Kindura: Repository services for researchers based on hybrid clouds

Simon Waddington, Jun Zhang, Gareth Knight, Mark Hedges, Jens Jensen, Roger Downing[1]

Abstract

The paper describes the investigations and outcomes of the JISC-funded Kindura project, which is piloting the use of hybrid cloud infrastructure to provide repository-focused services to researchers. The hybrid cloud services integrate external commercial cloud services with internal IT infrastructure, which has been adapted to provide cloud-like interfaces. The system provides services to manage and process research outputs, primarily focusing on research data. These services include both repository services, based on use of the Fedora Commons repository, as well as common services such as preservation operations that are provided by cloud compute services. Kindura is piloting the use of the DuraCloud[2], open source software developed by DuraSpace, to provide a common interface to interact with cloud storage and compute providers. A storage broker integrates with DuraCloud to optimise the usage of available resources, taking into account such factors as cost, reliability, security and performance. The development is focused on the requirements of target groups of researchers.

Keywords: repositories; cloud infrastructure; digital preservation; storage broker, DuraCloud, iRODS.

Introduction

Preservation of research data has become the subject of intense interest for research administrators as both research funders and journals are increasingly mandating the long-term retention of data to support research findings. In the UK, a number of research councils such as EPSRC and NERC published policies in 2011 that require the preservation of research data for a minimum period of 10 years as a condition of funding (EPSRC Policy Framework on Research Data 2011; NERC Data Policy 2011).

In a broader context, the advantages of effective solutions for managing research outputs are widely recognised. They ensure that researchers' work can be found, made available and correctly attributed to its creators. Availability of data means that published results can be reproduced, and the data can be re-used in new research. For instance, climate data collected by Victorians are used today in climate modelling, alongside current data. Further, such metrological data can never be regenerated. Despite this, management of research outputs is often not the top priority for researchers, and is addressed using various ad hoc, fragmented and non-optimal approaches. Research outputs are typically stored on researcher PCs, portable drives or small dedicated server facilities within their departments. There is a considerable administrative overhead for researchers in maintaining such facilities, since they are typically not supported by the central IT organisation.

---

[1] Simon Waddington, Jun Zhang, Gareth Knight and Mark Hedges are based at the Centre for e-Research, King's College London, UK. Jens Jensen and Roger Downing are based at the Science and Facilities Research Council, UK.

[2] DuraCloud, http://www.duracloud.org/.

Little provision is made for backup or disaster recovery, placing at risk valuable data. Data is often not annotated in a systematic way, making interpretation and reuse by other researchers difficult.

At a time of stringent budget controls, IT departments of higher education institutions are increasingly stretched. Often, research outputs will need to be retained far beyond the funding period of the research project. The implication is that the institutions will, in many cases, be required to underwrite the long term preservation of research outputs.

Deploying new storage infrastructure often requires months of planning and procurement. At the same time, there are often pockets of underused storage at institutions that could potentially be reused for other purposes.

Commercial and private cloud infrastructure provide a potential alternative to in house storage that can reduce deployment times as well as making use of the economies of scale of large data centres. There is already considerable use of cloud resources within research communities. Typically this is funded by individuals paying for services with personal credit cards. This makes accounting for IT usage difficult and does not exploit economies of scale in using centrally procured resources. In addition, Service Level Agreements (SLAs) offered by commercial providers do not provide the required service availability and security guarantees for certain types of data.

Related work

Due to the increasing demands on institutions to provide open access to research outputs, there has been considerable interest in using cloud resources to provide back-end storage for digital repositories. A meeting organised by EduServe and JISC in London in 2010 entitled "Repositories and the Cloud"[3] explored some of these issues. On the one hand the cloud offers rapid provisioning and transparent costing. However, there are many issues such as security, quality of service and trust to be overcome. The trust issues surrounding use of cloud for storage of data, particularly sensitive personal data are explored in Ko et al. (2011). Fujitsu Research Institute (2010) performed a customer survey investigating perceptions of storing personal data in the cloud, which illustrates the considerable concerns of end users.

