Towards a Trustworthy Digital Infrastructure for Core Identities and Personal Data Stores

(Extended Abstract)

Thomas Hardjono
MIT Consortium for
Kerberos & Internet Trust
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
Email: hardjono@mit.edu

Dazza Greenwood
MIT Media Lab and Civics.com
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
Email: dazza@media.mit.edu

Alex (Sandy) Pentland
MIT Media Lab
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
Email: sandy@media.mit.edu

***Abstract*—The growth of the future digital commerce on the Internet is dependent on the use of trustworthy digital identities. This paper identifies the need for a *Core Identity Infrastructure* and identify several phases and sub-phases for this infrastructure. At the heart of digital identities is the concept of the *core identity* of an individual, which inalienably belongs to that individual. The core identity serves as the root from which emerge other forms of digital derived identities (called *personas*) that are practically useful and are legally enforced in digital transactions. An individual must have the freedom to choose to deploy one or more digital personas on the Internet, each tailored to the specific aspect of that individual's life and each carrying varying degrees of legal enforceability as relevant to the usage context of that persona. This paper also outlines the role of the Core Identity Infrastructure to support the use of *Personal Data Stores* as defined by the *MIT OpenPDS* project. The ability for a user to own and control his or her personal data through deployment of a PDS represents a key requirement for the future of the digital commerce on Internet.**

Index terms: Digital identities, core identity, persona, identity federation, personal data store.

I. Core Identities & Persona: The Digital Identity Dilemma

The United States is in the midst of a fundamental transformation into a digital economy and information society, and whether the result of this transition will be successful or create new problems and vulnerabilities will depend significantly upon solving the digital identity dilemma. The United States has (wisely) rejected adoption of a national identity card or equivalent non-card-based national identity systems. However, lack of an interoperable and secure identity infrastructure is creating serious friction and in some cases harmful economic and legal distortions that are inhibiting the evolution toward a networked world.

A. An Identity Ecosystem

An identity ecosystem enshrining user-control and mass-market use is now emerging based on the culmination of economic, political and social forces. The National Strategy for Trusted Identities in Cyberspace (NSTIC) reflects this trend. However, as currently configured, existing business models, legal instruments and technical implementations are insufficient to support this type of identity ecosystem. This is because something is missing: an architecture for individual ownership of and primacy over one's own core identity. With such a core identity, it is possible for multiple aliases, accounts and attributes to be authenticated and authorized in a reliable, privacy enhancing and scalable manner. To this end, a core identity infrastructure provides a way for each person to own their single underlying root identity (i.e. the "Core Identity") and to bind several "persona" to that Core Identity without the need for other parties to access the Core Identity or be aware of any other persona [1]. With this approach, government issued identity credentials such as driver licenses, passports, professional licenses, birth certificates and the like (i.e. a "CivicID") can be leveraged through a persona designed to reflect such identifiers. In this way, a person can choose to connect a "CivicID" with, for example, a bank account, to more easily facilitate high risk or heavily regulated financial transactions, but to use a different persona that is not linked to a "CivicID" for purposes of, say, participating in an election discussion or a creative writing group.

One key inhibitor to digital identity has been the infeasibility of provisioning high assurance credentials at a mass-scale. The digital identity authentication solutions used in health care differ from those in banking and finance and from those in employment and workplace system and so on, through every sector and aspect of online life. While mass-scale biometric solutions are somewhat further in the future, a physical token (such as a phone or a smart-card) can be used much sooner to provide a second factor of authentication - if only an adequate system architecture and scalable approach were used to meet extant business, legal and technical requirements. The cost and inconvenience for users to have different higher assurance physical tokens for multiple systems has been a major block. The need is to allow a user to leverage one physical token tied to a Core Identity so that higher assurance authentication can be amortized across many downstream persona, accounts

and services, without necessarily giving up the Core Identity unique information to every requesting party or other counterparty. The Core Identity infrastructure, using widely accepted "claims- based" federated identity approaches, provides "authentication as a service". This means the cost of higher assurance authentication can be spread across many relying parties in many economic sectors and parts of society, while also creating a basis for truly user- centered, user-controlled and user-owned Core Identity to ensure privacy, civil liberties and individual autonomy.

The potential business models for Core Identity service providers and Persona providers (specializing in personalization, privacy and preferences services for a unified user experience across many sites and systems) are many fold. Among other models, using the Core Identity infrastructure, it is possible for a token holder to authorize a "Data Broker" to have access to the attributes and other transactional information related to many otherwise unconnected personae, so that the broker can provide a rich picture of the token holder to researchers, marketers and others willing to pay for access to such complete and consent based identity information and potential marketing or other offers. In addition, in the event of identity theft or fraud, the Core Identity infrastructure makes it quite easy for a token holder to quickly re-establish their ownership of a given Core Identity and to execute the revocation and re-issuance of any affected Persona, including the necessary updating of linking for all relying parties to the updated subsequent Persona. In this way, the cost to relying parties (e.g. merchants, service providers, etc) and the sometimes catastrophic costs to individual victims of identity theft and fraud can be largely avoided and contained. There are many Geo-location, aggregated demand, social and other business models made possible by token-based Core Identity as well, all supported by a common Core Identity infrastructure.

Finally, while the Core Identity architecture underlying the infrastructure is designed to prevent unauthorized linking of a Persona to a Core ID or linking of one Persona to another Persona of the same person absent that person's consent, the system does allow for a process by which a warrant from court of law may be honored and such linking would be authorized. However, the system does not permit mass-scale linking of this nature, and instead requires a warrant for each person for whom such information is sought. In this way, the needs for law enforcement, national security and the like can be met while preventing unconstitutional and illegal abuses such as have occurred many times over the past years and decades.

B. Supporting Pseudonyms and Anonyms

Identity is fundamental to the digital economy and more broadly to the information society that is emerging all around us. And yet, the architectural and design principles for identity are currently unsettled at best, and at worst are in a state of confusion or conflict. As the basic business, legal and technical architectures and infrastructure for online transactions and activities are being set via legislation, business models, technical standards and through basic social expectation and implicit agreement the role and shape of identity remains a wildcard. To make the transition to the online era, an identity system based on accepted principles is necessary.

