

Decentralized Identities for Self-sovereign End-users (DISSENS)

Martin Schanzenbach¹, Christian Grothoff², Hansjürg Wenger², Maximilian Kaul¹

Abstract: This paper describes a comprehensive architecture and reference implementation for privacy-preserving identity management that bucks the trend towards centralization present in contemporary proposals. DISSENS integrates a technology stack which combines privacy-friendly online payments with self-sovereign personal data management using a decentralized directory service. This enables users to be in complete control of their digital identity and personal information while at the same time being able to selectively share information necessary to easily use commercial services. Our pilot demonstrates the viability of a sustainable, user-centric, standards-compliant and accessible use case for public service employees and students in the domain of retail e-commerce.

We leverage innovative technologies including self-sovereign identity, privacy credentials, and privacy-friendly digital payments in combination with established standards to provide easy-to-adapt templates for the integration of various scenarios and use cases.

Keywords: privacy; decentralization; digital sovereignty

1 Introduction

The cryptography for secure, privacy-preserving online payments has been known since Chaum's ground-breaking work [3] around 1990. While such digital payments can in theory enable anonymous one-click purchases, in practice online services sometimes require contextual personal information from their customers, the most basic example being a shipping address for delivery.

Furthermore, some of the contextual information may require verification. Today, online services commonly validate contact information like phone numbers or e-mails by sending test messages, burdening users with an extra processing step. Verification becomes even more complex if shops need to satisfy regulatory requirements involving age verification. But not only minors might be interested in supplying false credentials, as governments and other institutions sometimes subsidize eligible citizens such as students, pensioners and the disabled.

While we accept that online services may to some degree have a legitimate need for their users' personal information, we contend that the contemporary practice of registering user

¹ Fraunhofer AISEC, München, Germany, firstname.lastname@aisec.fraunhofer.de

² Bern University of Applied Sciences, Biel/Bienne, Switzerland, firstname.lastname@bfh.ch



accounts prior to receiving online services is an unnecessary and obsolete burden on users. Additionally, with legislation such as the General Data Protection Regulation (GDPR) user information and accounts pose a significant liability for service providers, and setting up and maintaining accounts and associated credentials with the multitude of service providers in existence today is a usability nightmare.

The latter is partially addressed through the use of third-party identity provider services (IdPs). Using an IdP is practical for online services as it delegates the task of verifying and storing personal information while reducing the number of accounts users have to manage. However, the resulting centralization further harms the users' ability to control their data as it becomes easily linkable by the IdP across various life domains. Collecting user data that tracks users' movements across services is a business model that strikes at the core of the right to informational self-determination.

1.1 Transforming the role of the IdP

What is actually needed is a technology stack which gives users a maximum amount of personal control over their digital identities and at the same time relieves service providers from the liability of user data, including payment information. Self-sovereign identity (SSI) management is a way to replace IdPs with a user-centric, decentralized mechanism where data- and access-control are fully under the control of the data subject. In SSI, users are free to create many a priori unlinkable pseudonyms, and ideally they decide which of their attributes they share with which service providers. SSI thus enables users to exercise the right to be forgotten by deleting pseudonyms, which is only meaningful in systems where users can have multiple concurrently useful identities. Attributes can generally be set freely for the various pseudonyms by each user.

As a result, the role of IdPs is transformed: instead of controlling the subject's data, IdPs in SSI are certification bodies that provide attestations for the authenticity of certain attributes. Here, information disclosure can be further minimized using privacy-friendly credentials [2]. Given the possibility of IdPs offering competing certification services, users can pick and choose sets of attested attributes from any number of IdPs. This concept is natural, as the authority over a user's email address is likely different from the authority over the user's employment status, for example.

1.2 Objectives

Our objective is to enable seamless use of online services where users generally do not need any accounts, while improving on the best-case user experience that platforms with accounts can provide. Instead of remembering login information with usernames and passwords, multi-factor authentication, and disclosure of personal information when yet another service

provider’s database is leaked online, users should only authorize to share relevant personal details *while* necessary. GDPR-compliant services would then resolve those personal details only at the time of need without ever storing them, thereby minimizing the chance of inadvertently compromising private information of their users.

