Opportunities and Challenges of CREDENTIAL

Towards a Metadata-Privacy Respecting Identity Provider

Farzaneh Karegar[1], Christoph Striecks[2], Stephan Krenn[2], Felix Hörandner[3], Thomas Lorünser[2], and Simone Fischer-Hübner[1⋆]

[1] Karlstad University, Karlstad, Sweden
{firstname.lastname}@kau.se
[2] AIT Austrian Institute of Technology GmbH, Vienna, Austria
{firstname.lastname}@ait.ac.at
[3] Graz University of Technology, Graz, Austria
felix.hoerandner@iaik.tugraz.at

Abstract. This paper summarizes the results of a workshop at the IFIP Summer School 2016 introducing the EU Horizon 2020 project CREDENTIAL, i.e., Secure Cloud Identity Wallet. The contribution of this document is three-fold. First, it gives an overview of the CREDENTIAL project, its use-cases, and core technologies. Second, it explains the challenges of the project's approach and summarizes the results of the parallel focus groups that were held during the workshop. Third, it focuses on a specific challenge—the protection of metadata in centralized identity providers—and suggests a potential architecture addressing this problem.

Keywords: Metadata privacy ◇ Identity provisioning ◇ Data sharing

1 Introduction

With increasing mobility and Internet use, the demand for digital services has increased and already reached critical and high assurance domains like eGovernment, eHealth, and eBusiness. Those domains have particularly high security and privacy requirements, and services are harnessed with various novel mechanisms for securing access. Handling all the different authentication and authorization mechanisms requires user friendly support, which can efficiently be provided by digital identity management (IdM). Due to business mergers and acquisitions as well as the increasing number of cloud applications, IdM is currently experiencing a paradigm shift, moving away from company-internal custom-tailored IdM system towards non-standard fragmented authentication situations. However, under the given change, many current solutions fall short with respect to security, privacy, or usability. Therefore there exists a strong demand to delegate the management of multiple credentials, as well as traditional corporate

⋆ This work is being performed within the H2020 EU project CREDENTIAL under grant agreement number 653454. The ordering of the authors was determined using a fair dice.

identity and access management (IAM) functions, like single sign-on (SSO), to a cloud-based service.

The transformation in the IdM world goes hand in hand with the tremendous shift to cloud computing that has shaped the information and communications technology (ICT) world during the last years. IdM has not remained unaffected in this respect. By now, numerous IdM systems and solutions are available as cloud services, providing identity services to applications operated both in closed domains and in the public cloud. This service model is referred to as Identity (and Access) Management as a Service (IDMaaS). Popular examples for cloud IDMaaS providers are big companies from the sectors of social networks, search engines, business solutions, or online retailers. They offer their user identity base for authentication and identification at various services. However, for increased usability, identity services should cover more than login and authentication. For instance, they could also serve as online password vaults, replacing local password managers for providing better portability and anytime-anywhere access to protected resources. Finally, more general online vaults retaining entire identity documents or personal files and records (e.g., OneDrive, Dropbox, tresorit) can also be considered as identity services. However, currently no satisfactory approaches allowing for the privacy-preserving storage and advanced sharing of identity data by cloud service providers exist.

The vision of CREDENTIAL is to fill this gap, and to develop a privacy-preserving solution for data sharing and identity provisioning. Users will be able to store identity data and other sensitive data such as health records in a cloud-based CREDENTIAL wallet such that confidentiality and privacy are upheld. In particular, the wallet provider will not be able to access the users' personal data, and can build its business strategy around this advantageous security property. If a user wants to share specific data with other users, or share identity information with a service provider in order to log on to a system, she will be guaranteed that after transmission the intended receiver of the data will be the only party capable of accessing the data items in plain text.

At the IFIP Summer School 2016 in Karlstad, Sweden, the CREDENTIAL project organized a workshop to present the project, raise awareness of the existing technologies and solutions, and to receive feedback and input from experts from different domains. To do so, the project's ambition, the used core technologies, and a representative use case from the eHealth pilot were presented to the audience. This was then followed by three focus groups to discuss different aspects and challenges of the CREDENTIAL approach. This paper summarizes the content and results of this workshop as well as subsequent findings that were inspired by those discussions.

1.1 Outline

This document is structured as follows. In Section 2 we give a short overview of the CREDENTIAL project and introduce the pilots and underlying core technologies. Section 3 then describes the challenges discussed during the workshop,

as well as the inputs and recommendations received from the participants. In Section 4, a special focus is put on privacy-related challenges introduced by a central identity wallet. We recap existing countermeasures to those issues in Section 5 and explain their shortcomings, before we describe a potential high-level architecture solving those problems in Section 6. Finally, we briefly conclude in Section 7.

