2014

Authentication and authorisation in entrusted unions

Ayed F. Dhouha
Thales DSC

Jan Camenisch
IBM Research GmbH

Tanya Ignatenko
Eindhoven University of Technology

Michael N. Johnstone
Edith Cowan University

Paul Koster
Philips Research Eindhoven

See next page for additional authors

DOI: 10.4225/75/57b660c8343d4

12th Australian Information Security Management Conference. Held on the 1-3 December, 2014 at Edith Cowan University, Joondalup Campus, Perth, Western Australia.

AUTHENTICATION AND AUTHORISATION IN ENTRUSTED UNIONS[1]

Ayed F. Dhouha[1], Jan Camenisch[2], Tanya Ignatenko[3], Michael N. Johnstone[4],
Paul Koster[5], Brigitta Lange[6], Milan Petković[3,5], Dieter Sommer[2], John Zic[7]
[1]Thales DSC, [2]IBM Research GmbH, [3]Eindhoven University of Technology, [4]Edith Cowan University Security Research Institute, [5]Philips Research Eindhoven, [6]NEC Laboratories Europe, [7]CSIRO
dhouha.ayed@thalesgroup.com, jca@zurich.ibm.com, t.ignatenko@tue.nl, m.johnstone@ecu.edu.au, r.p.koster@philips.com, brigitta.lange@neclab.eu, milan.petkovic@philips.com, dso@zurich.ibm.com, john.zic@csiro.au

Abstract

This paper reports on the status of a project whose aim is to implement and demonstrate in a real-life environment an integrated eAuthentication and eAuthorisation framework to enable trusted collaborations and delivery of services across different organisational/governmental jurisdictions. This aim will be achieved by designing a framework with assurance of claims, trust indicators, policy enforcement mechanisms and processing under encryption to address the security and confidentiality requirements of large distributed infrastructures. The framework supports collaborative secure distributed storage, secure data processing and management in both the cloud and offline scenarios and is intended to be deployed and tested in two pilot studies in two different domains, viz, Bio-security incident management and Ambient Assisted Living (eHealth). Interim results in terms of security requirements, privacy preserving authentication, and authorisation are reported.

Keywords

Security, Authentication, Authorisation, Encryption

INTRODUCTION

There appears to be an emerging need to securely share information with multiple, often independent parties, across organisational or jurisdictional borders, e.g., in the case of world-wide outbreaks of communicable diseases. In order to achieve this sharing, those parties form virtual distributed (e.g., cloud-based) systems. However, the key problem with such distributed systems is that adequate authentication and identity management systems must be used, else the parties do not trust and therefore do not use the shared infrastructure. This is a non-trivial task, since those parties often use systems that belong to different security domains governed by different authorities and use different identity attributes that are utilised in their access control polices. Moreover, in distributed systems or in the cloud, control over the environment where applications are running and data are stored is abrogated in favour of cost or convenience. Furthermore, the privacy aspect of using such systems should not be ignored. Dealing with identity management systems in distributed environments, a user leaves behind a digital identity footprint on every application or service /she accesses. This information can be misused by the application or leaked to attackers. The identity management provider learns the user/application relationships, and also this information can be used to monitor the users or can leak to attackers, thus compromising privacy of the users.

Cloud-based services are an opportunity not to be wasted in terms of maximising value for any firm's IT budget. Such services seem attractive until issues of confidentiality or privacy are raised. A standard approach to the problem of confidentiality is to encrypt stored data. This approach is effective only if these data are not required to be processed in any way whilst in the cloud. In this case, the data must usually be decrypted (which will likely violate confidentiality) and stored temporarily whilst a database query is executed. To assure end-users of services that their privacy is maintained and that confidentiality is preserved, a means must be found to execute such queries on an encrypted dataset.

[1] This work has been partially funded by the EC via grant agreement no. 611659 for the AU2EU project.