Based upon history and a careful analysis of current design principles governing identity, it is possible to extrapolate some general boundary conditions, requirements and constraints for what an digital identity architecture and infrastructure would be. Specifically, there is a need in open and democratic societies for an identity system that supports the use of one or more pseudonyms in a wide range of interactions. For example, the copyright office specifically recognizes a pseudonym as a name that may file and own a copyright (see http://www.copyright.gov/fls/fl101.html) and there is a long history of Supreme Court cases, statutes and regulations that require acceptance of pseudonymity or anonymity (see: Anonymity and Encryption in Internet Commerce, chapter of the American Bar Association book Internet Law for the Business Lawyer, 2001, at: http://civics.com/wp-content/uploads/2012/01/Cryptonon.pdf). A general recitation of the basis and options for achieving anonymity and pseudonymity is available in the running bibliography at: http://www.freehaven.net/anonbib). However, this paper proposed an approach and architecture that is unique and, the authors believe, well tailored to the business, legal and technical needs of many stakeholders and tailored to form the basis of a broadly adoptable approach.

The Identity Commandments [1] published by the Jericho Forum establishes a number of fundamental principles regarding the creation, usage and management of digital identities in the Internet (see Appendix A). At the heart of these principles is the notion of the "core identity" and "persona" as key concepts that are necessary for any solution for a globally scalable identity ecosystem.

An identity ecosystem represents an additional infrastructure layer above the existing information technology infrastructure including the IP Internet infrastructure and the telecoms infrastructure. Today there is a general understanding of the importance of the IT and telecoms infrastructure as critical infrastructure not only for economy and commerce, but also for the survivability of the nation.

> *The concept of an identity infrastructure for the nation suggest that there will be both an important role for public sector leadership, investment and oversight as well as private sector innovation, implementation and user-facing services.*

The increase dependence today of citizens on the IT and telecoms infrastructure for their day-to-day activities points to the crucial need for an "identity infrastructure" that offers an ecosystem in which digital identities can be created, managed and destroyed in a practical manner. Such an identity ecosystem must support digital identities which maintain the privacy of the human person associated with the identity, and allows the human person to personalize their identity according to their needs [2], [3], [4].

C. Background: current state of digital identities

The recent NSTIC whitepaper [5] on the National Strategy for Trusted Identities in Cyberspace has recently provided a renewed drive for the development of an identity ecosystem for the US (and possibly internationally). The NSTIC strategy specifies four Guiding Principles to which the Identity Ecosystem must adhere:

- Identity solutions will be privacy-enhancing and voluntary.
- Identity solutions will be secure and resilient.
- Identity solutions will be interoperable.
- Identity solutions will be cost-effective and easy to use.

Specifically, the NSTIC strategy document calls for an identity ecosystem that will minimize the ability to link credential use among multiple service providers, thereby preventing them from developing a complete picture of an individual's activities online. Finally, service providers will request individuals credentials only when necessary for the transaction and then only as appropriate to the risk associated with the transaction. As a result, implementation of the FIPPs will protect individuals capacity to engage anonymously in cyberspace. Universal adoption of the FIPPs in the envisioned Identity Ecosystem will enable a variety of transactions, including anonymous, anonymous with validated attributes, pseudonymous, and uniquely identifiedwhile providing robust privacy protections that promote usability and trust." (emphasis added).

D. Goals of this paper

The goals of this paper are as follows:

- Propose a high-level blueprint as the basis of developing an architecture for a core identity infrastructure.
- Identify several clearly-defined phases for managing the creation, usage and archival of core identities within the infrastructure.
- Use each phase as a "black-box" within which to identify issues (Business, Legal & Technical) pertaining to the solutions that may implement that tasks/events occurring in that segment of the infrastructure.
- Propose the MIT OpenPDS model (see [4], [6], [7]) as the basis for personal data stores on the Internet, through which individuals can safely store and manage core identities and personas.

What this paper does not do:

- It does not provide a survey of identity technologies or standards.
- It does not introduce requirements, other than those put forward in [1] and those found in [8].
- It does not survey the area of privacy and privacy-preserving technologies.
- It does not propose any specific technologies (e.g. solutions, protocols) for implementing the technical requirement in each phase of the core identity infrastructure.

E. Terminology and Notation

In this paper we strive to re-use existing terminology in the space of digital identities. In particular, we have borrowed heavily technical terminology from Open Group's Jericho Forum [1] and borrowed legal terminology from the American Bar Association (ABA) Federated Identity Management Legal Task Force [8]. Additional technical terminology have also been borrowed from the OASIS Security Assertions Markup Language (SAML) [9], [10].

When discussing identifiers, we use the term *class* to denote the *degree of derivation* of a given identifier. For example, a Class 3 identifier means that it was derived from a Class 2 identifier. We will discuss one method for derivations below, while recognizing that other methods also exist.

The reader is directed to [11], [12], [13], [14] for more information regarding the legal foundations of digital signatures and identities.

II. Life-cycle of the Core Identity Infrastructure

In order to achieve a digital identity ecosystem with the properties defined in the Jericho Forum *Identity Commandments* [1], we believe that the first step is to develop an blueprint for an open architecture for the ecosystem that that supports existing products/services and which is based on the important notions of Core Identity and Personae.

A. Unlinkable Identities: Some Desired Properties

One of the key principles underlying the Identity Commandments is the notion that core identifiers must be protected (e.g. in storage and in-transit) to ensure their secrecy and integrity, and that a one-way linkage be used to connect a core identifier with the persona(s) that make use of the identifier.

In this paper we denote this "one-way linkage" through a more explicit process of "one-way obfuscation" (or simply"obfuscation"). In practical terms, this may mean applying the identity value as input into a strong cryptographic one-way hash function [15]. The aim of the one-way obfuscation is to prevent the receiver of a given identifier to deduce or mathematically derive the original identifier string that was input into the process. Other cryptographic approaches to unlinkability maybe also be used, but their details are outside the scope of this paper.