2 The CREDENTIAL Project

The overall vision of CREDENTIAL is to develop a user-centric cloud-based data storage and sharing platform, which enhances the user's privacy compared to current approaches and keeps the user in control, while retaining the benefits of cloud-based solutions. In order to achieve this, CREDENTIAL will employ advanced cryptographic mechanisms, such as proxy re-encryption [1] and redactable signatures [2]. The developed solution will follow state-of-the-art security and privacy by design principles. By using and extending well-established standards and protocols, we aim to not only apply the CREDENTIAL approach to a comprehensive cloud system but to also facilitate integration into existing solutions. In the following, we will first explain CREDENTIAL's basic technologies and architecture, and then highlight its application to three different domains, namely eGovernment, eHealth, and eBusiness.

2.1 Basic Technologies

CREDENTIAL uses the following two cryptographic mechanisms as a foundation to enable confidentiality, integrity, and authenticity from end-to-end during the sharing process of whole messages or subsets.

Proxy re-encryption, introduced by Blaze et al. [1], extends asymmetric encryption with the ability to transform a ciphertext c_A encrypted for party A into another ciphertext c_B of party B without revealing the underlying plaintext in an intermediate step. To enable this transformation, party A generates a re-encryption key $rk_{A \to B}$ from her private key sk_A and the public key pk_B of party B. As neither plain text nor decryption keys are exposed during re-encryption, this operation can be outsourced to a semi-trusted proxy. The technology is therefore well suited for end-to-end encrypted data sharing.

Redactable signatures, introduced by Johnson et al. [2], make it possible to black-out parts of a signed message and still verify the signature on the remaining parts; this is in contrast to plain digital signatures, where every bit flip in the message invalidates the signature. This redaction can be performed without access to the signer's private key. The technology is therefore well suited for realizing selective disclosure.

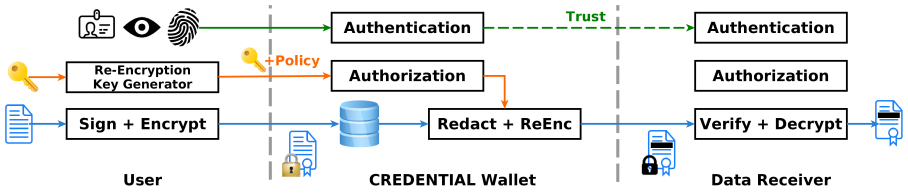

Fig. 1. CREDENTIAL's Basic Architecture

2.2 Architecture

CREDENTIAL's basic architecture integrates the above presented cryptographic mechanisms into three key actors: user, CREDENTIAL wallet, and data receiver. After outlining these actors, their interactions are described.

The *user* owns data that should be securely stored or shared with other participants. A client application is deployed in the user's domain to handle operations involving user's private key material, such as signing or generating a re-encryption key. This application should not be accessible or online when the intended receiver wants to access the data.

The *wallet* represents the central component of CREDENTIAL. This wallet is a data storage and sharing service deployed in the cloud yielding among others benefits such as constant availability, scalability, and cost effectiveness. A powerful identity and access management system performs multi-factor authentication and authorizes access to the stored data. With proxy re-encryption, the confidentiality of the stored and shared data is ensured even when deploying the wallet at an honest-but-curious cloud provider, as no plain data is exposed. Furthermore, once a re-encryption key is provided, the data can be shared with other participants even when the user or her client application are not available.

The *data receiver* might be a service provider or another CREDENTIAL user. It relies on data stored in or authentication assertions issued by the wallet. With this information, the data receiver reaches authorization decisions and performs arbitrary data processing.

A simple *data sharing process* highlighting the interaction among the individual components is shown in Figure 1. First, the user authenticates to the wallet to get permission to upload signed and encrypted data c_U. This data c_U is encrypted for the user herself to retain maximum control. To share encrypted data, the user generates a re-encryption key $rk_{U \to DR}$ towards a selected data receiver DR. This generation operation has to be performed in the user's domain, as the user's private key is involved. Along with this re-encryption key, a policy defining which data may be disclosed is transmitted to the wallet. Upon request of a data receiver, not-required data is redacted based on the policy. Then, using the re-encryption key $rk_{U \to DR}$, the user's redacted ciphertext c'_U is transformed into data encrypted for the data receiver c'_{DR}. Finally, the data receiver is able to decrypt the data c'_{DR} and verify the signature on the disclosed parts.

