D. Hühnlein, H. Roßnagel (Hrsg.): Open Identity Summit 2013

GI-Edition

Lecture Notes in Informatics

Detlef Hühnlein, Heiko Roßnagel (Hrsg.)

Open Identity Summit 2013

**10.–11. September 2013
Kloster Banz, Germany**

223

Proceedings

Open standards and interfaces as well as open source technologies play a central role in the current identity management landscape as well as in emerging future scenarios based on cloud computing for example. While there are already plenty of successful applications in which those techniques are used to guarantee the authenticity and integrity of entities, there are still many closely related areas which demand further research. The aim of the "Open Identity Summit 2013" is to link practical experiences and requirements with academic innovations. Focus areas of this event are research and applications in the area of Identity Management and Open Source with a special focus on Cloud Computing.

Detlef Hühnlein, Heiko Roßnagel (Hrsg.)

Open Identity Summit 2013

10.-11.09.2013
Kloster Banz, Germany

Gesellschaft für Informatik e.V. (GI)

Volume Editors
Detlef Hühnlein
    ecsec GmbH
    Sudetenstr. 16, D-96247 Michelau, Germany
    E-Mail: detlef.huehnlein@ecsec.de
Heiko Roßnagel
    Fraunhofer IAO
    Nobelstr. 12, D-70569 Stuttgart, Germany
    E-Mail: heiko.rossnagel@iao.fraunhofer.de

Chairs' Message

Welcome to the "Open Identity Summit 2013", which has been jointly organized by the Special Interest Groups BIOSIG, CRYPTO and PET and the regional chapter Upper Franconia within the German Computer Science Society (Gesellschaft für Informatik), the EU-funded FutureID Project, the Open eCard Project, the Federal Association for Information Technology, Telecommunications and New Media (BITKOM), the Competence Center for Applied Security Technology (CAST), the open source initiative "Deutsche Wolke", the European Association for eIdentity and Security (EEMA), the IDentity.Next Association, the Open Cloud Initiative (OCI), the OpenID Foundation, the Open Identity Exchange (OIX), the Open Source Business Alliance (OSBA), the Open Source Integration Initiative (OSII), the TeleTrusT IT Security Association Germany, the SkIDentity Project, which aims at providing trustworthy identities for the cloud, and last but not least the Trusted Cloud Program supported by the German government.

The international program committee performed a scientific review process according to the LNI guidelines with at least five reviews per paper and accepted less than 47 % of the submitted papers as full scientific papers.

Furthermore, the program committee has created a program including selected contributions of strong interest (further conference contributions) for the outlined scope of this conference.

We would like to thank all authors for their contributions and the numerous reviewers for their work in the program committee.

Kloster Banz, 10th September, 2013

Detlef Hühnlein
ecsec GmbH

Heiko Roßnagel
Fraunhofer IAO

Chairs

Detlef Hühnlein
ecsec GmbH, Germany (detlef.huehnlein@ecsec.de)

Heiko Roßnagel
Fraunhofer IAO, Germany (heiko.rossnagel@iao.fraunhofer.de)

Program Committee

John Bradley, Arslan Brömme, Bud Bruegger, Hartje Bruns, Christoph Busch, Victor-Philipp Busch, Roger Dean, Jos Dumortier, Torsten Eymann, Hannes Federrath, Arno Fiedler, Lothar Fritsch, Jens Fromm, Walter Fumy, Ulrich Greveler, Thomas Groß, Oliver Hinz, Olaf Herden, Gerrit Hornung, Moritz Horsch, Detlef Houdeau, Detlef Hühnlein, Jan Jürjens, Stefan Katzenbeisser, Andreas Kuckartz, Andreas Kühne, Herbert Leitold, Luigi Lo Iacono, Nils Magnus, Tarvi Martens, Wolf Müller, Anja Lehmann, Peter Lipp, Johannes Loxen, Holger Mühlbauer, Alexander Nouak, Axel Nennker, Sebastian Pape, René Peinl, Sachar Paulus, Henrich C. Pöhls, Marco von der Pütten, Kai Rannenberg, Volker Roth, Alexander Roßnagel, Heiko Roßnagel, Ahmad-Reza Sadeghi, Ivonne Scherfenberg, Johannes Schmölz, Jörg Schwenk, David Simonsen, Tobias Wich, Don Thibeau, Thomas Uhl, Thomas Wieland, Alex Wiesmaier, Klaus-Dieter Wolfenstetter, Xuebing Zhou, Jan Zibuschka, Frank Zimmermann

Hosts and Partners

- **BIOSIG – Biometrics and Electronic Signatures (www.biosig.org)**

  The special interest group "Biometrics and Electronic Signatures" (BIOSIG) within GI e.V. is dedicated to the fundamentals, methods, techniques, processes and implementations used to guarantee the authenticity and integrity of entities.

- **CRYPTO – Applied Cryptology (fg-krypto.gi.de)**

  The special interest group "Applied Cryptology" (CRYPTO) within GI e.V. connects users and researchers in the area of cryptology, whereas the scope of activities comprises the design, analysis, implementation and practical application of cryptographic systems.

- **PET - Privacy-Enhancing Technologies (fg-pet.gi.de)**

  The special interest group "Privacy-Enhancing Technologies" (PET) within GI e.V. aims at introducing and promoting privacy-enhancing technologies in science, industry and policy.

- **FutureID Project (www.futureid.eu)**

  The EU-funded FutureID project builds a comprehensive, flexible, privacy-aware and ubiquitously usable identity management infra-structure for Europe, which integrates existing eID technology and trust infrastructures, emerging federated identity

management services and modern credential technologies to provide a user-centric system for the trustworthy and accountable management of identity claims.

- **Open eCard Team (www.openecard.org)**

  The Open eCard Team is an open community, which aims at providing an open source and cross platform implementation of the eCard-API-Framework (BSI-TR-03112) and related international standards such as ISO/IEC 24727 and OASIS DSS through which arbitrary applications can utilize authentication and signatures with arbitrary smart cards.

- **BITKOM (www.bitkom.org/)**

  The Federal Association for Information Technology, Telecommunications and New Media (BITKOM) is the voice of the information technology, telecommunications and new media industry in Germany. BITKOM represents more than 1,700 companies, of which 1,200 are direct members and many of them are involved in working groups focusing on Identity Management, eID technology, Open Source Software and Cloud Computing for example.

- **CAST Forum (www.cast-forum.de)**

  The Competence Center for Applied Security Technology, (CAST) e.V. offers a variety of services in the field of secure modern information technology and is a contact for all questions regarding IT security.

- **Deutsche Wolke (www.deutsche-wolke.de)**

  The open source initiative "Deutsche Wolke" has been established as a network of renowned German and international organisations, which aims at establishing a federal cloud infrastructure for Germany.

- **European Association for eIdentity and Security (EEMA) – (www.eema.org)**

  For 25 years, EEMA has been Europe's leading independent, non-profit e-Identity & Security association, working with its European members, governmental bodies, standards organisations and interoperability initiatives throughout Europe to further e-Business and legislation.

- **Open Cloud Initiative (OCI) (www.opencloudinitiative.org/)**

  The Open Cloud Initiative (OCI) is a non-profit organization, which has been initiated to advocate open cloud computing. For this purpose it maintains a set of Open Cloud Principles (OCP) and uses them to determine whether a given product or service is compliant and therefore "Open Cloud", both by way of community consensus.

- **OpenID Foundation (www.openid.net)**

  The OpenID Foundation is an international non-profit organization of individuals and companies committed to enabling, promoting and protecting OpenID technologies. Formed in 2007, the foundation serves as a public trust organization representing the open community of developers, vendors, and users. OIDF assists the community by providing needed infrastructure and help in promoting and supporting expanded adoption of OpenID. This entails managing intellectual property and

brand marks as well as fostering viral growth and global participation in the prolif-eration of OpenID.

- **Open Identity Exchange (OIX) – (www.openidentityexchange.org)**

  The Open Identity Exchange (OIX) is a non-profit organization comprised of lead-ers from identity data-centric industry sectors including the internet (Google, Pay-Pal, etc.), data aggregation (Equifax, Experian, LexisNexis, etc.), and telecommuni-cations (AT&T, Verizon, etc.) driving the expansion of existing services and the deployment of new services. OIX develops and certifies trust frameworks, pre-negotiated sets of business, legal, and technical policies that provide identity service providers and relying parties with mutual assurance that their online transactions are reliable and repeatable. OIX is a global center of excellence for identity in next gen-eration digital transactions delivering domain expertise, joint research commis-sioned by competitors, and pilot projects to test assumptions in the real world.

- **Open Source Business Alliance (OSBA) (www.osb-alliance.de)**

  The Open Source Business Alliance – short OSB Alliance – is with more than 190 members Germany's largest network of enterprises and organisations, which devel-op or use open source software.

- **Open Source Integration Initiative (OSII) (www.osi-initiative.com/)**

  The Open Source Integration Initiative (OSII) brings together a range of open source software applications for use by businesses. It's an initiative by MFG Baden-Württemberg — Innovation Agency for ICT and Media — and the Open Source Business Alliance (OSB Alliance). The aim of OSII is to create a low-cost modular solution — a software stack — that meets the needs of many different operating processes.

- **SkIDentity Project (www.skidentity.de)**

  The SkIDentity Project aims at facilitating the use of electronic identity cards (eID) within existing and emerging cloud computing infrastructures in order to provide trustworthy identities for the cloud.

- **TeleTrusT – IT Security Association Germany (www.teletrust.de)**

  TeleTrusT is a widespread competence network for IT security comprising members from industry, administration, research as well as national and international partner organizations with similar objectives.

- **Trusted Cloud Program (www.trusted-cloud.de)**

  The Trusted Cloud Program is an initiative of the German Federal Ministry of Eco-nomics and Technology in which 38 companies and 26 academic institutions are collaborating in 14 projects in order to develop innovative, secure and legally valid technologies for trustworthy Cloud Computing.

Table of Contents

Open Identity Summit 2013 – Regular Research Papers

Open Identity Summit 2013 – Further Conference Contributions

Mobile Devices as Secure eID Reader using Trusted Execution Environments

Maximilian Stein

secunet Security Networks AG
Alt-Moabit 91c
10559 Berlin
Maximilian.Stein@secunet.com

Abstract: This work presents a prototype implementation of a smartphone as secure eID reader using NFC technology. The presented approach aims to reach a security level close to standalone smart card readers. This security level will be allowed by the means of a trusted execution environment (TEE) which allows strong isolation and separation for critical applications and provides trusted, not interceptable user input and output. The prototype supports the German eID (nPA) and follows the relevant guidelines.

1   Introduction

Mobile internet devices (smartphones, tablets) have become the omnipresent companion in the modern society. The capabilities and processing power of today's devices is enormous. Especially high-end devices featuring quad-core CPUs and high definition graphics can compete easily with mid-range PC systems while being smaller, more energy-efficient and portable. They can satisfy nearly all needs of an ordinary PC user like access to the internet, e-mail and social networks, music and video playback or other entertainment. In the long run mobile internet devices may replace the PC for such users and use cases completely.

National electronic ID cards are emerging slowly but surely. In Germany there are yet few citizens using their eID in online processes and there are still not many applications available. Nevertheless electronic IDs are believed to become more important and will be essential in future citizenship. To gain more acceptance from citizens though, it is important to provide low-threshold access to technology and knowledge for the usage of electronic IDs. The necessity of an additional, expensive reader device to make use of eID cards is not likely going to raise acceptance.

Since a few years mobile internet devices feature near-filed-communication (NFC) technology. Among others the NFC specification is based on the standard for contactless smart cards [ISO11]. Therefore NFC-devices are technically enabled to communicate with proximity cards like national eIDs. By this they can be used as card reader for eIDs and other smart cards. This has already been implemented for the German eID in [Hor11]

and [Mor12]. Both showed a proof of concept that the eID can be accessed using the PACE protocol through a NFC-enabled mobile phone (Nokia 6212 & Openmoko Neo FreeRunner customized smartphone). The security established through PACE depends on the secure input of the PIN on the device. The security in both approaches relies only on the assumption that the used mobile devices are trustworthy and no user input can be intercepted by a software of an attacker. However, since mobile devices gained popularity, more sophisticated attacks and malware for such devices emerged. For this reason smartphones and tablets have to be regarded as potentially untrustworthy and malicious.

The present work presents an approach to use a NFC-enabled mobile internet device as secure embedded reader for the German eID card by using a trusted execution environment (TEE). The remainder of this work is structured as follows. Section 2 briefly presents related work. In section 3 the basic principle of a trusted execution environment is described. Section 4 presents the current embedded smart phone reader implementation. Finally section 5 concludes this paper.

2   Related Work

Horsch [Hor11] implemented the eID application MONA as Java MIDlet on a Nokia 6212. It is capable of performing an online authentication with the German eID card. In [Mor12] Morgner implemented an embedded eID reader with PACE support on an Openmoko Neo FreeRunner customized smartphone with SHR Linux operating system. The implementation relies on OpenPACE [MO12], an open source implementation of PACE based on OpenSSL. An efficient implementation of the PACE protocol for mobile devices has been proposed in [WHB$^+$11]. Alternative solutions for the security concerns regarding the mobile use of eID and eSignature were proposed in [BHW12] and [BHWH11], respectively. An open source eID application for the German eID for Android devices is available through the Open eCard project [Ope12]. This app supports multiple methods to access the eID card. It is possible to use an external reader or the internal NFC interface, if available. The Governikus Autent PINApp [bre12] provides PIN management functionalities on Android devices for the German eID. The NFC Tag Info app [Hag13] is capable of reading electronic passports (eMRTDs) via basic access control (BAC) but does not provide PACE capabilities.

3   Trusted Execution Environment

A trusted execution environment (TEE) is a separate execution environment that runs alongside the Rich OS (i.e. regular mobile device OS). The TEE provides security services for the rich environment and isolates access to its hardware and software security resources from the Rich OS and its applications [Glo11].

Figure 1 depicts the TEE architecture as envisioned by the GlobalPlatform industry forum. It was designed for mobile and embedded devices but could be used for PCs as well if

all requirements are met. The three depicted TEE APIs in figure 1 were specified by GlobalPlatform in 2010 and 2011, respectively.

Figure 1: Architecture of the TEE as specified by GlobalPlatform

The security of the TEE relies on hardware extensions that help isolating the two environments. The hardware architecture to enable TEEs on ARM based mobile devices is the ARM TrustZone technology [ARM02][1]. The TrustZone extensions are integrated into the SoC and comprise an address space controller, memory adaptor, interrupt controller, reserved secure cache, and hardware keys. These features are available for ARM Cortex-A processors if and only if they were included by the SoC manufacturer. Secure operating systems can be implemented on top of the TrustZone hardware[2]. The physical CPU is presented as two virtual CPUs to the secure OS via TrustZone. One CPU dedicated to the rich environment and the other one to the trusted environment. In TrustZone terminology the rich OS application environment is referred to as Normal World (NWd) and the TEE as Secure World (SWd). The secure operating system controls the virtualisation, the security extensions, and provides the TEE.

A TEE can host multiple trusted applications. These applications are executed in the trusted environment where the TEE guarantees isolated execution between different trusted applications, protection, confidentiality and integrity. Trusted application binaries are stored in the file system as cryptogram and they are verified by the TEE each time before their execution. The root of trust for the TEE is established at boot time through a chain of trust: a hardware anchor verifies the boot loader which in turn verifies the TEE loader which verifies the TEE system image.

Current TEEs based on ARM TrustZone were MobiCore by Giesecke & Devrient and Trusted Foundations by Trusted Logic Mobile. However, both decided to merge their products in a joint venture named Trustonic[3] together with ARM. The TEE developed

---

[1] Other hardware architectures with similar features are for example Aegis, XOM, and SecureCore.

[2] TEEs can be implemented through pure software virtualisation as well (XenARM, SIVARM), but lack the additional security through hardware support.

[3] http://www.trustonic.com/about-us/who-we-are

by Trustonic is called <t-base. Sierraware implemented the SierraTEE and SierraVisor TEE solution which is freely available under the GNU GPL v2 license for the Samsung Exynos 4412 and nVIDIA Tegra 3 SoCs. So far Giesecke & Devrient's MobiCore was integrated in the Samsung Galaxy S3 and the Galaxy Note II. Since Samsung is hardware integrator and device maker partner of Trustonic it can be expected that Samsung will integrate <t-base in upcoming high-end devices too.

4   Implementation

The prototype device used for the implementation is a Samsung Galaxy S3 (GT-i9300) NFC-enabled smartphone running Android 4.1.1. The device combines all necessary components for the use of an eID card in one entity. Table 1 shows the analogy of components in the mobile eID reader system.

Original Component	Counterpart in Mobile Scenario
Host Computer	GT-i9300 NWd with Android
eID Application	Android App (e.g. Open eCard)
eID Reader Hardware	GT-i9300 SWd virtual CPU
eID Reader Firmware	Trusted Application (Trustlet)
eID Reader Driver	Trustlet Connector

Table 1: Analogy of components in the embedded eID reader system

A regular smartphone has the same capabilities of using an eID securely as a regular PC. It requires a smart card reader with a PIN pad, that is connected via the USB interface, and it needs to run an eID application. The already mentioned Open eCard project provides such an open source eID application for the Android OS. It is capable of using external smart card readers via USB and the internal NFC interface.

The here described eID reader implementation is an embedded smart card reader, which consists of a firmware part and a driver part. So far, this is identical to regular standalone eID readers. The difference is, that the firmware of a standalone reader resides inside the reader hardware. As shown in table 1, the reader hardware in this approach is physically the same as the host computers hardware, therefore it is called an embedded eID reader.

The smartphone is split up into two virtual devices by the TEE. The eID application resides in the so called normal world with the Android OS. The embedded reader firmware resides in the secure world. By this, the embedded reader can be treated as if it had its own separated hardware. The implemented prototype can be categorised as seen in table 2. The depicted categorisation for the reader categories Cat-S and Cat-C is taken from [BSI13] and shows the properties an eID reader has to have to be categorized as *standard reader* (Cat-S) or *comfort reader* (Cat-C). The prototype currently implements a *standard reader* with an additional display. So it can be categorized as Cat-S with additional functionality. Regardless of certification issues, the prototype can be enhanced in future to implement

	Cat-S	Cat-S$^+$	Cat-C
Interface to the host computer	✓	✓	✓
Contactless interface according to ISO/IEC 14443	✓	✓	✓
Contact interface according to ISO/IEC 7816			✓
PIN pad (secure PIN entry) with PACE support	✓	✓	✓
Display (2x16 alpha-numeric characters)		✓	✓
Qualified signature with contact cards			✓
Qualified signature with contactless cards (e.g. identity card)		✓	✓
Firmware update	✓	✓	✓

Table 2: Overview of Smart Card Reader Categories (source: [BSI13])

the properties of a signature terminal. In this way it implements the same properties as a *comfort reader* only without a contact interface. As it is unlikely that smartphones will be equipped with contact interfaces for smart cards, embedded readers like the presented prototype will only be capable to implement the properties that are presented here as Cat-S$^+$.[4]

The system architecture of the embedded reader and the associated components is depicted in figure 2. The shown eID reader **Trustlet** represents the reader firmware. The

Figure 2: Architecture for a device using the eID Reader Trustlet

reader driver is implemented through the so called **Trustlet Connector**. The Trustlet Connector implements the PC/SC IFD handler interface through which it can be accessed by any application, that is PC/SC aware. Generally the Trustlet Connector provides the security services of the Trustlet to regular applications in the NWd. Any information that is proccessed inside the Trustlet will not be accessible for NWd applications unless provided through the Trustlet Connector. The PIN Pad API, the Secure Display API, and the Cryptographic API are provided by the TEE itself. The PIN Pad and Secure Display API together provide a trusted user interface which is immune to interception and manipulation by other software.

---

[4]Please note that no new reader category is proposed here. The name Cat-S$^+$ is only chosen to symbolize the enhanced functionality compared to a Cat-S reader.

4.1 Trustlet

The embedded eID reader firmware is implemented as trusted application – a so called **Trustlet** as depicted in figure 2. The TEE is capable to isolate the execution of Trustlets and grants access to a secure user interface. In this way it is guaranteed that: (i) all processing results and the execution of the Trustlet itself is safe to interception and manipulation by malicious software (neither Android Apps nor other Trustlets), and (ii) a PIN can be entered directly inside the protected environment.[5]

The Trustlet implements the PACE protocol by using the internal cryptographic API that is provided by the TEE. The NFC interface is not (yet) available for Trustlets due to lack of driver support by the TEE. However, the GlobalPlatform TEE roadmap shows that additional peripheral devices like the NFC interface will be included in future versions of the specification. Currently the APDUs are transmitted from the inside of the TEE to the eID card through the Trustlet Connector via the Android NFC API. However, all secret information and processes of the PACE protocol – the PIN, key material and key generation – are isolated inside the TEE. The APDUs are transmitted encrypted through secure messaging between the endpoints Trustlet and eID. The security of this solution relies on the security of the PACE protocol. This is because the interface between the Trustlet and the eID can be assumed to be as insecure as the air interface in a regular PACE establishment process. The encrypted APDUs are interceptable from the NWd, because they are forwarded by the Trustlet Connector and the Android NFC API. Since the PACE protocol for key-agreement has been proven to be secure [BFK09], the implemented transmission of APDUs via the normal world can be considered to be secure as well.

It is intended to implement extended access control (EAC v2) in the future to use the smart-phone as signature terminal for the German eID as well[6]. The mobile eID reader can reach a security level comparable to a physically separated standalone card reader device through the hardware backed detachment of NWd and SWd. A security and conformity certification according to the technical guidelines of the German Federal Office for Information Security [BSI13, BSI11] seems possible at the moment. However, this highly relies on the certifications for TrustZone hardware implementations and TEE systems, which is a future challenge.

4.2 Trustlet Connector

To access the eID Reader Trustlet from any application in the NWd, a counterpart is required – the so called **Trustlet Connector**. In the prototype implementation the Trustlet Connector is an Android app that bundles two native C/C++ libraries, the native NFC wrapper and the actual Trustlet Connector library.

The Trustlet Connector is a standalone Android application with access to the Android NFC API. This is necessary because the Android NFC API is only accessible from Android

---

[5]As of the writing of this paper the secure UI functionality is not yet available, see section 4.3
[6]The certification of such a solution poses a greater challenge than the actual implementation.

applications with the appropriate Android permissions. It can not be accessed directly by native C/C++ libraries or executables. The Android app "catches" the eID when it is placed on the device and provides access to it for the Trustlet Connector library. This is achieved by providing transmission methods to the native NFC wrapper library via JNI. This native wrapper provides a RPC server for the actual Trustlet Connector library. It should be mentioned, that the NFC interface can not be used to actively poll for contactless cards and to initialize connections manually. This functionality is encapsulated by the Android NFC framework. Therefore apps can only wait for the NFC API to notify them, once an eID – or some other NFC tag – is available. This implies, that it is impossible for the embedded reader to power down the NFC interface or reset the connection with a present card. The eID has to be moved from the device and then replaced manually to achieve this.

The Trustlet Connector library implements the PC/SC IFD handler interface. The IFD handler is the driver of the embedded reader for the PC/SC interface. The IFD handler is loaded by the PC/SC daemon, which makes the reader available to any PC/SC aware application. In a regular PC environment this would be enough to provide access to the embedded reader. But there exists no standard implementation of PC/SC for Android. Therefore applications that rely on PC/SC may include the PC/SC library and daemon locally. This is the case for the Open eCard App. In this case, the location of the driver for the embedded reader has to be provided to the PC/SC daemon by a configuration file, to make the reader available for the application. For the prototype, the corresponding class of the Open eCard App was customized and a configuration file for the embedded reader was added. The Trustlet Connector library resides in the data folder of the Trustlet Connector app, the configuration file only points to the location of the driver library.

The Trustlet Connector library further contains an interface for the communication with the Trustlet. This interface uses the systems TEE driver to communicate with the Trustlet. The functional interface between the Trustlet and its Trustlet Connector can be defined freely. Basically both components have access to a shared buffer and are able to notify each other, if the content of this buffer has changed. Through this buffer, a RPC interface was implemented to allow the Trustlet to execute specific functions of the Trustlet Connector and vice versa. It should be noted that the shared buffer is not protected in any special way, nor does it reside inside the TEE. Therefore no unencrypted secret information should be written to it.

In short the Trustlet Connector app works as follows. When a contactless card is placed on the NFC interface, the Trustlet Connector app will be notified by the Android event manager and can be chosen to handle the event by the user. It will check if the present card is a German eID and start a background service that will listen to the native NFC wrapper library. The NFC interface can now be accessed via the native wrapper RPC server. If a PC/SC daemon has already loaded the Trustlet Connector library, the IFD handler will be informed, that there is a card present at the NFC interface. A PC/SC aware application is now able to use the embedded eID reader.

The Trutlet Connector library is the hub of the implementation. It handles the function calls from the PC/SC interface, from and to the Trustlet and to the native NFC wrapper.

4.3 Discussion

It has to be noted that the security of the overall implementation is highly depending on the secure and correct implementation of the TEE and the proper implementation of the Trustlet itself. For example it is crucial to implement the interface between the Trustlet and the Trustlet Connector very carefully. Since the Trustlet Connector resides in the NWd, it could be replaced by a malicious Trustlet Connector which tries to read secret information from the Trustlet by manipulating pointer locations or input data.

The implemented reader was successfully tested to be usable by the Open eCard Android app as external reader via PC/SC. The integration of the secure reader into an eID application requires some effort as there is not yet a default smart card reader interface for the Android OS.[7] The integration via PC/SC allows the usage of the embedded reader on unrooted off-the-shelf devices, because no special access rights like for the USB interface are required for the application.

As of the writing of this paper the author is not aware of any off-the-shelf smartphones with support for extended length APDUs. This is the same for the prototype device. Therefore the prototype only allows PIN management functionalities for the German eID and is not capable to perform an online authentication.

As of the writing of this paper the trusted UI functionality is not yet available, but is expected to be ready soon. The prototype therefore contains workarounds. In its current form it can not be regarded as secure or trustworthy because the trusted user interface, especially the secure PIN entry is the main feature that prevents unauthorized access to the eID. As a workaround, for each of the two UI APIs an Android Activity was implemented to simulate a PIN Pad and a "Secure" Display, respectively. When the Trustlet normally would access the secure APIs, it currently calls the Trustlet Connector to start the appropriate Android Activity and to either return the entered PIN or display the given certificate holder information.

5 Conclusion & Future Work

The present work showed the possibilities of using mobile internet devices as trustworthy and secure card reader for eIDs. It further gave a short introduction to trusted execution environments for mobile platforms as specified by the GlobalPlatform industry forum. The general concept and advantages of a TEE were described. It was presented how an embedded eID reader was implemented on an unmodified (but rooted) Samsung Galaxy S3 using a TEE. With the given implementation it is possible to use the eID in a mobile scenario, meaning that it is accessed by applications residing on the mobile device itself. Subject of future work is the usage of the smartphone eID reader as external reader for PC systems. Furthermore the conformity and security certifications of the embedded reader implementation pose the next steps in this working field.

---

[7]An implementation of a smart card reader API has been achieved by the SEEK for Android project, but is not applicable due to system manufacturer restrictions

References

[ARM02] ARM Ltd. TrustZone® technology. `http://www.arm.com/products/processors/technologies/trustzone.php`, 2002.

[BFK09] Jens Bender, Marc Fischlin, and Dennis Kügler. Security Analysis of the PACE Key-Agreement Protocol. In Pierangela Samarati, Moti Yung, Fabio Martinelli, and ClaudioA. Ardagna, editors, *Information Security*, volume 5735 of *Lecture Notes in Computer Science*, pages 33–48. Springer Berlin Heidelberg, 2009.

[BHW12] Johannes Braun, Moritz Horsch, and Alexander Wiesmaier. iPIN and mTAN for Secure eID Applications. In MarkD. Ryan, Ben Smyth, and Guilin Wang, editors, *Information Security Practice and Experience*, volume 7232 of *Lecture Notes in Computer Science*, pages 259–276. Springer Berlin Heidelberg, 2012.

[BHWH11] Johannes Braun, Moritz Horsch, Alexander Wiesmaier, and Detlef Hühnlein. Mobile Authentisierung und Signatur. In Peter Schartner and Jrgen Taeger, editors, *D-A-CH Security 2011: Bestandsaufnahme, Konzepte, Anwendungen, Perspektiven*, pages 32–43. syssec Verlag, sep 2011.

[bre12] bremen online services GmbH & Co. KG. Governikus Autent PINApp. `https://play.google.com/store/apps/details?id=de.bos_bremen.android.autent.pinapp`, 2012.

[BSI11] BSI – Federal Office for Information Security. Technical Guideline BSI TR-03105 Part 5.2: Test plan for eID and eSign compliant eCard reader systems with EAC 2, 2011.

[BSI13] BSI – Federal Office for Information Security. Technical Guideline BSI TR-03119: Requirements for Smart Card Readers Supporting eID and eSign Based on Extended Access Control, 2013.

[Glo11] GlobalPlatform. The Trusted Execution Environment, White Paper. `http://www.globalplatform.org/documents/GlobalPlatform_TEE_White_Paper_Feb2011.pdf`, 2011.

[Hag13] NFC Research Lab Hagenberg. NFC TagInfo. `https://play.google.com/store/apps/details?id=at.mroland.android.apps.nfctaginfo`, 2013.

[Hor11] Moritz Horsch. Mobile Authentisierung mit dem neuen Personalausweis (MONA). Master's thesis, Technische Universität Darmstadt, Darmstadt, 2011.

[ISO11] ISO/IEC. ISO 14443: Identification cards – Contactless integrated circuit cards – Proximity cards, 2011.

[MO12] Frank Morgner and Dominik Oepen. OpenPACE: Crypto library for the PACE protocol. `http://openpace.sourceforge.net/`, 2012.

[Mor12] Frank Morgner. Mobiler Chipkartenleser für den neuen Personalausweis: Sicherheitsanalyse und Erweiterung des „Systems nPA". Diploma thesis, Humboldt-Universtität zu Berlin, Berlin, 2012.

[Ope12] Open eCard Project. `https://www.openecard.org`, 2012.

[WHB+11] Alex Wiesmaier, Moritz Horsch, Johannes Braun, Franziskus Kiefer, Detlef Hühnlein, Falko Strenzke, and Johannes Buchmann. An efficient mobile PACE implementation. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*, ASIACCS '11, pages 176–185, New York, NY, USA, 2011. ACM.

Using Trusted Execution Environments in Two-factor Authentication: comparing approaches

Roland van Rijswijk-Deij[1,2] and Erik Poll[1]

[1]Radboud University Nijmegen, The Netherlands
`{rijswijk,erikpoll}@cs.ru.nl`
[2]SURFnet bv, Utrecht, The Netherlands

Abstract: Classic two-factor authentication has been around for a long time and has enjoyed success in certain markets (such as the corporate and the banking environment). A reason for this success are the strong security properties, particularly where user interaction is concerned. These properties hinge on a security token being a physically separate device. This paper investigates whether *Trusted Execution Environments* (TEE) can be used to achieve a comparable level of security without the need to have a separate device. To do this, we introduce a model that shows the security properties of user interaction in two-factor authentication. The model is used to examine two TEE technologies, Intel's IPT and ARM TrustZone, revealing that, although it is possible to get close to classic two-factor authentication in terms of user interaction security, both technologies have distinct drawbacks. The model also clearly shows an open problem shared by many TEEs: how to prove to the user that they are dealing with a trusted application when trusted and untrusted applications share the same display.

Keywords: trusted execution environment, Intel Identity Protection Technology, IPT, ARM TrustZone, two-factor authentication

1 Introduction

Two-factor authentication, based on *"something the user knows"* and *"something the user has"*, is a mature technology that has been around for a long time[1]. Classic two-factor authentication technologies[2], based on one-time password or challenge/response algorithms, have favourable properties from a user interaction perspective. Figure 1 shows an abstract model for user interaction in classic two-factor authentication. It shows a security token on the left and the user's regular device (e.g. laptop, tablet, ...) on the right. The model clearly shows the strict physical separation between the trusted environment (token) and the untrusted environment (laptop, etc.). Whenever a user interacts with one of the two devices it is always clear whether they are dealing with a trusted device.

Smart cards (also often used as authentication tokens) are an exception to this model. Most cards lack a display and a means to input data[3]. This means that the user has no (or only

---

[1]For example, the first RSA SecurID token was introduced in 1987.
[2]For a comprehensive overview of two-factor authentication solutions we refer to [vRvD11].
[3]There are exceptions, e.g. NagraID cards http://www.nidsecurity.com/

Figure 1: User interaction in classic two-factor authentication

a weak) assurance about the integrity of the data displayed on their screen and about the confidentiality of data entered and sent to the card (e.g. their PIN).

One solution to this problem is to use smart card readers with an integrated keypad for PIN entry and a display. These provide more assurance that the user is interacting directly with the card. A downside is that this requires the reader to be a separate device; this is less attractive because of cost and users needing a reader everywhere they use their card. It also precludes using the readers commonly integrated in modern laptops and smart phones.

In this paper we discuss a different approach to user interaction in two-factor authentication: the use of a *Trusted Execution Environment* (TEE). We investigate if the security model of classic two-factor authentication can be approached for smart cards without the burden of requiring a separate trusted card reader with its own I/O. To do this, we explain what we mean by a Trusted Execution Environment in section 2 and introduce two examples, one from Intel and one from ARM. We then show abstract models for user interaction using these two approaches to a TEE. The paper ends with a comparison of these two approaches and the classic two-factor model and gives directions for future research.

Our contribution   We introduce a conceptual model for user interaction with Trusted Execution Environments, which we apply to two concrete TEE technologies (Intel IPT and ARM TrustZone). We show that the model enables us to reason about the security aspects of the interaction between the user and a TEE. The model also clearly illustrates the open problem of how the user can ascertain that they are really dealing with a trusted application on a display that is shared between trusted and untrusted applications.

2   Trusted Execution Environments

Many definitions for a TEE are influenced by the Trusted Computing Group's (TCG) point of view, as the TCG-specified *Trusted Platform Module* (TPM)[4] is the most pervasive approach to trusted computing currently on the market.

---

[4]http://www.trustedcomputinggroup.org/developers/trusted_platform_module

Vasudevan et al. [VOZ⁺12] provide a more technology-neutral description, describing a set of features that enable trusted execution[5]. These can be summarised as follows:

- **Isolated Execution** – ensures applications execute completely isolated from and unhindered by others and guarantees that any code and data is protected at run-time.

- **Secure Storage** – protects persistently stored data (e.g. cryptographic keys) belonging to a certain application from being accessed by other applications.

- **Remote Attestation** – enables remote parties to ascertain they are dealing with a particular trusted application on a particular TEE.

- **Secure Provisioning**[6] – enables communication by remote parties with a specific application on a specific TEE while protecting integrity and confidentiality.

- **Trusted Path**[6] – a channel for the user to input data to the TEE and for the TEE to output data to the user; the channel protects against eavesdropping and tampering.

The remainder of this section examines two TEE technologies, Intel's *ITP* and ARM's *TrustZone*.

2.1 Intel Identity Protection Technology (IPT)

It is hard to find technical documentation about IPT. The only public documentation consists of marketing materials and high-level white papers [Int12, Car12, Smi11]. Careful reading of these, however, paints a picture of what IPT is. Intel markets IPT as a number of applications; we describe these below based on Intel's documentation.

One-time Passwords (OTP)   The IPT OTP application resembles OTP tokens sold by vendors such as RSA (SecurID) and Vasco (DigiPass). Intel provides a basic implementation based on the OATH time-based OTP algorithm [MMPR11]. Several vendors of classic OTP solutions have also ported their OTP algorithms to IPT (see [Int12], p. 8).

PKI   In [Int12] Intel claims that the PKI application[7] introduces hardware protection for RSA keys. The IPT PKI application integrates with Windows applications using a Cryptographic Service Provider (CSP) provided by Intel for Microsoft's CryptoAPI. This is similar to how PKI-enabled smart cards are usually integrated in Windows applications.

---

[5]Note also that a TEE is much more than just a TPM, which would fulfill only some of the features listed.
[6]Note: secure provisioning is I/O with a *remote* party, a trusted path is *local* secure I/O with the user.
[7]Intel sometimes refers to IPT with PKI as *Platform Embedded Asymmetric Token* (PEAT) (e.g. [Smi11]).

Protected Transaction Display (PTD)   PTD is not really an application but rather a feature that supports IPT applications. In documentation Intel describes how this feature can be used to secure PIN entry by the user. The "How It Works" video on Intel's website also shows PTD being used for confirming transactions (e.g. of a bank transfer).

NFC   Intel also includes NFC as one of the technologies under the IPT umbrella, but insufficient information is available for us to make any claims about NFC and its relation to IPT, so we have chosen to ignore it in our discussion.

2.1.1   Architecture

Figure 2: IPT abstract architecture (for a detailed explanation see §2.1.1)

Figure 2 shows an abstract architecture of IPT. It shows the different components identified in Intel's documentation and what environment these components belong to. The paragraphs below provide more detail on each component. Notably absent in this architecture is a trusted path for user input, this is discussed in more detail in section 4.

Management Engine   The Management Engine ❶ (ME) appears to be the core of IPT. Based on the naming of the ME it is very likely that Intel re-uses the ME included in their Active Management Technology (AMT)[8]. Assuming this is the case, the ME runs on a separate CPU (an ARC4 RISC processor, shown as ❷ in Figure 2) that runs the Nucleos Real-time OS[9]. IPT applications run as applets ❸ on a Java VM ❹ inside the ME.

Secure Storage ❺   The OTP and PKI application rely on secure storage for key material. It proves difficult to determine if a single subsystem fulfills this function. For OTP Intel [Int12] mentions that one-time passwords are based on a machine-specific key generated by the Intel chipset, but there is no indication of how and where this key is stored. For PKI they [Car12] mention that keys are stored on the hard drive and are wrapped with - what

---

[8]A technology for remotely managing systems, for instance desktop systems in a large enterprise (http://en.wikipedia.org/wiki/Intel_Active_Management_Technology)

[9]http://www.mentor.com/embedded-software/nucleus/

Intel calls - a Platform Binding Key. All operations on keys then take place in hardware where the key is unwrapped before use. The documentation does not explicitly state this, but it seems likely that the underlying technology used for this is (similar to) a TPM.

Display ❺ It is unclear how the secure display feature integrates with the rest of the system. The examples [Int12, Car12] show that the untrusted OS "sees" black boxes where trusted content is rendered on the screen. This implies that IPT either relies on memory protection for the graphics frame buffer that prevents the untrusted OS from accessing protected parts of the frame buffer, or that the trusted environment has its own frame buffer that is overlaid on frame buffer data from the untrusted OS. It is highly likely that this feature only works with an integrated graphics processor that is part of the chipset.

IPT platform middleware ❼ Communication between applications running in the regular OS on the main CPU and IPT applications in the ME requires some sort of channel. Intel has middleware components that provide such a channel to applications.

Applications that run in the IPT ME can be installed at will. This requires a conduit for installing applications into the ME, a role also performed by the IPT platform middleware.

Attestation and secure provisioning A system with IPT can perform remote attestation to prove that the IPT implementation is genuine using the *Enhanced Privacy Identifier* (EPID) scheme [BL07]. IPT can also set up a mutually authenticated secure channel with the issuer of the attestation identity using Intel's SIGMA protocol [WL11]. This mutually authenticated secure channel can, for instance, be used for secure provisioning.

Developing for IPT As already mentioned, Intel works with independent software vendors to port their OTP solutions to IPT. This implies that there is a software development kit available for IPT. We inquired with Intel as to the availability of an SDK. Intel indicated that such an SDK exists, but that access to the SDK requires a contract with Intel.

IPT and TEE requirements Intel does not market IPT as a TEE. The architecture described above, however, when combined with the description of IPT applications and features in section 2.1, aligns well with the five requirements for TEEs introduced in section 2. Based on this we think that the underlying technology of IPT must be viewed as a TEE.

2.2 ARM TrustZone

ARM offers a technology platform that is similar in its applications to IPT, called TrustZone. Where IPT currently seems to be mostly geared towards use in PC or server class systems, ARM TrustZone is aimed at system-on-a-chip (SoC) architectures used in mobile devices such as smart phones and tablets. This section provides a high-level overview of TrustZone, mostly based on [ARM09].

2.2.1 Architecture

ARM specialises in providing designs for (parts of) so-called Systems-on-a-Chip (SoCs). This is reflected in the TrustZone architecture. The core of TrustZone is a "two worlds" paradigm, with a *normal world* and a *secure world*. This concept shows up all through the architecture. At the hardware level the two worlds are separated on the system bus. What is in effect a special 33rd address line on the bus determines whether bus transactions are part of either one of the worlds. Devices connected to the bus set this address line during a read or write action to indicate whether they are operating in the normal or the secure world. The bus mediates access from bus masters to slaves such that a secure master may access both secure as well as normal slaves whereas a normal master may only access normal slaves and will trigger a bus error if it attempts to access a secure slave.

ARM has also created extensions to its CPU cores called ARM Security Extensions. These allow a single CPU core to run both normal world software and secure world software. Figure 3 shows an abstract model of the Security Extensions. Switching between the two security worlds is managed by the *monitor*, a process that runs in the secure world. The monitor process can be entered by a number of triggers, either programmatically (by executing a special instruction) or by a number of hardware triggers such as interrupts.

Figure 3: Security Extensions abstract model

2.2.2 Software and TrustZone

ARM does not directly provide any software to execute in the secure world. Developers of systems based on ARM IP either have to develop their own solutions or can choose to use existing secure micro kernels like MobiCore from Trustonic[10]. Trustonic has recently certified that its secure μ-kernel implementation meets the Global Platform Trusted Execution Environment specifications[11,12].

There are also efforts to create open source secure μ-kernels that use the capabilities of TrustZone. Especially worthwhile are the efforts of IAIK (part of the TU Graz). In [Win08] they propose a framework for secure applications on top of TrustZone by executing a modified Linux kernel in the secure world. They also propose an open source development environment for TrustZone [WWPT12] and their own μ-kernel on top of a cheap development board with a Samsung SoC [Win12].

---

[10]http://www.trustonic.com/about-us/who-we-are/

[11]http://globalplatform.org/specificationsdevice.asp

[12]http://www.trustonic.com/news/release/trustonic-is-first-to-qualify-a-globalplatform-compliant-tee/en

2.2.3 TrustZone and TEE requirements

The list below revisits the requirements for a TEE from section 2 and examines how Trust-Zone meets these requirements and where additional effort by SoC designers is required:

- **Isolated Execution** – the ARM Security Extensions allow separation of a CPU core into a secure and a none secure world. That in itself is insufficient to provide isolated execution; a secure μ-kernel that supports isolated execution and a memory management unit in the SoC that supports memory protection are also required.

- **Secure Storage** – TrustZone does not include any means for secure storage. Adding something like a Secure Element or a TPM to the SoC design can address this.

- **Remote Attestation** – TrustZone does not provide remote attestation capabilities. This requirement can be fulfilled by introducing a *Mobile Trusted Module* (MTM) [EK07], implemented in hardware (SE/TPM) or in software (in the secure μ-kernel).

- **Secure Provisioning** – Again, this is not explicitly specified as a part of TrustZone, but would most likely be implemented in the secure world μ-kernel.

- **Trusted Path** – Establishing a trusted path is addressed explicitly in TrustZone. In section 3.2 of [ARM09] ARM explains how the bridge between the peripheral bus and the system bus can be used to secure interaction with peripherals like a keyboard. In the example system design in the same document ARM also makes suggestions how the same can be achieved for the display.

3 Related work

Much of the research into trusted execution focuses on aspects of TPMs and cryptographic means to support trusted execution (e.g. attestation). Specific references are not provided as it is easy to find entries into the large body of work around this topic.

Section 2 already references the work by Vasudevan et al. In addition to providing a good definition for a TEE, they argue that TEE facilities are mostly not available to application developers for various reasons, and give recommendations on how to improve this situation. Zhou et al. [ZGNM12] outline an approach for establishing a trusted I/O path between the user and an application on commodity x86 hardware by proposing modifications to the system's I/O architecture.

Finally, there are two implementations of authentication tokens that mimic the behaviour of a PKI-enabled smart card inside a TEE. Brasser et al. [BBFS12] demonstrate a token running on the user's PC on top of Intel TXT. Tamrakar et al. [TEL$^+$11] take a different approach and emulate a smart card on a smart phone that can interact with a PC as if it were a real smart card.

4 Models for secure user interaction using TEEs

In section 1 we introduced an abstract model for user interaction in classic two-factor authentication (Figure 1), which shows the clear, physical, separation between the trusted and the untrusted environment. In this section we construct similar models based on Intel IPT and ARM TrustZone as TEEs. The models clearly illustrate how IPT and TrustZone differ from the classic approach and also highlight the common issue shared by any approach using a TEE: how to convince the user that they are interacting with a TEE. Note that we do not address securing communication between a TEE and a smart card; existing secure channel solutions provide sufficient means to achieve this.

Figure 4: Models for user interaction

4.1 Intel IPT

Based on the features Intel markets under the IPT umbrella (see section 2.1) we have constructed the model shown in Figure 4a. The model shows the trusted environment in gray, the untrusted environment (i.e. the normal OS) in white and components that are in a sense part of both worlds in interleaved gray and white.

The model clearly shows the weakest link in the chain when using IPT: user input does not flow through a trusted path. This is best illustrated by how Intel implements its Protected Transaction Display feature. For PIN entry, the software running in the trusted environment randomises the layout of the PIN entry pad. This is done to prevent applications running in the regular operating system from recording mouse clicks to steal the PIN.

The display at the top of the model is shaded to indicate that it contains content from both

the trusted as well as the untrusted environment. We assume that merging of secure and non-secure elements on the display takes place under supervision of the secure environment (although this is not explicitly stated in the available Intel documentation).

4.2 ARM TrustZone

Figure 4b shows a similar model for ARM TrustZone. Because ARM TrustZone is a set of building blocks and not a stand-alone technology, we have made assumptions (reflecting the most desirable situation that can be created using TrustZone) about the specific configuration, namely

- there is a trusted path to the display, e.g. as suggested in section 3.2 of [ARM09];

- all user input goes through a TrustZone-aware peripheral bus;

- there is a *Memory Management Unit* (MMU) that supports protected memory separation between the secure and normal world.

Under these assumptions the model shows that a fully trusted path can be created all the way from user input to output on the display. The model reflects that there may be multiple implementation options for a trusted display; the display may show either content exclusively from the secure world or the normal world (indicated by "switch" in the model), or it may show a mix of the two just like Intel IPT (indicated by "merge" in the model).

4.3 Local attestation

The models highlight that IPT and Trust-Zone share a common issue: the display is used for communication by both the trusted and the untrusted environment. This makes it hard for users to ascertain whether they are dealing with a trusted application or not. In fact, all trusted execution environments that allow direct user interaction have this problem.

To remedy this situation the trusted environment will need to provide some form of proof to the user that the data displayed belongs to the TEE and can be trusted. Section 2 mentions remote attestation (prov-

Figure 5: Local versus remote attestation

ing to remote parties they are dealing with a genuine application and TEE). In keeping with this naming we will call proving trustability to the local user *local attestation*. Figure 5 shows the relation between local and remote attestation.
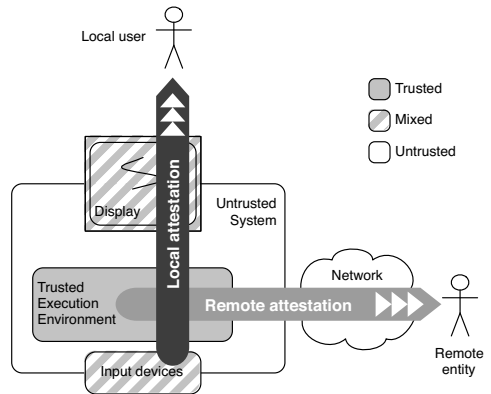
There are a number of approaches to implementing local attestation. One approach is to set the colour of the title bar of application windows such that all the windows belonging to a single application have the same colour (this approach is taken by Qubes OS[13]). The colour is set by a secure layer of the OS. This approach, however, does not stop malicious applications from spawning windows with content similar to a trusted application. Another approach is *personalisation* of the trusted environment with something specific to the user (e.g. a picture of their family). This personal item is then shown every time the TEE uses the display. The problem with this approach is that it is vulnerable to phishing. The user can, for instance, be tricked into thinking they are reconfiguring their trusted environment and unwittingly submit their personal item to a malicious application. There are also proposals for using a separate trusted device that the user can rely on to perform local attestation of a TEE (e.g. [Toe09, MPSvD07]). Finally, a truly convincing solution is using a hardware indicator on the device that shows the status of the TEE. An example could be an LED that only lights up when the TEE is active. Texas Instruments has submitted a patent application for this [CD02]. Note that this only works well if the entire display is controlled by the TEE.