There are at least three benefits to using one-way obfuscation process from a "seed" identifier into a "derived" identifier:

- *Privacy-preserving*: By using a derived identifier in transactions, the owner of both identifiers is accorded with some degree of anonymity (albeit limited) and therefore privacy in their transactions. The use of one-way cryptographic obfuscation allows multiple (almost unlimited number of) Class N identifiers to be derived from a Class N-1 identifier.
- *Indirect Verifiability*: Given a derived identifier (call it Class N identifier), the "issuer" of the derived identifier can respond to queries from Relying Parties about the provenance of the derived identifier. In this case the issuer

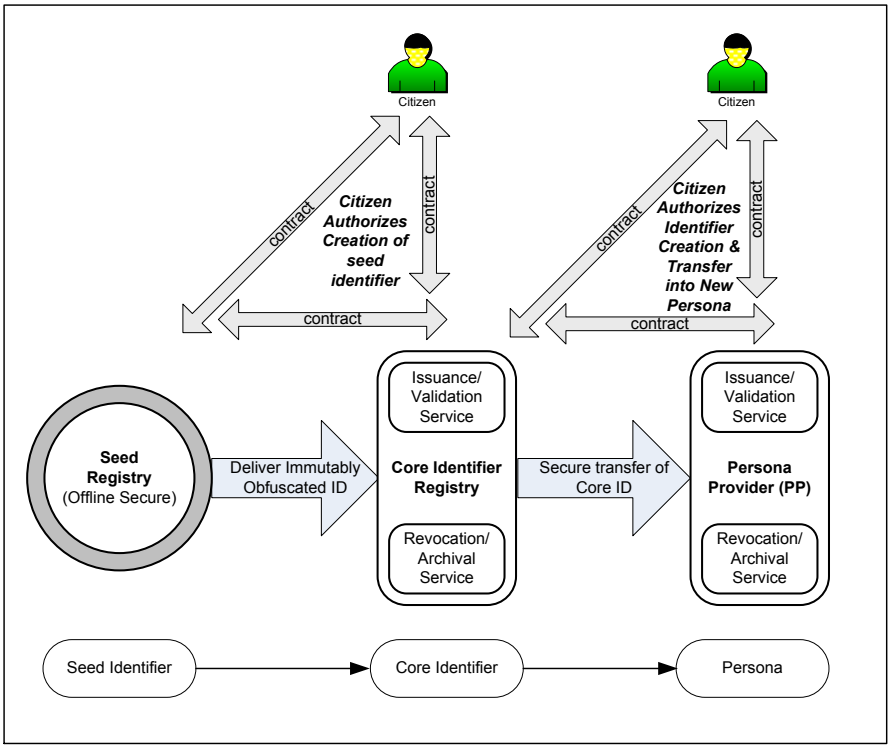

Fig. 1. Overview of the Core Identity and Persona Issuance Process

is the entity that performed the one-way obfuscation on behalf of the owner of the Class N-1 identifier. In other words, the Relying Party can always query the issuer of the derived identity about the validity of the derived identifier in a transaction. A derived identity (with verifiability) has the added benefit that the identifier can be used to establish verified anonymous (or pseudonymous) attributes.

- *Partial Recoverability*: In the case that the owner of a Class N identifier loses his/her identifier or that it was "stolen" (e.g. online identify theft) the owner can legally instruct the issuer of the Class N identifier to "revoke" the identifier. In this case a revocation of Class N identifier means that henceforth (from the given date/time of the revocation) the identifier is rendered legally invalid and that henceforth the owner ceases to bear legal responsibility for the identifier. Since a Class N identifier was derived from a Class N-1 identifier, the owner can either derive a new Class N identifier from that same Class N-1 identifier, or make use of other pre-existing (pre-manufactured) Class N identifiers.
- *Support for archiving*: When the owner of a Class N identifier decides to retire a persona (that was created based on that Class N identifier), the owner can archive the persona and deploy other (or generate other) identifiers. For example, the owner can use his/her Class N-1 identifier to generate a new Class N identifier that is distinct from the retired one. As such, the retirement and archiving of a given Class N identifier does not impede the owner from using other personas.

B. Core Identity and Personas

In this paper we distinguish between basic or *Base Personas* and *Principal Personas*. When a new unique core identifier is brought into use within the ecosystem, it is represented within the ecosystem within a basic (or plain) persona data structure. The term *base* persona is used to denote the fact that no externally-sourced attributes have been associated with the base persona. As such a base persona has very limited usefulness to a Relying Party.

- *Base Persona*: a digital data structure enveloping a core identifier, without any associated attributes.
- *Principal Persona*: a base persona associated with one or more attributes.

A base persona denotes the existential nature of the persona in the digital space (and by extension the human person in the real-world legally owning the base persona and the identifier embedded within). A useful analogy maybe a newly born baby (base persona) who exists in the real world, but who may yet have any attributes (name, social security number, email address) associated within it. Once some information or data is associated with the baby (e.g. registration in the government birth registry), the baby becomes more like a principal persona. As such, there is (a) the baby (existential) and (b) there are

some data associated with the baby. The combined (a) and (b) is what we refer to as a principal persona. When a base persona is later associated or linked to one or more attributes, it becomes a principal persona.

C. Pseudonymous Personas and Anonymous Personas

An important requirement for a core identity infrastructure is the support for pseudonyms (pseudonymous personas) and anonyms (anonymous personas) [16]. More specifically, it means that the infrastructure must allow the citizen to choose to create and deploy pseudonymous personas and anonymous personas, and allow these citizens to transact with these forms of personas as means to safeguard their privacy. The acceptance (or rejection) of these type of personas must then be decided by the Relying Party (i.e. the other party) in the transaction.

In order to ground in reality the notion of pseudonymous and anonymous personas, we provide the flowing broad definition of these forms of personas:

- *Pseudonymous personas*: a pseudonymous personae is one in which both its owner-principal and the issuing Persona Provider (a) have the capability to disclose the binding (link) between the pseudonymous personae and the core identifier used in the persona; and (b) have entered into a legal agreement to limit such disclosure to only parties authorized explicitly by the owner principal.
- *Anonymous personas*: an anonymous persona is one in which only the owner-principal has the capability to disclose the binding (link) between the pseudonymous personae and the core identifier.