Eliminating the online service provider’s data liability should include eliminating the need to store sensitive payment data. Instead of having to outsource payments or to satisfy the non-trivial requirements of the Payment Card Industry Data Security Standard, payments should be processed without revealing any personally identifiable information about the consumer to anyone. As a result, the need for users to enter personal information, credit card data or to pass authorization procedures is completely eliminated — except for logging into their own computer.

We also want to satisfy the practical requirement that back-office business processes related to fulfillment might be executed when the user is not directly interacting with the service. Thus, even though the service is not itself storing a user’s personal data, authorized services must be able to access it even if the user is offline. Finally, we want to maximize compatibility to existing standards where possible, to minimize the migration cost.

1.3 Technical contribution

We integrated the self-sovereign identity management system re:claimID [8] with the privacy-preserving GNU Taler payment system in a pilot application based on WooCommerce, a widely used online shop based on WordPress. Our Free Software pilot demonstrates how contemporary cumbersome and privacy-unfriendly online shopping processes can be eliminated, improving the user experience, privacy and security of all parties involved. Our reference scenario includes credentials issued by the partners Fraunhofer AISEC and BFH for employees and students, respectively. The Web shop can thus distinguish between “trusted” attributes that are certified by these two IdPs and “untrusted” attributes that are freely chosen by the users.

Our system demonstrates how businesses, users and credential authorities can easily integrate using standard protocols such as OpenID Connect (OIDC), without requiring any a priori registration at Big Tech gatekeepers such as Apple, Google, Facebook or Amazon. The existing OIDC standard already accommodates identity aggregation features necessary in SSI use cases. The resulting architecture enables citizens and businesses to avoid vendor lock-in to those major platform providers without sacrificing usability or privacy. In our system users remain in direct control of their personal data. This is an alternative to the contemporary privacy-unfriendly approaches which either call for governments or Big Tech to “safeguard” personal information storage or access control.

Incidentally, we show that SSI systems can be built without distributed ledger technology (DLT), challenging the common wisdom that DLTs are crucial in this domain.

1.4 Unique benefits

This approach offers significant benefits over existing solutions built using other SSI systems such as Sovrin ³ or sertto (formerly uPort) ⁴.

No gatekeepers; No vendor lock-in: Our approach is completely open to issuers and does not impose any registration restrictions (such as registration fees) in order to define domain specific credentials. Further, our system does not impose a consortium-based governance model — which tend to eventually be driven by commercial interests and not consumer interests. Our design enables all participants in the ecosystem to participate without prior onboarding while at the same time being offered full transparency and control regarding their personal data and processes involved. Finally, we try to integrate our technology stacks as much as possible with existing standards in order to facilitate transitioning.

Support for non-interactive business processes: At the same time, unlike the SSI systems cited above, our technology offers a way to access user information without online interaction with the user. Offline access of shared identity data is a crucial requirement in almost any business process as such processes often occur after direct interaction with the user. For example, customer information such as billing addresses are required in — possibly recurring — back office billing processes which occur well after interaction with a customer.

Scalability and sustainability: Finally, both re:claimID as the SSI system as well as Taler do not suffer from the usual predicament Blockchain-based systems find themselves in: Both systems do not require a decentralized, public ledger. This eliminates the need for consensus mechanisms, which do not scale and are ecologically unsustainable. In fact, we employ decentralization only where it provides the most value and use more efficient technology stacks where needed: re:claimID builds on top of the GNS, which makes use of a DHT, an efficient ($O(\log n)$) peer-to-peer data structure. For payments, GNU Taler uses centralized infrastructure operated by audited and regulated exchange providers and facilitates account-less end-to-end interactions between customers and services where all parties have $O(1)$ transaction costs.