To further complicate matters in terms of user trust, most software is insecure, according to Shostack and Stewart (2008). This could be because, as Wysopal et al. (2007) note, security requirements are often omitted from requirements specifications altogether. Recent events, such as the publication of the "shellshock" vulnerability, do nothing to persuade users that software systems of any type are to be trusted. Therefore, the way forward is to provide users with a system that offers privacy-preserving authentication, testable and correct authorisation and homomorphic encryption to preserve confidentiality.

This paper describes an authentication and authorisation framework that provides those properties of privacy-preservation, authorisation and homomorphic encryption. To address the problems described above, we have initiated an EU FP7 project called AUthentication and AUthorisation for Entrusted Unions (AU2EU), which is a joint collaboration between seven EU partners and five research institutes and universities in Australia.

THE PROBLEM SPACE

Multiple parties involved in a distributed system often belong to different security domains governed by different authorities and use different identity attributes that are utilised in their access control polices. Therefore, there is a need for a framework that enables mechanisms for matching, mapping and applying of different roles, attributes, and access policies, when a party needs to use resources from different domains.

Another important requirement surfaces when (even before engaging in any collaboration), interested parties need to be able to assess reliability of the potential collaboration. This can be achieved by assessing trust indicators of the involved devices, platforms and services. Thus the desired infrastructure should support provisioning of trust metrics and signals by assessing trustworthiness, privacy, and security mechanisms of a service, platform or a device in question (e.g., interpretation of privacy policy, assessment of security, reputation, certification, data provenance etc.).

When a decision to collaborate is made, and initial trust levels are established, authentication processes that involve identities pose the next challenges. Dealing with identity management systems in distributed environments, any user leaves behind a digital identity footprint on every application or service s/he accesses. This poses privacy threats, since information can be misused by identity management providers, applications and attackers. Thus protection of credentials and assurance of unlinkability, while achieving a high level of security, privacy, mobility, usability and reduced cost, are other requirements for a superior authentication and authorisation infrastructure. On the other hand, to assess authentication claims coming from various parties involved in collaboration, assurance of claims should be designed and integrated into any framework supporting cross-domain collaborations.

From an authorisation point of view, regulation of access to information and attributes brings forward the requirements to guarantee privacy and confidentiality of information, controlled disclosure of attributes and information, and to provide information on proof of attributes and validity of information (data provenance, origin authentication) in distributed systems. This can be achieved by incorporating cryptographic policy enforcement, digital user consent and again assurance of claims mechanisms. Moreover, since in distributed systems one cannot control the complete environment where applications are running and data are stored, processing of (sensitive) data in the encrypted domain also plays an important role in addressing the aforementioned requirements.

Finally, all security and privacy mechanisms involved in a large distributed infrastructure should be composed in such a way to provide a seamless integrated operating architecture. This architecture guarantees security and privacy claims supported with each individual security mechanism across multiple jurisdictions and security domains. From a practical perspective, the corresponding platform should be easily scalable and allow for integration of additional systems, allowing parties and users to join and leave collaborations as needed. Another practical aspect that needs to be taken into account is efficiency – especially important in mission-critical applications.

THE FRAMEWORK

In order to design a joint eAuthentication and eAuthorisation framework (henceforth, "the framework"), which satisfies the concepts mentioned above, we take the XACML-based authorisation work of Rissanen (2010), and combine it with the ABC4Trust (Camenisch et al., 2011) authentication architecture as well as with the TDL (Trust in Digital Life) system described by Bjones (2012). The TDL system defines the core generic building blocks needed to build large-scale distributed authentication architectures.

The creation and use of the framework provides for:

- Cross-domain and jurisdictional collaborations, supporting different identity/attribute providers and organisational policies and guaranteeing privacy, security and trust;
- Advancement of the state-of-the-art with assurance of claims, trust indicators, policy enforcement mechanisms and processing under encryption techniques to address specific security and confidentiality requirements of large distributed infrastructures;
- Implementation that supports collaborative secure distributed storage, secure data processing and management in both the cloud and offline scenarios;
- Deployment in two pilots on bio-security incident management and collaborative services in Australia and on eHealth and Ambient Assisted Living in Europe; and
- Validation of practical aspects such as scalability, efficiency, maturity and usability.