Neither IPT nor TrustZone provide a clear way to perform local attestation. The examples in Intel's documentation seem to indicate that they hope to achieve this with consistent branding; from a security perspective that has no use, though, since it is trivial for an attacker to observe this branding and to falsify it. TrustZone itself does not address local attestation, but online demonstration videos suggest that Trustonic's MobiCore supports personalisation.

5   Conclusions and future work

Classic two-factor authentication has very desirable security properties but also has practical problems. Users may forget their security token or may lack the infrastructure to use their token (for instance when the token is a smart card that requires a reader). Zooming in on smart cards we already outlined that their security properties are less favourable since they commonly lack a secure display and trusted input device.

We wanted to examine if Trusted Execution Environments can provide secure user interaction similar to classic solutions. It would be particularly interesting if TEEs can also be used to secure interaction with a smart card (given the less favourable properties of a smart card when compared to classic security tokens). To illustrate this we introduced an abstract model for user interaction. We described two TEE technologies (from Intel and ARM) and applied the same abstract model to these two TEE technologies. When we look at how the models for these TEEs compare to the classic model we can conclude that they can approach the classic model up to a certain extent. They do, however, both have significant drawbacks when compared to the classic model. Intel IPT has a serious issue where there is no trusted input path for the user to enter data. ARM TrustZone requires careful selection of the right components by the system-on-a-chip designer that puts the parts of

---

[13]http://qubes-os.org/

the TEE together to guarantee that it can be trusted. An added disadvantage of TrustZone is that - unlike IPT - it does not come with a dedicated software implementation, further complicating the choices for designers of a TrustZone-based TEE. Finally, both technologies share a common issue, which is how to prove to the user that they are dealing with a trusted application.

It is clear then that these technologies cannot provide a drop-in replacement for classic two factor authentication solutions. This does not mean they do not have their benefits. The convenience of a built-in two-factor authentication solution, such as e.g. Intel IPT can offer, makes it much easier to deploy the solution, thus lowering the threshold for using something that is more secure than the age-old username/password paradigm. Note that a TEE is effectively an embedded smart card, a fact that is capitalised upon by Intel IPT and by the two examples mentioned in section 3. Furthermore TEEs could be leveraged to secure interaction with the user when using smart cards, thus improving the security properties of smart cards when used as a two-factor authentication token. This would also mean that no special secure card reader is required and the built-in smart card readers that appear in more-and-more laptops, tablets and smart phones can be used.

Finally, we note that it proved hard to find detailed public documentation about the specific technologies we investigated, particularly about Intel IPT. Although we feel that this did not impact the conclusions of our research unduly, this is worrisome from a security perspective; public scrutiny is essential for a good understanding and acceptance of these kinds of technologies.

Future work   A consortium of partners[14] is currently working on a privacy-friendly authentication technology implemented on smart cards called IRMA[15]. One of the open issues in the project is secure user interaction (both for showing and confirming transaction details and for secure PIN entry). We would like to investigate if a TEE can help solve this issue, which motivated the current paper.

Another question for future research concerns the problem described in Section 4.3: what are alternatives for the personalisation approach that are less likely to be phished?

Finally, it would be worthwhile to investigate and compare the size of the *Trusted Computing Base* (TCB) for IPT and TrustZone-based TEEs, as their security to a large extent depends on the size of the TCB.

References

[ARM09]   ARM Ltd. ARM Security Technology - Building a Secure System using TrustZone Technology, 2009.

[BBFS12]   F.F. Brasser, S. Bugiel, A. Filyanov, and A. Sadeghi. Softer Smartcards - Usable Cryptographic Tokens with Secure Execution. In *Financial Cryptography and Data Security*, vol. 7397 of *LNCS*, pp 329–343. Springer, 2012.

---

[14]TNO (http://www.tno.nl), SURFnet (http://www.surfnet.nl) and SIDN (http://www.sidn.nl)

[15]https://www.irmacard.org/

[BL07]   E. Brickell and J. Li. Enhanced Privacy ID: A Direct Anonymous Attestation Scheme with Enhanced Revocation Capabilities. *IEEE Transactions On Dependable And Secure Computing*, 9(3):21–30, 2007.

[Car12]   P. Carbin. Intel Identity Protection Technology with PKI (Intel IPT with PKI) Technology Overview, 2012.

[CD02]   B. Cornillault and F. Dahan. Secure Mode Indicator for Smart Phone or PDA, 2002.

[EK07]   JE Ekberg and M. Kylänpää. Mobile Trusted Module (MTM) - an introduction. Technical report, Nokia, 2007.

[Int12]   Intel. Deeper Levels of Security with Intel Identity Protection Technology, 2012.

[MMPR11]   D. M'Raihi, S. Machani, M. Pei, and J. Rydell. RFC 6238 - TOTP: Time-based One-Time Password Algorithm, 2011.

[MPSvD07] J.M. McCune, A. Perrig, A. Seshadri, and L. van Doorn. Turtles all the way down: Research challenges in user-based attestation. In *Proceedings of HotSec*. USENIX Association, 2007.

[Smi11]   N. Smith. Identity Protection Technology (presentation). In *2011 Kerberos Conference*, Cambridge, MA, 2011. Intel.

[TEL$^+$11] S. Tamrakar, JE Ekberg, P. Laitinen, N Asokan, and T Aura. Can hand-held computers still be better smart cards? In *INTRUST 2010*, vol. 6802 of *LNCS*, pp 200–218. Springer, 2011.

[Toe09]   R. Toegl. Tagging the turtle: local attestation for kiosk computing. In *Advances in Information Security and Assurance*, vol. 5576 of *LNCS*, pp 60–69. Springer, 2009.

[VOZ$^+$12] A. Vasudevan, E. Owusu, Z. Zhou, J. Newsome, and J.M. McCune. Trustworthy Execution on Mobile Devices: What security properties can my mobile platform give me? In *Trust and Trustworthy Computing*, vol. 7344 of *LNCS*, pp 159–178. Springer, 2012.

[vRvD11]   R.M. van Rijswijk and J. van Dijk. tiqr : a novel take on two-factor authentication. In *Proceedings of LISA '11: 25th Large Installation System Administration Conference*, pp 81–97, Boston, MA, 2011. USENIX Association.

[Win08]   J. Winter. Trusted computing building blocks for embedded linux-based ARM trustzone platforms. In *Proceedings of the 3rd ACM workshop on Scalable trusted computing - STC '08*, pp 21–30. ACM Press, 2008.

[Win12]   J. Winter. Experimenting with ARM TrustZone – Or: How I Met Friendly Piece of Trusted Hardware. In *2012 IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications*, pp 1161–1166. IEEE, 2012.

[WL11]   J. Walker and J. Li. Key Exchange with Anonymous Authentication Using DAA-SIGMA Protocol. In *INTRUST 2010*, vol. 6802 of *LNCS*, pp 108–127. Springer, 2011.

[WWPT12]   J. Winter, P. Wiegele, M. Pirker, and R. Toegl. A Flexible Software Development and Emulation Framework for ARM TrustZone. In *Trusted Systems*, vol. 7222 of *LNCS*, pp 1–15. Springer, 2012.

[ZGNM12]   Z. Zhou, V.D. Gligor, J. Newsome, and J.M. McCune. Building Verifiable Trusted Path on Commodity x86 Computers. In *2012 IEEE Symposium on Security and Privacy*, pp 616–630. IEEE, 2012.

Unlinkability Support in a Decentralised, Multiple-identity Social Network

Simon Thiel[1,] Fabian Hermann[1], Marcel Heupel[2], Mohamed Bourimi[2]

[1]Fraunhofer IAO
Nobelstraße 12
70569 Stuttgart, Germany
surname.name@iao.fraunhofer.de

[2]Universität Siegen
Hölderlinstr. 3
57067 Siegen, Germany
name@wiwi.uni-siegen.de

Abstract: Providing support for unlinkability in a decentralized, multiple-identity social network is a complex task, which requires concepts and solutions on the technical as well as on the user-interface level. Reflecting these diverse levels of an application, this paper presents three scenarios to impede the linkability of multiple identities in decentralized social networking. Solutions cover a communication infrastructure which allows referencing to multiple identities; analysis of user content and sharing history to present linkability warnings; and user interface means allow for a privacy-ensuring management of partial identities. The di.me userware research prototype of the EU FP7 funded digital.me (di.me) is introduced to show the integration of the solutions accordingly.

Introduction

Social networking and personal information management are closely connected fields for end-user applications: personal content, and information describing the person, is organised within diverse applications and services, and shared or disclosed when communicating and collaborating with others. However, in the Social Web, privacy and data protection are an issue often debated and criticised by data protection commissioners and citizens. Several technological trends and initiatives aim to empower users to have more control over personal data, e.g. emerging implementations of *decentralised social networks* [YL09]. Diverse initiatives base on the approach to provide personal servers as collection centres for the user's data (e.g. FreedomBox[1], Friendica[2], Cunity[3], VRM[4]).

---

[1] http://freedomboxfoundation.org
[2] http://friendica.com/
[3] http://www.cunity.net/
[4] http://cyber.law.harvard.edu/projectvrm/Main_Page

The di.me platform for decentralised social networking

The European funded project digital.me[5] (di.me) adopts the approach of decentralisation for social networking.

The research prototype "di.me userware", developed as the major outcome of the project and published as open source,[6] provides decentralised social networking and privacy-enhanced social functionalities based on a semantic model as its key features. A running prototype environment is currently hosted by the di.me consortium, and tested within the evaluation phase of the project.[7]

Figure 1: di.me userware network with di.me servers
hosting 1 or up to N Personal Services (PSs) for users

The di.me userware is subdivided into three packages: (1) the di.me server, containing the main functionality, (2) the di.me client, providing user interfaces (UI) for accessing the di.me server and (3) the di.me cloud, incorporating the necessary infrastructure required for setting up a decentralized network of di.me servers.

The decentralized approach is realized as a network of di.me server nodes communicating via HTTPS REST interface (REST-API). Each server is able to transparently host *Personal Services (PS)* for a number of users (Figure 1). The di.me PS represents the virtual personal node of a user.

This flexible solution allows for different hosting setups: a user may either use a single PS on a self-hosted server or apply for a user account as a tenant on a server provided by a trusted third party.

---

[5] http://www.dime-project.eu
[6] https://github.com/dime-project/meta
[7] For the evaluation environment see http://www.dime-project.eu

On the application level, the PS provides networking functionality. Messaging, document and profile sharing is supported between di.me PSs and also by use of external communication channels, e.g. by sending messages to twitter. Personal information from other sources (e.g. LinkedIn, Facebook, etc.) can be integrated by service adapters and semantic data representation standards. Based on this, the di.me userware provides pro-active functionalities [SC12] to the user, such as trust warnings, merge recommendations and situation detection. The case of recommendations to avoid content-based linkability is presented below in this paper.

At an architectural level, the di.me userware has been built upon a multi-layered approach native to dynamic web applications, providing a decoupled component schema that benefits future scalability requirements. For secure information management, the reference implementation provides a rich subset of semantic models, in compound with access control, processing components (such as the Sesame framework), and higher-level access APIs.

Like for PS-to-PS communication, also the UI clients connect to the PS accessing the di.me server's REST-API. The API establishes a generic way to access the user's PS. This supports clients with different functionality scope and running on various operating systems. Within the scope of the digital.me project a web-client and a client for Android mobile phones have been developed and are included in the OS publication.

Multiple partial identities

Many (centralised and decentralised) social networks offer advanced privacy settings, allowing for the filtering of information shown to contacts. With these settings – often at the level of the UI – the users can adjust which parts of their identity information are shown to others. While this can be considered as support of partial information sets linked to a root identity, di.me supports partial identities [PH10] which are potentially fully distinct information sets that can be shown to communication partners. In the process of sharing information, a partial identity can be used to control the specific set of personal data to be shared with individual contacts or groups. Such identities might become linkable to each other or to the person in real-life and could therefore possibly threaten users' privacy by revealing more information than intended. Consequently, the provision of *unlinkability* support is an essential feature for the di.me userware. Following the definition of [PH10], we define unlinkability of two items as the inability of an attacker to decide if they are related or not. In the case of di.me, such items are e.g. profiles which are indirectly representing an identity of a di.me user, and their attributes, in particular unique attributes like e.g. email address.

Unlinkability support in di.me

The following sections describe the unlinkability support in di.me, based on multiple identities. We focus on the following three different scenarios where identity linkability may be impeded with different technical means:

(1) Scenario 1: Linking (partial) identities to each other or to the person's root identity[8] by analysing the technical communication protocol (e.g. IP address discloses a physical location). This reduces the anonymity possibly leading to the revelation of the "real-life identity".

(2) Scenario 2: Linking a digital (partial) identity to the person's root identity because of disclosed information by the person him-/herself or by others[9] (e.g. the user's real-life or email address, his/her current geo-location, etc.).

(3) Scenario 3: Linking different (partial) identities to each other, because the same information is contained or shared via those identities (e.g. the same document shared under two pseudonyms).

The next section describes the requirements background of di.me with focus on the threats analysis. Based on this, the following sections detail di.me concepts and solutions for the scenarios of unintended linking of identities. Di.me combines solutions on the network and application communication level, as well as user recommendations and UI means to support the user.

Requirements background and threats analysis

A further requirements-driven analysis mainly based on the comparison of existing social networking led to the identification of five high-level requirements categories (cf. [TB12]) candidate to demonstrate innovation: (Category 1; **C1**) Integrated Personal Information Management, (**C2**) Secure, Privacy-respecting Sharing of Personal Information, (**C3**) Intelligent User Support with Context Sensitive Recommendations and Trust Advisory, (**C4**) Transparent Multi-Platform UI, and (**C5**) Integration of Existing Services. In the focus of this paper are the categories **C2** and **C3** by considering interdependencies to the other categories. The di.me open trust, privacy, and security infrastructure fulfill the major security goals, namely, *Confidentiality*, *Integrity*, and *Availability* also acronymised as the CIA Security Triangle. According to Santen in [S06], from the three major goals, confidentiality "*is one of the most practically difficult to achieve whereas integrity and availability can be achieved by means of standard software engineering techniques.*" Confidentiality goes beyond protecting the content of messages to protecting communication relationships in general (e.g. by means of anonymisation, i.e. pseudonimity), as this could reveal a lot about involved parties in such communications, i.e. the identities of senders and receivers allowing for different forms and degrees of linkability. Further, Santen also states that the interpretation of confidentiality (and the other protection goals) depends on application circumstances and scenarios to be supported: They can be classified by defining the attacker model within a threats analysis. Such model can be defined by answering the question "who may gain information and who must not". Thereby it is important for the stakeholders (i.e. all involved parties) to have some idea of who might assume the role of an attacker and

---

[8] For simplicity we use the concept of a root identity reflecting unique attributes of the respective person's real-life. In di.me the root identity is defined by the superset of the established partial identities.

[9] For instance a contact disclosing in a status update where s/he is and with whom by using real-life attributes („I m with Bob and Marry in Rome"). Such information could lead to linking these real-life attributes to used pseudonyms, e.g., if represented on a map functionality offered by the social network by using pseudonyms representing partial identities of Bob and Marry and their contact

what kind of behaviour (malicious or not) to expect from an attacker by performing a threats analysis. In this respect, correlations among protection goals and stakeholders (e.g. end-users, provider(s), and legislative)[10] could lead to conflicts, which is a classical multilateral security concern [R00]. As di.me supports multiple-identities it is crucial to integrate unlinkability support within. Linkability as non-functional requirements (NFRs) may conflict with other competing NFRs such as providing collaboration awareness[11] (in the UI) or negatively affecting user experience (in terms of performance penalties by using anonymity networks). The security requirements and threats analysis with respect to tasks of **C2** and **C3** was carried out by following the AFFINE methodology [BB10]. This enforces the early consideration of multilateral security requirements along with other (N)FRs by involving all stakeholders, negotiating, and aligning their potentially conflicting interests in the design[12] and development process.[13]

The requirements for our scenarios 1, 2, and 3 are addressed with a set of approaches solving specific linkability problems and implemented within di.me's open trust, privacy, and security (TPS) infrastructure. In the following, these solutions are summarized by showing how the TPS infrastructure enables di.me users to securely use and share personal data by considering respective threats analysis.

Avoiding linkability on a network and application communication level (scenario 1)

Derived from the first scenario of linking information items or persons, an important technical requirement is that the IP of the di.me server hosting the PS of a user must not be revealed as part of the communication protocol. Otherwise, the number of potential owners of an identity can be drastically reduced by the potentially low number of users hosted on a single di.me server.

Although the application of a PS is bound to a specific di.me server, the role of the server is transparent for the communication between two di.me PSs. Therefore, the information about the physical location of a PS is not required on this level of communication (e.g. when connecting for exchange of information, for sending of liveposts or for sharing). However, to hide a di.me server's location in the communication flow, when accessing a foreign PS, is a non-trivial task. Based on three main iterations, a solution for this has been developed for the di.me system (see also [BH12], [FH12], and [SB13] respectively):

---

[10] This is also the case in di.me since different parties from academia, research, and industrial fields are involved

[11] Social, group, and workspace awareness answering „who" is collaborating with „whom", „where", and „when"

[12] The solution's design process compromises consideration of an attacker model and threat analysis

[13] Santen begun the motivation of his work by citing from Viega and McGraw 2001 who stated that "*Bolting security onto an existing system is simply a bad idea. Security is not a feature you can add to a system at any time*". He further argues that "*the discipline of "Security Engineering" is far from mature today, and that, in practice, it still is not an integral part of the engineering processes for IT systems and software is based on the fact that security awareness results from reports on attacks – and not from the latest security feature that would make an application even more secure than it already was before*"

- Addressing of identities by use of a unique ID: at the level of the platform design as an internal reference to a specific identity was introduced, the Service-Account-ID (SAID).
- Hidden resolving of the SAID within a di.me proxy layer
- Concealing network communication flow by using the Tor[14] network

Figure 2: di.me proxy layer hides IP address of the PS providing a resource.

As alternative to an initial realisation relying on the Tor network, in which the attacker model does not trust any party, in the di.me environment [BH12], a di.me proxy layer, acting at the same time as an anonymity network has been designed (Figure 2). This flexible solution enables di.me to support end-users as well as business anonymity needs for (decentralised) social networking scenarios. It allows tailoring the anonymity degree according to needed privacy degree in the respective scenarios as a means for meeting multilateral security. A parallel usage of the Tor is still possible.

In order to actually contact another person's PS the corresponding SAID needs to be resolved. Therefore, a specific di.me DNS has been developed, allowing the translation the unique identifier to a network address, which can be either an IP, a forwarding proxy or a Tor-Onion address. As described in [FH12], it is also possible to be reachable over

---

[14] https://www.torproject.org

different channels at the same time. Therefore, it is possible to use the direct IP address for communication with close friends, while an anonymous/pseudonymous identity uses, for example, the Tor network.

A second important requirement is that IDs used internally as references to shared items must not be re-used when sharing via different identities. Otherwise, a receiver that is receiving shared items from two apparently independent senders is able to link these identities by comparing the IDs, even when the shared content is not unique. To resolve this, a concept of masquerading of shared internal IDs was developed: IDs used internally are mapped to anonymised IDs used for sharing as a single identity only. For further communication between PSs and the used name services and proxies, the anonymous credential system (anonymity at application layer) allows for balancing some linkability risks and threats as described in [PB13]. At the technical level, we leveraged *idemix* [CL00] along with OAuth for showing how other users could be retrieved within the di.me environment from a user's PS. The described technical solution is agnostic from the underlying social networking protocol used for enforcing authorisation in the respective server resources (e.g. OAuth). Special focus was also put on the support of mobile devices for future identity management scenarios since those devices are still have restricted anonymity support at the network as well as the application level.

User recommendations to avoid content-based linkability (scenarios 2 and 3)

Even though di.me is a decentralized solution unwanted information disclosure and linkability issues cannot be not completely avoided (accidental or intentional).[15] Santen [S06] points out that also user errors, in particular disclosing linkable information, may be used by attackers. The approach followed within di.me consists of (1) increasing the awareness of users by warning them during risky information sharing activities; and (2) engineering the system to securely process data and to help detect potential threats, e.g. when aggregating contact's information.

To sustain unlinkability of multiple identities of a single user, an approach is to control the potential linkability of partial identities. Therefore, the di.me userware analyses profiles and published information in order to find identifiability sets (a set of attributes identifying an individual user within a set of other users [PH10][Br05][Cl94]) and provides recommendations to the user based on that. This technique can also be applied even before information is shared in order to warn the user about potential privacy risks (e.g. when two contacts receiving different information might be the same person). In both cases, the di.me userware utilises techniques like semantic matching and semantic lifting to analyse textual information (see [CS12], [BR12] and [HB13]). This is used on the one hand to detect similarities in profiles of contacts and trigger a merge recommendation. On the other hand, profiles are compared and the user is warned accordingly, when linkability risks occur.

---

[15] By the users themselves or by others, e.g. by third parties (s. above example or someone disclosing information about his/her contacts)

Further, semantic analysis is applied on text and status messages written by the user to detect privacy threats because of shared information. For this, messages are decomposed into named entities and matched to identify persons, places, activities, etc. This text-analysis can be used to present privacy-enhancing recommendations to users. A prototype implementation shows warnings to the user that potentially sensitive information about third persons is being shared if contact names together with place or activity information is contained in a written text.

User-awareness on partial identities in the UI (scenario 3)

Many studies (cf. [CG06] and [KF10]) show that the UI plays a central role in handling privacy preferences and interpreting privacy notifications in threat situations like intentional or accidental information disclosure. Within a system offering multiple partial identities, the UI is a central mean to support the user's understanding of the segregation of identities [AW13], their distribution in the social network, and for avoiding undesired linkability (scenario 3). The di.me UI shall foster the user's awareness of multiple identities, the privacy preferences, and sharing history, and offer means to manage and control them. For representing identities within a user's PS, a UI object "profile card" has been chosen [HS13]. By selecting a profile card, e.g. for sharing an information item via it, the user selects the identity information shown to the recipient, *and* the SAID representing the identity on the network level. The decision to combine the selection of the SAID with the profile card was taken in order to reduce the UI complexity (for a discussion of usability and test results see [HS13]).

In the di.me approach the system supports the user in selecting the appropriate profile card (and this way implicitly the identity) to be used for sharing or communicating. Heuristics for that cover several rules, like suggesting profile cards already known to a recipient, profile cards already used for sharing a particular information item, or – based on di.me's recognition of contexts [SC12] – profile cards related to a current situation or sharing context. However, complex cases cannot easily be covered by heuristics. E.g., when multiple recipients are selected, no single profile card may be identified as sharing identity.

For such cases, di.me provides linkability warnings based on the sharing history: The system shows warnings that a selected profile card was never shared to a recipient before, or that a profile card (and other information item) will be shared outside the usual groups.

Summary and Outlook

The requirements for supporting unlinkability scenarios in a decentralised, multiple-identity social network comprise efforts on the technical level, the level of the UI and pro-active support e.g. in terms of recommendations and warnings. For three scenarios, di.me implements approaches to impede the linkability of multiple identities: A proxy layer as communication infrastructure is combined with SAIDs which allows reference

to identities independent from the IP address of the corresponding PSs. As result of this, analysing the technical communication protocol to link shared information to a root identity is inhibited. On the level of user content, di.me analyses the messages and profile information to find identifiability sets. Based on this, warnings about potentially critical content are presented. Warnings are triggered when information is being shared in order to make the user aware of potentially unintended linkability of partial identities and the root identity. Finally, the UI is designed to avoid privacy and linkability risks: Based on the UI object "profile cards", representing the user's partial identities, the user shall be enabled to control, and manage partial identities. To further support the user avoiding unintended disclosure of identities, warnings based on the sharing history are provided.

For the di.me userware as decentralised social network, these solutions form an integrated approach to avoid linkability of the offered multiple identities. To evaluate the approach, a prototype has been developed and implemented within a testing environment. The current version comprises of support for SAID resolving, sharing and communication using profile cards, and warnings based on sharing history and context. Further solutions, e.g. additional context-based heuristics for user recommendations, shall be incorporated and tested within the main demonstrator and the open source project. Recently started evaluation activities offer the prototype to a larger group of test-users and aim at gathering usability results as well as general feedback to the acceptance of the presented solutions for privacy-ensuring social networking. While preliminary results on the general concepts appear promising, further results on linkability advisory and other specific features are pending.

Acknowledgement

References

[AW13] Angulo, J., Wästlund, E.: Identity Management through "Profiles": Prototyping an Online Information Segregation Service. In Lecture Notes in Computer Science. Human-Computer Interaction. Users and Contexts of Use (pp. 10–19). Berlin Heidelberg: Springer (2013).

[Br05] Brands, S.: A primer on user identification. The 15th Annual Conference on Computers, Freedom and Privacy, Keeping an Eye on the Panopticon: Workshop on Vanishing Anonymity, Seattle. 2005.

[BB10] Bourimi, M., Barth, T., Haake, J., Ueberschär, B., Kesdogan, D.: AFFINE for enforcing earlier consideration of NFRs and human factors when building socio-technical systems following agile methodologies, in Human-Centred Software Engineering, ser. Lecture Notes in Computer Science, R. Bernhaupt, P. Forbrig, J. Gulliksen, M. Larusdottir, Eds. Springer-Verlag, 2010, vol. 6409, pp. 182–189.

[BH12]   Bourimi, B., Heupel, M., Westermann, B., Kesdogan, D., Planaguma, M., Gimenez, R., Karatas, R., Schwarte, P.: Towards Transparent Anonymity for User-controlled Servers Supporting Collaborative Scenarios. In Ninth International Conference on Information Technology: New Generations (ITNG), pages 102–108, April 2012.

[BR12]   Bourimi, M., Rivera, I., Scerri, S., Heupel, M., Cortis, K., Thiel, S.; Integrating multi-source user data to enhance privacy in social interaction. In Proceedings of the 13th International Conference on Interaccion Persona-Ordenador, INTERACCION '12, pages 51:1–51:7, New York, NY, USA, 2012.

[CL00]   Camenisch J., Lysyanskaya, A.: An efficient system for non- transferable anonymous credentials with optional anonymity revocation, in Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques: Advances in Cryptology, ser. EUROCRYPT '01. London, UK, UK: Springer-Verlag, 2001, pp. 93–118.

[Cl94]   Clarke, R.: Human Identification in Information Systems: Management Challenges and Public Policy Issues, Information Technology & People, Vol. 7 Iss: 4, pp.6 – 37. 1994.

[CS12]   Cortis, K., Scerri, S., Rivera, I., Handschuh, S.: Discovering semantic equivalence of people behind online pro_les. In Proceedings of the 5th International Workshop on Resource Discovery (RED 2012), 2012.

[CG06]   Cranor, L., Garfinkel, S.: Security and Usability. O'Reilly Media, Inc. (2005)

[FH12]   Fischer, L., Heupel, M., Bourimi, M., Kesdogan, D., Gimenez, R.: Enhancing Privacy in Collaborative Scenarios Utilising a Flexible Proxy Layer. In International Confernce on Future Generation Communication (FGCT). IEEE Computer Society, 2012.

[HB13]   Heupel, M., Bourimi, M., Scerri, S., and Kesdogan, D.: Privacy-preserving concepts for supporting recommendations in decentralized OSNs. In Proceedings of the 4th international workshop on Modeling Social Media, MSM '13, New York, NY, USA, 2013.

[HS13]   Hermann, F., Schuller, A., Thiel, S., Knecht, C., Scerri, S.: The di.me User Interface: Concepts for Sharing Personal Information via Multiple Identities in a Decentralized Social Network. In Lecture Notes in Computer Science. Human-Computer Interaction. Users and Contexts of Use (pp. 29–38). Berlin Heidelberg: Springer. (2013)

[KF00]   Krontiris, I. Freiling, F.: Integrating people-centric sensing with social networks: A privacy research agenda. In Pervasive Computing and Communications Workshops (PERCOM Workshops), 2010 8th IEEE International Conference on, pages 620 –623, 29 2010-april 2 2010.

[PH10]   Pfitzmann, A. Hansen, M.: A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management (Version v0.34). 2010. Retrieved from http://dud.inf.tu-dresden.de/Anon_Terminology.shtml (Last access: July 2013).

[R00]   Rannenberg, R.: Multilateral security a concept and examples for balanced security, in Proceedings of the 2000 workshop on New security paradigms, ser. NSPW '00. New York, NY, USA: ACM, 2000, pp. 151–162.

[S06]   Santen, T.: Security Engineering: Requirements Analysis, Specification, and Implementation. Habilitation, Fakultät Elektrotechnik und Informatik, Technische Universität Berlin (2006).

[TB12]   Thiel, S., Bourimi, M., Gimenez, R., Scerri, S., Schuller, A., Valla, M., Wrobel, S., Fra, C., Hermann, F.: A requirements-driven approach towards decentralized social networks. In Future Information Technology, Application, and Service, volume 164 of Lecture Notes in Electrical Engineering, pages 709–718. Springer-Verlag, 2012.

[SB13]   Schwarte, P., Bourimi, M., Heupel, M., Kesdogan, D., Gimenez, R., Wrobel, S., Thiel, S.: Multilaterally secure communication anonymity in decentralized social networking. To appear in IEEE Xplore as part of the proceeding of: 10th International Conference on Information Technology: New Generations (ITNG 2013).

[SC12]   Scerri, S., Cortis, K., Rivera, I., Hermann, F., Bourimi, M.: di.me: Context-Aware, Privacy-Sensitive Management of the Integrated Personal Information Sphere. In 9th Extended Semantic Web Conference (ESWC2012). 2012. Retrieved from http://data.semanticweb.org/conference/eswc/2012/paper/project-networking/372/html (Last access: July 2013).

[YL09]   Yeung, C., Liccardi, I., Lu, K., Seneviratne, O., Berners-Lee, T.: Decentralization: The Future of Online Social Networking. W3C Workshop on the Future of Social Networking. 2009. Available at http://www.w3.org/2008/09/msnws/papers /decentralization.pdf (Last access: July 2013).

Secure Hardware-Based Public Cloud Storage

Bernd Zwattendorfer[1], Bojan Suzic[2], Peter Teufl[2], Andreas Derler[3]

[1]E-Government Innovationszentrum (EGIZ)
bernd.zwattendorfer@egiz.gv.at

[2]A-SIT – Secure Information Technology Center – Austria
{bojan.suzic, peter.teufl}@a-sit.at

[3]Graz University of Technology
andreas.derler@student.tugraz.at

Abstract: The storage of data on remote systems such as the public cloud opens new challenges in the field of data protection and security of the stored files. One possible solution for meeting these challenges is the encryption of the data at the local device, e.g. desktop, tablet, or smartphone, prior to the data transfer to the remote cloud-based storage. However, this approach bears additional challenges itself, such as secure encryption key management or secure and effective sharing of data in user groups. Including an additional encryption layer and security checks may additionally affect the system's usability, as higher security requirements and a group sharing workflow increase general overhead through the complete organization of processes. To overcome such issues, we propose a solution which is based on highly secure and attack-resistant hardware-based encryption applied through the use of the Austrian citizen card public key infrastructure. As the citizen card infrastructure is already deployed and available to a wide population, the service overhead and additional requirements of our proposed solution are lower in comparison to other approaches, while at the same time synergistic and networking effects of the deployed infrastructure facilitate its usage and further potentials.

1 Introduction

The quantity of digital information increases steadily as businesses improve processing and managing of information by digitizing and structuring them. Additionally, the amount of data stored by private users is boosted by high quality multimedia files and constantly declining storage prices. The way of accessing data ought to be independent of location and device, especially with the rising popularity and broader usage of mobile devices such as smartphones and tablets. These factors contributed to increased demand for storage capabilities e.g. for archiving or backup purposes. From that point, many subjects identified public cloud storage services as adequate or optimal means to lower costs and increase service flexibility and potential by outsourcing data storage and

providing file synchronization across multiple clients. Popular examples of such public cloud storage services are e.g. DropBox[1] or Google Drive[2].

While insensitive information and data can simply be stored on such public cloud providers, security and confidentiality plays an inevitable role if sensitive data needs to be stored in the cloud. Most cloud providers cannot easily fulfill such requirements, as the providers usually are able to inspect the stored data. Even if the cloud provider encrypts the data and stores it in encrypted format, the provider is always in possession of the decryption key.

To still be able to store sensitive data securely and confidentially in the cloud, some cloud providers offer solutions where data is encrypted on client-side prior to its transfer to the cloud. We introduce such solutions briefly in Section 2. However, most of those solutions have the drawback that the encryption and decryption process relies on software-based keys, which are stored on the respective client device and under some conditions could be accessible by unauthorized parties. To bypass security issues raised with that approach, we propose a solution which uses a hardware-based key pair kept on a smart card to protect data stored in the cloud. Our solution therefore relies on the Austrian citizen card, which represents the official eID in Austria [HKR+08]. The usage of the Austrian citizen card has the advantage that it is based on a solid and independent Public-Key-Infrastructure (PKI). Hence, data can be practically encrypted for each Austrian citizen and securely stored and shared in the cloud. In this paper, we present the implementation of this approach and compare it with existing solutions.

2 Related Work

As importance of security and privacy concerning cloud storage services increased, several designs enhancing these properties have been proposed. In this section, we firstly introduce two different designs for cloud storage, namely "cloud storage services" and "encryption middleware". Secondly, we describe related work in the area of encryption middleware designs, as our proposed solution also fits into this design approach. Finally, the related work will also serve as a basis for our evaluation in Section 5.

Cloud storage services usually consist of a user-friendly client application and server-side software to store data. Some of these services also provide a web interface and a service API. The aim of cloud storage services is to provide cost-effective solutions for users to store and backup their data remotely, which should be easily accessed by different clients. Variable amount of storage can be bought as packages, while limited space is available for free. All information is redundantly stored in different places in order to increase availability of files. The client application creates a specific folder in the user's home directory. File actions within this directory trigger automatically the syncing to the cloud storage. Furthermore, if available, files can be accessed and managed through a web interface. Typical features of cloud storage services are backup,

---

[1] https://www.dropbox.com
[2] https://drive.google.com

synchronization, and sharing. Typical implementations of cloud storage services are DropBox, Google Drive, Microsoft SkyDrive[3], Wuala[4], or SugarSync[5].

Encryption middleware describes an encryption interface between client and cloud storage provider, with the purpose to ensure security and confidentiality, independent of the cloud storage provider. As many users doubt the security features of cloud storage providers, encryption middleware tries to resolve this issue. It provides an additional security layer in the form of client-side file encryption, which is performed before files are uploaded to the cloud storage service. This process involves management of required encryption keys, which are required for the en/decryption process.

In the next sub-sections we briefly describe the encryption middleware implementations Boxcryptor[6], CloudFogger[7], and Viivo[8]. An evaluation of these solutions as well as of our proposed approach is given in Section 5.

2.1 Boxcryptor

Boxcryptor is available for multiple platforms, e.g. Windows, Mac OS X, iOS, and Android. Boxcryptor provides support for the cloud storage services Dropbox, SugarSync, Microsoft SkyDrive, and Google Drive. A basic version of BoxCryptor is offered for free. Additionally to this free version, Boxcryptor can be purchased in an unlimited version, which enables filename encryption. Storage is managed in volumes, where each volume is mapped to a specific cloud storage service. Copying files into a volume invokes encryption and the encrypted file is copied into a corresponding subfolder of the cloud storage service directory. For example, copying files into a volume mapped to DropBox will store the encrypted files into a Boxcryptor specific subfolder of the DropBox folder.

2.2 CloudFogger

CloudFogger is freely available for Windows, Mac OS X, Android, and iOS platforms. Supported cloud storage services are DropBox, SkyDrive and Google Drive. Users need to specify which cloud storage services they wish to protect, with the option of disabling protection for subfolders. Protected cloud storage service directories can be accessed and manipulated as usual. However, before uploading files to the cloud storage, CloudFogger encrypts each file and uploads the encrypted file instead.

2.3 Viivo

Viivo is a free product and available for iOS, Android, Mac OS X, and Windows platforms. As of April 2013, DropBox is the only supported cloud storage service. When copying files into the Viivo folder within the user's home directory, it causes the encrypted versions of those files to be stored into a specific subfolder of the DropBox

---

[3] https://www.sugarsync.com
[4] http://www.wuala.com
[5] https://www.sugarsync.com
[6] https://www.boxcryptor.com
[7] http://www.cloudfogger.com
[8] http://www.viivo.com

directory, which are subsequently uploaded to Dropbox servers. The opposite way around, encrypted files added to the DropBox subfolder are decrypted automatically and consequently stored in the Viivo home folder.

3 Citizen Card Encrypted (CCE)

The following two sub-sections explain the concept of the Austrian citizen card and the Citizen Card Encrypted (CCE) software, which takes use of the Austrian citizen card functionality for encrypting and decrypting data.

3.1 The Austrian Citizen Card Concept

The Austrian citizen card [HKR+08], the official eID in Austria, constitutes a core element within the Austrian e-Government concept. The main aim is to facilitate electronic communication processes between citizens and public authorities. Moreover, by the help of the Austrian citizen card such electronic communication processes can be accelerated and secured at the same time.

In general, the term "citizen card" is more seen as a concept rather than a card. The Austrian e-Government Act [EGovG], which defines the Austrian citizen card in legal terms, emphasizes especially its technology neutrality and its independence of technical components. Due to declared technology neutrality, different implementations are possible and do already exist for the citizen card. Currently, the most dominant citizen card implementation in Austria is a smart card. For instance, each Austrian citizen gets issued a health insurance card (*e-card*), which can easily be activated to use citizen card functionality. Nevertheless, another emerging citizen card technology is based on mobile phones. In this implementation, a server-side hardware security module stores the citizens' secret keys, which can be activated by the use of the citizen's mobile phone.

In general, the most important functionalities of the Austrian citizen card, as regulated in the Austrian e-Government Act, are (1) citizen identification and authentication, (2) generation of qualified electronic signatures and (3) data encryption and decryption.

By using the Austrian citizen card, citizens can be uniquely identified and securely authenticated at governmental or private sector online applications. Additionally, the Austrian citizen card contains a qualified signature certificate according to the EU Signature Directive [EP95]. Hence, electronic signatures created with an Austrian citizen card are legally equivalent to handwritten signatures. Besides this signature certificate, an additional key pair is stored on the card, which can be used for the secure encryption and decryption of data. Thereby, the public encryption keys of every Austrian citizen are available through a central LDAP directory. Hence, data can be encrypted for each Austrian citizen and stored confidentially. In the remainder of this paper, we focus on the encryption and decryption functionality of the Austrian citizen card only.

3.2 The CCE Software

The CCE (Citizen Card Encrypted Software) is a platform-independent and open source software developed by A-SIT (Secure Information Technology Center – Austria). The

software is available through the JoinUp platform, a service initiated and supported by European Commission[9]. CCE especially supports the public authorities demanding high data security and easy and flexible data management. Basically, CCE allows for the encryption and decryption of arbitrary data and the management of files or directories both for single and multiple users.

For file and directory encryption and decryption CCE relies on hardware-based keys, which are stored on the Austrian citizen card. However, also software-based keys can be used within CCE. Particularly the use of the Austrian citizen card enables a highly secure and confidential data exchange since the required keys are stored in hardware and thus cannot be read out by an application. CCE currently supports the smart card-based implementation of the Austrian citizen card only, as no encryption and decryption functionality is provided by the mobile phone signature at the moment. However, other smart card implementations can be easily integrated by implementing an application interface for a particular implementation.

CCE relies on the well-known and established S/MIME [RT04] standard as container format for storing data. S/MIME is also widely integrated in several e-mail clients for encrypting e-mails. In the following, we briefly explain main features of the CCE software.

- *Smart card as secure decryption unit*

  The CCE software supports the use of smart cards to decrypt the keys used in S/MIME containers. The process of decryption is directly carried out on the smart card, initiated by the user entering a personal PIN.

- *Support of group encryption*

  Files and directories can be encrypted for multiple users, which can be organised in a group-like hierarchy. The management of groups is handled manually by the users on their own. However, the support of multiple users also allows for the inclusion of appropriate backup keys.

- *Support of the Austrian PKI infrastructure*

  Asymmetric public key encryption facilitates encryption procedures of users and groups. The public keys of recipients are hence publicly available through the Austrian PKI infrastructure by querying the central LDAP directory. Nevertheless, CCE also enables the integration of arbitrary PKI infrastructures (e.g. from an enterprise context), which can be done by extending its open-source application interface to support the new infrastructure.

4 Architecture and Implementation

In this section we explain the architecture and implementation of our smart card-based approach for storing data securely and confidentially in the public cloud.

---

[9] http://joinup.ec.europa.eu/software/cce/description

4.1 Architecture

For our solution the CCE software has been extended in order to be able to store data also at public cloud providers and not only on the local storage. Citizens can thereby select between different cloud storage services where data should be stored. The current implementation supports the providers DropBox and Google Drive.

Fig. 1 illustrates our architecture for secure encryption and decryption of data by using the Austrian citizen card functionality and storing the encrypted data in the public cloud. In this architecture, in fact three different entities are involved: (1) the citizen who wants to store some file or directory securely in the public cloud, (2) the Austrian citizens the files or directory should be encrypted for and, (3) the public cloud provider where the encrypted files will be stored.

Figure 1: Architecture for securely storing data in the public cloud using the Austrian citizen card

Fig. 1 also illustrates the encryption process using CCE and subsequently the process of storing the encrypted data in the public cloud. In a first step (Step 1), the citizen selects the files and directories she wants to store securely and confidentially in the cloud. In the next step (Step 2), the citizen selects one or more other persons (Austrian citizens) the chosen files or directories should be encrypted for. If citizens' encryption certificates are not known by CCE yet, they can be queried from the central LDAP directory[10]. In this directory, all public certificates of every Austrian citizen registered in the system are stored. Before starting the encryption process, the validity of the encryption certificates of the selected persons is checked. Finally, in Step 3 the data are encrypted for the intended citizens and transferred to the selected public cloud provider. Authentication credentials for accessing the public cloud provider need to be provided during the

---

[10] The querying of the external LDAP service is not necessary if the users have exchanged the certificates, e.g. using email or by using organizational certificate store. It is also possible to include own LDAP server.

configuration and setup of CCE. During the data transfer, the credentials are retrieved from the CCE configuration and provided to the public cloud provider automatically.

The decryption process is similar to the encryption process; hence the decryption process will not be illustrated. In the decryption process, the encrypted data are downloaded from the public cloud into the local file system by the user. Afterwards, the data are decrypted by using CCE and invoking the citizen's citizen card. Now, the citizen is able to inspect the plain data.

4.2 Implementation

For supporting public cloud storage as an option, CCE had to be amended and extended accordingly. In particular, emphasis was put on flexible adding of additional public cloud providers besides DropBox and Google Drive. For adding an additional cloud provider, the server communication with the cloud provider and its configuration management needs to be implemented. Hence, the modular internal architecture of CCE allows for an easy implementation of new providers.

The creation of a new public cloud provider configuration requires a smart card because the smart card is linked to credential information necessary to access cloud provider services. The credential information for the cloud provider is thereby encrypted by the affiliated smart card, stored in the local file system, and assigned to the corresponding person. Hence, an automatic mapping between smart card and cloud provider authentication credentials is achieved. The advantage of this approach lies in the fact that cloud specific authentication data need to be entered once during configuration; it is then accessed automatically during each subsequent cloud data transfer.

In details, configuration of authentication credentials for cloud provider access is as follows. Authentication at the cloud provider is based on the authorization protocol OAuth[11] for both cloud providers DropBox and Google Drive. Required authentication tokens of OAuth are ascertained during the configuration of a new cloud provider in CCE. This requires the input of the authentication credentials from the user, which in turn adds CCE as trusted cloud application and gives CCE access to the user's cloud account. Subsequently, CCE receives an access token from the cloud provider for the secure access to the cloud storage. According to the OAuth protocol, this access token can be continuously used for cloud provider authentication, so that additional provision of user authentication credentials is not required anymore.

To store data confidentially, users are able to select their desired storage location. The default location is the local file system, whereas users are now able to also store encrypted data at different cloud providers, which are linked with their citizen card. During data upload, saved cloud provider credentials are decrypted by using the user's smart card and are used for cloud provider authentication.

Besides extending the pure CCE application, integration into the operating system's file system has been implemented too. In this case, users are able to copy files into a specific

---

[11] http://oauth.net

folder of the personal HOME directory and files are then automatically encrypted and transferred to the cloud. When moving files into this specific folder, the CCE wizard starts automatically. Recognition of moved or newly created files in this specific folder is implemented using WatchServices[12], which observes file system operations. Using the CCE wizard, not only files can be automatically encrypted but also desired recipients can be selected. For distribution of encrypted files the existing mechanisms of the respective cloud provider can be used.

5 Evaluation

In this section we evaluate encryption related features and functionalities of middleware implementations for cloud storage in terms of encryption and data sharing.

5.1 Boxcryptor

Boxcryptor **encrypts** files using the AES-256 encryption algorithm. The encryption scheme is volume specific, where all files inside one volume are encrypted with its particular key. This volume-specific key is generated randomly, encrypted with the master key derived from the user's password, and placed in the volume's root. Therefore, in this approach encryption keys are derived from the user's password, which may be leaked through phishing attacks, caught by Trojans, or accidentally published to vicious third parties. Another disadvantage of Boxcryptor's approach is the fact that filename encryption is performed only in the unlimited and retail version of the software. The standard and free version of the software does not obfuscate filenames, which poses additional security risk and information channels for attackers.

Sharing in Boxcryptor is possible only for entire volumes mapped to a specific provider. In order to gain volume access, it is required for the user to share the password, which is not considered as a highly secure practice.

5.2 CloudFogger

During a new account creation on the CloudFogger service, a user specific RSA key pair is generated locally on the user's device. The private key is then **encrypted** with a user provided password, using AES-256 and uploaded together with the public key to CloudFogger servers. In this approach, the encrypted private key information is always downloaded and decrypted with the user's password locally on the user's device, allowing access to protected files. This way, CloudFogger is never able to gain knowledge of private key or password information, making it possible for the user to consume the service on different devices. Each file is individually encrypted using AES-256 whereby AES-keys are encrypted with the user's public key and embedded in the file. Due to file encryption based on user passwords, phishing and Trojan attacks, as well as password leaking, are viable threats to the security of this approach.

As AES-keys are embedded directly in each of the encrypted files, they can easily be **shared** with other subjects. For such purpose, embedded AES-keys files are encrypted

---

[12] http://docs.oracle.com/javase/7/docs/api/java/nio/file/WatchService.html

with the public keys of invitees[13]. This allows the invitee to locally decrypt shared files with her private key. Sharing can be handled independent of the underlying cloud storage services. However, all participants are required to be registered to CloudFogger.

5.3 Viivo

Similarly as for CloudFogger, RSA key pairs of the users are created locally during the process of account registration. Both public and **encrypted** private keys are stored on Viivo servers. The encrypted private key is downloaded on the user's client device and decrypted by providing the corresponding password. Moreover, each file in the system is encrypted using the AES-256 encryption algorithm, whereby AES-keys are encrypted with the private key associated with the user. As the encryption approach of Viivo is basically similar to the one of CloudFogger, they both share similar disadvantages from the security perspective. Having the encryption keys derived from user passwords, attacks ranging from phishing and Trojan attacks to information leakage are possible for both of the approaches. As all the keys and files depend on one master user password, its leakage may render the whole service and system unusable.

The **sharing** of files with others is performed by inviting the respective user, which has to manually allow sharing of particular files. Creating a share invokes generation of new AES-keys for all files in the share. These keys are in then encrypted with the public key of every invitee. Then, the encrypted keys are sent by the inviting user to each invitee.

From the user's perspective, sharing of a file stored on DropBox is done in two activities. Firstly, the file has to be shared through the DropBox sharing mechanism. Secondly, the sharing of specific files has to be allowed by the invitee through the Viivo interface. In contrary, when access to shared files is revoked, the shared files are not re-encrypted. Instead, new keys are created. New keys ensure that newly created files are no longer accessible by the previously invited user.

5.4 CCE

CCE uses a slightly different approach for file **encryption** than other evaluated solutions. Instead to create RSA key pairs for new users each time they register, and store them on (potentially insecure) local storage prior to the encryption, CCE relies on the existing Austrian citizen card PKI infrastructure. This way, it uses independent, third-party smart card and secure hardware based encryption.

The containers in CCE, which can hold files and directories, are encrypted with AES symmetric keys. These keys are further encrypted using the public RSA key of the Austrian citizen card, taking the public RSA key of each user being allowed to access the container. The containers itself are stored in S/MIME format, which is compatible with a broad range of other applications, including popular e-mail clients. For the decryption of encrypted files, the Austrian citizen card in the form of smart cards is used. This presumes that encryption keys are encrypted with the user's public key and are decrypted in the smart card, using the securely stored private key.

