User Centricity in Healthcare Infrastructures

Matthias Quasthoff, Christoph Meinel

Internet Technologies and Systems
Hasso Plattner Institute
PF 900460,
14440 Potsdam, Germany
matthias.quasthoff@hpi.uni-potsdam.de
meinel@hpi.uni-potsdam.de

Abstract: The introduction of electronic national identity management solutions is accompanied by cheers and bravos from the contributors, but also by doubts and worries from data protection experts. Helping at clearing the view on the actual features and properties of the planned German electronic healthcare infrastructure, we have mapped the parts of the healthcare infrastructure to a taxonomy on user centricity in identity management. That mapping opens perspectives in two directions. On the one hand, a clear analysis of the security of identity data can be reached. On the other hand, future applications building upon the electronic healthcare infrastructure will benefit from that mapping and the abstraction layer it suggests. We show that an additional layer of abstraction will bring huge and complex identity management systems forward to a better understanding of their opportunities and threats and a fairer discussion leading to a better user acceptance.

1 Introduction

Governments introducing electronic identity cards often have to deal with rigid regulations on data protection. Healthcare data in particular receive special protection, e.g. in Germany. That special protection led to the design of an electronic healthcare infrastructure, granting patients the ownership on their healthcare data. From a formal point of view healthcare data is just a form of identity-related data.

For the planned German electronic healthcare infrastructure, detailed specifications and proposals are already available ([5], [12] etc.) One of the purposes of the German electronic healthcare infrastructure including the identity cards is to act as an enabler for future applications building upon the infrastructure [10]. Those applications could use the security protocols and products in place. Yet, to meet all legal requirements and allow for consistent implementation, the specification documents provide a very technical view on the architecture.

For rapid development of future applications it would be thoroughly helpful to build upon an abstract identity management framework. That abstract framework would help to decide which parts of the healthcare infrastructure to use or not use. As identity management plays a major role in any electronic healthcare infrastructure, we chose to abstract from detailed requirements and specifications to the German telematics solution and relate this solution to the user centric identity management paradigm [2].

1.1 User Centricity in Healthcare Identity Management

User centricity is a recent paradigm for various aspects of IT systems. The term addresses problems arising from certain features of established computer systems that have become too complex, too unintuitive or too far away from the actual user needs. Typical aspects of IT systems that can be user centric are e.g. software development processes, user interface design, and besides some others, identity management. User centric or contextual [1] approaches oppose “design centric” approaches, where users or customers are much less or not at all involved in design or development decisions.

IT systems designed in a user centric manner have higher chances on good usability and high user acceptance. User centricity in identity management systems does not only refer to design processes leading to better usability, customer satisfaction or something similar. It is rather focused on processes and information in the actual production systems. Finding user centricity in federated identity management (FIM) a fuzzy term, Bhargav-Spantzel et. al. [2] created a taxonomy for user centricity in identity management. Their understanding of user centric identity management systems that should provide stronger user control and privacy fits well to the upcoming national digital identity management systems and particularly those in the healthcare sector.

As patients are the actual users of an electronic healthcare infrastructure, these systems will at least partially be user-centric. One of the main reasons for the introduction of these electronic healthcare infrastructures is the wish for improvement of healthcare services and the reduction of costs through streamlined processes and efficient provision of healthcare data to health professionals. On this side of the system, health professionals like medical practitioners and pharmacists are also users of the system. These streamlining and efficiency goals are typical goals of user centric engineering. Thus, it makes perfect sense aiming at user centric healthcare infrastructures, keeping in mind that there will be at least two classes of users—patients, different types of health professionals, and maybe also administrative personnel within the health insurance funds.

1.2 Contribution

This article aims at providing a foundation for the consistent development of future application based on existing or arising healthcare infrastructures. Thus, in Section 3 we will map the user centricity taxonomy given in [2] to the objectives in electronic healthcare infrastructures, obeying the different classes of users with their respective objectives. We will also show which parts of the infrastructure need an additional layer of abstraction to allow for a comprehensive and efficient analysis.