2.3 Pilot Scenarios

The CREDENTIAL technologies and architecture are showcased by pilot scenarios from three different domains. A more detailed description of the scenarios can be found in Hörandner et al. [3].

eGovernment. In the eGovernment pilot, the focus lies on identity management to authenticate citizens and assess their eligibility for a service, based on – often sensitive – identity attributes. This identity management is considered an instance of CREDENTIAL's data sharing process via standardized identity protocols such as SAML [4] or OpenID Connect [5]. In such a protocol, the service provider (i.e., the data receiver) triggers the process by requesting user authentication and identity attributes from the identity provider (i.e., the CREDENTIAL wallet). Concerning authentication, we will not only integrate national eID solutions, but also look at a broader context and enable cross-border authentication according to the eIDAS regulation [6]. The wallet prompts the user for consent as well as a re-encryption key, and then selectively discloses re-encrypted attributes to the service provider.

eHealth. The eHealth pilot is concerned with a data sharing between patients, doctors, and further parties, in particular in the context of Type 2 Diabetes. Namely, the developed components will allow patients to record their health data (blood sugar level, weight, blood pressure, etc.) using external mobile devices. The data measured on these devices will be collected by a CREDENTIAL eHealth mobile app, which remotely stores this data in the CREDENTIAL wallet. The user can then define who is allowed to access which parts of this medical data, to share specific parts of the measurements, e.g., with the family doctor, diabetologist, nutritionist, or personal trainer. Based on the data they see, they can then provide recommendations back to the patient. This remote data sharing functionality between patients and doctors is of key importance also in other telemedicine applications, as the patients' privacy is respected while the doctor still obtains high authenticity guarantees for the received data.

eBusiness. The eBusiness pilot showcases how easy the privacy offered by existing solutions can be enhanced through the integration of modular libraries implementing CREDENTIAL's technologies. Encrypted mails are a requirement for many companies to protect their data and inventions, but they also represent a significant challenge when employees are temporary unavailable, e.g., because of vacation. Currently, employees have to expose their private key material so that a substitute member of staff can still read and answer incoming mail. In contrast, with proxy re-encryption, an employee generates a re-encryption key for his or her deputy before leaving, with which the mail server is able to translate incoming mail during the absence. Advanced re-encryption schemes even allow one to program expiration dates into the re-encryption keys, thereby further reducing the required trust assumptions.

3 Summary of Focus Groups

In this section, we summarize the outcome of three parallel focus groups that were held at the end of our workshop with CREDENTIAL project members and participants of the IFIP Summer School. In these groups, we discussed open research challenges in regard to the tradeoff between privacy, efficiency, and usability, end user trust, and adoption factors that we identified for the CREDENTIAL project. Possible approaches and solutions in regard to how these challenges could be addressed were discussed.

3.1 Focus Group 1: Privacy Challenges in CREDENTIAL

In the first focus group, we discussed privacy issues and the technologies to solve them, as well as the effects they may have on usability and efficiency. In particular the discussion was about the need of metadata privacy, i.e., whether information like access patterns, file sizes, or access rights need to be considered as sensitive or not. The following lines summarize the questions and discussions:

Which metadata needs to be protected, which may be leaked? It was discussed that it is not possible to generally say what kind of metadata is sensitive because it depends on the use cases, to what extent we want to trust the central identity wallet, and possible leakages from other sources the wallet might be able to access. If your file includes health data and the wallet already knows it, the frequency of accessing the file may divulge the severity of your disease. Knowing about the access rights of the files also reveals a lot about the type of the document. It was also suggested that we shall apply anonymization techniques on the receiver side and we should review the regulations related to the health domain to ascertain what kind of metadata is sensitive due to regulations. Furthermore, we should define some performance and privacy goals to see the added-value that we may gain by utilization of cryptographic techniques to hide the metadata.

May the wallet be allowed to see access patterns? In general, participants discussed that we should define objective decision criteria to decide on the assumptions we want to or we can make. Those criteria should not only be security-related, but also cover acceptable running times on well-defined devices or required costs to realize the solution. This is because exerting enhanced cryptographic tools like oblivious transfer (OT) or private information retrieval (PIR) would massively increase the requirements on the computational capabilities of the server. Depending on the available resources, this might render the entire system impracticable or too expensive.

How could access and behavior patterns be hidden from the wallet? Specifically, this question aimed at receiving feedback on whether splitting the identity providers into two components and making a non-collusion assumption would be a viable alternative to cryptographic "sledgehammers" like PIR. Alternatively, one could think about routing all data through the user's device

so as not to enable the identity provider (IdP) learn which service provider is currently contacted.