The aforementioned activities contribute to increased trust, security and privacy, which in turn shall lead to the increased adoption of (cloud-based) critical infrastructures and collaborative delivery of services dealing with sensitive data. Fig. 1 depicts the key elements of the framework.

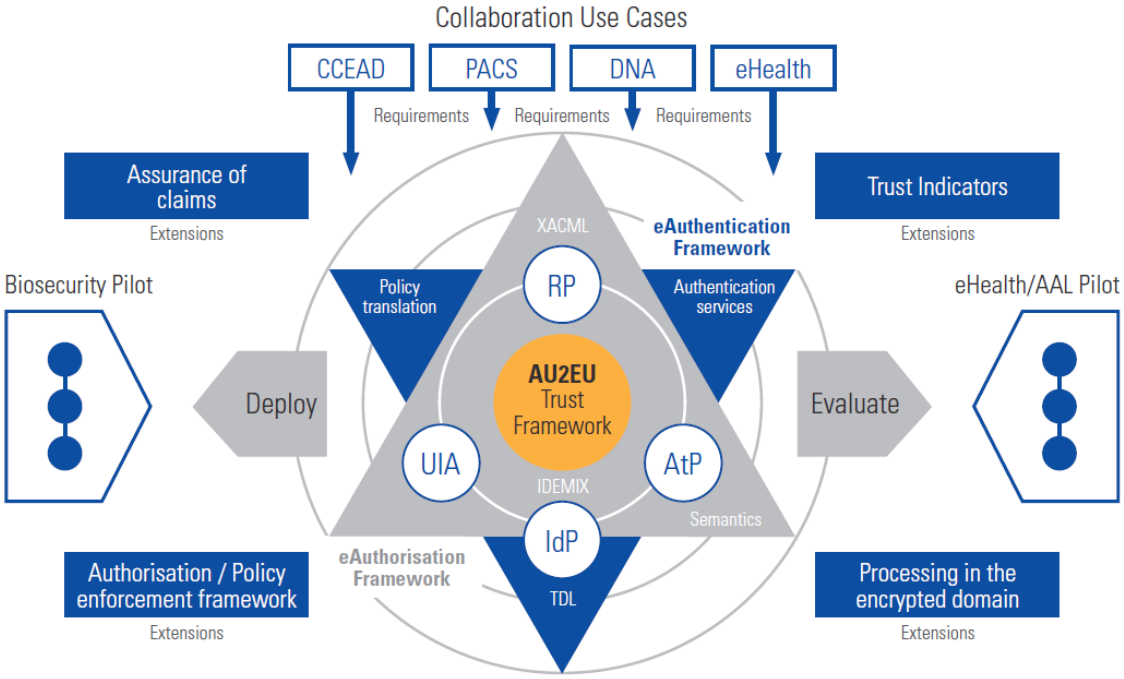

Figure 1: Conceptual Model of the Platform.

ANALYSIS AND DISCUSSION

The main goal of this project is to build a joint eAuthentication and eAuthorisation infrastructure by incorporating privacy-preserving cloud-based authentication service and techniques for unifying attributes and authorisation policies of different security domains. Moreover, we develop and integrate the technologies for assurance of claims, trust indicators, cryptographic policy enforcement and mechanisms to perform operations under encryption. In this section we describe the approach to develop the aforementioned solutions. Fig. 1 captures the structure of our approach to achieving our objectives and creating the proposed joint eAuthentication and eAuthorisation infrastructure.

We start with the analysis of the identified uses cases for cross-organisational collaborations that deal with sensitive data and critical infrastructures and derive from them the key requirements for our framework and reference architecture. Next, to create a collaborative system with distributed resources, we design the framework by taking as a starting point XACML authorisation and extending it with mechanisms for policy and attribute mapping between different security domains, as well as with policy translation needed to support distributed collaborative environments. The elements of the Authentication framework (see Fig. 1 - relying party - RP, attribute provider - AtP, identity provider - IdP and user identity agent - UIA) and XACML architecture of the Authorisation framework are integrated into a joint framework. To support user privacy, as already

mentioned, we bring in the ABC4Trust and TDL eAuthentication systems and deploy digital user consent and Identity Mixer technologies to realise it.