Current work on cloud computing has been influenced by grid computing. The iRODS[4] grid-based storage system is gaining wide acceptance in data-intensive research to provide back-end storage for research data. Integration of the Fedora Commons repository with iRODS has already been investigated by a number of authors including Pcolar et al. (2011), and Ashenbrenner and Zhu (2009). Preservation and management of research data using iRODS was investigated in Hedges et al. (2009).

DuraSpace are expected to launch a commercial cloud storage service[5] in 2011. The service is based on the DuraCloud middleware, which itself runs in the cloud. Institutions can replicate existing repository collections to the cloud, for example for backup purposes. In contrast to our approach, there is no brokerage between internal and external storage so institutions will need to manually appraise their content to determine if data is to be replicated to the cloud.

---

[3] Repositories and the Cloud, http://www.jisc.ac.uk/events/2010/02/repcloud.aspx.
[4] iRODS: Data Grids, Digital Libraries, Persistent Archives, and Real-time Data Systems. http://www.irods.org.
[5] DuraCloud solutions, http://www.duracloud.org/solutions.

Ruiz-Alvarez and Humphrey (2011) describe an XML schema for characterising cloud storage services in order to enable selection of an appropriate provider based on criteria entered by the user.

Objectives

Kindura pilots the use of hybrid cloud storage infrastructure to provide repository services to researchers. For internal storage, the project is deploying the iRODS storage system. The project is developing management middleware to provide improved and efficient storage management. This approach also enables generation of business intelligence information by provisioning of a centralised service for allocation of storage resources, both in-house and external. Making use of cloud storage via a hybrid cloud approach enables elastic extension of in-house storage resources.

We provide access to data preservation and curation tools that enable ingest of data at the scale of multiple terabytes, making use of elastic cloud computing resources and storage provisioning. The project is focusing on the preservation of research outputs, which include both documents and datasets, although the overall approach is suitable for other data types such as administrative data.

By providing a single point of access to multiple storage services, data security is increased by providing automated data replication across multiple providers. Such providers can be in-house and external, and spread across multiple geographic locations.

We also aimed to demonstrate cost savings in combining internal and external storage resources more efficiently.

User requirements

The Kindura project has worked with researchers in a number of disciplines requiring processing of data in order to understand their data preservation requirements. These include researchers in environmental science, financial mathematics, humanities and biomedical sciences. Research data is typically either numerical, image or textual information, and may comprise single large datasets or a large number of small datasets.

When considering a transition from local to centrally managed repository storage, researchers were particularly concerned about retaining a degree of control of their data. Thus researchers should be able to ensure, for instance, that particularly valuable data such as source datasets that cannot be regenerated are adequately backed up. There was particular concern about the costs and payment models for longer term storage. Understanding the contractual and legal contracts (SLAs) offered by commercial cloud providers is a concern for researchers without access to specialist advice. Hence faculties such as biomedical sciences that are handling sensitive personal data are often unwilling to make use of external storage providers due to the potential risks of unintentional disclosure and breaches of EU data protection legislation. Research groups handling large volumes of data were concerned both about the network transfer speeds in storing and retrieving data across the internet as well as the data transfer charges imposed by commercial cloud providers.

Architecture

The architecture of the Kindura system is illustrated in Figure 1. The Kindura pilot system provides a back-end storage system to the Fedora Commons[6] repository system. At the top level, a User

[6] Fedora Commons, http://www.fedora-commons.org.

Interface is provided to enable the researcher to deposit collections of data as well as searching and retrieving collections from the repository. The interface also allows users to view cost information relating to their data collections as well as information regarding the storage tier at which their collections are held. Storage is classified into tiers ordered by access speeds ranging from low-latency discs to tape stores. One storage provider may offer multiple storage tiers.

The Repository Layer comprises an authentication and authorisation module to enable access by registered users, an Object creation and Ingest module to enable creation of repository objects and an Object store, which is implemented as a Fedora repository instance.