2 Electronic Healthcare Infrastructure

Typically, the modernization of healthcare systems involves the transition from paper-based to electronic processes. Two large parts of that transition are the introduction of both a widespread digital identity management and a comprehensive document management solution. The identity management would follow the national regulations on topics like qualified electronic signatures [11]. The document management systems would handle healthcare-related data like patient records, prescriptions and medical images. As all of those healthcare documents contain identity information, the document management system must be part of any work on identity management in healthcare environments.

Administrative documents and prescriptions are the typical healthcare documents that patients are involved with in paper-based environments. Additionally to them, in the German healthcare infrastructure patients will also “own” medical records they would not be aware of in a paper-based environment. To find a solution that meets legal and ethical constraints on the one hand and does not inhibit efficient processes in the healthcare system on the other, it is important to know what objectives are essential in an electronic healthcare infrastructure.

2.1 Healthcare Applications

Identity Management System. As many processes in healthcare directly relate to identity information from patients or health professionals an identity management system is a mission critical part of such a healthcare environment. As healthcare data is often protected by national regulations, the identity management system needs to meet those restrictions as well. For the different participants in the healthcare system, there are different requirements. For example, patients typically act as customers of both health professionals and their health insurance funds.

Usually legislation puts more responsibility on the health professional side leading to more complex processes regarding e.g. signatures on their side compared to the patient side. As a starting point, it seems reasonable to clearly separate the identity management for the different types or roles of participants. In the long run one would expect those different types of identities to converge, i.e. putting the patients' health insurance information on some generic identity card and putting the health professionals' authorization on their identity card. In the case of separation by role, it may well be that for each role different properties of the user centricity taxonomy are fulfilled.

Document Management System. Bringing electronic processes to the healthcare systems, many different types of documents need to be managed in an IT infrastructure. Those documents cover *electronic prescriptions*, being issued by medical practitioners for patients and being invalidated by pharmacists, *medical data* of different kinds such as diagnoses, medical images or ECG data forming the patient's personal health record, and *administrative documents* like referrals or billing information. Along with the likely separation of the different roles in the identity management system, there are clear scenarios attached to each of the document types, each with different requirements to the identity management system. E.g. in many cases the anonymity of a patient is more important compared to the anonymity of the patient's health professional.

Services, Service Bus and Telematics Infrastructure. All the objectives aforementioned must be realized in a context providing directories, storage capabilities, connectivity, healthcare-related functions and so on. Instead of just mapping the user centricity taxonomy onto the identity management system, it is also necessary to apply the taxonomy to the general infrastructure with its underlying service bus, especially when design decisions are bound to a weak trust model.

2.2 User Centricity Taxonomy

As a basis for the user centricity taxonomy, two trust models are defined. A *Weak Trust Model* [2], in which the identity provider is trusted, i.e. it is expected to behave according to the protocols and friendly. For healthcare applications it is important that this model expects the identity provider also to not try gain advantage from the information it controls from the different parties. On the contrary there is the *Strong Trust Model*. In this model, applications and protocols must be designed allowing the identity provider being under control of the attacker. Both models do not make assumptions on friendly behaviour of services or users.

The choice of trust model is also influenced by another design choice: whether identity tokens carried in messages are *user-generated* or *issuer-generated*. Issuer-generated tokens do not allow all privacy properties in strong trust model environments [2].

2.2.1 Identity Management System Properties

As first of four identity management system properties, Bhargav-Spantzel et. al. mention the *user-chosen identity provider* property. A system fulfilling this property must allow each user to choose which identity provider to trust in and register with. The *policy specification and enforcement* property requires exact policy management that other properties can rely upon. *Assurance support* is a property related to policy enforcement allowing users to react upon e.g. policy infringements and thus adjust their trust in the system. *Auditing*, being the last property, needs to be designed such that it does not tamper the other properties.

2.2.2 Transaction Properties

To guarantee transactions being *context bound* it is necessary to not by mistake accept transaction messages in any other transaction. Transactions are *unlinkable* when they do not expose information about their end points suitable for linking like, e.g. a user id. *User consent* is fulfilled if the user is informed about the actual effect of a transaction and agrees to the execution.

2.2.3 Identity Information Properties

Identity information is often required to be *confidential*, i.e. no unauthorized person should have access to that information even if it is released to another person. *Integrity* is another property providing assurance the identity information has not been altered on the way from the identity provider to the verifier. *Verifiability* is related to integrity and user consent allowing the user to see what information the identity provider releases. That property leads directly to *selective release* and *conditional release*. While selective release would enable the user to release only desired amounts of identity information, conditional release would guarantee that certain identity information would not be released before some conditions are fulfilled.