We also extend the described framework with novel advanced solutions for cryptographic policy enforcement, trust indicators, assurance of claims and mechanisms for processing data in an encrypted form to address specific reliability, confidentiality and privacy requirements of distributed collaborations. The resulting infrastructure will be deployed in two pilots, which are described below, one on bio-security incident management in Australia, and another one on collaborative services for eHealth and AAL in Europe. Using these pilots, we evaluate the security, usability, scalability and flexibility of our framework for distributed collaborations. The two pilots also allow us to analyse and unify the framework for joint Australian and European collaborations taking into account legal, procedural and business differences.

Framework

In this section we describe the framework. First we focus on the fundamental concepts required for creating a sustainable trustworthy large-scale infrastructure as identified by the Trust in Digital Life (TDL) alliance (Bjones, 2012). These concepts are:

1. *Composable architecture:* The architecture to deliver advanced scenarios builds on traditional federated identity and authorisation architectures. This provides a tiered approach to handling the more complex scenarios and privacy requirements of the different entities that may be involved. This leads to greater flexibility in the application, and the ability to consume information on identity and access control policies from multiple sources in an interoperable way.
2. *Open to technology and standards evolution:* The components of the architecture must communicate using open standards. The architecture itself should not be bound to a single protocol, as over time protocols change and depending on particular needs different protocols might be required.
3. *Attributes remain with the source of the data:* To help reduce security and privacy risks, personal data should reside with the data owner rather than be moved elsewhere. Including multiple attribute providers allows the data to be kept separate from the identity provider.
4. *User consent/privacy preferences:* Electronic user consent mechanisms should allow a user to engage in different services and make different privacy choices related to use, communication and disclosure of his/her information to different services and involved parties. Also, from the collaborating party or service provider point of view, it is beneficial if the user consent/privacy preferences can be directly used together with their organisational security policies as it enables a uniform policy enforcement framework for governing access to sensitive data.
5. *Privacy:* The architectures should ensure privacy protection of the user, supporting such concepts as (un)Traceability, characterising the degree to which a provider of digital identity information (a claims provider) can know where this information is used (at a relying party); (un)Linkability from relying party to claims provider; (un)Linkability between relying parties; and (non)Disclosure: the degree to which specific pieces of sensitive information are disclosed for a given transaction.
6. *Correctness and accountability:* When designing a privacy aware architecture it should also openly state the requirements for correctness of personal information and accountability in case of its incorrectness.

Governance of the resources in the distributed systems requires an authorisation framework to support attributes and policies belonging to different security domains. Therefore in our framework we deploy novel mechanisms for semantic mapping to translate policies and attributes to the required authentication claims that can be verified in our authentication framework. To guarantee strong authentication and privacy at the same time, we use idemix technology (IBM Research Zurich, 2010) for attribute-based authentication as a building block of our authentication architecture. We integrate this technology as a cloud-based service in our architecture to further support privacy as well as ease of use for collaborating parties.

Advancing the state-of-the-art

In the following we describe in more detail four topics with the cutting-edge functionalities that we develop in the AU2EU project and therefore advance the state-of-the-art.

Assurance of claims: Access control and identity management in distributed environments, like large-scale authentication infrastructures in the cloud, involves verifying claims that come from multiple parties. However, assessing such claims can be very challenging. One way to increase assurance of an identity-claim is the use of multiple-factor authentication that involves claims coming from different, independent parties. However, a

novel alternative approach can also investigate how to solve this issue by introducing device health claims (characterising the device trustworthiness). These claims provide information about the state of the user's device with respect to the user's privacy.