---

[13]Invitees – persons having access rights on the file

The advantage of this approach is that the private keys are never loaded into the computer system, nor can they be directly accessed or read. Instead, they are contained in the smart card and operations involving them are executed on smart card hardware only when necessary conditions are met (e.g. PIN-based authentication). However, CCE is not limited on the use of Austrian citizen cards only. It can support other PKI infrastructures or smart card implementations, or can rely on software-based keys too.

The **sharing** of files in CCE is performed in two steps. First, the user selects intended recipients during the encryption process. CCE encrypts files for these users by encrypting and storing the symmetric key in the container for each particular user, using her public key. The public key of the user can be stored locally or retrieved from the public LDAP directory of the Austrian citizen card PKI. Furthermore, it is possible to encrypt files for not previously known or contacted users, where prior key exchange or establishing of contact is not necessary. In the second step, the user enables access to the underlying cloud storage for intended users and performs upload of the encrypted containers or synchronization with the local directory with the containers.

The credentials to access remote cloud services in CCE are stored in secure manner. They are encrypted and stored in a local XML file. In order to enable access to remote cloud services, the user has to insert her smart card used during credentials initialization. This approach prevents the leakage of the cloud credentials to unauthorized third parties.

5.5 Summary

In the previously presented evaluations and based on summarized comparison in Table 1, we demonstrate the advantage of our CCE-based solution. From the security perspective, our solution relies on hardware-based encryption, where the private key used for decryption never leaves the smart card. This case does not require the usage of a master password and consequent key derivation as it is the case for BoxCryptor, CloudFogger, and Viivo. The CCE approach is prone to phishing and keylogger attacks[14], which may render the complete system unusable in the case the master password is compromised. From the broader sharing and usability view, our solution's advantage lies in the fact that it relies on a third-party public PKI infrastructure.

The Austrian citizen card is available to all persons living in Austria as a part of several implementations, including bank cards or social health insurance cards, which are basically most widely deployed. As the software is open-source, the support for other PKI infrastructures can be easily implemented by extending the application interface.

Thus, the existing, already present and widely deployed infrastructure is used without incurring additional costs or overhead. That enables the exploration of networking effects, as there are already many users having their public certificate enabled and reachable through the public LDAP endpoint. Users can simply encrypt and exchange encrypted files between each other without the necessity to maintain prior contacts and engage in secure key exchange.

---

[14] If PIN-Pad smart card readers are used.

Feature \ Middleware	BoxCryptor	CloudFogger	Viivo	CCE
Encryption				
AES support	√	√	√	√
RSA support	-	√	√	√
Volume-based encryption	√	-	-	-
File-based encryption	-	√	√	√[15]
Container-based encryption	-	-	-	√
File names securely stored	-	√	√	√
Software keys supported	√	√	√	√
Hardware keys supported	-	-	-	√
User master password derivation based encryption	√	√	√	-
Encryption keys stored locally	-	-	-	√
Encryption keys stored remotely	√	√	√	-
Phishing attack prone	√	√	√	-
Keylogger attack prone	√	√	√	-
Sharing				
Prior key exchange necessary	√	√	√	-
Public LDAP Key discovery	-	-	-	√
Encryption for unknown users	-	-	-	√[16]
Volume based sharing	√	-	-	-
File/directory based sharing	-	√	√	√
Feature \ Middleware	BoxCryptor	CloudFogger	Viivo	CCE
Relies on third-party PKI	-	-	-	√[17]
Multiplatform support	√	√	√	√[18]

Table 1: Comparison of middleware encryption and sharing features

[15] Possible by encrypting one file per container
[16] Using public LDAP endpoint for searching/browsing recipients
[17] Using the Austrian citizen card public key infrastructure. Can be extended with a private PKI.
[18] The mobile versions of the software are not publicly available, however, they are currently in the development phase (iOS and Android platforms)

6 Conclusion and Further Work

The storage of data in the public cloud is becoming popular and a widely used scenario. As the cloud market is intensively growing and providing many new and innovative service and integration solutions, it can be expected that the necessity to store personal or business data in the public cloud will grow even further in nearly future. However, storing data in remote public cloud systems brings new challenges and security risks.

In our work, we focused on data confidentiality and general security aspects of cloud storage in the multi-group and multi-provider scenario. For such purpose we extended the file encryption tool CCE, which acts as an encryption middleware on the local user computer. This extension includes the support for data encryption and sharing via public cloud services. Furthermore, we analyzed publicly available middleware encryption solutions, compared their features, and provided an overview of the features of these tools. Based on this evaluation, we demonstrated that our solution provides significant advantages in terms of data security and resistance to several popular attack techniques. As the proposed solution is based on already deployed and widely used infrastructure, it requires minimal costs or overhead in order to be applied.

There are several directions which could be taken for further development of our work. Currently, there are two versions of the software for iOS and Android in development. We plan to integrate them with the desktop solution presented in this work. Another possible task for the future is the integration of the tool into the web browser, which may provide additional quality in user experience and broader platform support. In addition, we evaluate possibilities to provide cloud storage redundancy at the middleware level, meaning to store data distributed on different cloud services. Finally, we try to integrate our solution in the user's operating system, so that the data can be visible, accessed, and manipulated directly at the operating system level.

References

[ABC01] Abraham, N.; Bibel, U.; Corleone, P.: Formatting Contributions for LNI. In (Glück, H.I. Hrsg.): Proc. 7th Int. Conf. on Formatting of Workshop-Proceedings, New York 1999. Noah & Sons, San Francisco, 2001; S. 46-53.

[Ez99] Ezgarani, O.: The Magic Format – Your Way to Pretty Books, Noah & Sons, 2000.

[EGovG] Federal Act on Provisions Facilitating Electronic Communications with Public Bodies (The Austrian E-Government Act - E-GovG) StF: BGBl. I Nr. 10/2004

[EP95] Data Protection Directive 95/46/EG, EU Parlament, Official Gazette Nr. L 281 from 23/11/1995 P. 0031 – 0050

[HKR+08]A. Hollosi, G. Karlinger, T. Rössler, M. Centner: Die österreichische Bürgerkarte, Version 1.2, 2008, http://www.buergerkarte.at/konzept/securitylayer/spezifikation/aktuell/

[RT04] B. Ramsdell, S. Turner (2004): Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Message Specification", RFC 3851, 2004, http://www.ietf.org/rfc/rfc3851.txt

An extensible client platform for eID, signatures and more

Tobias Wich[1] · Moritz Horsch[2] · Dirk Petrautzki[3] · Johannes Schmölz[1]
Detlef Hühnlein[1] · Thomas Wieland[3] · Simon Potzernheim[3]

[1] ecsec GmbH, Sudetenstraße 16, 96247 Michelau,
{tobias.wich,johannes.schmoelz,detlef.huehnlein}@ecsec.de

[2] TU Darmstadt, Hochschulstraße 10, 64289 Darmstadt,
horsch@cdc.informatik.tu-darmstadt.de

[3] Hochschule Coburg, Friedrich-Streib-Str. 2, 96450 Coburg
{petrautzki,thomas.wieland,potzernheim}@hs-coburg.de

Abstract: The present paper introduces an extensible client platform, which can be used for eID, electronic signatures and many more smart card enabled applications.

1 Introduction

Against the background of various electronic identity (eID) card projects around the globe there have been numerous initiatives in the area of research, development and standardization of eID cards, smart card middleware components and related services. Nevertheless, whenever a new eID project emerges, new software is often developed from scratch. This happens despite all similarities of the systems and requirements. The present paper introduces a modular and extensible client platform, which can be extended for the use with eID, electronic signatures and many other smart card related applications. The design of this extensible platform is a refinement of the architecture of the Open eCard App [HPS⁺12], which in turn is based on the eCard-API-Framework (BSI-TR-03112) and its integrated international standards, such as ISO/IEC 24727 [ISO08a, ISO08b] and OASIS Digital Signature Services [Dre07]. The design and implementation of the platform has been based on previous work [Hor11, Pet11] and realized as a joint effort of industrial and academic experts within different projects, such as ID4health[1], SkIDentity[2], FutureID[3], and Open eCard[4].

The remainder of the paper is structured as follows: Section 2 provides an overview of the proposed client platform. Section 3 describes the extension points of the client platform. Section 4 presents the design of the add-on framework and its mechanisms to dynamically load missing functionality. Section 5 closes the paper with an outlook on the next steps and future development.

---

[1] See http://www.id4health.de.
[2] See http://www.skidentity.de.
[3] See http://www.futureid.eu.
[4] See http://www.openecard.org.

2 Overview of the extensible architecture

The proposed client platform is aligned to the eCard-API-Framework (BSI-TR-03112) which integrates major international standards (e.g. [ISO08a, ISO08b, Dre07]) in order to provide a common and homogeneous interface for a standardized usage of different smart cards. The architecture depicted in figure 1 is designed to separate the overall functionality of an eID application in suitable components, reuse of common modules and to provide means for expandability. The modular approach and the platform-independent implementation of the core modules in Java allow the Open eCard App to be used on various computing platforms, such as desktop systems running on Windows, Linux and Mac OS X as well as mobile systems running Android for example.

Figure 1: Extensible architecture of eID-client-platform

The components of the extensible eID platform are described in the following:

Interface Device (IFD) The IFD provides a generalized interface for communication with arbitrary card terminals and smart cards according to ISO/IEC 24727-4 [Fed12b, ISO08b]. It abstracts from specific interfaces and physical properties like contactless interfaces. Furthermore it provides expandability for the integration of secure channel establishment protocols which protect the communication between the eID client and the smart card.

Event Manager   The Event Manager is responsible for managing card terminal and card events. It periodically asks the IFD for the current status of terminals and cards and determines changes like the connection and disconnection of card terminals and smart cards by comparing status reports over different time periods. Furthermore, the Event Manager performs the card recognition to determine the type and the functionality of the card as explained in section 3.2.

Service Access Layer (SAL)   The SAL provides a generic interface for common smart card services according to ISO/IEC 24727-3 [ISO08a, Fed12c], which allows to manage data that is stored on the card for example. In detail, the SAL comprises Connection Services, Card Application Services, Named Data Services, Crypto Services, Differential Identity Services and means for accessing card application services in an authorized manner. Furthermore, the SAL provides an interface for integrating arbitrary authentication protocols, which provides expandability without changing other parts of the implementation (see section 3.3).

Dispatcher   The Dispatcher provides a centralized communication component for handling incoming and outgoing messages.

Add-ons   Add-ons provide additional functionality to the basic eID platform. Signature functionality and PIN Management, for instance, can be realised as an add-on to provide additional functionality and allow customisation. The Add-on Registry provides a service to search and retrieve add-ons. Such a registry can, e.g., be realised based on the Java Network Lauching Protocol (JNLP) [Her11]. After an add-on is loaded, the Add-on Manager takes over the management of the add-on instances and enforces the compliance with the defined security policy by a sandbox mechanism.

Bindings   The Binding component comprises modules for message transport. The components implement a particular protocol like HTTP or SOAP to transmit messages from external applications to the client.

Crypto   The Crypto component encapsulates common cryptographic functions, which are used by other components. It is based on the Bouncy Castle crypto library [The] which makes it easy to port it to platforms without support for the full Java Cryptography Architecture (JCA) [Orab], such as Android for example.

Graphical User Interface (GUI)   The GUI component provides an abstract framework to develop user interfaces and interactions. This allows the exchange of GUI implementations and therefore providing platform-specific GUI implementations, while leaving the other components unchanged.

3 Extension Points

This section describes the extension mechanisms of the eID platform, which allows enhancing the application's functionality on different levels. In detail, it allows adding arbitrary protocols to the IFD and SAL component, supporting various card terminals and smart cards as well as enhancing the application functionality by add-ons.

In general, we use the term *add-on* to describe a software component which enhances the functionality of the basic eID platform. Furthermore, we distinguish between *plug-in* and *extension*.

Plug-ins depend on the context in which the user uses the application. Performing an authentication to a service using a particular smart card, for instance, requires a plug-in which is capable of providing such functionality. Subsequently, plug-ins require a communication with bindings to interact with external applications and services. Furthermore, we distinguish between IFD, SAL and application plug-ins, which are described in detail in the following sections.

Extensions are independent from the context. Moreover, they are directly integrated into the user interface and can be executed by the user. For instance, an add-on that provides a PIN change functionality for smart cards is classified as an extension.

3.1 IFD Plug-ins

The IFD provides a generalized interface for communication with arbitrary smart cards and card terminals. It also can be extended by plug-ins, i.e. protocols which perform a user authentication and/or establish a secure channel between a smart card and a card terminal to protect the communication from being eavesdropped.

Each protocol must have a unique identifier in form of a URI. The URI must be associated with the actual implementation as described in section 4.1. In addition, each protocol plug-in must implement the IFD Protocol Interface and must define protocol-specific `AuthenticationProtocolData` used in the `EstablishChannel` call[5] and corresponding response message.

The Password Authenticated Connections Establishment (PACE) protocol is one example of a protocol which is executed in the IFD layer. It is a password-based protocol that performs a user authentication, based on a PIN, and establishes a Secure Messaging channel (cf. [ISO]) to ensure that only the legitimate user can use the card and that the communication is encrypted and integrity protected. The details of the PACE-protocol are specified in BSI-TR-03110 [Fed12a].

ISO/IEC 24727-4 Interface An IFD-protocol will be executed by an `EstablishChannel` IFD API call. The function call includes a `SlotHandle` to address an established con-

---

[5]See `http://ws.openecard.org/schema/ISOIFD-Extension.wsdl`.

nection and a protocol-specific extended `AuthenticationProtocolData` element.

Java Interface   The `IFDProtocol` interface defines functions for IFD protocols (cf. figure 2). Each protocol must implement the `establish` function that executes the protocol. The function gets as input an `EstablishChannel` request that includes protocol-specific data. The parts which are necessary to communicate with the eID application are handed over to the implementation in the `init` function. The context `ctx` contains the user consent implementation, which allows a protocol to perform user interaction, e.g. to receive PIN entries. In addition, the interface specifies the functions `applySM` and `removeSM` to apply and remove Secure Messaging. The `establish` function returns an `EstablishChannelResponse`. The `IFDProtocolFactory` provides a factory class which also proxies the protocol interface. The usage of the class name decouples the actual loading of the class and prevents execution of plug-in code outside of the sandbox.

Figure 2: IFD-Protocol-Interface UML diagram

3.2   CardInfo Files

In order to support a broad range of smart cards, the eID platform supports CardInfo files (CIF) according to [ISO08a]. A CIF is an XML file that describes the data structure and the functionality of smart cards in a standardized way. Besides the abstract definition of the card, it also contains information how to recognize the specific card type.

To provide a sophisticated recognition of smart cards it is prudent engineering practice to construct a decision tree based on the set of available CIFs (cf. [Wic11]). While the construction of the tree could be performed by the eID application on demand, this task is better performed by a central CardInfo repository, which performs the construction and only distributes the decision tree (cf. [Fed12e]). To make the eID application capable of recognizing new smart cards, only the corresponding CIFs and an updated version of the decision tree have to be added.

3.3 SAL Plug-ins

The SAL provides a generic interface for common smart card services comprising different services, such as the Crypto Services and the Differential Identity Services. The SAL can be extended by plug-ins, which provide implementations of protocols for the Crypto Services and the Differential Identity Services [Fed12d, Section 4] as required for the use of specific signature cards and electronic identity cards for example.

The plug-in concept is quite similar to the one that is used in the IFD layer (cf. section 3.1). Each SAL protocol must define a unique identifier (URI). In contrast to the IFD, the SAL supports protocols with multiple steps and allows the definition of more sophisticated user interfaces including a sequence of interaction steps to represent information dialogues and general user consents.

One example of a SAL protocol is the Extended Access Control (EAC) protocol which is used for the authentication with the german eID card. The protocol-specific messages are specified in [Fed12d, Section 4.6].

ISO/IEC 24727-3 Interface A protocol execution is triggered by invoking an action within the Crypto Services or Differential Identity Services API (cf. [ISO08b, section 3.5 and 3.6]). The functions includes an `AuthenticationProtocolData` element, which is extended in a protocol-specific manner.

Figure 3: SAL Protocol Interface UML diagram

Java Interface Each protocol must implement the `SALProtocol` interface. A convenience abstraction which works for the common protocol flows is realized in the class `SALProtocolBaseImpl`. An internal data object is used for the exchange of data between the different protocol steps. A protocol step is represented by the `ProtocolStep` interface which defines a `FunctionType` defining a Crypto or Differential Identity Service and a `perform` function to execute the step. The control of the application flow is performed automatically after being triggered by incoming Crypto or Differential Identity Service requests. The instantiation is performed through the `SALProtocolFactory`

similar to the IFD protocols explained in section 3.1.

3.4 Application Plug-ins

Application plug-ins provide a mechanism to add additional functionality to the eID application with which external applications can communicate. Depending on the type of the underlying binding, this could be a browser, a PKCS#11 module or even a remote application.

Protocol bindings realize the connection to the external world. While a broad variety of transport protocols could be supported, the most obvious choices are HTTP and SOAP, as they are stipulated by [Fed12d, Section 3.2] for example. Given the properties of the activation mechanism, HTTP and SOAP, as well as similar transport protocols, the abstract requirements for a protocol binding are given as follows: A protocol binding must support

1. a request-response semantic,

2. a mapping mechanism to identify the appropriate plug-in for a request,

3. messages comprising a body, named parameters and attachments,

4. an error delivery mechanism, and

5. a redirect semantic.

Figure 4: Application-Plug-in-Interface UML diagram

Figure 4 shows the interfaces and the data model of the application plug-ins. On the plug-in side it is easy to see that all properties are fulfilled. The interface `AppPluginAction` provides an `execute` function with a strict data oriented semantic, meaning no callback

code can be injected for asynchronous responses. The second property is fulfilled by a named identification of the action which is discussed in detail in section 4.1. The data structures for body and attachments can be seen on the right side of the diagram. Named parameters have no particular ordering and no special type so a string of characters can represent either key and value. These three elements form the input parameters of said `execute` function and are part of the result. The body element carries exactly one DOM node. This representation has the advantage that it can carry strings as well as more complex XML elements. That makes it suitable to provide the content of a SOAP body, JSON data converted to an XML representation or string based entities. Attachments are included to transport binary files. The data structure is modelled to support the most important features of MIME messages such as Multipart MIME messages (cf. [FB96, Section 5.1]). The fourth and fifth requirement are fulfilled by providing predefined response codes and auxiliary data for the specific type of action. In case of an error, a localized message may be attached to the result. A redirect needs a redirect target value in the auxiliary data. It is up to the receiving application how to interpret and perform the redirect. The open character of the auxiliary data makes it easy to add new capabilities for further use cases to the bindings without the need to change the Application Binary Interface (ABI) of the interface.

While different transport protocols (e.g. *HTTP on localhost*, *LiveConnect*, *SOAP*) may be used to realize bindings for the different add-ons (e.g. *eID Activation*, *Status*, *Signature PKCS#11*) we will explain the general concept using the example of a signature plug-in with the localhost binding according to BSI-TR-03112-7 [Fed12d] in the following.

Given the containers parameters, body and attachments, the plug-in can define its interface. A signature plug-in can be modelled in two ways. Either via an RPC-style interface where the properties of the plug-in are transported in parameters, or via an OASIS DSS [Dre07] like interface where the properties are transported as a structured object in the body.

Suppose the variant with the simple parameters is used, the following HTTP request (listing 1) and response (listing 2) messages can be modelled. The simple model might be desirable when the signature functionality is limited to a few base cases and thus the full OASIS DSS capabilities are not needed.

```
1  POST /signature?signatureType=XAdES&cardType=... HTTP/1.1
2
3  Content-Type: multipart/form-data; boundary=AaB03x
4
5  --AaB03x
6  Content-Disposition: form-data; name="files"; filename="data.xml"
7  Content-Type: text/xml
8
9  <?xml version="1.0" encoding="UTF-8" standalone="yes"?>
10 <Data xmlns="myns">to be signed</Data>
11 --AaB03x--
```

Listing 1: RPC-Style Sign Request

In order to sign to a document, at least the signature type and the data to be signed is required. To take away the responsibility of the user to select a signing entity, e.g. a specific

```
1   HTTP/1.1 200 OK
2
3   Content-Type: Multipart/mixed; boundary=AaB03x
4
5   --AaB03x
6   Content-Disposition: attachment; name="files"; filename="data.xml"
7   Content-Type: text/xml
8
9   <?xml version="1.0" encoding="UTF-8" standalone="yes"?>
10  <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
11    <SignedInfo>...</SignedInfo>
12    <SignatureValue>...</SignatureValue>
13    <KeyInfo>...</KeyInfo>
14    <Object Id="dataId">
15      <Data xmlns="myns">to be signed</Data>
16    </Object>
17  </Signature>
18  --AaB03x--
```

Listing 2: RPC-Style Sign Response

smart card, this information may be given as well. The parameters `signatureType` and
`cardType` as given in listing 1 line 1 represent the latter choices. The document itself
is included as a named part shown in line 5 ff. in the HTTP body. The representation
as `multipart/form-data` according to [RLJ99, Section 17] has been chosen so that
typical browsers can issue requests easily. Named parts can be matched to the attachment
type of the interface as well as to the body. To resolve the ambiguity, the body can simply
be an attachment with a special name value, but other schemes may be allowed as well to
capture other communication patterns.

```
1   POST /signature?cardType=... HTTP/1.1
2
3   Content-Type: application/xml
4   Content-Length: ...
5
6   <?xml version="1.0" encoding="UTF-8" standalone="yes"?>
7   <dss:SignRequest xmlns:dss="urn:oasis:names:tc:dss:1.0:core:schema">
8     <dss:InputDocuments>...</dss:InputDocuments>
9       <dss:Document>
10        <Data xmlns="myns">to be signed</Data>
11      </dss:Document>
12    <dss:OptionalInputs>
13      <dss:SignatureType>urn:ietf:rfc:3275</dss:SignatureType>
14    </dss:OptionalInputs>
15  </dss:SignRequest>
```

Listing 3: OASIS DSS-Style Sign Request

A more sophisticated data exchange for a signature plug-in is shown in listing 3. The
example uses OASIS DSS `SignRequest` messages to specify what kind of signature
should be performed and what should be signed. The signing entity is chosen as in the
previous example. The example also shows that the request is nearly identical to a SOAP

request, so the parameters can be mapped by either the localhost binding or a SOAP binding.

3.5 Application Extensions

Extensions enhance – similar to plug-ins – the basic eID platform and provide additional functionality, but they do not depend on the context in which the eID application is used. Further, extensions are included into the user interface and can be started directly by the user. Similar to application plug-ins, the `AppExtensionAction` interface, as shown in figure 5, contains an `execute` function. However, this function does not have any parameters nor does it have a result. Therefore, it cannot be used with a binding and only be triggered manually.

Figure 5: Application Extension Interface UML diagram

4 Add-on Framework

4.1 Add-on Anatomy

Add-ons are described by the data model shown in figure 6. This model is the representation of the XML structure of an add-on's manifest file. It contains general information such as the name, the textual description and configuration entries for changeable settings of the add-on, and its contained actions which represent the interfaces shown in section 3. The settings are saved in an add-on specific storage location and are loaded as Java properties by the add-on framework. Each action has one or more entries which identify it unambiguously. The IFD and SAL protocol plug-ins are identified by their protocol URI, whereas the application extensions and plug-ins are identified by the add-on id and action id, or resource name respectively. A reference to the action class makes it possible for the framework to find and load the implementation dynamically.

Based on the add-on manifest, bundles can be formed which can be integrated into the base application with zero configuration overhead on the user side. The structure of a bundle is largely dictated by the Java archive (JAR) file specification [Oraa]. A single JAR file

Figure 6: Add-on Description data model UML diagram

bundles the add-on and all dependent libraries. The manifest describing the add-on must be present in the `META-INF` directory with the name `addon.xml`.

4.2 Secure Retrieval and Execution

When a request message is received, the `AddonRegistry` (cf. figure 7) can be consulted to retrieve an applicable add-on for the requested resource. If an applicable add-on is found, it's JAR file will then be downloaded and a `ClassLoader` for subsequently loading the plug-in in a secure manner is returned. The `ClassLoader` will then be used in the factory responsible for the plug-in's type to load the class files.

Furthermore, a custom security policy implementation is set in the JRE and will therefore automatically be consulted every time a security relevant operation (e.g. reflection, classloader creation, filesystem access etc.) is performed. This policy allows to differentiate between signed add-ons, add-ons from a trusted origin and add-ons from an untrusted origin. Depending on the trust level, the add-ons may be granted different privileges.

By the use of privileged actions in a `AccessControler.doPrivileged()` call, trusted add-ons are permitted to call functions of the eID application that themselves do security relevant operations which the add-on would otherwise not have the appropriate rights for and therefore would fail.

Figure 7: Plug-in Manager and Registry UML diagram

5 Conclusion

The new add-on mechanism of the eID application proposed in the present paper provides an extensible framework which makes it easy to build tailormade eID and similar smart card based applications without re-developing basic functionality again from scratch. The proposed platform provides a set of well defined extension points and the initially provided modules ensure that existing installations can be utilized without modifications. With an App-Store like distribution method, it will be easy for third party vendors to provide their own add-ons. Paired with restrictive security measures, the App-Store model does not sacrifice the security and privacy of the user.

References

[Dre07]   Stefan Drees. Digital Signature Service Core Protocols, Elements, and Bindings, Version 1.0. OASIS Standard, 2007. `http://docs.oasis-open.org/dss/v1.0/oasis-dss-core-spec-v1.0-os.pdf`.

[FB96]   N. Freed and N. Borenstein. Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types. RFC 2046, November 1996. `https://www.ietf.org/rfc/rfc2046.txt`.

[Fed12a]   Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik). Advanced Security Mechanism for Machine Readable Travel Documents - Extended Access Control (EAC), Password Authenticated Connection Establishment (PACE), and Restricted Identification (RI). Technical Directive (BSI-TR-03110), Version 2.10, 2012. `http://docs.ecsec.de/BSI-TR-03110`.

[Fed12b]   Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik). eCard-API-Framework – IFD-Interface. Technical Directive (BSI-TR-03112), Version 1.1.2, Part 6, 2012. `http://docs.ecsec.de/BSI-TR-03112-6`.

[Fed12c]   Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik). eCard-API-Framework – ISO24727-3-Interface. Technical Directive (BSI-TR-03112), Version 1.1.2, Part 4, 2012. `http://docs.ecsec.de/BSI-TR-03112-4`.

[Fed12d]   Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik). eCard-API-Framework – Protocols. Technical Directive (BSI-TR-03112), Version 1.1.2, Part 7, 2012. `http://docs.ecsec.de/BSI-TR-03112-7`.

[Fed12e]   Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik). eCard-API-Framework – Support-Interface. Technical Directive (BSI-TR-03112), Version 1.1.2, Part 5, 2012. `http://docs.ecsec.de/BSI-TR-03112-5`.

[Her11]   A. Herrick. JSR 56: Java Network Launching Protocol and API. Maintenance Release 6, 2011. `http://jcp.org/en/jsr/detail?id=56`.

[Hor11]   Moritz Horsch. MONA – Mobile Authentication with the new German eID-card (in German). Master-Thesis, Technische Universität Darmstadt, 2011. `http://www.cdc.informatik.tu-darmstadt.de/mona/pubs/201107_MA_Mobile%20Authentisierung%20mit%20dem%20neuen%20Personalausweis%20(MONA).pdf`.

[HPS$^+$12] Detlef Hühnlein, Dirk Petrautzki, Johannes Schmölz, Tobias Wich, Moritz Horsch, Thomas Wieland, Jan Eichholz, Alexander Wiesmaier, Johannes Braun, Florian Feldmann, Simon Potzernheim, Jörg Schwenk, Christian Kahlo, Andreas Kühne, and Heiko Veit. On the design and implementation of the Open eCard App. In *Sicherheit 2012*, GI-LNI, 2012. `http://subs.emis.de/LNI/Proceedings/Proceedings195/95.pdf`.

[ISO]   ISO/IEC 7816. Identification cards – Integrated circuit cards – Part 1-15. International Standard.

[ISO08a]   ISO/IEC. Identification cards – Integrated circuit cards programming interfaces – Part 3: Application programming interface, ISO/IEC 24727-3. International Standard, 2008.

[ISO08b] ISO/IEC. Identification cards – Integrated circuit cards programming interfaces – Part 4: API Administration, ISO/IEC 24727-4. International Standard, 2008.

[Oraa] Oracle Inc. JAR File Specification. `http://docs.oracle.com/javase/6/docs/technotes/guides/jar/jar.html`.

[Orab] Oracle Inc. Java Cryptographic Architecture (JCA). `http://www.javasoft.com/products/jdk/1.2/docs/guide/security/CryptoSpec.hml`.

[Pet11] Dirk Petrautzki. Security of Authentication Procedures for Mobile Devices (in German). Master-Thesis, Hochschule Coburg, 2011.

[RLJ99] Dave Raggett, Arnaud Le Hors, and Ian Jacobs. HTML 4.01 Specification. W3C Recommendation 24 December 1999, 1999. `http://www.w3.org/TR/html401/`.

[The] The Legion of the Bouncy Castle. Bouncy Castle API. `http://www.bouncycastle.org/`.

[Wic11] Tobias Wich. Tools for automated utilisation of Smart-Card Descriptions. Master-Thesis, Hochschule Coburg, 2011.

Service providers' requirements for eID solutions: Empirical evidence from the leisure sector

Michael Kubach, Heiko Roßnagel, Rachelle Sellung

Fraunhofer IAO,
Nobelstr. 12
70569 Stuttgart
firstname.lastname@iao.fraunhofer.de

Abstract: Although eID technology has undergone several development cycles and eID have been issued to citizens of various European countries, it is still not as broadly used as originally expected. One reason is the absence of compelling use cases besides eGovernment. Current Research focuses mainly on the needs of the user and technical aspects. The economic perspective is often disregarded. This is especially the case for the service providers that play a fundamental role in the adoption of the technology. The requirements of these stakeholders certainly have to be considered in the development of viable business models. So far, however, little empirical evidence on these requirements exists. We therefore performed a survey-based empirical analysis in two industries from the leisure sector to gain first insights into this topic. Results show that the service providers in our sample don't see a pressing need to change their currently used authentication method. However, they think that certain eID features could be valuable for their services. Our analysis of the hurdles showed that there is no ultimate reason that keeps service providers from implementing the eID technology.

1 Introduction

eID (Electronic Identity) infrastructures have been implemented with various strategies and expectations in many European Union Member States. In some cases, Governments have issued eID cards on a large scale basis to their population [Eu12]. However, the actual usage lacks behind the original expectations and eIDs are still not used on an everyday basis [RZ12]. One of the reasons for this is the lack of real-world use cases and applications that are perceived as beneficial by users [Ro09]. Currently, most research efforts are focused on the technical aspects of the technology [St13] However, creating a technology that is only shaped by the technical aspect will not bring long term success. It is necessary to create a well-rounded product including the economical and societal aspects as well [ZR12]. In [ZiRo12] an economic analysis has illustrated how the relationships between users and relying parties are significantly influenced by indirect network effects. In detail, the indirect network effects lead to a common problem found in multi-sided markets, "the chicken or egg" problem [CJ03]. When there is a lack of a user adoption or a user base, the motivation for service providers to assist or implement a new product or service is insignificant. On the other side [ZR08] argue that "for a user to gain meaningful reduced sign-on capabilities across the web, a system has to be widely adopted, and its underlying protocol implemented by a wide range of service

providers". Therefore, the existence of promising services is an essential requirement for the success of eIDs. This requires co-operation by the service providers, which have to perceive a benefit in adopting eID technology. One of our main motives in this paper is to derive an empirical investigation in order to grasp a better understanding of the motives and needs and to identify potential roadblocks for the service providers. Our results will only show first insights of these perspectives, as our results will only reflect two industries in the leisure sector. As this side of the two sided marked has been mostly disregarded, we think that this examination can nevertheless serve as a starting point for further work. The article is organized as follows. Section 2 outlines the economic perspective on the eID technology in greater detail. In section 3 we describe our empirical analysis before presenting the survey results. We conclude in section 4.

2 An economic perspective on the eID technology

In this paper we take an economic perspective on eIDs and look at the service providers as important stakeholders for the success of the eID technology. Therefore it is important to look at the structure of the market first. The challenge eID solutions are facing is that the market is multi-sided. According to [Ev03] and [Ha07], "A market is said to be two-sided if firms serve two distinct types of customers, who depend on each other in some important way, and whose joint participation makes platforms more valuable to each. In other words, there are indirect network externalities between the two different customer groups." There are three major actors in the market for eIDs: the end-user, the service provider or relying party and the Identity Provider. We now see a multi sided market, where the success of the Identity Provider depends on the amount of users of relying parties. Furthermore, the relying parties themselves benefit, if the Identity Provider has a large installed user base. In addition, the attractiveness of the Identity Provider increases with the amount of available services that are perceived as beneficial by the users. This can result in a positive feedback and thus an exponential growth once a critical mass has been reached. However, all this can also happen the other way around, resulting in a negative feedback [ZiRo12] [MR99]. When no services are supporting eID, the usefulness for the user is presumably low. And when no users have adopted the product yet, service providers' motivation to implement it is quite minimal. Empirical analysis for the Alexa Top 300 websites seems to support this model for the relationship between relying parties and Identity Providers [LM12]. In order to utilize the full potential of eIDs, the technology needs to be adopted on a wide basis. As it is a multi-sided market, this will only be achieved if all participating parties perceive a benefit in adopting the technology. For the user-side, [RZHM14] have shown in their experimental analysis that there is indeed a willingness to pay for Identity Management Solutions. They find, however, that the acceptable price varies heavily from 3 to 48 Euros per year depending on the individual's psychographic and demographical aspects. The price aside, the authors also examine the importance of several other issues that are of potential importance to the users. Interestingly, they conclude that sophisticated privacy and security features are not valued by prospective users as much as suggested by previous research. As they focus on the user-side the perspective of the service providers is largely neglected. Our argumentation illustrates the significance of the service providers being just as important as the other stakeholders in the market for eID technology. Thus,

it seems to be reasonable to shed light on the motives and needs of the service providers in regard to the success of this technology – especially as there is so far a lack of research in this area.

3 Empirical analysis

3.1 Study design

We chose the method of a short quantitative survey to gain insight into a so far relatively unstudied field. Additionally, to integrate qualitative and quantitative methods as recommended by Gable [Ga94] for Information Systems Research, we are conducting qualitative semi-structured interviews with stakeholders in the research area to gain a deeper understanding. As the qualitative part of the research project isn't yet completed, we present only the quantitative results at this point.

Our target populations are the adult entertainment industry and the tourism industry. There are several reasons why we chose these two industries. The adult entertainment industry is quite heterogenic, consisting of numerous small and medium size enterprises as well as a few larger companies. As the vast majority of the companies are privately held, no reliable data on the exact value of the adult entertainment exists. Nevertheless, various sources speak of a quite substantial market, generating several billion dollars in revenue in the United States alone (the only market for which such figures can be found) [Do08] [Da13]. In Western Europe, the online market for adult content generated 540 million euros of revenue in 2009 [Bu09]. Moreover, users of adult content have proven to be very open to innovations in the past. Adult entertainment companies have pioneered innovations in the online and offline world such as VCRs, DVDs, or the internet itself. It has also been stated that the market demand for adult entertainment products is the main driver behind the success or failure of new technologies [Da13], [An07] ,. In the past, consumers of adult content have shown a high willingness to pay as well [AKTC06] [Co98]. For the study of the eID technology especially relevant is the specific aspect that users of adult content are very sensitive in terms of their privacy [CD01], making it a promising field for privacy enhancing credentials [HZPD09]. Thus, the inclusion of this industry into the target population of the survey can be justified. The second industry in the sample is the tourism industry. In various countries all over the world this is the leading industry sector. It is dominated by small and medium sized enterprises [WR04]. The tourism industry constitutes a natural use case for cross border identification and authentication services. This is because in many cases the transactions involve the demand for services outside the borders of the home country of the person requesting the service. Moreover, electronic, mobile and especially personalized location-based services have been becoming more and more popular in the tourism industry over the last years and provide the basis for a range of novel applications and business opportunities to service providers [RZ12]. Existing work has shown how information intermediaries are a suitable infrastructure component for offering services to an installed base of travelers, when they provide information from various sources in a concise manner. The availability of attractive services that are compliant with privacy

settings is according to this argumentation one factor that determines customer satisfaction [Bu98]. Examples could be mobile hotel reservation services provided by local authorities to promote the local tourism sector or other location-based services provided. For these services some form of reliable authentication or identification is necessary [RZ12] [SMR09] As this can be made possible with the eID technology, the tourism industry is a quite suitable sector for this survey. To maximize the response rate and reduce nonresponse bias, the short survey was designed according to various recommendations that can be found in the literature on surveys in the corporate context [Di07] [HH04] [RL98]. Potential respondents were found on two International trade fairs in late 2012 and early 2013 and in online databases. Questionnaires were distributed on the fairs as paper versions and as a link to a digital version via e-mail. Follow-up e-mails were used to raise the response rate. Through this approach we received a total of 56 usable questionnaires. The data were analyzed using SPSS. The profile of the respondents included in our survey is shown in Figure 1 to Figure 3. The vast majority of the companies' headquarters of the respondents is located in Europe. More than half of them (55 percent) are located in Germany. About two thirds of the companies offer their services internationally (45 percent globally, 25 percent all over Europe). About one third is focused only on the German market. This makes clear that it can be assured that the eID technology is available for most (if not all) of the companies in the sample, as Germany and other European countries have already issued eIDs to their population [Eu12].

Figure 1: Country statistic of sample companies

Figure 2 shows the size characteristics of the sample companies. Our data reflects what has been written before about the structure of both industries. Very small and small companies constitute the largest groups in our sample.

Figure 2: Size statistics of the sample companies

While 42 percent of the companies employ less than 10 employees and about two thirds of the companies achieve just 0 to 2 million Euros in turnover, only 10 percent of the companies have more than 250 employees and 12 percent of the companies achieve more than 50 million Euros in turnover. The distribution into the focus sectors "Tourism" and "Adult Entertainment" are quite balanced. This will later allow us to compare the results from both sectors. The hierarchical position of the respondents permits us to see them as key informants with sufficient expertise and insight into the topics in question. The key informant approach is a well-established method for conducting survey-research [HKRS12]. We can conclude that for a preliminary study the sample is relatively balanced and suitable to give us first insights into a relatively unexplored topic.

Figure 3: Industries and functional areas represented

3.2 Study results

In the following, respondents' profiles in terms of the frequencies in which they encounter user errors and the types of errors are discussed first. Then, results on the authentication methods used are presented. Afterwards, we compare the requirements for eID solutions and the hurdles to their implementation. Finally we show the costs they expect to arise from the transition to eID technology. To get a first impression of the empirical evidence for the research field, data was analyzed using frequency statistics and for the Likert-type scale items using analysis of means. Moreover, the results of the two focus sectors were compared. We show where this revealed substantial differences.

Figure 4: Frequency of (user generated) errors in user authentication management[1]

Problems in user authentication management don't seem to be a major difficulty for the respondents (Figure 4). Only about 12 percent agree or strongly agree to the general statement „We often encounter problems in our user authentication management". Focusing more on user generated errors, e.g. incorrect user address, only about nine percent of the respondents agree to the given statement. Judging from these results, it appears that user authentication management doesn't seem to be a major issue in our sample industries.

Figure 5: Different types of problems encountered due to user errors

Looking at the different types of user errors we notice certain differences between the sample industries (Figure 5). All in all, typing errors seem to be the most important cause of error, especially for the tourism industry where almost two thirds of the cases report this problem. However, for the adult entertainment industry this problem is a major problem as well, with almost half of the cases reporting it. The second biggest cause of error is password loss. Here we can recognize an even larger contrast between the two industries, more than half of the respondents from the adult entertainment industry report this problem, whereas only 13 percent of the respondents from the tourism industry do so. The next two types of problems are interpretation error and false information, both of which are reported from about one quarter to about one fifth of the respondents. Therefore, they seem to be of some relevance, but not of major relevance in this context. The striking differences between the two sectors with the first two types of errors can be explained through the consideration of the authentication methods that are

---

[1] In order to facilitate the readability and as the results don't vary significantly, Figure 4 shows the numbers for both industries combined.

used for the services provided (Figure 6). Here we see that passwords are only rarely used in the tourism industry (14 percent of the cases) whereas most services can be used without any necessary authentication (91 percent of the cases). In the adult entertainment industry, however, we find the opposite situation. The authentication with username and password are possible or even required for about two thirds of the services, while only about a quarter can be used without any authentication. It comes as no surprise that for firms that don't require the authentication with username and password, lost passwords don't pose a major problem. Therefore, as this authentication method is more widely used in the adult entertainment industry as compared to the tourism industry, the former industry sees lost passwords as a bigger problem. Login with Facebook is somehow common in the tourism industry, with little more than one quarter of the services in the sample offering this method. However, it is not surprising that no respondent from the adult entertainment industry offers this method for his service, as customers might be hesitant to connect an adult entertainment website with the social network of their friends and family. One fifth of the adult entertainment services offer the possibility to authenticate with a national eID card, whereas in the tourism sample this method is non-existent. Here we might see that the adult entertainment industry is indeed very open to new technologies (see above). The same might apply to the methods Google ID / Open ID that are supported by 13 percent of the services from the adult entertainment industry in the sample.

Figure 6: Authentication methods for services provided

The results above show that eID solutions are not supported by a substantial number of services in the sample (only 20 percent in the Adult entertainment Industry which corresponds to 6 percent of the valid total cases). Moreover, the service providers in our survey do not report to have lots of problems with their currently offered authentication methods. Therefore, to contribute to a broader success of eIDs, it could be valuable to assess what eID-features they would see as beneficial in order to develop these features further and promote eIDs by highlighting these aspects. Results from the part of the survey, that assessed the importance of certain eID features for the service providers' services, again show some differences between the two industries in the sample (Figure 7). Unsurprisingly, having a certified age verification is of very high importance (average score of 6.4) for the respondents from the adult entertainment industry. Here the eID technology could be of great value. For the tourism industry, this feature doesn't have that much relevance (average score of 4.9), but still seems to be desirable.

Respondents from both industries regard the possibility to certify the user address as quite important (average scores of 6.1 and 5.8). The certification of user age and user address appear to be the most important features for the respondents. Next, we asked for the importance of the feature that allows offering services to the user anonymously. This doesn't seem to be as important for the respondents, somehow more to the adult entertainment industry (average score of 4.9 as compared to the tourism industry (4.3). One could have expected that this feature would be of higher relevance in the adult entertainment industry. The feature that allows obtaining certified information about users' attributes is only of medium importance to the respondents of both industry sectors (average scores of 4.5). Of greater importance is the possibility to achieve higher transaction security, which can be achieved with eIDs. This feature is even a bit more important to the respondents from the tourism industry (average score of 5.6) than to those from the adult entertainment industry (average score of 5.1).

Figure 7: Importance of eID features for service providers' services

The outsourcing potential that is provided through eID solutions is the feature with the lowest importance by the respondents. For the adult entertainment industry (average score of 2.9) even less than for the tourism industry (average score of 3.7). More important, but still not of great importance, seems to be the possibility to approach an already established user network. For this feature, as for the next two, no substantial differences between the two industries can be found (adult entertainment industry 4.5, tourism industry 4.9). About the same importance scores achieves the possibility to reduce user support costs (adult entertainment industry 4.9, tourism industry 4.5). Finally, as the last feature the ability to enable legally binding transaction was assessed. Although one could expect this feature to be of high importance for the respondents, it achieves only fairly mediocre scores (adult entertainment industry 4.8, tourism industry

5.0). The last paragraphs made clear that the respondents indeed think that eIDs can provide some useful features for their services. At the same time, the diffusion of the eID technology remains low. Therefore, it seems to be reasonable to look at the hurdles to the implementation of eID technology (Figure 8). This shows even greater differences between the two focus sectors than when looking at the features. The respondents from the tourism industry score the hurdles higher than the respondents from the adult entertainment industry. At first we asked for the familiarity with eIDs. On average, respondents from both industries claim to have a medium knowledge. The level of unfamiliarity with the technology in the adult entertainment industry (average score of 3.6) is, however, substantially lower than in the tourism industry (average score of 4.7). Maybe this can be connected to the next finding that shows that the tourism industry perceives no need to change its authentication services (average score of 5.6), whereas this score is considerably lower for the adult entertainment industry (average score of 4.7).

Figure 8: Hurdles to the implementation of eID technology

The adult entertainment industry might be more open to innovations (see above), be better informed about the valuable features of eIDs (see study results above), or due to its specific need for age verification be unsatisfied with the currently applied authentication methods. Nevertheless, it has to be noted, that the overall need to change the authentication service is only mediocre to low. Looking at the expectation that the respondents have concerning the inconvenience of the transition to eIDs we see as well

that this hurdle is scored substantially higher in the tourism (average score of 4.5) compared to the adult entertainment industry (average score of 3.4). High costs are also seen as a bigger problem in the tourism industry (average score of 4.9) compared to the adult entertainment industry (average score of 3.5). In terms of the usability, both industries have the same impression of eIDs (average score of 3.8). Thus, a lack of usability doesn't seem to be an essential problem for neither of the industries in the sample. Compared to that, the additional hardware that is needed by the users is seen as a somewhat higher hurdle. However, with scores of 4.4 (tourism) and 4.5 (adult entertainment industry) this shouldn't be decisive. Next, we wanted to see whether the service providers expect the users to demand the support of eIDs. The respondents of the tourism industry are a bit more skeptical towards the user demand than the respondents from the adult entertainment industry. The average score for the item that says that there won't be a high demand is 4.7 for the tourism and 4.3 for the adult entertainment industry. This shows that the end users are not seen as a major driver for the diffusion of the eID technology. Another interesting aspect could be, if the respondents feel not comfortable with the eID technology. The results show, however, that this isn't a substantial hurdle for the average respondent from both industries (3.8 for the tourism and 3.6 for the adult entertainment industry). However, in total 20 percent of the respondents partially or strongly agree to the statement that they are not comfortable with the technology. So partly, for some respondents, the negative attitude towards eIDs could play a role. Especially for the respondents from the tourism industry an uncertainty about regulations regarding eIDs seems to be a problem for the adoption of this technology (average score of 5.2). Interestingly this aspect is of much less importance for the respondents from the adult entertainment industry (average score of 3.2). The fear of a reduced ability to acquire information about users isn't a major issue for the respondents (3.8 for the adult entertainment industry and 3.6 for the tourism industry). Respondents also don't fear a reduction of the quality of user data respondents (3.7 for the adult entertainment industry and 2.5 for the tourism industry). Finally, we asked for the expected costs for the transition to eIDs. Here we could not find any substantial differences between the two industries. The vast majority of the respondents (about 61 percent) did not have any idea of the costs. The second largest group (about 16 percent), however, expected costs of about 3.000 – 10.000 EUR. Judging from our experience with past projects, we would say that this is price range is relatively realistic.

Figure 9: Expected costs for the transition to eIDs

4 Conclusion

Our empirical analysis of the requirements that service providers in the tourism and adult entertainment industry have for eID solutions has revealed several issues. The service providers in our sample don't see a pressing need to change their currently used authentication method to support eID technology – which seems understandable, as they don't encounter many authentication problems at the moment. However, they think that certain eID features could be valuable for their services. In particular, for the adult entertainment industry the possibility to obtain certified information is seen as a plus. An increase in transaction security is also seen positively. Thus, further development and marketing of the eID technology, as well as corresponding sample business cases should especially focus on these features. Our analysis of the hurdles showed that there is no ultimate reason that keeps service providers from implementing the eID technology. Only a few aspects from the tourism industry stand out. Respondents from this industry don't really see a need to change their services and are uncertain about the regulations in the eID context. They also fear that the transition might be inconvenient and expensive. All in all, the adult entertainment industry is more positive towards the eID technology. This fits quite well, as this industry is often an early adopter of new technologies. Focusing use cases and marketing on service providers of this industry when trying to introduce eIDs could therefore maybe help to establish the technology. The positive feedback resulting from the indirect network effects through an established base of users and service providers could then help to solve the "chicken and egg problem" and encourage the diffusion to other industries. As with any research approach our findings are subject to certain limitations. Noteworthy to this quantitative survey study are especially the comparatively small sample size, and the sample selection, which might lead to a possibly biased sample. The representativeness of the sample for the population of the two industries might therefore be reduced. However, we made the structure of our sample transparent and this study is only a first empirical analysis of the subject. The industry focus also has to be highlighted as a possible limitation, as the findings can't or only to a certain extend get generalized to other industries. So while these limitations are legitimate concerns given the chosen research methodology, they should not impact this papers ability to give a first empirical insight into the subject. Moreover, they open possibilities for further research. It would certainly be valuable to extend the analysis to other industries and to extend the sample size. It also seems advisable to contrast the quantitative findings of this study with qualitative data, generated through in-depth interviews. The results of this survey can serve as a basis for this approach.

