

Using iRODS to manage, share and publish research data: Yoda

Ton Smeele
Utrecht University ITS/RDM
Heidelberglaan 8, Utrecht,
The Netherlands
a.p.m.smeele@uu.nl

Lazlo Westerhof
Utrecht University ITS/RDM
Heidelberglaan 8, Utrecht,
The Netherlands
l.r.westerhof@uu.nl

ABSTRACT

Researchers face challenges when they want to manage, share and publish their research data.

In 2014, Utrecht University commenced development of Yoda, a research data management system to meet these challenges. More recently features have been added to facilitate researchers to describe, deposit and publish research data in compliance with the FAIR principles.

Yoda deploys iRODS as its core component, customized with approximately 10,000 lines of rules and extended with a graphical user interface. It accommodates multiple metadata schemas to support varying requirements across fields of science.

We will discuss Yoda design principles related to the new features and their iRODS realization.

Keywords

Research data management, iRODS, FAIR, metadata.

INTRODUCTION

As the amount, complexity and use of digital research data grows, so does the need to manage data properly. The requirements for research data management may vary during the research data life cycle [1]. For instance, during creation, processing and analysis phases sensitive data typically is shared only between project members while the data sharing and reuse phases imply the need for publication of metadata to the research community at large. Other requirements such as compliance with law and institute policies have impact on the entire data life cycle.

Many research data management tools are focused only on the last phases of the data life cycle. Our approach is to support data management *during* research so that researchers can benefit from services immediately and data management integrates smoothly with existing research processes. In this paper we discuss key requirements and principles related to the design of Yoda and their realization using iRODS [2].

CHALLENGES

Need to safely collaborate on research data

An internal survey held in 2014 at Utrecht University amongst 3200 researchers revealed that approximately two third of the researchers work with research data which is at least partly sensitive. Half of the researchers that already work with sensitive data expect this fraction to increase over time.

European, national and institutional funders stimulate international and cross discipline research cooperation. Therefore while data may be sensitive, often there is a need to collaborate on data between research consortium members that

originate from different institutes and countries. An example is the Consortium on Individual Development [3], a longitudinal research project in the Netherlands that seeks to understand why most children develop well yet not all children manage to do so.

Need powerful tool yet easy to use

Some disciplines such as Bioinformatics tend to prefer to use the command line e.g. Linux bash to access their data. They may have to manage files with sizes in the order of many gigabytes for which fast and resumable data transfer is a key requirement to minimize the impact of poor network connections. Large data files may also imply a need to have multiple physical storage units act as one logical unit. The iRODS iCommands [4] are well suited to their needs.

Yet most disciplines demand a graphical user interface with intuitive access, informally summarized during a Yoda requirements for I-Lab [5] workshop as “We need Dropbox [6] on steroids”. This workshop also revealed the wish to be able to revert to a previous revision of a data file e.g. in case of accidental deletion.

Need to comply with data policies

For reasons of reproducibility, Utrecht University’s research data policy [7] requires researchers to retain data underpinning research output for a period of at least 10 years after conclusion of the research. This is in line with national [8] and European [9] guidelines. Patent, privacy or otherwise sensitive data necessitates appropriate measures to protect against unauthorized access.

All the above policies require research data to be protected against loss as a result of hazards or data corruption. Replication across multiple geographically distributed datacenters is adopted as a strategy against data loss by the EUDAT project’s B2SAFE services [10].

Need support for data beyond the research project

Research data often has value beyond the scope of a research project. A UKRDS survey [11] reveals that duration of the data life cycle tends to vary per discipline where less than 20 years is common for business & management studies while e.g. geography data does not appear to expire. Therefore stewardship of data is unable to depend on the existence of and funding by the research project. This underpins the importance of well described data packages so that entitlements and reuseability of the data can be assessed at any time. The FAIR principles [12] provide researchers with guidance on the required level of data description.

The availability of reusable data can benefit new research and is promoted via initiatives such as the European Open Science Cloud [13].

Flexible cost coverage for data storage and data management is needed to support changes in stewardship and to support reuse based billing.

SOLUTION

Communities and data compartments

Yoda stores data in compartments and access to the compartment is granted only to users that are a member, a viewer or a data manager. Members have edit rights on data while viewers have read-only access to the data. Multiple data compartments are jointly referred to as a community¹. Each community can have one or more data managers.