For a comprehensive discussion and comparison of re:claimID and other SSI systems we refer to [8].

2 Background

Our architecture builds on three core technologies, which are introduced in this section.

³ <https://www.sovrin.org>, accessed 2021/02/12

⁴ <https://www.sertto.id>, accessed 2021/01/12

2.1 Classical identity providers

Identities are usually managed through the use of so-called directory services. Traditionally, directory services based on the X.500 protocol family or derivatives such as LDAP are used for identity and organizational information management. While it may not be intuitive at first what name systems such as DNS have to do with identities, efforts such as NameID show how name systems can be used by users to manage pseudonyms and personal information in a decentralized directory.

In most cases, access to identity information through an identity provider is realized using an authorization and authentication service such as SAML or OIDC. Behind every identity provider, there is either a directory service which is used to supply identity information or another identity provider as part of a federation.

In our use case, we presume the existence of identity provider services which follow the traditional approach outlined above. We will show how these services can be used to provision certified attributes in a user-centric SSI architecture.

2.2 re:claimID and the GNU Name System

re:claimID [8] is a self-sovereign identity management system developed by Fraunhofer AISEC which uses the GNU Name System (GNS) [9] as a directory service to store credentials and provide authorized parties with access to identity data even when the respective user is temporarily offline.

GNS is an alternative name system developed in the context of the GNUnet project. It offers secure and decentralized directory storage for identity and personal data. Encrypted name data in GNS is stored in a distributed hash table (DHT). To decrypt the information, one must know a *label* and a *public key*. It should be noted that *these* public keys can be shared secrets and are not exposed by the GNS protocol. GNS enables privacy-preserving name resolution because the peers serving answers do not inherently learn anything about the information they are helping process. GNS is fully decentralized: each user is free to define their own root zone which serves as their trust anchor.

In re:claimID, users manage their identities without the need for a single IdP while at the same time retaining the ability to securely and selectively disclose identity information with other services. Complete transparency, decentralization and elimination of intermediaries are the core value offerings of re:claimID. re:claimID also offers an OIDC compatibility layer which facilitates integration with existing standards-compliant client software.

2.3 GNU Taler

GNU Taler is a privacy-preserving payment system using blind signatures to protect the identity of payers. Using blindly signed digital coins to create digital cash was first proposed by Chaum [3]. Taler extends Chaum's work by providing an efficient mechanism to obtain unlinkable change while preserving the income transparency properties of Chaum's original design [5]. The resulting payment system is expected to be generally compatible with financial regulation and neutral with respect to central bank's ability to implement monetary policies [4].

Taler allows consumers to authorize payments using a single click in their Taler wallet. Two-factor authentication is implicit from the consumer having physical control over the wallet's computing device and the ability to unlock it. Before DISSENS, this account-less single-click shopping experience did not work for physical goods where the shop needs a delivery address, or even for digital goods if they require age verification.

3 Architecture

To create a broadly applicable one-click online shopping experience without the need to create and maintain accounts where personal information is persisted with the shop outside of the control of the user, we combine the decentralized, SSI and personal data management system re:claimID with the privacy-preserving GNU Taler payment system as illustrated in Figure 1. Identity providers (such as sovereign states and academic institutions) issue credentials to users. Users manage identities with associated third-party attested credentials and self-attested attributes in a self-sovereign fashion. Users publish their encrypted identity data to the GNS, which serves as a decentralized directory. They can then selectively disclose sets of attributes to relying parties.

3.1 Third party identity providers

The design of re:claimID does not require the existence of third party identity providers. Users may self-attest and self-sign attributes. In fact, this can be considered the default and is probably sufficient for many use cases. However, as elaborated in Section 1, there are use cases where a relying party may require an assertion from a trusted third party. We assume that such trusted third parties operate OpenID Provider (OP) services. OPs are assumed to have issued unique user identifiers in the form of email addresses to facilitate the discovery of the OP service through the use of OIDC Discovery [6]. This requirement could in theory be relaxed, but OP discovery through this method is the simplest form of discovery within the standard. The OPs are connected to the respective institution's directory, either directly (e.g. LDAP) or through federation (e.g. another upstream OP). The backend architecture of the IdP is irrelevant as long as it exposes OIDC endpoints.