References

[AKTC06] Ang X, Kwan WY, Teo CP, Chua C. Why do people pay for information goods: a study of the online pornography industry. AMCIS 2006 Proc [Internet]. 2006; Available from: http://aisel.aisnet.org/amcis2006/9

[An07] Angell IO. Ethics and Morality: a business opportunity for the Amoral. J Inf Syst Secur. 2007;3:3–18.

[Bu98]   Buhalis D. Strategic use of information technologies in the tourism industry. Tour Manag. 1998 Oct;19(5):409–21.

[Bu09]   Bundesministerium für Wirtschaft und Technologie. 12. Faktenbericht 2009: Eine Sekundärstudie der TNS Infratest Business Intelligence. Berlin: Bundesministerium für Wirtschaft und Technologie; 2009.

[CJ03]   Caillaud B, Jullien B. Chicken & Egg: Competition among Intermediation Service Providers. RAND J Econ. 2003;34(2):309.

[Co98]   Coppersmith J. Pornography, Technology, and Progress. Icon. 1998;4:94–125.

[CD01]   Cronin B, Davenport E. E-rogenous zones: Positioning pornography in the digital economy. Inf Soc. 2001;17:33–48.

[Do08]   D'Orlando F. The Market for Pornography in Italy: Empirical Data and Theoretical Considerations. 2008;

[Da13]   Darling K. What Drives IP without IP? A Study of the Online Adult Entertainment Industry [Internet]. Cambridge; 2013. Available from: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2198934

[Di07]   Dillman DA. Mail and Internet Surveys: The Tailored Design Method. Hoboken: John Wiley & Sons; 2007.

[Eu12]   European Commission. Electronic identification, signatures and trust services: Questions & Answers. Brussels: European Commission; 2012.

[Ev03]   Evans DS. The Antitrust Economics of Two-sided Markets. Yale J Regul. 2003;20(2):235–94.

[Ga94]   Gable G. Integrating case study and survey research methods: an example in information systems. Eur J Inf Syst. 1994;3(2):112–26.

[Ha07]   Hagiu A. Merchant or two-sided platform? Rev Netw Econ. 2007;6(2):115–33.

[HH04]   Hague P, Hague N, Morgan C-A. Market research in practice: a guide to the basics. London, Sterling: Kogan Page; 2004.

[HKRS12] Homburg C, Klarmann M, Reimann M, Schilke O. What Drives Key Informant Accuracy? J Mark Res. 2012;49(4):594–608.

[LM12]   Landau S, Moore T. Economic tussles in federated identity management. First Monday. 2012;17(10).

[MR99]   Mahler A, Rogers E. The diffusion of interactive communication innovations and the critical mass: The adoption of telecommunications services by German banks. Telecommun Policy. 1999;23(10/11):719–40.

[RL98]   Rogelberg SG, Luong A. Nonresponse to Mailed Surveys: A Review and Guide. Curr Dir Psychol Sci. 1998;7(2):60–5.

[RZHM14] Roßnagel H, Zibuschka J, Hintz O, Muntermann J. Users' willingness to pay for web identity management systems. Eur J Inf Syst. 2014;forthcomin:1–35.

[HZPD09] Roßnagel H, Zibuschka J, Pimenidis L, Deselaers T. Facilitating the Adoption of Tor by Focusing on a Promising Target Group. Identity Priv Internet Age Nord 09. Springer; 2009. p. 15–27.

[RZ12] Roßnagel H, Zibuschka J. eID in Leisure Time Activities: Results from the SSEDIC Stakeholder Consultations in the Leisure Sector. Stuttgart; 2012.

[Ro09] Roßnagel H. Mobile qualifizierte elektronische Signaturen. Gabler; 2009.

[SMR09] Scherner T, Muntermann J, Rossnagel H. Integrating value-adding mobile services into an emergency management system for tourist destinations. 17th Eur Conf Inf Syst [Internet]. 2009. p. 1–13. Available from: http://aisel.aisnet.org/ecis2009/327/

[St13] STORK 2.0 - Secure idenTity acrOss boRders linKed 2.0 [Internet]. 2013. Available from: https://www.eid-stork2.eu/

[WR04] Werthner H, Ricci F. E-commerce and tourism. Commun ACM. 2004;47(12):101–5.

[ZR08] Zibuschka J, Roßnagel H. Implementing Strong Authentication Interoperability with Legacy Systems. In: Leeuw E, Fischer-Hübner S, Tseng J, Borking J, editors. Policies Res Identity Manag [Internet]. Springer US; 2008. p. 149–60. Available from: http://dx.doi.org/10.1007/978-0-387-77996-6_12

[ZR12] Zibuschka J, Roßnagel H. A Structured Approach to the Design of Viable Security Systems. ISSE 2011 - Secur Electron Bus Process Highlights Inf Secur Solutions Eur 2011 Conf. Wiesbaden: Vieweg+Teubner; 2012. p. 246–55.

[ZiRo12] Zibuschka J, Roßnagel H. Stakeholder Economics of Identity Management Infrastructures for the Web. Proc 17th Nord Work Secure IT Syst Nord 2012. Karlskrone, Sweden; 2012.

An Open eCard Plug-in for accessing the German national Personal Health Record

Raik Kuhlisch[1] · Dirk Petrautzki[2] · Johannes Schmölz[3] · Ben Kraufmann[1]
Florian Thiemer[1] · Tobias Wich[3] · Detlef Hühnlein[3] · Thomas Wieland[2]

[1] Fraunhofer FOKUS, Kaiserin-Augusta-Allee 31, 10589 Berlin, Germany
{raik.kuhlisch,ben.kraufmann}fokus.fraunhofer.de

[2] Hochschule Coburg, Friedrich-Streib-Str. 2, 96450 Coburg, Germany
{petrautzki,thomas.wieland}hs-coburg.de

[3] ecsec Gmbh, Sudetenstraße 16, 96247 Michelau, Germany
{johannes.schmoelz,tobias.wich,detlef.huehnlein}@ecsec.de

Abstract: An important future application of the German electronic health card (elektronische Gesundheitskarte, eGK) is the national Personal Health Record (PHR), because it enables a citizen to store and retrieve sensitive medical data in a secure and self-determined manner. As the stored data is encrypted with an eGK-specific certificate and retrieving the encrypted data is only possible after TLS-based authentication, the citizen needs to use a so called "PHR Citizen Client", which allows to use the eGK for strong authentication, authorization, and decryption purposes. Instead of building such an application from scratch, this paper proposes to use the Open eCard App and its extension mechanism for the efficient creating of a PHR Citizen Client by developing an Open eCard Plug-in for accessing the German national Personal Health Record.

1 Introduction

Usually the implementation of smart card based applications is fairly challenging, because it requires a profound knowledge of several technologies: communication with diverse card terminals, reading and writing data on various smart card models and performing cryptographic operations for authentication, signature and encryption. Furthermore, such a smart card based application should remain usable and secure if new types and generations of smart cards and terminals as well as operating systems come into market.

Due to several emerging application domains around e-government and e-health in Germany based on the electronic identity card (neuer Personalausweis, nPA) or the electronic health card (elektronische Gesundheitskarte, eGK), providing smart card based applications is a key objective for many software vendors. Against this background an important question is how software vendors can provide smart card based

applications *without* focusing on smart card technology and security infrastructures but mainly on their application-specific expertise.

In 2005 the Federal Government of Germany launched the eCard strategy [HoSt11] – a plan that aims at harmonizing different smart card based governmental IT projects. Based on this preparatory work, the specifications developed by the German Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik) [BSI-TR-03112] and first-hand experiences from projects in the context of the electronic health card, the project ID4health was started. The project group of ID4health develops a comprehensive architecture for federated identity management for the public health sector. Part of this architecture is a set of secure services that allow for using the electronic health card and the electronic health professional card. These services rely on the platform independent open source implementation of the eCard-API framework – the Open eCard App [HPS+12].

In order to demonstrate the feasibility of the ID4health approach, the German national Personal Health Record (PHR) was chosen. In this paper we present our approach on how to securely access electronic health information about individual patients using the electronic health card and the extensible Open eCard platform.

The rest of the paper is structured as follows: In Section 2 we briefly introduce the German national PHR according to the health card act, highlight the key aspects, and give a short overview of its architecture and security aspects. In Section 3 we briefly describe the extensible Open eCard platform including the architecture, the interfaces, and sketch how the extension mechanism works. In Section 4 we present our approach to integrate the Open eCard App with a PHR application for the citizen to fulfill the security requirements of accessing a PHR without implementing the security mechanisms in the PHR-specific desktop or mobile application. In Section 5 we discuss our solution and provide an overview of the next steps.

2 German national Personal Health Record

In 2009 the German Federal Ministry of Health initiated an R&D project on the design and prototypical implementation of a German national PHR. Project participants were the Fraunhofer Society, the Federation of German Clinical Research Networks (TMF), and the University Medical Center Göttingen (UMG). Associated contractors have been the German Hospital Federation (Deutsche Krankenhausgesellschaft, DKG), the German Medical Association (Bundesärztekammer, BÄK), and the Society for Telematic Applications of the Healthcare Card (gematik) which designs and provides the infrastructure for health data services usable with the health card. Currently, the R&D project is in its second phase, covering the three-year period from 2012 to 2014.

Major requirements of the special kind of an electronic health record (EHR) included advanced security and privacy mechanisms and a generic platform semantics that allow for operating various patient-centric e-health applications within the PHR. Another

prerequisite was to strongly focus on innovative care scenarios, stable business models, and better patient participation.

2.1 Focus areas

The German health card act (§ 291a SGB V) mentions the electronic health card as an enabler for patient records. The PHR project assesses implementation options and puts the focus on use cases and scenarios that do not solely rely on the defaults of the health card act. That is why additional use cases stretching beyond the health card act with a particular focus on enabling medical research are elaborated, too. The patient record is regarded as a platform for patient centric applications (medication plan, diary etc.). Security aspects play an important role in order to ensure privacy and patient rights. The last thing put into focus is the information model that must support semantic data selection for different applications.

2.2 Key concepts

The PHR project pursuits several key concepts that are introduced subsequently.

2.2.1 Power of disposal of the citizens' data

The basic idea of the PHR is that it is understood as a "tool of the citizen" (the project does not speak of a patient since it is expected that the citizen can store its own data that is outside of a running treatment). The citizen does not regularize the flow of data between systems of health service providers, but initializes and controls it. Essentially, the citizen acts both as a source and target of flows of data and is furthermore a beneficiary of his medical data. The approach taken by the project members does not define a virtual integration of the data, but rather a transmission of copies from medical data into the PHR.

2.2.2 Platform for target group specific record systems

Every citizen has specific needs with regard to a concrete definition of his record. Needs result from his current living situation and mirror in expectations of record features. Vendors should be able to provide target group specific records systems and product variants. This in mind, the PHR has to have support of different interaction and communications patterns (i.e., synchronously and asynchronously communication – an online PHR supports solely synchronously patterns whereas an offline one must support asynchronous communication comparable to email) as well as different authorization mechanisms (ad hoc, in advance, interactively). This offers maximum design flexibility for vendors.

2.2.3 Binding of online records and decentralized data storage

It was recognized that a citizen might want to store data on own decentralized storage devices (e. g., USB flash drives) or rely on secure online storage. In order to meet the individual requirements of the citizens, the project provides a solution for both the binding of online records and decentralized storage devices. The difference is abstracted away from the primary system of the health care provider that accesses the PHR.

2.2.4 Communication of normalized information objects

Besides normalized interfaces, medical data objects need to be interoperable, too. So the last key concept defines a model for a unified handling of different content structures based on Semantic Signifiers developed by the Retrieve, Locate and Update Service (RLUS) standard [RLUS]. This functional standard developed by the Object Management Group (OMG) was chosen to be used on top of a service-oriented security architecture. The introduced level-based concept for step-by-step implementation of normalized content structures (e.g., for a medication summary or a vaccination certificate) relies on the Clinical Document Architecture (CDA) by Health Level Seven (HL7) [HL7 CDA] for the interoperable processing through primary and record systems. Fallback documents (e.g., PDF/A as a standardized version of a Portable Document Format (PDF)) are supported, too.

2.3 Architectural overview

As a design rule, the project idea separates application services from security aspects. This enables reuse of functionalities in other application contexts and provides flexibility for operator models. Security services comprise a central authentication service, a local authorization service, and PKI services. Application services are the mailbox service for communications initiated by the citizen and the communication service that might cache requests from health service providers for offline PHRs.

On the patient side RLUS services encapsulate existing PHR systems (e.g., Microsoft's HealthVault) or simple medical data storage devices (e.g., USB cards). The PHR Core System and the Citizen Client are non-normative and their design/implementation is solely in the responsibility of the respective operators.

Figure 1: Architectural Overview of German national PHR system

The exchange of medical data is encapsulated in two container formats: Request objects are used to formulate specific information needs. This need is logically expressed as a dedicated object to address the significant decoupling of systems of health service providers and citizens. Request objects must be interpreted by the primary systems. Contrary, a provisioning object represents a logical container for any information provided in the communication flow between the health service provider systems and the PHR. A provisioning object can be regarded as a response to a previously asked information request (synchronous and asynchronous) submitted either by the health service provider or the citizen.

From the health service provider perspective the RLUS operations are safeguarded by SAML assertions [SAML] that carry authentication and authorization data. The architecture allows for both pushing and pulling XACML policies [XACML]. Pushing of policies is used for ad hoc authorization with an electronic health card. In this case an XACML policy is generated on demand by the citizen and pushed in a SAML assertion to the RLUS services within the SOAP security header.

2.3.1 Semantic signifiers

For interoperability reasons, classification and description of the (medical) data must be available in order to ensure the processing in primary systems. A PHR must identify

requested data and must be able to transfer local data into interoperable return types. This is due to the fact that the exchange format can differ from the format stored in the PHR – the exchange format may not be exactly the same as the one stored in the PHR. Search requests and filters should point to the specific information model depending on the content type. These demands feature the so-called semantic signifiers specified by RLUS that have an analogy to a WSDL document for a web service description. They describe an underlying information model instead of services. A semantic signifier's essential elements include its name, its description, and a normative data structure that describes instances of it – for example, implementation guidelines, schemas, and specifications for validation.

From the implementation independence view there are no explicit requirements for data storage and information models. However, each vendor must implement functionality for the mapping between internal and external structures defined by semantic signifiers. A further advantage is that any contents and structures might be described. Conversely, each semantic signifier must be explained in detail in order to prevent unchecked distribution.

2.3.2 Capability list

Similar to the capability object that is stored in the inbox of an OBEX (Object Exchange Protocol) server in order to provide capability services for Bluetooth clients, the capability list contains information about the features the PHR provides. This XML-based representation of a capability list includes address information with an Endpoint Reference according to WS-Addressing [WS-Add], a supported public key to be used for hybrid encryption of provisioning objects according to XML Key Management Specification (XKMS) [XKMS], supported communication patterns, and supported semantic signifiers. The capability list is digitally signed by the issuing service.

3 Open eCard App

The Open eCard App is a fully modular and platform-independent implementation of the eCard-API framework [BSI-TR-03112] which can be further extended by plug-ins and add-ons. The architecture of the Open eCard App is depicted in Figure 2.

Access to card terminals, smart cards and the functions they make available is provided by the following three components of the Open eCard-Framework:

- The *Interface Device (IFD)* module provides a generalized interface to access arbitrary card terminals and smart cards and an interface for password-based authentication protocols (e.g., PACE [BSI-TR-03110]). It is specified in part 6 of [BSI-TR-03112] and part 4 of [ISO24727].

- The *Event Manager* monitors the IFD for incoming events (e.g., insertion or removal of smart cards) and triggers the card recognition. Registered components are notified by the Event Manager if an event occurs.

- The *Service Access Layer (SAL)* manages all available smart cards and executes all kinds of protocols (e.g., the Extended Control Access protocol [BSI-TR-03110]). It is specified in part 4 of [BSI-TR-03112] and part 3 of [ISO24727].

Figure 2: Extensible architecture of the Open eCard App

The Open eCard platform provides different extension points to support newly issued smart cards, new protocols and new application specific functions. A detailed description of the architecture as well as the plug-in mechanism of the Open eCard platform can be found in [WHP+13].

The PHR plug-in described in the next section will basically utilise the SAL for all operations that include access to the electronic health card, support the establishment of TLS channels with client authentication using the health card and provide a simple application interface which allows to access the plug-in via the localhost binding. In this binding, which has been introduced in part 7 of [BSI-TR-03112] for purposes of eID-activation, a function is called by performing an HTTP GET request to http://127.0.0.1/ (localhost) and parameters may be provided in URL-encoded form. See [WHP+13] for more information on application plug-ins and corresponding binding mechanisms.

4 Approach

In order to demonstrate the feasibility of using the Open eCard platform for health care applications, we consider the following scenario as depicted in Figure 3:

Figure 3: Solution scenario for using the Open eCard App to access the PHR system

A citizen wants to access his PHR via his desktop computer or mobile device. The citizen has his electronic health card and a smart card reader, which is connected by means of USB for example.

The PHR Core System is operated in a secure environment by a PHR provider. The environment can be accessed via an RLUS interface over a secure point-to-point data channel. This data channel is established by a mutual authentication which is performed using the PHR provider's certificate and the X.509-certificate (C.CH.AUTN) deployed with the citizen's health card.

4.1 Architecture

The linkage of a citizens PHR application with the PHR Core System is realized by using the Open eCard platform and developing a plug-in which

1. uses the smart-card functionality and cryptographic features of the framework and

2. provides a PHR-specific application interface, which can be used by application services.

The following layer model depicts the relationship of the logical components that ensure the secure communication of PHR applications and PHR core system using the extended Open eCard App.

Figure 4: Accessing the PHR Core System with the extended Open eCard App

The system comprises three main components:

1. PHR Application, which provide tailormade user interfaces

2. Open eCard Framework with PHR specific plug-in

3. PHR Core System, which in particular contains the repository service.

The PHR Core System provides the PHR client with personal medical data. The PHR client is realized by a PHR-specific plug-in for the Open eCard platform. It has an

interface following the RLUS standard [RLUS] and exchanges messages via a mutually authenticated TLS channel.

Third-party applications can utilise the PHR plug-in to store and retrieve information from the PHR Core System. For this purpose the PHR plug-in provides two operations: the `getInformationObject` function for information retrieval and the `putInformationObject` function for information storage. Both functions are available through the localhost binding of the Open eCard framework.

The information retrieval from the PHR Core System works as follows:

A third-party application calls the `getInformationObject` function with a parameter identifiying a certain semantic signifier that classifies the requested information. If the electronic health card is inserted in the card terminal the personal identification number is requested in order to read a special container (EF.Verweis) on the card that stores the service reference to the PHR Core System. Otherwise the citizen will be prompted to insert the smart-card. The service reference includes a provider ID as well as a record ID. The provider ID can be used to determine the service location of the PHR Core System. The plug-in establishes a mutually authenticated TLS channel to the PHR Core System using the C.CH.AUTN certificate, which is an X.509 certificate stored on the health card carrying a pseudonym of the citizen. The capability list of the PHR is initially requested which includes the supported record types (i.e., semantic signifiers) as well as a the public key of the record instance which can be used to protect symmetric document keys that in turn encrypt information objects submitted by the PHR plug-in to the PHR Core System. Afterwards, the specified content, addressed by a parameter of the `getInformationObject` function, is requested from the PHR Core System using the RLUS LIST operation which returns the requested data as an encrypted provisioning object. The plug-in decrypts the retrieved provisioning object by means of the health card's key material and returns it to the calling application as a Base64-encoded string.

To store information in the PHR Core System, the following steps are performed:

A third-party application invokes the `putInformationObject` function with the Base64-encoded information object and a semantic signifier. The plug-in decodes the information object and generates a symmetric document key. In order to secure the information object, the plug-in encrypts it by means of that document key. The document key is secured using the public key of the capability list (which was retrieved by a former call). Finally, both the (symmetrically) encrypted information object and the associated encrypted (symmetric) data key are used to create a provisioning object which is embedded as a payload in a RLUS PUT message and sent to the PHR Core System in order to store the information object.

The Open eCard App implements all cryptographic components and integrates smart card terminals, which may or may not include a display and key pad for secure PIN entry. The PHR plug-in mirrors a much simpler version of the PHR Core System

interface to the PHR client-application and provides high-level operations for encryption and signing based on the smart card currently in use.

5 Summary and conclusion

The present paper briefly described how the extensible Open eCard platform can be used for securely accessing sensitive medical information, which is stored in the German national Personal Health Record according to § 291a SGB V. This demonstrates not only the feasibility of the ID4health approach but also underlines the power of the extensibility of the Open eCard platform. As the supported extension mechanisms are both independent of the used smart card and the implemented application logic, it may be expected that this approach can easily be carried over to entirely different application scenarios.

Future work will include the support of record initialization (including key generation and secure transmission) as well as the investigation of the complementary use of the identity card for health professionals in consideration of signing and decrypting medical data. This is very useful when the Open eCard plug-in is deployed in primary systems of health service providers. Such scenarios require e.g. the issuing of authorization assertions that are digitally signed by the electronic health card (named ad-hoc authorization). Hence, the access control system of the PHR Core system can verify the eligible use of the patient data by health professionals.

References

[BSI-TR-03110]   Bundesamt für Sicherheit in der Informationstechnik (BSI): Advanced Security Mechanism for Machine Readable Travel Documents – Extended Access Control (EAC), Password Authenticated Connection Establishment (PACE), and Restricted Identification (RI), Technical Directive TR-03110, Version 2.10, 2012

[BSI-TR-03112]   Bundesamt für Sicherheit in der Informationstechnik (BSI): eCard-API-Framework. Technical Directive TR-03112, Version 1.1.2, Part 1-7, 2012

[HL7 CDA]   ISO/HL7 27932:2009: Data Exchange Standards - HL7 Clinical Document Architecture, Release 2

[HPS+12]   Hühnlein, D.; Petrautzki, D.; Schmölz, J.; Wich, T.; Horsch, M.; Wieland, T.; Eichholz, J.; Wiesmaier, A.; Braun, J.; Feldmann, F.; Potzernheim, S.; Schwenk, J.; Kahlo, C.; Kühne, A.; Veit H.: On the design and implementation of the Open eCard App, In Sicherheit   2012,   GI-LNI   (2012). http://subs.emis.de/LNI/Proceedings/Proceedings195/95.pdf

[HoSt11]   Horsch, M.; Stopczynski, M.: The German eCard-Strategy, CDC Report,   TU   Darmstadt,   February   2011, https://www.cdc.informatik.tu-darmstadt.de/en/publications-details/?no_cache=1&tx_bibtex_pi1%5Bpub_id%5D=TUD-CS-2011-0179

[ISO24727]   ISO/IEC 24727: Identification cards – Integrated circuit cards programming interfaces, Part 1-6, International Standard, 2008

[RLUS]   OMG: Retrieve, Locate, and Update Service (RLUS) Specification, Version 1.0.1., 2011.

[SAML]   Cantor, S.; Kemp, J.; Philpott, R.; Maler, E.: Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V2.0., OASIS Standard, 15.03.2005, http://docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf, 2005

[TLS]   Dierks, T.; Rescorla, E.: The Transport Layer Security (TLS), Protocol Version 1.2. Request For Comments – RFC 5246. http://www.ietf.org/rfc/rfc5246.txt, August 2008

[WHP+13]   Wich, T.; Horsch, M.; Petrautzki, D.; Schmölz, J.; Hühnlein, D.; Wieland, T.; Potzernheim: An extensible platform for eID, signatures and more, in the present proceedings

[WS-Add]   Gudgin, M.; Hadley, M; Rogers, T.: Web Services Addressing 1.0 – Core, 2006, http://www.w3.org/TR/ws-addr-core/

[XACML]   Moses, T.: XACML 2.0 Core: eXtensible Access Control Markup Language   (XACML)   Version   2.0,   http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf

[XKMS]   Hallam-Baker, P.; Mysore, S. H.: XML Key Management Specification (XKMS 2.0), http://www.w3.org/TR/xkms2/

Selective LDAP Multi-Master Replication [*]

Thomas Bauereiss[1], Stefan Gohmann[2], Dieter Hutter[1], and Alexander Kläser[2]

[1] German Research Center for Artificial Intelligence, Bibliothekstr. 1, D-28359 Bremen, Germany, {thomas.bauereiss|hutter}@dfki.de
[2] Univention GmbH, Mary-Somerville-Straße 1, D-28359 Bremen, Germany, {klaeser|gohmann}@univention.de

Abstract: LDAP directory services are widely used to store and manage information about the assets of organisations and to ease the administration of IT infrastructure. With the popularity of cloud computing many companies start to distribute their computational needs in mixed-cloud infrastructures. However, distributing an LDAP directory including sensitive information to partially trusted cloud servers would constitute a major security risk. In this paper, we describe an LDAP replication mechanism that allows for a fine-grained selection of parts of an LDAP directory tree that are replicated to another server using content-based filters, while maintaining the availability and performance advantages of a full multi-master replication. We discuss sufficient conditions on replication topology and admissible operations such that the replication mechanism provides eventual consistency of selectively replicated data.

1 Introduction

Directory services with LDAP interface are widely used to store and manage information about infrastructure and assets of organisations. Multi-master replication (MMR) mechanisms are readily available for directory services, providing high availability and eventual consistency of directory data on different servers via optimistic replication. However, existing MMR mechanisms provide only limited options for configuring a master to replicate only selected parts of the LDAP directory tree. Besides full replication, typically only the division of the LDAP directory into disjoint subtrees is supported. However, with the popularity of cloud computing many companies start to distribute their computational needs in mixed-cloud infrastructures. Groupware, typically supporting the collaboration between employees, are deployed in the cloud to realise a highest level of availability while data and programs constituting the assets of a company should be kept on in-house installations. To ease administration there is a need for uniform mechanisms to administer the in-house installations as well as the various installations in the cloud. A naive approach would distribute a common LDAP directory to all the individual installations causing a major security hole as this common LDAP directory contains also the information needed

---

94

to access in-house installations and thus to access the assets of the company. Consequently, the security of these assets would depend on the security of the LDAP directory stored in the cloud.

Therefore, we aim to develop more flexible mechanisms for selective multi-master replication, giving organisations the ability to select which parts of the directory to replicate to which (cloud) server, while maintaining the advantages of full MMR. In this paper, we describe such a mechanism where each LDAP master has an associated view on the LDAP directory defined in terms of a set of LDAP filter expressions. We discuss sufficient conditions on replication topology and admissible operations such that the replication system provides eventual consistency of replicated data.

The rest of this paper is structured as follows. In Section 2, we describe an example application scenario for our replication mechanism. In Section 3 we give a formal model of LDAP directories, operations and filters. We describe the replication mechanism in Section 4 and define its consistency guarantees in Section 5. Section 6 describes related work in the area of optimistic replication. Section 7 concludes the paper with a summary and directions for future work.

2 Example scenario

As a typical application scenario for infrastructure and identity management based on LDAP directories, consider a large organisation with several local branches in different cities. There is a central LDAP master server at the organisation's headquarters that hosts the full LDAP directory tree, so that the top level management has an overview over the organisation. In addition, every local branch office has an LDAP master of its own that only receives and maintains information about its own employees and its used infrastructure.

Such a scenario is typically implemented by modelling the local branches as organisational units (OUs) and configuring the branch masters to replicate their respective OU information. This is possible with existing implementations of LDAP directory servers, e.g. Active Directory, provided that the directory is partitioned into disjoint subtrees.

Now consider the situation where two branches start a joint project and require a groupware server where the employees working on the project can organise appointments and share documents. In order to avoid having to operate additional infrastructure, they decide to set up the groupware server at a cloud provider. There are several possible approaches:

- Set up a groupware server in the cloud without connection to the enterprise's identity management. The disadvantage is that manual creation of accounts is required and there is no password synchronisation.
- Set up an identity provider for single sign-on at a central enterprise server and configure the groupware server to make use of it. However, only server software that supports the chosen framework for single sign-on can be used in this case.
- Set up an LDAP directory server on the groupware machine and configure it to replicate identity data for employees working on the project. Replicating the whole LDAP directory of the enterprise or even of both branches is not acceptable for

performance and security reasons. Replicating only the data of employees in the project using a content-based filter is only supported in read-only (slave) mode in existing LDAP implementations. However, this implies a limited availability if the connection between the groupware server in the cloud and the LDAP server in the enterprise network is temporarily broken. If a user then wants to change his address information, for example, or wants to define a text for an absence notification on the groupware server, then this would fail. If write access is required during a login process, for example due to a mandatory password change, then even the login fails.

In this paper, we describe a mechanism for selective multi-master replication that allows one to specify which parts of an LDAP directory to replicate based on its content. For example, the organisation described above could specify the employees that work on the joint project and should get access to the groupware server by assigning a corresponding value to the "project" attribute of their LDAP entries. The replication component on the groupware server can then be configured to replicate those and only those LDAP entries. Any application with LDAP support can then be configured to read and possibly write this data. Modifications, for example the change of a password by a user, are replicated back to those LDAP masters that can see that part of the LDAP directory tree, e.g. the central LDAP master and the local master of the user's branch. We envision that such a replication mechanism can give organisations more flexibility and control over their replication setup according to their organisational structure and security requirements.

3   Formalisation of LDAP structures

3.1   LDAP, Schemata and Filters

In the spirit of [WL02] we rephrase the notions of LDAP schemata, directories (instances), and filters in a formal way, in order to be able to reason about the consistency of replication later on. We start with an LDAP schema specifying the ontology of an LDAP tree.

Definition 1. An *LDAP schema* $\mathcal{L}$ is a tuple $\langle \mathcal{C}, \mathcal{A}, \mathcal{T}, \mathsf{req}, \mathsf{opt}, \mathsf{type} \rangle$ where $\mathcal{C}$ is a set of classes; $\mathcal{A}$ is a set of attributes for the classes with $\{oc, dn\} \subseteq \mathcal{A}$, where oc and dn denote the "object class" and "distinguished name" of entries, respectively; $\mathcal{T}$ is a set of types for the attributes; $\mathsf{type} : \mathcal{A} \to \mathcal{T}$ maps each attribute to its type; $\mathsf{req} : \mathcal{C} \to 2^{\mathcal{A}}$ maps each class to its required attributes such that $\forall C \in \mathcal{C}.\ \{oc, dn\} \subseteq \mathsf{req}(C)$; and $\mathsf{opt} : \mathcal{C} \to 2^{\mathcal{A}}$ maps each class to its optional attributes such that $\forall C \in \mathcal{C}.\ \mathsf{req}(C) \cap \mathsf{opt}(C) = \emptyset$.

Definition 2. An *LDAP* $\mathsf{L} = \langle N_\mathsf{L}, E_\mathsf{L} \rangle$ is a forest where each node $N \in N_\mathsf{L}$ is labelled by its class C_N and a set I_N of pairs (a, v) where $a \in \mathcal{A}$ and v is a value of type $\mathsf{type}(a)$. Each edge in E_L is labelled by a DN pair (a, v) such that each node $N \in N_\mathsf{L}$ is uniquely determined by the sequence $(a_0, v_0) \dots (a_n, v_n)$ of labels on the path from its root to itself.

Definition 3. An LDAP L *complies* to an LDAP schema $\mathcal{L} = \langle \mathcal{C}, \mathcal{A}, \mathcal{T}, \mathsf{req}, \mathsf{opt}, \mathsf{type} \rangle$, or is an $\mathcal{L}$-LDAP for short, iff for all $N \in N_\mathsf{L}$ it holds that

(i) $\forall a \in \mathsf{req}(C_N).\ \exists v \in \mathsf{type}(a).\ (a, v) \in I_N$, and

(ii) $\forall (a,v) \in I_N.\ a \in \mathsf{req} \cup \mathsf{opt} \wedge v \in \mathsf{type}(a)$.

We use filters to define views on a particular LDAP. In particular, each filter consists of a Boolean expression controlling which parts of the LDAP are visible in the corresponding view. These Boolean expressions operate on the existence or value of selected attributes entries and combine them with the help of Boolean junctions to complex expressions. The following definition specifies the language for building such Boolean expressions.

Definition 4. Let $\mathcal{L} = \langle \mathcal{C}, \mathcal{A}, \mathcal{T}, \mathsf{req}, \mathsf{opt}, \mathsf{type} \rangle$ be an LDAP schema. The set $\mathsf{Expr}_{\mathcal{L}}$ of $\mathcal{L}$-LDAP expressions is the smallest set satisfying

$$
\begin{aligned}
(a = *) &\in \mathsf{Expr}_{\mathcal{L}} && \text{if } a \in \mathcal{A}, \\
(a\ op\ t) &\in \mathsf{Expr}_{\mathcal{L}} && \text{if } a \in \mathcal{A}, t \in \mathsf{type}(a) \wedge op \in \{=, <, >, \leq, \geq\}, \\
F_1\ R\ F_2 &\in \mathsf{Expr}_{\mathcal{L}} && \text{if } F_1, F_2 \in \mathsf{Expr}_{\mathcal{L}} \wedge R \in \{\wedge, \vee, \rightarrow\}, \text{ and} \\
\neg F &\in \mathsf{Expr}_{\mathcal{L}} && \text{if } F \in \mathsf{Expr}_{\mathcal{L}}.
\end{aligned}
$$

Given the values of some attributes A as a set I of attribute-value pairs, an evaluation function $eval_I \colon \mathsf{Expr}_{\mathcal{L}} \rightarrow \mathsf{bool}$ is defined as usually.

In the following we present the formal definition of an LDAP filter. LDAP filters are the main building blocks to define views on LDAPs and thus to determine those parts of an LDAP to be replicated and maintained in a restricted master.

Definition 5. Let $\mathcal{L} = \langle \mathcal{C}, \mathcal{A}, \mathcal{T}, \mathsf{req}, \mathsf{opt}, \mathsf{type} \rangle$ be an LDAP schema. An $\mathcal{L}$-*filter* is a tuple $\langle p, s, \mathsf{A}, expr \rangle$ such that p is a sequence of DN-pairs, $s \in \{\mathsf{base}, \mathsf{one}, \mathsf{sub}\}$, $expr \in \mathsf{Expr}_{\mathcal{L}}$, all attributes occurring in $expr$ are contained in A, and for all $C \in \mathcal{C}.\ \mathsf{req}(C) \subseteq \mathsf{A}$. Given an $\mathcal{L}$-LDAP L and a $\mathcal{L}$-filter F then a node $N \in N_{\mathsf{L}}$ is *in the focus* of F iff

1. $p = \mathsf{Path}(N)$ and $s = \mathsf{base}$,
2. $\exists (a, v).\ p \circ (a, v) = \mathsf{Path}(N)$ and $s = \mathsf{one}$, or
3. $\exists p'.\ p \circ p' = \mathsf{Path}(N)$ and $s = \mathsf{sub}$.

A node $N \in N_{\mathsf{L}}$ is *accessible* wrt. F iff it is in the focus of F and $eval_{I_N}(expr) = true$.

The application $F(N)$ of a filter to a node N is A if N is accessible and the empty set else.

Definition 6. An $\mathcal{L}$-*view* is a set $\mathcal{V}$ of $\mathcal{L}$-filters. A node $N \in N_{\mathsf{L}}$ is *in the focus* of $\mathcal{V}$ iff it is in the focus of some $F \in \mathcal{V}$. It is *accessible* in $\mathcal{V}$ iff it is accessible wrt. some $F \in \mathcal{V}$. The view $\mathcal{V}(N)$ of a node N is the union of all applications of filters in $\mathcal{V}$ to N.

Using LDAP filters to define the visibility of an LDAP in external masters means that changing the attributes of an object may also change its visibility and thus its accessibility in the cloud. This results in the problem to evaluate a filter in the cloud but having only restricted access to attributes of an object. A simple approach is to require that attributes used in filter expressions have to be a subset of the filter attributes. A more sophisticated approach would be to simplify the filters with respect to the attributes that are not replicated to the cloud. In general however, this results in filter rules that are individual to each object of the replicated LDAP, which is not feasible in practice.

Definition 7. Let L be an $\mathcal{L}$-LDAP and $\mathcal{V}$ be an $\mathcal{L}$-view. Then, $\mathcal{V}$ *induces* an $\mathcal{L}$-LDAP $\mathcal{V}(\mathsf{L})$ on L by

1. an isomorphism $\zeta \colon N'_\mathsf{L} \to N_{\mathcal{V}(\mathsf{L})}$, where $N'_\mathsf{L} = \{N \in N_\mathsf{L} | N \text{ is accessible wrt. } \mathcal{F}\}$,
2. there is an edge (a, v) between $\zeta(N), \zeta(N') \in N_{\mathcal{V}(\mathsf{L})}$ iff there is an edge (a, v) between $N, N' \in N_\mathsf{L}$, and
3. $C_{\zeta(N)} = C_N$ and $I_{\zeta(N)} = \{(a.v)|(a.v) \in I_N| \ a \text{ is accessible in } N \text{ wrt. } \mathcal{V}\}$ hold.

3.2 Operations

In this section we are concerned with manipulating an LDAP or one of its views. The main question is to find appropriate conditions that allow us to relate modifications of the view on an LDAP to corresponding modifications on the LDAP itself. The main requirements to this setting are 1. that the modification on the global LDAP is uniquely determined by the modification on the view and 2. that each modification of the view that results in a consistent state corresponds to a modification on the global view that also results in a consistent state. In order to make this precise, we introduce the notion of admissibility of operations.

Definition 8. A *basic operation* on an LDAP L is one of the following operations:

1. modify (also add or remove)[1] a possibly multi-valued attribute a in a node $N \in N_\mathsf{L}$
2. insert or delete a node in L, or
3. rename a node N in L.

Definition 9. Let L be an $\mathcal{L}$-LDAP and $\mathcal{V}$ be a $\mathcal{L}$-view. A basic operation op is *admissible* on $\mathcal{V}(\mathsf{L})$ iff $\mathcal{V}(op(\mathsf{L})) = op(\mathcal{V}(\mathsf{L}))$ holds. A basic operation op with $op(\mathsf{L}) \neq \mathsf{L}$ is *visible* in $\mathcal{V}(\mathsf{L})$ iff $\mathcal{V}(op(\mathsf{L})) \neq \mathcal{V}(\mathsf{L})$ holds, and *invisible* on $\mathcal{V}(\mathsf{L})$ otherwise.

For simplicity, we assume that complex operations are broken up into a set of basic operations, such that each basic operation is either completely visible (i.e. admissible) or completely hidden in a view defined by a view $\mathcal{V}$. In particular, modifications of multiple attributes of an entry are broken up into operations modifying a single attribute each. The filtering of an operation then reduces to a binary decision whether the operation is visible in a view or not, and we avoid having to alter the content of operations when filtering. This will be useful when we define operation-based selective replication in Section 4.

The modification of an attribute a of a node N is admissible on $\mathcal{V}(\mathsf{L})$ if $a \in \mathcal{V}(N)$. The deletion of N is admissible on $\mathcal{V}(\mathsf{L})$ if N is accessible wrt. some $F \in \mathcal{V}$. The insertion of a node in the $\mathcal{V}(\mathsf{L})$ corresponds to the insertion of N in L with the exception that we allow the further insertion of default attributes, so called I_0 for N not accessible to $\mathcal{V}(\mathsf{L})$. In all other attributes N and $\zeta(N)$ coincide. The insertion is admissible (wrt. a preset I_0) if N is in the focus of some $F \in \mathcal{V}$ and $\forall a \in \mathsf{A}_N \setminus DOM(I_0). \ a \in \mathcal{V}(N)$ and there is no

---

[1]We assume that an attribute does not exist in a node if it has no associated value. Hence, insertion and removal can be regarded as special cases of adding or removing values of an attribute. Also, the replacement of attributes as defined by the LDAP standard can be modelled by removing all attribute values known at the time of submission of the operation, and adding the new values.

other node N' in $\mathcal{V}(\mathsf{L})$ with the same path as N. The renaming of a node N to a path p' is admissible if the insertion of N with its attributes is admissible at path p'.

3.3 LDAP conflicts

Many LDAP operations are commutative, e.g. the modification of different attributes of a node or the insertion of nodes at different paths. In some cases, however, the concurrent submission of operations in a multi-master LDAP system can lead to conflicts. If we assume that operations refer to nodes using a unique identifier, then two concurrent operations are in conflict in the following cases:

- both are modifications of an attribute v of the same node N, and there is a value v that is added by one operation and deleted by the other;
- both are insertions or renamings of nodes at the same path; or
- one is a deletion of a node N and the other refers to N, but is not a deletion.

For these conflicts, we aim to perform immediate automatic resolution in some deterministic way so that the repositories are always in a state that is consistent with schema and application constraints, while at the same time recording conflict so that the conflicts can be properly resolved manually.

In order to detect conflicts, we first have to determine concurrency of updates. For this purpose, LDAP masters propagate basic operations enriched with additional metadata. We denote such an enriched update as $('update', op, m, t, H)$ where op is an operation submitted at master m at (local) time t, and H is the set of all updates known to m at the time op was submitted.[2] An update $a = ('update', op, m, t, H)$ *happened before* an update $b = ('update', op', m', t', H')$, denoted $a \rightarrow b$, iff $a \in H'$. Two updates are concurrent, denoted $a \leftrightarrow b$, iff $a \nrightarrow b \wedge b \nrightarrow a$.

Conflicts are detected by checking if concurrent updates make conflicting changes to an entry. The typical conflict resolution strategy for the modification of attributes is the "Last Writer Wins" strategy, where operations are ordered using timestamps and the newest operation simply overwrites older, conflicting operations. For naming conflicts, we rename the nodes that are moved or created by operations that are dominated by a conflicting operation according to a deterministic naming scheme. For conflicts where a deleted node is concurrently modified, we can either copy the node to a lost-and-found area of the directory tree, or simply ignore the modification.

Overall, a consequence of deterministic conflict resolution is that all concurrent updates commute. In [SPBZ11] it has been shown that strong convergence for full replication easily follows from the commutativity of operations. With selective replication, however, a master might know only a subset of updates, so if only one of two conflicting updates is visible to a master, it cannot perform conflict resolution on its own. In the case of LDAP replication, this affects naming conflicts, as illustrated by the following example.

Example 1. Consider, for example, the insertion of a node N with path $(ou, sales)$,

---

[2]In the actual implementation, we use a compact representation for this set, such as version vectors [SS05].

$(cn, john)$ and attribute $(project, A)$ into an LDAP $\mathcal{V}(\mathsf{L})$ with

$$\mathcal{V} = \{\langle(ou, sales), \mathsf{sub}, \{ou, cn, project\}, (project = A)\rangle\}$$

If L already contains a node with the same path and attribute $(project, B)$, then the insertion of N causes a conflict that cannot be seen by a master with restricted view $\mathcal{V}(\mathsf{L})$.

This means that a master with restricted view cannot, in general, exclude the possibility of inter-view conflicts for node insertions or renamings. One approach to solve this problem is to provide restricted masters with additional information about hidden nodes that are in the focus of one of its filters, e.g. by replicating a dummy node for each hidden node to the restricted master.

A more general solution is to inform the restricted master about the result of conflict resolution by replicating an explicit conflict resolution operation. When a master receives an update from another master with restricted view $\mathcal{V}$, and it detects that the update is in conflict with another update that is not admissible for $\mathcal{V}$, then it generates a conflict resolution update that contains an operation op' with the effect of the resolution that can be propagated back to the restricted master.

Consider again the conflict in Example 1. Assuming we resolve naming conflicts by renaming all but one affected nodes to unique and deterministic names, the conflict resolution operation is in this case the renaming of the relative DN part of N to $(cn, john + m + t)$, where m is the identifier of the restricted master and t is the timestamp of the insertion operation of N, and $+$ is a special concatenation symbol that can only be introduced by LDAP masters, not by users submitting operations. Hence, the new name of N is unique and deterministic. The renaming update is then propagated to all masters where N is visible. Masters with insufficient view to resolve the conflict themselves will apply the update and reach a state consistent with full masters, while for masters that have already resolved the conflict themselves, the update will have no further effect, because renaming a node to a name it already has is redundant.

4 Replication mechanism

We now describe in more detail an operation-based replication mechanism that incorporates our considerations from above. First, we introduce some notation for the state of a selective multi-master LDAP system. We denote the *state* of an $\mathcal{L}$-LDAP master m as a tuple $\langle \mathsf{L}, \mathcal{H}, \mathcal{Q} \rangle$ comprising an $\mathcal{L}$-LDAP L, a sequence $\mathcal{H}$ of updates that have been applied already, called the history of the master, and a sequence $\mathcal{Q}$ of updates that have been received by other masters, but not yet applied, called the queue. We assume that the updates in the queue are additionally annotated with the master from which they were received. The state of a master evolves as it processes updates submitted by users or received by other masters. We denote the state of m at step k (i.e. after the k-th update) as $m(k) = \langle \mathsf{L}(k), \mathcal{H}(k), \mathcal{Q}(k) \rangle$, where $m(0) = \langle \emptyset, \emptyset, \emptyset \rangle$, i.e. masters are initially empty.

A selective multi-master system then consists of a set of masters that communicate with each other at possibly irregular intervals. In existing full-replication LDAP systems, even-

tual delivery of operations is ensured by creating a replication topology in the form of a connected graph, i.e. there is a communication path between any two master servers. In the case of selective replication, we have to additionally take into account the views of the master servers. We have to avoid the loss of information that would occur when all paths between two masters m and m' go through masters m'' with a view that is smaller than both the views of m and m'. In order to guarantee that there is always at least one path without information loss, we require that the topology contains a spanning tree such that views always monotonically increase along a path towards the root:

Definition 10. A *selective multi-master system* $\mathcal{M} = \langle M, (\mathcal{V}_m)_{m \in M}, G \rangle$ consists of

- a set M of LDAP masters, with at least one full master $m_{root} \in M$,
- a family $(\mathcal{V}_m)_{m \in M}$ of $\mathcal{L}$-views $\mathcal{V}_m$ for every master m, with $\bigcup_{m \in M} \mathcal{V}_m \subseteq \mathcal{V}_{root}$,
- a replication topology represented as a connected, directed graph $G = (M, E)$ such that $\forall (m, m') \in E. \, \mathcal{V}_m \subseteq \mathcal{V}_{m'}$ holds and for all master $m \in M$ there is a path from m to m_{root}.

We assume that every master m will always eventually propagate relevant updates to every adjacent master m', i.e. we assume liveness of communication. The propagated updates will then eventually appear in the queue of m' in the correct order, i.e. we assume causal delivery. An update is relevant for $\mathcal{V}_{m'}$ either if it is admissible for $\mathcal{V}_{m'}$ at the time of submission, or if it becomes admissible for $\mathcal{V}_{m'}$ afterwards, because the affected node and attributes have been moved into the view by another operation in the meantime. For example, if a node N is in the focus of a filter $F \in \mathcal{V}_{m'}$ with attribute set A, and an operation changes attributes of N such that the filter expression of F becomes true, then all updates affecting attributes of N in A retroactively become relevant for $\mathcal{V}_{m'}$. Formally, we define the subsequence of updates in a history $\mathcal{H}_m(k)$ that are relevant for a view $\mathcal{V}$ as

$$\mathcal{V}(\mathcal{H}_m(k)) = [u \in \mathcal{H}_m(k) \, | admissible(u, \mathcal{V}(\mathsf{L}_{m_u}(k_u))) \vee$$
$$\exists k' \leq k.u \in \mathcal{H}_m(k') \wedge admissible(u, \mathcal{V}(\mathsf{L}_m(k')))]$$

where m_u is the master where u was submitted and k_u the step of m_u at which it was submitted. Such a history filtering is monotonic in the sense that $u \in \mathcal{V}(\mathcal{H}_m(k'))$ implies $u \in \mathcal{V}(\mathcal{H}_m(k))$ for all $k \geq k'$, i.e. the history filtering only grows with increasing k. We consider two history filterings equivalent, denoted $\mathcal{V}(\mathcal{H}) \equiv \mathcal{V}(\mathcal{H}')$, if both contain the same set of updates, but concurrent updates possibly occur in a different order.

There are two types of local state transitions for a master. Either the state transition is caused by an operation that has been submitted by a user, or it takes an update coming from another master out of its queue and applies it to its LDAP. In the second case, it might also be necessary to generate conflict resolution updates for masters with insufficient view. The effects of the two kinds of state transitions are as follows:

1. If a user submits an operation op at m at step k and op is admissible for $\mathcal{V}_m(\mathsf{L}_m(k))$, then $\mathsf{L}_m(k + 1) = op(\mathsf{L}_m(k))$ and $('update', op, m, t, \mathcal{H}_m(k))$ is appended to the history, where t is a current local timestamp.

