

Trusted Cells: A Sea Change for Personal Data Services

Nicolas AnCIAUX^{1,2}, Philippe Bonnet³, Luc BouganIM^{1,2},
Benjamin Nguyen^{1,2}, Iulian Sandu Popa^{1,2}, Philippe Pucheral^{1,2}

¹ INRIA Paris-Rocquencourt
Le Chesnay, France
<Fname.Lname>@inria.fr

² PRISM Laboratory
Univ. of Versailles, France
<Fname.Lname>@prism.uvsq.fr

³ IT University of Copenhagen
Copenhagen, Denmark
phbo@itu.dk

ABSTRACT

How do you keep a secret about your personal life in an age where your daughter's glasses record and share everything she senses, your wallet records and shares your financial transactions, and your set-top box records and shares your family's energy consumption? Your personal data has become a prime asset for many companies around the Internet, but can you avoid -- or even detect -- abusive usage? Today, there is a wide consensus that individuals should have increased control on how their personal data is collected, managed and shared. Yet there is no appropriate technical solution to implement such personal data services: centralized solutions sacrifice security for innovative applications, while decentralized solutions sacrifice innovative applications for security. In this paper, we argue that the advent of secure hardware in all personal IT devices, at the edges of the Internet, could trigger a sea change. We propose the vision of *trusted cells*: personal data servers running on secure smart phones, set-top boxes, secure portable tokens or smart cards to form a global, decentralized data platform that provides security yet enables innovative applications. We motivate our approach, describe the trusted cells architecture and define a range of challenges for future research.

1. INTRODUCTION

With the convergence of mobile communications, sensors and online social networks technologies, we are witnessing an exponential increase in the creation and consumption of personal data. Paper-based interactions (e.g., banking, health), analog processes (e.g., photography, resource metering) or mechanical interactions (e.g., as simple as opening a door) are now sources of digital data linked to one or several individuals. They represent an unprecedented potential for applications and business.

Until now, the enthusiasm for new opportunities has thwarted privacy concerns. Nevertheless, the risk of a backlash is growing as new devices and new services bring us closer to the dystopias described in the science fiction literature. This risk is well documented and the nature of the solution is consensual: it is necessary to increase the control that individuals have over their personal data [11,9,12]. The World Economic Forum even formulates the need for a data platform that allows individuals *to manage the collection, usage and sharing of data in different contexts and for different types and sensitivities of data* [13].

Unfortunately, none of the solutions available today can be used to implement this vision. Centralized solutions, including emerging cloud-based personal data vaults management platforms¹, trade security and protection for innovative services. At best, such approaches formulate sound privacy policies, but none of them propose mechanisms to automatically enforce these policies [1]. Even TrustedDB [3], which proposes tamper-resistant hardware to secure outsourced centralized databases, does not solve the two intrinsic problems of centralized approaches. First, users get exposed to sudden changes in privacy policies. Second, users are exposed to sophisticated attacks, whose cost-benefit is high on a centralized database.

Decentralized solutions are promising because they do not exhibit these intrinsic limitations. However, existing decentralized solutions sacrifice functionality or usability for security. Many examples are discussed in [8]. Other examples include the PDS vision [2] or the FreedomBox [4]. In PDS, a personal data server is embedded in a tamper-resistant portable token to hold the personal data of a user, but the sharing of data is cumbersome (since the tokens are mostly disconnected) and the range of personal services is limited (since the tokens have extreme resource constraints). FreedomBox aims at providing a software platform that interconnects groups of individuals that trust each other, thus drastically limiting the range of services it can support.

We argue that the advent of secure hardware embedded in all forms of personal devices, at the edges of the Internet, will trigger a sea change. Recently, AMD announced that it will incorporate a secure Trust Zone-based² ARM processor on its chips to be included into smart phones, set-top boxes and laptops. Such secure tamper-resistant microcontrollers provide tangible security guarantees in the context of well-known environments³. We can now imagine that whenever you take a picture, your smart phone securely contacts the personal services of all individuals in the frame of the picture, and automatically blurs the face of those who request it. We can also imagine that the GPS tracker in your son's car gives him detailed turn-by-turn guidance, but hides those details to local government, only delivering road-pricing results.

In this paper, we propose the vision of **trusted cells**, i.e., personal data servers running on secure devices to form a decentralized data platform. We illustrate how trusted cells can be used in the context of an application scenario, describe the trusted cells architecture and discuss requirements and challenges for future research.