It was discussed that the information about access patterns are definitely sensitive and should be hidden from the IdP or disguised. It is therefore valuable for users' privacy to avoid the IdP from learning access patterns in a usable way. The participants discussed that—depending on the concrete context—they tend to prefer to trust their own mobile phones rather than identity providers, meaning that information should be routed through the user's device. However, this approach would typically require key management and/or heavy computations on the user's side, which might not always be possible. Furthermore, this approach would contradict the idea of a fully cloudified solution, and therefore might work for identity provisioning but not for data sharing in general, as the data owner might not be online when the receiver wants to access the data.

3.2 Focus Group 2: Establishing Trust in CREDENTIAL

In the second focus group, we discussed how to build the user's trust into the solution developed within CREDENTIAL, as this represents a key factor for adoption. The following lines summarize the questions and discussions:

Which aspects could be used to build trust? It was suggested that both the ability as well as motivation of the entity deploying and maintaining the CREDENTIAL solution has to be clearly stated. The competence assessment can be delegated to another trusted party, for example a consumer agency, which issues certificates or simplified results as icons. Also, the motivation of the deploying party is an important factor, which is influenced by this party's business plan but also legal consequences of not complying, for example, to the General Data Protection Regulation (GDPR) [7]. Furthermore, a gradual approach of first experiencing the system with non-sensitive data and later expanding to further use cases might help users to gain familiarity and overcome trust obstacles.

How should users be informed about potential risks, which still remain even after applying the CREDENTIAL approach? It was discussed that potential risks should be explained honestly and concisely, as users are then able to make informed decisions and are more willing to accept consequences if they occur. This information could be provided in a layered approach, where non-technical scenarios are first presented, for example as cartoons, but further details are also available. However, a balance has to be found where users are sufficiently informed but are not scared away so that they use less secure alternatives which do not explain their potential weaknesses with the same level of detail.

Should users be offered detailed access records for their data? Providing access logs to users would be very beneficial, as these records can not only be used to detect abuse, but they also inform the user about progress and involved people, for example when sharing data for a hospital visit. Having a detailed log might also be helpful in law suits or for quality assurance reasons.

3.3 Focus Group 3: Technology Adoption and Applications

The third focus group discussed challenges and strategies for promoting the adoption of CREDENTIAL technology by individuals and organizations and formulated hypotheses for questions which the project could investigate in its future research on adoption factors:

How can users be convinced to change to CREDENTIAL technology? It was suggested that adoption of a CREDENTIAL application could be promoted by framing and marketing it as a "Secure App", similarly to the Swedish BankID solution [8] that is basically perceived as a secure and trustworthy authentication solution by Swedish citizens. Moreover, additional functionalities should be provided and marketed, which CREDENTIAL could in contrast with other solution providers offer in a secure and privacy-friendly way. For instance, additional data sharing applications could be offered enforcing data minimization, e.g., for social sharing of pictures and posts, or sharing of confidential documents by companies. If CREDENTIAL can promise compliance with the GDPR for its applications, this can be another important adoption factor for industry and government.

Who should pay for the adoption of CREDENTIAL technology? The reason for this question was that protecting the privacy of end-users might harm the business model of currently free data sharing or identity provisioning services, as other sources of income (e.g., personalized advertisements) would not be possible any more. It was discussed that in Sweden and other countries, there are efforts spent by the government to offer patients more treatments at home and to keep them away from hospitals as long as possible, as this has been proven to be also a more cost-efficient solution. The CREDENTIAL eHealth solution should therefore not only provide benefits for the patients, but should also enable cost savings for Health insurances and the public Health Care system, which should therefore also have incentives to pay for it. Other options could be non-monetary payment solutions such as novel "tagging-based payment systems", where users could use the service basically for free, in return for getting every X minutes a picture in which they should tag people or objects.

4 Why Metadata-Privacy Matters

As already briefly discussed in Section 3.1, hiding metadata is an essential requirement for a privacy-preserving identity provisioning and data sharing platform. This section is dedicated to defining metadata in more detail and to explaining some of the potential privacy risks related to metadata.

Metadata defines and describes the data and is often referred to as "data about data". It is the information about a particular data set which may describe how, when, and by whom it was received, created, accessed, and/or modified and how it is formatted. Metadata is not limited to the files. It also encompasses the records of interactions (e.g., a history of login times) or simple facts about individuals (e.g, an account number or mailing address) [9]. Metadata comes from a variety of sources; it can be created automatically by a computer, supplied

by a user, or inferred through a relationship with another document to serve various purposes and functions, including enhancing the editing, viewing, filing, and retrieval of documents [9]. It opens the possibility for search engines to index the information. It promotes clearer understanding, better data management, and efficient use and reuse of information.

