Government Online

— A Credentica White Paper —

February 13, 2007

Credentica Inc.
www.credentica.com

About Credentica: Credentica provides innovative software products to protect identity-related information in Internet, mobile, and smart-card applications. Our products leverage patented cryptographic technology to provide multi-party security, scalability, and privacy. We work in collaboration with independent software vendors, systems integrators, and equipment manufacturers to ensure smooth product integration. We also provide specialized consultancy services. For more information, see www.credentica.com.

Abstract: Government Online, also known as e-government, refers to the electronic delivery of government services over the Internet and other computer networks. The current priorities for most governments are personalized online service delivery and electronic data sharing. Both require governments to roll out digital identity and access management solutions. The solutions of the leading vendors offer much of the functionality needed to implement Government Online. However, they fail to meet critical requirements with regard to security, privacy, availability, and autonomy. In this white paper we examine these unmet requirements and show how to meet them by using Credentica's innovative security technology.

1 Introduction

Government Online, also known as e-government, refers to the electronic delivery of government services over the Internet and other computer networks. Its primary objectives are to provide citizens and businesses with improved access to government services, to cut costs and improve productivity, and to improve participation in democratic processes.

Many governments around the world have already established an online presence by making information, forms, and non-personalized services available online. The current priorities for most governments are personalized online service delivery and electronic data sharing across government departments. The latter ability can eliminate the need for citizens and businesses to repeatedly provide the same data (such as address changes) and can also help government to detect entitlement frauds.

Both priorities require governments to roll out digital identity and access management solutions. The solutions of the leading vendors offer much of the functionality needed to implement Government Online. However, they fail to meet a number of critical requirements with regard to security, privacy, availability, and autonomy.

With regard to privacy, avoiding a unique identifier for each citizen is far from sufficient. Indeed, the identity and access management solutions of the leading vendors meet this requirement but, if adopted as is, would give government the ability to electronically link and trace all citizen actions in real time; this would have unprecedented repercussions for civil liberties and democracy.[1] Responsible governments must protect citizens even in the face of collusions involving corrupt insiders. Another motivation for this multi-party security requirement is that insider powers can also be abused by computer hackers, viruses, and malware that manage to gain insider status.

In Section 2 and Section 3, respectively, we examine the shortcomings in the context of Government Online of two otherwise excellent industry solutions: federated identity management and Microsoft's Windows CardSpace initiative. Following this we describe in Section 4 how to meet all of the unmet requirements by using Credentica's innovative security technology.

[1] Historical evidence shows that one of the first actions by oppressive governments is to resort to identification technologies to eliminate political dissidents, minorities, and other targets. See, for instance, "National Identification Systems: Essays in Opposition," Carl Watner and Wendy McElroy (editors), McFarland & Company, ISBN 0786415959, January 2004.

2 Federated Identity Management

With identity federation, services do not authenticate users themselves but delegate this step to trusted parties that have already established authenticated relations with these users.