A Data Migration layer is provided to determine where data collections should be stored based on the metadata entered by the user. Data migration covers two activities:

1. Determining storage requirements when content is first uploaded.
2. Periodically reviewing and migrating content to a lower storage tier at appropriate points in the content lifecycle.

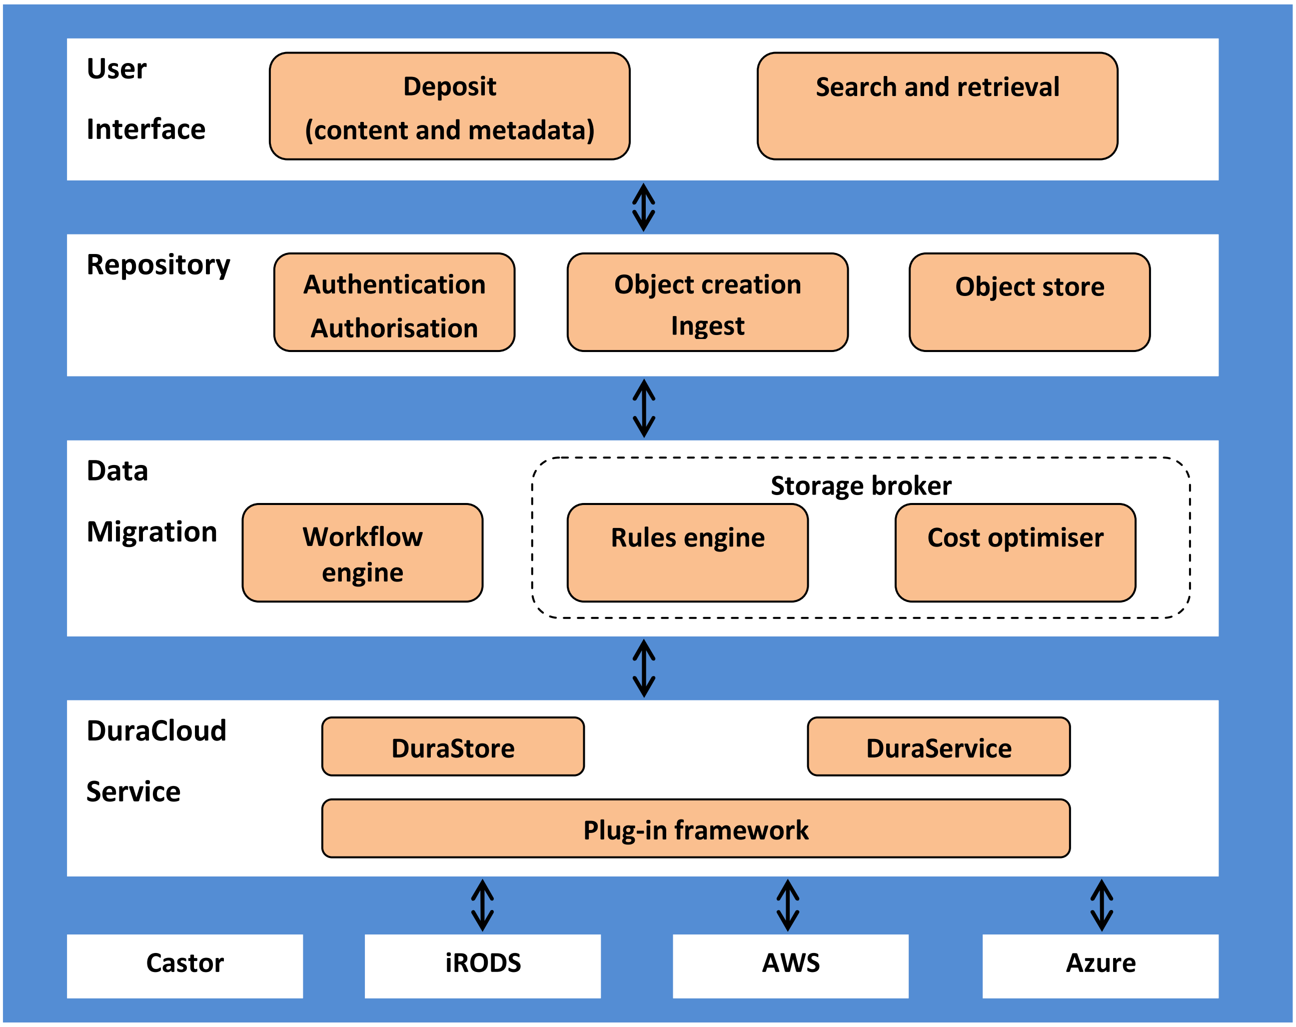

Figure 1: Kindura system architecture

The Ingest and Migration processes are run as workflows, for which we are using Bonita Open Solution[7], an open source Business Process Management application. The Kindura Storage Broker comprises a Rules Engine and Cost Optimiser. The brokerage, ingest and migration processes are discussed in more detail in the following section.

The DuraCloud application provides the DuraStore and DuraService web services. DuraStore provides a REST API to enable interactions with cloud providers through a common interface. A number of cloud storage providers have been integrated with DuraCloud "out-of-the box" such as Amazon S3 and Microsoft Azure. The project is integrating internal storage providers into this framework based on iRODS. The CASTOR tape store provides backup storage for iRODS.