However, despite the benefits of collecting and analyzing metadata, all too often metadata lacks the privacy protections afforded to the content [10] and limited privacy protection may expose sensitive information to the wrong people and present significant risks. Possible inferences from metadata can invade a user's privacy in the same way as sensitive personal information obtained from the content [11]. Literally, the emergence of new technologies manifests how sensitive metadata can be: how friend lists can reveal a person's sexual orientation or political views, purchase histories can identify a pregnancy before any visible signs appear, and location information can expose individuals to harassment for unpopular political views, theft, or even physical harm [10]. Thus to protect the privacy of personal data we should consider the necessity of minimizing the amount of accessible metadata.

As mentioned previously, all the contents stored in the CREDENTIAL wallet are encrypted and the wallet does not have access to the plain-text. Ostensibly, the idea of the encrypted content implies the elimination of privacy issues. Nonetheless, metadata of encrypted content can still leak lots of personal and sensitive information not originating from the content itself but from its properties or generated information during content management to the CREDENTIAL wallet. Ergo, we want to illustrate which threats to the users' privacy still remain when metadata is available in plain-text. We only consider a simple design of the cloud-based IDMaaS system in order to give a comprehensive overview of the possible problems with honest-but-curious adversary model for the CREDENTIAL wallet.

In the following sections, we analyze the consequences of metadata privacy leakage and access to metadata by the wallet as an IdP and file hosting service.

4.1 Metadata for Identity Providers

When users select CREDENTIAL to be authenticated for different services, CREDENTIAL wallet as an identity provider can infer lots of information from the communication flow as explained in the following.

Behavioral pattern. Albeit by design the wallet cannot see the plain attributes which service providers request to authenticate the users, it is aware of when, from where, and how frequent the users connect to service providers and what kind of attributes are required. As an IdP, CREDENTIAL can glean the information about the IP address and the devices from which users connect to different services and it can guess the location. Knowing the IP address, the wallet could further infer general usage patterns, such as the user's online times or working habits (when the user connects from which device). Furthermore, even in the case that the IP address is disguised (e.g., using Tor) the accessed services might

reveal some information about the user's physical location, as one often uses regional or at least national services, e.g., for online banking or governmental tasks.

Moreover, the type of the service provider could leak information about the user's gender (some services are mainly used by a specific sex) or income situation (prices of services). Based on the attributes they require to authenticate users (e.g., if the user is over 18), CREDENTIAL can learn about the user's age range withal. The user's age range can also be derived from kind of service providers they use (e.g., pension-related services). The broader the range of authentication and usage of service providers, the more CREDENTIAL can learn about the users' attitudes and personal lives (e.g., websites related to special political parties or communities in society). Also very sensitive information might be leaked through access frequencies (e.g., addictions).

Behavioral pattern combined with background information. While CREDENTIAL can learn a lot directly from analyzing when, from where, how frequently and to whom the users communicate, it can also combine this kind of information with some background information it has extracted about the user and learn more. For example, if the wallet already knows that a user shares health data frequently with another user (probably a doctor), a request from an online pharmacy website to authenticate the user, reveals that the illness is in a new stage or user's health has not improved because the patient has to buy new medicine or repeat the previous ones.

4.2 Metadata for Data-Sharing

Similarly to the case of identity provisioning the CREDENTIAL wallet infers lots of information about stored content when acting as a file hosting service. The size, structure, modification time and access rights of a file may reveal information that the user originally intended to hide. In the following we give examples for each of these properties.

Size, name and type. Even if the type of an object was not known to the wallet, the size of the cipher-text would act as an indicator of the content type of the stored object (e.g., text, image, and video). Type or size of an object when combined with the object's name (especially if it is assigned by the users and not the system) divulge more information about the content. Additionally, consider a scenario in which some users (patients) should upload and share their blood pressure collected over three months with another user (doctor). Users benefit from the same kind of health devices (sphygmomanometers introduced by doctors using CREDENTIAL) and health applications (applications for health devices) to collect the necessary information and send it to the wallet. Blood pressure is collected over a specific period so the number of records in each file is probably the same and due to the fact that they are collected from the same kind of health devices and apps, they have some common metadata like size, type and access rights defined by the data owner. Correspondingly, the wallet can guess that the users have the same disease and the type of illness may also

be derived from the identity of the one with whom the files are shared. Doctor's identity or field of expertise might be revealed previously from some background data.