2.1 Architecture overview

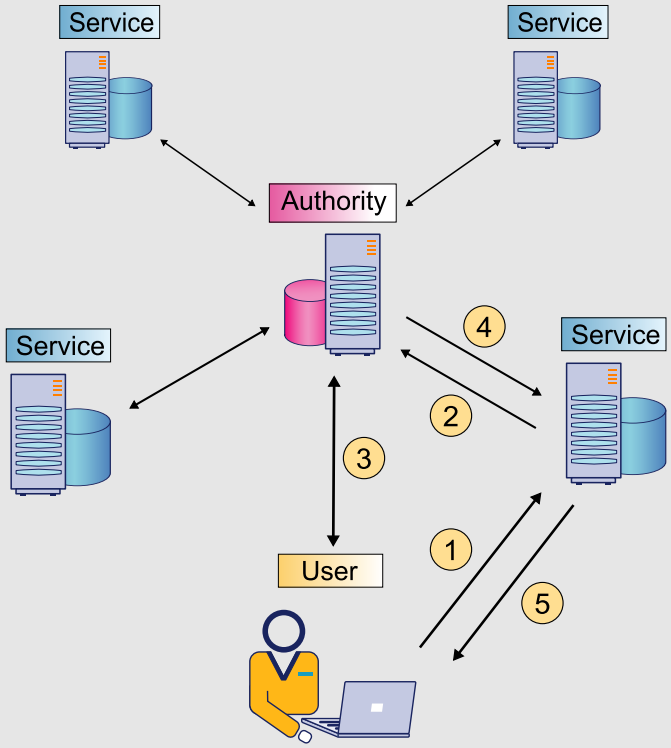

Figure 1: Federated identity

Upon receiving an online access request from a user, a service asks a central server (the "authority") if it has already authenticated the user in the current session. If not, the user is asked to authenticate to the authority by means of a password or a stronger authentication credential (which the user has established with the authority via an enrollment process). If authentication is successful, the authority provides an "authentication assertion" to the service, including a user identifier (or "handle") to enable the service to personalize its interaction with the user (e.g., giving the user access to account data or retrieving entitlement data to make an authorization decision). Since the user needs to authenticate only once to the central authority for the duration of each online session, the user enjoys a single sign-on experience regardless of how many services are visited.

The services and the authority need not reside within a proprietary network: they may reside in different organizational domains and the communications may take place over open networks such as the Internet. Conventional cryptographic techniques (such as Kerberos tickets or X.509 certificates) can be used for source authentication, integrity, and message encryption. All communications between a service and the authority can be directed through the requesting user (e.g., via HTTP redirection) or handled via Web service calls, and user session protection can be handled using SSL/TLS and a secure web browser.

The specifications of the Liberty Alliance, an industry consortium dedicated to federated identity, enable a user to have a different "alias" with each service; as long as the authority maintains a mapping of all of the user's aliases to their corresponding services, the above approach works just fine. There is no need for the user to register all aliases at enrollment time with the authority; federated identity enables the authority to generate and link a user's aliases on the fly by involving the user in the linking process as he or she goes about accessing services.

In much the same manner, a service could ask the authority to provide an "attribute assertion" pertaining to a user (i.e., account information held by the authority on the user or information derived from it). A service can also request attribute assertions pertaining to a user from another service that deals with the same user, even if the two services do not share a common user identifier; hereto the service submits its request to the authority, which knows the mappings of all of each user's aliases to their corresponding services and can therefore facilitate any data sharing.

2.2 Shortcomings for Government Online

Consider the application of federated identity management in the context of Government Online; the services are government services, users are citizens or businesses, and the authority is a government authentication service. The following shortcomings become apparent:

Surveillance capabilities Government would have the capability to electronically link and trace all citizen actions in real time; it would be trivial to centrally collect all electronic information streams from citizens to service providers. The resulting governmental surveillance capabilities would be unprecedented.

Impersonation Government would have the capability to impersonate targeted users. Consider, by way of a Canadian example, the implications of provincial government service providers relying on the federal revenue authority as the provider of authentication assertions about citizens: the federal government could access any account information at the provincial level of any targeted citizen by issuing authentication assertions in that citizen's name to its online investigators.

Systematic lock-out The authority can deny targeted citizens access to services by providing incorrect authentication assertions.

Wrong account linking The authority could cause assertion flows between accounts that do not belong to the same citizen; the authority could cause one user's criminal record indication to flow to another user's account to discredit that user or cause service denial.

The same capabilities are also highly problematic for government service providers, particularly in multi-jurisdictional settings. Service providers would be none the wiser; frauds originating from within the authority can be performed in an undetectable manner.

Carving out multiple service groups, each with their own central authority, does not help: it contravenes the very benefits that Government Online is seeking to attain. The only way to enable single sign-on and data sharing across the resulting silos is for the authorities link all their "master" user identifiers, which would only exacerbate the privacy and security problems.