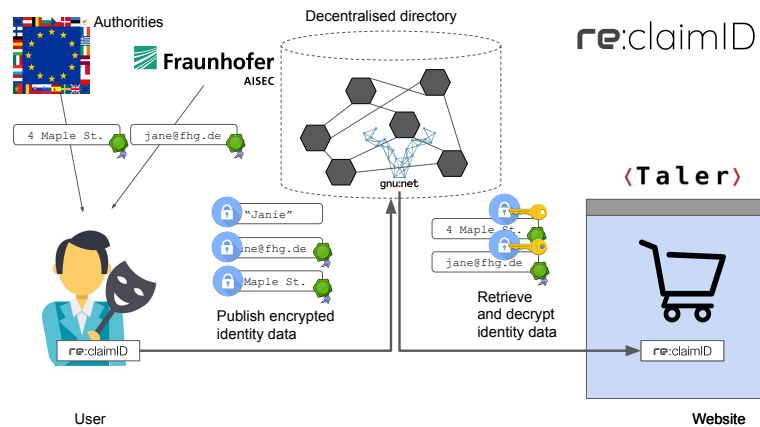


Fig. 1: The DISSENS architecture.

An important caveat for the OP setup is that it must allow re:claimID to request authorization to access user information. For this, the OPs must either have pre-configured a client for re:claimID, or allow for dynamic client registration for the re:claimID redirect URI as a public client. The re:claimID redirect URI is not a central server, but points the browser to the re:claimID web extension which forwards any authorization request to the re:claimID service running on the local device.

In our design, we optionally allow IdPs to provide credentials in a non standards-compliant fashion to re:claimID: Our new API exposed at the IdP allows re:claimID to retrieve a credential which allows the user to selectively disclose *subsets* of the attributes attested in the credential.

3.2 User identity management

In our architecture, users can manage any number of pseudonyms with attributes containing their personal information on their own devices. Through re:claimID, attributes basically become records in a zone of GNS. The encrypted attribute data is published in a DHT. When a user interacts with an online service, they can provide the online service a re:claimID “ticket” [8] which grants the service access to subsets of the attributes of a user’s pseudonym. The ticket can be piggybacked within the authorization code of an OIDC authorization flow. It contains the necessary information to find the attributes in GNS. With it, the service can access the data when it is needed, even if the user is offline at the time. It also allows the service to retrieve updates in case the user changes their information. The user can choose to revoke the grant at any time, henceforth preventing the store from resolving attribute data.

From the user’s perspective, interacting with re:claimID is equivalent to interacting with other OPs, with the addition that each user is able to manage multiple pseudonyms. Users can freely add or remove attribute data associated with any of their pseudonyms. They are expected to interact with re:claimID through a browser extension which communicates with a locally running client.

Users can import certified attributes by authenticating with identity providers and extracting certified attributes using the OIDC protocol [6]. To do so, users first add an attribute with an email address registered with the third party OP to one of their pseudonyms. The OIDC Discovery protocol [6] allows re:claimID to discover the relevant OIDC endpoints. If necessary, re:claimID will attempt to dynamically register an OIDC client at the OP. The user can then initiate an OIDC authorization code to allow re:claimID to retrieve an ID token asserting information about the user.

The user may then add new attributes to their pseudonym and select attribute values from the contents of the ID token. When the user shares attributes which are backed by an ID token, re:claimID will provide the token to the RP as part of the OIDC protocol: As an OIDC OP, it will issue an ID token and provide access to a userinfo endpoint. The payloads will include references to ID tokens issued by third parties as defined in the OIDC specification [6] as “Aggregated Claims”. The standard defines how JWTs such as ID tokens can be used within an aggregated claim.