Note that we have defined pseudonymous and anonymous personas in terms of requirements instead of technological implementations. In the last two decades, there have been a number of proposals for pseudonyms and anonyms based on cryptographic algorithms or protocols, notably in the area of electronic voting and in digital cash. Some solutions have been implemented, and have some have seen some limited deployment [17], [18], [19]

III. Phases in the Core Identity Infrastructure

In order to more easily enable the development of a core identity infrastructure, we have grouped the tasks into four main phases. The phases of the infrastructures are summarized as follows (see Figure 1). We believe that such an infrastructure would necessitate the market creating a number of new core identity service providers.

A. Phase 1: Seeding and Base Persona Creation

In this phase a *seed identifier* value is generated and is kept in a secure facility. The seed-identifier value becomes the basis for generating the core identifier belonging to a person (e.g. through one-way obfuscation). It is the core identifier (not the seed-identifier) that will be used to derive other identifiers, which in-turn will be used within deployed personas (see Figure 2).

The process of generating the seed identifier involves three parties, and must be governed by a legal trust framework that is specifically tailored for the seed generation process. The trust framework must be agreed to and be binding on all three parties.

The three parties are as follows:

1) *Human citizen*: The citizen must partake in and provide explicit consent to the creation of the seed identifier that will be legally owned by the citizen. In some architectures, the citizen may be the entity generating the seed.
2) *Authoritative Seed Issuer*: The role of the Issuer is to legally bind a citizen's seed value to the human person as a legal person. This entity must validate the legal identifiers of the citizen and validate supporting documents provided by the citizen. Note that an Issuer may offer additional services, including generating a seed on behalf of a citizen.
3) *Seed Registry*: The Seed Registry is the entity that is responsible for keeping the seed value of the citizen in a secure fashion. The Registry plays an important role when the citizen seeks to recover his core identifier (e.g. when lost or stolen) by way of re-starting the persona generation cycle. A citizen would need to provide proof of legal identity (e.g. birth certificate, passport, etc) to the Registry and the Issuer in order to generate new derived identities and new personas.

There are a number of possible architectures that can be used as the basis of addressing Phase 1, ranging from self-generated seed values to institution-issued seed values. The technological implementation of this phase is outside the scope of the current paper. A number of cryptographic approaches using zero-knowledge protocols can be used to implement the creation of the seed. However, there are a number of technical requirements of the seed and seed-generation process:

- *Randomness and global uniqueness*: The seed value must be globally unique (i.e. among the current future population of human persons on the planet). As such, we recommend the seed to be cryptographically random (using the best random-number generator known today). For example, the seed could be a 128-bit or 256-bit true (or near true) random number.
- *Authoritative trustworthy source*: The entity (or source) that computed the seed value (for a citizen) must be authoritative and legally recognized. This implies legal trust frameworks and accreditations underlying the operations of the Seed Issuer and Seed Registry.
- *Legal assignment and privacy*: Once generated, the seed must be legally designated to a human person as a legal owner. In the same way that property is owned by a human person, the seed is henceforth owned by the person [3]. However, this ownership is constrained in that the seed as a legal property is not transferable to other legal entities.
- *Secrecy of seed*: The seed must never be used directly or referred to directly, either within digital transactions or

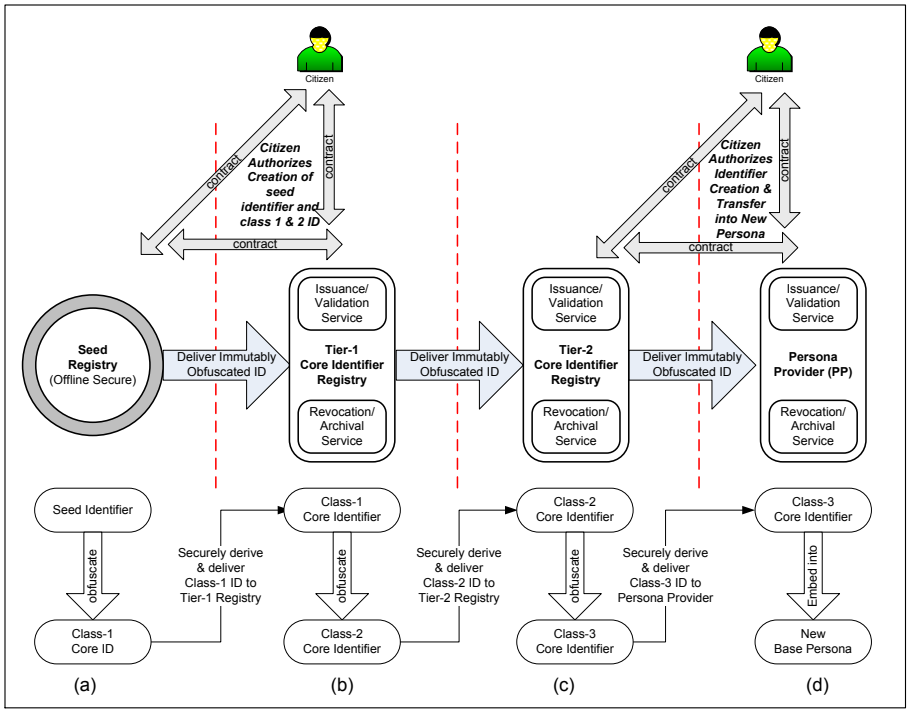

Fig. 2. Phase 1: Seed and Base Persona Creation using 3rd degree core-identifier

even within publicly accessible legal documents.

- *Secure storage (off-line)*: Due to the sensitivity of the seed, it must be kept off-line and in secure fashion. The level of protection given to a seed value must be at least equivalent to the protection given by Certificate Authorities (CA) to their own certificate-signing cryptographic keys [20], [21]. For example, the seed values could be stored off-line within tamper-evident Hardware Security Module (HSM) cards within physical vaults.
- *Privacy of binding*: The binding (legal and technical link) between a seed and its owner (human citizen) must be maintained as private information by entities involved in this transaction.

Figure 2 attempts to illustrate the process of deriving a core identifier from a seed. In Figure 2, an example of three (3) levels of derived identifiers are shown, where each level is denoted as *Class* of identifier. Thus, Figure 2 shows that a Class 3 identifier was cryptographically derived from a Class 2 identifier, which in-turn was derived from a Class 1 identifier (which itself was derived from the secret seed). It is the Class 3 identifier in Figure 2 that will be embedded (used) within the base persona data structure.