This article is published under a Creative Commons Attribution License (<http://creativecommons.org/licenses/by/3.0/>), which permits distribution and reproduction in any medium as well allowing derivative works, provided that you attribute the original work to the author(s) and CIDR 2013.

6th Biennial Conference on Innovative Data Systems Research (CIDR '13) January 6-9, 2013, Asilomar, California, USA.

¹ These include Personal (<http://www.personal.com>), My personal vault (<http://www.mypersonalvault.com>), or Mydex (<http://www.mydex.org>).

² <http://www.arm.com/products/processors/technologies/trustzone.php>

³ The adoption of a standard API for secure micro-controllers [5] and the availability of an open source embedded secure operating system based on it (Open Virtualization) now enable higher level services.

2. MOTIVATION

Alice lives in France with Bob and their two children. Their house is now one of the 35 million households equipped with a Linky power meter. The power meter reports once a day to the distribution company, a certified time series of readings for verification, billing and network operation [6]. Alice and Bob have installed an energy butler app on their secure home gateway, a trusted cell managing all smart appliances in their home and storing their data. That award-winning app relies on external feeds from their utility and local weather prediction, as well as a feed of readings received every second from the Linky⁴, to control their heat pump and the charge of their electrical vehicle. This app minimizes overall load on the distribution network and saves them 30% on their bill. In addition, Alice is engaged in a social game (a follow-up to simpleEnergy.com) where she competes with some friends on their energy savings, reducing consumption by 20%.

At the 1Hz granularity provided by the Linky, most electrical appliances have a distinctive energy signature. It is thus possible to infer from the power meter data which activities Alice and Bob are involved in at specific points in time [7]. How do Alice and Bob configure the home gateway trusted cell to preserve privacy while preserving the benefit of their applications? They have a shared account on this trusted cell. Bob, Alice and their children have agreed that they do not want to fully disclose all their activities to each other. They rather have access to 15 min aggregates via a visualization app – at that granularity one cannot detect specific activities, but it is still possible to infer a daily routine. At the same time, daily statistics feed their social game, monthly statistics are delivered to the distribution company and time series at required granularity are securely exchanged with other trusted cells in their neighborhood to achieve consumption peak load shaving.

None of this data leaves the trusted cell application unless it is accessed via a predefined set of aggregate queries. The trusted cell guarantees that no malware can tamper with the data. If the trusted cell gets stolen, an elaborate attack would need to be mounted to break the secure hardware and get access to their personal data.

This scenario can be easily transposed to different types of personal data like GPS traces, Internet traces, mobile phone data, bills, pay slips, photos as well as health, administrative or scholar records. We classify the data that could be managed with trusted cells, based on how and who actually produces it:

- (1) *Data produced by smart sensors* installed by companies in the user's home (e.g., power-meter, heat sensor) or in the user's environment (e.g., user's car GPS tracking box for a PAYD application) on which the user has full or shared ownership, externalizing aggregated data. Users may opt-in for small-scale sharing (e.g., local traffic optimization) or larger-scale sharing (e.g., social games or traffic optimization).
- (2) *Data produced or inferred by external systems* (e.g., purchase receipt obtained by near field communication or medical data sent by the hospital or labs). Small-scale sharing allows the user to optimize her buying habits or to compare her medical treatment with people having the same disease. Larger-scale

⁴ In France, such a short-range radio link is a requirement from the regulation authorities. In other countries, the data from a smart meter might not be directly accessible. In the US for example, the Green Button initiative allows customers to obtain online the smart meter data collected by their utility (<http://www.greenbuttondata.org/>)

sharing brings public health insights (e.g., epidemiological study cross-analyzing diseases and alimentation).

- (3) *Data authored by the user herself* (e.g., a photo, a mail, a document) on which she has complete ownership. Small-scale sharing benefit is obvious here. Larger-scale sharing of partial data (e.g., photo location only, number of exchanged mails) is undoubtedly a source of precious information (e.g., most interesting places on Google maps).

3. TRUSTED CELLS ARCHITECTURE

What personal data services actually run on a trusted cell? How do these services allow a user to control whom she shares her secrets with? How do applications access these services? What kind of guarantees do trusted cells offer about the security of the data they manage? We obviously do not aim at answering those questions fully in this paper. Our goal here is to draw the contours of an architecture based on *Trusted Cells* interconnected via an *Untrusted Infrastructure*.