Allowing citizens to be anonymous vis-à-vis the authority does not help: the issuance of strong authentication credentials by the authority becomes problematic, and the authority cannot verify user qualifications, cannot limit the number of identities a user can establish with each service,

cannot prevent unauthorized transfers by users of identity assertions about themselves, cannot deny fraudulent users access to services, and so forth. Even worse, the measure is ineffective, since a user's access requests remain linkable; thus, if a user's "real name" can be identified in any single access event (e.g., by a service) then all of that user's past and future actions become traceable.

In addition to these shortcomings the following concerns for governments and citizens alike arise:

Transferability Citizens could transfer (copies of) authentication and attribute assertions about themselves to other parties. Namely, the only thing that prevents a user from transferring assertions to other users is a verification that the same IP address is used to request and use the assertion, regardless of how well-protected the user's credentials for authenticating to the authority may be; all it takes for savvy users is to proxy the request for an assertion. While this may not be problematic for enterprise use cases, in the context of Government Online it would enable citizens to obtain services and credentials to which they are not entitled.

Denial-of-service and availability concerns Involving a central authority each time an authenticated interaction between a citizen and a government service must be established does not scale well, may be costly, and gives rise to availability and denial-of-service risks.

Identity theft risks An identity thief that manages to impersonate a citizen in its interactions with the authority can impersonate that citizen across all services. Today, the natural segmentation caused by the fact that most government services are their own "identity silos" prevents identity theft from spilling over to other domains.

In Section 4 we will explain how to overcome all of these shortcomings by using Credentica's innovative security technology.

3 Windows CardSpace

3.1 Architecture overview

Windows CardSpace is a client application that gives users access to an "identity selector." Similar to federated identity management, Windows CardSpace enables users to transfer assertions (called "identity claims") from "identity providers" to "relying parties." By building user control capabilities into the identity selector and providing client-side cryptography, Windows CardSpace improves the involvement and security of users in the sharing of information about them:

- The identity selector enables the user to select which assertions to present. Users can also select which identity provider to use to meet relying party requirements for assertion presentation, and can withhold the identities of relying parties from identity providers.
- Client-side cryptography enables the user to perform a cryptographic challenge–response protocol with relying parties when presenting assertions. Different cryptographic keys can be used for different interactions, to prevent unwanted linking capabilities. Keys are generated and managed in a secure subsystem of the operating system, for improved security against viruses and malware.

The most visible benefit of Windows CardSpace is improved protection against attacks by phishers and computer viruses on the authentication credentials of online users.

To encourage independent third-party implementations, Microsoft has opened up the Web services protocols needed to interact with Windows CardSpace. Microsoft is also enabling the simplistic

"your URL is your identifier" approach of the "lightweight identity" community (which is advancing sign-on and profile sharing protocols for blogging and social networking) to benefit from the security features of Windows CardSpace and to enable discovery of identity providers.

3.2 Shortcoming for Government Online

Prior to releasing Windows CardSpace, Microsoft promoted seven "laws of identity"[2] that "explain the successes and failures of digital identity systems." The first four laws are in essence privacy guidelines. They express the need to give users control over the release of information about them, to minimize the disclosure of identity information, to avoid unnecessary involvement of additional parties in interactions between users and relying parties, and to avoid unwanted linking and tracing capabilities. Windows CardSpace meets these privacy guidelines only in a weak threat model. While this may be appropriate for many consumer applications, responsible governments seeking to implement Government Online must address a much stronger threat model.

The shortcomings of Windows CardSpace in the context of Government Online are, in fact, almost identical to those of federated identity:

Impersonation, lock-out, and availability concerns Assertions that are issued by identity providers (as opposed to self-generated assertions) cannot be stored locally for future use: users must retrieve all requested assertions on demand from their identity providers when interacting with relying parties. This gives rise to the same security, scalability, and availability concerns as those described in Section 2.