Although the idea of using cryptographically derived identifiers is hardly new [16], what is challenging today is implementing an infrastructure which can realize this process of identity derivations, whilst maintaining the secrecy and privacy requirements of the individual citizen. Efforts such as the ABC4Trust [22] project in the European Union represents promising progress.

B. Phase 2: Creation of Principal Personas

The primary aim of this phase of the infrastructure is to make new base personas available and to support the process of associating attributes to these new base personas. (See Figure 3). For this phase to be realized, we proposed the creation of a new category of core identity service providers, namely the *Personal Data Store Providers* (PDSP). We believe that a rudimentary version of the PDSP category exist today in the form of online cloud-based file storage providers (e.g. DropBox, iCloud, etc). However, we also believe that a more sophisticated design and architecture for personal data stores will be need to ensure the safety of identities and personas, and for the provable privacy of the user's personal data. We will discuss this topic further in Section IV.

The PDSP entity has a number of functions besides storing base personas:

- Secure online storage of Principal Personas, which are base personas associated with one or more attributes issued by an Attribute Provider.
- Secure storage of copies of attributes issued an Attribute Provider.
- Secure storage of personal data items belonging to the citizen (eg. medical records, tax records, etc).
- User-centric sharing of his or her resources with other users (or organizations) based on the user's consent [23], [24].

- Providing Validation Service End-Points to external entities, where they can validate any received Principal Personas. In the case of attribute validation, the PDSP may route queries (via back channels) to the issuing Attribute Provider.
- Archiving of personas (including Base Personas and Principal Personas) and of the attributes associated with these personas over time.

We believe the PDSP category of providers is valuable because it also provides the citizen with a engagement platform from which he or she can begin to request attributes from Attribute Providers (ATP).

A citizen may purchase (or obtain freely) validated attributes from an attribute provider. For example, the citizen Alice could request from her phone company that her phone number be associated (linked) to her base persona. The phone company as an attribute provider creates an attribute with clear references or pointers to the Alice's base persona structure. The phone company becomes the authoritative source of that attribute for the given persona belonging to the citizen.

The attribute provider also becomes the authoritative validation point of the data (i.e. phone number) contained in the attribute data structure.

C. Phase 3: Transactions using Personas

In this phase, the persona with its various attributes is used by its owner in transactions with a relying party (eg. online shop).

When transacting with a Relying Party, the citizen can then choose to deliver one or more persona attributes from her account at the PDSP. For example, if Alice needs to prove to a relying party that she is over 18 years old or that she possesses a valid credit-card, she can forward one or more signed attributes (together with the persona she is using) to the Relying Party.

When a Relying Party seeks to validate the claims contained in a persona attribute, the Relying Party must look at the *Attribute Provider Issuer* information (within the persona attribute), and query the Issuer by forwarding the persona attribute in question. When attribute provider receives a query about one or more of the attributes it has issued, the attribute provider must respond with an answer (e.g. "Valid", "No Longer Valid", "Revoked" or "Archived").

D. Phase 4: Retirement, Archiving and Digital Death

In this phase the owner of a persona and its associated attributes performs a retirement of his/her persona and archives the persona for future needs (e.g. future audit purposes).

The field of "digital death" today represents a new frontier of research for citizens on the Internet. Although a number of service providers are beginning to provide support for resource transferability in the situation of the owner's death, such support is still very rare currently. Furthermore, legislation is unclear or even non-existent in the case of digital assets belonging to an individual.

We believe a new category of service providers may emerge, whose main service would involve intelligent archiving of personal data in a manner that observes the last wishes of the individual and which can provide benefits to succeeding generation of citizens on the Internet.

IV. Core Identity and Personal Data: OpenPDS

A. Background on OpenPDS

The *OpenPDS* project at the MIT Media Lab is a groundbreaking project that seeks to provide consumers on the Internet with a dynamic personal data store (PDS) [4], [6], [7]. The OpenPDS is an open-source Personal Data Store (PDS) enabling the user to collect, store, and give access to their data while protecting their privacy. Users can install and operate their own PDS, or alternatively users can operate an OpenPDS instance in a hosted environment.

We use the term "dynamic" here to denote that fact that the PDS does not only contain static data but also incorporates the ability to perform computations based on policy and is user-managed or user-driven. In a sense, the OpenPDS can be considered a small and portable *Trusted Compute Unit* belonging to an individual.

Figure 4 illustrates one possible deployment mode of OpenPDS while figure 5 provides a high level architecture of OpenPDS in a PDSP-hosted scenario. The user (owner) of a PDS has established his/her PDS within a hosted environment at a *Personal Data Store Provider* (PDSP). The user remotely manages the PDS by selecting a number of sources of his/her personal data. These sources include his/her own mobile devices (e.g. GPS location data), various social networks, government data sources and other sources of personal data.

Since the user is the legal owner of his/her personal data, the user can deploy a user-centric access control management systems (such as the UMA protocol [23] based on OAuth2.0 tokens [24]).

As mentioned previously in Section III-B, the Core Identity Infrastructure identifies the need of a new breed of providers to operate/host individual personal data stores. We refer to such a provider as the *PDS Provider* (PDSP). This is shown in Figure 3. As mentioned above, the PDSP is expected to provide a number of functions to support the lifecycle of a user's PDS:

- Secure storage of identifiers, Principal Personas, Base Personas and Attributes of a user within the user's PDS.
- Secure storage of personal data belonging to the user (Figure 4).
- Support the import/export of a PDS in the cases where its owner wishes to relocate his/her PDS to another PDSP.
- Archiving of a PDS of the user, in the case that the user seeks to retire a given PDS or where the user becomes deceased.