Data compartments are implemented as a set of iRODS groups so that authorizations can be role based. Note that in iRODS access rights granted on an object are equivalent to the sum of access granted to each group membership (in addition to access granted to the user specifically) so this conveniently fits with our approach.

¹in the Yoda portal a community is listed as a “category”

Given a compartment named X belonging to community C then (full) members are listed in the iRODS group `research-X` and viewers are listed in the iRODS group `read-X`. Datamanagers are listed in a group `datamanager-C`.

Data compartments consist of a research “workspace” and an accompanying “vault”. The vault is a space where data is deposited that must remain immutable forever for reasons of demonstrable data integrity. While data objects in the workspace are editable for full members and readable for viewers and data managers, the data in the vault is read-only for all roles. Researchers may opt to deposit data in the vault. Subsequently the data manager approves the deposit and the Yoda system takes care of the actual deposit action. Hence even the data manager cannot directly write to the vault.

Storage resource strategy

Each iRODS storage resource is tagged with a tier identifier that indicates its relative cost factor. Yoda administrators can edit this tag. Hence a community can be invoiced for the total size of data stored on Yoda based on the tier rate times the sum of data object sizes for each storage tier. This information is collected by a cron job and annotated to the iRODS group object. The Yoda statistics module shows users their current and rolling 12-months use per tier consolidated on a group and community level.

We use composable resources in a hierarchy. All user accessible resources are underneath resource `irodsResc`. The data stored here is replicated asynchronously to a top level resource `irodsRescRep1` located in another datacenter to protect it against loss. This is complemented by a revision strategy to maintain previous versions of a data object.

The `irodsResc` resource distributes data objects randomly among its children, the actual leaf storage resources. Each leaf has a pass-thru parent resource to allow an administrator to switch the resource on and off for reads and/or writes. This facilitates easy maintenance on resources. Resources can be decommissioned by replicating their data to another resource and then trimming the original resource.

Ease of use: graphical user interface

One of the requirements of Yoda is a graphical user interface with intuitive access. We meet this requirement through a combination of a network disk access to transfer and process data and a web portal to authorize access and to manage the data. The network disk interface Davrods [14] is written in C and supports WebDAV protocol class 2. The portal is written in PHP, JavaScript and HTML. Using the iRODS PHP library [15], the graphical user interface of Yoda is connected to the back office. It is strictly presentation layer only, all business logic is present in the back office. The business logic is implemented using iRODS rules and microservices. This facilitates that Yoda functions and policies are applied regardless of user interface.

Collaboration: revisions and folder locking

To satisfy the wish to be able to revert to a previous revision of a file and collaborate safely as a group in the research workspace, we have implemented folder locking and revisions.

Folder locking

To prevent modification of a folder in the research workspace, a researcher can use the lock function. This will catch attempts to write data to or to remove data from that folder until the folder is unlocked again. The same mechanism is used to prevent data modification while it is being approved and copied to the vault. We have opted not to use iRODS Access Control Lists (ACL) for this purpose. ACLs can be modified by the object creator which is not desirable for system-initiated lock actions.

Folder locks are implemented through metadata on the folder, its children and its parents. The lock consists of an iRODS Attribute-Value-Unit (AVU) triplet with `org_lock` as attribute name and the name of the collection that the lock applies to as its value. Whenever a folder is locked an AVU is set on every child recursively. In addition it is set on all parent folders up to and including `/rodsZone/home/research-{groupName}`.

Revisions

For each new file and upon a file modification Yoda creates a revision. This is a timestamped backup of the new or modified file in the revision store. The revision store is a system collection.

Not all revisions are kept, only a predefined number of revisions is kept per time bucket. A time bucket is a time offset from now into the past. Each revision strategy has a predefined set of time buckets and minimum number of revisions stored in those buckets.

Metadata handling

Metadata is registered in an XML [16] file together with the data in a data package. This XML file complies with a metadata schema chosen by the community. Researchers can register the metadata through the web portal or upload a metadata XML file directly through WebDAV. The second scenario is suitable for experiment software that can produce a metadata file automatically.