Authorisation/policy enforcement framework: When multiple distributed organisations have to collaborate and form a single (virtual) organisation, the problem arises on how to unify and create a common authorisation framework given the fact that each of the involved organisations has its own security domain, policies that are put in place and roles or attributes used in their systems. Therefore there is a need for a framework that can incorporate various cross-domain policies and attributes (including user digital consent) and perform policy enforcement. The latter in particular is challenging, especially when cloud-based systems are used whose access control system is itself not trusted. Then policy enforcement needs to be done using encryption of sensitive information according to an attribute-based policy in such a way that only users with certain attributes (e.g., roles) can access the information. Additionally, the framework needs to support complex policies (context-based or via assurance of claims) as well as infringement management or compliance assurance (e.g. a break-the-glass scenario).

Trust indicators: Large-scale authentication infrastructures build on traditional federated identity architectures. One of the advantages that this gives is a tiered approach to handling complex scenarios and privacy requirements of different involved entities. However, it also means that a user, who engages a service built based on such infrastructure, needs to be able to evaluate trustworthiness of such services that rely on a chain of various building blocks. Consider, for example, an eHealth service that manages patient medications, diet, appointments with healthcare practitioners and needs to use the data provided by the patient's hospital and Electronic Medical Records (EMR). Thus, there is a need for an evaluation technique that gives an aggregated trust indicator of a service by quantifying the security and privacy mechanisms used by the aforementioned service, interpreting and quantifying its privacy policies, and evaluating and quantifying its integrated reputation.

Operations in an encrypted domain: When sensitive data, e.g. patient DNA or bio-security incident data, need to be accessed and processed by various parties in the distributed collaborative systems, restricting access to only partial data is a difficult task, since processing and extracting partial information from the data often requires access to the whole dataset. Policy enforcement mechanisms alone cannot guarantee this fine-grained level of access control. Processing in the encrypted domain that builds on homomorphic encryption techniques (Fontaine and Galand, 2007), secure multi-party computation (Damgård et al., 2009; Damgård et al., 2010; Damgård et al., 2012) and code-based security (Finiasz and Sendrier, 2009; Barthe et al., 2009; De Cristofaro et al., 2012) suggest possible solutions. It would allow a party to engage into applications dealing with sensitive data and obtain required answers without revealing those data.

The Pilot Studies

Critical to the wide adoption of the framework is the deployment of the research developed in the project in two pilots. This effectively bridges the gap between research and practice, facilitating technology adoption by the market.

Fig. 1 shows that there are two pilot studies, one in the area of the planning, and execution of Biosecurity incident management, and a second pilot in the area of eHealth, specifically in Ambient Assisted Living. As the pilots are in different domains, it is expected that their differing requirements will provide an effective coverage of the functions of the framework as well as a legitimate test of the security and robustness of the framework.

Bio-security Incident Response Pilot: The pilot on bio-security incident management considered in the AU2EU project is an important mechanism to sustain Australia's agricultural productivity and related export trade that heavily relies on nation's disease free status. When an animal disease outbreak happens, the Consultative Committee on Emergency Animal Disease (CCEAD[2]) is formed to collaboratively analyse current research and diagnostic information and discuss the strategies for dealing with the emergency outbreak. Organisationally, the CCEAD represents a dynamic, distributed collaboration with strict access and security protocols. This collaboration involves multiple individuals, locations and groups that represent different sectors and states, amongst whom sensitive information must be shared in a timely manner, to facilitate the best evidence-based decisions. It is important to realise that despite committees such as the CCEAD having carefully defined and agreed upon protocols, each member of the committee has its own jurisdictions, policies and procedures for the control of sensitive information, as well as corresponding identity management and security systems. Therefore, a properly automated Authorisation and Authentication infrastructure is required that can incorporate the various identity management systems so as to: (i) ensure authentication with high levels of

[2] http://www.daff.gov.au/animal-plant-health/animal/committees/ccead

assurance; (ii) manage and enforce cross-domain access policies; and (iii) guarantee integrity of the information, access to specialised equipment, facilities and services in such a dynamic distributed system, even after the incident has been resolved and the formal collaboration has been terminated. The information shared and discussed during the incident needs to be curated properly, with a rich set of provenance information and access control mechanisms put in place, such that any subsequent incident can draw upon the historical information following strict security, privacy and control of information protocols.

