

Progress Towards a NASA Earth Science Reuse Enablement System (RES)

James J. Marshall¹, Robert R. Downs², Chris A. Mattmann^{3,4}

¹*INNOVIM
NASA Goddard Space Flight Center
Mail Stop 614.9
Greenbelt, Maryland 20771 USA
James.J.Marshall@nasa.gov*

³*Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109, USA
mattmann@jpl.nasa.gov*

²*Center for International
Earth Science Information Network
Columbia University
Palisades, NY 10964 USA
rdowns@ciesin.columbia.edu*

⁴*Computer Science Department
University of Southern California
Los Angeles, CA 90089, USA
mattmann@usc.edu*

Abstract

A Reuse Enablement System (RES) allows developers of Earth science software to contribute software for reuse by others and for users to find, select, and obtain software for reuse in their own systems. This paper describes work that the NASA Earth Science Data Systems (ESDS) Software Reuse Working Group has completed to date in the development of an RES for NASA.

Keywords: NASA, Earth science, software reuse

1. Introduction

The National Aeronautics and Space Administration (NASA) Earth Science Data Systems (ESDS) Software Reuse Working Group (WG) was chartered to enable and facilitate reuse of software assets within the Earth science community. To this end, the WG conducted two survey studies of the community, in 2004 and 2005, to gauge their experiences and practice with software reuse [1].

The primary difference between these two studies was the target audience – in 2004 the survey instrument was only sent to government employees and members of the WG, but in 2005, it was sent to members of the larger community of Earth science data systems, including developers, academia, and industry. Both studies returned near-identical results. In particular, the WG found that the Earth science community could benefit from the availability of a system that describes and provides access to Earth science software and applications that are available for reuse [1]. As a result, the WG submitted the following recommendations to NASA [2]:

1. NASA should establish an effective mechanism for dissemination of reusable assets within the Earth science community.

2. Based on the conclusions of a technology evaluation, NASA should implement a reuse enablement system (RES).

NASA responded by requesting that the WG conduct a study of existing systems to determine whether any operational systems or existing software platforms could be leveraged to implement the RES.

Over the past five years, the WG has investigated these issues and has performed a series of studies and related efforts to meet the challenges of software reuse that are faced by the Earth science data systems development community. Initially, use cases were collected to identify the requirements for a system that could facilitate the reuse of software and related components. Next, a trade study was conducted to investigate the availability of systems and services that could address the community's requirements for software reuse. These requirements included: (1) facilitating access to software for potential reuse; and (2) the ability to catalog Earth science software assets and to enable their discovery. A prototype RES was developed and a plan was created for testing the system. In addition, policies were drafted for managing and operating a RES.

This paper describes major steps completed to develop the prototype RES system, enumerates our experience to date, and discusses the possibility of developing a distributed set of reuse enablement systems – an option for which NASA Headquarters has indicated a preference.

2. Use Cases and Requirements

Development of the proposed RES began in 2004 when the WG identified sixteen starting use cases. In 2006 the use cases were formally included as part of the supporting documentation for the proposed RES [3].

The use cases were used to create a set of requirements for the RES. These were also formalized as part of the RES supporting documentation in 2006, resulting in a total of fifty-four requirements (a few of which are shown in the rows of Table 1) [4].

In 2007, the requirements were revised for clarity, but otherwise were left unchanged. During 2008, the WG performed some consistency comparisons between various RES documents, and revised the RES requirements based on these assessments.

3. Trade Study

Following the development of the use cases and requirements in 2004 and in accordance with the direction provided by NASA in response to the WG's original RES recommendation, the WG performed a trade study of several operational systems [5]. These systems were evaluated in terms of the functionality and services offered to support the community of Earth science data systems and software developers. The trade study indicated that no current system met the software reuse needs of the community and concluded that an RES system should be developed.

4. Architecture Study

Based on the results of the completed trade study, the WG continued by conducting an architecture study to identify an approach for the creation of the RES [6]. The architecture study identified the open source XOOPS content management system as a platform that could be adopted and tailored to establish the RES.

As an open source offering, XOOPS is customizable by adopters to meet their specific needs. The content management functionality of XOOPS in particular suited the RES as it afforded software developers the ability to describe and contribute software components for review and reuse by other members of the community.

Also, the platform offered capabilities for administrative users to manage the system. Of the candidate technologies, XOOPS met or partially met the most (45) RES requirements of the technologies studied, leaving only 9 unmet requirements. Additionally, due to its flexibility, the WG estimated that it would require around 8.12 months of effort required to implement the RES, an improvement by a factor of 4 overall in the amount compared to Savane; GCMD and GForge were not estimated due to their lack of applicability to the RES requirements.

5. Prototype RES

In 2007, as the results of the architecture study became clear and were being finalized, the WG used the study's conclusion to begin prototyping the proposed RES.