B. Hosted OpenPDS: Some Requirements

There are a number of inherent features of the OpenPDS architecture that warrants careful technical design for its implementation. Many of these requirements emerge from

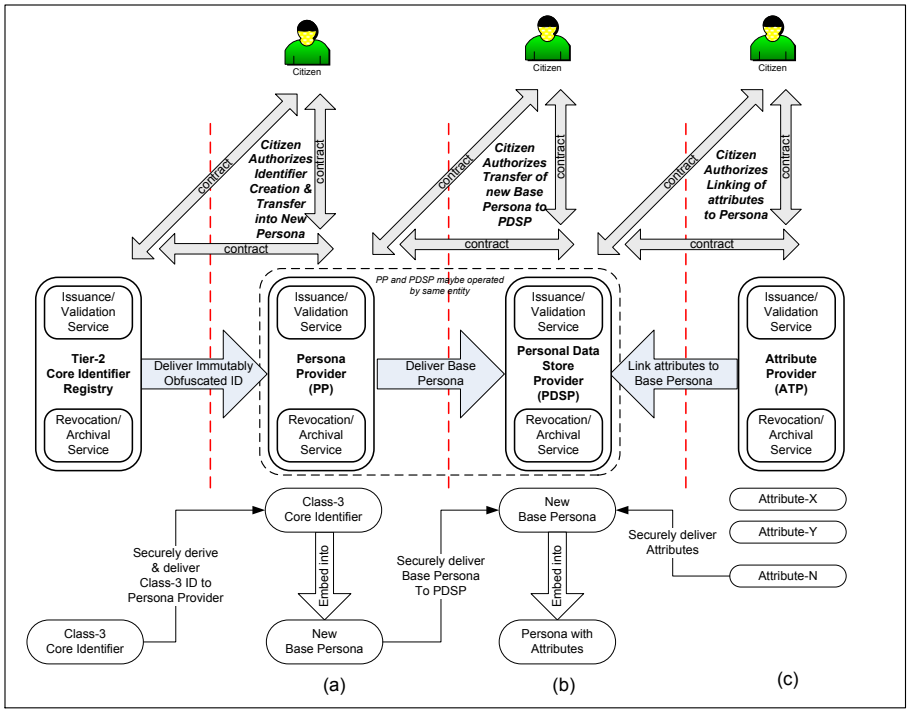

Fig. 3. Phase 2: Creation of Principal Persona using 3rd degree derived core-identifier

the fact that the PDS will hold core identifiers and personas belonging to the user (in addition to personal data items obtained from various sources).

In the following, we list a number of fundamental requirements for a provider-operated (hosted) PDS model, where a PDS Provider runs instances of OpenPDS within a virtualized computing environment. In this operational model it is useful to view the OpenPDS instance as a *virtualized resource container*, which embodies data processing and storage capabilities, as well as well-defined APIs by which the OpenPDS interacts with the virtualized environment supporting it. A distinct API set must also be defined for the OpenPDS to interact with external data sources and external data consumers/readers.

Some key security and privacy requirements are as follows (see [25], [26], [27]):

- *Unambiguous identification*:
  A given OpenPDS instance must be unambiguously identifiable. That is, there must be some means – both externally and from within the OpenPDS itself – to distinguish one OpenPDS instance from another. This requirement is key in that it is a precondition for the addressability of OpenPDS instances within a cloud provider's virtualization infrastructure and the addressability of OpenPDS instances on the Internet generally.

- *Operates unhindered*:
  A given OpenPDS instance must be able to operate unhindered. That is, the OpenPDS as a compute unit must be able to boot-up, execute and close-down without direct or indirect influence or interference by other processes that are co-located on adjacent virtualization stacks (i.e. multi-tenanted).

  The cloud provider's virtualization infrastructure is assumed (by definition) to serve the needs of its customer (namely the OpenPDS owner). We assume that the service level agreements (SLAs) and other contractual agreements will provide some degree of social and business trust between the OpenPDS owner and the cloud provider.

- *Truthful attestations*:
  Related to the previous requirements of unhindered operations is the need for the OpenPDS to be able to correctly report its internal status truthfully to its owner – who may be remotely running the OpenPDS instance. In other words, the OpenPDS must be able to produce signed *attestations* regarding its current status. Such reporting must occurred unhindered. These attestations must be continuously generated and logged by the OpenPDS as part of its security audit and tracking requirements.

There are a number of new and emerging trustworthy computing technologies that can be used to achieve the security requirements of the OpenPDS. For example, a hardware-based root-of-trust could be used (on the user side) to ensure that a given OpenPDS is bound cryptographically to the hardware

of the user (either to the device of the user or to a portable hardware dongle like a USB token). Widely available technologies like the *Trusted Platform Module* (TPM) hardware [26] can be used as the starting to point to achieve secure/trusted boot of the OpenPDS atop the cloud provider's virtualization stack. The use of the TPM chip on a USB token within cloud environments have been proposed in [27].

Cloud providers themselves who are committed to providing secure computing environments to their customers could employ tamper-resistant hardwares in their servers in order to provide a root-of-trust for the virtualization stacks in their platform. Prior research work in the use of hardware-based root-of-trust for a virtualization stack has been reported in [28].

V. Service Oriented PDS

Work is currently being undertaken at MIT in collaboration with Denmark Technical University to develop a standardized Open Architecture for individual identity and data sharing systems. The initiative features a generic model system comprised of a platform, a services layer and a standard interface for connection of external apps and services. All implementations of the model system include a user account system that is designed to interoperate with OpenID Connect, and a core system service ensure the provision of a PDS for every account holder upon enrollment. A basic design principle is that component integrate together through standard REST interfaces. Another basic design principle is that key roles, relationships, rights and obligations among participants in the system are supported and reflected in business, legal and technical integrated architecture and operating models.

There are two initial reference implementations of the Model System:

1) Sensible-Data: This is a designed to conduct computational social science research studies. An earlier version of the system, called Sensible-DTU, was the predecessor of the technical layer of MITs Model System. The next production deployment is a 1,000 Android phone study by the Denmark Technical University including the freshman class.
2) PublicEnterprise: This is designed to facilitate business creation and management services with US state government agencies. The initial pilot partner is the State of Kansas and the reference implementation is called the Kansas Business Center. A working life-cycle prototype of this system was demonstrated to the State of Kansas by MIT on April 24, 2013 and a larger pilot is expected to be conducted this summer.