Surveillance capabilities Identity providers know when users are interacting with relying parties, what assertions they are asked to present, and when they present them. Furthermore, in collusion with relying parties it is trivial to trace all assertion presentations to their issuance (either by comparing issuing and presentation times or by linking the provider's signatures on the assertions). In effect, each time a user uses Windows CardSpace to sign on to an account or to transfer identity-related information, the user is unwittingly linking ("federating") his accounts at the provider and the relying party.

Transferability Users can transfer (copies of) assertions about themselves to other parties. This is more difficult to accomplish if assertions are bound to keys that reside in a secure subsystem of the operating system that cannot be accessed by users themselves. However, this approach raises new privacy concerns that are very similar to those that are holding back the adoption of Trusted Computing. Namely, if a secure subsystem is relied on as a black box with regard to its own computer user, then privacy claims (such as the claim that cryptographic keys are randomly generated and are unknown to outside parties) cannot be publicly verified.

No selective disclosure Users cannot selectively disclose attribute information contained in assertions that have been issued to them; they can merely choose between disclosing or not disclosing the entire contents of a particular assertion. A CardSpace assertion will typically contain multiple attribute statements.

To protect the privacy of citizens participating in a society-wide identity and access management infrastructure, it is not sufficient for data subjects to be a "choke point" for the flow of identity data about them. At its worst, user-centrism does nothing for data subjects but greatly extend the reach of unintended cross-domain sharing of identity data about them, resulting in the creation of a

[2] See http://www.identityblog.com/stories/2004/12/09/thelaws.html.

common identifier with each and every user-centric data transfer; in this scenario, the data subject is in essence contributing to "super-federation." Once previously unlinked accounts are linked, the data subject is powerless: from here on, organizations can reliably exchange data about the user directly among themselves.

4 Credentica's Technology

In this section we describe how to simultaneously meet all the security, privacy, availability, and autonomy requirements of Government Online by using Credentica's innovative security technology.

4.1 Overview of ID Tokens

At the heart of the answer is our ID Token™ technology. An ID Token is a protected assertion that is issued to a user (or its agent) and subsequently (possibly at a later time) presented to a relying party. An ID Token can contain any kind of attribute information that is bound to a key pair. ID Tokens cannot be forged or modified, cannot be stolen through eavesdropping or phishing, and cannot be replayed by legitimate verifiers. They can be revoked prior to expiration and can be either software-only or hardware-bound (for improved security). Furthermore, relying parties can capture user-authenticated transcripts that prove their interactions with ID Token users. The power of ID Tokens goes far beyond PKI certificates and other conventional authentication technologies:

- Issuers can protect ID Tokens through software-only DRM-like techniques against unauthorized uses by their own users (such as transferring, discarding, and reuse).
- Issuers can cryptographically bind ID Tokens to trusted modules (such as smart cards or Trusted Computing chips) that can enforce third-party security policies throughout the entire life cycle of the ID Tokens. A single low-cost device can protect arbitrarily many ID Tokens.

What makes ID Tokens truly unique is that they provide these powerful multi-party security features while protecting, by design, the privacy and autonomy interests of users and relying parties:

- Attribute information contained in one or more ID Tokens can be selectively disclosed in response to unanticipated requests from relying parties. Users can also prove that attribute information in their ID Tokens is not blacklisted without disclosing the information itself.
- In contrast to conventional authentication technologies, the use of an ID Token does not leak any information that others could exploit to link or trace user activities; the degree of traceability or linkability is determined solely by the attribute information that users disclose.
- Transcripts are as unlinkable and untraceable as the ID Tokens from which they originate, and user-disclosed information can be censored from them without destroying their verifiability.
- Trusted modules cannot be programmed in a manner that enables third parties (including issuers and verifiers in collusion) to bypass any of the privacy properties of ID Tokens.