The Bio-security Incident Response Pilot will carry out a series of committee meetings and evaluate their effectiveness and productivity improvement over current practices and procedures. The committees investigated are expected to be similar in structure and membership to the CCEAD, see Fig.2, whose members are located at research facilities and organisations across the Australian states of Queensland, NSW, and Victoria.

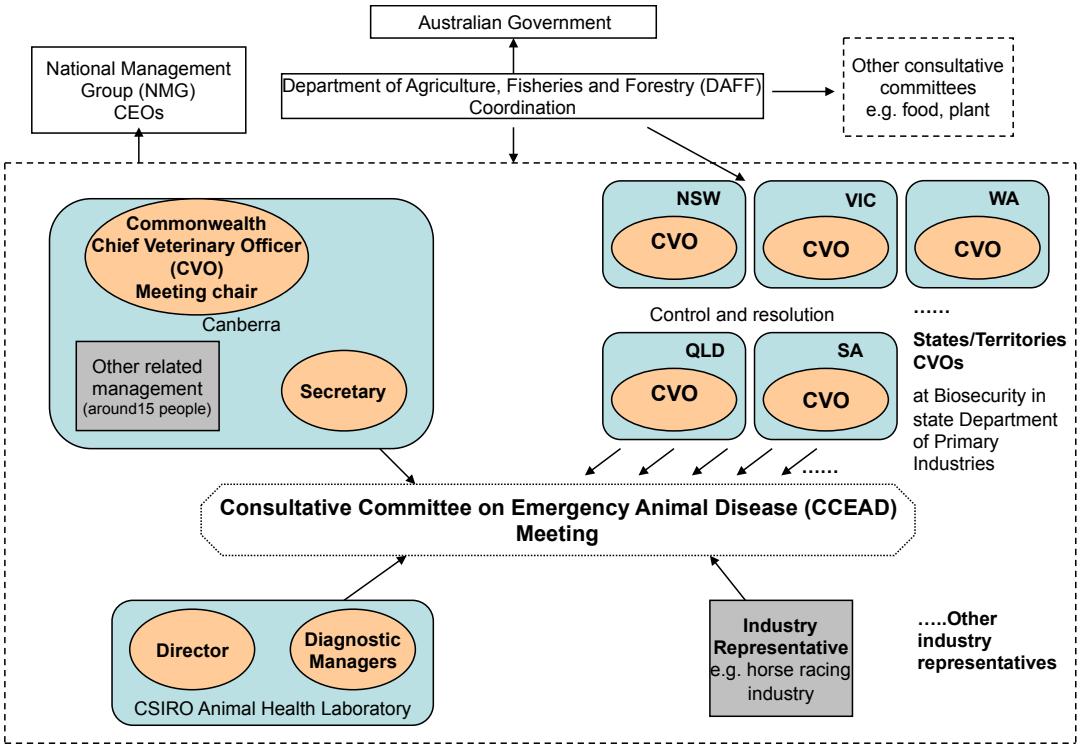

Figure 2 High level view of relationships between CCEAD members

The pilot will be based on the operational collaboration infrastructure developed by CSIRO. The CSIRO Collaboration Platform is currently used to discuss research results and operational matters between partner organisations of various committees. Thus this collaboration infrastructure is the kernel upon which the pilot software systems developed in the AU2EU project will be deployed and evaluated.

Collaborative Services for eHealth and Ambient Assisted Living Pilot: The framework will also be deployed in real-life conditions for a specific instance of a collaborative eHealth and Ambient Assisted Living (AAL) service.

