

Dynamic federations: storage aggregation using open tools and protocols

Tuesday, September 18, 2012 11:20 AM (20 minutes)

Description of the work

The Dynamic Federations project considers HTTP/WebDAV and NFS4.1-based storage elements and makes them able to cooperate through an architecture that properly feeds the redirection mechanisms that they are based upon, thus giving the functionalities of a “loosely coupled” storage federation.

One of the key features is to use standard clients (provided by OS'es or open source distributions, e.g. Web browsers) to access an already aggregated system; this approach is quite different from aggregating the repositories at the client side through some wrapper API, like for instance GFAL, or by developing new custom clients.

The challenge, here undertaken by the providers of dCache and DPM, but pragmatically open to other Grid and Cloud storage solutions, is to build such a dynamic system while being able to accommodate name translations from existing catalogues (e.g. LFCs), experiment-based metadata catalogues, or stateless algorithmic name translations, also known as “trivial file catalogues”.

Other technical challenges are latency and scalability, and the ability to create worldwide storage federations that are able to redirect clients to repositories that they can efficiently access, for instance trying to choose the endpoints that are closer or applying other criteria.

The project is in an advanced demoable status, and the inclusion in the EMI distribution will be one of the next short term steps.

Wider impact of this work

The features of a loosely coupled, high-performance federation of open-protocols-based storage elements will open many possibilities of evolving the current computing models without disrupting them, and, at the same time, will be able to operate with the existing infrastructures, follow their evolution path and add storage centers that can be acquired as a third-party service.

Printable Summary

A number of storage elements and catalogues now offer standard protocol interfaces like NFS 4.1/pNFS and WebDAV, for access to their data repositories, in line with the standardization effort of the European Middleware Initiative (EMI). Here we report on work that exploits the federation potential of these protocols, making it possible to build a system that offers a unique view of the storage and metadata ensemble and the possibility of integration of other compatible resources such as those from cloud providers.

Such so-called storage federations of standard protocols-based storage elements give a high performance unique view of their content that does not need to index the content of all the storage, thus promoting simplicity in accessing the data they contain and offering new possibilities for resilience and data placement strategies.

Primary authors: Mr DEVRESSE, Adrien (CERN); Mr ALVAREZ AYLON, Alejandro (CERN); FURANO, Fabrizio (CERN); KEEBLE, Oliver (CERN); FUHRMANN, Patrick (DESY); ROCHA, Ricardo (CERN)

Presenter: FURANO, Fabrizio (CERN)

Session Classification: New middleware products

Track Classification: Resource Infrastructure services (Peter Solagna: track leader)