The Sensible-Data and PublicEnterprise reference implementations of the MIT Model System are being actively iterated with partners in agile rounds of development and testing and feedback. A user-facing dashboard integrated with the account and federated identity services is a key feature needed to ensure meaningful user-centered control and management of identity and data sharing. Registration of services by the platform allows for the platform to provision and utilize a variety of services. The connectors to third-party applications such as Facebook, LinkedIn and Dropbox are designed to play simply and re-usably at both the service and the platform layers. In this way individual account holders can use the same dashboard to display and add, modify or revoke identity or data sharing grants of authorization. This facet of the integrated business, legal and technical codification is called "terms of authorization" and represent a simplified and combined contractual and technical grant of authorizations via the dashboard or at the transactional interfaces with services or apps. The reader is directed to `http://eCitizen.MIT.edu` for further details on this MIT Model System initiative.

VI. Conclusions

The goal of this paper has been to propose a high-level blueprint for a *Core Identity Infrastructure*. We believe such an infrastructure must have the ability to support identities and *Personas* (as defined in [1]) through its support of personal data stores (PDS) as defined by the MIT OpenPDS. Recognizing that planning and building such an infrastructure is not a trivial task, we have identified several phases within which core identifiers and personas can be created, managed, stored in a PDS and be archived or destroyed.

There are a number of key take-aways from this paper:

1) The need for a core identity infrastructure: We argue that an infrastructure to support the establishment and use of core identities and personas is needed in order to provide equitable access to data and resources on the Internet.

2) The need for privacy-preserving Personas: Personas are needed which are legally bound to core identifiers belonging to the individual. We see personas as a means to achieve individual privacy through the use of derived identifiers. These derived identifiers paves the way for a user to create anonymous and pseudonymous personas. Its also allows various attributes provider entities (ATP) to issue *Verified Anonymous Attributes*, which in-turn support the establishment of *Limited Liability Personas*. There exist already a number of strong cryptographic schemes and protocols to achieve anonymous and pseudonymous identifiers which can be the basis for personas.

3) The crucial role of the core identity infrastructure to support personal data stores: We argue that the privacy-preserving features of core identities and personas fully satisfy the data privacy requirements of *Personal Data Stores* as defined by the *MIT OpenPDS* project. The ability for an individual to own and control his or her personal data through deployment of a PDS represents a key requirement for the future of the digital commerce on Internet.

4) The economic need for PDS Providers (PDSP): Currently there are a number of storage providers that offer "data

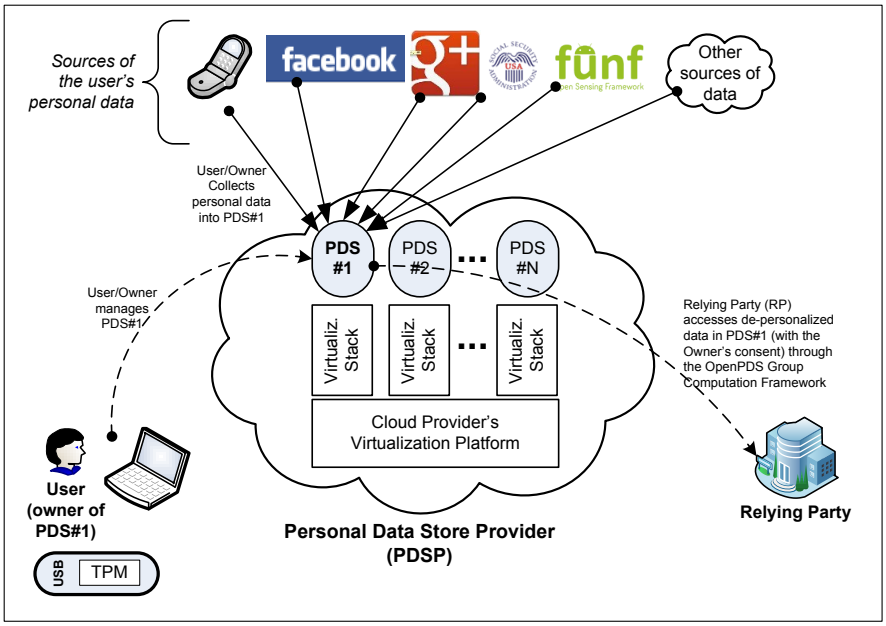

Fig. 4. OpenPDS Deployment Scenario

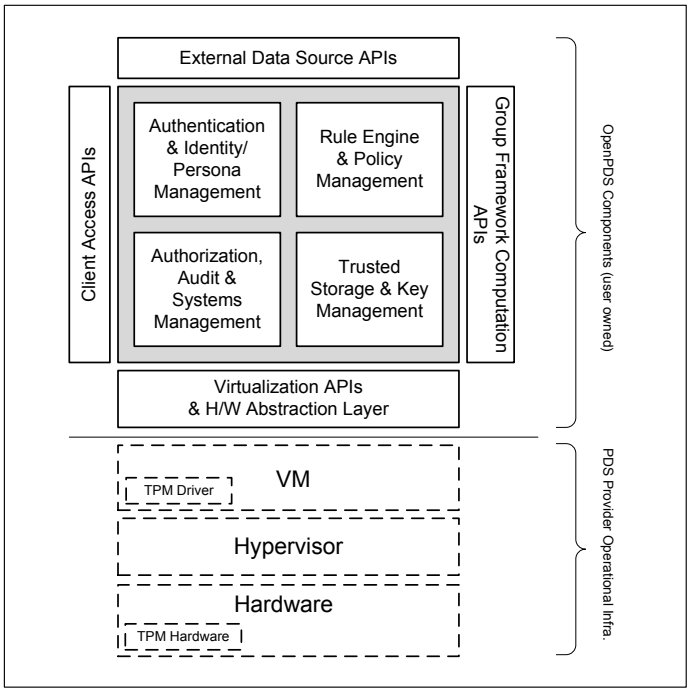

Fig. 5. High Level Architecture of OpenPDS Components

storage in the cloud" for individual consumers. Although these systems are very useful for day-to-day usage by the ordinary consumer, there remains the issue of privacy related to the mining of data residing within these cloud-based data stores. In the majority of cases, the terms of service are tilted against the consumer's interests. We believe the MIT OpenPDS design allows for a new breed of providers to emerge who will support consumer

privacy, while at the same time allow the consumer to optionally partake in various data mining and exploration schemes in a privacy-preserving manner.