DuraService provides a set of services which can be deployed and used for a variety of purposes, primarily to process the content which has been loaded into DuraCloud storage. This includes content transformation, data replication and bit-integrity checking.

Implementation

Brokerage process

The Kindura brokerage process is illustrated in Figure 2.

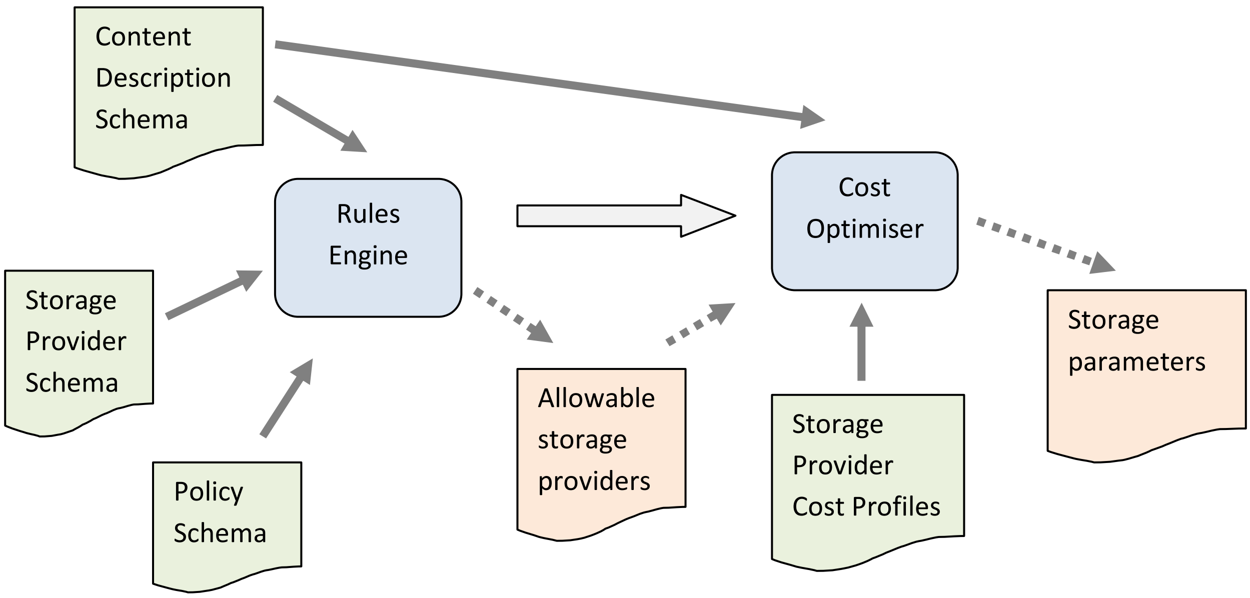

Figure 2: Storage brokerage process

The Rules Engine determines the allowable storage locations and number of replications required for a particular data collection of data. This is based on three types of input. The Content Description Schema is based on descriptive metadata entered by the user as well as automatically extracted metadata such as file sizes, which describe the collection to be stored. The Storage Provider Schema describes the attributes of the individual storage providers, including information such as geographic location and access speed. Individual storage providers can be sub-divided into tiers reflecting

[7] Bonita Open Solution, http://www.bonitasoft.com/.

different storage attributes. The Policy Schema represents institutional policies to be applied when storing data and can be used for performing additional configuration.

