Identity and privacy in the digital age

Robin Wilton

Federated Identity
Chief Technology Office
c/o Sun Microsystems
Guillemont Park, Camberley, Surrey
GU17 9QG, UK
E-mail: robin.wilton@sun.com

Abstract: Organisations and data subjects alike are all facing up to the new challenges of online identity and privacy. For all the activity in this field, there is still much confusion and argument over basic terms and concepts. Technologists, systems architects, business people and policy-makers, for example, all describe privacy in different and sometimes incompatible terms and yet, for any given system to work, the views of all these different stakeholders must be successfully translated into technical implementation and coherent policy enforcement.

This paper sets out a simple data model to describe the different types of identity data and relates this to typical system architectures and then to privacy as an information systems policy objective. The aim is to provide a consistent framework for thinking about identity data so we can understand how privacy arises from the correct interaction of data and policy.

Keywords: privacy; digital identity; data model; credentials; consent; disclosure; privacy policy.

Reference to this paper should be made as follows: Wilton, R. (xxxx) ‘Identity and privacy in the digital age’, *Int. J. Intellectual Property Management*, Vol. X, No. Y, pp.000–000.

Biographical notes: Robin Wilton is a Corporate Architect for Federated Identity at Sun Microsystems. He also represents Sun Microsystems as the Co-chair of the Public Policy Expert Group (PPEG) of the Liberty Alliance. He has worked in the IT industry for over 22 years in systems engineering, technical support, consulting and product programme management. He is, essentially, a ‘translator’ between the technology vendor and its customer communities – including commercial and public sector customers, policy-makers, standards bodies and academia. He is a frequent speaker at conferences; his recent engagements including the OECD Workshop on Privacy and Digital Identity (Trondheim), the NetID Conference 2007 (Berlin) and the Digital Identity Forum 2007 and lectures to the Royal Holloway College MSc (Computer Security) stream and the Graduate School of Ethics, Hokkaido University. In 2007–2008, he ran a worldwide series of Privacy Summits for the Liberty PPEG; the programme has run sessions in Berlin, Brussels and Washington, DC, with further events planned in London and Basel in 2008.

1 Introduction: what is 'identity' in the context of 'digital identity'?

Questions of personal identity have exercised philosophical minds for centuries: 'Is the person who wakes up in the morning the same person as went to sleep the previous evening?'; if so, does the same hold true for someone who does not fall asleep, but lapses into a coma or suffers loss of memory? We will come back to one philosophical definition of identity in a moment, but for the time being, let us consider whether the concept of 'identity' is any clearer in the digital world.

Most of us who interact with computers have, more or less consciously, come to assume that there is such a thing as a 'digital identity', and that we probably have one or more of them. For instance, anyone who has an e-mail address has, in principle, a consistent way to ensure that e-mails reach them and not some other person. Anyone who logs on to a server or application probably has a user-ID and password which ensures that they identify themselves before gaining access. Someone who has a blog or a personal website may well expose enough information about themselves to provide personal identification.

So it certainly seems as though we have 'digital identities'... and yet each of the examples above might actually be a lot weaker, as an indication of 'identity', than it seems. To see why, let us briefly go back to the 1600s and Leibniz. He defined identity in terms of whether one thing can be distinguished from another; if object A shares absolutely every characteristic of object B, including its shape, extent, position in time and space, then A and B are identical: they have the relationship of identity.

This may seem rather abstruse, but is a very useful principle when considering online identity systems, particularly because of their heavy reliance on credentials. When a credential is presented in support of an authentication request, it is essentially being used as evidence that the person presenting the credential now 'has the relationship of identity with' the person to whom it was issued at some point in the past. I use the term 'person' here in order to maintain the link between this and the notion of personal identity... but clearly this model of identity can be applied to any entity to which credentials have been assigned (such as a software application, a device or a web service).

So one definition of 'identity' in the online world is 'the relationship of identity between a person at enrolment time, and a person at authentication time'.

By this definition, identity is something other than a 'snapshot' of who someone is: it is part of a chain of events from enrolment and credential issue through to credential presentation and authentication (and eventually to the expiry of the credential or the person). I use the phrase 'chain of trust' to stand for this over-all process of credentials and their lifecycle. Once one has the idea that digital identity is a process, rather than a state, the integrity of this process (and therefore the usefulness of the identity assertions it supports) can be seen to depend on a number of factors:

- How reliable were the initial processes of registration, verification and enrolment?
- How hard is it to duplicate or alter the credentials used?
- How strong is the verifiable link between the credentials themselves and the person presenting them?
- Correspondingly, how strong is the link between the authentication step and the entity carrying it out?

- How reliable and timely is the process of revoking credentials when they are no longer valid?
- In short, how easy or hard is it to subvert any of the steps of the process?