One problem with the concept of aggregated claims in combination with JWTs is excessive disclosure of information. While anonymous credentials such as Camenisch-Lysyanskaya [2] using blind signature schemes such as BBS+ can be used to enable selective disclosure of attributes, such signature schemes are not explicitly defined for the OIDC Aggregated Claims standard. For DISSENS, we have implemented⁵ our own C library that supports non-interactive zero knowledge proofs for this use case using pairings on a BLS12-381 curve [1]. With this extension, users can selectively disclose attributes from the credential without invalidating the issuer’s signature. Consequently, the system can selectively disclose only the required information to parties that support this extension.

3.3 Service providers

We expect that service providers offer their services through an HTTPS presence on the Internet and utilize an e-commerce platform which supports the use of OIDC identities. This implies that the service is set up with an OIDC relying party (RP) configuration. In particular, the OIDC RP component exposes an OIDC “redirect URI” which is used within a standard authorization code flow. In order to initiate such a flow, the service provider must first register its OIDC RP as a client at the OP service.

⁵ <https://github.com/Fraunhofer-AISEC/libpabc>

In re:claimID, an OIDC RP client is registered by creating a public-private key pair, which establishes a GNS identity that is henceforth identified by its public key. This public key also serves as the OIDC client ID. In order to register the HTTPS redirect URI(s) of its client, the service provider publishes this information in the client's GNS zone. This requires the private key of the identity. The resulting GNS records cryptographically bind the client ID to the client's redirect URI in a way that it is verifiable by the users' re:claimID instances.

From the perspective of the service, interacting with re:claimID is equivalent to interacting with any other OP with the exception that the OIDC endpoints are split between user and service: The client initiates an OIDC authorization request to the user with an authorization redirect to the user's local OIDC Authorization Endpoint. The user's re:claimID instance ensures that the provided redirect URI matches any of the URIs registered in GNS for the respective client ID. Upon authorization, the service receives user information by exchanging the user-provided authorization code at its local OIDC Token Endpoint. This exchange yields an ID Token as well as an access token for use at the local Userinfo Endpoint. The latter allows the service to retrieve fresh and up-to-date information directly from the GNS directory as the code contains the re:claimID "ticket" necessary for such a query [8].

Services that require certified attributes need to be configured with the public keys of certification bodies for the respective types of attributes. The establishment of trust in the third party IdP is out of scope of the OIDC standard. We assume that the relying party is configured a priori with a list of trusted institutions. This configuration would typically be a list of domains corresponding to the respective institution's IdPs. This allows the relying party to verify any ID tokens provided by the user through re:claimID as part of an Aggregated Claim, typically through the use of a JSON-Web-Key-Set (JWKS) provided by the OP. Alternatively, the relying party could be pre-configured with a key that can be used to verify the signatures of trusted OPs. Note that the method by which the relying party retrieves the key material is independent of the issue of trust: The key material is a means to verify the authenticity of the attestation. Trust establishment into the key material must be done out of band as part of a conscious selection process by the relying party. In order to facilitate trust establishment into a number of entities within a certain consortium or group, trust anchors provided through an X.509 PKI in combination with JWKS are viable.

4 Usability survey

We conducted a small usability survey on a pilot setup of DISSENS with seven participants to obtain insights into the usability of the system and derive directions for future improvement. The participants were between 18 and 54 years old and ranged from undergrad and grad students to professionals.

4.1 The pilot setup

We setup an OP service connected to the LDAP of the Bern University of Applied Sciences. This enables us to certify attributes based on real-world data about students and employees. The OP was configured with a client for the re:claimID web extension as elaborated in Section 3.1. Users were asked to authenticate using their existing credentials. The OP attested the information found within the directory. The resulting aggregated claims are then stored in the user's GNS zone from where they can be shared via re:claimID with Web sites.