Stealing protection and *sharing prevention* are both about misuse of other users' identity information, but from different perspectives. Stealing protection would protect a user's identity information from theft by e.g. hackers or viruses. On the contrary, sharing prevention would prevent users from giving away their credentials. Sharing is quite usual in today's IT environments because users tend to prefer short-run convenience. However, sharing is in the way of other high-level identity management properties.

Revocation is needed when identities are redescribed or destroyed. Depending on the type of identity management system and validity period of identity tokens, revocation can be more or less difficult to manage and propagate. The last property on identity information Bhargav-Spantzel et. al. mention is *portability*, meaning that users could use their identity information in different places with different devices.

2.2.4 Composite Properties

Composite properties at least partially depend on the aforementioned properties and can be seen as more high-level or abstract. The *attribute security* property assuring all-time correctness of user attributes depends on sharing prevention, confidentiality, verifiability, revocation, integrity, and stealing protection. On the user side, *attribute privacy* would enable them to keep control over their attributes building upon e.g. confidentiality, policies or user-chosen identity provider.

As another property, attribute security and policy specification and enforcement enable for *service protection*. *Non-repudiation*, building on stealing protection and also related to sharing prevention, is fulfilled when transactions can be linked to the entities involved. To be in line with unlinkability, the property is also related to conditional release.

Accountability is a more abstract property building upon non-repudiation but being also related to policy specification and enforcement. Accountability partially conflicts with the *anonymity* property. Anonymity can be reduced to *conditional anonymity*, thus allowing for accountability when a well-defined condition is met. Anonymity is a special case of *data minimization*, the higher-level goal behind selective release and unlinkability. Interestingly, data minimization not only puts more control over identity data on the user side, but also reduces identity providers' costs for regulatory compliance [2].

Policy specification and enforcement forms the basis for the *privacy policy, obligations, and restrictions* property. This is needed for embedding the identity management system in actual business environments. The *user in the middle* property states that either the users' computer systems or the users themselves have to participate in transactions. Systems fulfil the *notification* property if the user is informed about transactions including their identity information. Notification would build open auditing mechanisms and possible policies.

3 User Centricity in German Healthcare Applications

In this section we map the models and properties from the previous section to the healthcare infrastructure. As parts of the specification for the German electronic healthcare infrastructure are quite extensive, and as said before, very technical or not yet complete, we cannot decide for all properties in detail to what extent they will be fulfilled in the infrastructure. Rather, we indicate in what part of the specification further examination needs to be done.

Following the user centricity taxonomy, the Gematik documents [12] specify an architecture following the *Weak Trust Model* [2] meaning the identity provider is being trusted. If one considers the health records a part of the actual identity information, the architecture follows at least partially the *Strong Trust Model*. For legal reasons, some tasks like the storage of health records has to happen in a way that no unauthorized personnel like system administrators could have access to parts of or full health records.

Depending on the level of abstraction, the German electronic healthcare infrastructure will contain *issuer-generated* and *user-generated* tokens. The certificates on the cards—both card verifiable certificates and the certificates for the qualified electronic signature—will be issuer-generated. In the transactions, also tokens generated on the users' cards will be used.

3.1 Identity Management System Properties

The *choice of identity provider* is partially on the users' side. Patients will have one or two identities. One mandatory identity is that of a health insurant. This identity is provided by the respective health insurance fund. Thus, for the patients this choice is not fully free in the sense of the taxonomy, but is influenced by insurance rate, service etc. A second, optional identity can be provided by one of the accredited certification authorities, according to the German legislation.

However, there are rigid regulations in place for the accredited providers of certification services. The certificates are required to contain enough information to fully identify individuals. Thus, the choice of identity provider would not allow for much difference in whether they fulfil different properties from the user centricity taxonomy. A further investigation on the user centricity properties of the German qualified electronic signature is needed. The health professional's identity providers are defined by their different professional organisations for pharmacists, dentists, medical practitioners etc. Hence, their freedom of choice is similarly limited as the patients' choice.