The term 'identity' is also commonly used in the sense of 'digital footprint'. Here, online identity is analogous to personal identity as a social construct. That is, in our social lives, a person's identity is a reflection of those things which are generally known about them by the people with whom they interact. Many people may know many different things about someone; some of those things may be shared knowledge, some may not. To the extent that they are public, they reflect what someone is 'seen as'.

In our online lives, our 'network identity' or 'digital footprint' is represented by all those pieces of information about us which accrete in various ways as we interact online. Obvious examples are:

- the user profile assigned to an employee, which is used to control their access to online resources
- the personal and payment details a consumer discloses to an online retailer.

Some other examples are:

- the 'browsing behaviour patterns' an online retailer collects as you use their website
- the data held about you by credit reference agencies and disclosed to third parties
- the audit trail of vehicle data collected by a road charging scheme.

This data is by nature highly distributed; it is usually collected and held by entities other than the data subject; it is used with varying degrees of the subject's knowledge and consent. As identity data, its collection and processing may well be covered by data protection and/or privacy laws if these exist in the jurisdiction in question – but the global, distributed and interconnected nature of today's online services mean that this can end up being of little practical consequence. The law and its enforcement may, quite simply, be failing to keep up with changes in the technology of identity. Jeffrey Robinson, an investigative writer who has studied identity theft in the context of money-laundering, sums it up as follows:

> "As long as we persist with a C17th notion of national sovereignty, an C18th judiciary and C19th law enforcement, the C21st will belong to organised crime."

It is a chilling prospect, dramatically expressed, but it also serves to highlight some of the real issues concerning the balance of technical and policy controls over identity data in the networked world.

We will return later to the idea of 'network identity' and its management, but first let us look at two more detailed models for digital identity data of the kind I described above.

2 Two models for identity data and its use

As the preceding section probably indicates, digital identity management is still a relatively young and evolving field. As such, it is all too easy even for a roomful of experts (perhaps especially a roomful of experts...) to spend inordinate time resolving confusion over terms, and arriving at a common vocabulary. The problem is compounded when one tries to express identity-related concepts in terms which will convey the same meaning to policy-makers, technologists, users and privacy advocates.

The following models are presented as a step towards generalisable concepts for the most common elements of an identity data-set. They are based on work carried out over the last 18–24 months with a wide range of identity-related stakeholders; I do not claim any credit for originality – rather, I am grateful to all those who have contributed their ideas and experience to this effort.

2.1 Model 1: the 'onion'

This model has proved to be useful in clarifying some of the different types of identity data, and how they relate to the concept of credentials described earlier. It also reveals an item of identity-related data which is not only often overlooked, but needs special treatment if the system as a whole is to work properly.

Figure 1 The 'Onion' model (1)

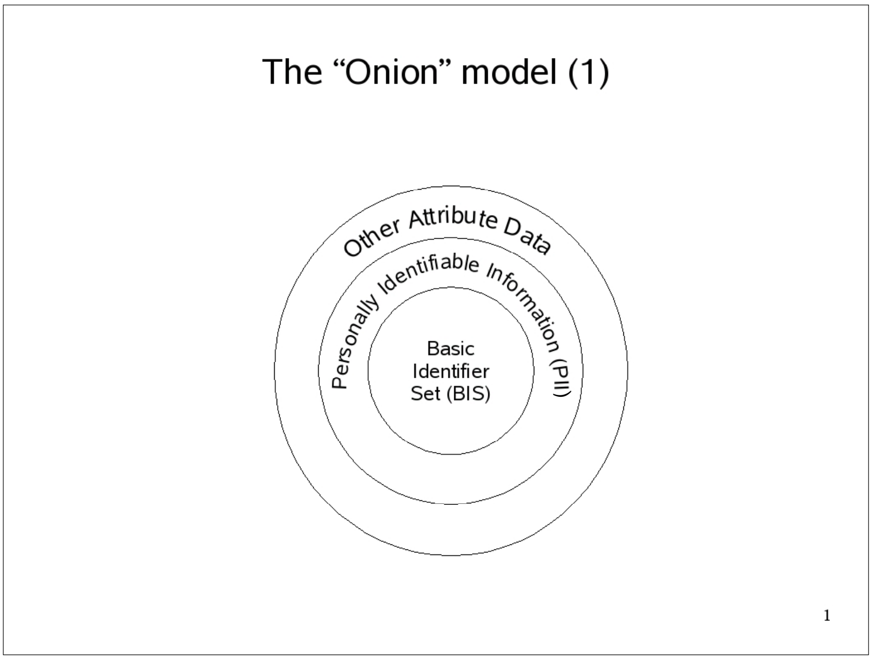

The "Onion" model (1)

Other Attribute Data

Personally Identifiable Information (PII)

Basic Identifier Set (BIS)

1

The model starts with three concentric rings of identity data; the boundaries of these rings are not immutable, but the starting point set out here serves to cover enough cases to be useful. At the core is the so-called 'Basic Identifier Set' (BIS), which is what most governments, for instance, consider to be the minimum set of data items necessary to establish the uniqueness of a given person. Typically, this will consist of:

- given name, family name (*modulo* cultural conventions in this area)
- gender

- date of birth
- place of birth.