5) The deliberate approach of the MIT Model System initiative is to provide a method for enterprises, government and other groups to use OpenPDS within existing business models, roles, relationships and does not require or assume any particular changes in business or legal arrangements. The method of provisioning a PDS to every account holder can be justified in the pilots and prototypes without distracting form the key purpose of the system. Similarly, the implementer of such a system can start using federated identity services in a way that is compatible with existing identity accounts and yet will catalyze independent business models and revenue opportunities over time. As Identity and PDS services start to emerge and survive, account holder would be able to point the identity account to the web address "end-point" where their identity and/or their PDS existed. A bridge from current to future potential models and systems is thereby established.

ACKNOWLEDGMENTS

We thank Stephen Buckley from the MIT Kerberos and Internet Trust (KIT) Consortium for his on-going support. We thank Henrik Sandell from the MIT Media Lab for the various technical insights regarding the OpenPDS open-source project at MIT. We are also grateful to John Clippinger from the ID3 Institute for Data Driven Design for his valuable insights into the potential crucial role that OpenPDS may play in the quantified-self Internet of the future.

REFERENCES

[1] The Jericho Forum, "Identity Commandments," The Open Group, 2011, available on www.opengroup.org.

[2] A. Pentland, "Reality Mining of Mobile Communications: Toward a New Deal on Data," in *The Global Information Technology Report 2008-2009: Mobility in a Networked World*, S. Dutta and I. Mia, Eds. World Economic Forum, 2009, pp. 75–80, available on http://hd.media.mit.edu/wef_globalit.pdf.

[3] World Economic Forum, "Personal Data: The Emergence of a New Asset Class," 2011, available on http://www.weforum.org/reports/personal-data-emergence-new-asset-class.

[4] Y. A. de Montjoye, S. S. Wang, and A. Pentland, "On the trusted use of large-scale personal data," *IEEE Data Eng. Bull.*, vol. 35, no. 4, pp. 5–8, 2012.

[5] The White House, "National Strategy for Trusted Identities in Cyberspace: Enhancing Online Choice, Efficiency, Security, and Privacy," The White House, April 2011, available on http://www.whitehouse.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf.

[6] Y. A. de Montjoye, J. Quoidbach, F. Robic, and A. Pentland, "Predicting personality using novel mobile phone-based metrics," in *Social Computing, Behavioral-Cultural Modeling and Prediction (LCNS Vol. 7812)*. Springer, 2013, pp. 48–55.

[7] Y. A. de Montjoye, E. Shmueli, S. Wang, and A. Pentland, "openPDS: Regaining ownership and privacy of personal data," 2013, (Submitted for publication).

[8] American Bar Association, "Overview of Identity Management," ABA Identity Management Legal Task Force, May 2012, available on http://meetings.abanet.org/ webupload/commupload/CL320041/relatedresources/ ABA-Submission-to-UNCITRAL.pdf.

[9] OASIS, "Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0," http://docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf, March 2005.

[10] ——, " Glossary for the OASIS Security Assertion Markup Language (SAML) V2.0," http://docs.oasis-open.org/security/saml/v2.0/saml-glossary-2.0-os.pdf, March 2005.

[11] L. Brazell, *Electronic Signatures and Identities Law and Regulation (2nd ed)*. London: Sweet and Maxwell, 2008.

[12] S. Mason, *Electronic Signatures in Law (3rd ed.)*. Cambridge University Press, 2012.

[13] OASIS, "OASIS Privacy Management Reference Model and Methodology (PMRM) version 1.0 (OASIS committee specification draft 01)," March 2012, available on http://docs.oasis-open.org/pmrm/PMRM/v1.0/csd01/PMRM-v1.0-csd01.pdf,.

[14] C. Reed, "What is a signature?" *Journal of Information, Law and Technology (JILT)*, vol. 3, 2003.

[15] J. Pieprzyk, T. Hardjono, and J. Seberry, *Fundamentals of Computer Security*. Springer-Verlag, 2003.

[16] D. L. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Communications of the ACM*, vol. 24, no. 2, pp. 84–88, February 1981.

[17] S. Brands, "Untraceable off-line cash in wallets with observers," in *CRYPTO'93 Proceedings of the 13th Annual International Cryptology*. Springer-Verlag, 1993, pp. 302–318.

[18] ——, *Rethinking Public Key Infrastructures And Digital Certificates*. MIT Press, 2000.

[19] J. Camenisch and E. Van Herreweghen, "Design and implementation of the Idemix anonymous credential system," in *Proceedings of the 9th ACM conference on Computer and communications security*. ACM, 2002, pp. 21–30.

[20] S. Farrell and R. Housley, "An Internet Attribute Certificate Profile for Authorization," RFC 3281 (Proposed Standard), Internet Engineering Task Force, Apr. 2002, obsoleted by RFC 5755. [Online]. Available: http://www.ietf.org/rfc/rfc3281.txt

[21] C. Adams and S. Farrell, "Internet X.509 Public Key Infrastructure Certificate Management Protocols," RFC 2510 (Proposed Standard), Internet Engineering Task Force, Mar. 1999, obsoleted by RFC 4210. [Online]. Available: http://www.ietf.org/rfc/rfc2510.txt

[22] ABC4Trust, "ABC4Trust: Attribute-based credentials for trust," available on https://abc4trust.eu.

[23] T. Hardjono, "User Managed Access (UMA) profile of OAuth 2.0," Internet Engineering Task Force, draft-hardjono-oauth-umacore-06, December 2012, work in progress.

[24] D. Hardt, "The OAuth 2.0 Authorization Framework," RFC 6749 (Proposed Standard), Internet Engineering Task Force, Oct. 2012. [Online]. Available: http://www.ietf.org/rfc/rfc6749.txt

[25] Trusted Computing Group, "Trusted Computing Group Home," Web page.

[26] ——, "TPM Main Specification," Web Page, 2011.

[27] J. Zic and T. Hardjono, "Towards a cloud-based integrity measurement service," *Journal of Cloud Computing: Advances, Systems and Applications*, February 2013.

[28] S. Berger, R. Caceres, K. A. Goldman, R. Perez, R. Sailer, and L. van Doorn, "vTPM: Virtualizing the Trusted Platform Module," in *Security 06: 15th USENIX Security Symposium*, Vancouver, Canada, July-Aug 2006, available on www.usenix.org.