The German Red Cross, Heidelberg (DRK) is a major regional home emergency call and social service provider in the Rhine-Neckar Metropolis region. The DRK is collaborating with several regional emergency and social care providers, as well as third party health and home care service providers in order to deliver tailored social care services to their customers. In its Home Emergency Call Service Centre (HEC) in Heidelberg the DRK coordinates 24h/7d emergency response and home care services as well as 3rd party service procurement (e.g., Menu-Services, Assisted Mobility /Travel-Services, Housekeeping and Nursing Assistance, etc.) for customers equipped with a home emergency call system and connected via analogue telephone networks, broadband IP-networks, or mobile networks to the HEC, see Fig. 3. As a service provider the DRK has responsibility for the safety and confidentiality of sensitive personal data used to deliver their services, safeguard customer privacy, and strengthen the trust relationship with customers. Intrinsic security and safety of all data transfer, authentication, and authorisation mechanisms therefore are essential. Thus, a reliable and efficient architecture for access control is also essential.

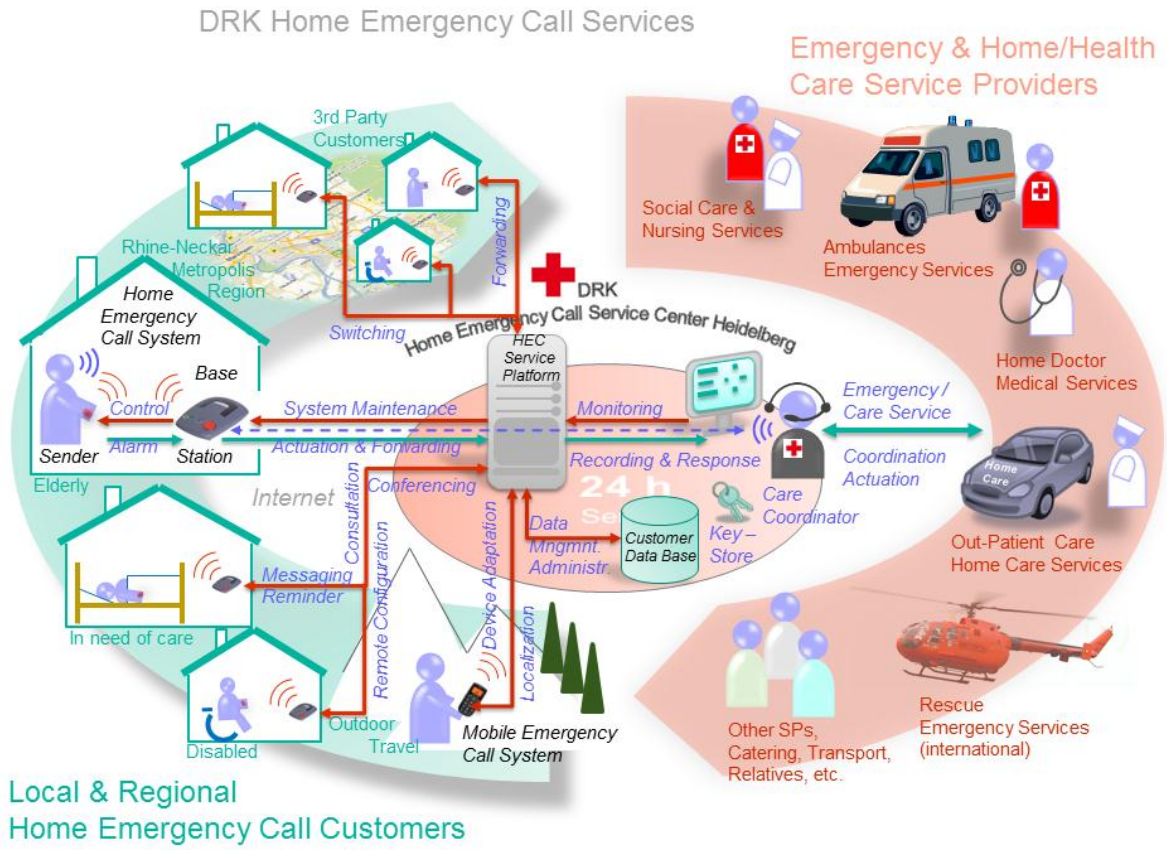

Figure 3 Overview DRK HEC Service and Collaboration Scenario

For this pilot, approximately 20 homes of DRK HEC customers will be equipped with AAL system infrastructure and hardware components, e.g., various sensors for in-home activity and status monitoring. Typically, customer data gathered from these installations is transferred to a dedicated server, where the data is analysed. Depending on the analysis results either an active or a reactive intervention by care service providers is triggered to assist the home customers. While the care coordination service provider (DRK) is the primary user of the data, there are situations when limited access to selected data has to be provided to third parties that are involved in the home-care/emergency service handling, e.g., doctors, home care service providers or care givers. The framework will facilitate the provision of flexible, efficient and personalised care services. The framework will be integrated with the existing HEC and AAL service infrastructure so that it can be tested and evaluated by real customers, in real-life settings.

CONCLUSIONS AND FURTHER WORK

This study explored the problem of trans-border authentication and authorisation. The complex nature of the policy claims assurance mechanism was revealed and techniques for representing trust indicators were articulated and discussed.