Trusted Cells: A trusted cell implements a client-side reference monitor [10] on top of secure hardware. At a minimum, the hardware must guarantee a clear separation between secure and non-secure software. We abstract a Trusted Cell as (1) a Trusted Execution Environment, (2) a tamper-resistant memory where cryptographic secrets are stored, (3) an optional and potentially untrusted mass storage and (4) communication facilities. Physically, a trusted cell can either be a stand-alone hardware device (e.g., a smart token) or be embedded in an existing device (e.g., a smartphone based on ARM's TrustZone architecture).

The very high security provided by trusted cells comes from a combination of factors: (1) the obligation to physically be in contact with the device to attack it, (2) the tamper-resistance of (part of) its processing and storage units making hardware and side-channel attacks highly difficult, (3) the certification of the hardware and software platform, or the openness of the code, making software attacks (e.g., Trojan) also highly difficult, (4) the capacity to be auto-administered, contrary to high-end multi-user servers, avoiding insider (i.e., DBA) attacks, and (5) the impossibility even for the trusted cell owner to directly access the data stored locally or spy the local computing (she must authenticate and only gets data according to her privileges).

In terms of functionality, a full-fledged trusted cell should be able to (1) acquire data and synchronize it with the user's digital space, (2) extract metadata, index it and provide query facilities on it, (3) cryptographically protect data against confidentiality and integrity attacks, (4) enforce access and usage control rules, (5) make all access and usage actions accountable, (6) participate to computations distributed among trusted cells. Basic (e.g., sensor-based) trusted cells may implement a subset of this.

Untrusted infrastructure: The infrastructure provides the storage, computing and communication services, which expand the resources of a single trusted cell and form the glue between trusted cells. By definition, the infrastructure does not benefit from the hardware security of the trusted cell and is therefore considered untrusted. We consider that the infrastructure is implemented by a Cloud-based service provider⁵.

In terms of functionality, the untrusted infrastructure is assumed to: (1) ensure a highly available and resilient store for all data outsourced by trusted cells, (2) provide communication facilities

⁵ A P2P infrastructure among trusted cells could be envisioned but would raise many technical issues of limited interest for this article.

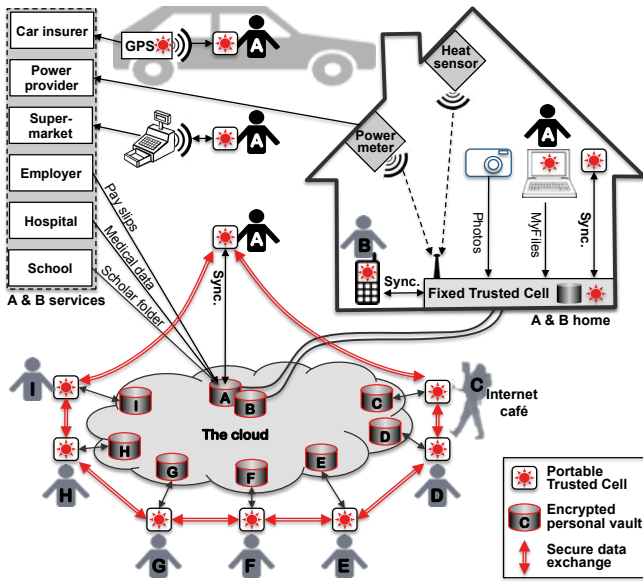


Figure 1: Alice (A) and Bob (B) are equipped with fixed and portable trusted cells, acquiring data from several data sources, synchronizing with their encrypted personal digital space on the cloud. Charlie (C) is travelling around the world and can securely access all his data from any (unsecure) terminal thanks to his portable trusted cell. All users equipped with trusted cells can securely share their encrypted data through the cloud.

among cells and (3) participate to distributed computations (e.g., store intermediate results), provided this participation can be guaranteed harmless by security checks implemented at the trusted cells side.

Figure 1 illustrates how trusted cells and the untrusted infrastructure can collaborate to implement scenarios meeting the privacy requirements stated above.

Threat model: In our context, the primary adversary is the infrastructure. The infrastructure may deviate from the protocols it is expected to implement with the objective to breach the confidentiality of the outsourced data. Integrity attacks (e.g., on data related to access control) must also be deterred since they may lead to subsequent confidentiality leaks. The infrastructure is assumed trying to cheat only if it cannot be convicted as an adversary by any trusted cell. Indeed, revealing a data leak (or a denial of service) in a public place would cause irreversible political/financial/legal damage to the service provider. Such adversaries are usually called malicious adversary having weakly malicious intents [14]. Trusted cells are themselves presumably trusted. However, even secure hardware can be breached, though at very high cost, so that one cannot exclude with certainty that a very small number of trusted cells be compromised. Hence, the trusted cells' cryptographic secrets must be managed in such a way that a successful attack on a (small set of) trusted cell cannot degenerate in breaking class attack. This is of utmost importance considering also that an individual succeeding in breaking her trusted cell could have effective malicious intents.