Metadata schemas

Each community in Yoda can configure its own metadata schema. Metadata can be entered and edited through a web form. A Yoda metadata form is constructed from three files (see Figure 1): an Extensible Markup Language (XML) file that contains the actual metadata, an XML Schema Definition (XSD) [17] file and a form elements XML file.

The XSD file lists requirements and restrictions that apply to the metadata. It describes properties such as input length and datatype of each metadata field and selectable options of List-Of-Values fields. While Yoda uses the XSD to validate the XML structure, it does not use the XSD to specify mandatory metadata fields required before a data folder can be submitted for inclusion into the vault.

The form elements XML file contains all presentation information relevant to the web portal metadata form and specifies which metadata fields are required when submitting a data folder into the vault. This file has been designed for the Yoda implementation and is not based on any standard apart from XML.

The Yoda metadata XML file holds the actual metadata as entered by the researcher in the research workspace.

In the research workspace, the XML metadata file is validated against the metadata schema XSD before it is loaded into the metadata form. If the metadata XML does not comply with the metadata schema, it will not be loaded into the metadata form. When a data folder is deposited to the vault, it is validated against the metadata schema XSD and against the form elements XML to check if all required metadata fields have been completed.

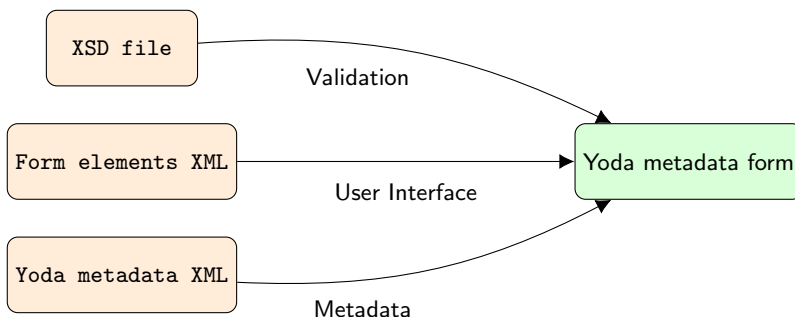


Figure 1. Construction of Yoda metadata form

Metadata indexing

The content of the XML metadata file is indexed in the iCAT database so that users can find data by its metadata. The submitted metadata forms are transformed with an Extensible Stylesheet Language Transformations (XSLT) [18] stylesheet into an XML containing AVUs. Subsequently this XML is loaded into the iCAT database using a microservice. Metadata is added to the data folder to make it searchable in the web portal.

Metadata mapping

When Yoda publishes a data package, the metadata stored in XML is transformed into several other formats using XSLT stylesheets. The first transformation creates a landing page for the publication. This page is published in a public directory. A second transformation generates a DataCite Metadata Schema 4.0 [19] compliant XML used for the DataCite Metadata Store. The last transformation creates an XML file used by the public OAI-PMH [20] service of Yoda which maps the Yoda metadata to Dublin Core Schema [21] and DataCite Metadata Schema 4.0.

Workflows

The two main workflows in Yoda are depositing a data package (see Figure 2) and publishing a data package (see Figure 3). We faced several challenges upon workflow implementation. First of all integrity: data packages in the vault must be protected from undesired changes. In case of desired changes we need traceability: it should be traceable who was responsible for each change. Furthermore we want interactive and scalable workflows. Users should not have to wait for tasks to finish and long running tasks should not block other processes.

To provide security and traceability for data packages in the vault, tasks in the vault require exclusive privileges. These privileges are granted to actors based on their role. To achieve interactive and scalable workflows, tasks with long or uncertain execution time must be executed asynchronously (e.g. file copy actions).

Workflow progress is registered as a state with the data package. The state of a data package is stored as metadata on the data package collection.

Actors

Yoda workflows involve three actors:

- **Researcher** Researchers initiate the workflows by submitting data folders for the vault and for publication.
- **Data manager** Data managers react on the actions initiated by the researchers, submitted data folders have to be approved by a data manager.
- **System** The system will copy and publish the packages after they have been approved by a data manager. Since tasks in the vault require exclusive privileges this is done by the system.