The middle ring consists of other personal data which is sometimes more variable than the BIS, but still useful to identify someone. 'Address' is a good example of this: someone's address usefully identifies them, but may change in the course of their lifetime. Physical characteristics such as height, hair/eye colour and complexion might also fall into this category.

The outer ring contains those data items which are identity-related but 'sector specific', such as personal data relating to my healthcare records, my tax return, my driver or vehicle licence and so on.

With the basic 'onion' model in mind, consider how credentials relate to it. A credential such as a passport or driving licence typically includes some items from each of the three rings: some or all of the BIS, some items of other personal data such as physical characteristics or address, and some sector-specific data such as entitlement to drive specific classes of vehicle, or visas indicating entitlement to enter a specific country.

At this stage, we have identified three basic types of identity data:

1 credentials

2 attributes

3 entitlements.

And also that, in fact, credentials are simply a way of encapsulating attributes and entitlements in a reliably-verifiable form, as part of the 'chain of trust' described earlier.

Let us also note some of the historical factors which have resulted in this kind of credential, and some shortcomings from which the digital equivalent need not suffer. Historically, a document such as a passport or driving licence had to be usable as a 'stand-alone' credential. That is, it needed to contain all the necessary information for the relying party to make a judgement about authenticity, accuracy and therefore entitlement. Some form of passport has existed in Europe for 500 years or so, so passports (and subsequently driving licences) had to be usable as credentials without the assumption of an online network to verify any associated attributes or entitlements.

One drawback of this is that a credential may reveal more about the individual than is required for a given authentication context. For instance, if you use your passport to prove that you are over 18, you might also reveal your name, date of birth and citizenship. If you use your driving licence, you might reveal your date of birth and address.

In an online system, it is no longer necessary for credentials to carry all this information around with them, or to disclose data which is not necessary for the context in question, provided it offers a way of linking to that data held somewhere else. It is also unnecessary for a date of birth to be disclosed in order to prove that someone is over a given age; a simple 'yes/no' response can be given to a question of the form 'is the bearer over 18?'.

Thinking back to the 'onion' diagram, one way to visualise this is that there is an opportunity for credentials to 'migrate' inwards through the rings; that is, instead of including data items from all three rings, credentials can realistically encapsulate just those data items which serve to uniquely identify the holder (such as the BIS), as long as this provides a way of linking to the rest of the holder's personal data, which may be held elsewhere.

In other words, the option now exists to make use of the distributed nature of networked computing, so as to allow much more flexible 'placement' of identity data of different types. This is valuable in terms of policy control, because it makes it possible to apply controls at the place where the data is held, rather than trying to enforce it wherever the credentials are verified.

It also makes it easier to enforce policy appropriate to the data in question, particularly when different sector-specific data items entail different policy controls. For instance, the fact that I am entitled to drive a Heavy Goods Vehicle may not be particularly sensitive in terms of personal privacy – so one set of data security policies may apply to my 'driver and vehicle' sector-specific data; on the other hand, data about my healthcare history and medical conditions may be very sensitive, so should benefit from a different and more stringent set of data security policies.

Graphically, one might think of this as the ability to divide the 'onion' into sector-specific wedges, and cater for discrete management policies by sector and data type. Thus, within a given 'wedge', assertions of identity ('the holder of this credential is Robin Wilton') may make one kind of data security policy appropriate, while assertions of other attributes ('the holder of this credential has been treated for Repetitive Stress Injury') may require quite different policy treatment.

Figure 2 The 'Onion' model (2) (see online version for colours)