Access rights and patterns. Metadata describing access rights granted to different objects reveals not only some extra information to the wallet but also some information to the people who have been granted access to a same file. For example, all users who have the read access to a file can see between whom the file is shared which may not be desirable for the owner. In addition, being aware of the fact that some users have shared medical or identity documents with the same user or the request to access some files initiated by a specific user is always accepted helps CREDENTIAL to infer the identity or occupation of the user with whom the files have been shared. Because that user seems to be trusted by the others and the characteristics of the files shared with her support the inferences that she may be a doctor or from government sector. Gradually learning about users' occupations, the fact that a file is shared between different people can reveal lots of sensitive information. A high amount of medical documents related to a user indicates frequent consultations of medical practitioners. A document shared among several medical practitioners may indicate an uncommon disease. On the other hand, certain access patterns by a gynecologist may leak more positive news than regular accesses by an oncologist.

More interestingly, the user's behaviors not only progressively contribute to the information CREDENTIAL collects in the user's profile but also add to what CREDENTIAL knows about the people who have some common characteristics with the user (e.g., they all share some documents with a doctor). As a consequence, requesting to authenticate to a service provider like a psychiatric hospital by one of the users can reveal the type of common disease they share together.

Date of modification. Monitoring the cipher-text for changes reveals possibly sensitive information considering the fact that type and location are also among visible metadata. The modification history can for example tell something about the frequency of a user's illness updates, the intensity of her illness or monitoring her illness activity, or other general usage patterns.

Structure. The structure of an encrypted object may be needed to selectively download or insert some parts to an object or to enable redacting some parts. But it can also result in leaking extra information. For example, a folder is shared between different users but not all of them have the same access rights for all objects in the folder. Inferring from the data structure that there are more objects in the folder than what a user has access to, the user learns the exact numbers of objects that the owner of the folder has made hidden for her.

5 Hiding Metadata

In this section, we briefly recap some important solutions from the cryptographic literature that address the metadata problem and offer partial solutions by hiding certain types of metadata. Unfortunately, for many of the proposed solu-

tions, only inefficient instantiations are known and therefore the given techniques are not entirely suitable for real-world usage. Nevertheless, solving the hiding-metadata problem with efficient instantiations is an active and vibrant area of research.

Private information retrieval. In a private information retrieval (PIR) protocol, a user is able to query database records such that the database server is unable to tell which of the records the user accessed during that query and which not. The cryptogaphic literature offers a variety of proposed instantiations which can be categorized at least in single-server and multi-server variants in the information-theoretic or computational setting with different communication and computation complexity overhead, see, e.g., [12,13]. Unfortunately, in such systems, each database record has to be accessed by the database server for each user query, since otherwise the server would trivially gain at least some information about the query which renders the schemes inherently inefficient (if the server does not need to touch a record, it can not have been queried by the user). Additionally, within the PIR setting, a user might be able to obtain any record of the user's choice; hence, the PIR paradigm lacks some hidden access control mechanisms which, however, can be overcome with the following Hidden Access Control Policies (HACOT) technique.

Oblivious transfer with hidden access control policies. In the work of Camenisch et al. [14], the term Oblivious Transfer with HACOT was coined and the most efficient instantiation by now (and also a real-world implementation thereof) was given by Camenisch et al. [15]. Within their setting, each database record is associated with some hidden access control policy while any user is able to query the database in an anonymous fashion. The proposed solution uses attribute-based encryption and the user can obtain a database record if and only if the attribute(s) associated with the user's secret key matches the policy of the database record. In particular, the HACOT technique does not reveal to the database provider what database record was being accessed, what the policy of that record is, who accessed the record, and whether the access was denied or granted. (This holds even in a strong adversarial model, where a malicious database server is allowed to collude with the key issuer.) Nevertheless, the database server can monitor the amount of total accesses. Unfortunately, the given HACOT real-world implementation from [15] is quite inefficient in a large-scale setting.

Oblivious random-access memory. All described techniques above only consider reading of database records. If however one wants to deal with oblivious write operations, one has to consider oblivious random-access memory (ORAM). The technique and a scheme was proposed by Goldreich and Ostrovsky [16]. A more efficient ORAM instantiation (based on the Goldreich-Ostrovsky system) was given by Pinkas and Reinman [17]. Unfortunately, as with the other described techniques above, ORAM is still considered to be quite impractical for real-world large-scale usage.

Anonymous credentials. Specifically in the case of identity provisioning, anonymous credentials can be used to authenticate a user while only revealing the absolutely minimum necessary information. Such schemes have been envisioned by Chaum [18,19], and several practically efficient instantiations exist, with the most prominent ones being IBM's identity mixer [20,21] and Microsoft's UProve [22]. In such a system, the user receives a credential on its attributes, and can dynamically choose which attributes to reveal to a service provider, while not leaking any information about undisclosed attributes. In those schemes, the authentication process is done fully locally, and thus there is no need for a central IdP and also the related metadata-privacy problems disappear. However, these constructions are inherently non-cloudified, as the attributes are required in the plain for performing authentication, contradicting the ambition of a solution not requiring any key material, dedicated software, or heavy computations on the user's side, cf. Section 3.1.

