

Onedata global data storage

Wednesday, 21 May 2014 16:20 (20 minutes)

Increasing number of powerful computing environments causes more difficult problems than increase of overall data volume. Such aspects as variety of data and the processing speed required to use the large amount of information have to be addressed. Different requirements of multiple groups of users make it necessary to provide different storage systems and services that process data. It is usually not possible to provide it all in one site so many users have to process their data in multiple data centers. However, data management in such distributed environment is too complicated for many of them.

The authors present a novel solution, called Onedata, that simplifies users work in organizationally distributed environments by provision of a uniform and coherent view on all data stored at the storage systems in all data centers where user work. It also supports work in groups, data sharing/publication and serves data efficiently. However, simplification of the system from the user point of view results in increase of number and difficulty of management tasks that have to be done by administrators or automations. Hence, Onedata provides functionalities that also simplify administrators work: automatic rule-based data management, the infrastructure state monitoring data gathering and visualizing and data protection from unauthorized access.

This paper describes Onedata in terms of its architecture, current implementation status and presents exemplary use cases.

Wider impact and conclusions

The increase of availability of computing environments results in increase of the number of less technically advanced users. The tasks of user may be executed in one or many sites depending on availability of needed storage solutions and services. However, data management in distributed environment is too difficult for less technically advanced users. They expect that data access will be simple using one tool, preferably based on standard POSIX, even when many sites are involved.

Onedata simplifies data management and provides useful functionalities such as support for work in groups and data publication. Installation of Onedata in data and computing centers should not only simplify work of current users but also attract new ones. Hence, Onedata also provides functionalities that simplify administrators work to help them to cope with the growing number of users.

Description of work

When Onedata is installed in a site, it virtualizes logical namespace used to access data and coordinates data access of processes that are executed at all worker nodes. The processes access data through virtual file system implemented with FUSE, which translates a logical name to actual data location at a storage system. To do the translation, the filesystem communicates with the Data Management component. In order to provide high-performance, filesystem always tries to operate on the data locally, as in many cases worker nodes are connected within a site with a shared storage. In order to provide high-throughput the Data Management component is deployed at a cluster and based on efficient, scalable technologies (Erlang, NoSQL) that allow handling of large number of requests simultaneously. To provide load balancing and high availability inside the Data Management component, an advanced method of requests routing, that includes control over the DNS, was designed and implemented.

Onedata supports also data management from the outside of the site. It provides packages that allow installation of virtual file system at user PC and the Web-based GUI. Additionally, a fully functional REST API allows direct interfacing from third party applications.

Onedata instances installed in many sites are able to cooperate on the basis of administrator-defined rules so the user does not see any barriers. When the process in one center needs data located in another, Data Management components of both sites cooperate to provide the data as efficiently as possible, e.g., the data may be copied using many hosts and many channels at the same time. If the administrators of all sites agree on the advanced rules, Onedata is also able to automate complex data management between sites, e.g., the data may be migrated to the site where it is used most frequently.

Primary author: DUTKA, Lukasz (CYFRONET)

Presenter: DUTKA, Lukasz (CYFRONET)

Session Classification: New data management solutions for EGI

Track Classification: Requirements and solutions for data management and computing (Track Leaders: B. Konya, H. Heller, S. Tarkoma)