Table 1. Trade study results for three systems studied

Req./Feature	GCMD	GSFC OSS	Ames OSS
<i>Domain</i>	Earth science	Earth and space science	General science
<i>Type of Assets</i>	Data sets, data services	Open source packages	Open source packages
<i>Register User</i>			
<i>Contribute / Update Assets</i>	**	*	*
<i>System Feedback</i>	**	**	**
<i>Automatic Notifications</i>	***	*	*
<i>Discovering Assets</i>	Hierarchy, Search	List	List
<i>Register Asset Usage</i>		***	***
<i>Provide Asset Reviews</i>		*	
<i>Monitoring Feedback</i>	*	**	*
<i>Secure Login or Registration</i>	N/A	No	No
<i>Catalog or Repository</i>	Catalog	Both	Both
<i>Operation Support</i>	Large	Small	Small
<i>Technology</i>	RSYNC, Zope, CVS, Linux, Java, and others	PHP	JavaServer pages

XOOPS was obtained, installed, configured, and modified to enable it to function in the manner desired for the RES, resulting in a prototype of the system. A few of the main features of the prototype include: a menu bar for logged in users, a site-only search box, user-selectable notifications, average asset ratings with the number of votes for each, a menu of options specific to individual assets, and comment areas for users to provide textual feedback, including reviews on individual assets.

The original plans for the RES included populating a centralized system with content recommended by members of the WG [2]. Access to the RES and its content was initially limited to members of the WG during a period in which formal tests could be conducted, prior to offering wider access to the RES.

Planning for the next set of Earth missions [14] has included deploying a prototype RES within each new mission site. This subject is discussed further in Section 8 below.

6. Test Plan

The Software Reuse WG recently completed a test plan for the proposed RES, designed to ensure that any

prototype or operational system that is created would meet the requirements previously identified and documented by the WG [11]. This test plan assumes that the system is created according to the findings of the architecture study (i.e., XOOPS with selected modules and customization forms the base of the system) and includes detailed descriptions of tests that can be used to evaluate how an implementation of the system meets the original fifty-four requirements that were initially identified for the RES.

While performing some consistency checks between the test plan and the system requirements, the potential for improvement was identified and some of the requirements were modified slightly for clarity. Similar to previous experiences where the WG contributed to the reuse of systems by others [12], the experiences of others could identify recommendations that can be used to improve the RES for subsequent implementations.

7. Policies

Recognizing that the system requirements and use cases may not be sufficient for managing instances of the RES, the WG has developed policies for the operation and maintenance of the system. The initial version of this document has been completed recently [13].

Based on the initial requirements for the RES and reviews of the prototype system, an initial set of thirty policies have been established for managing an RES. These were developed by the WG through iterative discussions of scenarios for conducting various tasks related to the discovery, contribution, selection, and retrieval of software assets while using the RES. Similarly, these policies reflect potential uses of the RES by administrators to manage the system for use by software contributors and by software adopters. In addition, the policies cover issues such as copyright and intellectual property rights, restrictions on use, privacy, and security. The WG is aware that additional reviews of the document may be necessary to ensure that the policies comply with all legal codes and applicable standing policies of NASA.

8. Centralized RES or Distributed Systems

The WG's original recommendation to NASA was for the RES to be a single, centralized system [3–5]. The intention was to provide one place for Earth science software developers to go in order to find reusable software assets. This would help break down some major barriers to software reuse within the community as identified in surveys conducted by the WG [1]. However, as indicated in Section 5, recent direction from NASA Headquarters [14] has indicated that new missions could benefit from the implementation of a set of distributed

systems, run on a per-mission basis. These implementations could begin with the upcoming decadal survey missions, which have been recommended by the National Research Council [15]. The WG has considered this alternative, and the two approaches are discussed further in this section.

8.1. Centralized RES

The WG's original vision of a centralized RES was partially based on the results of surveys of the community of Earth science software developers [1]. These surveys revealed barriers to reuse, including that many people did not practice software reuse either because they did not know reusable assets existed or because they did not know where to locate reusable assets. Additionally, they identified a catalog/repository system for reusable Earth science assets as important.

The recommendation to establish a system came primarily from the latter point and would provide a location for such assets, while the intention to have a single, centralized system came from the former point and would facilitate finding such assets. If there were one place software developers could visit to obtain reusable software assets, it would help break down the barriers noted above. Some of the benefits of a centralized system, apart from providing one-stop shopping, are that it would be domain-specific (focused on assets relevant to Earth science) and operated and maintained by one entity.

8.2. Distributed Systems

As an alternative to a centralized RES, the WG is also considering the possibility of having a distributed set of systems, each functioning as a smaller RES. Some of the advantages of this option are that it allows for area-specific systems (ones designed for specific sub-domains within Earth science) and the systems are generally smaller making them easier to operate and maintain. In terms of control, missions have expressed the desire to maintain the RES on their own physical hardware, leveraging their own system administrators, and software asset curators, who are required to record information about reusable software assets.