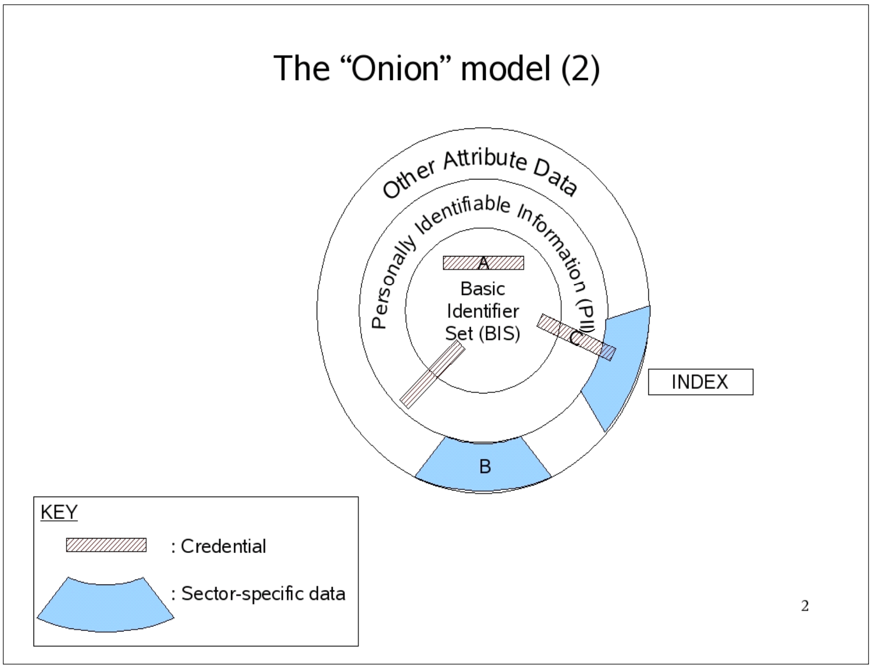

The "Onion" model (2)

Other Attribute Data
Personally Identifiable Information (PII)
A
Basic Identifier Set (BIS)
C
INDEX
B
KEY
: Credential
: Sector-specific data
2

To recap: the 'onion' model makes it possible to identify different types of identity-related data (the three rings), and the role and position of credentials relative to those rings (the option for credentials to 'migrate' inwards towards the centre of the onion). It also allows one to distinguish between different sector-specific subsets of identity data, so that they can be segregated and treated accordingly (the 'wedges'). As we will see later, that segregation and treatment may rely on technical and non-technical means; the challenge often lies in getting the balance right between technology and policy.

Developing the 'onion' model, we also identified another kind of identity-related data item which exists in almost every credential system, and which requires specific management controls – which are often missing. That is: for every system of credentials (managing a wedge of the onion), each individual's credentials are associated with an index value of some kind. For instance, passports and driving licences have an index number which uniquely identifies them in the system. This index to an individuals records may, but does not necessarily appear on their credentials.

3 The role of the index

Let us look at three examples which illustrate different ways of treating this data item, and the implications of each:

1 the 'Social Security Number' approach

2 the policy-based approach

3 the technology-based approach.

Please note that although I do refer to the specific countries which give rise to these examples, I make no judgement as to the over-all effectiveness of national identity systems in those countries. I am offering a very selective view of a specific characteristic in each national case to illustrate a wider principle which will still only be a part of the whole picture. That principle may be beneficial in a given country's case, but counteracted by other things which do not work well; conversely, it may seem flawed, but in fact be counteracted by other beneficial factors which I have not mentioned.

3.1 The social security number approach

The US Social Security Number (SSN) is an index into the Social Security system and its various repositories of citizen data. It was introduced for a specific purpose, and US law prohibits it use as an authentication credential for other uses. Nevertheless, it is now widely requested, for instance, in support of applications for bank accounts, utilities and so on. It is an index which has been treated as if it were a credential, when the intent was that it should be neither widely used nor widely published. As a result, SSNs are easier to 'steal' than was the intent, and once stolen can be used to do more damage than was ever intended (by being used in spurious applications for bank accounts, credit and so on).

In the long term, this both devalues the SSN as a credential, and also undermines its intended use as a reliable unique identifier for individuals in the Social Security system. The abuse of the SSN is now so widespread amongst fundamentally well-meaning organisations (such as utilities and financial institutions) that it is hard to see how this trend could be reversed... and yet, because it is still probably the most pervasive identifier in the USA, the temptation is there to use it as the index to a proposed national identity system. It is hard to see how this could be achieved without thereby building into the new system all the problems and issues which currently exist regarding the widespread general use of the SSN.

3.2 The policy-based approach

The Norwegian government has implemented a scheme based on a national identity number, but their policy is that these should not be revealed by default. Public sector applications wishing to make use of that number as an index to citizens' sector-specific records must apply for specific legal permission to do so.

This system seeks to protect the citizen's index number, but primarily through policy measures. Of course, enforcement will rely on the auditability of the system and its use, which one can expect to be technologically mediated – but the basic assumption is that, if an application, agency or government body is found to have misused its access to the index number, the remedy of first recourse would be a legal one.

3.3 The technology-based approach

The Austrian national electronic identity scheme has been widely publicised, often as an example of the application of technical measures to help implement the desired policy. The policy starts from a stance in which privacy and data protection apply to the sharing of citizen data amongst public sector bodies, and that national identity systems should be designed in such a way that the national data commissioner is able to exercise appropriate control over such data sharing.

To this end, the Austrian scheme uses a 'national' index number which is used to produce a number of sector-specific identifiers by means of a one-way cryptographic algorithm. The intention of this is that a sector-specific identifier cannot be used to work out either the 'national' index number for that citizen, or any of their other sector-specific identifiers. The only way to find out, for example, which health system identifier corresponds to a given tax system identifier is to apply to the data commissioner to be given the linkage.