As said before, the specifications by Gematik are very detailed and so are huge parts of the legislation concerning that healthcare infrastructure. Hence, these are the parts where *policy specification and enforcement* do happen: the policies are fixed mostly in the regulations and along with the enforcement in the technical specifications of the architecture that has been designed so far. Also, the identity cards will have support for *auditing* as all users will be able to retrieve complete transaction logs with the help of their digital identity cards. Parts of the infrastructure in design only deal with mechanisms like many separate virtual private networks and PKIs [11] preventing anything from behaving against the detailed specifications. Those technologies are expected to allow for *assurance support*.

3.2 Transaction Properties

On the technical level, many of the identity cards' transactions belong to well-established cryptographic protocols that have a clear mapping to security objectives like confidentiality or integrity and also identity management properties like context bound transactions. As *context bound transactions* are necessary for some other properties discussed in later sections and the property is well known with regards to cryptographic protocols, it is to be expected that this property holds in the infrastructure. Yet, formulating a complete layer of abstraction for identity management on top of the detailed specification would help finally clarifying that question.

Unlinkability is also at least partially covered by standard cryptographic protocols implemented on the identity cards. As not all parts of the final infrastructure are in place, a final examination of this property is not possible. Additionally, for some transactions involving health professionals' identity information, unlinkability is not desired. For the patients unlinkability is definitely required, at least during access to their personal health records and for the creation of statistical information.

Another thing thoroughly discussed during the legislation process was *user consent*. The specification clearly states when patients and health professionals need to present their identity cards, enter PINs or use similar means to declare their intention.

3.3 Identity Information Properties

Again, the well-known properties *confidentiality*, *integrity* and *verifiability* are handled with mature cryptographic protocols. The definition of the abstract layer and the deeper analysis on the certification authorities, that also work using asymmetric cryptography, would bring final results on these properties. Yet, all needs for these properties can already be found in the architecture. Identity information is *stealing protected*, as in most scenarios, two factor authentication [11] with card presentation and PIN entry is required. As some parts of the infrastructure like the one storing the patients' health records are not fully specified, examination of those has to happen at a later point in time. The identity information is also partially *sharing protected*, as private keys are bound to the identity cards, which have a photograph printed on top. If participants follow the protocols, the risk caused by sharing is reduced, albeit not eliminated. Also, as most transactions involve card-based interaction, pooling of credentials by collecting multiple cards is highly unlikely.

Revocation is not an issue for most scenarios of the healthcare infrastructure. Most services are expected to communicate with a centralized directory service allowing for card revocation. However, the requirements to the infrastructure expect it to work during loss of connection to the Internet. During these for the time being rare events, revocation will definitely be broken.

During the discussion of the accredited certification authorities, it was suggested that digital certificates according to the regulations couldn't fulfil all possible properties. The *selective release* of identity information is limited by the granularity of signed entities [8], [7]. So, the identifying information from both patients and health professionals can just be released in a coarse-grained manner. But from the design proposals for the infrastructure storing patients' records [4], it seems that those records can be released on a per-entity basis. Also, for some scenarios *conditional release* is defined by the specification, allowing access to certain information stored on the eGK just after presentation of HBA or even PIN entry.

3.4 Composite Properties

Most of the user centricity composite properties map to design objectives in the Gematik documents. *Attribute security* is essential for the orchestration of an electronic healthcare infrastructure. Yet, attribute security is an abstract property. Its parts could be analyzed at hand of the technical specification. The complete property could only be looked at after additional layers of abstraction had been introduced to the identity management of the proposed infrastructure.

Service protection, *accountability* and *non-repudiation* have not been discussed in an abstract form so far, as not all parts are fully specified. On the other hand, these properties are part of the regulations on accredited certification authorities [11]. An abstract identity management layer comprising the electronic healthcare infrastructure would help analyzing that area. *Data minimization* and *attribute privacy* are objectives of data protection experts. While these properties were considered during the design of the infrastructure, the goal has not been fully reached, e.g. due to the inflexible nature of fully identifying certificates or concerns regarding service protection and accountability.