These privacy properties hold unconditionally, in the strongest imaginable threat model: issuers and relying parties cannot learn even a single bit of information beyond what can be inferred from user-disclosed attribute information at presentation time. This holds even if they collude from the outset (indeed, they may be the same entity) and have unlimited computing resources at their disposal in a coordinated attempt to deviate from ID Token protocols and analyze the resulting protocol data flows. In other words, a user's privacy does not depend on any cryptographic infeasibility

assumptions nor does it require trust in the honest behavior of any other participant; each user need merely trust the proper behavior of his or her own computing device. Trustworthy operating systems, open code, code certification, open market availability, and anti-virus and anti-spyware software can all contribute to the trust that users can place in their own devices.

By leveraging the privacy properties of ID Tokens, applications can support the full privacy spectrum from unconditional anonymity to persistent pseudonymity to omni-directional identification.

We will now show how to apply this technology to meet the requirements of Government Online.

4.2 Secure single sign-on

This section explains how a citizen, Alice, can securely authenticate to government services online in such a manner that:

- Alice enjoys a single sign-on experience, regardless of how many services she accesses;
- Government services can authenticate Alice without needing to involve any other party;
- Phishing attacks and replay attacks by service providers are ruled out;
- No linking or tracing powers are created, neither with regard to timing analysis nor with regard to data flow analysis; and
- Alice can be prevented from transferring (copies of) assertions about herself.

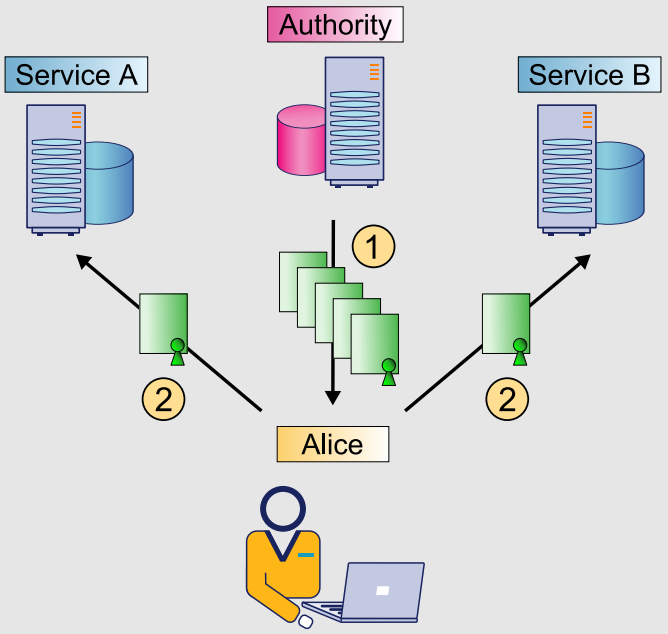

Figure 2: Hooking ID Tokens up to accounts

In an enrollment phase, Alice's computer obtains several ID Tokens from a government authority. The authority protects the ID Tokens against any unauthorized manipulations. Alice's computer sanitizes each ID Token by randomizing any information that would otherwise give rise to unwanted linking and tracing powers by service providers and the authority. The result of this sanitization

is that the authority never gets to see the actual ID Tokens it issues to Alice; from a privacy perspective, each of Alice's ID Tokens is the equivalent of a randomly self-generated number.

When Alice subsequently accesses a government service for the first time, either in the same session or at a later time, her computer transmits a fresh ID Token to the service provider. Alice's computer uses a different ID Token with each government service, and maintains a mapping of all of her ID Tokens to their corresponding services. Each service provider associates the ID Token it receives from Alice with its account information on her. If a services has already established a legacy account relation with Alice, then she must authenticate to the service one last time under the legacy ("silo") authentication method.

Because Alice's ID Tokens are the equivalent of randomly self-generated numbers, they are unlinkable and untraceable in and of themselves. Consequently, service A, service B, and the authority do not gain any linking, tracing, and profiling powers over Alice. This privacy guarantee holds even if the service providers and the authority were to actively collude from the outset. Note that this authentication approach does not make government services anonymous or pseudonymous where they previously were not; instead, it prevents unwanted new correlation powers.

