

The Charon File System: A cloud-of-clouds infrastructure for Biobanks data storage and sharing

Friday, 26 September 2014 09:25 (20 minutes)

During the last years several cloud-of-clouds (or multi-cloud) storage systems have been proposed with the objective of minimizing trust on cloud providers, decreasing costs and improving performance [RACS, DepSky]. Such systems range from archival storage [RACS], object stores [DepSky], key-value stores [SPANStore] and even full-fledged file systems [SCFS].

In this talk we will present Charon, an experimental cloud-of-clouds file system designed to be a storage and sharing infrastructure for integrating European biobanks. Charon is the core component for building federations of Biobanks using the BiobankCloud PaaS, a bioinformatics processing, storage and interconnection platform being developed in the BiobankCloud FP7 project (<http://www.biobankcloud.com>).

Charon's main objective is to build a data-centric (or "servless") infrastructure that enable federated biobanks to share large volumes of data and metadata (related with samples and studies [MIABIS]) respecting security and performance constraints. Furthermore, we want to give authorized bioinformaticians a "dropbox-like" experience when accessing biobanks datasets.

Besides the obvious scalability issues (related with the number and size of files kept in the system), there are other important requirements and design principles that guide the development of Charon:

1. A truly servless design: we want to minimize the operational effort required for maintaining the shared infrastructure by implementing the whole system at the client-side, relying only on widely available cloud services.
2. Dependable metadata storage: all the file system and biobank-specific metadata [MIABIS] must be widely available to authorized users despite possible failures on communication and cloud providers.
3. Flexible data location: due to the legal, performance and criticality constraints, shareable data must be stored either in the edges (file system clients), in a single cloud or even in a cloud-of-clouds.
4. Efficient read/write and read/read sharing: the system must be as efficient as possible when reading data created by others and, at the same time, consistency issues and write-write conflicts must be managed automatically by the infrastructure.

Differently from our previous work on cloud-of-clouds (CoC) file systems [SCFS], which required coordination servers for controlling data sharing, Charon relies only on cloud services such as Amazon S3, without requiring any dedicated process other than the ones running on the file system clients (i.e., biobank servers and bioinformatician desktops). Furthermore, the system manages file metadata and data in a different ways. The former is encapsulated in namespace containers (shared or private, depending on the directory sub-tree visibility) that are stored in the CoC, while file data can be kept at different locations. When a file is created, users can specify if it will be maintained locally, in a single cloud provider or in multiple cloud providers (CoC). An improved version of DepSky [DepSky] is used for ensuring data is stored in an efficient, secure and dependable way in the CoC (by combining erasure codes, secret sharing and Byzantine-quorum replication).

An interesting novel aspect of Charon is its concurrency control algorithm. Although we give up strong consistency for achieving low latency, we still need concurrency control to avoid write-write conflicts on shared files. Such control is provided by a new design for a CoC lease control based on existing cloud services. More specifically, we devised a compositional lease algorithm that uses fail-prone lease objects implemented using appropriate services offered by individual cloud providers. Currently, we implemented efficient lease objects for Windows Azure, Amazon Web Services, Rackspace and Google App Engine. Our approach significantly improve lock latency when compared with other data-centric coordination protocols.

References

[RACS] Hussam Abu-Libdeh, Lonnie Princehouse, and Hakim Weatherspoon. RACS: a case for cloud storage diversity. In Proc. of the 1st ACM symposium on Cloud computing. 2010.

[DepSky] A. Bessani, M. Correia, B. Quaresma, F. Andre, and P. Sousa. DepSky: Dependable and secure storage in cloud-of-clouds. ACM Transactions on Storage, 9(4), 2013.

[SPANStore] Zhe Wu, Michael Butkiewicz, Dorian Perkins, Ethan Katz-Bassett, and Harsha V. Madhyastha. SPANStore: cost-effective geo-replicated storage spanning multiple cloud services. In Proc. of the 24th ACM Symposium on Operating Systems Principles. 2013.

[SCFS] A. Bessani, R. Mendes, T. Oliveira, N. Neves, M. Correia, M. Pasin, and P. Verissimo. SCFS: a shared cloud-backed file system. In Proc. of the 2014 USENIX Annual Technical Conference, 2014.

[MIABIS] Loreana Norlin, Martin N Fransson, Mikael Eriksson, Roxana Merino-Martinez, Maria Anderberg, Sanela Kurtovic, and Jan-Eric Litton. A minimum data set for sharing biobank samples, information, and data: Miabis. Biopreservation and Biobanking, 10(4). 2012.

Primary author: BESSANI, Alysson (University of Lisbon, Faculty of Sciences)

Co-authors: Mr MENDES, Ricardo (University of Lisboa); Mr OLIVEIRA, Tiago (University of Lisboa); Mr COGO, Vinicius (University of Lisboa)

Presenter: BESSANI, Alysson (University of Lisbon, Faculty of Sciences)

Session Classification: EGI-GEANT Symposium: Technology