We also created a pilot web shop based on the popular WooCommerce platform ⁶. As the platform is based on WordPress, we enabled a third party OIDC plugin and configured a re:claimID OIDC client as outlined in Section 3.3. The plugin initially did not support the aggregated claims feature of the OIDC standard, which is why we implemented this functionality ourselves for the use case and offered our patches upstream ⁷. We also implemented a WooCommerce payment extension that added the option to pay with the GNU Taler payment system.

As this survey was conducted during a COVID-19 lockdown, we deviated from our original plan to provide pre-configured workstations and instead opted for the distribution of a virtual machine with a minimal Debian buster image as a base. This image runs on GNU/Linux host systems using QEMU virtualization. In this image, the required GNUet peer-to-peer software is installed and automatically started. The GNUet REST service gives the host access to the GNUet REST API, allowing the re:claimID extension running in the browser to access the GNUet DHT. We also pre-installed the required plugins ("GNU Taler Wallet" and "re:claimID") in the browser of the virtual machine.

The advantage of this approach is a small and quickly downloadable VM image without a full desktop installation. The disadvantage is that testers needed a working GNU/Linux desktop system to run the setup, which substantially reduced the number of participants we could recruit.

The instructions given to the survey participants are reproduced in the appendix A.

4.2 Results

The results of our SUS questionnaire are promising. From the responses we calculated the following SUS scores, in ascending order: 62,5 (x2), 70 (x1), 77,5 (x2), 80 (x1) and 87,5 (x1). This results in a median score of 77,5 and an average of 73,9. SUS scores above 68 are considered above average ⁸.

⁶ <https://woocommerce.com>, accessed 2021/01/02.

⁷ <https://github.com/oidc-wp/openid-connect-generic/pull/255>, accessed 2021/2/12

⁸ <https://www.usability.gov/how-to-and-tools/methods/system-usability-scale.html>, 2021/02/12

Four participants gave qualitative feedback in the form of free text remarks. One noted that the interfaces were a bit “old fashioned”. Another remark noted that it was necessary to enter credentials on a domain foreign to the institution. This was due to our setup being deployed on an unrelated test domain and would likely not happen in a real world deployment. It shows the importance of domain name familiarity when deploying authentication services.

Another participant noted that it was not clear to them which attributes were required to be attested from the third party and which attributes were acceptable to provide “self-attested” along with the pseudonym. This is actually a shortcoming of the OIDC standard, as it does not offer a way for the relying party to convey this information in a claims request. A related issue was that it was unclear of the relying party would accept attributes issued by other IdPs than the one which was supposed to be used. Both issues could be resolved by having the relying party provide additional metadata along with the claims request. Hence, this may be an area for improvement of the standard.

5 Future work and conclusion

While the GNU Name System, re:claimID and GNU Taler have been subjected to extensive standardization work, the same cannot be said for the DHT. Existing RFCs on DHTs do not cover necessary features for GNS, such as in-network block validation and revocation flooding. Thus, additional standardization work is called for to avoid referring the GNUet source code serving as an inadequate substitute for protocol documentation.

Our usability survey provided test subjects with a pre-configured virtual machine image that included some heavy-handed workarounds for problems at the network layer. Today, most computers access the Internet from behind one or more layers of network address translation, possibly further restricted by firewalls and other obstacles. This limits the usable deployment of applications requiring end-to-end connectivity. Modern protocols like WebRTC/ICE work around these issues using centralized infrastructure providing STUN and TURN services to help consumers that experience connectivity problems. For properly decentralized networking, this centralized infrastructure should be replaced by peers that have found ways to traverse their NAT restrictions helping other peers do the same. This requires integrating STUN/TURN signaling into the peer-to-peer network layer.

Further interface improvement will have to be combined with work on the OIDC standard, as the idea of SSIs in combination with trusted attribute issuers are still in its infancy in that regard. While we did consider the use of the recent W3c DID standard [7] together with OIDC, its features do not directly address the shortcomings identified in the OIDC protocol.