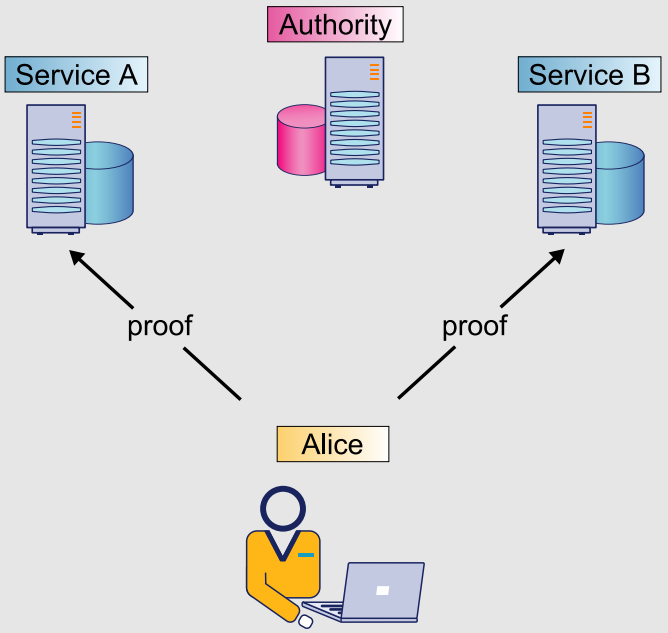

Figure 3: Single sign-on

In subsequent visits to a government service, Alice's computer authenticates using the ID Token that the service provider associated with her account. To this end, Alice's computer generates a cryptographic proof-of-knowledge of a private key corresponding to the ID Token. The private key itself is never revealed and cannot be learned even by malicious verifiers; this ensures that the authentication process cannot be phished and cannot be replayed by compromised service providers. The service provider can verify Alice's proof without consulting the authority or any other party.

To thwart content manipulation attacks subsequent to signing on using an ID Token, Alice can use the same ID Token to digitally sign instructions and other interaction-related data. (Alice can also use her ID Tokens non-interactively to sign documents, forms, and any other kind of data.) This

enables her to dispute fraudulent account operations by insiders of the service provider as well as fraudulent access to her account by authority insiders; they will not be able to come up with the corresponding signed proofs.

To protect against local attacks on Alice's ID Tokens by viruses and malware, the authority can electronically bind her ID Tokens at issuing time to a previously issued trusted module. This module may take the form of a tamper-resistant computing device (such as a smart card or a USB key with a CPU), a tamper-resistant chip (such as a Trusted Computing chip), a software-only emulation thereof (such as code running in a secure subsystem of the operating system or a "virtual" smart card that relies on code obfuscation techniques), or an online service out of Alice's reach (which could play the same role for many users). A resource-constrained module suffices to achieve high performance for any number of ID Tokens: all computationally expensive operations, in particular modular exponentiations, can be securely offloaded. Furthermore, any ID Token-related information (other than a portion of the private key common to all of Alice's ID Tokens) is stored and managed by Alice's own computer by default. At the same time, trusted modules cannot be programmed in a manner that enables third parties (including insiders collusions) to bypass any of the privacy properties of ID Tokens; in particular, trusted modules cannot leak any information through their responses. Furthermore, unless enabled by Alice, her trusted module cannot learn any attribute information in the ID Tokens it helps to protect and cannot learn any interaction-related data between Alice and the outside world.

At issuing time, the authority can optionally enrich Alice's ID Tokens with attribute information (e.g., her province of residence) to enable government services to make more informed decisions. At presentation time, Alice's computer can selectively hide any irrelevant attribute information that the ID Token may contain. This capability can also be exploited by the authority (with Alice's awareness and cooperation) to encode a unique number into all of Alice's ID Tokens; this number will in effect be invisible since Alice will never disclose it when using her ID Tokens, but it invisibly links all of Alice's accounts. This invisible linkage can be leveraged by services for non-transferable data sharing, as we will explain shortly.