The prototype rules engine has been implemented using the Java-based Jess[8] application that uses an enhanced version of the Rete algorithm[9] to process rules. Jess provides a lightweight and efficient tool for implementing storage rules.

Content description schema

In order to assign content items or collections to specific content providers, a metadata schema has been developed to represent content attributes required by the storage broker. These attributes have been classified into a number of different categories:

- *Ownership*: The ownership and usage rights of content collections can restrict where the content is stored and the operations that can be performed on it. For example, financial mathematics researchers at King's make use of commercial stock market data, which is supplied under a licence that does not permit data to be stored outside the institution.
- *Protective marking:* This describes the sensitivity of data, whether the data contains personal data and export classifications. In this category, we also consider the business need for retaining the data. In order to simplify the task of data entry for users, we have defined a universal classification of documents and data into fixed categories.
- *Usage:* The likely or intended usage of the content will influence the choices of storage location. Content intended for frequent and real-time access will need to be stored on low latency storage devices. Network speeds also play a major role in determining the portability of data. Initially users enter the projected access rate of their data collection. After an initial period, this estimate is replaced by the actual access rate generated by monitoring tools.
- *Provenance:* The provenance of the data includes such factors as whether experimental data is a source dataset, an intermediate output of an experiment or a final dataset.

In order to minimise the user workload in creating metadata, we have developed tools that can classify content in certain predefined formats such as data from experiments, as well as capturing the context in which the data was created. Many researchers use a predetermined directory structure to store their experimental data, which is the simplest method that can be applied to extract relevant metadata. This also enables us to use Fedora content modelling to capture the structure of the experiment in the repository. For content modelling, we have followed work done by the Hydra[10] project. Hydra specifies a set of guidelines[11] for creating Fedora objects that narrows down the huge range of possibilities. Experiments are modelled as a collection of Fedora objects in FOXML before being ingested. Metadata is entered at the collection level. A collection may correspond to a set of one or more experiments. A more sophisticated approach that we are

---

[8] Jess Rules Engine, http://www.jessrules.com/jess/index.shtml.
[9] Rete algorithm, http://en.wikipedia.org/wiki/Rete_algorithm.
[10] Hydra project, http://projecthydra.org
[11] Hydra objects, content models and disseminators, https://wiki.duraspace.org/display/hydra/Hydra+objects%2C+content+models+%28cModels%29+and+disseminators.

investigating uses the results of the JISC funded BRIL[12] project, which enable us to capture and annotate files from experiments that are captured in a mapped folder on a user's PC.

Cost optimisation

Commercial cloud storage providers such as Amazon S3 and Microsoft Azure provide explicit pricing information for use of their storage resources. Although there is no standard pricing model, typically providers charge for both for the storage itself, usually per GB, as well as transferring data into and out of the storage. Pricing models are also available for compute resources.

Pricing of internal storage is generally more complex as it depends on the costs of the hardware itself including periodic upgrades and replacements due to failures, staff costs for administration, as well as the maintenance and utility costs of the facilities. The Kindura project has developed a methodology for pricing of internal storage, in order to provide a common baseline for cost comparison. Costing issues for Kindura were discussed in Jensen et al (2011).

The Cost Optimiser uses linear programming techniques to select the most cost effective storage locations for data, based on the user and policy requirements. Information regarding the storage and network transfer costs of the storage providers is captured in the Storage Provider Cost Profile XML schema. It also uses cues such as how long the data is required in a particular resource, and whether it can be deleted when it is no longer required.