Recently, Krenn et al. [23] suggested the first anonymous credential system based on proxy re-encryption, where the authentication can be outsourced to a semi-trusted IdP, having the advantage that neither specific software nor heavy computations are required on the users' side. However, re-inventing the central IdP also re-introduces all the metadata problems, which are not further addressed in their work.

6 Metadata-Hiding Identity Provisioning

As discussed in the previous section, existing cryptographic mechanisms are either not suited or too inefficient for usage in a large-scale central identity provider. In the following we will therefore present a potential high-level architecture of a metadata-privacy respecting IdP. The main idea is to split the identity provider into two components. A first component communicates with the user, whereas a second component communicates with the service providers. Assuming that those two entities do not collaborate, we can circumvent many of the discussed metadata problems while not negatively impacting the usability of the overall system.

In the following we assume that a user U received an authentication credential in form of a signature on encrypted attributes from an issuer, and stored it to the first part of the identity provider, i.e., IdP_1, together with a re-encryption key from its personal public key to the public key of the second part of the identity provider, i.e., IdP_2 (note here that this computation requires access to the user's secret key). The used encryption scheme needs to support at least two subsequent proxy re-encryption operations. Furthermore, we assume that the user wants to log in to a service provider SP, which has already been accessed by some user previously. If this was not the case, the SP and the two identity providers jointly perform in a multi-party computation protocol, at the end of which IdP_2 learns the required re-encryption key from its own public key to that of SP. This protocol is necessary because for obvious reasons neither IdP_1 nor

IdP_2 must learn the secret key corresponding to IdP_2's public key, but rather the key needs to be appropriately shared between them.

The following protocol shows the message flow for a login process, assuming that every sent message is supplemented by a zero-knowledge proof of knowledge showing the correctness of the performed computations [24].

- In a first step, U contacts SP (via some anonymized network layer, e.g., using Tor as an anonymization service). The two parties agree on a random session identifier *sid* and the set of attributes to be disclosed.
- Next, the user encrypts the indices of the required attributes, *sid*, and the identity of SP under the public key of the second part of the identity provider, IdP_2, yielding a ciphertext c.
- The user then contacts (and potentially authenticates itself towards) the first part of the identity provider, IdP_1. It hands c to IdP_1.
- In a fourth step, IdP_1 re-encrypts a re-randomized version of U's encrypted credential to the public key of IdP_2, and forwards the result together with c to IdP_2.
- Next, IdP_2 decrypts c and computes a presentation token for SP. It therefore re-encrypts the requested attributes for SP, thereby redacting the remaining attributes in order to show the validity of the underlying credential. This can, e.g., be done similar to Krenn et al. [23].
- The service provider contacts IdP_2 and hands over *sid*.
- If the received *sid*'s are equal, IdP_2 forwards the presentation token to SP, which grants the user access in case that the presentation token is valid.

The above approach requires that the two parts of the identity provider do not collude. In this case, the following privacy requirements can be given. On the one hand, IdP_1 only learns that U is authenticating to some service provider, but it neither learns which attributes will be disclosed nor the identity of the service provider. Therefore, no sensitive information is leaked to IdP_1. On the other hand, IdP_2 only learns that some user is authenticating itself towards a specific service provider, but it neither learns the user nor the plain values of the revealed attributes. Again, IdP_2 therefore does not learn critical information about a specific user if the number of users is sufficiently large.

From a computational point of view, the costs of our conceptual architecture should be comparable to previous single IdP solutions such as, e.g., Krenn et al. [23], while giving far higher privacy guarantees. Security-wise the above protocol achieves similar guarantees as [23] by assuming essentially passive adversaries in the sense that the identity providers try to learn as much information about the user, but do not actively behave maliciously.

7 Conclusion

The workshop took advantage of the different expertise present in the focus groups. Experts from different domains discussed the opportunities and challenges of the CREDENTIAL approach towards a privacy-preserving IDMaaS and

data sharing solution. The discussions covered various aspects such as challenges related to user trust, or potential business models and applications of the resulting systems. Special attention was paid to the necessity of metadata-privacy and the related technical obstacles. This report summarizes the results of the focus groups. Furthermore, it presents a potential high-level architecture of a metadata-privacy respecting identity provider which arose based on the inputs from the focus groups.

A first step for future work will include the formal modeling, specification, and security and privacy analysis of this proposed concept architecture.