In principle re:claimID could not only be used to support direct certification of attributes from trusted identity providers, but also hierarchies of certification bodies. For example, a country could certify its universities, and certified universities could then issue certificates for students. Similarly, a country may certify its regions, and then certified regions would

certify information about their inhabitants. GNS already supports the concept of SDSI-style “attribute-based delegation” (ABD) for identity attributes. Consequently, it is a natural extension for re:claimID to support such attributes [8]. However, the complexity of a decentralized resolution mechanism must still be evaluated in future studies.

We also note that this work proposes to use certification to realize features like age-restricted payments, but such certifications can be socially problematic as they could give too much power to certification authorities that might be better vested with legal guardians. It is conceivable to integrate age-restrictions into Taler by tagging coins with age restrictions upon withdrawal, effectively providing the power to impose such restrictions to the guardian of the bank account. An open challenge in this context is preserving such age restrictions when Taler renders unlinkable change.

Acknowledgments

The DISSENS project is funded by NGI_Trust program as project number 2.11.

References

- [1] S. Bowe. *BLS12-381: New zk-SNARK Elliptic Curve Construction*. <https://electriccoin.co/blog/new-snark-curve/>, accessed 2021/02/12. Mar. 2017.
- [2] J. Camenisch, M. Drijvers, and A. Lehmann. “Anonymous attestation using the strong diffie hellman assumption revisited”. In: *International Conference on Trust and Trustworthy Computing*. Springer. 2016, pp. 1–20.
- [3] D. Chaum, A. Fiat, and M. Naor. “Untraceable Electronic Cash”. In: *Advances in Cryptology — CRYPTO’ 88: Proceedings*. Ed. by S. Goldwasser. New York, NY: Springer New York, 1990, pp. 319–327.
- [4] D. Chaum, C. Grothoff, and T. Moser. “How to Issue a Central Bank Digital Currency”. Accepted, under pre-publication political review by Swiss National Bank directorate.
- [5] F. Dold. “The GNU Taler System: Practical and Provably Secure Electronic Payments”. PhD thesis. University of Rennes 1, 2019.
- [6] OpenID Foundation. *OpenID Specifications*. <http://openid.net/specs/>. 2021.
- [7] D. Reed, M. Sporny, D. Longley, C. Allen, R. Grant, and M. Sabadello. *Decentralized Identifiers (DIDs) v0.11*. <https://w3c-ccg.github.io/did-spec>. 2018.
- [8] M. Schanzenbach. “Towards Self-sovereign, Decentralized Personal Data Sharing and Identity Management”. Dissertation. München: Technische Universität München, Dec. 2020.
- [9] M. Schanzenbach, C. Grothoff, and B. Fix. *The GNU Name System*. <https://www.ietf.org/archive/id/draft-schanzen-gns-02.html>, accessed 2021/01/02. Oct. 2020.

A Usability survey participant instructions

Study: Privacy-friendly online shopping

Welcome to our user study! Please carefully read this introduction before following the instructions on the next pages to ensure that you have a high-level idea of what is going on.

In this study you will use the decentralized authentication service “re:claimID” in combination with the privacy-friendly payment system “GNU Taler” in order to do some online shopping.

“re:claimID” functions just like other popular so-called “Social Logins” such as “Login with Google” or “Login with Facebook”. It allows you to create profiles using nicknames (pseudonyms) to log in and log out. Your personal information is stored locally on your device and encrypted in a decentralized directory from which it can be shared with websites. Additionally, “re:claimID” allows you to *import* (certified) attributes from other identity providers, such as the BFH.

“GNU Taler” allows you to withdraw money from your bank account and store it as digital cash in your local wallet. The digital cash can then be used to purchase goods and services online. Unlike other payment systems, Taler respects your privacy and does not require any personally identifiable information for the payment process. For this usability study, you will of course not actually order some real product or pay with actual money. Instead, we will use a demonstration shop and a demonstration currency called “KUDOS”.