Data deposit workflow

Researchers can deposit their data folder from the research workspace into the vault. To start the data deposit workflow a researcher submits a data folder for inclusion into the vault. The data manager will check if the submitted data folder with metadata complies with institutional, faculty and other policies. If the data folder complies with policies the data manager accepts the data folder for the vault. If the data folder does not comply with policies the data manager rejects the data folder. The researcher can now change the data folder before resubmitting it to the vault. After a data folder is accepted for the vault the system will deposit a copy of the data folder in the vault. A data package is created and the data folder content is included in an **original** sub folder. The license chosen by the researcher is automatically added to the data package in the vault.

FAIR data publication workflow

Once secured in the vault, a data package can be published. The researcher initiates a data publication via the Yoda web portal by clicking the **submit for publication** button. The data manager accepts or rejects the request

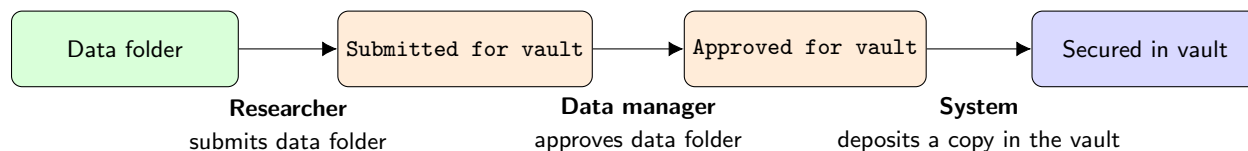


Figure 2. The workflow to deposit a data package into the vault

depending on whether the data package complies with publication policies. If the publication is rejected the researcher can make the necessary changes before submitting the data package again. After the data manager has accepted the publication request, the system will publish the data package. The system mints and registers a persistent Digital Object Identifier [22], submits the data package metadata to DataCite, publishes a landing page for the data package and makes the metadata available through OAI-PMH. These actions make the data package findable, accessible and reusable in line with the FAIR principles. Third parties can harvest the metadata of the published package through OAI-PMH. If the data package is classified as “open”, the contents of the data package will be made publicly accessible as well.

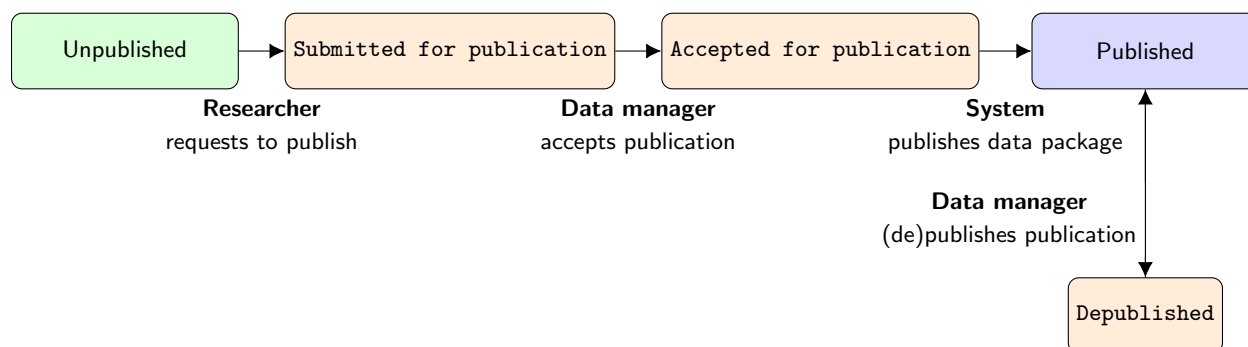


Figure 3. The workflow to publish a data package

FUTURE WORK

Going forward we intend to extend Yoda with functions to manage collections of data packages (e.g. manage membership and import, export and distribute data). Also we aim to support other data types such as software and experiment configurations.

CONCLUSION

The current functions of Yoda facilitate institutional researchers to manage and share sensitive data in line with policies during and after research. Yoda allows researchers to comply with grant requirements on data management and open data.

ACKNOWLEDGEMENTS

The development of Yoda would not have been possible without the input, support and executive sponsorship of researchers and IT management. In particular we would like to thank Chantal Kemner (Dynamics of Youth research) and her team. Chantal has provided the requirements for the initial Yoda functions that help to manage sensitive lab data and she trusted us to deliver on our promises. Also we thank Vincent Buskens (Institutions for Open Societies research) and his team. Vincent helped us to shape the data sharing and data publication functions. Finally we thank Carolien Besselink (IT Services) for her sustained executive sponsorship. Carolien put research data management on the IT Services agenda and facilitated the development of Yoda.