Thus, there are specific technical measures built into this system to ensure that indices are managed in a way which enhances privacy rather than undermining it.

The following diagram illustrates the role of the national index number relative to the sector-specific ones.

So, in summary, it is important to be aware of the index as a piece of identity data, and of the implications of confusing the index with the credential or with other attributes. There is no single 'right answer' about whether indices should be protected through technical or non-technical means. As I will suggest below, a balance of the two is likely to be needed.

Figure 3 The ‘Silo’ model (1) (see online version for colours)

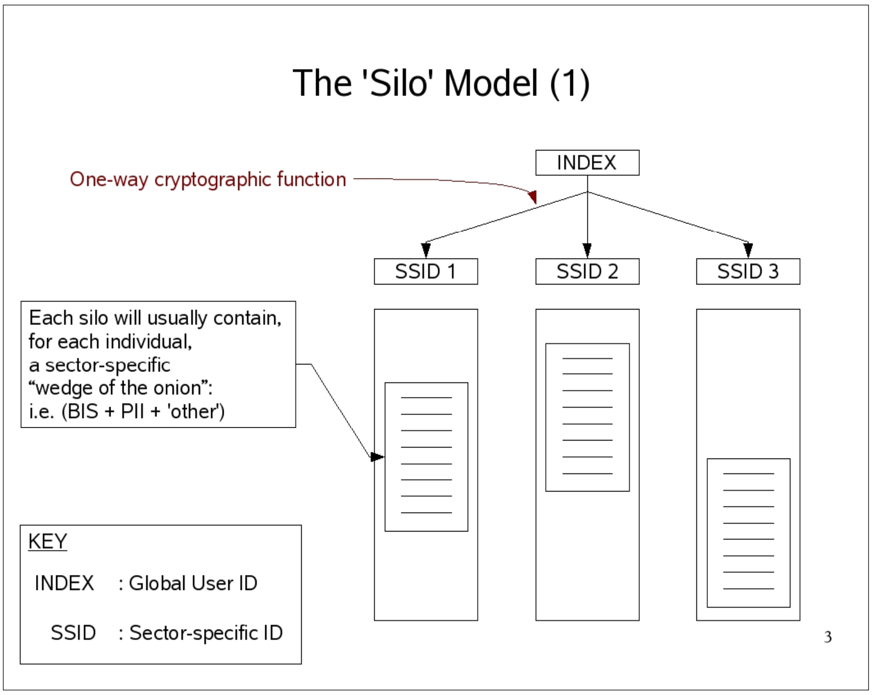

The 'Silo' Model (1)

One-way cryptographic function

INDEX

SSID 1

SSID 2

SSID 3

Each silo will usually contain,
for each individual,
a sector-specific
“wedge of the onion”:
i.e. (BIS + PII + 'other')

KEY

INDEX	: Global User ID
SSID	: Sector-specific ID

3

4 Model 2: balancing technical and non-technical measures

In the previous section, I used a ‘silo’ diagram to illustrate the role of the index relative to other pieces of identity data. In this section, we will look at how that principle relates to the ‘onion’ model described earlier, and how technical and non-technical factors interact in multi-system identity schemes.

First, let us consider how the ‘onion’ model applies to these silos of identity data. I noted that the complete set of an individual’s distributed identity data often consists of subsets, each relating to a particular context (health, tax, employment, *etc*.) and each consisting of some combination of credentials, personally-identifiable information and attribute data, and index. These context-specific subsets are what is held in each silo of the diagram below.

Going back for a moment to the Austrian national identity example, one can see that what their use of cryptographic mechanisms does is impose a one-way relationship between the primary index and the sector-specific ones. However, this does not, in itself, create any technical or policy barrier to the exchange of data between one silo and another.

Figure 4 The 'Silo' model (2) (see online version for colours)

