

Securing the iRODS Metadata Catalog for Digital Preservation

Gonçalo Antunes¹ and José Barateiro²

¹ INESC-ID, Information Systems Group, Lisbon, Portugal

² LNEC - Laboratório Nacional de Engenharia Civil, Lisbon, Portugal
goncalo.antunes@tagus.ist.utl.pt, jbarateiro@lnec.pt

Abstract. Digital preservation is the ability to retrieve, access, and use digital objects through time, while ensuring the authenticity and integrity properties of these objects. Data grids represent a model of storage systems designed for data management and sharing, which concept also has been proposed for digital preservation. However, since data grids are not specifically designed for this purpose, they present weaknesses that have to be handled. This poster will present a set of services to address a problem in the metadata catalogue of the iRODS data grid, strengthening that platform for digital preservation purposes.

Keywords: Digital Libraries, Digital Preservation, Data Grids.

1 Introduction

The creation of OAIS by the CCSDS [1] acknowledges the importance of the preservation of scientific data to scientific research (*e-Science*), where research may span several decades and produce large quantities of data that is crucial to preserve. Moreover, preservation of scientific data may be mandatory by law, as it is the case of the Portuguese Dam Safety Legislation [2], which establishes that the National Laboratory of Civil Engineering of Portugal¹ is responsible of keeping an updated archive of data concerning dam safety.

Data grids offer distributed computation and storage of massive data sets, and allow sharing and collaboration, which are crucial for *e-Science*. iRODS², is an open-source data grid system that has been proposed for digital preservation purposes [3]. An important component of the iRODS data grid system is the metadata catalogue, which maintains crucial information like the physical location of files and nodes, state of operations and processes, etc. This metadata catalog is supported by a centralized database system, which represents a point of failure when used in the context of a preservation scenario. If a failure, disaster or intentional attack affects the metadata catalog, it can cause a total data loss, even if the data stored on the other nodes remain intact. Previous work summarizes the threats that can affect the metadata catalogue [4].

In order to reduce the risk of threats to the iRODS metadata catalog, we propose the following set of extensions as new services: (*i*) a Backup Service, which exports

¹ <http://www.lnec.pt>

² <http://www.irods.org>

the contents of the metadata catalogue database to a structure of XML files and replicates them to other nodes of the data grid; and (ii) a Recovery Service, a service to recover the metadata catalogue database in case of failure.

This work has been done in the scope of two projects: the international project SHAMAN³, and the national project GRITO⁴.

2 The Backup Service

The backup service is intended to be sensitive to aspects such as the workload of the system, in order to not interfere with the normal functioning of other regular processes that might be running on the system.

Accordingly, the measure of the current workload conditions has to take place early in the process. We consider the CPU and memory used at a determined time by the node running the backup service as a workload measure. To make the backup policy less intrusive we took the decision of implementing an incremental backup policy. Figure 1 presents the UML activity diagram of replicating the metadata catalogue to other nodes in the data grid.

The process starts with the evaluation of the conditions for copy which assesses the workload of the node running the backup process (the central node) by measuring the CPU and memory in use. If the conditions are met, the process proceeds with the evaluation of the type of backup that must be done. If the conditions are not met, a log is registered (to notify the administrator) and a new backup process is scheduled. If the process is running for the first time, a full backup of the catalogue is made, including the schema of the database. The actual policy determines that the replicas

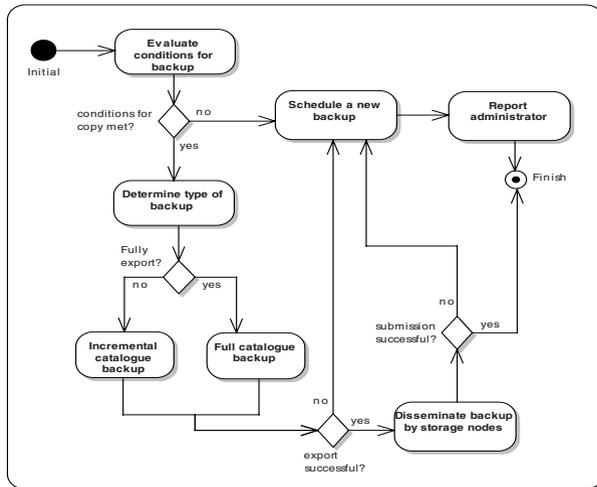


Fig. 1. High-level activity diagram of the Backup service

³ <http://www.shaman-ip.eu>

⁴ <http://grito.intraneia.com>

are disseminated through all of the nodes in the system using built-in replication features. Along with those replicas, a list containing the physical path and hostname of each of the replicas is created in a folder in the host file system, therefore outside the grid virtual system. The reason for this is that, in case a disaster affects the metadata catalogue, the metadata elements with the location of the recovery files will not be lost. So, by creating this list with the recovery files, we can bypass the catalogue and recover from the disaster.