To prevent Alice from transferring the ID Tokens she hooks up to her accounts, they can be bound to a trusted module. Alternatively, the authority can discourage Alice from transferring ID Tokens by encoding attribute information into them that is confidential to Alice, such as her credit card number, a password, or a secret key; it is impossible to use an ID Token without knowing all of its attribute contents, even if these contents are not disclosed at presentation time.

4.3 Data sharing between unlinked accounts

We now explain how government services that do not know Alice under a common identifier can securely share information on Alice, in such a manner that:

- Government services can accept assertions without needing to involve any other party;
- Phishing attacks and replay attacks by service providers are ruled out;
- No unnecessary linking or tracing powers are created, neither with regard to timing analysis nor with regard to data flow analysis;
- Alice can be prevented from transferring (copies of) assertions about her, without having to bind these assertions to trusted modules; and

- Alice is involved and given partial control over the release of information about her, and has the ability to selectively disclose attribute information in a highly granular manner.

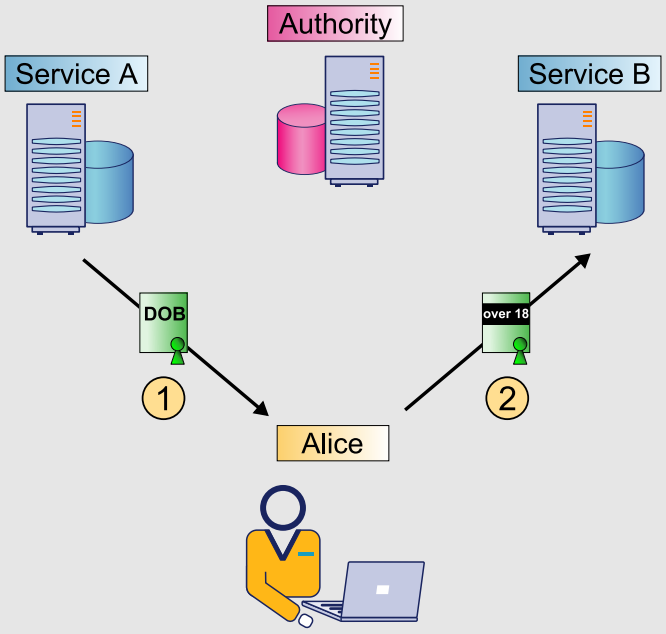

Figure 4: Data sharing

Service A makes an "identity assertion" about Alice based on its account information on her; for example, service A may make an assertion regarding Alice's residency and birth date. Instead of sending the information directly to service B (which would not know to which account it pertains), service A issues it in the form of an ID Token to Alice, protecting it against unauthorized manipulations. Alice's computer sanitizes the ID Token by randomizing any information that would otherwise give rise to unwanted linking and tracing powers.

When Alice subsequently accesses service B, either in the same session or at a later time, her computer sends the ID Token in a manner that selectively discloses only the minimal assertion information needed. For example, if service A has asserted Alice's birth date, Alice can disclose to service B only the fact that she is over 18 years of age, without revealing even a single bit of extra information. Service A and service B (even when they collude from the outset with the authority) do not gain any linking, tracing, and profiling powers over Alice beyond what they can learn from the disclosed assertion itself; in the example, all that they can learn is that the individual accessing service B must be one of all the adult users of service A. Service B can verify on its own the origin of the ID Token and its integrity.

To prevent correlations via timing analysis, Alice simply obtains multiple unlinkable "copies" of the same ID Token from service A.