A *privacy policy* is not fully specified. The difficulties arise, as nobody knows yet what identity information will be available for release in investigations on e.g. insurance fraud. It is also unclear if the infrastructure will allow for an electronic equivalent to house searches. The *notification* property will be fulfilled by the auditing facilities that will be log-based. *Anonymity* is needed for statistical records, which are not yet fully specified. There will be a pseudonym certificate on the health insurance cards, but it will be issued by the health insurance funds. Thus it will be part of a weak trust model, as the patients would need trust the privacy statement of the insurer.

The last property, *user in the middle*, is fulfilled, as the eGK and HBA are needed for all transactions involving the users' identities. As cryptography happens directly on the cards, this property will be hard to break.

4 Conclusion and Future Work

We found the user centricity taxonomy by Bhargav-Spantzel et. al. providing a solid foundation for the formulation of abstraction layers on existing identity management systems. Such layers are needed for both the planned German electronic healthcare infrastructure and the qualified electronic signatures regulations. E.g. SWOT analyzing these abstract models according to the taxonomy will give hints on the security, usability and scalability of the underlying identity management systems.

There already exist models for PKI environments that can serve as a basis for the model of the accredited certification authorities. We will create a new model for the healthcare infrastructure and other electronic national identity management projects. The taxonomy mapped on the healthcare infrastructure will allow for an easy abstraction. It would then be interesting mapping the abstract model to other existing identity management architectures e.g. like that specified in the WS-Trust standard [9].

A sound abstraction model for the German electronic healthcare infrastructure would further enable application developers to build their identity management on top of that layer instead of copying the mechanisms suggested by the healthcare infrastructure. As next step after the definition of the abstract layer, a guide on how to achieve the desired properties in the healthcare infrastructure needs to be created.

The guide could be followed by a reference implementation to be used by the future application. Separating application logic and identity management would thoroughly ease security analysis on new applications. Clearer analysis will then result in more secure applications with less vulnerabilities and a higher user acceptance.

References

[1] Hugh Beyer, Karen Holtzblatt: Contextual Design: Defining Customer-Centered Systems, Morgan Kaufmann, 1998

[2] Abhilasha Bhargav-Spantzel, Jan Camenisch, Thomas Gross, Dieter Sommer: User Centricity: A Taxonomy and Open Issues. DIM'06, November 3, 2006, Alexandria, Virginia, USA

[3] Jörg Caumanns: Der Patient bleibt Herr seiner Daten, in Informatik-Spektrum Volume 29 No. 5, pp323–331. October 2006

[4] Jörg Caumanns, Herbert Weber, Arne Fellien, Holger Kurrek, Oliver Boehm, Jan Neuhaus, Jörg Kunsmann, Bruno Struif: Die eGK-Lösungsarchitektur, in Informatik-Spektrum Volume 29 No. 5, pp341–348. October 2006

[5] Die Spezifikation der elektronischen Gesundheitskarte. Teil 2: Anwendungen und anwendungsspezifische Strukturen, Version 1.2.1. gematik, 7.9.2006

[6] Die Spezifikation der elektronischen Gesundheitskarte. Teil 3: Äußere Gestaltung, Version 1.2.1. gematik, 7.9.2006

[7] Fachkonzept Versicherungsstammdatenmanagement (VSDM), Version 1.1.0. gematik, 19.9.2006

[8] Festlegungen zu den X.509 Zertifikaten der Versicherten, Version 1.2.0. gematik, 2.10.2006

[9] C. Kaler, A. Nadalin, et al.: Web Services Trust Language (WS-Trust) Version 1.1, May 2004. At http://msdn.microsoft.com/ws/2004/04/ws-trust/.

[10] Jan Neuhaus, Wolfgang Deiters, Markus Wiedeler: Mehrwertdienste im Umfeld der elektronischen Gesundheitskarte, in Informatik-Spektrum Volume 29 No. 5, pp332–340. October 2006

[11] Christoph Meinel, Matthias Quasthoff: Identity Management in Telemedicine – The German Health Insurance Card. In Proc. XV Wintercourse of the CATAI, La Laguna, Tenerife, Spain, 2007, pp. 50–57

[12] Bruno Struif (Ed.): German Health Professional Card and Security Module Card. Part 2: HPC Applications and Functions. Version 2.1.0, 21.2.2006

[13] Bruno Struif (Ed.): German Health Professional Card and Security Module Card. Part 3: SMC Applications and Functions. Version 2.1.0, 21.2.2006