2. Otherwise, the master dequeues the first update $u = ('update', op, m', t', \mathcal{H})$ from its queue. If u is already known to m, i.e. $u \in \mathcal{H}_m(k)$, or if u is a conflict resolution

update for a conflict that has already been resolved locally, or if u is not admissible for the view of the master from which it has been received, then the update is ignored and the state remains unchanged for $k+1$. Otherwise, the update is appended to the history and applied to the LDAP, i.e. $\mathsf{L}_m(k+1) = op(\mathsf{L}_m(k))$. If op causes a conflict, the master then determines whether it is necessary to generate a conflict resolution update: If there is a master m' with $\mathcal{V}_{m'} \subseteq \mathcal{V}_m$ and the conflict resolution is visible but op is not admissible on $\mathcal{V}_{m'}(\mathsf{L}_m(k))$, then m generates a conflict resolution update r for u and appends it to its history.

We have now defined both the communication behaviour as well as the local state transitions of masters in a selective multi-master system. In the next section, we discuss the consistency guarantees that such a system provides.

5 Eventual consistency with respect to views

In [SPBZ11], strong eventual consistency (SEC) for full replication is defined in terms of eventual delivery, strong convergence, and termination of operations. In this section, we adapt the definitions of these notions for the case of selective replication. Eventual delivery then means that an update that is submitted at a master m eventually reaches a master m' if and only if it is relevant for $\mathcal{V}_{m'}$. The constraints on the replication topology of a selective multi-master system, combined with liveness, are sufficient to ensure eventual delivery.

Theorem 1. *A selective multi-master system $\mathcal{M}$ provides eventual delivery with respect to views, i.e. if an operation op is submitted at a master m at step k, then the corresponding update $u = ('update', op, m, t, \mathcal{H}_m(k))$ eventually reaches a master m' if and only if it is relevant for $\mathcal{V}_{m'}$:*

$$\exists K_{m'}, K_m > k. \; \forall k_m > K_m, k_{m'} > K_{m'}. \; (u \in \mathcal{V}_{m'}(\mathcal{H}_m(k_m)) \Leftrightarrow u \in \mathcal{H}_{m'}(k_{m'}))$$

Proof. This easily follows from the replication topology, liveness of communication, correct history filtering during communication, and monotonicity of history filtering. □

For strong convergence with respect to views, we require that equivalent knowledge of two masters with respect to a common subview implies equivalent states when filtered for that view:

Definition 11. A selective multi master system provides *strong convergence with respect to views* if for all masters m and m' and for every view $\mathcal{V}$ that is a subview of $\mathcal{V}_m$ and $\mathcal{V}_{m'}$:

$$\forall k, k' : (\mathcal{V}(\mathcal{H}_m(k)) \equiv \mathcal{V}(\mathcal{H}_{m'}(k'))) \Longrightarrow \mathcal{V}(\mathsf{L}_m(k)) \equiv \mathcal{V}(\mathsf{L}_{m'}(k))$$

In order to show that our replication mechanism provides strong convergence, we first show a lemma establishing that for each individual master, applying the updates it knows that are relevant for a view to an empty LDAP results in a state equivalent to the master's actual state filtered for that view.

Lemma 1. *Let $m(k)$ be the state of a master in a selective multi-master system, $\mathcal{V} \subseteq \mathcal{V}_m$ a view, and $\mathsf{L}_{\mathcal{V}(\mathcal{H}_m(k))}$ the LDAP that results from successively applying the operations in $\mathcal{V}(\mathcal{H}_m(k))$ to an empty LDAP. Then $\mathcal{V}(\mathsf{L}_m(k)) = \mathcal{V}(\mathsf{L}_{\mathcal{V}(\mathcal{H}_m(k))})$ holds.*

Proof. By induction on k. The base case $k = 0$ trivially holds, as $m(0)$ is initialised with empty LDAP and history. In the induction step, we perform a case distinction on the type of local state transition from k to $k + 1$.

1. Assume a user submits an operation op at step k. If op is invisible on $\mathcal{V}(\mathsf{L}_m(k))$, then the filtered history and state remain unchanged, and the conclusion follows from the induction hypothesis. If op is admissible on $\mathcal{V}(\mathsf{L}_m(k))$, then $\mathcal{V}(\mathcal{H}_m(k + 1))$ results from appending an update containing op at the end of the history, and the conclusion follows from admissibility of op and the induction hypothesis. If op changes attribute values such that a set A of formerly invisible attributes of the affected node N are now visible according to the filters in $\mathcal{V}$, then op is admissible on $\mathcal{V}(\mathsf{L}_m(k + 1))$, and $\mathcal{V}(\mathcal{H}_m(k + 1))$ results from $\mathcal{V}(\mathcal{H}_m(k))$ by appending an update containing op and possibly interleaving a sequence U of updates that have been made admissible by op. The updates in U are exactly those that affect the attributes A of N and that are not yet contained in $\mathcal{V}(\mathcal{H}_m(k))$. They are independent from and effectively commute with all operations in the filtered history from the previous step that affect other nodes and attributes, and they happen after or concurrently with updates affecting attributes A of N in $\mathcal{V}(\mathcal{H}_m(k))$ due to causal delivery. Hence, applying U and op to $\mathsf{L}_{\mathcal{V}(\mathcal{H}_m(t))}$ leads to a state where the values of the attributes A of N are consistent with $\mathsf{L}_m(k + 1)$, while the consistency of other nodes and attributes visible in $\mathcal{V}$ follows from the induction hypothesis.
2. Assume an update u received from another master is dequeued from $\mathcal{Q}_m(k)$ and applied at step $k+1$. If the update is admissible or invisible on $\mathcal{V}(\mathsf{L}_m(k))$ or changes visibility, then the conclusion follows by the same arguments as in the case of local submission. In addition, however, it is now possible that u is in conflict with an update u' in $\mathcal{H}_m(k)$. If the conflict resolution is visible, but u is not admissible on $\mathcal{V}(\mathsf{L}_m(k))$, then m also generates a conflict resolution update that is admissible on $\mathcal{V}(\mathsf{L}_m(k))$ such that $\mathcal{V}(\mathsf{L}_m(k + 1))$ includes the visible effects of the conflict resolution, and the conclusion again follows as above. $\qquad\square$

Theorem 2. *If concurrent operations commute, then a selective multi-master LDAP system $\mathcal{M}$ provides strong convergence.*

Proof. This is a direct consequence of Lemma 1: For any m, m', k, and k', if $\mathcal{V}(\mathcal{H}_m(k)) \equiv \mathcal{V}(\mathcal{H}_{m'}(k'))$ for $\mathcal{V}$ with $\mathcal{V} \subseteq \mathcal{V}_m$ and $\mathcal{V} \subseteq \mathcal{V}_{m'}$, then $\mathsf{L}_{\mathcal{V}(\mathcal{H}_m(k))} = \mathsf{L}_{\mathcal{V}(\mathcal{H}_{m'}(k'))}$ by commutativity of concurrent operations. Hence $\mathcal{V}(\mathsf{L}_m(k)) = \mathcal{V}(\mathsf{L}_{m'}(k'))$ by Lemma 1. $\qquad\square$

Overall, our replication mechanism in combination with the restrictions on replication topology and admissibility of operations provides both eventual delivery and strong convergence with respect to views. Since we assume termination of operations, we can say that it indeed provides strong convergence in the sense of [SPBZ11], adapted for selective replication.

6 Related work

There is a large body of related work on replication, both in theory and practice, in various settings and with different performance and consistency guarantees. An overview can be found in [CBPS10]. Multi-master replication is an instance of optimistic replication [SS05], where any replica can accept modification operations without waiting for consensus with other replicas. Modifications are propagated from time to time, detecting and resolving any conflicts due to concurrent conflicting modifications. Existing implementations of LDAP directory servers typically support selective replication, but only in slave mode or with limited options for defining which parts of the directory to replicate. To the best of our knowledge, there is no existing support for selective LDAP multi-master replication that allows to define the visible parts of the directory using content-based filters.

In [SBKH05], an abstract formalism for consistency in replicated systems is presented. Partial replication is discussed based on the assumption that the replicated data is partitioned into a set of disjoint databases, with every master replicating a subset of these databases and every database having a primary master. In this paper, we discuss the concrete case of selective LDAP replication and the possible dependencies and conflicts between views. Our definition of eventual consistency includes a notion of eventual delivery with respect to views, and therefore goes beyond the Mergeability property of [SBKH05].

In [RRT$^+$09] a replication platform is presented where devices can select the items they replicate (out of a set of independent items) using content-based filters, similar to our LDAP filter expressions. Also, the Eventual Filter Consistency property of Cimbiosys is similar to our Eventual delivery with respect to views. However, the paper does not discuss dependencies between items or conflicts between updates.

An interesting recent development are "Conflict-free Replicated Datatypes" (CRDTs) [SPBZ11]. These are data types that satisfy certain sufficient conditions for a given definition of eventual consistency. For example, in the case of operation-based replication, the main condition is that all concurrent operations commute, i.e. there are no conflicts. The authors of [SPBZ11] give several examples of non-trivial CRDTs for data structures such as sets, where conflicts are avoided by designing operations for commutativity with the help of additional metadata. In a sense, our work is both a generalisation of the notion of eventual consistency of [SPBZ11] to the case of partial replication, as well as a specialisation to LDAP directories as the data type.

7 Conclusions

In this paper, we presented a mechanism for selective replication of LDAP directory trees together with sufficient conditions to guarantee eventual consistency of replicated data. We are currently working on a prototype implementation of a replication component using the mechanism described in this paper. It is layered on top of a local LDAP server at each master in the replication topology, and is responsible for the communication between masters, enforcing the constraints described above. This includes checking the admissi-

bility of operations submitted by users, propagating correctly filtered operation histories to connected masters, and ensuring that the replication topology satisfies the conditions of Definition 10. The implementation effort also includes work on practical aspects that we were not able to discuss here due to space constraints, e.g. a garbage collection mechanism that allows masters to purge old updates from their histories.

We will evaluate our prototype by integrating it with the Univention Corporate Server (UCS), which is a Debian-based GNU/Linux distribution that allows administrators to manage infrastructure, services and user accounts using tools based on an underlying LDAP directory. We plan to release our prototype as open-source software so that it can be evaluated and applied by others.

Opportunities for further research include the formal verification of correctness and security properties of our replication mechanism with the help of a theorem prover such as Isabelle/HOL [NPW02] in the spirit of works such as [IROM06]. Research in this direction might also lead to a more general theory for selective optimistic replication with eventual consistency for datatypes other than LDAP directory trees.

References

[CBPS10] Bernadette Charron-Bost, Fernando Pedone, and André Schiper, editors. *Replication - Theory and Practice*, volume 5959 of *LNCS*. Springer, 2010.

[IROM06] Abdessamad Imine, Michaël Rusinowitch, Gérald Oster, and Pascal Molli. Formal design and verification of operational transformation algorithms for copies convergence. *Theoretical Computer Science*, 351(2):167–183, February 2006.

[NPW02] Tobias Nipkow, Lawrence C Paulson, and Markus Wenzel. *Isabelle/HOL: a proof assistant for higher-order logic*, volume 2283 of *LNCS*. Springer, 2002.

[RRT⁺09] Venugopalan Ramasubramanian, Thomas L. Rodeheffer, Douglas B. Terry, Meg Walraed-Sullivan, Ted Wobber, Catherine C. Marshall, and Amin Vahdat. Cimbiosys: a platform for content-based partial replication. In *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, NSDI'09, page 261–276, Berkeley, CA, USA, 2009. USENIX Association.

[SBKH05] Marc Shapiro, Karthikeyan Bhargavan, Nishith Krishna, and Teruo Higashino. A Constraint-Based Formalism for Consistency in Replicated Systems. In *Principles of Distributed Systems*, volume 3544 of *LNCS*, page 900. Springer, 2005.

[SPBZ11] Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. Conflict-Free Replicated Data Types. In *Stabilization, Safety, and Security of Distributed Systems*, volume 6976 of *LNCS*, pages 386–400. Springer, 2011.

[SS05] Yasushi Saito and Marc Shapiro. Optimistic replication. *ACM Comput. Surv.*, 37(1):42–81, March 2005.

[WL02] Fang Wei and Georg Lausen. A Formal Analysis of the Lightweight Directory Access Protocol. In *Conceptual Modeling for New Information Systems Technologies*, volume 2465 of *LNCS*, pages 306–319. Springer, 2002.

Approaches and challenges for a single sign-on enabled extranet using Jasig CAS

Florian Holzschuher, René Peinl

Institute of Information Systems – Hof University
Alfons-Goppel-Platz 1
95028 Hof
florian.holzschuher2@iisys.de
rene.peinl@iisys.de

Abstract: In this paper we describe our experiences with setting up a single sign-on enabled intranet with open source software and then making it accessible over the internet using a reverse proxy. During this process, we encounter several issues. We describe those, discuss possible solutions and present our final setup.

1 Introduction

Companies today often have a multitude of software systems running in their internal networks to support their business processes. For employees the ease of use increases, the more integrated these systems are. However, apart from Microsofts server systems that provide out-of-the-box single sign-on (SSO) in Windows domain environments, it is still not common to have even this basic integration, although Gartner called SSO as part of identity and access management (IAM) solutions a "must have for enterprises of all sizes and industries" [Wi+03] already in 2003. Especially in open source settings things seem complex since there are a large number of technological choices and no clear market leader, so that in many cases only authentication against a central LDAP directory is configured instead of SSO.

Our goal was to create an SSO-enabled extranet setup, making as few changes to the software used as possible. In our example, we connected Apache Rave, XWiki, Alfresco and Zarafa to Jasig CAS in order to provide single sign-on, based on accounts taken from a local LDAP directory service. This way, users only have one global account and only need to log in once per session, granting them access to all connected systems through their web browser. All systems do also share user profile information and report user activities to Apache Shindig in order to centrally display them in the Rave portal.

Jasig Central Authentication Service (CAS) was chosen as it is a relatively widely adopted open source authentication system, supporting multiple authentication protocols and methods. The rest of the paper is organized as follows. We first describe different SSO technologies and give an overview of some open source implementations. The we

present our test setup. Afterwards we discuss general issues with the reverse proxy setup and specific issues with SSO, before we conclude with a discussion of results.

2 Single sign-on technologies

SSO can be applied in different scenarios that have different levels of complexity. The probably easiest case is given when all applications are running on the **Intranet** are using the **same runtime** environment like a Java application server and are **prepared for pluggable authentication** [Cl02] like using JAAS in Java, or SAP applications inside the Netweaver Application Server [Bo10]. In this case the container is managing authentication and authorization anyway, so it is quite easy to switch the container from the usual LDAP authentication to a central identity provider like CAS. Ideally, you don't have to make any changes on the application side. Pseudo SSO, using client-side technology to store passwords for server applications is not considered here [PM03]. SSO becomes more complicated if you are considering **multiple runtime environments** like one application running on Java, another on PHP and a third one on ASP.NET for example. You have to either find applications supporting authentication standards of the identity provider (see below) or an identity provider that supports all those (e.g., CAS). An additional level of complexity is added, if you are running your applications in an **extranet** setup [UB09], using a **reverse proxy** to relay and rewrite client requests that address a single host, to the multiple machines running the applications (see figure 1). The reverse proxy could be used to pre-authenticate the requests, so that only authenticated users are directed to the single applications [Ha+12]. Since CAS does not support all functionality with Apache Web server[1], we chose CAS filters inside Apache Tomcat running the applications.

A lot of research has been conducted on even more complex **federated scenarios** that enable users across organizations to access to access applications without additional login [CP07], [Ch09]. This requires one identity provider per organization and an established trust relationship, so that security tokens issued by one identity provider are trusted by all the others.

SSO can further be achieved using different authentication protocols. These are ideally transparently managed by the SSO system.

Kerberos is the predominant standard for SSO in Windows environments [Pr11]. In contrast to **NTLM**, which is the default authentication protocol for Windows, it is able to transfer credentials not only from client to server applications, but also down the road to further systems used by the service provider (e.g. the database). This feature is called delegated authentication. Kerberos can be used in Linux environments as well, although it is not trivial to setup the whole stack consisting of DNS (e.g. Bind), a certificate authority (e.g. OpenSSL), a directory service (e.g. OpenLDAP) and the core component key distribution center (KDC; e.g. Heimdal).

The Security Assertion Markup Language (**SAML**) is mainly used for authenticating against Web services. However, version two includes the Web Browser SSO profile

---

[1] https://wiki.jasig.org/display/CASC/Client+Feature+Matrix

(SAML SSO) designed to authenticate a user against Web applications [Hu+05]. It's an XML-based protocol designed to be included in transport means like SOAP over HTTP and already implemented by some large application providers like Google [Ar+11]. Besides service-oriented architectures (SOA) it is mainly discussed for cloud scenarios [Ce+10]. More recently, SAML is frequently accompanied by **XACML**, in order to provide attribute-based access control, which is a more general form of role-based access control [VDB07]. Emerging from public Web sites like social networks, **OpenID** was proposed as a means to use an identity from one identity provider for accessing other services [Io12]. However, in the internet OpenID currently suffers from reluntance of relying parties [Su+10] and lack of trust from end users [Su+11]. For authorization, OpenID is often accompanied by **OAuth 2.0** [SB12]. **OpenID connect** is a recent development in this area trying to better harmonize both parts [Bo12].

Finally, in the open source area, a multitude of SSO providers is available, each with tested compatibility to a number of different open source systems. A selection of well know open source SSO providers is briefly compared in table 1 and discussed below.

Table 1: comparison matrix for open source SSO solutions

	Jasig CAS	**JOSSO**	**WSO2 Id Server**	**Open AM**
Latest version	3.5.2 (22.02.13)	2.3.0 (31.08.12)	4.1.0 (11.02.13)	10.1.0 (20.02.13)
License	Jasigs own open source license	LGPL	APL v2	CDDL 1.0
Protocols	CAS, OAuth, OpenID, SAML, Kerberos	SAML, NTLM	OAuth, OpenID, XACML, SAML, … (18+),	OAuth, SAML, Kerberos
Authenti-cation backends	JAAS, LDAP, AD, Radius, JDBC, X.509, Negotiate (Kerberos)	JAAS, LDAP JDBC, two factor auth with WiKID, X.509	LDAP, AD, JDBC, Cassandra	LDAP, AD, two-factor auth with HOTP, Negotiate (Kerberos)
Runtimes	Tomcat or other Servlet 2.4 container	JBoss, Tomcat, Websphere, Geronimo, Jetty	WSO2 Carbon server	Tomcat, JBoss
Agents	Spring, MS IIS, JEE, Apache 2.2, PHP, PAM	Apache 2.2, PHP 4+, MS IIS, Liferay, Alfresco, phpBB, Coldfusion, Spring	None found	Apache 2.4, MS IIS, Sun Web Server, JBoss, Glassfish, Tomcat, Websphere, Web Logic

Shibboleth is the implementation of the Internet 2 consortium and specifically designed for federated scenarios [Ch09]. It uses SAML messages with digital signatures in order to improve trustworthiness (ibid.). It allows protecting the user's identity by using random pseudonyms for every session. Since both federation and anonymity are not required in our scenario, we did not consider Shibboleth.

Jasig CAS (Central Authentication Service) is a SSO system using its own protocol (also named CAS). However, it also supports SAML and included support for SAML version 2.0 in CAS v 3.5.1 dating in October 2012 by updating to OpenSAML 2. It also includes support for OAuth 2.0 and can act both as an OAuth client and delegate authentication to other OAuth servers like Facebook or Google, as well as an own OAuth server to directly authenticate OAuth clients. The basic architecture of CAS is similar to the Kerberos model [Io12]. The CAS protocol also supports ticket proxying, which is similar to the Kerberos' delegated authentication. Starting with CAS version 3, it does also support single logout [WY10]. We chose CAS due to its direct support of Liferay, Moodle and Mule.

Other open source SSO systems include **JOSSO** [AFG06], a completely Java based identity provider that also supports PHP and dotNET service provider and has a nice graphical tool to configure SSO scenarios, the **WSO2 Identity Server** [SFB10], which is especially interesting when using the family of WSO2 infrastructure products as well as the successor of the Sun OpenSSO framework **Forgerock OpenAM** [Th11]. We plan on testing some of these in future work. Especially OpenAM in conjunction with the Open Identity Gateway seems a promising alternative for our scenario.

3 Basic setup

Our setup consists of a single machine with an external IP running an Apache web server that acts as a reverse proxy, a single machine with Jasig CAS and several machines running our service providers, all of which are accessible through web interfaces. All machines are located inside a DMZ behind two firewalls, one towards the internet and one towards our internal network (see figure 1).

In contrast to common patterns [So03], we do not separate the proxy from the other machines by an additional firewall, but only use it as a gateway for terminating the SSL connection [Ma99]. The reverse proxy is using a signed certificate, only allowing HTTPS connections and redirecting any unencrypted calls. The Apache Tomcat instance running Jasig CAS is also SSL-enabled to allow for secure ticket validation, but is using a self-signed certificate. We are using AJP for connecting to Tomcat-based applications. However, that didn't prove much better than a normal http connection (see section 4.2). We did also consider nginx as a replacement for Apache httpd since it is optimized for reverse proxy scenarios and provides an easy to use caching mechanism. Another alternative worth testing would be to use a specialized SSO gateway like the Forgerock Open Identity Gateway [BCN12]. It promises to enable SSO for those 30% of typical Web applications that do not work with the usual SSO filters or agents.

For all connected software systems running inside Apache Tomcat we used the authentication, validation, request wrapping and single sign-out filters provided by CAS.

These filters together redirect unauthenticated users to the SSO login page, validate incoming tickets and store the authenticated user in their respective sessions. The Apache web server used for the PHP-based Zarafa server is using a CAS authentication module (mod_auth_cas) which also redirects users, evaluates tickets and sets the logged in user for requests.

figure 1: system architecture of the extranet scenario

Yet, unless software is prepared for reading the provided session information, an authentication plugin is required, telling the service provider which user is currently logged in. Moreover, the first time a user logs in, a new local user account may have to be created, preferably using user data from the local LDAP server. A random local password should be set at that point to avoid empty local passwords, especially if local login cannot be disabled completely. Since we are using open source software exclusively, we were able to implement suitable plugins for almost all systems we wanted to include. However, in most cases configuration was enough and no programming was necessary. All systems are also connected to an LDAP server in order to retrieve additional user information like full name and email address from it.

We also conducted some tests with Android-based smartphones and tablets and found out that Chrome on Android behaves in the same way as its Desktop counterpart.

4 General challenges

Running our services through a reverse proxy and with single sign-on filters caused several problems, not all of which could be solved completely.

4.1 Platform problems

In general, the whole setup has more layers than a normal intranet setup, reducing performance noticeably. We addressed this by keeping rewriting to a minimum and only including paths in the SSO filtering that we were certain needed direct protection or an automatic redirect to a login page. When applicable we were able to achieve slightly better performance by using the binary Apache JServ Protocol (AJP) to connect application servers to the proxying web server, instead of a normal HTTP connection. It did also enhance performance to use nginx instead of Apache and enable its caching. However, the login process is not affected by this caching and nginx needs a plug-in for AJP support instead of supporting it natively.

Some applications had problems with being **accessed using HTTPS** in their external URL while the reverse proxy accessed them using HTTP. We were able to fix this by setting the appropriate parameters in all applications and placing redirects in our web server configuration. When using an HTTP reverse proxy, CAS' login page displays a **warning message about an unencrypted connection** that will not support SSO, but works anyway. We could eliminate this warning by using even an unencrypted AJP connection. Supplying all application and web servers with certificates and reconfiguring the reverse proxy to use SSL may also solve these protocol-related problems, but will result in a more complex setup with slightly lower performance.

In terms of usability, we found that unless directly connected via Spring Security, the applications' **logout buttons did not work with SSO**. They may terminate the application's own session, but with the SSO session still active, the user is immediately logged in again. This is especially confusing since CAS provides a dedicated single log-out (SLO) filter and the client feature matrix states that SLO should work out of the box for all clients, except Spring [Fr12]. It is also tough to correct this behavior programmatically. Although most systems provide an interface to create an authentication plugin, overriding the logout action is usually not available. In any case the question whether a user only wants to log out of a single application or terminate the whole SSO session is still open.

Beginning with Java 7, some **SSL warnings are treated as errors** and cannot be easily circumvented inside applications. We found that if the name in the certificate and the URL don't match, an **"unrecognized_name" error** is detected, which can be fixed by including all possible external server names as aliases in the web or application server's configuration. In this context, further problems can be caused by faulty DNS and domain name configuration, causing further **name mismatches on reverse lookups**.

4.2 Rewriting

Depending on which parameters are used for their generation, web pages delivered by proxied web services can contain **incorrect URLs** referencing other resources. This is caused by the Tomcat server detecting its own machine's external address, incorrectly specified external hosts and contexts differing from the local context inside Tomcat.

Since we were trying to run all services under sub paths of our external host to avoid URL collisions, we kept experiencing faulty redirects and incorrect links. One of our approaches was to use the Apache web server's rewriting functionality. It can be used to correct URLs in headers, links, cookies and references to and within other resources such as JavaScript and CSS.

As this approach did not produce consistently satisfactory results, we resorted to replicating the sub path structure on all web and Tomcat application servers. This way, we solved most problems concerning URLs and redirects and only had to manually correct some URLs in the applications' resources and configuration and move static resources to their new location if needed. The drawback of this method is having to create **static contexts in Tomcat**, with a hash tag in its name as a delimiter to denote the sub path. This was still necessary when using AJP instead of http as described below.

While using **AJP could largely solve the problem** of incorrect links and faulty redirects, we found that some applications' resources like CSS files and especially links within JavaScript code were still wrongly referenced. Therefore, additional rewrites or the sub path replication mentioned above were needed in this setup as well. We determined that this problem is based on the fact that applications, when started, use their local context information to generate resource links which are incompatible with the differing context used by the web server. We correctly configured the proxyName and proxyPort attributes in the AJP connector but still had those problems. The most notable benefit of AJP was, that CAS was no longer complaining about the unsecured connection between the proxy and the service providers when the "secure" option in AJP is set to true.

Another interesting configuration option in the AJP connector is called "tomcatAuthentication" and causes the **authentication to already be performed on the reverse proxy** instead of the connected tomcat application servers. This configuration looks similar to the Forgerock Open Identity Gateway solution. However, in our test, we could not perceive any notable differences compared to authentication using Tomcat, especially regarding performance. The configuration might become a bit easier though. Finally, you can configure encrypting the connection between Apache httpd and Tomcat by configuring a pre-shared key using AJP's "requireSecret" option.

4.3 Service accessibility

Direct access to applications using their own **local administrative accounts**, did also prove to be a challenge. After the initial setup, those accounts are often the only way to properly configure applications and delegate permissions to other users. However, the

filters used to protect services by redirecting unauthenticated users to the CAS login page are blocking access to the applications' own login mechanisms. We found no way to enable both SSO as a default method and still provide local access for admin users as a fallback. The only option looking promising in CAS is called **gateway mode**. This mode is attaching tickets to the request for already authenticated users and passing unauthenticated request through to the service provider. However, this mode requires larger changes to the service providers in order to start the SSO session.

One approach we took was **creating users with the same ID in the LDAP directory**, preserving for example the user's administration rights in the application. Alternatively, one could manually add or override these rights in the applications' account database for existing LDAP accounts, since the administration interface is unavailable. Of course, this approach is not necessarily suitable for production environments.

This could be circumvented by **storing per-application rights in the LDAP directory**, which some applications offer as an option. But this would require all necessary schemas to be incorporated into the directory's structure and additionally writing plug-ins for applications that don't already support this approach. Moreover, we reckon a dedicated authorization system might be a more elegant solution when dealing with a greater number of systems. Though, this will require more and possibly more complicated plugins, for which there is even less predefined support from applications.

Another way around this problem is a more **sophisticated authentication chain**, checking several authentication possibilities before redirecting the user and offering an opt-out functionality for the SSO mechanism. This way at least users knowing their full URL could still log in. For this to work, the SSO login page would need an additional opt-out button which redirects the user to the original page with an additional parameter. This parameter would then be detected by a modified SSO filter and disable or modify the redirect to allow a login.

Theoretically, one could also **disable the automatic redirects** altogether, making the user choose between a button to log in locally and one to log in using SSO, which triggers the redirect. This modification would be needed for each individual application and would make the fully automated sign-on procedure semi-automatic. To minizime the usability trade off made by this approach, at least within the comapny the SSO login page could be set as the browsers' starting page, offering the user to log in at the beginning of each session while still leaving the option to opt out.

Furthermore, services will also be used from within a company's network and concerning performance it would be desirable to **access them directly**, bypassing the reverse proxy. But we found that some applications need to have their external URL specified, which in our example would be pointing to the reverse proxy. Thus accessing applications directly can cause inconsistent web pages being generated, with resources being referenced internally as well as externally or possibly with an incorrect URL.

5 Specific problems

With Jasig CAS we encountered the problem that the **certificate** used for its Tomcat server needs to have the **subject alternate name** set correctly. While we could fix this for our self-signed certificates by generating them accordingly, CAS would not work with the existing certificate used by our reverse proxy server, since it lacks this parameter. To generate a suitable certificate using Java's keytool, Java version 7 is needed, so one needs to be careful when doing so as still many applications only work properly when using Java version 6.

When trying to connect our Apache Rave portal to its back-end, Apache Shindig, we realized that it was not easily possible to **authenticate against CAS** and maintain a SSO session, i.e. simulate a user, **from Java code**, to access the service protected by the CAS filters. The SSO setup we chose is designed to be used from a web browser and although the protocol is documented, we could not find or develop a connector that can establish a usable SSO session from java code. Similar problems will occur when communication between individual systems is required. Again, this could be resolved using a more sophisticated authentication chain, allowing other login methods to pass through without triggering redirects. We should also note that direct login from code using a username and a password is discouraged by the CAS developers, so our failure to maintain a session may be the desired behavior.

Otherwise **CAS' ticket proxying functionality** may provide a solution, also using a simulated client with a service account. This solution gives applications the possibility to request a proxy granting ticket for a logged-in user that it can use to request further tickets to be consumed and validated by other applications. This way, authenticated server-to-sever communication is possible without impersonating a user, assuming proxy tickets are accepted.

This is still not fully sufficient in our case as we also need to have server-to-server communication when there currently is no user logged in and background processes are firing events. Anyway, ticket proxying is more likely to work in our case since it only requires a single call to CAS from the user, causing a redirect to the calling service with the ticket needed to start proxying.

To enable this functionality we would need trusted, encrypted connections between the servers concerned, further HTTP service endpoints, capable of being validated by CAS and receiving proxy granting tickets and handling proxy tickets. Especially in case we wanted to use tickets from a real user session, we would also have to modify the applications' security systems, storing CAS session information so that it can be used from any part of the application requiring server-to-server communication.

Lastly, we tried to make some of our services searchable using **Apache Solr** with an unmodified ManifoldCF instance as a crawler. This also failed due to problems with the SSO session. As with our manual approach, we did manage to authenticate against CAS using built-in functionality but then failed to actually use the session for crawling. The

crawler was being redirected back to CAS after each request, even though it specifically supports application server sessions.

6 Conclusion

Our experiences show, that adding only a little more complexity by introducing a reverse proxy leads to several issues with SSO in a real world scenario. Some of the issues described are application or SSO system specific, some others are only basic challenges like specific settings in the SSL certificate and a few are issues by design, like single sign-out.

We found that the difficulty of including an application into the whole setup depends strongly on the software design, which frameworks were used in what way and how well modifying an application's configuration, code and plugin capabilities are documented. For example using Spring Security offers generic interfaces for adding SSO support, which make integration rather easy in case you are familiar with Spring configuration. But we also found that in some cases an individual preparation for SSO plugins can be more suitable and easier to configure – possibly easier to handle in a more sophisticated authentication chain.

Our biggest problem with the reverse proxy setup was handling the context switch between application and web server. Many applications are able to detect requests through a reverse proxy or can be configured accordingly, dealing with protocol and external hostname changes. But we found that hardly any application will work normally when placed under a different path in the web server compared to the application server it is running on. This must be considered bad code design, since it is no problem to query the current context information like hostname and context path from the application server. However, using hard coded paths at least in some areas of the application (especially JavaScript) seems still the default, based on our tests. Maybe this is a negative side-effect from the current trend to port application code from the server to client-side JavaScript.

In general, we can conclude that a basic SSO-enabled extranet with CAS can be created with a reasonable amount of work, given well-prepared applications. Yet creating a well-rounded, high quality working environment will require extensive modifications to many applications' authentication systems and partially to the IT infrastructure around them.

References

[AFG06] Agostino Ardagna, C., Frati, F., & Gianini, G. (2006). Open Source in Web-Based Applications: A Case Study on Single Sign-On. International Journal of Information Technology and Web Engineering (IJITWE), 1(3), 81-94.

[Ar+11] Armando, A., Carbone, R., Compagna, L., Cuellar, J., Pellegrino, G., & Sorniotti, A. (2011). From Multiple Credentials to Browser-based Single Sign-On: Are We More

Secure? In: Future Challenges in Security and Privacy for Academia and Industry (pp. 68-79). Springer Berlin Heidelberg.

[BCN12] Bryan, P., Craig, M., Nelson, J. (2012): Guide to OpenIG 2.1.0. Forgerock. 17.05.2012 http://docs.forgerock.org/en/openig/2.1.0/gateway-guide/index/index.html

[Bo10] De Boer, M., Essenpreis, M., Garcia-Laule, S., Raepple, M. (2010):Single Sign-on mit SAP – Lösungen für die Praxis. Galileo Press, Bonn

[Bo12] Boyd, R.. Getting Started with OAuth 2.0. O'Reilly Media, Incorporated, 2012.

[CP07] Camenisch, J., & Pfitzmann, B. (2007). Federated identity management. In: Security, Privacy, and Trust in Modern Data Management (pp. 213-238). Springer Berlin Heidelberg.

[Ce+10] Celesti, A., Tusa, F., Villari, M., & Puliafito, A. (2010). How to enhance cloud architectures to enable cross-federation. In IEEE 3rd International Conference on Cloud Computing, pp. 337-345

[Ch09] Chadwick, D. W. (2009). Federated identity management. In: Foundations of Security Analysis and Design V (pp. 96-120). Springer Berlin Heidelberg.

[Cl02] De Clercq, J. (2002). Single sign-on architectures. In Infrastructure Security (pp. 40-58). Springer Berlin Heidelberg.

[Fr12] Fritschi, J. (2012): CAS Client Feature Matrix. Jasig. 04.04.2012. https://wiki.jasig.org/display/CASC/Client+Feature+Matrix

[Ha+12] Haron, G.R., Maniam, D., Sadasivam, V., Loon, W.H. (2012): Re-engineering of Web Reverse Proxy with Shibboleth Authentication. International Conference for Internet Technology and Secured Transactions, 10-12 Dec. 2012, London

[Hu+05] Hughes, J., Cantor, S., Hodges, J., Hirsch, F., Mishra, P., Philpott, R., Maler, E. (2005): Profiles for the OASIS Security Assertion Markup Language (SAML)V2.0. OASIS standard. http://docs.oasis-open.org/security/saml/v2.0/saml-profiles-2.0-os.pdf

[Io12] Ionita, M. G. (2012). Secure Single Sign-On using CAS and OpenID. Journal of Mobile, Embedded and Distributed Systems, 4(3), 159-167.

[Ma99] Maier, P. Q. (1999): Implementing and supporting extranets. Information systems security, 7(4), 52-59.

[PM03] Pashalidis, A., & Mitchell, C. J. (2003). A taxonomy of single sign-on systems. In Information Security and Privacy (pp. 249-264). Springer Berlin Heidelberg.

[Pr11] Pröhl, M. (2011): Kerberos – Single Sign-On in gemischten Linux/Windows Umgebungen. d.punkt Verlag, Heidelberg

[So03] Sommerlad, P. (2003): Reverse proxy patterns. In European Conference on Pattern Languages of Programming, EuroPLoP 2003.

[SFB10] Steuer Jr, K., Fernando, R., & Bertino, E. (2010). Privacy preserving identity attribute verification in windows cardspace. In Proceedings of the 6th ACM workshop on Digital identity management (pp. 13-16). ACM.

[SB12] Sun, S. T., Beznosov, K. (2012). The devil is in the (implementation) details: an empirical analysis of OAuth SSO systems. In: Proceedings of the 2012 ACM conference on Computer and communications security (pp. 378-390). ACM.

[Su+11] Sun, S. T., Pospisil, E., Muslukhov, I., Dindar, N., Hawkey, K., & Beznosov, K. (2011): What makes users refuse web single sign-on? an empirical investigation of OpenID. In Proceedings of the Seventh Symposium on Usable Privacy and Security (p. 4). ACM.

[Su+10] Sun, S. T., Boshmaf, Y., Hawkey, K., & Beznosov, K. (2010, September). A billion keys, but few locks: the crisis of web single sign-on. In Proceedings of the 2010 workshop on New security paradigms (pp. 61-72). ACM.

[Th11] Thangasamy, I. (2011) OpenAM. Packt Publishing, Olton

[UB09] Ullrich, M., & Rieger, F. (2009). Brancheninitiative Single Sign-On: Der sichere, einheitliche und einfache Zugang zu den Extranets der Versicherer wird Realität. In: Maklerverwaltungsprogramme der Zukunft: Ein Ausblick auf zukünftige IT-Systeme zur Unterstützung von Versicherungs-und Finanzvertrieben, 179.

[VDB07] Vullings, E., Dalziel, J., & Buchhorn, M. (2007). Secure Federated Authentication and Authorisation to GRID Portal Applications using SAML and XACML. Journal of Research and Practice in Information Technology, 39(2), 101-114.

[WY10] Wang Y., Jia, Z. (2010): The Application Research of Single Sign Out Model Based On CAS. International Conference on Computer Design and Appliations (ICCDA 2010)

[Wi+03] Witty, R. J., Allan, A., Enck, J., & Wagner, R. (2003). Identity and access management defined. Research Study SPA-21-3430, Gartner.

Identity management in cloud computing in conformity with European Union law? – Problems and approaches pursuant to the proposal for a regulation by the European Commission on electronic identification and trust services for electronic transactions in the internal market

Stephan Sädtler[1]

Chair for Public Law, IT-Law and Legal Information Technology,
held by Prof. Dr. Gerrit Hornung, LL.M.,
University of Passau
Innstr. 39
94032 Passau
stephan.saedtler@uni-passau.de

Abstract: On 4 June 2012, the EU Commission submitted a draft of a regulation on "electronic identification and trust services for electronic transactions in the internal market" [EC12][2]. Due to its impact onto the infrastructure of the new German identity card (nPA) it is subject to fierce criticism, particularly from Germany. This essay seeks to address that criticism and to discuss potential approaches, amongst others that of the research project „SkIDentity – Trusted Identities in the Cloud" of the „Trusted Cloud" programme[3], whilst also addressing accompanying questions of law in the context of identity management in cloud computing.

1 Introduction

Data protection and data security in the sector of information technology – and especially in cloud computing – have become a continuous issue due to the rapid technological development and the accompanying variety of applications of IT-systems. With the constant increase of online-based data processing in nearly all areas of life and business and the corresponding potential risks, the demands for security have steadily increased. Driven by that demand, security technology has considerably improved. One accomplishment of that development in Germany was the introduction of the electronic

---

[1] Stephan Sädtler works as a research assistant at the University of Passau and is a certified specialist lawyer for IT-Law. The essay was originally written in German. The translation was poduced by Ray Migge, a student and lecturer for English constitutional law at the University of Passau, to whom the author feels greatly indepbted. The essay is part of the research project "SkIDentity – Trusted Identities for the Cloud", sponsored by the Federal Ministry of Economics and Technology (fundig plan # 01MD11031).
[2] Hereafter also reffered to as eIAS-R-D.
[3] See http://www.trusted-cloud.de/de/1645.php.

identification (eID) via the new identity card (nPA), as its underlying infrastructure is regarded as highly secure and effectively balances the interests of the user including a high level of data protection and those of the recipient regarding the authenticity of the data (regarding the eID-concept of the nPA see [RHS08][Bo10][Mö11]). Most of the member states have also issued electronic identification means that may prove suitable to strengthen trust in online applications. Currently the implementation of applications deemed secure often fails due to a lack of acceptance of such technologies, which, amongst other factors, is often caused by the significant financial and technical costs for service providers. Mere national approaches to online-applications – often in a cross-border context – are subject to several disadvantages. In the light of that, the efforts aiming at harmonizing the framework for electronic identification whilst respecting the principle of technology neutrality (see [EC12, recital 21]) and specifically the Commission's draft regulation on electronic identification and trust services for electronic transactions in the internal market, appear comprehensible and reasonable/sensible.

In Germany, however, the draft has been the object of justified criticism rather than approval (see [Ho12][SpRo13]; for criticism in the area of trust services see [RoJo13]). Constructive amendments to the proposed regulation are imperative, as calls for improvements have legitimately been raised. As the regulatory aims of the draft are deemed predominantly sensible and sound, it cannot stop there: in the interest of a mutual approach, it is rather necessary to find adequate technical solutions pursuant to the objectives of the European Union legislation.

This essay seeks to cover valid points of criticism along with a discussion on potential technological solutions in the context of cloud computing. It will be limited to the respective provisions on electronic identification in Chapter II, which is independent of the provisions on trust services in connection with electronic signatures, and the accompanying general provisions in Chapter I.

2 General Content of the Provisions on Electronic Identification

Central element of the regulation on electronic identification is the requirement for online service providers in Art. 5[4] to adhere to the principle of mutual recognition and acceptance of electronic identification means which will be notified following a notice to the European Commission in accordance with Art. 7 and the provisions on an independent procedure in Art. 6. Adherence to the principle of mutual recognition and acceptance is compulsory only in so far as an electronic identification by electronic identification means and authentication for an individual online service is required by domestic law or domestic administrative practice. Whilst the far-reaching implications of this provision are formulated unambiguously, the requirements as to application and notification remain largely unclear, as will be shown in the following.

---

[4] Art. without reference to a law or regulation are those of the eIAS-R-D.

2.1 Scope of Application

2.1.1 National eIDs

The first alternative of Art. 2 (1) restricts the regulation's applicability to electronic identification which is provided by, on behalf or under the responsibility of a Member State. It is complemented by Art. 6 (1)(a), which requires the implementation of electronic identification means by, on behalf or under the responsibility of an individual Member State.

Whilst the nPA is within the scope of application as it is issued by the Federal Republic of Germany, other identification means, e.g. electronic cards of the telematics infrastructure[5] of the health care services or the identification authentication according to § 6 De-Mail-G are difficult to define as being within the scope of application. At best, they could fall within the second and third alternative of Art. 6 (1)(a), i.e. "[…] issued […] on behalf of or under the responsibility of the notifying member state […]".

The application of Art. 6 (1)(a) to the electronic health data card appears reasonable as it is based on a legal requirement in accordance with § 291a SGB V; however, responsibility lies not with the state but with health insurance funds and companies, which in accordance with § 291a (7) and § 291b SGB V have entrusted the German company "Gesellschaft für Telematikanwendungen der Gesundheitskarte mbH (gematik)" with their responsibility. Indicative of public responsibility is the supervision and authorisation by the German Federal Ministry for Health and Social Insurance (e.g. § 291b (2)). Furthermore, the shareholders of gematik are largely publically financed. Including the private sector in the process of the issuance of electronic identification means does not preclude an assumption of a public responsibility (see [EC12, recital 14]). A serious counter-argument can be found, though, in the wording of the explanatory memorandum to the regulation: "Most EU Member States have introduced some form of electronic identification system" (see [EC12, explanatory memorandum, 3.3.2]). It suggests that the underlying presumption of the draft was the existence of a single national main identification scheme in the context of the regulation per individual Member State.

Even more questionable is whether the identification verification in accordance with § 6 De-Mail-G is within the scope of application of the draft. De-Mail services do not operate on behalf of the Member States. Although these services are subject to accreditation pursuant to § 17 De-Mail-G, a responsibility by the state as outlined in the draft must nonetheless be dismissed, as the accreditation does not entail any liability by the state. The wording of the draft provides a further argument in favour of such result as it requires the issuance of identification means as a pre-condition for notification of electronic identification schemes in Art. 6 (1)(a).

It is fair to conclude, that merely the nPA falls unambiguously within the scope of application of the eIAS-R-D, whilst it remains questionable whether that is the case for

---

[5] http://www.gematik.de/cms/de/egk_2/egk_3.jsp.

other German electronic identification schemes. That would be changed by the amendment proposed in the draft report by the European Parliament Committee on Industry, Research and Energy on 4 April 2013, which suggested at least a change to the wording of Art. 6 (1)(a) to "[…] either issued by the Member state, or issued by another entity as mandated by the Member State or issued independently but recognised by the notifying Member State […]" [EPC13, p. 12] and Art. 2 (1), thereby covering eID-schemes which are merely recognised by Member States. Both, the health data card and the De-Mail verification scheme indubitably fall within that category. However, it would thwart all efforts to notify merely the main identification scheme of any individual Member State.

2.1.2 Restriction to Public Services?

The recitals of the regulation give the impression that only public online services shall be subject to the principle of mutual recognition and acceptance. Recital 11 primarily refers to "[…] cross-border online services offered by the Member States […]". Services provided by the private sector are explicitly excluded as it refers to only those services provided by Member States. Furthermore, recital 9 manifests the aim to overcome obstacles in interactions with public authorities. Pursuant to recital 14, the decision on whether the private sector may be involved in the issuance of electronic identification means shall be left to the individual Member States. Apart from the imprecise wording of that provision it seems to contradict the aim, proclaimed in the very same recital, of diminishing the discrimination between public and private sectors. The wording of Art. 5 itself however does not contain any such restriction as it includes all online services requiring an electronic identification for access (hence, also private applications, for which such an identification is required by law). Merely the referral to the "administrative practice" is directed towards the public sector. Whether that is applicable also to the first alternative remains unclear. Should the regulation seek to target the public sector only, a clarification as to that effect is indispensable.

The very same applies to the proposed requirement of an electronic identification by electronic identification means and authentication. The strict wording would lead to an applicability of the regulation only where the use of online services without an electronic identification and authorisation is excepted. The aforementioned draft report by the European Parliament Committee on Industry, Research and Energy would substitute the term "required" by "available". That alteration appears advisable in the light of the foregoing. However, it would extent the scope of application significantly.

2.2 Conditions for Notification in Art. 6 (1)(d)

Pursuant to sentence 1 and 2 of Art. 6 (1)(d) electronic identification schemes shall be eligible for notification only on the premise of the notifying Member State guaranteeing the availability of an authentication possibility online, at any time and free of charge, enabling the validation of personal identification data whilst refraining from imposing specific technical requirements on relying parties established outside of the notifying Member State's territory. According to recital 15, that provision shall "rule out any specific national technical rules requiring non-national parties to […] obtain specific

hardware or software to verify and validate the notified electronic identification." Such restrictions of specific technical requirements do not apply to the user (holder) of the identification means [EC12, explanatory memorandum, 3.3.2].

As suggested previously, the driving force for that approach is the principle of interoperability of the various existent schemes. Obliging service providers to implement various differing eID-infrastructures would entail immense technical and financial costs, the avoidance of which could be achieved only by refraining from cross-border transactions. Small businesses and minor public authorities would simply not be able manage these costs.

However, the obligation on Member States to provide possibilities to validate person identification data free of charge leads to the question of who will assume responsibility for the costs of that infrastructure. Insofar as it is intended that costs shall be passed on to the users, the concept might fail due to a lack of acceptance by the users. That problem could at least theoretically be solved by public funding. However, it appears unlikely that safe infrastructures could be established without any specific technical requirements for service providers. In a secure infrastructure, a service provider will only be able to read received data with an appropriate software. It appears the principle of interoperability has been given unacceptable precedence over the principle of security.

3 Consequences for the nPA

In the literature, the draft faced fierce criticism for the requirements as the notification in Art. 6 (1)(d) would basically represent the end for electronic identification of the nPA [Ho12, p. 634] or would at least be in stark contrast to the data protection friendly concept of the nPA (critical also [RoJo13, p. 68] [SpRo13, p. 147 et seqq.]). In fact, the nPA infrastructure involves considerable specific technical requirements and financial costs for service providers: According to § 18 (4) PAuswG, the specific bilateral relationship of the nPA-infrastructure requires a service provider according to § 2 (3) PAuswG, to use a valid authorisation certificate to be able to read the data of the German identity card. The certificate is issued on the basis of an authorisation by the contracting authority (Vergabestelle für Berechtigungszertifikate) in the Federal Office of Administration (Bundesverwaltungsamt). It is to be issued if the requirements in § 21 (2) PAuswG are met, which principally serve the principle of data protection and ties the issuance of the certificate to a pre-validation of the service provider and its object of business (see in detail [Mö11][Bo10, p. 3365 et seqq.]).

The service provider requires a specific technical infrastructure comprising hard- and software to be able to read the data. The considerable costs of the infrastructure must be borne by the service provider. As a service provider aiming at the issuance of an authorisation certificate is classified as a relying party pursuant to Art 6 (1)(d), the nPA-infrastructure imposes specific technical requirements on it. The basic concept of the nPA-infrastructure thereby does not meet neither requirement of the regulation: it does not provide possibilities of authentication and validation free of charge nor does it

refrain from imposing specific technical requirements on relying parties. Hence, it is not notifiable under Art.6 (1)(d).

