Journal of Physics: Conference Series

Storage quality-of-service in cloud-based scientific environments: a standardization approach

Storage quality-of-service in cloud-based scientific environments: a standardization approach

Paul Millar[1], Patrick Fuhrmann[1], Marcus Hardt[2], Benjamin Ertl[2], Maciej Brzezniak[3]

[1]Deutsches Elektronen-Synchrotron DESY, IT - Scientific Computing, Notkestrasse 85, 22607 Hamburg, Germany
[2]Karlsruher Institut for Technology, Steinbuch Centre for Computing (SCC), Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany
[3]Poznan Supercomputing and Networking Cente, Noskowskiego 1, 61-704 Poznan, Poland

E-mail: `paul.millar@desy.de, patrick.fuhrmann@desy.de, hardt@kit.edu, benjamin.ertl@kit.edu, maciekb@man.poznan.pl`

Abstract. When preparing the Data Management Plan for larger scientific endeavors, PIs have to balance between the most appropriate qualities of storage space along the line of the planned data life-cycle, its price and the available funding. Storage properties can be the media type, implicitly determining access latency and durability of stored data, the number and locality of replicas, as well as available access protocols or authentication mechanisms. Negotiations between the scientific community and the responsible infrastructures generally happen upfront, where the amount of storage space, media types, like: disk, tape and SSD and the foreseeable data life-cycles are negotiated. With the introduction of cloud management platforms, both in computing and storage, resources can be brokered to achieve the best price per unit of a given quality. However, in order to allow the platform orchestrator to programmatically negotiate the most appropriate resources, a standard vocabulary for different properties of resources and a commonly agreed protocol to communicate those, has to be available. In order to agree on a basic vocabulary for storage space properties, the storage infrastructure group in INDIGO-DataCloud together with INDIGO-associated and external scientific groups, created a working group under the umbrella of the Research Data Alliance (RDA). As communication protocol, to query and negotiate storage qualities, the Cloud Data Management Interface (CDMI) has been selected. Necessary extensions to CDMI are defined in regular meetings between INDIGO and the Storage Network Industry Association (SNIA). Furthermore, INDIGO is contributing to the SNIA CDMI reference implementation as the basis for interfacing the various storage systems in INDIGO to the agreed protocol and to provide an official Open-Source skeleton for systems not being maintained by INDIGO partners.

1. Introduction

Storage Quality-of-Service and the associated data life-cycle for different storage solutions take on greater significance while preparing the Data Management Plan for larger scientific endeavors. Especially the expenses for different storage solutions become more important with the increasing amount of data. Large amounts of data have to be stored cost-effectively and in most cases for a time frame of 10 or more years. On the other hand the data should be accessible at low latencies for simulations and processing. Current existing storage solutions cover most of the requirements to store scientific research data, in particular with the growing amount of cloud-based storage solutions. However, there is a lack of common definitions of Quality-of-Service (QoS) and data lifecycle (DLC) terms and standardized communication and management protocols.

This paper presents the ongoing efforts to standardize QoS and DLC terms and definitions and the development of standardized approaches for QoS and DLC management in the INDIGO-DataCloud project. The INDIGO-DataCloud project is a European research project funded by the European Commissions Horizon 2020 Framework Program. The standardization efforts for QoS and DLC terms and definitions are conducted in collaboration with the Research Data Alliance (RDA). The development of standardized communication and management protocols is performed in close collaboration with the Storage Network Industry Association (SNIA) and based on the Cloud Data Management Interface (CDMI) specification.

The reminder of the paper is organized as follows: Section 2. provides an introduction to Quality-of-Service and data life-cycle management and the SNIA Cloud Data Management Interface. Section 3. describes the implementation and standardization efforts by the INDIGO-DataCloud project. At the end in Section 4. the conclusions and outlooks are provided.

2. Quality-of-Service and Data Life-Cycle Management

Standardized QoS and DLC terms and definitions can not only simplify the preparation of scientific data management plans, but also used for better comparison of different storage solutions from different storage providers. With a better comparability costs can be saved both on provider sides and consumer side. As an example, users might agree that their data is stored on more cost-effective tape drives provided that the data can be retrieved on-demand at a requested access latency. Moreover, standardized QoS and DLC terms and definitions are essential for automated processes, for example brokering of storage between autonomous systems.

The INDIGO-DataCloud project addresses these requirements by defining a common vocabulary for QoS and DLC attributes and their values, based on existing use-cases from different scientific communities. Another goal of the project is to analyze existing protocols and standards and to extend them if needed, to allow a standardized approach for QoS and DLC management on a rang of various storage solutions.