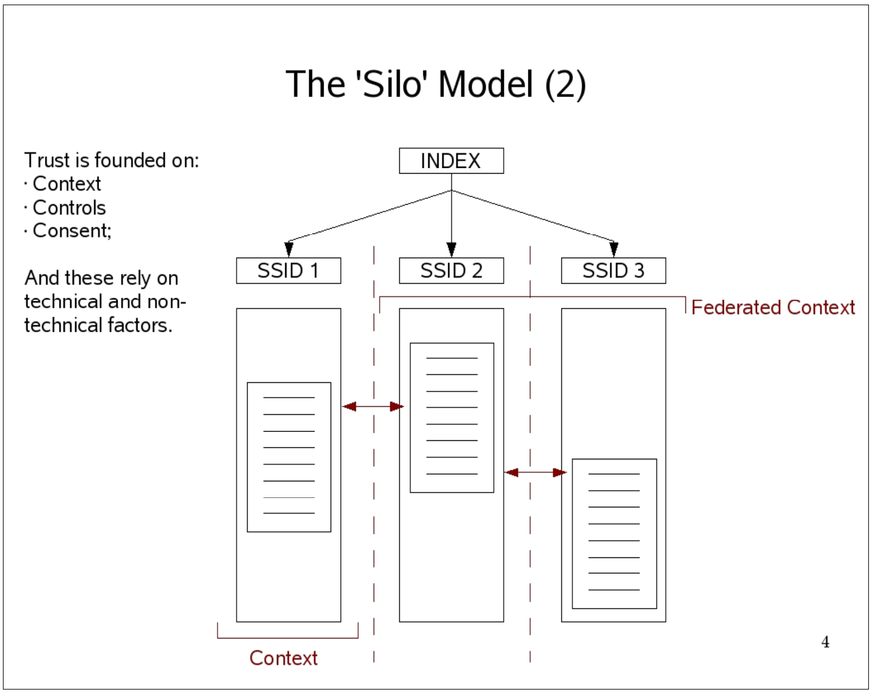

Let us make some hypothetical assumptions (not based on any factual knowledge of the Austrian system's internals). Let us assume that:

- the silos include (or are) databases, directories or equivalent repositories of citizen data
- in some circumstances, an individual has access (legitimate or otherwise) to the data in more than one silo
- that individual knows some of a citizen's attributes in one silo or context, but not their primary or (other) sector-specific identifiers.

Under these circumstances, whether or not the individual knows the citizen's primary or sector-specific identifiers, a search of each database on attributes such as name, phone number and/or address will be enough to link the sector-specific records of that citizen.

In other words, technological protection of the index numbers (through cryptographic separation) is not sufficient to rule out linkability of the citizen's various records across sector-specific silos. Again, I stress this is a hypothetical example, and is not based on knowledge of the internals of the actual Austrian system. The system may well include other measures to mitigate this risk.

5 Relating the 'silo' model to federated identity systems

5.1 A simple e-government example

The 'silo' model can be used to provide an interesting perspective on federated identity systems. There are other ways to describe federated systems – one being to talk in terms of 'circles of trust', and the various roles within such a circle: identity provider, service provider, attribute provider, relying party and so on. However, let us consider the 'silo' model with some added features, using a hypothetical e-government example.

So far, I have referred to each silo as a context: for instance, the data relating to an individual in terms of unemployment benefits, tax payment or other services. Federation is, essentially, the process of creating a new context which encompasses more than one silo. In the context of unemployment benefit pure and simple, it might never be appropriate to exchange data with the tax system. However, in the context of preventing benefit fraud one might wish to cross-check the two and ensure that someone paying tax on employment income is not simultaneously claiming to be unemployed.

I have deliberately constructed a (hypothetical) 'law enforcement' example, because it slightly simplifies things by removing the otherwise very important notion of user consent. We will come back to that in a moment. For the time being, though, our new, federated context establishes the conditions, rules and controls subject to which data may be shared between what are otherwise two discrete systems – each with its own governance framework and access control rules.

There would be two principal elements to the federation:

1 the governance framework which specifies how it should work, the responsibilities of the parties involved, what data can be exchanged, and what should be done if something goes wrong

2 the technical means by which the federation is effected and managed.

To illustrate the balance between these technical and non-technical elements, let us look at a slightly different example.

5.2 A commercial example

In this example (also hypothetical), two commercial service providers, a bank and an insurance company, surmise that they share many of their customers. They therefore decide to offer a federated system in which customers can log in to one service provider's online systems and thence transparently be offered services from both providers' systems. For this to happen without the user having to re-authenticate or re-type personal details, the two companies put the technical means in place to share the data.

They realise that there are regulatory constraints on the sharing of some of this data, so they consider ways of ensuring that they have the customer's consent. This factor was missing from the previous example. One approach they might take is to ensure that whenever a data exchange takes place, the user is shown what information is being shared, and invited to show consent. That consent might be given in varying degrees of complexity:

- it might be sought every time, or just the first time – with an option to revisit the decision subsequently

- it might be ‘directional’; that is, the user might be given the option for data to go from their insurer to their bank but not the other way. This might reflect some inherent difference in the relationships the user has with the bank and the insurer
- it might be specific to particular data items, or a blanket authorisation to ‘share such data as is necessary to carry out the business in question’.

One problem with building a system which assumes the user is always to be involved in the data-flow is that the user may sometimes want to initiate these kinds of transaction without being online at the time (for instance, by post or by phone). It may be impractical to offer the same range of consent options as set out above.

It may also be costly for the service providers to implement such a ‘granular’ system, and complex for users to choose all the various options which reflect their actual preferences.

For whatever reason, the service providers might instead decide that the first step in offering their customers a federated service is to get them to sign up to a set of terms and conditions which governs any subsequent exchanges of data once and for all. In practice, this is how a lot of electronic banking systems work: the user is not explicitly asked to show consent to the execution of online transactions, but signs up to a contract under which the transactions executed online are assumed to have the same legal weight as those the customer might execute in person at a branch office.

