

A ROOT OF TRUST FOR THE PERSONAL CLOUD

By Nicolas Anciaux, Benjamin André, Philippe Pucheral and Paul Tran-Van

We are witnessing an exponential accumulation of personal data on central servers: data automatically gathered by administrations, companies and web sites but also data produced by the individuals themselves and deliberately stored in the cloud for convenience (e.g., photos, agendas, raw data produced by smart appliances and quantified-self devices). Unfortunately, there are many examples of privacy violations arising from abusive use or attacks, and even the most secured servers are not spared.

The **Personal Cloud** paradigm has recently emerged as a way to allow individuals to manage under their control the collection, usage and sharing of data in different contexts and for different types and sensitivities of data, as requested by the World Economic Forum¹ [1]. Initiatives like Blue Button and Green Button in the US, MiData in Great Britain and MesInfos in France give body to this paradigm by returning personal data retained by companies and administrations to individuals in a practical way. This user-centric vision illustrates the gravity shift of information management from organizations to individuals. However, at the time individuals recover their sovereignty of their data, they also inherit the burden of organizing this personal data space and more importantly of protecting it against all forms of attacks and loss.

While a whole bunch of research work is tackling the organization and semantic exploitation of the user's workspace, much less attention has been paid so far on the security issues. Existing security solutions range from encrypting the personal data with an independent tool (e.g., TrueCrypt or BoxCryptor) before storing it on a cloud provider or relying on decentralized storage (like OwnCloud, SandStorm, SeaFile, Unity or OpenPDS personal cloud platforms to cite a few), letting the individual decide about the place(s) where to store his data. In both cases, data is usually encrypted with keys which in turn are encrypted with the user's password. Hence, the password is de facto the root of security and the Achilles heel of these systems since all data are definitely lost if the user forgets his password or can be leaked if the password is compromised. This is a critical point considering that a personal cloud is expected to concentrate the complete digital estate of an individual. The sharing functionality among users is also very limited, implemented either by manually exchanging links to files or by compiling the access control within the encryption (e.g., by encrypting the file key with the recipient's public key) making access control static and cumbersome to manage.

We promote a different approach, called Secure Personal Cloud, where the root of security is made of secure hardware and the access control policies are easily extracted from the personal cloud itself, with minimal user's interaction. Our solution is based on the tight integration of a Personal Cloud platform (Cozy) with a trusted component (PlugDB). Cozy is an open-source solution developed by Cozy Cloud. It is able to gather personal data from multiple sources and data providers, to organize it in a document database (CouchDB) and to synchronize it with multiple devices of the same user. PlugDB is a secure database engine embedded in a low-cost tamper-resistant hardware device (i.e., a combination of smartcard

¹ The World Economic Forum. Rethinking Personal Data: Strengthening Trust. May 2012.

technology and Flash storage in a USB form factor token) [3]. It is able to store data/metadata associated to Cozy documents, query them, manage access control rules and encrypt/decrypt data [2]. PlugDB is a research prototype from INRIA. Combining these two components allows to set up the architecture depicted in Figure 1.