2.1. QoS and DLC Terms and Definitions

One of the biggest problems for standardized QoS and DLC terms and definitions is the ambiguity of attributes and their values to describe storage solutions. Table 1 illustrates this issue using the example of four QoS attributes, access latency, durability, data rate and costs. Each attribute by itself and also its respective values require a more precise definition to compare them across different storage solutions. An abstract description as seen in Table 1 allows a relative comparison of the attributes, however new storage media or combinations of storage media can not be added easily to the comparison. Also automated processing and comparison of these attributes can not be performed in a straightforward manner.

Table 1. Comparison of different QoS attributes for different storage media

	Tape	HDD	SSD	HDD & HDD	Tape & HDD
Access latency	high	medium	low	medium	medium
Durability	ok	medium	unclear	acceptable	ok
Data rate	ok	ok	medium	ok	ok
Costs	very low	acceptable	very high	medium	medium

One goal of the INDIGO-DataCloud project is therefore to develop a common vocabulary for

IOP Conf. Series: Journal of Physics: Conf. Series **898** (2017) 062043 doi:10.1088/1742-6596/898/6/062043

QoS and DLC terms and definitions in collaboration with partners from industry and academia. For this reason a working group has been established as part of the Research Data Alliance (RDA).

In general, data life-cycle management can be seen as changes of QoS attributes over time. Figure 1 shows a typical data life-cycle of research data with possible administrative interactions and associated QoS attribute changes.

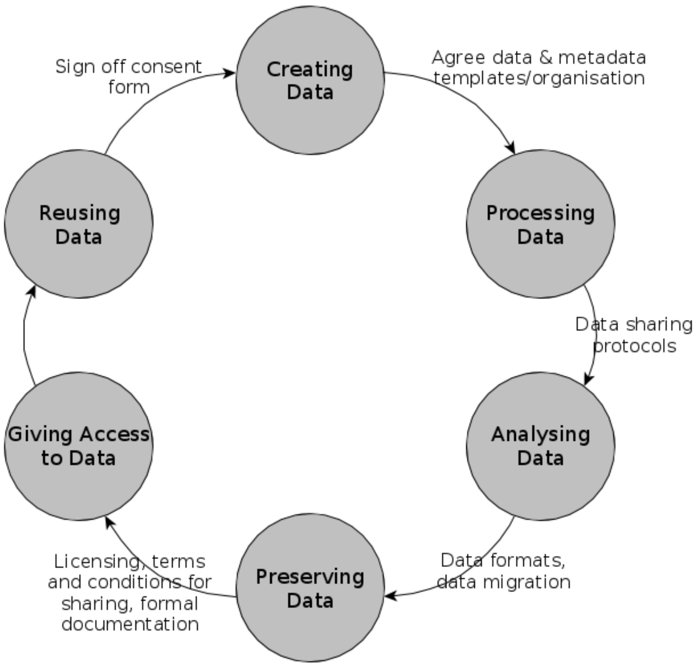

Figure 1. Data Life-Cycle management example [1]

As an example, the interactions illustrated in Figure 1 could trigger following changes of the data's QoS attributes.

- Change of access latency by migrating data from disk to tape
- Change of number of copies for archiving data
- Change of ownership and access rights for data sharing and data reuse
- Change of cost model e.g. pay-as-you-go, pay-in-advance

Because of the close relation between DLC management and QoS attribute changes, the focus of the presented approach towards standardization is generalized on QoS attribute changes.

2.2. Protocols and Standards for QoS and DLC Management

Essential considerations towards a standardized protocol for QoS and DLC management have been introduced by several partners in the grid computing environment with the Storage Resource Manager Interface (SRM) specification [2] and its implementations called Storage Resource Managers (SRMs). The SRM specification aims to abstract and unify different storage storage solutions and technologies via a standardized interface.

Based on the initial SRM considerations, the Cloud Data Management Interface (CDMI) specification [3] provides additional functionality to integrate cloud-based storage solutions from the perspective of industrial storage providers.

CDMI has been specified from partners of the Storage Networking Industry Association (SNIA) to provide uniform and standardized storage in cloud-based environments, independent of the underlying storage technology. The interface allows further to administer and manage information related to stored data, for example concerning data security, accounting or QoS and DLC management.

The Cloud Data Management Interface can be used to create, retrieve, update and delete data objects in cloud environments, as well as to explore available capabilities of connected storage systems. The interface is HTTP REST based [4] and offers therefore high compatibility with a great number of client and server applications.

The CDMI specification in its current version 1.1.1 [5] allows users to query container and data objects for their associated QoS and DLC attributes and also to request changes of specific attributes. All attributes are part of the container and data objects' related meta-data.

Currently an extension to the CDMI specification is under development, the Capabilities Selection Extension [6], that allows users to request changes of groups of QoS and DLC attributes, so called capabilities. These groups of attributes allow storage providers to specify specific QoS profiles that can be applied to data objects.