To prevent Alice from transferring ID Tokens, either the assertions are cryptographically bound to a trusted module or their issuers encode transferring disincentives into them. A third method is available that does not require the binding of ID Tokens to a trusted module and that offers much stronger protection than the encoding disincentives. Namely, if the authority has encoded into all of Alice's ID Tokens a unique number (which Alice never discloses), then service providers

can leverage the fact that all of Alice's accounts are invisibly connected via this number. Namely, a service provider can issue ID Tokens in such a manner that they will contain the invisible number in Alice's ID Token that has been hooked up to her account; Alice, in turn, can prove that the ID Tokens she presented contain the same invisible number as that which has been encoded into the ID Token that has been hooked up to the destination account.

4.4 Other benefits

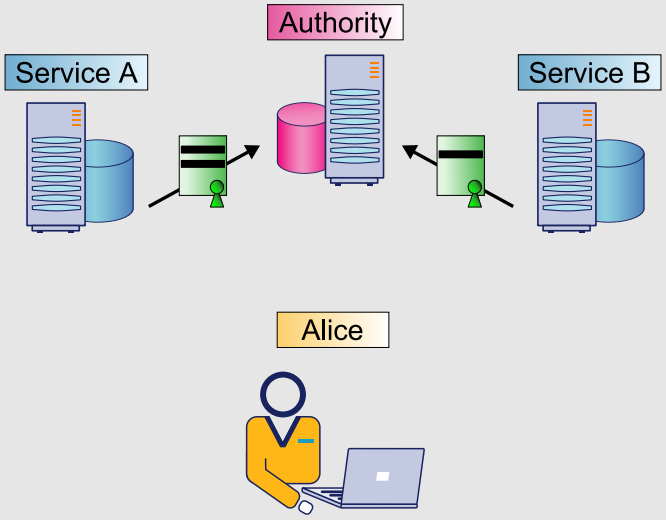

Figure 5: User-authenticated audit trails

For improved accountability, government services can collect non-repudiable audit trails that are not privacy-invasive. Namely, whenever Alice presents an ID Token a so-called presentation transcript can be captured. Presentation transcripts are user-authenticated; in effect, they are signed with the private key that corresponds to the ID Token. At the same time, transcripts are as unlinkable and untraceable as the ID Tokens from which they originate, and user-disclosed attribute information can be removed from them without destroying their verifiability. This latter censoring capability enables service providers to protect their own autonomy, security, and privacy interests in cases where transcripts need to be forwarded to auditors. In case of a dispute, censored data can be uncensored by the service provider as well as by Alice.

Another benefit of the ID Token technology is that it enables a group of government services to deny access to Alice in case she abuses any one service in the group, even though they do not know Alice under a common identifier. Suppose that Alice commits a fraud at service A and it should be possible to deny her access to service B. Service A provides service B with a unique random revocation number that is associated with the ID Token that Alice uses to authenticate to service A. Service B subsequently asks each and every access requestor to submit a cryptographic proof that the revocation number is not identical to attribute information that the authority has encoded into their ID Tokens at issuing time. This proof can be created only if the statement is true. This cross-domain revocation capability does not impinge on user privacy: cryptographic proofs provided by non-revoked users reveal no linking or tracing information, users must cooperate at issuing time in to encode revocation numbers into their ID Tokens, and blacklists of revocation numbers are visible to all users.

5 Conclusion

It is a formidable challenge for governments to reconcile a vertical privacy structure with a horizontal service delivery imperative. While adopting off-the-shelf identity and access management solutions may be appealing for economic or time-to-market reasons, doing so would be highly inappropriate: it would give government the power to conduct real-time mass surveillance, to impersonate individuals wherever they go, and to systematically lock individuals out of services. In order to move forward with Government Online without sliding down a slippery slope toward a digital dictatorship, it is important for government to adopt technological innovations that provide multi-party security and privacy by design.

Credentica has implemented a Government Online prototype on PCs and mobile phones that demonstrates many of the features described in Section 4; please contact us if you would like to arrange for a showcase.