4. REQUIREMENTS AND CHALLENGES

We identify five major requirements for the user to actually control how the data entering her personal digital space is collected, protected, shared and finally used.

Controlled collection of sensed data: The targeted user(s) should be the unique recipient(s) of raw sensed data and would accept externalizing only aggregates by opting in/out for selected applications/services.

At home, the power meter continuously pushes raw measurements to Alice's and Bob's trusted cell gateway, while a certified aggregated time series is sent to the power supplier company and aggregates for a social game are pushed to the Cloud every day. Similarly, the tracking box installed on Alice's car is a trusted cell delivering aggregated GPS data to her insurer and raw data to her trusted cell smartphone that she will synchronize with her personal space for further use when back home. Hence, adding a trusted cell to a sensor, allows defining e.g., the frequency and or precision of the data that should be externalized, thus leading to a *trusted source* both for the user (in terms of privacy preservation) and the provider (in terms of certification of the output data).

Related challenges: Co-design is a primary issue to allow the definition of affordable sensor-based trusted cells. Low-cost is indeed a prerequisite to the generalization of trusted sources, capable of securely filtering and aggregating stream-based spatio-temporal data with tiny hardware resources. Some trusted sources being weakly connected to the Internet; asynchrony problems must also be addressed. Finally, the combination of data streams from multiple sources, each being separately harmless, may generate new privacy risks that must be carefully tackled.

Secure private store: All data must be made highly available, resilient to failure and protected against confidentiality and integrity attacks. Accessing this data from any terminal, including those outside the user's ownership sphere (e.g., internet café), should leave no trace of the access.

Cryptographic techniques (i.e., encryption, hashing, signatures) are used to protect trusted cell's data, keeping cryptographic keys in their tamper-resistant memory. The data is then stored in the Cloud and potentially cached in the trusted cell local mass storage. At a minimum, trusted cells keep locally extended metadata: access information, indexes, keywords, and cryptographic keys. Metadata should be sufficient to allow performing queries before accessing the Cloud to retrieve the data of interest. Cryptographic keys never leave the trusted cells tamper-resistant memory. Hence a trusted cell can be used to get securely data from any (untrusted) terminal it is connected with.

Related challenges: Designing an intuitive HCI for managing this bunch of heterogeneous personal data (data modeling, data integration, querying) is a major challenge. Besides, a significant amount of data and metadata is likely to be embedded in some trusted cells and may need to be queried efficiently. While it does not seem a major issue in powerful trusted cells (e.g., a smart phone), it appears much more challenging when facing low-end hardware devices like secure tokens (e.g., a microcontroller with tiny RAM, connected to NAND Flash chips or SD cards, possibly with energy consumption constraints). Whatever their complexity, trusted cells should also be designed to support self-tuning, self-diagnosis and self-healing to minimize the management burden put on the trusted cell owner.

Secure sharing: The user can decide to keep her data private or share it with other users or group of users under certain conditions (e.g., time, location). Under which model the access control policies are actually defined is an open issue, but not the main concern of this paper. However, we insist that the user must get a proof of legitimacy for the credentials exposed by the participants of a data exchange and must trust the evaluation of the exchange conditions (if any).

Practically, sharing data means sharing the associated metadata (so that the recipient user can get the referenced data in the Cloud), the cryptographic keys (so that her trusted cell can decrypt them) and the sticky policy (so that her trusted cell can enforce the expected access control rules). Hence, thanks to its security properties, including the protection against illegitimate actions of the recipient user, the recipient trusted cell can enforce all the conditions appearing in the access control rules (user's credential, contextual conditions).

Related challenges: Again, an intuitive HCI for defining the access control policies and simple modes of operation must be devised. The trusted cells themselves may be a source of simplification (e.g., integration of biometric sensors to automatically authenticate users, automatic production of certified credentials safely computed on a trusted cell, definition of default policies by trusted third parties – e.g., citizen associations – which could be automatically selected depending on a computed individual's profile). Also, secret management is at the heart of any sharing protocol between trusted cells (i.e., at this level a secret is a cryptographic key) and must be carefully designed (e.g., class-breaking attacks must be prevented, master secrets must be restorable in case of crash/loss of a trusted cell).