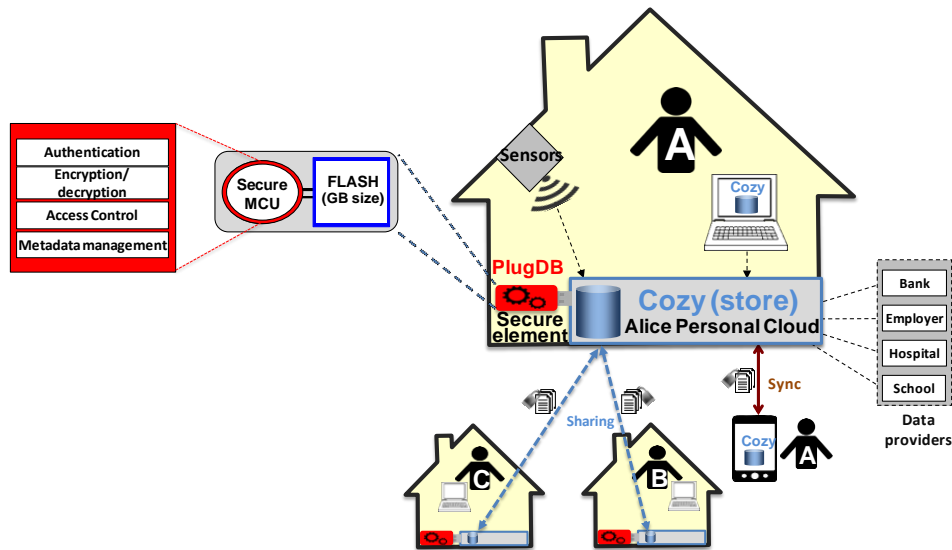


Figure 1: Secure Personal Cloud Platform

This architecture works as follows. Heterogeneous data issued by external sources (e.g., banks, employers, hospitals, commercial web sites) and by personal appliances (e.g., quantified-self devices, smart meters, domestic sensors, cameras) are all transformed into documents. Document metadata is extracted and stored in PlugDB and document themselves are encrypted by PlugDB before being integrated in the Cozy store. Encryption keys never leave the secure sphere of PlugDB. Once encrypted, documents can be safely stored in the cloud for resiliency reason or exchanged among users. Hence, PlugDB acts as a doorkeeper protecting the complete digital environment managed by Cozy. At the time an application/user asks for a Cozy document, the request is routed to PlugDB which evaluates the access control policy and decrypts the requested document if the authorization check succeeds. Thanks to PlugDB tamper-resistance, sticky policies can be safely enforced when documents are shared among individuals. Encryption keys and access/usage control rules are transferred along with the shared document and are enforced at the recipient side, with no way for the recipient user to tamper with the processing on his own platform. The implementation of a proof-of-concept of this architecture is partly supported by the SECSi French PIA (Programme Investissement Avenir) project.

Important scientific, technological, societal and legal issues are raised by such Secure Personal Cloud architecture. Notably, our objective is to enable the execution of distributed queries linking the personal data of several individuals with the guarantee that neither the result of the query, nor the observation of all intermediate steps of the execution discloses any information about a particular individual [4]. In other words, this means being able to implement Privacy-by-Design big data treatments on personal data. Another important challenge is to ease the declaration and administration of the access control policies by the individuals. Our hope is that the Secure Personal Cloud approach will provide a credible alternative to the systematic centralization of personal data on servers and will pave the way for new privacy-by-design architectures.

Reference:

- [1] S. Abiteboul, B. André, D. Kaplan, 'Manage your digital life in your personal info management system', *Communications of the ACM*, 2015, 58 (5), pp.32-35.
- [2] N. Anciaux, S. Lallali, I. Sandu Popa, P. Pucheral. 'A Scalable Search Engine for Mass Storage Smart Objects'. *Proc. of the 41th International Conference on Very Large Data Bases (VLDB)*, Hawaiï, PVLDB 8(9): pp 910-921, September 2015.
- [3] N. Anciaux, L. Bouganim, P. Pucheral, Y. Guo, L. Le Folgoc, S. Yin, "MILO-DB: a Personal, Secure and Portable Database Machine", *Distributed and Parallel Database Journal (DAPD)*, 32(1), 2014.
- [4] C. To, B. Nguyen, P. Pucheral. 'Private and Scalable Execution of SQL Aggregates on a Secure Decentralized Architecture', *ACM Transaction on Database Systems (TODS)*, to appear (<http://tods.acm.org/>).

Links:

PlugDB: <https://project.inria.fr/plugdb/>

Cozy Cloud: <https://cozy.io/fr/>

Please contact:

Benjamin André

Cozy Cloud, France

E-mail: benjamin@cozycloud.cc

Philippe Pucheral

University of Versailles & INRIA, France

E-mail: Philippe.Pucheral@inria.fr