An integrated eAuthentication and eAuthorisation framework was introduced. We discussed the challenges and problems that have surfaced when designing a distributed eAuthentication and eAuthorisation infrastructure, especially one which enables trustworthy secure collaborations. Furthermore, we have presented our approach to extend the framework by advancing the state-of-the-art in four different domains: assurance of claims, policy enforcement, trust indicators and operations in an encrypted domain.

The next stage of the project involves testing the authentication and authorisation system and formally verifying the claim mechanisms, thus assuring trust in the system and fostering adoption in the marketplace.

REFERENCES

Barthe, G., Grégoire, B. and Béguelin, S.Z. (2009). "Formal certification of code-based cryptographic proofs," Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp. 90–101.

Bjones, R. (2012). "Architecture serving complex Identity Infrastructures", Trust in Digital Life, p. 21.

Camenisch, J., I. Krontiris, A. Lehmann, G. Neven, C. Paquin, K. Rannenberg, H. Zwingelberg, (2011). Architecture for Attribute-based Credential Technologies", ABC4Trust, Available at https://abc4trust.eu/index.php/pub/results/107-d21architecturev1

De Cristofaro, E., Faber, S., Gasti, P. and Tsudik, G. (2012). "Genodroid: are privacy-preserving genomic tests ready for prime time?", WPES 2012, pp. 97–108, 2012.

Damgård, I., Nielsen, J.B. and Wichs, D. (2009). "Universally Composable Multiparty Computation with Partially Isolated Parties," Theory of Cryptography, LNCS 5444, pp.315–331.

Damgård, I., Ishai, Y. and Krøigaard, M. (2010). "Perfectly Secure Multiparty Computation and the Computational Overhead of Cryptography," Advances in Cryptology - EUROCRYPT 2010, LNCS 6110, pp. 445–465.

Damgård, I, Pastro, V., Smart, N. and Zakarias, S. (2012). "Multiparty Computation from Somewhat Homomorphic Encryption," Advances in Cryptology - CRYPTO 2012, LNCS 7417, pp. 643–662.

Finiasz, M. and Sendrier, N. (2009). "Security Bounds for the Design of Code-Based Cryptosystems," Advances in Cryptology - ASIACRYPT 2009, LNCS 5912, pp. 88–105.

Fontaine, C. and Galand, F. (2007). "A Survey of Homomorphic Encryption for Nonspecialists," EURASIP Journal on Information Security.

IBM Research Zurich. (2010). "Specification of the Identity Mixer Cryptographic Library Version 2.3.1".

Rissanen, E. (2010). "eXtensible Access Control Markup Language (XACML) Version 3.0", OASIS standard, 2010, Available at: http://docs.oasisopen.org/xacml/3.0/xacml-3.0-core-spec-cs-01-en.pdf

Shostack, A. and Stewart, A. (2008). *The New School of Information Security*. Upper Saddle River, NJ: Addison Wesley.

Wysopal, C., Nelson, L., Dai Zovi, D. and Dustin, E. (2007). *The Art of Software Security Testing*. Upper Saddle River, NJ: Addison Wesley.