In this last example, then, the balance has swung from technically-assured to contractually-assured consent – though of course the technology still has to be sufficient to show that a transaction executed online can reasonably be assumed to have originated from the customer in question.

5.3 Recap: elements of a federated identity system

To summarise: the silo model may not accurately describe every identity system, but it allows us to identify and consider the essential elements:

- how the data in a given silo corresponds to the kinds of information we identified earlier using the ‘onion’ model
- the role played by any sector-specific or over-arching index value, and the specific management measures appropriate to that
- what enables information transfer between the silos... whether technically or non-technically mediated
- conversely, what (technically or otherwise) *prevents* inappropriate information transfer between the silos
- the role of ‘context’ in defining what happens to data either in a given silo or when it is exchange between silos
- the importance of consent, and a distinction between law-enforcement and consensual data exchange
- the balance between technology and contract in different options for consent-seeking.

6 Trust and privacy

Having now used the ‘onion’ and ‘silo’ models to identify some of the essential elements and functions of a federated identity system, we have a basis on which to consider the abstract but important aspects of trust and privacy.

Trust in a federated system is likely to depend on aspects we have already encountered; specifically, context, controls and consent.

Context is important because it has so much to do with the expectations we have of a given system. Most of the ‘silos’ or services I have referred to so far have been ones with some kind of contractual or statutory foundation – e-commerce or e-government respectively. Those often set a pretty clear context within which we expect our online activity to be bounded. However, there are many other kinds of online activity where that contextual background is less clear or even non-existent. Here are two examples:

1 Recent news stories have served to raise the question of what liability, if any, an internet movie-hosting service might have to some of the content which is posted on it. Would YouTube bear some responsibility if some of its members post videos showing bullying or violent assault? Apart from the question of the content, does the service provider bear responsibility for the behaviour which it is being claimed such content might normalise, encourage or provoke? What is the context of expectation and regulation (if any) in which YouTube contributors and viewers are acting?

2 If I consent to publish my personal information via a social networking site, what expectations might I have about how it is processed? If, as a result of my membership or the information I disclose, I suffer some form of identity fraud, what recourse do I have and against whom?

Some more subtle risks are also starting to emerge; it seems that some social networking sites are the ideal environment in which to propagate various forms of malicious software. It is increasingly possible to create ‘mash-ups’ which enhance the basic functions of a service by adding new features which appeal to a particular group of users. Those mash-ups are often written and contributed by individual subscribers, as opposed to the owners of the service in question. Although not a social networking service, Google Maps is probably the best-known ‘host’ for mash-ups: for instance, there is a service which combines Google’s mapping with user reports of wi-fi availability, to produce a site which will show you where the available wi-fi hot-spots are in whichever part of the map you select.

So in a social networking site, if someone sends round a message saying “I’ve just written a mash-up which will let you know when any of your buddies is in the same town” it is increasingly normal to take that at face value. ‘Just add this plug-in to your browser’; ‘just turn on this option in your subscriber preference profile’, and so on. In the social networking context, it might also be the case that when you ‘enable’ the attractive mash-up, you might also unknowingly granting it access to the personal details you have stored there, and which you thought were being protected by the privacy preferences you expressed in your profile.

In other words, the context in which you thought you were acting might be being subverted or over-ridden without your knowledge.

Controls are important because they are the means by which we ought to be able to govern the way in which our information is used – within or between 'silos'. As noted earlier, a 'context' might be extended to cover multiple silos – for instance, in a federated system – in which case the controls over data exchange between silos are what reassure us that there is porosity where we want data-sharing, and impermeability where we want data privacy. The controls may be technical or non-technical in nature, but non-technical controls (such as contractual agreements) will need to be underpinned by enough technology to show that they are working.

Just take a moment to think: when was the last time you disclosed your personal details to a website which advertised a privacy policy? Did you read the policy? If you accepted it (with or without reading it), do you have any ideas about how you would tell if the site in question were abiding by it or not? Putting a non-technical control in place may be relatively simple; auditing and enforcing it may pose some very knotty problems. The forensics of data privacy are still a relatively untested art – outside the intelligence services, one might conjecture. We will come back to this topic later when we look at privacy in more detail.

Consent is important because it signifies that we know why we are disclosing information, that we have some expectation about what will be done with it, and that we have some level of confidence that one will match the other. Those concepts are the basis for most Data Protection legislation: why was this data collected? What are you using it for now? Are those two things compatible?

Thinking back to the first of the two examples described above, one can imagine that the victim of a violent assault might not readily consent to having a video of it posted on the internet by a third party – unless, of course, that contributed to identification and prosecution of their assailant.

On the other hand, the consent model in the second example (voluntarily making use of a malicious mash-up circulated via a social networking site) might be a little more ambiguous. The user consented to some degree of disclosure, and although they did not knowingly consent to the malicious disclosure, some degree of consent might be considered implicit anyway in a social networking site.

