



Contribution ID: 37

Type: **not specified**

Quality of Service in storage and the INDIGO-DataCloud project.

Thursday 15 October 2015 09:20 (20 minutes)

The pressure to provide cheap, reliable and unlimited cloud storage space in the commercial area has provided science with affordable storage hardware and open source storage solutions with low maintenance costs and tuneable performance and durability properties, resulting in different cost models per storage unit. Those models, already introduced by WLCG a decade ago (disk vs tape) are now offered by large companies like Amazon (Bock Storage, S3 and Glacier) or Google (Standard, Durable Reduced, Cloud Storage Nearline). Industry appliances or software stacks (e.g. HPSS) offer similar storage properties for your locally installed storage.

However, other than with SRM for WLCG, those offered storage quality properties don't follow a common description or specification and are hard to compare programatically. Moreover they aren't even close to a common way of been negotiated between the requesting client and the providing storage technology, which would be a prerequisite for federating different public and private storage services.

To fill this gap, the INDIGO-DataCloud project is proposing a process to agree on common semantics in describing QoS attributes in storage in a consistent way, independently of the used API or protocol.

The process involves gathering uses-cases from scientific communities and creating working groups in international organisations, like RDA, OGF and SNIA to further discuss possible solutions with other interested parties like EUDAT and EGI.

In a second step, based on feedback received, INDIGO will propose an implementation of the defined semantics to steer quality of service in storage as an extension to an existing industry standard, e.g. CDMI. As a proof of concept INDIGO will implement the proposed solution in storage systems used within the INDIGO project, like dCache, StoRM and some typical industry products, as a reference for other systems.

We are presenting our work at HEPiX, in order to receive feedback from HEP communities on our plans and because we regard the outcome of our work as beneficial for HEP in general and for WLCG experiments in particular, as with the proposed end of the usage of SRM in WLCG, experiments are again left alone with the necessary steering of data quality of service attributes (e.g. disk, tape or both) at the different storage endpoints.

This presentation will report on the goals achieved so far and on our next steps.

Length of presentation (max. 20 minutes)

20

Authors: DONVITO, Giacinto (INFN-Bari); HARDT, Marcus (Karlsruhe Institute of Technology); FUHRMANN, Patrick (Deutsches Elektronen-Synchrotron Hamburg and Zeuthen (DE)); MILLAR, Paul

Presenter: FUHRMANN, Patrick (Deutsches Elektronen-Synchrotron Hamburg and Zeuthen (DE))

Session Classification: Storage and Filesystems

Track Classification: Storage & Filesystems