We could also replicate the recovery files outside the grid virtual system, but using the grid to store the files we can use the replication mechanisms it provides and so ensure the preservation of the backups by applying to them to the same preservation policies as for the other preserved objects (including the audit processes).

3 The Recovery Service

The recovery service is illustrated in Figure 2. After setting up a new central node, the administrator can initiate the recovery process by specifying the location (full path) of the backup lists and the hostname of one of the nodes containing the backup files.

If the backup files do not exist (they might had been affected by the failure or disaster), the process fails and the administrator has the chance to give another location path or host. If the backup files exist, the service accesses the list and retrieves the most recent full backup replica, parsing the document from the beginning to the end. When it finds the most recent backup file, it tries to access its first listed replica. If none of the replicas of the full backup is accessible, the process fails. Otherwise, the process continues with the recovery until it concludes.

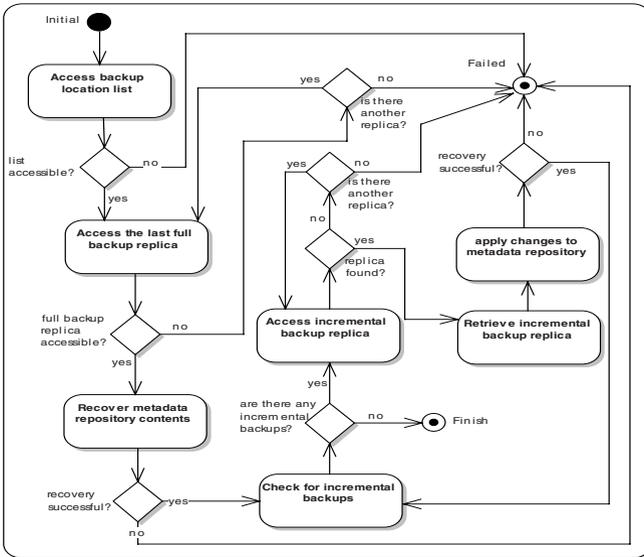


Fig. 2. High-level activity diagram of the Recovery service

After this full recovery, the process checks for incremental backups. As the incremental backups stand for the history of the changes applied to the records in the catalogue, the recovery process applies the changes to the catalogue in the correct order until it reaches the state prior to the failure.

4 Conclusion

In this poster, we present a set of services with the objective of better adapting the iRODS data grid system to preservation purposes. The developed services may present a conceptual solution to some of the problems that may arise from the use of a data grid for preservation, independently of its implementation details, as long as the grid makes use of a metadata catalogue.

We implemented the proposed services using the extensibility properties of the iRODS system, as it allows the addition of new functionalities through the development of micro-services, which are small and well-defined functions/procedures (written in the C language), to execute a determined micro-level task. Users and administrators can chain those micro-services in order to create Actions (macro-level functionalities), also named Rules. The two services were implemented through the definition of Rules.

Future work will consist in testing and refining the developed services so that the gap that separates data grids from digital preservation is filled.

References

1. CCSDS: Reference Model for an Open Archival Information System (OAIS) - Blue Book. National Aeronautics and Space Administration (2002)
2. RSB. Dam safety regulation, decreto-lei (344), Diário da República, Lisbon (October 15, 2007) (in Portuguese)
3. Hedges, M., Hasan, A., Blanke, T.: Management and preservation of research data with iRODS. In: Proceedings of the ACM First Workshop on Cyberinfrastructure: information Management in eScience, CIMS 2007, Lisbon, Portugal, November 09, 2007, pp. 17–22. ACM, New York (2007)
4. Barateiro, J., Antunes, G., Cabral, M., Borbinha, J., Rodrigues, R.: Using a Grid for Digital Preservation. In: Buchanan, G., Masoodian, M., Cunningham, S.J. (eds.) ICADL 2008. LNCS, vol. 5362, pp. 225–235. Springer, Heidelberg (2008)