Acknowledgement. The authors would like to thank the workshop participants for their valuable comments and inputs during the focus groups. In particular, we want to thank Anna Klughammar for her introductory presentation on the eHealth pilot, Charlotte Bäccman and John Sören Pettersson for moderating the second focus group, and Karl Koch for providing us with his extra notes for the third focus group.

References

1. Blaze, M., Bleumer, G., Strauss, M.: Divertible Protocols and Atomic Proxy Cryptography. In Nyberg, K., ed.: EUROCRYPT '98. Volume 1403 of LNCS., Springer (1998) 127–144
2. Johnson, R., Molnar, D., Song, D., Wagner, D.: Homomorphic Signature Schemes. In Preneel, B., ed.: CT-RSA 2002. Volume 2271 of LNCS., Springer (2002) 244–262
3. Hörandner, F., Krenn, S., Migliavacca, A., Thiemer, F., Zwattendorfer, B.: CREDENTIAL: A Framework for Privacy-Preserving Cloud-Based Data Sharing. In: Availability, Reliability and Trust – SECPID@ARES 2016. (2016) 742–749
4. Cantor, S., Kemp, J., Philpott, R., Maler, E.: Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2 .0 - Errata Composite. Technical report, OASIS (2009)
5. Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., Mortimore, C.: OpenID Connect Core 1.0. Technical report, OpenID (2014)
6. European Commission: Regulation (EU) No 910/2014 of the European Parliament and the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market. Official Journal of the European Union **L257/73** (2014)
7. European Commission: Regulation (EU) No 679/2016 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Official Journal of the European Union **L119/59** (2016)
8. BankID: Swedish BankID. (https://www.bankid.com/en/) accessed: 2016-10-27.
9. Breen, M.: Nothing to hide: Why metadata should be presumed relevant. Kansas Law Review, Kansas Law Review Inc. **2/1/2008: vol. 56(2).** (2008)
10. Conley, C.: Metadata: Piecing together a privacy solution. **2/1/2008: vol. 56(2).** (2014) Available at SSRN:https://ssrn.com/abstract=2573962 or http://dx.doi.org/10.2139/ssrn.2573962.

11. Greschbach, B., Kreitz, G., Buchegger, S.: The Devil is in the Metadata - New Privacy Challenges in Decentralised Online Social Networks. In: PerCom 2012, IEEE Computer Society (2012) 333–339
12. Chor, B., Kushilevitz, E., Goldreich, O., Sudan, M.: Private information retrieval. J. ACM **45** (1998) 965–981
13. Kushilevitz, E., Ostrovsky, R.: Replication is NOT Needed: SINGLE Database, Computationally-Private Information Retrieval. In: FOCS '97, IEEE Computer Society (1997) 364–373
14. Camenisch, J., Dubovitskaya, M., Neven, G., Zaverucha, G.M.: Oblivious Transfer with Hidden Access Control Policies. In Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A., eds.: PKC 2011. Volume 6571 of LNCS., Springer (2011) 192–209
15. Camenisch, J., Dubovitskaya, M., Enderlein, R.R., Neven, G.: Oblivious Transfer with Hidden Access Control from Attribute-Based Encryption. In Visconti, I., Prisco, R.D., eds.: SCN 2012. Volume 7485 of LNCS., Springer (2012) 559–579
16. Goldreich, O., Ostrovsky, R.: Software Protection and Simulation on Oblivious RAMs. J. ACM **43** (1996) 431–473
17. Pinkas, B., Reinman, T.: Oblivious ram revisited. In: Advances in Cryptology – CRYPTO 2010: 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings, Springer Berlin Heidelberg (2010) 502–519
18. Chaum, D.: Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. Commun. ACM **24** (1981) 84–88
19. Chaum, D.: Security Without Identification: Transaction Systems to Make Big Brother Obsolete. Commun. ACM **28** (1985) 1030–1044
20. Camenisch, J., Herreweghen, E.V.: Design and Implementation of the *idemix* Anonymous Credential System. In Atluri, V., ed.: CCS 2002, ACM (2002) 21–30
21. Camenisch, J., Lysyanskaya, A.: An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In Pfitzmann, B., ed.: EUROCRYPT 2001. Volume 2045 of LNCS., Springer (2001) 93–118
22. Paquin, C., Zaverucha, G.: U-prove Cryptographic Specification v1.1 (Revision 2). Technical report, Microsoft Corporation (2013)
23. Krenn, S., Salzer, A., Striecks, C.: Attribute-based credentials on encrypted attributes. unpublished manuscript (2016)
24. Bellare, M., Goldreich, O.: On defining proofs of knowledge. In Brickell, E.F., ed.: CRYPTO '92. Volume 740 of LNCS., Springer (1992) 390–420