Likewise, each RES can be customized to meet the specific needs of the mission in which it has been implemented. Providing capabilities for a mission to contribute software in a manner that has been customized for their use will facilitate adoption of the RES and population of its contents. Enabling other missions to access reusable software within a distributed network will allow software to be shared when it has been prepared and adequately tested for reuse by others.

The adoption of multiple instances of the RES also will provide independent opportunities to examine the potential for the RES to serve the needs of specific

missions. Suggestions from the individually customized instances of the RES will contribute to the design and architecture of the RES, which the WG can further enhance as new versions that will be available to others.

An effort is underway to standardize the software packaging data model for assets cataloged and stored within each RES. This effort will allow for RES data to be exported on a mission-by-mission basis into a centralized RES, providing a global view of assets captured for all contributing missions. The WG is currently working with the Soil Moisture Active/Passive and the Ice, Cloud, and land Elevation Satellite 2 decadal survey missions to help them install, set up, and configure instances of the RES for their use. The WG is also pursuing opportunities to help other missions, such as the Deformation, Ecosystem Structure, and Dynamics of Ice decadal survey mission and the Orbiting Carbon Observatory 2 mission, with their reuse efforts and encouraging them to consider hosting instances of the RES. These collaborations and the development of a distributed set of systems have just begun.

9. Conclusions and Future Work

The WG has gone from the results of surveys of the Earth science community's reuse practices identifying the need for a domain-specific catalog/repository of smaller-sized reusable assets through a series of studies and development efforts to produce a Reuse Enablement System (RES) that is beginning to be adopted by the upcoming Earth science decadal survey missions. The WG will continue to work with its members, the Earth science community, and upcoming missions to refine and improve the RES for future users of such systems.

10. Acknowledgments

The authors would like to acknowledge the funding and support received from NASA, especially the support from Frank Lindsay, the Earth Science Data Systems Working Group Program Manager; the support from NASA contract NAS5-03117 for Robert Downs; and the support from NASA for other members of the Software Reuse Working Group. This effort is also partially supported by the Jet Propulsion Laboratory, managed by the California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

11. References

- [1] J.J. Marshall, S.W. Olding, R.E. Wolfe, and V.E. Delnora, "Software Reuse Within the Earth Science Community," *Proc. IGARSS*, 2006, pp. 2880–2883.
- [2] J.J. Marshall, V. Delnora, R. Downs, R. Gerard, S. Olding, S. Samadi, R. Wolfe, N. Casey, and S. Falke, "Earth Science Software Reuse Enablement Systems", *Eos Trans. AGU*, 86(52), Fall Meet. Suppl., Abstract IN23B-1220.
- [3] NASA ESDS Software Reuse WG. (2006). "Reuse Enablement System (RES) – Use Cases". <http://www.esdswg.com/softwarereuse/Resources/res/RESUseCases.pdf/view>
- [4] NASA ESDS Software Reuse WG. (2007). "Reuse Enablement System (RES) – Requirements". <http://www.esdswg.com/softwarereuse/Resources/res/RESRequirements.pdf/view>
- [5] NASA ESDS Software Reuse WG. (2005). "Reuse Enablement System (RES) – Trade Study Report". <http://www.esdswg.com/softwarereuse/Resources/res/RESTradeStudy.pdf/view>
- [6] NASA ESDS Software Reuse WG. (2008). "Reuse Enablement System (RES) – Architecture Study". <http://www.esdswg.com/softwarereuse/Resources/res/RESArchStudy.pdf/view>
- [7] XOOPS CMS. (2010). XOOPS CMS (Content Management System). <http://www.xoops.org/>
- [8] The Gna! People. (2010). Savane. <http://gna.org/projects/savane/>
- [9] National Aeronautics and Space Administration. (2010). Global Change Master Directory (GCMD). <http://gcmd.nasa.gov/>
- [10] GForge Group. (2010). GForge. <http://gforge.org/>
- [11] NASA ESDS Software Reuse WG. (2010). "Reuse Enablement System (RES) – Test Plan". http://www.esdswg.com/softwarereuse/Resources/res/RESTestPlan_v1.0.pdf/view
- [12] R. Gerard, R.R. Downs, J.J. Marshall, and R.E. Wolfe, "The Software Reuse Working Group: A Case Study in Fostering Reuse", *Proc. IEEE IRI*, 2007, pp. 24–29.
- [13] NASA ESDS Software Reuse WG. (2010). "Reuse Enablement System (RES) – Policies". http://www.esdswg.com/softwarereuse/Resources/res/RESPolicies_v1.0.pdf/view
- [14] J.J. Marshall, R.R. Downs, and S. Samadi, "Relevance of Software Reuse in Building Advanced Scientific Data Processing Systems", *Earth Science Informatics*, 2010, in press.
- [15] National Research Council, *Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond*, National Academies Press, Washington, 2007.