Based on the current CDMI specification and the work in progress capabilities extension, the INDIGO-DataCloud project has specified and implemented CDMI for QoS and DLC management [7], integrating a variety of different storage solutions.

3. Design and Implementation

INDIGO-DataCloud utilizes CDMI to manage QoS and DLC related attributes grouped into specified capabilities for data objects. This control path for the data's meta-data is independent from the data path (upload/download). The data path is established via common data transfer protocols as for example WebDAV, SFTP or GridFTP. For supported storage technologies INDIGO focuses on popular and widely used solutions in scientific communities like dCache, Ceph, GPFS, Gemss+TMS, StoRM and HPSS.

A subset of technical attributes have been chosen as supported QoS attributes by all storage providers. These attributes are grouped into capabilities and exposed to the user as groups abstracting the technical aspect of the attributes, while maintaining a per attribute comparison between different storage solutions.

For the first release of the INDIGO-DataCloud project three important QoS attributes have been selected by the involved partners. For the second release more attributes will be added and discussed in close collaboration with the dedicated QoS terms and definitions RDA working group.

- **Data redundancy** - number of complete copies
- **Geographic placement** - restrictions on the geographic regions where the object is permitted to be stored
- **Latency** - desired maximum time to first byte, in milliseconds

Based on this selection a multitude of QoS and DLC related management tasks can already be achieved, for example staging and archiving of data. Table 2 shows a possible definition of the respective capabilities classes.

As an example, a requests of change for a data object from the capabilities class `TapeOnly` to `DiskAndTape` as specified in Table 2 would correspond to a staging process, where data is temporarily moved to storage media with fast data access, for example for more efficient data processing tasks.

A reverse change would correspond to an archiving process, where data is moved from more expensive storage media to more cost-effective storage media, typically from SSD or HDD to Tape.

IOP Conf. Series: Journal of Physics: Conf. Series **898** (2017) 062043 doi:10.1088/1742-6596/898/6/062043

Table 2. CDMI notation for example Capabilities-Classes

CDMI Capabilities-Class	CDMI Capabilities-Attribute
`/cdmi_capabilities/dataobject/TapeOnly`	`"cdmi_data_redundancy": 2, "cdmi_geographic_placement": [ "DE" ], "cdmi_latency": 50000, "cdmi_capabilities_allowed": ["/cdmi_capabilities/dataobject/ DiskAndTape"]`
`/cdmi_capabilities/dataobject/DiskAndTape`	`"cdmi_data_redundancy": 3, "cdmi_geographic_placement": [ "DE" ], "cdmi_latency": 0, "cdmi_capabilities_allowed": ["/cdmi_capabilities/dataobject/ TapeOnly"]`

Theses two classes and operations are exemplified by the INDIGO-DataCloud CDMI-QoS Management Interface implementation, which has been released in June 2016. Additional storage back-ends will be integrated towards the second release end of March 2017.

The implementation is divided in two components, a common server component and a respective storage back-end component as described in the following sections.

3.1. CDMI-QoS Server

The implementation of the CDMI-QoS server component is based on the CDMI reference implementation provided by SNIA and written in Java. The reference implementation has been ported to a Spring Boot application for an improved deployment and extended functionality.

Information about storage resources can be retrieved via CDMI by issuing the specified HTTP REST requests to the server. The server looks up the corresponding meta-data and if necessary queries the appropriate back-end module to retrieve updated information from the storage system directly. The meta-data are kept in a Redis NoSQL database for caching and persistence across different server instances.

The server component is fully integrated within the INDIGO-DataCloud Authentication and Authorization Infrastructure (AAI) based on OpenId Connect. User and client information are passed on directly to the respective back-end module in the context of the Java Authentication and Authorization Service (JASS).

3.2. CDMI-QoS Backend Module

The selection of initial supported storage back-ends is based on the use-case requirements from scientific partners in the INDIGO project and includes dCache, Ceph, GPFS, Gemss+TMS, StoRM and HPSS. Each storage back-end can be integrated with the CDMI-QoS server via its own back-end module, where all back-end modules follow a common interface definition, abstracting the storage system specific interactions for the server component.

The implementation of the back-end modules is exemplified with the implementation of the HPSS back-end module in the following section.

3.3. CDMI-QoS for HPSS

Figure 2 outlines the example architecture of QoS and DLC management via CDMI in combination with the High Performance Storage System (HPSS). HPSS is a flexible, scalable and policy-based hierarchical Storage Management System by IBM [8] and is provided by the Karlsruher Institute of Technology (KIT) as a partner in the INDIGO project.

CDMI is utilized for the control-path supporting meta-data queries and requests for QoS and DLC attribute changes. The server has been configured to use the HPSS back-end module that itself communicates via a HTTP REST interface with the HPSS storage system.