REFERENCES

- [1] S. Higgins, “The DCC curation lifecycle model,” *International Journal of Digital Curation*, vol. 3, no. 1, 2008.
- [2] A. Rajasekar, R. Moore, C.-y. Hou, C. A. Lee, R. Marciano, A. de Torcy, M. Wan, W. Schroeder, S.-Y. Chen, L. Gilbert, *et al.*, “irods primer: integrated rule-oriented data system,” *Synthesis Lectures on Information Concepts, Retrieval, and Services*, vol. 2, no. 1, pp. 1–143, 2010.
- [3] C. Kemner, “Consortium on individual development.” <http://www.individualdevelopment.nl>, 2014. Visted on 2018-05-31.
- [4] “iRODS documentation.” <https://docs.irods.org>. Visted on 2018-05-31.
- [5] “Institutions for Open Societies.” <https://www.uu.nl/en/research/institutions-for-open-societies>. Visted on 2018-05-31.
- [6] “Dropbox.” <http://www.dropbox.com>. Visted on 2018-05-31.
- [7] E. Stiekema, “University policy framework for research data Utrecht University,” 2017.
- [8] Association of Universities in the Netherlands (VSNU), “The Netherlands Code of Conduct for Academic Practice.” [http://www.vsnul.nl/files/documenten/Domeinen/Onderzoek/The_Netherlands_Code%20of_Conduct_for_Academic_Practice_2004_\(version2014\).pdf](http://www.vsnul.nl/files/documenten/Domeinen/Onderzoek/The_Netherlands_Code%20of_Conduct_for_Academic_Practice_2004_(version2014).pdf), 2014. Visted on 2018-05-31.
- [9] “Guidelines on FAIR Data Management in Horizon 2020.” http://ec.europa.eu/research/participants/data/ref/h2020/grants_manual/hi/oa_pilot/h2020-hi-oa-data-mgt_en.pdf, 2018. Visted on 2018-05-31.
- [10] EUDAT, “B2safe.” <https://eudat.eu/services/userdoc/b2safe>. Visted on 2018-05-31.
- [11] N. Beagrie, R. Beagrie, and I. Rowlands, “Research data preservation and access: The views of researchers,” *Ariadne*, no. 60, 2009.
- [12] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne, *et al.*, “The fair guiding principles for scientific data management and stewardship,” *Scientific data*, vol. 3, 2016.
- [13] “European Open Science Cloud.” <http://ec.europa.eu/research/openscience/index.cfm?pg=open-science-cloud>, 2018. Visted on 2018-05-31.
- [14] C. Smeele and T. Smeele, “Davrods, an Apache WebDAV Interface to iRODS,” in *iRODS User Group Meeting 2016 Proceedings*, pp. 41–47, iRODS Consortium, Dec. 2016.
- [15] DICE-UNC, “Prods php irods client library.” <https://github.com/DICE-UNC/irods-php>. Visted on 2018-05-31.
- [16] T. Bray, J. Paoli, C. M. Sperberg-McQueen, E. Maler, and F. Yergeau, “Extensible markup language (xml).,” *World Wide Web Journal*, vol. 2, no. 4, pp. 27–66, 1997.
- [17] P. V. Biron, A. Malhotra, W. W. W. Consortium, *et al.*, “Xml schema part 2: Datatypes,” 2004.
- [18] World Wide Web Consortium (W3C), “XSL Transformations (XSLT) Version 2.0.” <https://www.w3.org/TR/xslt20/>, 2007.
- [19] DataCite Metadata Working Group, “DataCite Metadata Schema for the Publication and Citation of Research Data. Version 4.0.” <http://doi.org/10.5438/0013>, 2016.
- [20] Open Archives Initiative, “The Open Archives Initiative Protocol for Metadata Harvesting. Version 2.0.” <http://www.openarchives.org/OAI/2.0/openarchivesprotocol1.htm>, 2015.
- [21] Dublin Core Metadata Initiative, “The Dublin Core Metadata Element Set.” <http://www.ietf.org/rfc/rfc5013.txt>, 2007.
- [22] I. ISO, “26324: Information and documentation-digital object identifier system,” 2009.