Secure usage and accountability: Usage control usually refers to $UCON_{ABC}$ [8]: obligations (actions a subject must take before or while it holds a right), conditions (environmental or system-oriented decision factors), and mutability (decisions based on previous usage)⁶. Again, defining appropriate usage control policies for trusted cell applications is an open issue.

Similarly to access control rules, usage control rules can be implemented as sticky policies so that they are made cryptographically inseparable from the data to be protected. Hence usage control rules will be enforced by any trusted cell downloading data and cannot be bypassed by the recipient user. Regarding accountability, the recipient trusted cell can maintain an audit log, encrypt it and push it on the Cloud to the destination of the originator trusted cell.

Related challenges: Many challenges are common with secure sharing. However, trusted cells hold the promise of new usages and new usage controls. For example, trusted cells could be parameterized so that any personal data produced by a trusted source linked to an individual A and referencing individual B be submitted for approbation to B's trusted cell before being integrated to A's digital space.

Shared Commons: Privacy has also a collective dimension in the sense that preserving one's privacy should not hinder societal benefits (e.g., census, epidemiologic releases, global queries). A trusted cell user is thus expected to participate to global treatments assuming her data suffers appropriate transformations (e.g., anonymization, output perturbation) depending on the trustworthiness of the recipient(s) and the expected usage of the data/query. When data needs to be transformed before being delivered, the recipient trusted cell implements the transformation on its own if possible (e.g., filtering, local data perturbation) or in collaboration with other trusted cells if the transformation requires a collective action (e.g., anonymization, global data perturbation). In the latter case, the computation may be implemented in a pure Secure Multi-Party fashion or may require the participation of the untrusted infrastructure (e.g., to store intermediate results).

⁶ For instance, a photo could be accessed ten times (mutability), in the course of 2012 (condition), informing the owner of the precise access date (obligation).

Related challenges: Such large scale computations may lead to atypical distributed protocols combining security and performance requirements in an asymmetric context made on one side of a very large number of highly secure, low power and weakly available trusted cells and on the other side of a highly powerful, highly available but untrusted infrastructure. Hence, the trusted cells architecture can be seen as a massive untrusted interconnection of trusted co-processors.

5. CONCLUSION

We proposed the trusted cell architecture, a vision reconciling individual's privacy with innovative acquisition and sharing of personal data. This vision is based on the premise of ubiquitous and open secure hardware. Trusted cells enforce access and usage control at the edges of the Internet, and thus constitute a sea change with respect to personal data management. This vision undoubtedly opens a set of exciting challenges that must be explored by the database community.

6. ACKNOWLEDGEMENTS

This work has been partially funded by the French ANR KISS project.

7. REFERENCES

- [1] R. Agrawal, J. Kiernan, R. Srikant, Y. Xu: Hippocratic Databases. VLDB 2002: 143-154
- [2] T. Allard et al.: Secure Personal Data Servers: a Vision Paper. PVLDB 3(1): 25-35 (2010)
- [3] S. Bajaj, R. Sion: TrustedDB: a trusted hardware based database with privacy and data confidentiality. SIGMOD Conference 2011: 205-216
- [4] FreedomBox: <http://freedomboxfoundation.org/>.
- [5] Global Platform Device Technology. Trusted Execution Environment Internal API Specification. Version 1.0. December 2011.
- [6] S. Katzenbeisser and K. Kursawe, Privacy and Security in Smart Energy Grids, *Dagstuhl Seminar 1151*, 2011
- [7] H. Lam. A Novel Method to Construct Taxonomy Electrical Appliances Based on Load Signatures. IEEE Transactions on Consumer Electronics, 2007.
- [8] A. Narayanan, V. Toubiana, S. Barocas, H. Nissenbaum, D. Boneh: A Critical Look at Decentralized Personal Data Architectures CoRR abs/1202.4503: (2012)
- [9] H. Nissenbaum, Privacy in context: Technology, policy, and the integrity of social life," *Stanford Law Books*, 2010.
- [10] J. Park and R. Sandhu, "The $UCON_{ABC}$ usage control model," *ACM Trans Information System Security*, vol. 7, no. 1, pp. 128-174, 2004.
- [11] A. Pentland et al. Personal Data: The Emergence of a New Asset Class. World Economic Forum. January 2011.
- [12] S. Petronio, Unpacking the paradoxes of privacy in CMC relationships: The challenges of blogging and relational communication on the internet, *In Computer-mediated communication in Personal Relationships*, 2011.
- [13] The World Economic Forum. Rethinking Personal Data: Strengthening Trust. May 2012.
- [14] N. Zhang, W. Zhao: Distributed privacy preserving information sharing. VLDB 2005.