That leads to the much criticised and contradictory result that institutions that accept the nPA will also have to accept the electronic IDs of other Member States although they do not provide the same level of data protection (as those IDs are covered by the principle of mutual recognition in Art. 5), whilst the nPA would not have to be accepted by other Member States (as it is not notifiable under Art. 6). The consequence is that the more secure an ID is, based on specific technical components, the less probable its notifiability is. That would defeat the proclaimed objective of creating trust in online services.

4 Approaches to the Problem

4.1 "Gateway-Approach" by the STORK-Project[6]

The contradiction of Art. 6 (1)(d) to the proclaimed aim of enhancing trust in electronic identification and authentication could be resolved by differentiating between specific technical requirements on the one and general technical requirements in the sense of requirements generally applicable on the other hand. Such generally applicable requirements could be defined by the EU Commission, which would receive the authority in Art. 8 (3) to pass delegated legislation on technical minimum standards. That interpretation of the draft and Art. 6 (1)(d) would allow the existence of a homogeneous scheme based on a high level of protection. The technical harmonisation of various eID-schemes within the EU will be as difficult, though, as would be a safe scheme without specific technical requirements.

From a technological perspective, the factually sole solution would be an intermediate institution independent of any relying party, which would coordinate all schemes and would make obsolete the utilisation of differing system components by service providers. This so-called gateway-approach or proxy-approach was developed by the STORK[7] research project, whose main component was a central Gateway in each member state. Should the draft of the regulation have meant to provide for such system [Be13], it would have required a more precise wording that would have had to regulate the specification of general criteria. The aforementioned proposal would provide such system and would eliminate the provision regarding cost free authentication and validation possibilities.[8]

Apart from that though, the approach faces major objections due to data protection concerns. Technically, it would be possible for the intermediate institution to collect and store all user identification data and information on its specific use. That would allow a

---

[6] See https://www.eid-stork.eu/; in this context also [Be13].
[7] Another approach developed by the STORK-project is the middelware-approach, which would intends the setup of a middleware-software at the service provider. However, as it would also involve specific technical requirements, it neiter would be covered by the current requirements of Art. 6 (1)(d).
[8] According to [ECP13, p. 16] requirements shall be admissable, which have been defined by the Commission in a special procedure.

single institution to create a comprehensive user profile. It must be kept in mind, though, that the use of external identity providers is not uncommon: Even the technical guideline on the eID-service of the nPA-infrastructure explicitly allows outsourced eID-services (see [BSI12, 2.4.2]). The external service provider is responsible for the reading, authenticating and forwarding of nPA data including the result of the validation process to the actual service provider. The external provider thereby manages the authentication certificate of the original service provider. The difference to a central identification provider is that the external provider is not responsible for the entire identification management of a Member State. Furthermore, the external service provider is providing the data processing services pursuant to § 11 BDSG and thereby is subject to a duty to comply with the controlling service provider's instructions.

4.2 SkIDentity-Project

The research project SkIDentity could provide some relief as it might assist in ensuring that the nPA-infrastructure meets the notification requirements whilst overcoming the disadvantages of a central gateway-approach: The project aims at bridging cloud-applications and safe electronic identification means, such as the nPA and the German health data card. It seeks to overcome hurdles for small and medium sized businesses and local authorities, such as the lack of adjustment of cloud-service infrastructures to the specific needs of eIDs and the resulting complications, such as technical compatibility issues, unsolved questions of data protection and questions of law (see [HH+11, p. 297]). Integral part of the SkIDentity-infrastructure is a so-called identity-broker, which would connect the various identification services with the cloud-service providers. Whilst the identity provider would process the actual authentication, the identity-broker merely bundles these services and makes them available to cloud-service providers in a single interface, a so-called cloud-connector. Such single interface makes specific technical requirements in the sense of differing requirements obsolete. The applicability of the concept is not limited to cloud services but can be extended to any internet service.

From a legal perspective, the identity-broker is to not be understood as a natural or legal person. It could be managed by an identity provider as well as by a fourth entity, independent from user, cloud-service provider and identity provider. Thereby, a system would be established that could be used by a cloud-service provider as a relying party without specific technical requirements and free of charge and at the same time remove all data protection concerns about a centralized identity provider by separating the identity provider from the institution that communicates with the user and the cloud-service provider.

5 Questions of Law

Nonetheless, the SkIDentity-project also raises questions of law, which need to be addressed in the context of technical design that conforms to legal requirements.

5.1 Personal Identity Card Law

It remains questionable whether the SkIDentity-infrastructure can be reconciled with the strict requirements of the German personal identity card law. It depends largely on whether an entity independent of the cloud-service provider and involved in the identity management process of reading identification data will be able to obtain an authentication certificate pursuant to § 21 PAuswG for the purpose of identity transfer. It would need to meet the requirements in § 21 (2) PAuswG and § 29 PAuswV. It is questionable whether the entity would make business-related transmissions according to § 21 (2)(1) Nr. 2 part of its objects of business, as § 21 (2)(1) Nr. 2 PAuswG renders it, unlawful.

The determination of business-related transmissions as an exclusion characteristic is based on § 29 BDSG (see [BT08, p. 43]), which by way of example determines advertisements, the practice of credit agencies and address trading to be such transmissions and therefore addresses services whose object of business is the commercialisation of the value of information of data. The exclusion of such transmissions shall prevent that the electronic identification scheme be used as a tool to collect data for address pools or other business entities dealing with data, thereby diminishing the trust of citizens in electronic identification schemes, as the use of such schemes could, for example, lead to an increase of unwanted promotional mailings (see [BT08, p. 43]).

Even though it involves the transmission to cloud-service providers, the aim and actions of a potential identity-broker are different, as the identity-broker would also act in the interest of the users. The object of the transmission would be authentication and not the commercialisation of the data, which makes it fundamentally different from § 29 BDSG and § 21 (2)(1) Nr. 2 PAuswG. Furthermore, § 21 (2) PAuswG must be read in the light of the right to informational self-determination. As far as the freedom and rights of the user are duly taken into account within an infrastructure, it must be rendered admissible as long as it does not compromise the security and safety of the infrastructure.

The same interpretation must be applied to § 29 (1) Nr. 1 PAuswV, according to which the reading of data performed for third parties is prohibited due to data security and protection concerns. In the light of the informational self-determination, such restriction should not be applicable to the provision of data to the owner of IDs. Originally, the section included such a restriction but recently had been amended to exclude ID owners from the restriction.[9]

It follows, that it certainly is possible to design a SkIDentity-infrastructure that would involve an identity-broker managed by an independent party with the aim of independently analysing nPA-data, and which would conform to the requirements of the personal identity card law. As a precondition, the user must retain exclusive control over his personal data without compromising the safety and reliability of the infrastructure. Technical assistance by the provider of broker services does not affect the owners' control. Legal literature proposed an example scenario involving an online data-safe

---

[9] See Art. 2 Nr. 3 PassVuaÄndG of 20.02.2013, BGBl. (2013) I, p. 330 (Nr.10).

provided by the identity-broker, in which nPA data and the authentication result could be stored and within which the ID owner could independently manage the stored data (see [Mö11, § 21, para. 15]). Should data be transferred to cloud-service providers using this method, the exclusion characteristics of the personal identity card law would not be applicable. Neither would it be rendered an evasion of the system of authentication certificates (in a similar context discussed by [Sch10, p. 55]) as the entity responsible for the management of the data would still require such certificate. The cloud-service provider would also still be bound to the general data protection law. The legal relationship between the service provider and the owner of the certificated furthermore could be subjected to civil agreements, barring the service provider e.g. from § 29 BSGD activities. This technical approach is not limited to the nPA-infrastructure but could be applied to a variety of eIDs.

5.2 Further Requirements of the eIAS-R-D

Nonetheless, the imprecise wording of the draft leaves some fundamental questions regarding the requirements of the eIAS-R-D unanswered: Should the term "specific technical requirements" be read strictly or should there not follow a clarification to allow general technical requirements (defined, for example, by the Commission), factually no eID-scheme would be notifiable.

Further clarification is also needed regarding the "relying parties". As identity provider and identity-broker are neither the final recipient nor end-user of the data it would be reasonable to not classify them as relying parties for the purpose of the regulation. Should that be seen different – e.g. because an identity provider or broker under certain circumstance should be liable to the cloud-service provider and therefore must be able to rely on the hard- and software of the user – the notification requirements would not be met by the nPA-regulation as it cannot be designed to exclude specific technical requirements for the identity provider and for the broker provider.

It remains questionable whether that would even sufficiently ensure a possibility for authentication and validation pursuant to Art. 6 (1)(d). Besides, the provision would require the cooperation of the ID owner and the owner's approval of the involvement of another identity.

5.3 Adequate Level of Confidence

Furthermore, it must be ensured technologically that the level of security of and confidence in the infrastructure matches that of a bilateral relationship. That could be accomplished by making the relationship between the owner of the certificate and the cloud-service provider similar to that between the owner of the certificate and the external eID-service (see [BSI12, 2.4.2]). That question is of special relevance where the authentication entails specific legal consequences, e.g. in § 3a (2) VwVfG. The provision was amended due to the E-Government-Initiative and in its new version provides for a substitution of the written form by filling out electronic forms (see [BT12, p. 13]). So far, the written form could only be substituted by using a qualified electronic signature pursuant to the signature law – as is still the case in civil law transactions pursuant to

§ 126a BGB. It would have to be evaluated in how far the legal requirements would still be met in an extended nPA-infrastructure.

5.4 Downside to the Principle of Mutual Acceptance

5.4.1 Lack of Requirements regarding Data Protection and Security

As the aforementioned approach focussed on the notifiability of eID-systems, its implementation left unsolved existent problems regarding the mutual acceptance of foreign notified authentication means: Primarily, a lack of requirements for data protection and security persists. Accepting the nPA in EU-Member States whilst requiring Germany to accept as equivalent to the nPA such identification means with a lower level of security does not render German nor European legal transactions any more secure than they have been so far. However, the approach taken in the regulation might be helpful: Art. 6 (1)(e) makes Member States liable for the unambiguous attribution of the person identification data pursuant to Art. 3 (1) (which is a requirement for notification pursuant to Art. 6 (1)(c)) and for the provision of an authentication possibility pursuant to Art. 6 (1)(d). Potential liability for failure is a distinct incentive for Member States to ensure a high level of security as every Member State seeks to avoid liability (compare also [Bo13]). How effective that approach will be depends on the interpretation of the provisions on liability. The characteristic of unambiguous attribution has been construed narrowly (see [SpRo13, p. 144 et seqq.]). That provision could also, however, be construed in a wider sense, thereby assuming liability for any data leaks. That interpretation would lead to a significantly higher level of data protection and security. However, the afore-criticised exclusion of specific technical requirements for relying parties contradicts an assumption of comprehensive liability. Nonetheless, the liability approach could prove to be an effective measure. However, it would require more precision and reconciliation with the requirements of Art. 6 (1)(d).

5.4.2 Other eIDs

Applying the regulation to the identification means used in the telematics infrastructure of the German health data card would take that infrastructure ad absurdum: Notified identification means of other Member States would have to be accepted within that infrastructure. The evidence suggests that it is practical to exclude identification means used in individual sectors from the scope of application of the regulation. It must also be considered, that the German health data card, although it does represent a means of authentication, is an integral part of the telematics infrastructure that was designed to enhance the use of the electronic health data card. Apart from authentication, the card can serve the function of storing various other health data. Requiring another identification means would diminish the health data card's central role in the telematics infrastructure – apart from the questions as to data security that would be raised by such a requirement.

6 Conclusion

The imprecise and in part contradictory wording of the eIAS-R-D raises various questions with an impact onto safe identity management in cloud computing by electronic identification means. They are of particular relevance for the nPA-infrastructure leading to justified criticism of the draft. The German health data card scheme is also imperilled by the regulation.

As it appears that the decision on passing the regulation has already been made, it is imperative to develop solutions that are reconcilable with the intentions and aims of the regulation. According to the current wording, only the approach of an intermediate entity for the management of identities appears to be a viable solution. As a central gateway approach is rendered questionable because of data protection concerns, the implementation of an eID-broker mediating between various eID-services at least theoretically appears to be the better option. Although that concept appears reconcilable with the content and rationale of the regulation, further amendment of the wording is necessary. The proposal by the European Parliament Commission on Industry, Research and Energy aiming at the elimination of the requirement to provide said services free of charge and at the modification of the technical requirements provides a first and valid starting point. Moreover, further questions must be addressed in the context of designing technology in conformity with the law. Further, there remain imperfections as to the security and protection of data, which could generally be addressed by a concept of Member State liability. It remains to be seen whether the critical voices will be heard during the forthcoming deliberations.

References

[Be13]   Bender, J.: at, Sichere Identifizierung und Vertrauensdienste in Europa. Recht und Technik für sichere elektronische Transaktionen, Stuttgart, 02/03.05.2013.

[Bo10]   Borges, G.: Der neue Personalausweise und der elektronische Identitätsnachweis, NJW 2010, p. 3334-3339.

[Bo13]   Borges, G.: at, Sichere Identifizierung und Vertrauensdienste in Europa. Recht und Technik für sichere elektronische Transaktionen, Stuttgart, 02./03.04.2013.

[BSI12] Bundesamt für Sicherheit in der Informationstechnik (BSI): Technische Richtlinie eID-Server, V.1.6, BSI TR-03130, Bonn 20.04.2012.

[BT08]   Deutscher Bundestag: Drucksache 16/10489, Berlin 07.10.2008.

[BT12]   Deutscher Bundestag: Drucksache 17/11473, Berlin 14.11.2012.

[EC12]   European Commission: Proposal for a regulation of the European Parliament and of the Council on electronic identification and trust services for electronic transactions in the internal market, COM(2012) 238 final, Brussels 04.06.2012.

[EPC13] European Parliament Committee on Industry, Research and Energy: Draft Report on the proposal for a regulation of the European Parliament and of the Council on electronic identification and trust services for electronic transactions in the internal market, Brussels 04.04.2013.

[Ho12]   Hornung, G.: Brüsseler Angriff auf den neuen Personalausweis?, MMR 2012, p. 633-634.

[HH+11] Hühnlein, D.; Hornung, G.; Roßnagel, H.; Schmölz, J.; Wich, T.; Zibuschka, J.: SkIDentity – Vertrauenswürdige Identitäten für die Cloud. In: Schartner, P.; Taeger, J. (Ed.): D-A-CH Security 2011. Bestandsaufnahme, Konzepte, Anwendungen, Perspektiven, Klagenfurt 2011, p. 296-303.

[Mö11]   Möller, J.: §§ 18-21 PAuswG. In: Hornung, G.; Möller, J.: PassG PAuswG, Kommentar, München 2011.

[RHS08] Roßnagel, A.; Hornung, G.; Schnabel, C.: Die Authentisierungsfunktion des elektronischen Personalausweises aus datenschutzrechtlicher Sicht, DuD 2008, p.168-172.

[RoJo13] Roßnagel, A.; Johannes, P. C.: Entwurf einer EU-Verordnung über elektronische Identifizierung und Vertrauensdienste. Neue Regeln für elektronische Sicherheitsdienste, ZD 2013, p. 65-72.

[Sch10] Schulz, S. E.: Grundbegriffe des Identitätsmanagements. Postfach- und Versanddienst, Identitätsbestätigungsdienst und Dokumentensafes. In: Schliesky, U. (Ed.), Technikgestütztes Identitätsmanagement. Rechtsfragen und Lösungsvorschläge dargestellt am Beispiel der De-Mail und elektronischer Dokumentensafes, Kiel 2010.

[SpRo13] Spindler, G.; Rockenbauch, M.: Die elektronische Identifizierung. Kritische Analyse des EU-Verordnungsentwurfs über elektronische Identifizierung und Vertrauensdienste, MMR 2013, p. 139-148.

How to authenticate mobile devices in a web environment - The SIM-ID approach

Florian Feldmann, Jörg Schwenk

Horst Görtz Institute for IT-Security
Ruhr-University Bochum*
Universitätsstr. 150
44801 Bochum
florian.feldmann@rub.de
joerg.schwenk@rub.de

Abstract: With the advent of the iPhone AppStore and Google Play, the 'walled garden' approach of telecommunication companies to supply content to their customers using standard GSM/UMTS/LTE authentication has failed: Neither Google nor Apple, nor any other content provider on the mobile internet, uses the SIM card for authentication. This is mainly due to the fact that mobile telecommunication and internet architectures differ substantially.

In this paper, we propose several bridging technologies to fill this gap. We exemplarily show how to use SIM authentication for web-based Single-Sign-On protocols. Starting from simple password replacement in the authentication between User Agent (UA) and Identity Provider (IdP), we show how we can achieve strong channel bindings between all TLS channels and SIM based authentication.

1 Introduction

In many ways, today's smartphones can be regarded as fully operational computer systems, packing most features of desktop PCs from a few years ago in addition to extended communication functionality. This makes it possible to run applications similar to those of desktop PCs on these devices.

Many of these applications require communication with one or several internet servers providing a certain service (hence called "Service Provider", or *SP*). In most cases, these services require some form of authentication or authorization because certain information connected to these services may either be privacy restricted (e.g. personal data, the user does not wish to make publicly available) or legally restricted (e.g. certain company data which only employees of the corresponding company should have access to).

The most prominent type of authentication method nowadays is the username/password combination. This method, however, has some significant drawbacks: Passwords can be

---

spyed out by an attacker during transmission and weak passwords could even be guessed easily. Though password policies can be used to force users to choose strong passwords and data encryption can be used to protect passwords from being spyed upon, there is still the risk of computer viruses and trojans reading the authentication data directly from the user's computer (e.g. by monitoring user input) or mounting Man-in-the-Middle attacks on security critical connections.

On desktop PCs, antivirus software is relatively common these days. On mobile devices, however, antivirus software is not yet as common. Also several shortcomings of mobile devices (e.g. restricted display size or restricted input possibilities) make it even harder for users to detect malicious behaviour on their devices. Thus, on these devices it is not advisable to use username/password combinations for login procedures. Instead, we propose a novel variant of Single Sign-On (SSO) services, where a user authenticates himself to a trusted entity which in turn authenticates the user towards the desired service. The proposed procedure makes use of the authentication features already available in every mobile device, i.e. the authentication features of GSM/UMTS using the SIM card plugged into mobile devices.

Previous Works   The WebSIM approach [GKP00], proposed in 2000 by Guthery, Kehr and Posegga, shows how to implement web server functionality into a SIM card, thus rendering the SIM card accessible by internet applications. In [KPS+01], Kehr et al. enhance this work and define an internet authentication protocol using features of the SIM Application Toolkit [3GP07a] and WebSIM. However, their approach significantly differs from the Single Sign-On protocols currently used, thus, it cannot easily be applied to current SSO scenarios. Also, it does not provide any information on secure TLS bindings nor any other security against Man-in-the-Middle attacks. Further, their implementation of WebSIM only took into account regular mobile phones, e.g. mobile phones providing only basic telephony functionality like phone calls or SMS, and not the nowadays commonly used smartphones, which are much more powerful and also customizable by the user, e.g. by installing additional software applications.

The Generic Authentication Architecture (GAA) [3GP13] specified by 3GPP for UMTS shares some similarities to our approach, but does not take into account current web architectures and developments. Instead, it forces service providers to support the interfaces defined in GAA. Secure TLS bindings are mentioned in this specification, but it is not explicitly defined how to use them in this context.

In [3GP12b], a framework is defined by 3GPP which allows for the adaption of the aforementioned GAA into actual standardized SSO scenarios. This very closely resembles our approach, but also does not provide any further information on secure TLS bindings. This is exactly the gap our work tries to close.

2 Related Work

For our approach, we mainly combine three existing mechanisms to form a new and secure authentication method for mobile devices: We make use of the mobile device's inherent authentication feature using the SIM card, utilize this feature to enhance a Single Sign-On login procedure and secure this procedure by applying certain cryptographic bindings.

2.1 SIM card authentication

Each "Subscriber Identity Module", or SIM for short, features a unique and secret key K_{SIM} (chosen by the mobile service operator during SIM card creation) which it shares with the corresponding mobile service operator. Also, several algorithms (e.g. for authentication, key derivation or encryption) are available on the SIM card which make use of the secret key K_{SIM} and generate corresponding cryptographic responses when triggered by a certain input. Detailed information on SIM cards can be found in [3GP07b]. For our purpose, only one algorithm is of interest, namely the authentication algorithm specified in the UMTS standard [3GP12a].

The standard does not demand a specific cryptographic algorithm to be used as authentication algorithm, but rather only states its functionality leaving the actual implementation to the design decisions of the mobile service operators. In short, the algorithm takes as input a 128 bit random value $nonce$ and uses the secret key K_{SIM} to compute the distinct corresponding "signed response" $SRES$, which is also 128 bits in length. Note, that "signed" in this case is rather a descriptive name and does not pertain to an actual digital cryptographic signature. The authentication algorithm functions more like a Message Authentication Code (MAC). As both SIM card and mobile service operator know the secret key K_{SIM}, the algorithm can be used for a challenge/response protocol where the challenger checks the other party's response - it computes the expected result $XRES$ (by computing $auth(K_{SIM}, nonce)$) and compares it to the received value $SRES$.

2.2 Single Sign-On

A Single Sign-On scenario is comprised of three parties - a user agent *UA* wishing to authenticate himself to a service provider *SP* and a trusted third party called identity provider *IdP*. *UA* and *SP* do not share any secret information which could be used by *UA* to authenticate himself to *SP*. This especially means that *UA* does not have a username/password combination for *SP*. Both parties *UA* and *SP*, however, have some sort of trust established in *IdP*, i.e. *UA* has some option to authenticate himself to *IdP* and *SP* trusts certain security assertions issued by *IdP*.

Figure 1 shows a typical Single Sign-On scenario. First, the user agent *UA* tries to access some restricted information on *SP* (denoted by the "GET" command). Because *SP* does not initially know and trust *UA*, it issues a so called "Authentication Request" $Auth\_Req$

Figure 1: Single-Sign On

to *UA* containing a redirect to the corresponding identity provider *IdP*. This authentication request holds information about the issuer *SP* and possibly about the authentication methods accepted by *SP* to grant access to the requested ressource. *UA* forwards $Auth\_Req$ to *IdP*, which starts an "Authentication" procedure with *UA*.

This authentication could be done by a simple username/password combination. Due to the sensitive nature of an *IdP* (usually, *UA* can use the same login information to get access to multiple service providers) this, however, should be avoided. Some form of strong authentication should be used between *UA* and *IdP*. Examples for strong authentication include smartcards, one-time passwords or biometric data.

After *UA* has authenticated himself to *IdP*, *IdP* issues a "Token" to *UA* with a redirect message to the service provider who originally issued the authentication request. This Token must be integrity protected (most times this is done by *IdP* creating a digital signature over it) and should include information about the identity of *UA*.

Once *UA* has forwarded the Token to *SP*, *SP* can validate the signature and then use the information given in the Token about the identity of *UA* to grant access to corresponding resources.

Several frameworks allowing SSO systems to be built currently exist and are already widely in use (e.g. OpenID [RR06], OAuth [HL10], Facebook Connect [MU11] or the Security Assertion Markup Language (SAML) [CKPM05]). Our approach is very well suited to be used in conjunction with SAML, but can easily be adopted for any other SSO framework.

2.3   Secure Bindings

Several attacks are known which allow an adversary to gain Man-in-the-middle access to certain secure connections between parties or steal authentication Tokens from the user agent to use them to authenticate himself as *UA* to *SP* (examples can be found in [Kam08], [SSA+09], [Mar09]).

To counter these threats, secure bindings have been proposed, which can be used to cryptographicaly bind certain identity information to specific TLS/SSL connections or specific

communication partners.

Figure 2: SSO login procedure with tls-unique binding

2.3.1 TLS-unique

For our purpose the *tls-unique* binding proposed in RFC 5929 [AWZ10] is of specific interest. The idea of this binding is to take some information uniquely identifying a certain TLS/SSL connection and cryptographically bind it to the authentication information. The *Finished* messages are the last two messages sent from user agent to server, resp. server to user agent, during the TLS/SSL connection establishment and the first messages which are actually encrypted with the key material derived in the connection establishment. They contain a MAC over all messages previously sent in this connection establishment, thus uniquely identifying this specific TLS/SSL session (any other TLS/SSL session would be established with at least differing user agent and/or server nonces, resulting in different Finished messages). By default, the first Finished message of a connection is used to uniquely identify it for the purpose of *tls-unique* bindings.

Figure 2 shows a Single Sign-On login procedure as described in Section 2.2. Before forwarding the authentication request to *IdP*, *UA* extracts the first Finished message from the TLS/SSL session established between *UA* and *SP* (shown in red) and appends it to the authentication request. If the subsequent authentication procedure between *UA* and *IdP* results in a positive outcome, *IdP* will include the Finished message it received from *UA* into the authentication token (along with all other information for this token as described above). When *SP* receives the token and successfully validates its signature, *SP* also extracts the first Finished message from the TLS/SSL session established with *UA* (note that this must still be the same TLS/SSL session, so the session must be kept alive throughout the authentication procedure and no renegotiation is allowed). It then compares the extracted Finished message to the one included in the authentication token. Because an attacker mounting a Man-in-the-Middle attack on the TLS/SSL session between *UA* and *SP* will have different Finished messages in its connections to *UA* and *SP* respectively, *SP* will detect any MitM attacker at this point.

2.3.2   Strong Locked Same Origin Policy

The Same Origin Policy (SOP) implemented in most current web browsers protects user data by allowing e.g. cookies only to be sent to the same server which has stored the cookie in the first place. The "same server" is hereby denoted by the triplet of (protocol, domain name, port). The Strong Locked Same Origin Policy (SLSOP) [KSTW07] enhances this concept to cryptographically bind an SSO Token to the public/private key pair of the intended *SP*. A detailed analysis of SSO using SLSOP is given in [SKA11].

3   The SIM-ID Protocol

In order to enhance the security of SSO login procedures on mobile devices we include SIM card functionality into it. As shown in Section 2.2, Figure 1 shows a typical SSO login procedure between a Client *UA* (in this case a mobile device) and a Service Provider *SP*, utilizing a trusted Identity Provider *IdP* to establish the authentication between *UA* and *SP*. In our scenario, *SP* is most likely a web server requiring user authentication to provide a certain service. *SP* is assumed to have a public/private key pair pk_{SP}/sk_{SP} along with a corresponding certificate to check the validity of its public key. The user agent *UA* in this particular setup is a mobile web browser running on a mobile device. The web browser does not own any cryptographic keys, but has access to a SIM card sharing a symmetric secret key K_{SIM} with the corresponding mobile service provider. This mobile service provider also acts as Identity Provider *IdP* possessing its own public/private key pair pk_{IdP}/sk_{IdP} together with the shared key K_{SIM}. We assume that this *IdP* is trusted by *UA* and the Service Provider(s) *SP* associated with it. We also assume that *IdP* knows the correct public keys pk_{SP_i} of its associated Service Providers.

3.1   SIM-ID Authentication Towards Mobile Network Provider

In our proposal the mobile network operator will serve as Identity Provider in the SSO scenario. Thus, the mobile device requires a means to authenticate itself to the mobile network operator. The GSM/UMTS standards already provide authentication of a mobile device towards a base station, i.e. mobile network operator. In case of UMTS, this authentication is even mutual. This approach was originally intended for securing communications within GSM/UMTS networks, but can easily be adapted for use in internet (e.g. WLAN) connection scenarios.

As described in Section 2.1, UMTS authentication between mobile device and base station is performed by mutually sending a nonce to which the other partner replies with the value resulting from the authentication algorithm using the symmetric secret key K_{SIM}. As K_{SIM} is known only to the SIM card and the mobile network operator, this functions as an implicit authentication between the two parties.

We use exactly this authentication algorithm in conjunction with the tls-unique binding

described in Section 2.3.

Figure 3: SIM-ID protocol - Authentication Towards Mobile Network Provider

Figure 3 shows the resulting protocol. First, *UA* establishes a TLS/SSL connection with *IdP* and extracts the first Finished message FIN (according to the tls-unique binding). *UA* forwards FIN to the authentication algorithm on its SIM card, which computes and returns the signed response value $RES = auth(K_{SIM}, FIN)$. *UA* then sends an authentication request $Auth\_Req$, a nonce N_{UA} and the authentication value RES to *IdP* via the established TLS/SSL connection.

IdP now verifies the authentication value by also extracting the first Finished message FIN from the TLS/SSL channel, performing the same computation of $auth(K_{SIM}, FIN)$ and comparing this value with the received RES. To protect against an attacker impersonating *IdP*, *IdP* is also required to perform an authentication by computing the signed response $SRES = auth(K_{SIM}, N_{UA})$ and sending the result to *UA*. *UA* can now use the SIM card to check whether the expected response to his nonce N_{UA} equals the one received from *IdP* by computing $XRES = auth(K_{SIM}, N_{UA})$ and comparing $SRES == XRES$.

3.2 SIM-ID Authentication Towards Service Provider

With authentication between mobile network provider and mobile device already described in Section 3.1, we now concentrate on how to extend this authentication to a three-party-scenario.

Figure 4 shows the setup and the enhanced SSO protocol. Note that according to Section 3.1 we already assume a mutually authenticated TLS/SSL connection between *UA* and *IdP*, even though this is not explicitly shown in the figure.

We denote by $[m]sk_i$ a signature over message m created with the secret key sk_i of party i which is sent along with the message (thus, $[m]sk_i$ means we send the message m and its correponding signature in the same communication phase). Likewise, $\{m\}pk_i$ denotes

Figure 4: SIM-ID protocol

an encryption of message m with public key pk_i for intended recipient i (of course, in this case only the encrypted message $\{m\}pk_i$ is sent, not the plaintext m).

With the authentication request $Auth\_Req$ corresponding to the selected SSO scheme, *SP* sends a nonce N_S (the request, or at least the nonce, should be signed with the private key sk_{SP} of *SP* to protect integrity and authenticity of the nonce). When forwarding the authentication request to *IdP*, *UA* extracts the nonce N_S and sends it to its integrated SIM module. The SIM module then takes the nonce N_S as input and uses the authentication algorithm to calculate the corresponding response $RES = auth(K_{SIM}, N_S)$.

The response $s = RES$ can now be sent back to *SP* where it can be validated. If for some reason the value RES should be further used by both parties, e.g., by deriving a symmetric key from it, the value should obviously not be sent in the clear. Instead, it can be blinded using a cryptographic one way function g (e.g. a cryptographic hash function), thus sending $s = g(RES)$.

In the next step *UA* has to authenticate himself to *IdP* in order to receive an authentication token. In Section 3.1 we already described an authentication protocol between a mobile network operator and a mobile device. Technically, though, every mutual authentication protocol between *UA* and *IdP* could be used here.

After authentication between *UA* and *IdP* has been performed, *IdP* calculates an expected result $XRES = auth(K_{SIM}, N_S)$ using the symmetric key K_{SIM} corresponding to

the authenticated user. $XRES$ is encrypted with the public key of *SP*, pk_{SP}, and the resulting cipher text $\{XRES\}pk_{SP}$ is included into the authentication token (the format of this token is dependent on the actual SSO scheme selected, but like the nonce N_S in the authentication request before, this token should at least be signed using the private key of *IdP*, sk_{IdP}, to ensure integrity and authenticity). *IdP* sends the token to *UA*, who forwards it to *SP* (the Token and the authentication value s can also be sent together in one single message).

When *SP* receives the token, it is validated according to the chosen SSO scheme. Also, *SP* decrypts the encrypted value $\{XRES\}pk_{SP}$ and checks if $s == XRES$ (or checks if $s == g(XRES)$ instead, if the blinded version is used).

3.2.1 Including Bindings

As we have not yet assumed any TLS/SSL session between *UA* and *SP* - apart from the encryption of $XRES$ with the public key of *SP* (and of course the authentication procedure between *UA* and *IdP*, described in Section 3.1) no encryption is used - the protocol described thus far is prone to a simple Man-in-the-middle attack between *UA* and *SP*: An attacker could simply read all data sent from *UA* to *SP* and then steal the result value s as well as the token and use both to authenticate himself as *UA* against *SP*.

The viability of the proposed protocol becomes visible when SSL/TLS is used for the communication between *UA* and *SP* as well. After all, the connection between a mobile web browser and a corresponding web server providing access to restricted resources should most likely be encrypted to provide a secure transport of these resources.

Utilizing the idea described in Section 2.3 we can bind the authentication info of *UA* to the current TLS/SSL session, again by using the first Finished message FIN from the TLS/SSL session establishment. Only two minor adjustments have to be done to the protocol to include this form of session binding:

- When creating the blinded response s, the user agent also includes the first Finished message FIN by calculating $s = g(FIN \text{ XOR } RES)$

- Similarily, *SP* also takes the Finished message into account by checking if $s == g(FIN \text{ XOR } XRES)$

In Figure 5, the enhanced protocol is shown.

Note that the actual TLS-unique binding has the Finished message sent from user agent to *IdP*, where it is included into the token. The corresponding equivalent in our case would be to input FIN into the authentication function of the SIM card, providing $RES = auth(K_{SIM}, FIN)$ and having the *IdP* also compute this value as $XRES = F_{K_{SIM}}(FIN)$ (compare Section 3.1). However, the *SP* is not able to compute this value himself, and also there is no possibility for him to extract FIN from $XRES$. The *SP* does not gain any knowledge about the Finished message sent from user agent to *IdP* and thus it is not possible for *SP* to check whether or not the token is really bound to a specific TLS channel.

Figure 5: SIM-ID protocol with included TLS-unique binding

By using $s = g(FIN$ XOR $RES)$ on *UA* side and checking if $s == g(FIN$ XOR $XRES)$ on *SP* side, *SP* can be sure that the user agent really is the mobile device which it claims to be (because $XRES$ is confirmed by the trusted *IdP*) and also that it shares the current TLS channel with *SP* (by providing the correct Finished message).

Also note that by encrypting the value $XRES$ specifically for *SP*, we effectively included a server end-point binding of SLSOP type (compare Section 2.3.2).

4   Conclusions and Outlook

We have shown how to adapt the UMTS authentication procedure for use within internet scenarios. Our proposed solution makes use of secure cryptographic bindings to strengthen the mutual authentication between a user agent *UA* and a mobile service provider acting as identity provider *IdP*. By using these secure bindings, we can effectively prevent an attacker from being able to mount a Man-in-the-Middle attack on the communication between *UA* and *IdP*.

We have then extended our protocol to also apply to a three party Single Sign-On scenario, where *UA* wants to authenticate to a service provider *SP* without any prior shared secret between *UA* and *SP*. For this purpose we use a mutually trusted third party *IdP* to build

mutual trust between *UA* and *SP*. Our protocol is also robust against impersonation attacks and Man-in-the-Middle attacks between *UA* and *SP*.

Our protocol can be used with any existing Single Sign-On scheme. Service providers which already support Single Sign-On with the tls-unique channel binding, only need minor adjustments to their authentication procedures. Users can continue to use their current mobile devices and do not need any additional secure hardware.

The limitations of our protocol become evident when considering mobile malware present on the user's device: Though a simple keylogger would no longer be sufficient to spy out the user's username/password combination (this type of espionage in our protocol is, in fact, impossible, as we do not use username/password combinations), a powerful malware which has gained full control of the mobile device could use the SIM card to authenticate against arbitrary service providers. Once the user has unlocked use of his SIM card, there is nothing stopping a malware from obtaining the required authentication values from it, as the SIM card cannot distinguish between 'legal' (i.e. invoked by the user) and 'illegal' (i.e. triggered by the malware) authentication requests. In short, with our protocol, *IdP* and *SP* will be convinced that they both are communicating with a specific mobile device (or rather, a mobile device using a specific SIM card), but they cannot be sure whether the communication was conducted by the regular user or a mobile malware.

References

[3GP07a]   3GPP. 3GPP TS 11.1; Specification of the SIM Application Toolkit (SAT) for the Subscriber Identity Module - Mobile Equipment (SIM-ME) interface. Technical Report 11.11, 3rd Generation Partnership Project (3GPP), http://www.3gpp.org/ftp/Specs/html-info/1114.htm, June 2007.

[3GP07b]   3GPP. 3GPP TS 11.11; Specification of the Subscriber Identity Module - Mobile Equipment (SIM-ME) interface. Technical Report 11.11, 3rd Generation Partnership Project (3GPP), June 2007.

[3GP12a]   3GPP. 3GPP TS 33.102; 3G Security; Security architecture. Technical Report 33.102, 3rd Generation Partnership Project (3GPP), December 2012.

[3GP12b]   3GPP. 3GPP TS 33.980; Interworking of Liberty Alliance Identity Federation Framework (ID-FF), Identity Web Services Framework (ID-WSF) and Generic Authentication Architecture (GAA). Technical Report 33.980, 3rd Generation Partnership Project (3GPP), September 2012.

[3GP13]   3GPP. 3GPP TS 33.222; Generic Authentication Architecture (GAA); Access to network application functions using Hypertext Transfer Protocol over Transport Layer Security (HTTPS). Technical Report 33.222, 3rd Generation Partnership Project (3GPP), March 2013.

[AWZ10]   J. Altman, N. Williams, and L. Zhu. Channel Bindings for TLS. RFC 5929 (Proposed Standard), July 2010.

[CKPM05] Scott Cantor, John Kemp, Rob Philpott, and Eve Maler. Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0. Technical report, March 2005.

[GKP00]   Scott Guthery, Roger Kehr, and Joachim Posegga. How to turn a GSM SIM into a web server. In *Smart Card Research and Advanced Applications*, pages 209–222. Springer, 2000.

[HL10]   Eran Hammer-Lahav. The oauth 1.0 protocol. 2010.

[Kam08]   Dan Kaminsky. Black ops 2008: It's the end of the cache as we know it. *Black Hat USA*, 2008.

[KPS+01]   Roger Kehr, Joachim Posegga, Roland Schmitz, Peter Windirsch, et al. Mobile security for Internet applications. *Proceedings of Kommunikationssicherheit KSI'2001*, pages 27–28, 2001.

[KSTW07]   Chris K. Karlof, Umesh Shankar, Doug Tygar, and David Wagner. Dynamic pharming attacks and the locked same-origin policies for web browsers. Technical Report UCB/EECS-2007-52, EECS Department, University of California, Berkeley, May 2007.

[Mar09]   Moxie Marlinspike. More Tricks For Defeating SSL In Practice. *Black Hat USA*, 2009.

[MU11]   Marino Miculan and Caterina Urban. Formal analysis of Facebook Connect single sign-on authentication protocol. In *SOFSEM*, volume 11, pages 22–28, 2011.

[RR06]   David Recordon and Drummond Reed. OpenID 2.0: a platform for user-centric identity management. In *Proceedings of the second ACM workshop on Digital identity management*, DIM '06, pages 11–16, New York, NY, USA, 2006. ACM.

[SKA11]   Jörg Schwenk, Florian Kohlar, and Marcus Amon. The power of recognition: secure single sign-on using TLS channel bindings. In *Proceedings of the 7th ACM workshop on Digital identity management*, DIM '11, pages 63–72, New York, NY, USA, 2011. ACM.

[SSA+09]   Marc Stevens, Alexander Sotirov, Jacob Appelbaum, Arjen Lenstra, David Molnar, Dag Arne Osvik, and Benne De Weger. Short chosen-prefix collisions for MD5 and the creation of a rogue CA certificate. In *Advances in Cryptology-CRYPTO 2009*, pages 55–69. Springer, 2009.

Authentication on Mobile Devices for Business Application

Martina Müller (MM),* Fabian Zoller (FZ),† Ingo Pansa (IP),‡ Ansgar Gerlicher (AG)§

martina.mueller@ic-consult.at
fabian.zoller@ic-consult.de
ingo.pansa@ic-consult.de
gerlicher@hdm-stuttgart.de

Abstract: Identity management faces new challenges of protecting resources that are aces from different and maybe unknown devices. This is caused by the fact that employees bring their own mobile devices to their offices and work with them [Gar12]. Consequently users, programmer and companies have to face challenges that arise from mobile authentication: the need for accessing business application is based on the employees user identity. This paper describes a methodical analysis and evaluation of the current state of the art authentication methods. The resulting evaluation is the basis for a prototypical implementation of the best evaluated authentication methods on mobile devices. To test and confirm the theoretical architectures implemented on mobile devices a usability test has been made. A conclusion sums up the lessons learned and recommendations are made.

1   Introduction

Authentication in Business Environments is a vital part in securing business data and business processes. Different approaches have been utilized in the past decades; various efforts in implementing externalized authentication and authorization systems have been made. However, these approaches are all based on the assumption that the device that requests access to business application is controlled by the IT department of the companies. With the success of mobile devices, employees start to bring their own devices to the enterprise networks to access protected resources. Thus, new challenges in integrating these devices arise.

Mobile Authentication is an environmental part of working with private devices in a company and accessing resources that are protected. This creates new challenges for the security department. How to verify that a user is actually the one he or she claims to be? Additionally, access via mobile devices emerges the mobility problem to the ubiquitous used authentication methods that have been discussed [AZEH09, BM02, DYB11, Rae11].

---

*MM, iC Consult Austria / Kernstockgasse 8 / 2700 Wiener Neustadt
†FZ, iC Consult GmbH / Keltenring 14 / 82041 Oberhaching
‡IP, iC Consult GmbH / Keltenring 14 / 82041 Oberhaching
§AG, Hochschule der Medien / Nobelstraße 10 / 70565 Stuttgart

This article analyzes and evaluates those methods against the background of mobility focusing the interest on three stakeholders that are involved in IAM: users, company and programmers.

The analysis used the requirements usability and functionality, security, accuracy, expenditure and implementation effort. Furthermore, a rating matrix that adds the results together of the analysis generates a ranking. Authentication methods are evaluated by using distinctive criteria. Based on that the best rated authentication methods have been used to create prototypical implementations. Four functional prototypes on two operating systems (iOS and Android) were developed. Following this, the prototypes that included seven showrooms were tested by a group of 26 testers in a usability test. The conclusive evaluation of the test was necessary to verify or reject the initial analysis of the authentication methods.

1.1 Distinction

This paper focuses on the process of mobile authentication. Several prerequisites have been created. The focus for the authentication is set on strong authentication (involving two factors). The prototypes are fully functional, including front end and back end. The data distribution is considered to be secure and is not part of that paper.

2 Foundation for Mobile Authentication

For the analysis of the requirements for mobile authentication current authentication methods were used and taken into account. Biometrics are used to identify an individual by using certain unique physiological (face, finger, iris, retina, hand) or behavioral (voice, signature, keystroke) characteristics [vT05]. Knowledge based authentication requests the knowledge of the user (secret question, username, specified knowledge). The input can be textual, graphical or made by a gesture. The basis is the challenge-response model. Property based authentication is not intrinsically linked to an individual but describes the possession of e.g. a token or an NFC chip that needs to be verified. Location based authentication methods use the physical aspect of location (latitude, longitude and altitude) that are used by Location Based Services like GPS or WLAN. An additional option to exchange data for authentication in a secure way is the use of digital certificates, signatures and keys, respectively a public key infrastructure (PKI) that can include encryption. Mobile authentication can use several different factors to protect a resource. But which do have a good usability? The specific constraints of mobile devices need to be taken into consideration when comparing mobile authentication methods, in order to find a balance between security and usability.

3 Comparison of Mobile Authentication Methods

The selection of requirements is based on the assumption that there are three stakeholders with different interests that disperse from each other: user (usability, functionality and accuracy), programmer (implementation effort, accuracy, security) and the company (expenditure, accuracy and security).

To make those requirements understandable and reasonable they are subdivided into characteristics as it is presented in table 1.

R_n	Requirement	Characteristic 1	Characteristic 2
R1	Usability and Functionality	Access Time	Acceptability
R2	Security	Distinctiveness	Resistance to Attacks
R3	Accuracy	False Accept Rate	False Rejection Rate
R4	Expenditure	Purchase	Administration
R5	Implementation Effort	Software	Hardware

Table 1: Requirements and Characteristics

Access Time describes the time that elapses from opening the application until the process of authentication has been executed.

Acceptability indicates to what extent people are willing to accept an authentication system. Attention should be paid to intrusiveness, intuitive handling, overview and performance.

Distinctiveness describes the level of uniqueness and the level of differentiation of authentication input.

Resistance to Attacks describes the level of robustness against attacks like fraud, man in the middle or impersonation.

False Acceptance Rate (FAR) and False Rejection Rate (FRR) categorized by Moulton into error types I (FRR) that classify authorized users as imposters and error type II (FAR) that classify imposters as authorized users [Mou83].

Purchase describes the level of expenditures that must be calculated for acquisition only; including hardware like fingerprint scanner, smart cards sensors etc.

Administration describes the amount of work that needs to be calculated for creating a running authentication system with the purchased objects. Reference points are: capturing of data, creation of IDs, managing IDs, maintain the database and solving occurring problems.

Implementation Effort Software describes the amount of work units for learning the required skills like programming language, coding, testing, using frameworks and libraries.

Implementation Effort Hardware describes the amount of work units that are required for implementing hardware (if necessary) like sensors, server and infrastructure to the existing infrastructure.

Every characteristic is subdivided into fragmentations. They are rated upon literature. The ratings are divided into the following essential level: minimal (1), low (2), medium (3) and high (4). Each level refers to the corresponding mathematical value.

Characteristics are using the mathematical equivalent for generating the appropriate requirement. Following Gartner the combination of characteristics is structured in table 2.

	Level of First Characteristic (C1)				
	High	High	Medium	Low	Minimal
Level of Second Characteristic (C2)	High	High	Medium	Low	Low
	Medium	Medium	Medium	Low	Minimal
	Low	Low	Low	Low	Minimal
	Minimal	Low	Minimal	Minimal	Minimal

Table 2: Combination of Characteristics

That means that the output is the geometric mean of the two input values. The output is the square root of the product of the two characteristics.

$$\sqrt{C_1 * C_2}$$

For example: combining level medium (3) with level low (2), the output is 2,44 and therefore low. The mathematical equivalent is:

$$R = Requirement = \sqrt{C_1 * C_2} = \sqrt{3 * 2} = 2,44 \approx 2$$

The mathematical results are converted into the connected values. Requirement one, two and three have a positive scaling, while requirement four and five have a negative scaling. For that reason a meta scale, shown in table 3, was introduced.

Scaling	Level			
Positive Scaling (R1, R2, R3)	high	medium	low	minimal
Meta Scaling	highly advisable (4)	advisable (3)	satisfactory (2)	inadvisable (1)
Negative Scaling (R4, R5)	minimal	low	medium	high

Table 3: Scaling System

3.1 Selection of Methods for further Investigation

Rated by the five requirements (R1-R5), and after the transformation of the values the mathematical results vary from 2,4 to 3,6.

Methods that reach more than three have been taken into consideration for the prototypical implementation. Certificates, tokens, signatures, and key exchange form a set and therefore can be seen as one method. Results are shown in table 4.

Text Based Authentication and Credentials (3,2) are the most widely authentication technique being used [CDW04]. Due to their ease of implementation, cost and accessibility to multiple platforms they reach a high level. Graphical Password authentication has a high acceptance and security level. This authentication method was rated with 3,2.

Method	Authentication	Usability/Functionality	Security	Accuracy	Expenditure	Implementation Effort	Result
Biometrics							
Physiological	Face Recognition	satisfactory	satisfactory	advisable	satisfactory	highly advisable	2,6
	Finger Recognition	advisable	advisable	satisfactory	advisable	highly advisable	3,0
	Iris Recognition	advisable	advisable	highly advisable	inadvisable	highly advisable	3,0
	Retina Recognition	satisfactory	satisfactory	inadvisable	satisfactory	inadvisable	2,0
	Hand Geometry Recognition	advisable	advisable	advisable	satisfactory	highly advisable	3,0
Behavioral	Voice Recognition	highly advisable	highly advisable	advisable	highly advisable	highly advisable	2,4
Knowledge							
	Text Based Authentication	highly advisable	inadvisable	highly advisable	highly advisable	advisable	3,2
	Graphical Password Authentication	advisable	advisable	advisable	highly advisable	advisable	3,2
	Gesture Based Authentication	satisfactory	satisfactory	highly advisable	advisable	advisable	3,0
Property							
	Hardwaretoken	highly advisable	highly advisable	highly advisable	inadvisable	highly advisable	3,0
	Softwaretoken	highly advisable	highly advisable	highly advisable	satisfactory	highly advisable	3,4
	NFC	highly advisable	highly advisable	highly advisable	advisable	highly advisable	3,6
Location							
	GPS	satisfactory	advisable	highly advisable	advisable	advisable	3,0
	WLAN	advisable	highly advisable	highly advisable	highly advisable	advisable	3,6
Other							
	Certificates	highly advisable	advisable	highly advisable	advisable	advisable	3,4
	Signatures	highly advisable	advisable	highly advisable	advisable	advisable	3,4
	Key Authentication	highly advisable	advisable	highly advisable	advisable	highly advisable	3,6

Table 4: Evaluated Rating Matrix

Software Tokens (3,4) are eminently suitable for adding a second channel to the process of authentication can be created without user interaction. The support of security aspects is also given. Location based services like WLAN (3,6) and NFC (3,6) reach a high level due to the fact that they do not require pre-established user-agreement, key, distribution or communication overhead [Bao08]. Additional acceptability and access time had a high score. Digital Certificates (3,4), Signatures (3,4) and Key Exchange (3,6) protect confidentiality, authenticity and integrity by using the public key. The exchange of keys (data) by using certificates and signatures between entities is organized by a PKI. Once installed, the usability and functionality, security and functionality are high.