You will thus experience three features of the system:

- Import of certified attributes from the BFH LDAP into re:claimID.
- Use of the decentralized re:claimID IdP to share shipping address details with an online shop.
- Privacy-preserving payment using GNU Taler.

In this study we ask you to **create** a “re:claimID” pseudonym and **import** certified attributes into it from your BFH account. Further you will **withdraw** some digital cash from our fake bank and use your personal data as well as your digital cash to **buy** goods from our website using your pseudonym’s data to **provide** the shipping details.

Note that in reality, the **import** step would usually only be performed once, and the **withdraw** step only when the remaining wallet balance is too low.

Prerequisites:

- Linux Desktop OS able to run “qemu” VMs (tested using “Ubuntu 20.04.1 LTS”)
- Installed “qemu-system”, “qemu-system-x86” and “qemu-kvm”
- 6GB of free disk space and 8GB RAM

Please follow the steps below:

1. Download the virtual machine image from <https://nextcloud.bfh.science/index.php/s/dqHAES6eoiqo9oc>
2. Extract the VM image and the run script using "tar -xzf survey-setup.tgz"
3. Start the VM using the run script i.e. ". /run-survey-system.sh"
4. Choose option 1) “firefox” to start the preconfigured firefox web-browser inside the VM (X11 access on the host must be permitted!)
You have to provide the password according to the output of the script.
5. The browser should open on your desktop after a short time and present the page <https://bank.demo.taler.net/>. (The browser may be a bit slow due to the virtual machine setup with X-forwarding. That is an artifact of how we deploy the system for this test, not of reclaim:ID or Taler.)
 - a) Register an account (using any name you wish). You will receive 100 KUDOS as a starting balance.
 - b) Withdraw digital cash (20 KUDOS) from that account.
 - c) Your “Taler Wallet” should open and you have to “Accept fees and withdraw”.
 - d) Confirm your identity by solving a CAPTCHA (imagine having received an SMS with a TAN number).
 - e) Click on the Taler extension logo to check your wallets balance.
6. Go to the “wooshop” (second tab → “Shop”) or <https://woonslab.ch/shop>.
7. Put something in your cart (“In den Warenkorb”).
8. Open your cart (“Warenkorb anzeigen”).
9. Go to the checkout (“Weiter zur Kasse”). Do **not** yet fill in the form!
10. Instead, **click on “Login with Re:claimID”** to login with the decentralized Re:claimID identity provider. At this point you will be redirected to your local identity service:

-
- a) As you are a first-time user, you so far have no pseudonyms. **Click on “Add your first identity!”** to create a new pseudonym (if no “Add your first identity!” button appears, there is likely a network problem, try waiting a moment and/or go back to the checkout <https://woo.nslab.ch/kasse> and retry “Login with Re:claimID”).
 - b) **Enter a username of your choice** and click “Save”.
 - c) Before you can use the pseudonym, you need to fill in the attributes requested by the webpage. **Click on “Edit identity” and add your BFH email address** for “Email address”, e.g. “john@bfh.ch” or “john@students.bfh.ch”. Note that it must be a BFH email address, as only BFH currently allows Re:claimID to import attributes.
 - d) After the email attribute has been configured, you are given the option to import attributes from the BFH (LDAP Directory):
 - i. **Click on “Try import from Berner Fachhochschule”.**
 - ii. Log in with your BFH credentials and click “Authorize” to authorize the export of information to Re:claimID (that is, Free Software running in your own browser on your own system).
 - e) Certified attributes from the BFH should now appear in your profile. You may now add additional attributes, or ...
 - f) **Click on “Share information marked with ...”** at the bottom of the page to share the information from this pseudonym with the webshop.
11. If necessary, you may change or complete the imported attributes on the checkout page. Then, click pay to complete the transaction.
- a) Review the purchase in your Taler wallet and confirm the payment.
 - b) You may click on the Taler extension logo to check the change in your balance.

After completing the experiment, please complete a short survey:

<https://surveys.bfh.ch/index.php/XXXXXX?lang=en>

Thank you for your participation!