The data-path for upload and download of data is separated from the control-path and realized via either WebDAV, SFTP or GridFTP in this example. The HPSS system is exposed via NFS and mounted to the access machine as illustrated in Figure 2 allowing direct I/O with the HPSS system.

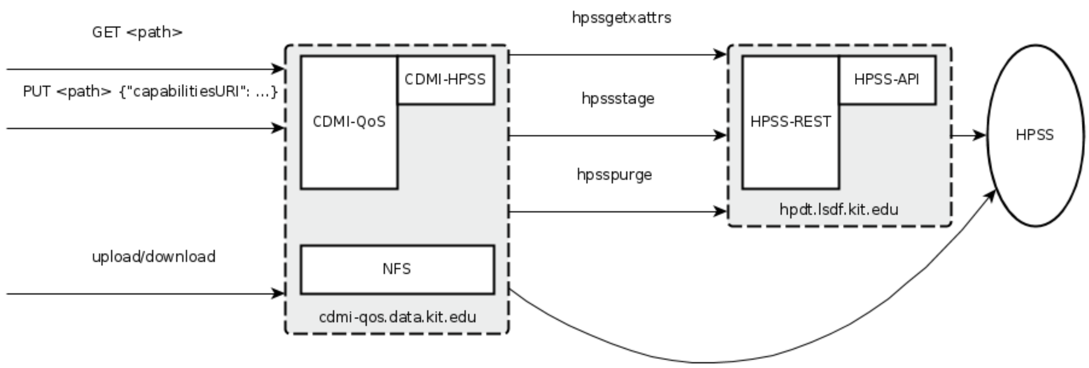

Figure 2. CDMI-QoS example architecture for the HPSS Storage-System

The presented staging and archiving operations can be triggered by a HTTP PUT request containing the respective CDMI capabilities-class as a reference. Listing 1 exemplifies the QoS change request from capabilities-class `DiskAndTape` to capabilities-class `TapeOnly`.

```
PUT /myObject HTTP/1.1
Host: cdmi-qos.data.kit.edu
Content-Type: application/cdmi-object
X-CDMI-Specification-Version: 1.1

{
  "capabilitiesURI":
    "/cdmi_capabilities/dataobject/TapeOnly"
}
```

Listing 1. HTTP PUT Request example to change capabilities-classes

The request outlined in Listing 1 is received from the CDMI-QoS server and passed on to the integrated HPSS back-end module. The HPSS back-end module in turn interacts via HTTP REST with a HPSS proxy server that can directly invoke HPSS API calls. The implemented stating and archiving operations are executed asynchronous, allowing the user to query the result of the operations at a later point in time.

4. Conclusion and Future Work

In this article we presented the ongoing work towards a standardized Quality-of-Service and data life-cycle management vocabulary as well as the design and implementation of a CDMI-based QoS management interface as part of the INDIGO-DataCloud projects.

With a subset of clearly defined QoS attributes and a common REST interface, the presented implementation allowed QoS and DLC management across a multitude of different storage technologies.

This has been exemplified with the staging and archiving operations available within the first release of the INDIGO-DataCloud CDMI-QoS implementation. These operations and their respective Quality-of-Services have been described with selected attributes by the involved partner.

Towards the second release end of March 2017 and beyond the duration of the project the participating partners are working on the integration of additional storage systems and the standardization of QoS and DLC terms and definitions that enable storage systems to support further common QoS and DLC operations. Additional attributes will be added and discussed in close collaboration with the dedicated QoS terms and definitions RDA working group.

The standardization approach is moreover the groundwork for automatized discovery and brokering of storage space based on QoS and DLC attributes.

Acknowledgment

The authors want to acknowledge the support of the INDIGO-DataCloud project (grant number 653549) funded by the European Commission's Horizon 2020 Framework Program.

References

[1] Corti L, Van den Eynden V, Bishop L and Woollard M 2014 *Managing and sharing research data: A guide to good practice* (Sage)
[2] Sim A and Shoshani A 2007 The storage resource manager interface specification, version 2.2 *CERN, FNAL, JLAB, LBNL and RAL* (Citeseer)
[3] Slik D 2010 Cdmi and cloud federation *Storage Developer Conference (SDC), SNIA, Santa Clara*
[4] Fielding R T 2000 *Architectural styles and the design of network-based software architectures* Ph.D. thesis University of California, Irvine
[5] SNIA 2015 Snia technical position: Cloud data management interface (cdmi) v1.1.1
[6] SNIA 2016 Snia working draft: Capabilities selection extension version 1.1c
[7] INDIGO 2016 Indigo working draft: Qos using cdmi version 0.2
[8] IBM 2016 Hpss high performance storage system URL `http://www.hpss-collaboration.org/`