4 Prototypical Implementation

4.1 Security in Mobile Operating Systems

This chapter discusses the two major mobile operating systems [The11].

Android and iOS. Both of them have strong application layer security models. Android and iOS are generating unique numbers[1] for an application during the installation process. These identifier remain the same until the application is deleted.

Android uses an application UID to enforce the permission for the application. For example accessing the camera. The granted permissions are only set during installation and cannot be changed later.

iOS handles the access to certain resources by using the Mandatory Access Control (MAC). The user can decide at runtime whether an application has access to a certain resource or not [Wat12].

Besides the mentioned application layer security, Android and iOS are also isolating applications from direct hardware access which is called sandboxing [Goo]. Android uses a service layer called Hardware Abstraction Layer (HAL) [Tan09] and in comparison to iOS using a system based on the TrustedBSD project [WFMV03].

4.2 System Components

The components of the system, as shown in figure 1, are divided into three parts, Authentication Back End (ABE), Information Storage (IS) and AuthApp. The ABE comprises a web server, application server and an authentication agent (AuthAgent). It furthermore contains the main business logic for the authentication process. The business data are located in the IS. In this case a basic directory services are used. The client application, AuthApp, is the main part where the user interacts with the system. It handles the communication to the ABE and presents different authentication methods to the user.

Figure 1: Complete Prototype Ecosystem

---
[1]Android: Unique Identifier (UID); iOS: Globally Unique Identifier (GUID)

4.3 Mobile Key

The Mobile Key showroom combines Near Field Communication (NFC) and credentials for the user authentication. AuthApp recognizes the NFC tag and reads the data on it. In case of the NFC authentication, a simple number is used as identification attribute. This number is stored on the NFC tag and in the users object in the IS. Once the data is read from the NFC tag, AuthApp requests the validation of the number. It sends an HTTP request to the ABE and processes the response. On the next step, the user enters his credentials and AuthApp performs an HTTP request to validate the credentials. Is the result also positive, the user is now properly authenticated.

4.4 Location Based

The Location Based showroom uses the BSSID from the connected wireless access point to authenticate the user in the first place. AuthApp gathers the BSSID and constructs a validation request. This request is send to the back end and is verified. The back end extracts the BSSID and checks if the BSSID is in the white list. Based on that result the back end returns âtrueâ or âfalseâ. After a successful BSSID authentication, an OTP generator is shown which requires a personal PIN to generate a code. This code is used to log in to a web service.

4.5 Mobile Desk

The Mobile Desk showroom uses certificates and graphical passwords to authenticate the user. In a previous step a suited device certificate is loaded on the mobile device. AuthApp uses the certificate and consults the back end service to validate the certificate. The next step is to enter a graphical password. AuthApp consults the back end server to validate the entered graphical password. When a positive match is found, AuthApp grants the access.

4.6 Quick Response Code and Credentials via Two Channels

This showroom uses the ability to separate username and password. The user has to use two physically independent systems. In this case, a desktop computer and a mobile device with camera.

On the desktop computer, the user types in his name in the Front End and generates a QR code. Then, he has to scan this QR code with the mobile device. The QR code is used to encode a URL. This URL is loaded in the mobile browser and the user can enter his password. If the correct password is entered, the desktop front end recognizes the successful login and grants access.

4.7 Front End Implementation

The Front End is where the user gets in touch with the system. It is browser based and uses the framework jQuery Mobile for the presentation. No device detection was used to distinguish between a mobile or desktop browser. For that reason, all devices (browsers) which access the front end will receive the same user interface. The interface is optimized for touch devices. Desktop user are still able to use the interface as normal due to the fact that a bigger button can still be clicked by a mouse pointer.

4.8 Back End Implementation

The Back End part of the system uses PHP as server side scripting language. The functionality, such as an authentication agent or communication to the directory services, are implemented in PHP.

There are several definitions about when a user is authenticated. One definition being when the credentials of the users returning an positive BIND to the directory service.

Another definition is comparing values, which are gathered as user input, with values from the directory. For example the showroom Mobile Desk uses a hash value which is generated in AuthApp by entering a graphical password. This hash is used to send a request to the Back End, validate the hash and response with a result.

4.9 AuthApp Prototype

AuthApp was developed based on previous analysis of the sufficient authentication methods on mobile devices. It is the central part of the mobile authentication system from a user perspective.

AuthApp has several authentication entry points gathering data from the user. The data is used to create an HTTP request which is send to the authentication back end. After receiving an HTTP response from the back end, AuthApp reacts by either continuing in the work flow or displaying an error message to the user.

4.10 Conclusion of Prototypical Implementation

This chapter discusses the implementation of different showrooms on iOS and Android. Each showroom differs from another, due to the combination of authentication methods. Thus, the Front End and Back End were basically the same. The Front End implementation was realized with jQuery Mobile while the Back End implementation uses PHP. A universal prototype that works with HTTP requests has been developed to combine the different demands of operating systems and showrooms.

All authentication methods used by the showrooms are available today. Some of them are more accepted than other. For example credentials are used by users for decades. On the other hand public-key authentication is far less common.

Additionally, a systems interface is used by users not by programmers! Due to this fact, it is vital to think about the users who are operating the system. Therefore, a detailed analysis of usability and testing is indispensable.

5 Evaluating the Usability of Mobile Authentication

The requirement Usability and Functionality has been selected to be verified by an usability test.

McLaughlin and Skinner have defined possible components of usability that have been used: confidence, ease of use, speed and understanding [MS10]. An additional component is the aspect of required background. Those components have been transformed into the following interrogations.

1. Is there a significant difference between the arithmetic average authentication time?

2. Is there a significant difference between the authentication time within the showrooms?

3. Is there a significant difference between the different operating systems focusing on the success quotient?

5.1 Used Methods

To obtain those aspired usability results different methods have been used. In order to receive a distinctive comparison, the testers filled in a personal questionnaire that also inquired their skills. A second method was observation, executed by the interviewer and the assistant. Problems and needed support were noted. Confidence and ease of use was verified by the standardized questionnaire AttrakDiff. To evaluate the understanding and the speed needed for a successful authentication time was measured by using a stopwatch. The tasks were explained in detail and an illustration was given as help. The task was considered as completed, when the tester has authenticated himself successfully. The group of participants amounted to 26 being divided into two groups, the target user group employees (14) and target user group students (12).

Considering Nielsen that the number of participants is 20 at the minimum the usability test produced reliable, replicable and applicable results [Nie94].

The testers have been mixed differently. The age ranged from 20 to 51, the tester group was male dominated and there was a majority of iPhone users present. Twelve of them were students while 14 of them were employees.

That corresponds with the described showrooms and therefore the results can be seen as reliable. Those are the consolidated results of the usability test. The results are subdivided into arithmetic average authentication time, authentication time showrooms and success quotient.

5.2 Arithmetic Average Time for Authentication

On average the fastest authentication could be done within 30 seconds using the iPhone and the showroom QR code. The second fastest authentication could be done with the Android (iPhone has no NFC) using the showroom Mobile Key (37 seconds). The showrooms Location Based enabled both iOS and Android users to authenticate within 39 seconds. The showroom Mobile Desk could be used to authenticate within 41 seconds, while using the showroom Mobile Desk (iPhone) the users needed 46 seconds to authenticate and 49 seconds to authenticate with the showroom QR code (Android). A reason for that difference may be the different frameworks the operating systems use for QR code recognition. This is displayed in figure 2.

Figure 2: Arithmetic Average Time for Authentication of all Showrooms

5.3 Authentication Time Showrooms

The average authentication time shown in figure 3 is independent from the operating system and lists each showroom separately. The fastest showroom was Mobile Key (37 seconds), followed by Location Based (39 seconds) and QR code (40 seconds). The slowest authentication was achieved with the showroom Mobile Desk (44 seconds). It can be assumed that authentication with Mobile Key was possible in 37 seconds, because the user interaction (place the NFC Card on the mobile device) is minimal. Location Based and QR code are similar. The time of 44 seconds for the showroom Mobile Desk may be a result of the unknown graphical password that determined incorrect inputs that let to reentering the password and a corresponding time.

Figure 3: Authentication Time Operating Systems

5.4 Success Quotient

Success quotient describes the number of testers, that were able to authenticate themselves in less than four minutes; retry was possible. The highest success quotient of the showrooms (on average and independent from the operating system) was achieved with QR code, 93 percent.

That may be caused to the fact that QR codes are ubiquitous and the use of NFC techniques is quite new to consumers.

Ignoring the individual showrooms the success quotient of the Android was 86 percent, while the success quotient of the iPhone was 91 percent. The fact that 70 percent of the testers had prior knowledge of using the iPhone should be consulted here.

6   Conclusion

With the assistance of the usability test the results generated by the rating matrix could be specified within the requirements usability and functionality. The following showrooms are advisable (with certain reservations) for a successful mobile authentication with business application: QR code, Mobile Desk and Mobile Key. Those showrooms include the following highly recommended authentication methods: credentials, graphical password, second channel, QR code , certificates and NFC.

With regards to stakeholder interests the showrooms must be looked at in particular; arrangements may be combined differently. The basis for that redefinition is directly interconnected with the standards of usability and functionality, security, accuracy, expenditure and implementation effort are set by each company itself. The rating matrix that was developed can be used for that redefinition by rating the significance of each requirement. In the case of accuracy is being considered as key performance, the results can be multiplied by two. A less important requirement can be multiplied with the factor 0,5. For this reason the developed rating matrix is a useful instrument to find the appropriate authentication methods that match with specific needs to face the challenges of mobile authentication.

References

[AZEH09]   Fadi Aloul, Syed Zahidi, and Wassim El-Hajj. Multi Factor Authentication Using Mobile Phones. *International Journal of Mathematics and Computer Science, 4 (2009), no.2, 65-80*, 2009.

[Bao08]   Lichun Bao. Location Authentication Methods for Wireless Network Access Control. 2008.

[BM02]   Nicky Boertien and Eric Middelkoop. Authentication in mobile applications. 2002.

[CDW04]   Art Conklin, Glenn Dietrich, and Diane Walz. Password-Based Authentication: A System Perspective. *Proceedings of the 37th Hawaii International Conference on System Sciences*, 2004.

[DYB11]   Mohammad Omar Derawi, Bian Yang, and Christoph Busch. Fingerprint Recognition with Embedded Cameras on Mobile Phones. March 2011.

[Gar12]   Gartner. Gartner Authentication Method Evaluation Scorecards, 2011: Assurance and Accountability. 2012.

[Goo]   Google Inc. Android Security Overview. visited 23.05.2012.

[Mou83]   R. T. Moulton. *Network Security*. Datamation, 1983.

[MS10]   Janice McLaughlin and David Skinner. *Developing Usability and Utility: A Compara-tive Study of the Users of New IT*. 2010.

[Nie94]   Jakob Nielsen. *Usability Engineering*. Morgan Kaufman, 1994.

[Rae11]   Jussi Raemaenen. Perceived security in mobile authentication. Master's thesis, Aalto University, School of Electrical Engeineering, August 2011.

[Sch04]   Jean Scholtz. Usability Evaluation. *National Institute of Standards and Technology*, 2004.

[Tan09]   Andrew S. Tanenbaum. *Modern operating systems*. Pearson Prentice-Hall, Upper Saddle River, NJ, 3. ed., pearson international ed. edition, 2009.

[The11]   The Nielsen Company. Generation App: 62 November 2011.

[vT05]   Henk C.A. van Tillborg. *Encyclopedie of Cryptography and Security*. Springer, 2005.

[Wat12]   Robert N. M. Watson. New approaches to operating system security extensibility. Technical Report UCAM-CL-TR-818, University of Cambridge, Computer Labora-tory, April 2012.

[WFMV03]   Robert Watson, Brian Feldman, Adam Migus, and Chris Vance. Design and Imple-mentation of the TrustedBSD MAC Framework. April 2003.

Common Criteria certified open source software – fact or fiction?

Tomas Gustavsson

PrimeKey Solutions AB
Andertorpsv 16
171 54 Solna, Sweden
tomas@primekey.se

Abstract: In 2012 the two open source projects CESeCore and EJBCA were Common Criteria certified [CCP], using open source tools and open source methodologies. As the actual software and its long term evolution is perhaps the most important result for most users, we will look at how certification, distribution and maintenance is managed. Can they be done in an open source way, and is certification always suitable?

The Common Criteria for Information Technology Security Evaluation (Common Criteria) is a standard for IT security certification defined by ISO/IEC 15408 [WP]. The Common Criteria provides trust that processes for specification, implementation and evaluation has been performed in a rigorous and standardized way. Recognized world wide and governed by national certification bodies, Common Criteria is used as requirement for procurement and use of security software in governments, banks and enterprises.

Common Criteria has been criticized for large costs and potential discrimination against Open Source Software [DW]. Given the rigorous system that Common Criteria enforces, how can open source software be certified, and maintained as certified? Drawbacks and benefits of a Common Criteria certification will be described, and how certification limits the maintenance of an open source project.

Common Criteria certified open source software – fact or fiction? After this presentation software developers will be able to determine if their open source project is suitable for Common Criteria certification, whilst software users will have a good idea if they should require certification.

References

[WP]   WikiPedia, Common Criteria, http://en.wikipedia.org/wiki/Common_Criteria
[CCP]   Common Criteria Portal, http://www.commoncriteriaportal.org/
[DW]   David A. Wheeler: Free-Libre/Open Source Software (FLOSS) and Software Assurance / Software Security, , December 11, 2006, http://www.dwheeler.com/essays/oss_software_assurance.pdf
[EJBCA] EJBCA.org website, http://www.ejbca.org/
[CESE]   CESeCore website, http://www.cesecore.eu/

The eID-Terminology Work of FutureID

Bud P. Bruegger, Moritz-Christian Müller

Identity Management
Fraunhofer IAO
Nobelstraße 12
70569 Stuttgart
bud.bruegger@iao.fraunhofer.de
moritz-christian.mueller@iao.fraunhofer.de

Abstract: The paper reports on the experience of the FutureID project in the creation and use of an eID terminology so far. A major part of work has reviewed the state of the art in eID Terminologies. Five existing terminologies have been compared and analyzed in detail to yield unexpected and surprising results. On this basis, FutureID has designed its approach for creation and use of an eID terminology that is currently being implemented in the project. It is hoped that the terminology, its approach, and the related infrastructure will constitute a general community resource, well beyond the scope and duration of the project.1 Section heading

1 Section heading

FutureID[1] – Shaping the Future of Electronic Identity – is a Collaborative Project of the EC's Seventh Framework Programme and is coordinated by Fraunhofer-Gesellschaft. It started in November 2012 for a duration of three years. The consortium is composed of 19 partners from 11 European countries, and combines the multi-disciplinary and complementary competence. The project builds a comprehensive, flexible, privacy-aware and ubiquitously usable identity management infrastructure for Europe. It integrates existing eID technology, trust infrastructures, emerging federated identity management services, and modern credential technologies. It creates a user-centric system for the trustworthy and accountable management of identity claims.

FutureID faces particularly difficult challenges in respect to terminology. Its work is often interdisciplinary and combines technical, legal, economical, and societal aspects. On top of this, FutureID's objective is to reach an unprecedented level of integration that comprises many different types of eIDs, federation protocols, trust infrastructures, platforms, and more, each of which often comes with its specific concepts and thus terms. In comparison, Stork, one of the most prominent eID interoperability projects,

---

[1] www.futureID.eu

decided to use a single federation technology for its whole infrastructure. This is the setting in which the terminology work described in the paper is embedded.

2 Purpose of Terminology Work in FutureID

In such an environment, the efficiency with which project partner can discuss, conceive, and implement various integration concepts depends largely on precise communications and a **common understanding** of the involved concepts. A common understanding is achieved when the meaning associated with words is the same for all participants. Hence, a well-defined terminology that describes the precise meaning of the most important concepts is of vital importance for efficiency. When different partners create components of a whole, different understandings can lead to serious problems when assembling these. A good terminology can go a long way to avoid such problems.

A terminology identifies and defines the concepts that we use to reason and communicate in the field of eIDs. eIDs are a young discipline in the process of moving from a collection of technologies to a more mature science. Our current position in this process is closely related to the concepts we use, the degree of consensus that exists on their meaning, and the degree to which concepts are defined by a single technology or are more general and thus applicable across technologies.

To design a strategy for dealing with terminology in a project such as *FutureID*, it is important to first understand what elements of an eID terminology already exist and at what level of maturity they are.

3 Review of the State of the Art in eID Terminology

To review the state of the art in eID terminology, existing terminologies and glossaries were identified, loaded on a MediaWiki, and compared and analyzed with ad-hoc scripts. The description of this work comprises the main part of this paper.

3.1 Previous Work on eID Terminology

First, links to **fourteen** terminologies were compiled and **seven** were selected. For each, an ad-hoc parser was written to extract terms and their definitions from the original format (mostly PDF) and load them on the MediaWiki. F**ive** were selected since they were comparable in scope and all focused on technical aspects of eIDs (see Table 1). The year col. indicates the range between first and last available version.

Terminology	Label	Year	Reference
Modinis-IDM Glossary	Modinis	2005	[Mo05]
Identity Management for eGovernment (IDEM)	IDEM	2005-2007	[HuAl07]
STORK Glossary	Stork	2008-2009	[Pi08]
U.S. IdM Task Force Glossary	US		[Id05]
ISO/IEC 24769-1	ISO	2011	[Is11]

Table 1: Overview of the Intersection of five Terminologies

The terminologies that were also parsed and loaded but excluded from this analysis were the Eurosmart Glossary [Eu13] and the glossary contained in the European Draft Regulation on Privacy Protection [Eu12].

3.2 Intersection Analysis of Five Existing eID Terminologies

Each of these terminologies was considered to be a set of terms and the intersection of all five terminologies was computed. Where synonyms were defined, they were treated like normal terms (see Table 2). It lists the total number of terms (incl. synonyms) defined by the terminologies, the number of "isolated terms", i.e., terms that occur only in a single terminology, and the percentage of isolated terms in each terminology.

Characteristics	Stork	IDEM	Modinis	ISO	US
Total terms	123	195	45	43	125
Isolated terms	67	132	4	18	72
Percentage	54%	68%	9%	42%	58%

Table 2: Overview of the Intersection of five Terminologies

It is evident, that the percentage of isolated terms is above 50% for all terminologies but Modinis, which was taken into account by Stork and IDEM, and partly by ISO. These latter two terminologies are also relatively small. So it is more likely to focus on the most important concepts that are shared by other terminologies.

Table 3 shows in how many terminologies each unique term is contained. Out of the 377 unique terms, only 10, i.e. a bare 2.7% are contained in all five terminologies. Even the terms contained in at least three out of the five terminologies are only 36, i.e., less than 10%. Consequently, 90.5% of all unique terms are contained by at most two terminologies, and of these 77.7% in only one.

Characteristics	No. Terms	Perc.	accum. No.	accum. Perc.
terms contained in all 5 terminologies	10	2.7%	10	2.7%
terms contained in only 4 terminologies	14	3.7%	24	6.4%
terms contained in only 3 terminologies	12	3.2%	36	9.5%
terms contained in only 2 terminologies	48	12.7%	84	22.3%
terms contained in only 1 terminology	293	77.7%	377	100.0%

Table 3: Detailed Intersection of five Terminologies

The terms contained in all terminologies are:
anonymity, attribute, authentication, context, credential, entity, federated identity, identifier, identity, identity management

The terms that are included in four terminologies are:
identification, pseudonym, enrolment, identity provider, relying party, access control, assertion, delegation, digital identity, principal, privacy, role, trust, trusted third party

The terms that are included in three terminologies are:
identity federation, verifier, partial identity, verification, personally identifiable information, characteristic, confidentiality, corroboration, identified entity, profile, registration, token

This intersection analysis illustrates that there is an unexpectedly low level of agreement on the key terms/concepts of identity management. This is not what would be expected from a mature field. It much rather indicates that the field is young and immature and that its maturation will be measurable in an increasing level of agreement. The hypothesis that the lack of agreement is caused by cultural differences between a Europan and a North American view cannot be sustained as shown in Table 4. Here, only the three European terminologies, that are even linked by their lineage, are intersected. While the situation improves some, the lack of agreement still subsists.

Characteristics	No. Terms	Perc.	accum. No.	accum. Perc.
terms contained in all 3 terminologies	26	9.2%	26	9.2%
terms contained in only 2 terminologies	28	9.9%	54	19.1%
terms contained in only 1 terminology	229	80.9%	283	100.0%

Table 4: Intersection of the three European terminologies

Further evidence for this is given in Table 5 where the agreement between pairs of terminologies is measured. Through the band, the overlap remains very low.

Terminologies	Common Terms	Percentage
Modinis-Stork,	30	21.7%
Modinis-IDEM	36	17.6%
Modinis-ISO	45	20.5%
Modinis-US	23	15.6%
Stork-IDEM	40	14.4%
Stork-ISO	15	9.9%
Stork-US	34	15.9%
IDEM-US	36	12.7%
IDEM-ISO	17	7.7%
ISO-US	22	15.1%

Table 5: Overlap between Pairs of Terminologies

With the help of Euler diagrams, graphical representations of intersections are particularly suited to understand the level of agreement of terminologies. The resulting diagrams are presented in the following figures 1 through 5. Each set of terms is represented by an ellipsis. Each subset of the intersection is labeled with the number of terms it contains. For example, the number under the terminology name represents its isolated terms.

Figure 1 shows the intersection of the European terminologies Stork, IDEM, and Modinis. Since Modinis was taken into consideration in both, the creation of the Stork and IDEM terminologies, it is not surprising that most of its terms were adopted. Only five terms remained unadopted by the other terminologies; they are:

identity management application, nym, privacy enhancing technology, proxy, unique identity

Since the Modinis Glossary is relatively small and was mostly absorbed in the other two terminologies (30 and 36 terms out of 45 for Stork and IDEM), it is less relevant than the other two to understand the relationship with the US and ISO terminologies.

The intersection in Figure 1 shows further, that apart from the common Modinis terms, only 14 additionally added terms are in common. This compares to 93 additional terms in Stork and 159 in IDEM. Surely, this strong divergence after adaptation of a common core of Modinis term is amazing.

Fig. 2 shows the relation of the two major European terminologies with that of the U.S. IdM Task Force. Again, a strong divergence of the three terminologies is evident. The hypothesis of possible cultural differences must be rejected also here, since the communality between the European terminologies is in all respect comparable to the communalities with the US terminology.

Fig. 3 shows that the major European terminologies have about the same level of agreement with the ISO terminology. The level of agreement seems even less that with the US IdM Task force. This may be partially explained by the smaller size of the ISO terminology.

Figure 1: Stork-IDEM-Modinis   Figure 2: Stork-IDEM-US   Figure 3: Stork-IDEM-ISO

Also partially explainable by the size of the ISO terminology is the conclusion from Figures 4 and 5 that the European terminologies have more communality with the US IdM Task Force one than with ISO. Also visible is that ISO has only a slightly higher level of agreement with the US terminology compared to the European ones. Figure 6 compared to Figure 1 illustrate that while Modinis has strongly influenced Stork and IDEM, it has much less communalities with US and ISO. This is probably explainable by the fact that the European terminologies have taken Modinis into account.

Figure 4: Stork-ISO-US   Figure 5: IDEM-ISO-US   Figure 6: Modinis-ISO-US

3.3 Conclusions on the State of the Art of eID Terminologies

The detailed analysis of communalities of the existing terminologies illustrates, that the level of agreement is unexpectedly low. Apart from Modinis, whose choice of terms show a strong overlap with Stork and IDEM, no clusters or families of terminologies with a higher level of agreement could be found.

This unexpected result has a strong impact on terminology work in the field of identity management. It is not possible to start from a consolidated set of core terms. A consensus decision on which the major concepts of identity management are has yet to be made. This is even more astonishing since the divergence is already present in the choice of terms, while semantic discrepancies between different definitions of the same term would have been expected to be more likely. Terminology work has to attempt to avoid creating yet another diverging glossary defined by a relatively small and closed group of experts. Therefore, for furthering the field, it is necessary to launch a collaborative process of consensus building that should attempt to reach out to as many experts as possible, integrating multiple disciplines and backgrounds.

4 The FutureID Approach to Terminology Creation and Use

A glossary *document* managed by an editor, as used by most pre-existing projects, is evidently an ill-suited choice for supporting a large-scale consensus process. FutureID therefore decided early on to use a collaborative semantic wiki (MediaWiki with a semantic extension) to support the terminology work.

A wiki-based approach makes it possible to load the existing terminologies, support analysis as the above, and to more easily see different terminologies side by side in order to compare them. For example, the current wiki makes it possible to see all related existing definitions of a given term in order to make more founded decisions about its semantics. A semantic wiki also allows to model properties and relations that would not be possible in a paper-based approach. This allows for both a richer expression of semantics, as also support of processes. For example, it is possible to capture the technologies for which a concept is relevant, state the community who agrees on a given definition, or model synonyms. A wiki can also support editorial processes such as assigning certain terms to certain editors or capture the status of a review process. The wiki's capability of easily annotating anything is an important support for discussions and consensus building.

Another aspect that the FutureID terminology work attempts to address is motivation. Terminology, surely when managed as a glossary document, is probably "boring" for most experts. Experience of past projects have demonstrated that is far easier to select and define terms, than getting them used in a consistent way by all project participants.

A machine readable management is therefore important in order to create tools and views that confront project partners with terminology and provide useful services. For example, *FutureID* has created a tool that automatically creates custom glossaries for project deliverables. A "motivational design" of the approach seems to be a major success factor for the creation of a valid terminology.

5 Conclusions and Future Work

While *FutureID* has managed to set up an initial infrastructure and set of tools for the creation and use of a terminology within the project, much of the work of deploying to participants them is only just starting. Social challenges and collecting experiences that identify new needs for tools and infrastructure are still ahead of us. We hope that *FutureID's* work will make a major step in the maturation of the eID terminology and thus the "science of eIDs". We also plan to gradually extend the community beyond the project itself and create an infrastructure that can remain as a resource beyond *FutureID's* duration.

References

[Eu12]	Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation), Article 4 'Definitions', page 41 (printed, 42 electronic),http://ec.europa.eu/justice/data-protection/document/review2012/com_2012_11_en.pdf, 2012.
[Eu13]	Eurosmart Glossary, http://www.eurosmart.com/index.php/glossary.html.
[HuAl07]	Xavier Huysmans and Brendan Van Alsenoy, IBBT, Identity Management for eGovernment (IDEM), Deliverable 1.3 Conceptual Framework, Annex I. Glossary of terms (v1.07), https://projects.ibbt.be/idem/index.php?id=161, 2007.
[Id05]	U.S. National Science and Technology Council (NSTC), Subcommittee on Biometrics and Identity Management, Identity Management Task Force Report, Annex K, Identity Management Glossary, http://www.biometrics.gov/documents/idmreport_22sep08_final.pdf, 2008.
[Is11]	ISO/IEC24760-1, Information technology — Security techniques — A framework for identity management —Part 1: Terminology and concepts, First edition, 2011-12-15, chapter 3: Terms and definitions, http://standards.iso.org/ittf/PubliclyAvailableStandards/c057914_ISO_IEC_24760-1_2011.zip, 2011.
[Mo05]	Common Terminological Framework for Interoperable Electronic Identity Management, Consultation paper, v2.01., https://www.cosic.esat.kuleuven.be/modinis-idm/twiki/bin/view.cgi/Main/GlossaryDoc, November 23, 2005
[Pi08]	Ana Piñuela (Ed.), STORK Glossary (v2.0), https://www.eid-stork.eu/index.php?option=com_processes&Itemid=&act=streamDocument&did=615, 2008.

Landscape eID in Europe in CY2013

Detlef Houdeau

Chairman Working Group eID of EUROSMART
c/o Infineon Technologies AG
Senior Director Business Development
Infineon Technologies AG
Am Campeon 1-12
D 85579 Neubiberg
detlef.houdeau@infineon.com

1 Executive Summary

CY 1998 the first electronic ID document in the public domain was launched in Europe. CY 2010 the EU Commission has published the roadmap of the Digital Agenda for Europe (DAE) and the roadmap for one single market, which should be in place by 2015. In June 2012 the EU Commission has published the proposal for a regulation on electronic identification as cornerstone for the growing of the European single market. This speech reviews all 16 existing electronic identity programs based on two factor authentication and analyses mainstreams on electronic identification, authentication and signature and these related services.

The content of the speech is related to a white paper of Eurosmart, which was published in July 2013.

Beside 16 national profiles, the following aspects are highlighted:

- Policies on national ID cards
- Minimum age of the card holder
- Travel function with the card
- Biometric data stored in the card
- Card lifetime and fee
- e-services with the card
- Mobile applications
- e-Signature
- Used technology platform

These 16 states are **Austria, Belgium, Czech Republic, Estonia, Finland, Germany, Ireland, Italy, Latvia, Lithuania, Monaco, Portugal, Spain, Serbia, Sweden and the Netherlands.**

Not Built On Sand - How Modern Authentication Complements Federation

Dr. Rolf Lindemann, Nok Nok Labs, Inc.

Nok Nok Labs, Inc.
4151 Middlefield Road
Palo Alto, California 94303, USA
rolf@noknok.com

Abstract: Even after 40 years of IT innovations, passwords are still the most widely used authentication method. They are inherently insecure. Neither users nor service providers handle passwords appropriately. On the other hand more than 1 billion Trusted Platform Modules (TPMs) and more than 150 million secure elements have been shipped; microphones and cameras are integrated in most smart phones and fingerprint sensors and Trusted Execution Environments (TEEs) are on the rise. There are better ways for authentication than passwords or One-Time-Passwords (OTPs).

The Fast Identity Online (FIDO) Alliance has been founded to define an open, interoperable set of mechanisms that reduce the reliance on passwords.

We explain how secure hardware in conjunction with a generic protocol can help overcoming today's authentication challenges and how this protocol can be used as a solid basis for federation.

Motivation

Passwords don't work: In 2007, the average user had 25 accounts, used 6.5 passwords and performed logins 8 times a day [FlHe07]. Today, things are much worse. An analysis of 6 million accounts showed that 10,000 common passwords would have access to 98.8% of the accounts [Trus10]. This basically means that only 1.2% of the users chose strong passwords. Even when looking at passwords for banking accounts only, it can be found that 73% of users shared their online banking password with at least one *non-financial* site [CSA10], which means that when the non-banking site gets hacked, the banking account is threatened.
"Account or service hijacking is not new. Attack methods such as phishing, fraud, and exploitation of software vulnerabilities still achieve results. Credentials and passwords are often reused, which amplifies the impact of such attacks." [CSA10]. It's not only about security. According to a recent study, more than 45% of the online transactions fail "Very Frequently" or "Frequently" due to authentication problems [Pone13].

Several proposals to replace passwords have been made. A good analysis can be found in [BHOS12].

Silos of Authentication: Current alternative technologies require their respective proprietary server technology. The current authentication architecture therefore consists of silos comprising the authentication method, the related client implementation and the related server technology.

Heterogeneous Authentication Needs: Authentication is used for electronically initiating high value money transactions and for accessing the personal purchase history in an online bookshop. The security needs are different.

Not all users are equal. A recent survey shows that more than two thirds of the participants in the study prefer authentication without sharing personal information, approx. 50% would accept use of a multi-purpose identity credential and 23% in the US and 40% in Germany would accept biometrics based authentication [Pone13].

The one authentication method satisfying all needs seems to be out of reach.

The FIDO Approach

We propose to (a) separate the user authentication methods from the authentication protocol and let an entity called FIDO Authenticator glue both together, and (b) to define an attestation method in order to attest the identity of the FIDO Authenticator to the relying party. Given this information, the relying party is able to infer the related assurance level (e.g. as defined in [BDN+13]). The assurance level can be fed into internal risk management systems. The relying party can then add implicit authentication methods as needed.

In the FIDO approach, standardized challenge response based cryptographic authentication schemes are used between the FIDO Authenticator (controlled by the user) and the FIDO Server (controlled by the relying party). The FIDO Authenticator can implement any user authentication method without requiring specific support in the FIDO Server and hence avoiding "silos" of authentication. Successful user authentication unlocks the relying party specific cryptographic authentication key.

The FIDO Protocol

The FIDO protocol supports the functions Discovery, Registration, Authentication and Transaction Confirmation.

The discovery enables relying parties to explore user authentication methods supported by the user's computer and hence handle heterogeneous client environments. The relying party can specify a policy for selecting FIDO Authenticators best suited for the specific purpose.
As part of the registration operation, the FIDO Authenticator generates a key pair specific to the relying party. The relying party binds the public key to a specific entity. This might be an existing user identity already present in the relying party's system or it

might be a user identity to be created. Using a dedicated key for each relying party enhances the user's privacy as two relying parties cannot link transactions to the same user. Storing only the public key at the relyiung party makes the FIDO protocol resililient to leaks from other verifiers.

The Authentication operation supports single or multiple FIDO Authenticators to be involved. Each FIDO Authenticator might be implemented to represent either simple or strong authentication / two factor authentication [ECB12]. The Authentication operation is used to establish an authenticated channel between the Browser / App and the relying party Web Server.

The Transaction Confirmation allows the user to approve and authenticate a particular well-defined transaction to the relying party. It is more secure as it doesn't rely on a Web Browser / App to not misuse an authenticated session.

This leads to the following reference architecture:

Fig. 1. FIDO Reference Architecture

The FIDO Authenticator is a concept. It might be implemented as a software component running on the FIDO User Device, it might be implemented as a dedicated hardware token (e.g. smart card or USB crypto device), it might be implemented as software leveraging cryptographic capabilities of TPMs or Secure Elements or it might even be implemented as software running inside a Trusted Execution Environment.

The User Authentication method could leverage any hardware support available on the FIDO User Device and hence avoid additional costs, e.g. Microphones ($\rightarrow$ Speaker Recognition), Cameras ($\rightarrow$ Face Recognition), Fingerprint Sensors, or behavioral biometrics [ObSa].

Attestation

The relying party is interested in estimating the risk of a transaction. This risk depends on the assurance level of the authentication (and other factors). The assurance level

depends on (a) the authentication method and (b) the certainty that the legitimate user controls the relevant portions of the client device. In the case of Transaction Confirmation, this could be limited to the FIDO Authenticator. In the case of Authentication it will also include the Browser / App or User Agent in general. Risk based authentication [Will06] methods try to estimate (b). Authenticator attestation provides a cryptographic proof of the FIDO Authenticator being used to the relying party, addressing (a). Trusted platform modules already support the concept of (pure) attestation [TCG08].

The FIDO Authenticator maintaines cryptographic authentication keys and performs the user authentication. The attestation provides a cryptographic proof of the Authenticator model to the relying party and hence allows the relying party to infer the assurance level from it.

FIDO and Federation

From a user's perspective, Federated Identity Management is a method that allows accessing privileged information across autonomous security domains after authenticating once. From an organization's perspective, it also "… allows organizations like enterprises and service providers to securely exchange user information across partners, suppliers and customers." [LaMo12]. InCommon is one example of successful real-world federation systems. SAML and OpenID Connect are examples for popular federation standards.

Federated Identity Management systems expect the user to authenticate to an Identity Provider (IdP). This user authentication method is relevant to the IdP, but not directly in the scope of current federation standards. Most IdPs still use password based authentication.

FIDO addresses this "first mile" authentication of the user to the IdP while leaving the user vetting up to it. FIDO protocol makes reliable information about the authentication assurance level available to the IdP ($\rightarrow$ attestation). Some of the federation standards[1] already support sharing this knowledge with the service provider. This enables IdPs to support heterogeneous authentication methods and it enables service providers to make informed decisions about the transaction risk.

References

[BDN+13]   William E. Burr, Donna F. Dodson, Elaine M. Newton, Ray A. Perlner, W. Timothy Polk; Computer Security Division, Information Technology Laboratory and Sabari Gupta, Emad A. Nabbus; Electrosoft Services, Inc., "Electronic Authentication Guideline," National Institute of Standards and

---

[1]   E.g. OpenID Provider Authentication Policy Extensions v1.0.

Technology (NIST), 2013.

[BHOS12]   Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano, "The Quest to Replace Passwords - A Framework for Comparative Evaluation of Web Authentication Schemes," in *Proceedings of IEEE Symposium on Security and Privacy*, Oakland, 2012.

[Burn13]   M. Burnett, "More Top Worst Passwords," 20 June 2011. [Online]. Available: http://xato.net/passwords/more-top-worst-passwords/. [Accessed 3 April 2013].

[CSA10]   Cloud Security Alliance, "Top Threats to Cloud Computing, v1.0," 2010.

[ECB12]   European Central Bank, "Recommendations for the Security of Internet Payments," Frankfurt am Main, 2012.

[FlHe07]   D. Florêncio and C. Herley, Microsoft Research, "A Large-Scale Study of Web Password Habits," Redmond, 2007.

[LaMo12]   S. Landau and T. Moore, "Economic tussles in federated identity management," 1 October 2012. [Online]. Available: http://www.firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/vie wArticle/4254/3340. [Accessed 7 February 2013].

[ObSa]   M. S. Obaidat and B. Sadoun, "Keystroke Dynamics Based Authentication," in *Biometrics. Personal Identification in Networked Society*, Kluwer Academic Publishers, pp. 213-229.

[Ping10]   Ping Identity, "The Primer: Nuts and Bolts of Federated Identity Management," 2010.

[Pone13]   Ponemon Institute LLC, "Moving Beyond Passwords: Consumer Attitudes on Online Authentcation - A Study of US, UK and German Consumers," 2013.

[TCG08]   Trusted Computing Group, "Trusted Platform Module (TPM) Summary," 2008.

[Trus10]   Trusteer, Inc., "Reused Login Credentials," New York, 2010.

[Will06]   Gregory D. Williamson, GE Money – America's, "Enhanced Authentication In Online Banking," *Journal of Economic Crime Management,* pp. Fall 2006, Volume 4, Issue 2, 2006.

An open source eID simulator

Holger Funke, Tobias Senger

holger.funke@hjp-consulting.com
tobias.senger@bsi.bund.de

Abstract: The German BSI started a project for an open source eID simulator. It will allow a wide range of personalisation, is more flexible than real cards and is free to use.

Background and goals

There is a rising need of test cards for developers of eID clients and companies which want to offer services by using the eID function of the German eID card. Today is difficult get test cards for new developers who want to evaluate the eID functions in their systems. Also for improvements and development of new protocols a open implementation of the eID functions would be helpful. Therefore the German BSI started a project together with HJP for an implementation of an open source eID simulator which provides all functionality of the German eID card.

We will implement all functions of the eID functions as described in BSI TR-03110 and provide all sources of this implementation to interested developers which wants to participate in improving the simulator or just use it for their development. Part of this project is also a "virtual card reader driver" which provides direct PC/SC access to the eID card simulator (see figure 1).

```
PC
  ┌──────────────┐      ┌──────────────┐
  │     eID      │      │     eID      │
  │  simulator   │      │    client    │
  └──────────────┘      └──────────────┘
     local socket          PC/SC
        virtual card reader driver
```

Figure 1: Link between eID client and eID card simulator

Later on we will start to port the simulator to mobile NFC devices (e.g. Android smartphones). Over the NFC interface the mobile device shall act in card emulation mode to simulate an eID card with all functions.

Authentication on high critical infrastructures using interoperable federated identities

Armin Lunkeit, Jürgen Großmann

OpenLimit SignCubes AG
armin.lunkeit@openlimit.com

Fraunhofer Institute for Open Communication Systems FOKUS
juergen.grossmann@fokus.fraunhofer.de

Abstract: The technical guideline TR-03109 divides between the roles of the SMGW technician and the Gateway administrator whereas the Gateway administrator gains full access to the SMGW and the service technician has only very limited access rights. In many scenarios the service technician will also need full access to the Smart Meter Gateway which means that he must be able to change its role. Federated identities can help to create a solution that keeps the strict role enforcement between service technician and Gateway Administrator. This article presents an approach on the background of the current Smart Grid development and identity technology adopting approaches used for the German national ID card. A short discussion pertaining threats and risks completes the discussion.

1 Smart Grid Infrastructures – German approach

Based on the European Directive 2003/54/EG the EU member states are required to introduce intelligent measurement devices for electricity. The German government regulates this requirement in the law on energy industry (Energiewirtschaftsgesetz – EnWG), especially in §21c, §21d and §21e. The ministry of economics (BMWi) requested the Federal Agency for Security in Information Technology (Bundesamt für Sicherheit in der Informationstechnik - BSI) to set up a technical guideline [3109-1] and Common Criteria protection profile [PP] addressing security and interoperability. The German smart grid approach requires a gateway component for data collection, consumption displaying and secure communication with meters, users and external entities. This component is called "Smart Meter Gateway (SMGW)". The Smart Meter Gateway itself is not a measurement device; it is a data aggregation and communication unit that protects the privacy, integrity and authenticity of the consumer data during local storage and network communication. A hardware security module is built into the Smart Meter Gateway for protection of key material and cryptographic operations. Three logical and physical distinct networks are defined for the Smart Meter Gateway:

- The Wide Area Network (WAN)

- The Local Metrological Network (LMN)
- The Home Area Network (HAN)

The connection in the Local Metrological Network is required for communication with external meters. The Smart Meter Gateway itself is not a metering unit for electricity or gas; it serves a data collection unit that is responsible for the secure transfer of the collected consumption data to the energy supplier for billing reasons. The connection in the HAN network provides a report of the values measured by the meter responsible for his household and also provides a transparent proxy channel into the WAN. This proxy channel is required for proprietary communication of local power suppliers (e.g. solar panels) with external entities. At least the WAN connection provides communication services with external entities (e.g. the power supplier delivering energy to the household).

Figure 1 Smart Meter Gateway and it's environment [PP]

As shown in figure 1, one of the external entities in the WAN is the gateway administrator (GWA) who is a trustworthy external entity with the ability to control and configure the SMGW. The gateway administrator is a technical role and responsible for configuration of tariffing profiles, certificate management on the Smart Meter Gateway, firmware updates etc. and has also access to several log information. It is clearly stated that the gateway administrator does not have access to user any measured consumption data. Figure 1 does also show the Service Technician who has access to the Smart Meter Gateway from the Home Area Network and is able to access the system log information. The service technician has no possibility to initiate a firmware update or similar

operations. Vendors of Smart Meter Gateways must proof the fulfillment of the technical and security requirements defined in [PP] and [3109-1] of their product by Common Criteria and technical guideline certifications which also check the strict role separation.

Today's energy infrastructures use Command and Control Centers which allow the central administration of devices in the energy network. It is highly probable that these central communication nodes will also be responsible for the management of the Smart Meter Gateway in the household. They will act as smart grid management nodes and execute the gateway administrator role. We claim that the service technician will be required to perform operations (e.g. configuration, initiation of firmware updates) which are not foreseen for this technical role. Personal maintenance of a smart device in the household by a service technician is a cost issue. In case that a problem is encountered on a Smart Meter Gateway that needs to be solved by trained maintenance personal, this maintenance personal must possibly be able to gain full access to the Smart Meter Gateway similar to the gateway administrator. Therefore this article focuses on the administrative access to the Smart Meter Gateway. We discuss secure authentication of an administrator on the Command and Control Center and the enrollment of security policies in order to secure the access to the Smart Meter Gateway via potentially unsecure wide area networks.

2 Technical Background

2.1 Electronic Identities and Authentication

One approach for authentication with a secure token is the use of the electronic citizen card (nPA). This electronic citizen card comes with a rich set of innovative authentication functionality, e.g. PACE (Password Authenticated Connection Establishment) and EAC (Extended Access Control). These mechanisms have been widely discussed in publications ([BDFK11], [3110]).

The basic communication model of the German citizen card in order to access identity information stored on the card contains a trusted remote terminal, typically called eID-Server. The citizen card is served by a local eID-Client like the AusweisApp or the Open eCard-App. The communication between the eID-Client and the eID-Server is standardized in the technical guideline TR-03112-7. Upon this protocol stack several communication models are used, e.g. SOAP and SAML. The use of the German identity card in a federated identity environment allows its use in typical scenarios like Single Sign On (SSO). Research projects like SkIDentity are dedicated to definition of secure and trusted identity exchange in the cloud and utilize a rich set of different authentication technologies in one federated identity environment. Especially SkIDentity addresses the issue of using several identity tokens. The approach of using the electronic citizen card might face difficulties due to legal reasons but other technologies adopting the EAC-approach are available [OLSC].

2.2 Administrative Access to the SMGW

The mutual authentication between the gateway administrator and the smart meter gateway utilizes digital certificates. A signed UDP packet is sent to the smart meter gateway by the administrator and the new mutual authenticated TLS channel is established. The assignment of the administrator role on that new connection is mainly based on the digital certificate used by the smart grid management node during the TLS handshake phase. Once authenticated as gateway administrator, the Smart Meter Gateway allows the full administrative access, e.g. configuration of communication routes or key management of the built in HSM (hardware security module). The communication between the gateway administrator and the Smart Meter Gateway is secured by a mutual authenticated TLS channel using elliptic curve cipher suites. The smart meter gateway therefore owns a private key which is protected by the built in HSM, the root anchor of the Smart Meter Gateway PKI and the administrator's certificate are also configuration items of the gateway.

3 Electronic Identities in the Smart Grid Infrastructure

The technology of the German citizen card and its adoptions (e.g. the truedentity technology of OpenLimit) make use of secure communication protocols and provide a secure and reliable authentication. The authentication process delivers an attested digital identity. Based on that digital identity the enrolment of security policies is undertaken: Rights are granted and limitations are enforced. Using this approach the administrative access to smart grid devices even from untrusted networks is possible: The electronic identity contains security attributes that are verified and attested by an Identity provider and the smart grid management node can rely on this information. The communication between the service technician and the smart grid management node is protected by EAC and TLS.

3.1 Adoption to Smart Grids

[PP] and [3109-1] define the roles of the gateway administrator and the service technician with different rights. The gateway administrator is a trustworthy entity that has full access to the Smart Meter Gateway. This includes the following:

- Configuration for measurement data, their processing and submission of electronic measurement data to external market entities
- Installation of firmware updates
- Configuration of access rights for external market entities within in the Smart Meter Gateway
- Configuration of the integrated security module
- Configuration of certificates in the Smart Meter Gateway

For privacy reasons the gateway administrator is not allowed to view current measurement data[1]. The service technician is only allowed to access system log information and diagnosis information. This enforces a strict separation between both roles. Beneath this role separation it must be taken into account that the service technician in person will need to change its role in order to act as a gateway administrator. This is required in order to install new communication profiles or firmware updates in case that the Smart Meter Gateway has encountered a problem is no longer working properly. In this case the service technician needs the possibility to access the smart grid management node in order to perform service operations that are only foreseen to be performed by the gateway administrator. Therefore an authentication method on the smart grid management node is required which enforces a reliable identification of the service technician in order to change his role. This role of the *natural person* is changed from service technician to gateway administrator.

3.2 Solution Concept

The previous chapter with a roughly description of the authentication framework we explained the technical interface for accessing electronic identity information. The service technician needs to authenticate itself with a secure token on the smart grid management node. Based on the authentication data this central system will gain access rights to the service technician so the service technician will change its role and now act as gateway administrator.

Figure 2 Solution concept

---

[1] Due to [PP] and [3109-1] this is only allowed for the service provides for billing reasons.

We introduce an authentication and communication gateway with the responsibility to authenticate the external entity and is moreover responsible to provide a secure remote communication channel that offers an appropriate level of security. Therefore an extension of the communication model described in the technical guideline TR-03112-7 is foreseen. The current standard describes the establishment of a TLS secured channel from the remote terminal to the client application and provides additional protection for the communication between the remote terminal and the chip with Secure Messaging keys agreed by using the elliptic curve version of Diffie-Hellman (ECDH). In our approach the secured channel (TLS and Secure Messaging Chanel) end in the client application. The client application makes use of this channel for transmission of commands and data from the external entity (service technician) to the authentication and communication gateway. Therefore an alternative token than the nPA is required, e.g. a chip card providing authentication keys for mutual TLS authentication.