Data lifecycle and migration

When moving data from working storage to a repository, appraisal and preservation processes should be applied to ensure the integrity of the data in order to enable retrieval and reuse. The DuraCloud application provides a set of common services for performing such tasks that are being deployed in Kindura.

When data has been moved to repository storage, a periodic reappraisal process is run to determine migration actions such as moving data from disc to tape or removing copies of data from certain cloud providers if their cost has become uncompetitive. The appraisal process uses information generated by usage monitoring to determine content suitable for migration.

An additional factor in the migration process is loss of service from a storage provider. This can happen if a commercial cloud provider ceases trading or there is a catastrophic failure at an internal storage facility such as a fire or flood. In this case, an additional copy of the data can be created in a different storage provider by using the replication services provided by Kindura.

Conclusions

Cloud storage provides an attractive option for increasing the scalability and flexibility of internal institutional repository storage. However, mechanisms are required to ensure legal compliance for sensitive data, data security and replication to ensure access in the event of loss of service or corruption. Kindura has implemented a storage broker to address these issues. The cost implications of using commercial pay-per-use providers also have an impact on the way that IT services are funde. In order to make effective comparisons with the cost of internal storage, we have created a common costing model across internal and external services. In the future, this may encourage

[12] BRIL: Biophysical Repositories in the Lab, http://www.jisc.ac.uk/whatwedo/programmes/inf11/digpres/bril.

institutions to move from “top-slicing” funding model for IT services to one based more heavily on actual usage.

Acknowledgements

We gratefully acknowledge funding for the Kindura project by the Flexible Service Delivery programme of the Joint Information Systems Committee (JISC) in the UK.

We are also grateful to colleagues, at King’s College London and the Science and Technology Facilities Council in helping us to understand requirements for preservation of research data and in evaluating prototypes.

References

Aschenbrenner A., Zhu B. (2009), “iRODS-Fedora Integration,” http://www.irods.org/index.php/Fedora.

EPSRC Policy Framework on Research Data (2011), http://www.legislation.gov.uk/ukpga/2000/36/contents.

Fujitsu Research Institute (2010), “Personal data in the cloud: A global survey of consumer attitudes” http://www.fujitsu.com/downloads/SOL/fai/reports/fujitsu_personaldata-in-the-cloud.pdf.

Hedges M., Hasan A., Blanke T. (2009) “Rule-based curation and preservation of data: A data grid approach using iRODS”, *Future Generation Computer Systems* 25, 4, pp. 446-452, 2009. http://dx.doi.org/10.1016/j.future.2008.10.003.

Jensen J., Downing R., Waddington S., Hedges M., Zhang J., Knight G. (2011) “Kindura - Federating Data Clouds for Archiving”, *Proc. Int'l Symp. on Grids and Clouds, 2011 (Proceedings of Science)*, to appear.

Ko R. K. L., Jagadpramana P., Mowbray M., Pearson S., Kirchberg M., Liang Q., Lee B. (2011) “TrustCloud: A Framework for Accountability and Trust in Cloud Computing”, *2nd IEEE Cloud Forum for Practitioners (IEEE ICFP 2011),* Washington DC, USA, July 7-8, 2011. http://www.hpl.hp.com/techreports/2011/HPL-2011-38.pdf.

NERC Data Policy (2011), http://www.nerc.ac.uk/research/sites/data/policy.asp.

Pcolar D., Davis D. W., Zhu B., Chassanoff A., Hou, C., Marciano R. (2010) “Conceptualizing Policy Driven Repository Interoperability (PoDRI)” (2010). In *Proceedings IRODS User Group Meeting 2010*, eds Moore R. W., Rajasekar A., Marciano R., Data Intensive Cyberinfrastructure Foundation, pp25-31.

Ruiz-Alvarez A., Humphrey M. (2011) “An Automated Approach to Cloud Storage Service Selection”. In *Proceedings of the 2nd Workshop on Scientific Cloud Computing (ScienceCloud 2010)*, USA. June 8, 2011.