Shaon, Juan Bicarregui, Catherine Jones, A Framework for Software Preservation, *International Journal of Digital Curation*, 5 (1), 2011

<http://www.ijdc.net/index.php/ijdc/article/view/148>

[2] Conway, E., Dunckley, MJ., Giaretta, D. and Mcilwrath, B. (2009): Preservation Network Models: Creating Stable Networks of Information to Ensure the Long Term use of Scientific Data, *Proc. Ensuring Long-Term Preservation and Adding Value to Scientific and Technical Data (PV 2009)*, Villafranca del Castillo, Madrid, Spain, 01-03 Dec 2009. Retrieved from:

http://epubs.cclrc.ac.uk/bitstream/4314/PV09_Conway_PNM.pdf

⁵ Plato: The Preservation Planning Tool (originally developed by the EU Planets project) - <http://www.ifs.tuwien.ac.at/dp/plato/intro.html>