This approach addresses one of the main issues of authentication: Even if the authentication is secure, the security of the consecutive communication depends on the provided security of the application utilizing the authentication service. Threat scenarios like session hijacking etc. are still relevant for web applications if authentication and communication security are not linked together. The presented approach combines authentication and offers reuse of the established secure communication channel so the channel is bound to the authentication based on the electronic identity.

This solution does not involve new exploitable interfaces and communication channels. The smart grid management node relies on the authenticity of the electronic identity and applies security policy for transmission of administrative commands initiated from a potentially untrusted network to the Smart Meter Gateway. The use of the EAC-secured channel between the service technician and the authentication and communication gateway in combination with the TLS provides a comprehensible security level.

3.3 Threats and Risks

Threat modeling analyzes the security of a system (hardware, software, networks) by utilization of assets, objectives, threats, attackers, vulnerabilities and countermeasures. This methodology is used by different IT-security frameworks, e.g. Common Criteria or the CORAS approach. The description of the threat model for the smart meter gateway is part of the protection profile [PP]. One important asset in [PP] is the privacy of the user's data. Neither the gateway administrator nor the service technician have permission to access any data that is specific for the customer. This includes consumption data as well as any other personal data. From the perspective of the smart meter gateway, our approach does not introduce any new threats to the smart meter gateway. The communication model of the smart meter gateway remains the same as required in the protection profile [PP] and the technical guideline [3109-1]. One of the essential threats to smart grid infrastructures is the model of central administration: We claim that the attacker in the network is more interested in manipulation of a smart grid management node than in attacking a single smart meter gateway. In our approach the authentication and communication gateway will be the object of interest for an attacker because it

offers the functionality of acting as a gateway administrator. This is attack scenario remains relevant even for the approach presented in [3109-1] with the difference that the smart grid management node would directly be offended. The presented approach does therefore not introduce a new threat to the smart grid on this level.

We claim that the most relevant threat that must be taken into account is a malicious service technician. The protection profile does explicitly assume a person in role service technician as a potential attacker so our solution must also deal with that assumption. This threat might not be mitigated by a single measure, several measures must be combined:

- The smart meter gateway must offer self protection of its IT-components (e.g. check of firmware integrity and authenticity).
- The security policies applied to the service technician with remote administrative access to the smart meter gateway must ensure that no malicious operations can be performed and access rights are limited
- The ID management system must offer the possibility of identity revocation in case that a particular service technician has been identified as attacker

Our ongoing research will be focused on the analysis of the remaining risks resulting from that approach using CORAS ([BBD+06]), which provides a model based method for risk assessment. CORAS includes a methodology, a formal language and tool support for the analysis and assessment of risks and consists of eight steps for the risk assessment process. We have planned to consider current research activities, e.g. the DIAMONDS [ITEA2] project with the goal not only to identify risks but also to provide model of how vulnerability will influence the security of the whole system.

4 Conclusion

We discussed a possible use case for federated identities on the example of accessing a smart grid administration node through an untrusted network in order to gain administrative rights to a person in role service technician. Our approach does not utilize privilege escalation to the role of the service technician; it offers an approach how a natural person can change its role from service technician to gateway administrator. The benefit of that solution is the binding of an electronic identity to a secure communication channel. An authenticated communication channel is established and the authentication protocol is based on the protocol used for the German citizen card. Therefore a new component – the authentication and communication gateway – is introduced which is responsible for provision of the authenticated, secure communication channel. This article discussed the approach on the example of the smart grid environment but it can easily be adopted to other scenarios with the same challenges. We identified that no new threats are introduced pertaining the smart meter gateway and its environment but we identified that the assumption of the malicious service technician must be taken into account. Measures will be required that mitigate this scenario. Currently we are working on the analysis of this scenario in order to provide a full risk assessment of this scenario.

References

[3109-1]   Bundesamt für Sicherheit in der Informationstechnik: Technische Richtlinie BSI-TR-03109-1, Version 1.0, Bonn 18.03.2013, https://www.bsi.bund.de/ContentBSI/Publikationen/TechnischeRichtlinien/tr03109/index_htm.html (Stand: 18.03.2013)

[3110]   Bundesamt für Sicherheit in der Informationstechnik (Ed.): BSI TR-03110 Technical Guideline Advanced Security Mechanisms for Machine Readable Travel Documents, https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/TechGuidelines/TR03110/TR-03110_v2.1_P1pdf.pdf, 2012

[3112-7]   Bundesamt für Sicherheit in der Informationstechnik (Ed.): Technical Guideline TR-03112-7 eCard-API-Framework – Protocols Version 1.1.2 28. February 2012,https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/Publikationen/TechnischeRichtlinien/TR03112/API/api1_teil7_pdf.pdf?__blob=publicationFile, 2012

[BBD+06]   Folker den Braber, Gyrd Brændeland, Heidi E. I. Dahl, Iselin Engan, Ida Hogganvik, Mass S. Lund, Bjørnar Solhaug, Ketil Stølen, Fredrik Vraalsen (Ed.): The CORAS Model-based Method for Security Risk Analysis, SINTEF, Oslo 2006, http://www.uio.no/studier/emner/matnat/ifi/INF5150/h06/undervisningsmateriale/060930.CORAS-handbook-v1.0.pdf, 2006

[BDFK11]   Jens Bender,Ozgur Dagdelen, Marc Fischlin, Dennis Kügler (Ed.): The PACE-AA Protocol for Machine Readable Travel Documents, and its Security, http://fc12.ifca.ai/pre-proceedings/paper_49.pdf, 2011

[HHR+11]   Detlef Hühnlein, Gerrit Hornung, Heiko Roßnagel, Johannes Schmölz, Tobias Wich, Jan Zibuschka: SkIDentity – Vertrauenswürdige Identitäten für die Cloud, http://www.ecsec.de/pub/2011_DACH_SkIDentity.pdf, 2011

[ITEA2]   ITEA2 – DIAMONDS: http://www.itea2-diamonds.org/index.html, 2013

[OLSC]   OpenLimit SignCubes AG: https://www.openlimit.com/de/produkte/truedentity.html, 2013

[PP]   Bundesamt für Sicherheit in der Informationstechnik: Protection Profile for the Gateway of a Smart Metering System, Version 1.2, 18. March 2013 https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/SmartMeter/PP-SmartMeter.pdf?__blob=publicationFile, 2013

Password Management through PWM

Menno Pieters

Consulting
Everett NL
Wiersedreef 5-7
3433ZX Nieuwegein, NL
menno.pieters@everett.nl

Abstract: There have been many initiatives around open source identity and access management, federated authentication standards and products for the web. All of these products and standards require a user store with credential information. Most often this is an LDAP directory. The most common type of credentials still is the combination of a username and a password.

Even though passwords have downsides and many alternatives to passwords exist [QRP], passwords are still here and probably will be for a long time. Passwords are forgotten and lost or expire due to password policies, requiring actions to reset or update passwords. People forgetting or losing their passwords is not just a problem for the people themselves, but also for your organization. Lost passwords result in cost and risk for your organization.

A password management system can help reducing these risks and cost. PWM is a feature rich password management web application for LDAP, written in Java and JavaScript and published under the GNU GPLv2 license.

PWM can help your organization by providing end user self service and supporting help desks in assisting their end users. The product has many features, including those that allow for better verification of the user's identity, enforcing secure password and detect brute force attacks.

The version currently under development has many new and useful features and lots of improvements. The presentation will show a short history of PWM and demonstrate how PWM's rich featureset can help your organization improving password management.

References

[QRP]   Bonneau, J.; Herley, C.; van Oorschot, P.C.; Stajano, F.: The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes, March 2012, http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-817.pdf.

Authentication and security integration for eCampus services at the University of Applied Sciences Harz using the German Electronic Identity Card/eID and eGovernment Standards

Hermann Strack

Fachbereich Automatisierung und Informatik
University of Applied Sciences Harz / Hochschule Harz
Friedrichstr. 57-59, D-38855 Wernigerode, Germany
hstrack@hs-harz.de

Abstract: A eCampus security shell architecture was developed and deployed to improve the security of existing university management systems (legacy UMS), integrating innovative eGovernment Standards e.g. the German Electronic Identity Card (GeID), the eGovernment Protocol OSCI and qualified Signatures (QES).

1 Problem and requirements

The challenge was to improve the security of an existing university management systems (legacy UMS/HIS), by satisfying of particular interoperability requirements (INTOP) and by integrating innovative eGovernment Standards e.g. the German Electronic Identity Card (GeID), the eGovernment Protocol OSCI [www.xoev.de] and qualified Signatures (QES). Especially, these security requirements should be satisfied: privacy and data protection, integrity, (multi factor) authentication. The additional INTOP requirements included particular boundary conditions and restrictions for the security implementations as follows: no changes of existing (legacy) UMS interfaces and GUI; no discrimination of applicants or students without GeID.

2 The eCampus security shell architecture

To achieve the above requirements and conditions, the following eCampus security components must be integrated in an additional security shell for the legacy UMS (as a sort of "security satellite systems"): the eCampus registry to store/check additional security credentials for users (e.g. GeID Pseudonyms, QES certificates, OSCI certificates); the eCampus Server to host additional eCampus secured applications; the eCampus Mediator as a trusted Security Gateway between OSCI based secure communications (incl. signed data) and the legacy http based web interfaces of the

legacy UMS (incl. OSCI client for signed/encrypted data transfer); a U-M-converter service to translate between public and confidential user id attributes in a trusted manner.

To achieve the above requirements and conditions, the following eCampus security components must be integrated in an additional security shell for the legacy UMS (as a sort of "security satellite systems"): the eCampus registry to store/check additional security credentials for users (e.g. GeID Pseudonyms, QES certificates, OSCI certificates); the eCampus Server to host additional eCampus secured applications; the eCampus Mediator as a trusted Security Gateway between OSCI based secure communications (incl. signed data) and the legacy http based web interfaces of the legacy UMS (incl. OSCI client for signed/encrypted data transfer); a U-M-converter service to translate between public and confidential user id attributes in a trusted manner.

Figure 1: An overview - eCampus architecture, eGov. components, examination data flow (Testat)

Figure 2: The eCampus security shell architecture, integrating GeID, OSCI, QES standards

References

[BKM+08] Bender J., Kügler D., Margraf M., Naumann I.: *Sicherheitsmechanismen für kontaktlose Chips im deutschen elektronischen Personalausweis*, DUD, 3/2008

[StBr12] Strack H., Brehm N., et al.: *eCampus – Services & Infrastrukturen für elektronische Campusverwaltung mit verbesserter Sicherheit auf Basis von eGov.-Standards/Komponenten*, eGovernment Review, 2012

[BMI13] BMI eGov. Init. (ed.): *Hochschule Harz - eID-Anwendungskonzept (eTestate)*, http://www.personalausweisportal.de

[EuCo12] European Commission (ed.): *Public Services Online, Centric eGovernment performance in Europe – eGovernment Benchmark 2012*: HS Harz, pp. 47

Cloud-based provisioning of qualified certificates for the German ID card

Marcel Selhorst, Carsten Schwarz

Bundesdruckerei GmbH
Oranienstr. 91, 10969 Berlin
marcel.selhorst@bdr.de, carsten.schwarz@bdr.de

Abstract: In November 2010 the German government introduced a new national ID card. The Bundesdruckerei GmbH was the responsible company for designing and producing the ID card including its highly sophisticated security features. Besides traditional means for visual identification, the card contains a wireless smartcard chip enabling online usage of the ID card. Thus citizens are now able to prove their identity, age or place of residence to an online service provider, e.g., through a web application. Additionally, the chip contains an inactive application for the generation of digital signatures based on elliptic curve cryptography (ECDSA) which - upon activation - can be used to digitally sign electronic documents (online as well as offline).

The Bundesdruckerei GmbH is currently the only party able to perform online post-issuance personalization of qualified electronic signature certificates on the ID card. In order to do so, a new web application called "sign-me"[1] has been developed enabling citizens to activate the signature application on the ID card. In order to diminish the technical challenges for the citizens, "sign-me" takes over the required steps of

- performing the required online identification of the citizen according to the German signature law by using the eID-application provided by the new ID card,

- generating a fresh signature key pair on the ID card,

- exporting the according public key to the certificate service provider "D-TRUST GmbH", the trustcenter of the Bundesdruckerei GmbH, which is then responsible for binding the citizen's identity to the generated signature key pair by issuing the according X.509-certificate, and finally

- storing the issued qualified certificate on the citizen's ID card.

This invited talk briefly introduces the German eID system and focuses on the organizational process as well as the infrastructure required for secure online issuance and management of the certificates. We will introduce the "sign-me" web application and show how citizens can activate the signature application on their ID card, how quickly it is possible to issue and store a qualified certificate on the ID card and how it can be used to finally sign documents. An outlook on envisioned further extensions of "sign-me" concludes the presentation.

---

[1] http://www.sign-me.de

Upcoming specifications from the OpenID Foundation

Henrik Biering[1] · Axel Nennker[2]

[1] Peercraft, Lergravsvej 53, 2300 Copenhagen S, Denmark,
hb@peercraft.com

[2] Telekom Innovation Laboratories, Winterfeldtstr. 21, 10781 Berlin, Germany.
axel.nennker@telekom.de

Abstract: The OpenID Foundation (OIDF), is an international non-profit organization of individuals and companies committed to enabling, promoting and protecting OpenID technologies. Currently OIDF is finalizing the third generation of OpenID Single Sign-On protocols under the brand name "OpenID Connect". In parallel with this effort OIDF has also launched Working Groups for solving other problems that arise when users interact with an ecosystem of interoperable service providers rather than a single service provider.

The presentation will cover the status, features, and benefits of OpenID Connect, Account Chooser, and the Backplane Protocol supplemented by feedback collected from various stakeholder groups.

1 Introduction

Formed in June 2007, the OpenID Foundation ("OIDF") serves as a public trust organization representing the open community of developers, vendors, and users. OIDF assists the community by providing standards and support for internet scale identity management and related technologies. This also entails managing intellectual property and brand marks as well as fostering viral growth and global participation in the proliferation of OpenID.

Currently OIDF is finalizing the third generation of OpenID Single Sign-On protocols under the brand name "OpenID Connect". In parallel with this effort OIDF has also launched the Account Chooser Working Group for solving the usability problems arising when a relying party supports multiple identity providers, and the Backplane Working Group which deals with the problems that arise when users interact with an ecosystem of interoperable service providers rather than a single service provider.

2 OpenID Connect

OpenID Connect is a suite of lightweight specifications that provide a framework for identity interactions via REST like APIs. A specific goal for OpenID Connect has been

to make it as easy as possible for relying parties to implement OpenID. The simplest deployment of OpenID Connect allows for clients of all types including browser-based, mobile, and JavaScript clients, to request and receive information about identities and currently authenticated sessions [Sa13a] [Sa13b]. The specification suite is extensible, allowing participants to optionally support more advanced features and encryption of identity data [Sa13e] [Sa13f], provider and user discovery [Sa13c], dynamic client registration [Sa13d], and advanced session management, including logout [Sa13g].

OpenID Connect performs many of the same tasks as OpenID 2.0, but does so in a way that is API-friendly. OpenID Connect also includes more robust mechanisms for signing and encryption allowing OpenID Connect to be used in scenarios requiring higher levels of assurance. Integration of OAuth 1.0a and OpenID 2.0 required an extension (the OpenID/OAuth hybrid). Being based directly on OAuth 2.0, OAuth capabilities are inherently built into OpenID Connect.

Additionally OpenID Connect supports propagation of both distributed and aggregated claims, and specifies a "self-issued" mode allowing a user to host his/her own Identity Provider while still being able to present trusted third-party claims to service providers.

3 Account Chooser

Account Chooser is a technique to improve the user experience for logging into a website. It produces a uniform and standardized UI to handle the use cases where a device is used by different users, where a single user has more profiles on a particular website, and in particular it solves the "Nascar Problem" [Me09] occurring when a new user wants to sign up at service provider supporting a large number of identity providers.

The Account Chooser will be implemented as a central service operated by OIDF at accountchooser.com, but may also be implemented locally by a service provider. Each method has its own distinct advantages and disadvantages.

The model is protocol agnostic and may in some cases improve usability on a website even if it does not support identity providers, or a website that only supports a single identity provider.

4 Backplane Protocol

Many websites on the Internet embed JavaScript applications into their web pages to provide social functionality such as single sign-on, commenting, sharing, polling, and chatting. As such applications are often developed and hosted by different vendors, they are effectively silos that cannot communicate with each other. This presents a significant problem because the user experience is disjointed and broken, which forces website operators to invest time and money to integrate these services through proprietary APIs.

The Backplane Protocol is a proposed open standard to solve this problem. Backplane Protocol is a secure framework for interaction between multiple, independent client- and server-side parties in the context of a browser session. The Backplane Protocol lets trusted applications share information. When placed together on a web page, Backplane-enabled applications share user identity and other information, seamlessly, regardless of their source. In essence, Backplane Protocol defines a message distribution system where messages are delivered securely, reliably, in order, and in real-time. When a user takes action in one app, the other apps will get the news using the Backplane Protocol.

5 Participation and timeline

The vast majority of OIDF's work is done in the Working Groups. A working group is focused on a specific problem, technology, or opportunity for which the members will deliver a document or series of documents, after which they may disband or create a revised charter for further work. The completion of a working group charter and subsequent disbanding of the group are viewed as a sign of success.

Membership of the Foundation is not required to participate in a working group but participants must agree to the IPR Policy by executing a Contribution Agreement and subscribe to the groups' mailing list. This allows anyone to participate in technology development while ensuring that the specifications remain freely implementable by anyone.

Each working group has one or more editors and a charter that the group is supposed to follow. When a specification is considered complete, an approvals process is initiated. First a review period followed by a vote among the OIDF members is conducted to approve an "Implementer's Draft" version of the specification. When sufficient feedback has been gathered and processed, a second review and vote is conducted to approve the specification as an official OIDF standard.

The OpenID Connect specification is presently entering the implementers draft review period and is expected to enter the final review period by fall 2013.

6 General Feedback

OpenID Connect has technically been designed to work in a variety of environments requiring different levels of security, identity assurance, and privacy. The Account Chooser proposal is expected to facilitate a smoother transition from local login to login via one or more identity providers.

Hence OIDF is currently soliciting feedback from both developer and business communities to determine how the new features of OpenID Connect and Account Chooser can be promoted to overcome the scepticism associated with current alternatives to local login, such as the previous OpenID versions, current government issued ID's, Facebook Connect, and various federation solutions.

References

[Sa13a] Sakimura, N.; Bradley, J.; Jones, M.; Medeiros, B.; Mortimore, C.: *OpenID Connect Basic Client Profile 1.0*, http://openid.net/specs/openid-connect-basic-1_0.html, 2013.

[Sa13b] Sakimura, N.; Bradley, J.; Jones, M.; Medeiros, B.; Mortimore, C.; Jay, E.: *OpenID Connect Implicit Client Profile 1.0*, http://openid.net/specs/openid-connect-implicit-1_0.html, 2013.

[Sa13c] Sakimura, N.; Bradley, J.; Jones, M.; Jay, E: *OpenID Connect Discovery 1.0*, http://openid.net/specs/openid-connect-discovery-1_0.html, 2013.

[Sa13d] Sakimura, N.; Bradley, J.; Jones, M.: *OpenID Connect Dynamic Client Registration 1.0*, http://openid.net/specs/openid-connect-registration-1_0.html, 2013.

[Sa13e] Sakimura, N.; Bradley, J.; Jones, M.; Medeiros, B.; Jay, E.: *OpenID Connect Standard 1.0*, http://openid.net/specs/openid-connect-standard-1_0.html, 2013.

[Sa13f] Sakimura, N.; Bradley, J.; Jones, M.; Medeiros, B.; Mortimore, C.; Jay, E.: *OpenID Connect Messages 1.0*, http://openid.net/specs/openid-connect-messages-1_0.html, 2013.

[Sa13g] Sakimura, N.; Bradley, J.; Jones, M.; Medeiros, B.; Agarwal, N.: *OpenID Connect Session Management 1.0*, http://openid.net/specs/openid-connect-session-1_0.html, 2013.

[Me09]   Messina, C.: *Does OpenID need to be hard?*, http://factoryjoe.com/blog/2009/04/06/does-openid-need-to-be-hard/, 2009

A Novel Set of Measures against Insider Attacks - Sealed Cloud

Hubert A. Jäger, Arnold Monitzer, Ralf O. G. Rieken, Edmund Ernst

Uniscon universal identity control GmbH
Agnes Pockels-Bogen 1, 80992 Munich, Germany
{hubert.jaeger, arnold.monitzer, ralf.rieken, edmund.ernst}@uniscon.de

Abstract: Security and privacy have turned out to be major challenges of the further Internet evolution in general and cloud computing, in particular. This paper proposes a novel approach to safeguard against previously unimpeded insider attacks, referred to as Sealed Cloud. A canonical set of technical measures is described, which, in conjunction, sufficiently complicate and thus economically prevent insider access to unencrypted data. This paper shows the advantages versus end-to-end encryption relative to communication services. Another application of the Sealed Cloud, referred to as Sealed Freeze, provides a seminal solution to privacy issues pertaining to data retention.

1   Introduction

For a long time, IT security concerns have focused on perimeter security, assuming the providers of software as a service (SaaS), clouds and cloud-based services to be trustworthy. However, data theft and privacy violation statistics [KKC+05], [HMKF10] reveal that at least every fourth harmful attack originates from within providing organizations. This data only confirms what many potential customers of SaaS and cloud based offers already sense regarding the data's security. Therefore, mission critical applications are not outsourced to cloud resources, and privacy preserving services have not been established on a significant scale, to date [CHP+09]. In other words, integrated security is absolutely essential as recently postulated by many IT security experts, e.g. [Eck09]. Is data created outside the cloud, then client encryption of this data provides basic security. However, is data to be generated within the cloud, the demand for a technical breakthrough protecting user data processed by providers is imperative.

The present proposal was elaborated within the framework of development of a Web privacy service [JM09], where, in an SaaS architecture, the data security exigence was extended to also consistently embrace the server components. Once this condition precedent was fulfilled, the resulting technical measures proved to equally solve the issue in general computing infrastructure.

Outline The remainder of this article is subdivided as follows. Section 2 gives account of previous work. The Sealed Cloud proposal is presented in Section 3. The advantages of the novel concept for communications and web privacy services as well as data retention technologies is elaborated in Section 4. Finally, Section 5 presents the conclusion.

2 Previous Work

In literature, there are several approaches as to how to secure computing infrastructure by employing Trusted Platform Modules (TPM), e.g. [DTM10] or [SGR09] for improved software integrity. In [GPC⁺03], a closed-box execution environment is used to protect the virtual machines against unauthorized access by an administrator. According to [BMe09], this method has not been implemented, yet.

These approaches secure the software's integrity and thus substantially restrict administrators' liberty to abuse infrastructure and data but do not fundamentally prevent access to unencrypted user data during processing. E.g., if the operation kernel of a processor fails or is provoked to fail, unencryped data is written to core dumps.

Similar ideas to clean up data as the ones presented in this paper, when perimeter security is surpassed, may be found in literature on tamper-proof hardware, e.g. [DBAS04].

The only somewhat comparable alternative to Sealed Cloud known to the authors to date is (fully) homomorphic encryption [Pai99], [SV10] and [Gen08]. However, this enabling technology (still in stage of research) discloses all meta data or connection data (i.e., who communicates with whom, how much and when) to the operators of these services. This is also valid for all end-to-end client encrypting services. Even if "mixing networks" (e.g. [SRG97]) are used to access an infrastructure, which computes on encrypted data, the operator can see which operations are dependent on each other. Thus, these alternatives do not meet secure cloud computing requirements to a sufficient degree.

Hence, in practice, controls as per ISO/IEC 27002, for example, which are integrated into elaborated information security management systems pursuant, e.g., to ISO/IEC 27001, are implemented on an organizational (yet not exclusively technical) level.

The following proposal to technically protect against insider attacks is a set of innovative technical measures yet employs off-the-shelf physical components only. It has been implemented for a concrete Web privacy service, and prototype development for generic use is ongoing.

3 Proposal

A processing infrastructure is assumed, hosting applications that process sensitive, critical or personal data.

Sensitive, critical or personal data is considered any data related to users or subject mat-

ter the users deal with using such applications and deemed worthy of protection against unauthorized access.

Unauthorized access is specified as any access of a party having no business directly related to the business logic of an application nor a legally justified access right.

Unauthorized parties are external attackers but may also be internal service or infrastructure operator staff, provided that no human interaction is needed for an application's business logic. Often, the user of an application and a legal entity are the only persons to have authorized access within the narrow scope of applicable law.

The following proposes a set of technical measures, aimed at protecting sensitive, critical or personal data from unauthorized access. It is essential that said protection of sensitive and mission critical application data be sufficiently effective by technical means only, i.e., it is paramount that potential impact of human deviance be minimized.

Basic Idea Due to the fact that current computing is normally only secured via state-of-the-art perimeter protection and in crucial cases, additionally protected by a comprehensive set of measures insuring software integrity, infrastructure administrators and the administrators of the hosted applications still have access to unencrypted sensitive, critical or personal data.

Of course, operators of such infrastructure and respectively implemented services are well aware of this weakness and tend to complement protection of unencrypted processing data via organizational, non-technical means, i.e., respectively defined processes and staffing with upright personell they deem trustworthy.

A good example of named full set of procedures is described in [MD08] or in the ISO/IEC 2700X standards. The aforementioned elaborates the best combination of technical, formal and informal measures, to maximize security.

In contrast, our proposal replaces this non-technical makeshift by commensurate key distribution and tailored data clean-up procedures.The latter measures, when combined with perimeter security and software integrity, can close contemplated gaps. Thus, with Sealed Cloud, no unencrypted processing data is easily accessible to unauthorized parties.

Key Distribution Let's assume that all data stored on persistent memory is encrypted. In order to avoid that this encrypted data is accessed by the operator of the infrastructure or the operator of the services in unencrypted form, it is necessary to either (a) use an encryption method, in which the operator (once the data is encrypted) is not able, in turn, to decrypt the information, e.g., asymmetric encryption, or (b) delete the encryption key from the processors memory, as soon as encryption is completed. The latter method is appropriate if the encrypted information is to be again used at a later point in time in unencrypted form.

These methods allow distribution of power among the various parties involved in an application's business logic.

The most straightforward use case consists of user data encryption in the database of the

service deployed in Sealed Cloud, with a key provided by the client of the application. If the data is again to be used in the cloud at a later point of time, no asymmetric key is used, and, consequently, the application has to delete the key, once the session or another unit representing the interaction with named data is completed.

A further use case comprises an application, which needs to provide access to specific data for a third party, e.g., when access of a business partner of the client is intentional, to ensure data access needed for partnership with the client organization. Such data can be encrypted in the Sealed Cloud with a business partner's public key, exported in encrypted form to the partner, and, once there, safely decrypted with the partner's private key.

Data Clean-up The database of the Sealed Cloud contains no unecrypted data. Pursuant to business logic, the key to said data is only available for the party owning it. However, unencrypted data is found in the persistent and volatile memory of the processing infrastructure alike. Planned access, i.e. planned maintenance to said memory, is inevitable, if one is to keep processing upright from an operational perspective. Unplanned access cannot be excluded either, since perimeter security can, in most cases, set off an alarm when intrusion is detected but not always prevent it effectively.

Data clean-up, as proposed here, implies that planned or unplanned access to the persistent or volatile memory is not possible until sensitive, critical or personal data has been deleted or reliably overwritten. This requires appropriate trigger signals, indicating to the data clean-up procedure, that planned access is requested, or unplanned access is immanent. Planned access postulates the creation of new trigger signals, whereas unplanned access can rely on perimeter security alarms as signals.

Implementation Figure 1 illustrates a sample implementation of the described set of measures. The cloud user's personal computers or other electronic devices are connected to Sealed Cloud, which is run by the cloud operator. The application software executed in Sealed Cloud was developed and produced by the application operator and has been examined and certified by one or multiple external auditors, before it was deployed in Sealed Cloud. All players' domains of control are indicated in Figure 1 with dashed lines, respectively. The structure of Sealed Cloud in this sample implementation is depicted in Figure 1 within the domain of the cloud operator. It consists of a so-called data clean-up area in the center (emphasized by two boldly printed "sealing" bows at the bottom and the top of the area) and the database and encrypted file system, as well as the peripheral seal and cloud control.

When the user connects to Sealed Cloud, an encrypted communication channel from the browser or any other application running on the user's personal computer or device is established to one of the application servers in the data clean-up area, pursuant to well-known standard procedures, e.g., secure socket layer protocol. The selection of the actual application server is performed by load distributing mechanisms, implemented within the routers and servers of the cloud control unit, which also hosts the state-of-the-art mechanisms for perimeter security, such as firewall and intrusion detection and prevention. It is worthy of mention that the necessary shared secret or private key for this encrypted connection is (for

Figure 1: A sample implementation of the canonical set of measures for a Sealed Cloud infrastructure.

the purposes of the Sealed Cloud) not known to the cloud operator but under the control of the external auditor, who deployed a non-reverse-engineerable software agent on each application server. For practial purposes, this can be approximated by code obfuscation [BGI+01]. Furthermore, each of these agents is individually produced for each respective application server, so that its execution is possible only on the individual server with the server's concrete TPM secrets.

The sensitive, critical or personal data is then processed in unencrypted form in the application server. For persistent storage, the data is encrypted with a key derived from the user's login credentials at the beginning of the session. The application software deletes these login credentials the instant the storage key is generated. External auditors focus on this deletion procedure, in particular. The data is then stored in the database in encrypted form. In the next session, the key which is necessary to read the data back from the database is again generated from the login credentials. At the end of each session, this derived key is also deleted. This procedure is also a main focus of examination through external auditors. The data encryption keys in the mass storage may be stored in the encrypted data, which, in turn, is stored in the database.

Access to the unencrypted data during processing within the data clean-up area is pre-

vented by the data clean-up method. The following illustrates this method as per implementation example in Figure 1: The sealing control unit monitors a comprehensive set of sensors and agents running on all system components, to detect access attempts to the Sealed Cloud infrastructure. In the event the data clean-up area is accessed without authorization, the affected application servers immediately stop operation and delete any unencrypted data. For the purpose of simplification, the data clean-up area of this implementation example contains volatile memory only. The deletion procedure is, e.g., brute forced by power-down of the affected application servers. This applies to both logical and physical attempts to access the data clean-up area. Such reversal of priorities, that privacy is ranked even higher than high availability requirements, lead to such system behavior. In the event of authorized access, e.g. for maintenance purposes, the same data clean-up mechanism is triggered only once access rights (authorization, system availability requirements, et al.) are confirmed by a state-of-the-art access control system.

When starting or restarting the application servers or other components of the Sealed Cloud system, the integrity of the system must be verified. A chain of trust must be established, embracing the full stack, from the server hardware to the highest application software layers, e.g., employing, in this implementation example, the TPMs as roots for the chains of trust.

Organizational Measures and Audit The user must be able to trust the Sealed Cloud operator and the application provider, i.e. that the system behaves as claimed and that both hardware and software in the system are trustworthy and execute only the specified functions. The implementation complexity needs to be limited by hierarchic structuring and encapsulation of the system modules, so that external auditors are able to understand and examine all components and, in particular, the critical focal points of an audit. Only then can external auditors issue certificates, thus providing the user an expert opinion, to justify trust in the operating parties. To further improve the coverage of examination by external auditors, they employ software agents, to dynamically observe system behavior and issue dynamic system integrity attestation for the user. Hence, despite the fact that the technical measures 'key distribution' and 'data clean-up' sufficiently complicate insider access to unencrypted processing data and therefore protect against insider attacks, organizational measures are needed, to secure a proper auditing and certification process by external auditors. That means that human integrity and processes are still important for the operation of the overall Sealed Cloud. However, this set of measures is, as illustrated in Figure 2, shifted to the second line of defence.

Core Principle The core principle underlying present proposal, is to implement a set of appropriate technical measures, to enforce the distribution of power between various parties. Such distribution of power (concerning the access to data), of course, only works, as long as no malicious coalitions are built between the various parties. The probability of such coalitions decreases, the less the external auditors depend on the operators and the more they depend on the users. This stipulates that no monopoly, neither for the operator nor for the auditor, is acceptable.

Figure 2: Organizational measures are shifted to the second line of defence.

Canonical Set of Measures   The presented set of measures is classified as canonical, because the entirety of technical measures, serving the purpose of protecting the unencrypted processing data, can be classified into the presented four categories "perimeter security", "software integrity", "key distribution" and "data-clean-up". Despite the various measures' dependency, each technical measure can be unambiguously categorized into one of the given groups of measures.

4   Applications

As mentioned in Section 1, the Sealed Cloud concept was elaborated, to develop a Web service designed to protect user privacy. The properties and a fundamental privacy advantage of such a service, in particular, compared to end-to-end-encryption, is described as a first application example in this section. The second application example was also developed in this connection. For cases with an obligation to court-ordered disclosure of data, e.g. connection data in telecommunications systems, stipulated the design of Sealed Freeze.

Web Privacy Services   Web privacy services empower the user to enjoy the opportunities of modern networking technology, while pro-actively maintaining user privacy alike. Sealed Cloud is an enabling technology, generating trust in web privacy services. The Web Privacy Service IDGARD [idg13] is the first privacy service to offer Sealed Cloud infras-

Figure 3: Sealed cloud also ensures connection data privacy

tructure. With a proxy function and additional measures as part of the application on the application servers, the source address and other identifying parameters of the user device can be disguised, to allow the user pseudonymous visits of websites. A certain number of users of such a service is necessary, for the individual user to be hidden among numerous fellow users. Further, Sealed Cloud can securely host user names and passwords safely, to provide for convenient and secure online authentication. Finally, shared storage allows a full range of communication services, such as e-mail, chat, file sharing, etc. The latter use case is illustrated in Figure 3. On the left-hand side of the figure, communication is depicted between users A-D via standard communication services. The connection data, i.e., who is connected with whom, when, and how much data is trafficked, is visible to the operator of the standard communication service. In contrast, a Sealed Cloud based communication service, as depicted on the right-hand side of Figure 3, does not disclose any of this connection data to the service operator.

Sealed Freeze   Vis-a-vis legal philosophy, aforementioned web privacy services ultimately ensure free democratic order. However, to prevent these services from degenerating to hiding places for criminals or terrorists, a method for authorized persons to be able to access connection data within a very restricted constitutional framework is imper-

Figure 4: Sealed Freeze based on Sealed Cloud technology: An approach to resolve privacy issues regarding data retention.

ative. Yet, the property that the operators, technically, have no access to this data, has to be held upright. Moreover, the strict rules of the tight constitutional framework of justified access should be enforced, technically.

Figure 4 depicts the basic concept of Sealed Freeze. Any relevant data acquisition and processing system, e.g. telecommunications networks, social networks or video surveillance systems, to name only a few, feature data acquisition devices and a system to transport the data to a storage area. With Sealed Freeze, a key store generates pairs of assymmetric keys, keeps them in volatile memory only, and provides the public key to the data acquisition devices. These encrypt the data to be retained block by block, each with a specific public key, respectively, and then forward the encrypted data to the storage area. In case court-ordered or other authorized persons are legally obliged to access the retained data, they can request the matching private keys from Sealed Freeze. The policy gate in Sealed Freeze will disclose the matching private keys only if the request fulfils the policy rules, as defined by lawmakers in advance and as programmed into the policy gate. The policy cannot be changed with retroactive effect, since all keys are deleted during deployment of a new policy. The policy can contain rules regarding a four-eyes principle, maximum storage duration, volume of disclosure, flows of disclosure within a given period of time, et al. The rule set within the policy can be chosen in a manner that no dragnet investigation is possible, because the number of private keys to be disclosed is limited. Through the rule defining that private keys be deleted after a specific amount of time, deadlines can be enforced, technically. Here, too, Sealed Cloud is the enabling technology that resolves

privacy issues.

5 Conclusion

The present proposal is a good example of an integrated security approach in information technology. By technical means, unauthorized access of any kind is effectually complicated and thus prevented efficiently. Unauthorized parties include the service amd infrastructure operators. The resultant Sealed Cloud therefore constitutes an unrivaled, trustworthy processing infrastructure for clients of hosted applications, as opposed to the user having to rely on the mere trustworthiness of the provider.

Present paper is a proposal, opening a field of research regarding the suggested measures' implementation options. Fields of interest are, in particular, software integrity in environments with virtual engines and approaches to reliable data clean-up in standard cloud application interfaces.

The Sealed Cloud prototype infrastructure is pursued by Uniscon GmbH, Fraunhofer Institute of Applied and Integrated Security, and SecureNet GmbH, and is co-funded by the German Ministry of Economy and Technology within the framework of the so-called Trusted Cloud initiative [BfWuTB10].

References

[BfWuTB10] Deutsches Zentrum für Luft-und Raumfahrt e.V. Projektträger im DLR Bundesministerium für Wirtschaft und Technologie (BMWi), Referat Entwicklung konvergenter IKT. Sichere Internet-Dienste – Sicheres Cloud Computing für Mittelstand und öffentlichen Sektor (Trusted Cloud). *Ein Technologiewettbewerb des Bundesministeriums für Wirtschaft und Technologie, http://www.bmwi.de*, 2010.

[BGI+01] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang. On the (Im)possibility of Obfuscating Programs. In J. Kilian, editor, *Advances in Cryptology – CRYPTO '01, volume 2139 of Lecture Notes in Computer Science*, pages 1–18. Springer, 2001.

[BMe09] G. Brunette, R. Mogull, and editors. Security Guidance for Critical Areas of Focus in Cloud Computing V2.1. *Cloud Security Alliance*, 2009.

[CHP+09] D. Catteddu, G. Hogben, A. Perilli, A. Manieri, A. Algom, J. Rhoton, M. Rohr, O. Biran, and R. Samani. Cloud Computing: Benefits, risks and recommendations for information security. *European Network and Information Security Agency (ENISA)*, 2009.

[DBAS04] Eric D, Bryant, Mikhail J. Atallah, and Martin R. Stytz. A Survey of Anti-Tamper Technologies, 2004.

[DTM10] W. Dawoud, I. Takouna, and C. Meinel. Infrastructure as a service security: Challenges and solutions. In Informatics and Systems (INFOS). *In Informatics and Systems (INFOS), 2010 The 7th International Conference on Informatics and Systems (INFOS)*, page 1 to 8, 2010.

[Eck09]   Claudia Eckert.   ITK-Kompendium 2010.   in: Marlene Neudörffer (Hrsg.), IT-Sicherheit der nächsten Generation – Herausforderungen und Entwicklungen, FAZ-Institut, September 2009.

[Gen08]   Craig Gentry. Computing Arbitrary Functions of Encrypted Data, 2008.

[GPC+03]   T. Garfinkel, B. Pfaff, J. Chow, M. Rosenblum, and D. Boneh. Terra: a virtual machinebased platform for trusted computing. *In Proceedings of the nineteenth ACM symposium on Operating systems principles, SOSP'03*, page 193 to 206, 2003.

[HMKF10]   L. Holmlund, D. Mucisko, K. Kimberland, and J. Freyre. 2010 Cybersecurity watch survey: Cybercrime increasing faster than some company defenses. *Carnegie Mellon University, Software Engineering Institute, CERT Program*, 2010.

[idg13]   www.idgard.de, 2013.

[JM09]   H. A. Jaeger and A. Monitzer. Device for generating a virtual network user. *Patent application WO 2010/084017*, January 22nd 2009.

[KKC+05]   M. Keeney, E. Kowalski, D. Cappelli, A. Moore, T. Shimeall, and S. Rogers. Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors. *Carnegie Mellon University, Software Engineering Institute, CERT Program*, 2005.

[MD08]   S. Mishra and G. Dhillon. Defining Internal Control Objectives for Information Systems Security: A Value Focused Assessment. In W. Golden, T. Acton, K. Conboy, H. van der Heijden, and V. K. Tuunainen, editors, *16th European Conference on Information Systems*, pages 1334–1345, Galway, Ireland, 2008.

[Pai99]   P. Paillier.   Public-Key Cryptosystems Based on Composite Degree Residuosity Classes. In Advances in Cryptology. *EUROCRYPT'99, LNCS, Volume 1592*, page 223 to 238, 1999.

[SGR09]   N. Santos, K. P. Gummadi, and R. Rodrigues. Infrastructure as a service security: Challenges and solutions. In Informatics and Systems (INFOS). *In Proceedings of the 2009 conference on Hot topics in cloud computing, HotCloud'09, Berkeley, CA, USA*, 2009.

[SRG97]   P.F. Syverson, M.G. Reed, and D.M. Goldschlag. Anonymous connections and onion routing. *Proceedings of IEEE Symposium on Security and Privacy, Oakland, CA*, pages 44–54, 1997.

[SV10]   N. P. Smart and F. Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. *In Proceedings of the Conference on Practice and Theory in Public Key Cryptography*, 2010.

GI-Edition Lecture Notes in Informatics

P-64   Peter Liggesmeyer, Klaus Pohl, Michael Goedicke (Hrsg.): Software Engineering 2005

P-65   Gottfried Vossen, Frank Leymann, Peter Lockemann, Wolffried Stucky (Hrsg.): Datenbanksysteme in Business, Technologie und Web

P-66   Jörg M. Haake, Ulrike Lucke, Djamshid Tavangarian (Hrsg.): DeLFI 2005: 3. deutsche e-Learning Fachtagung Informatik

P-67   Armin B. Cremers, Rainer Manthey, Peter Martini, Volker Steinhage (Hrsg.): INFORMATIK 2005 – Informatik LIVE (Band 1)

P-68   Armin B. Cremers, Rainer Manthey, Peter Martini, Volker Steinhage (Hrsg.): INFORMATIK 2005 – Informatik LIVE (Band 2)

P-69   Robert Hirschfeld, Ryszard Kowalcyk, Andreas Polze, Matthias Weske (Hrsg.): NODe 2005, GSEM 2005

P-70   Klaus Turowski, Johannes-Maria Zaha (Hrsg.): Component-oriented Enterprise Application (COAE 2005)

P-71   Andrew Torda, Stefan Kurz, Matthias Rarey (Hrsg.): German Conference on Bioinformatics 2005

P-72   Klaus P. Jantke, Klaus-Peter Fähnrich, Wolfgang S. Wittig (Hrsg.): Marktplatz Internet: Von e-Learning bis e-Payment

P-73   Jan von Knop, Wilhelm Haverkamp, Eike Jessen (Hrsg.): "Heute schon das Morgen sehen"

P-74   Christopher Wolf, Stefan Lucks, Po-Wah Yau (Hrsg.): WEWoRC 2005 – Western European Workshop on Research in Cryptology

P-75   Jörg Desel, Ulrich Frank (Hrsg.): Enterprise Modelling and Information Systems Architecture

P-76   Thomas Kirste, Birgitta König-Riess, Key Pousttchi, Klaus Turowski (Hrsg.): Mobile Informationssysteme – Potentiale, Hindernisse, Einsatz

P-77   Jana Dittmann (Hrsg.): SICHERHEIT 2006

P-78   K.-O. Wenkel, P. Wagner, M. Morgenstern, K. Luzi, P. Eisermann (Hrsg.): Land- und Ernährungswirtschaft im Wandel

P-79   Bettina Biel, Matthias Book, Volker Gruhn (Hrsg.): Softwareengineering 2006

P-80   Mareike Schoop, Christian Huemer, Michael Rebstock, Martin Bichler (Hrsg.): Service-Oriented Electronic Commerce

P-81   Wolfgang Karl, Jürgen Becker, Karl-Erwin Großpietsch, Christian Hochberger, Erik Maehle (Hrsg.): ARCS´06

P-82   Heinrich C. Mayr, Ruth Breu (Hrsg.): Modellierung 2006

P-83   Daniel Huson, Oliver Kohlbacher, Andrei Lupas, Kay Nieselt and Andreas Zell (eds.): German Conference on Bioinformatics

P-84   Dimitris Karagiannis, Heinrich C. Mayr, (Hrsg.): Information Systems Technology and its Applications

P-85   Witold Abramowicz, Heinrich C. Mayr, (Hrsg.): Business Information Systems

P-86   Robert Krimmer (Ed.): Electronic Voting 2006

P-87   Max Mühlhäuser, Guido Rößling, Ralf Steinmetz (Hrsg.): DELFI 2006: 4. e-Learning Fachtagung Informatik

P-88   Robert Hirschfeld, Andreas Polze, Ryszard Kowalczyk (Hrsg.): NODe 2006, GSEM 2006

P-90   Joachim Schelp, Robert Winter, Ulrich Frank, Bodo Rieger, Klaus Turowski (Hrsg.): Integration, Informationslogistik und Architektur

P-91   Henrik Stormer, Andreas Meier, Michael Schumacher (Eds.): European Conference on eHealth 2006

P-92   Fernand Feltz, Benoît Otjacques, Andreas Oberweis, Nicolas Poussing (Eds.): AIM 2006

P-93   Christian Hochberger, Rüdiger Liskowsky (Eds.): INFORMATIK 2006 – Informatik für Menschen, Band 1

P-94   Christian Hochberger, Rüdiger Liskowsky (Eds.): INFORMATIK 2006 – Informatik für Menschen, Band 2

P-95   Matthias Weske, Markus Nüttgens (Eds.): EMISA 2005: Methoden, Konzepte und Technologien für die Entwicklung von dienstbasierten Informationssystemen

P-96   Saartje Brockmans, Jürgen Jung, York Sure (Eds.): Meta-Modelling and Ontologies

P-97   Oliver Göbel, Dirk Schadt, Sandra Frings, Hardo Hase, Detlef Günther, Jens Nedon (Eds.): IT-Incident Mangament & IT-Forensics – IMF 2006

P-208 Ursula Goltz, Marcus Magnor,
Hans-Jürgen Appelrath, Herbert Matthies,
Wolf-Tilo Balke, Lars Wolf (Hrsg.)
INFORMATIK 2012

P-209 Hans Brandt-Pook, André Fleer, Thorsten
Spitta, Malte Wattenberg (Hrsg.)
Nachhaltiges Software Management

P-210 Erhard Plödereder, Peter Dencker,
Herbert Klenk, Hubert B. Keller,
Silke Spitzer (Hrsg.)
Automotive – Safety & Security 2012
Sicherheit und Zuverlässigkeit für
automobile Informationstechnik

P-211 M. Clasen, K. C. Kersebaum, A.
Meyer-Aurich, B. Theuvsen (Hrsg.)
Massendatenmanagement in der
Agrar- und Ernährungswirtschaft
Erhebung - Verarbeitung - Nutzung
Referate der 33. GIL-Jahrestagung
20. – 21. Februar 2013, Potsdam

P-212 Arslan Brömme, Christoph Busch (Eds.)
BIOSIG 2013
Proceedings of the 12th International
Conference of the Biometrics
Special Interest Group
04.–06. September 2013
Darmstadt, Germany

P-213 Stefan Kowalewski,
Bernhard Rumpe (Hrsg.)
Software Engineering 2013
Fachtagung des GI-Fachbereichs
Softwaretechnik

P-214 Volker Markl, Gunter Saake, Kai-Uwe
Sattler, Gregor Hackenbroich, Bernhard Mit
schang, Theo Härder, Veit Köppen (Hrsg.)
Datenbanksysteme für Business,
Technologie und Web (BTW) 2013
13. – 15. März 2013, Magdeburg

P-215 Stefan Wagner, Horst Lichter (Hrsg.)
Software Engineering 2013
Workshopband
(inkl. Doktorandensymposium)
26. Februar – 1. März 2013, Aachen

P-216 Gunter Saake, Andreas Henrich,
Wolfgang Lehner, Thomas Neumann,
Veit Köppen (Hrsg.)
Datenbanksysteme für Business,
Technologie und Web (BTW) 2013 –
Workshopband
11. – 12. März 2013, Magdeburg

P-217 Paul Müller, Bernhard Neumair, Helmut
Reiser, Gabi Dreo Rodosek (Hrsg.)
6. DFN-Forum Kommunikations-
technologien
Beiträge der Fachtagung
03.–04. Juni 2013, Erlangen

P-218 Andreas Breiter, Christoph Rensing (Hrsg.)
DeLFI 2013: Die 11 e-Learning
Fachtagung Informatik der Gesellschaft
für Informatik e.V. (GI)
8. – 11. September 2013, Bremen

P-221 Maria A. Wimmer, Marijn Janssen,
Ann Macintosh, Hans Jochen Scholl,
Efthimios Tambouris (Eds.)
Electronic Government and
Electronic Participation
Joint Proceedings of Ongoing Research of
IFIP EGOV and IFIP ePart 2013
16. – 19. September 2013, Koblenz

P-222 Reinhard Jung, Manfred Reichert (Eds.)
Enterprise Modelling
and Information Systems Architectures
(EMISA 2013)
St. Gallen, Switzerland
September 5. – 6. 2013

P-223 Detlef Hühnlein, Heiko Roßnagel (Hrsg.)
Open Identity Summit 2013
10. – 11. September 2013
Kloster Banz, Germany