Either way, abusing the consent of a data subject is a sure way to undermine their trust in the system – assuming they are able to find out that you did so (see the comment above on forensics).

7 Privacy online and in the 'real world'

Let us consider some characteristics of privacy in the online world, and ways in which it differs from privacy as we understand it in our 'offline' lives. Take the following two lists:

1 'Level one' list

- I am a systems architect
- I work for Sun Microsystems
- I am based in England
- I have a strong interest in identity, privacy and policy matters.

2 'Level two' list

- I like Scarlatti
- I never eat offal
- I oppose factory farming
- I would probably like books by Henning Mankell.

The first list contains the kind of information you might reasonably expect to gain from a brief conversation with me in person. The second list contains the kind of inferences which are made about me in practice by online retailers such as supermarkets and media retailers. They reveal a very different level of information about my preferences and values.

The second list would reflect a face-to-face conversation about attitudes, ethics and even beliefs. That is not the kind of interaction I am conscious of having with any online retailer, and yet I disclose information to those companies which is revealing to this relatively intimate extent.

I would like to draw three points from this example:

1 I may actually be unaware that I am disclosing this level of information in the second case. It may be being derived and inferred simply from input like the amount of time I spend on one web page rather than another, which links I click first on a page, or which pages I come back to repeatedly.

2 In the real world, I doubt that I would simply disclose a series of 'Level two'-type facts about myself in a rambling monologue; they would probably be part of a social exchange of information, in the course of which I would learn some equivalent things about the other person. In the online world that model does not apply. I do not find out anything about the retailer's personal preferences and values as a consequence of disclosing my own.

3 In a social context I would, over time, expect to form a view about who is discreet with personal information and who gossips. I would probably adjust my behaviour appropriately, particularly if I found that someone could not keep a secret. In the online environment it is often hard to make equivalent judgements. When an unsolicited e-mail arrives in your inbox it is usually impossible to determine where the sender got your e-mail address, for instance. If unexpected transactions start cropping up on your credit card statement, that will not tell you where the perpetrator got your payment details from.

Partly this is a problem of scale: online, you probably disclose your payment details to more retailers than you would entrust close friends with your personal secrets. Partly, though, it is also a problem of data forensics. I will conclude by looking at two ways in which that problem can be analysed.

7.1 'Privacy forensics' 1: beyond first disclosure

One of the problems which remain to be solved in the area of data privacy is this: assuming that you are able to both express and enforce a privacy preference when you disclose information to a counter-party, how can you be sure that that entity will not subsequently abuse the data? After all, you do intend them to see and understand the data in question – so ultimately, even if you protect it digitally, the recipient could simply turn away from their computer and write the data down.

The challenge is to make your privacy preference somehow stay attached to the data in question; this is sometimes referred to as 'sticky policy'. Because of the 'analogue copy' option just mentioned, any solutions to this are likely to be partial. One way to analyse the problem is to look at the two steps of 'privacy preference expression' and 'privacy preference enforcement' as two ends of a communication pipeline between yourself and the recipient. It might be possible to insist that there should be an enforcement point at the recipient's end as a pre-condition of disclosure, but it is probably unrealistic to expect that you can insist on a policy enforcement point at every subsequent entity to whom the recipient might disclose the information – so your preferences provide only limited protection.

One suggestion is to arrange things so that any recipient of your data can only gain access to it by bringing it back to a policy enforcement point over which you have control. For instance, you encrypt the data before disclosing it, and make the recipient come back to your policy enforcement server in order to be given access to the keys.

Digital watermarking may offer countermeasures to some other aspects of data disclosure – though again, if the data is such that an 'analogue' copy can be taken, it may be difficult to tell which original has been abused and therefore who disclosed it.

Ensuring the persistence of policy beyond first disclosure is a problem which shares many of the characteristics of Digital Rights Management (DRM) – except that in this case, we are aiming to protect our own data and privacy. Application of DRM-style technology might therefore enjoy greater adoption and end-user support in this case. However, as digital media piracy suggests, it would be unlikely to prevent abuse altogether.

7.2 'Privacy forensics' 2: separation of personas

> "One mans' silo is another man's vault."

In the earlier sections of this document I used several 'silo'-based examples to illustrate principles of identity data and identity federation. The term 'silo' is often used pejoratively to refer to a system or application in which lack of openness and flexibility in the architecture has negative results in terms of data-sharing, software reuse or extension, and so on. However, the examples also serve to illustrate how it may be highly appropriate to maintain functional separation between different context-specific sets of identity data.

Some of the personal data in one context is inappropriate to share with other contexts. The local council does not need access to my prescription records; the pharmacist does not need to know if I got a parking ticket last week. Again, in the 'real world', people discriminate when disclosing personal information, and often, do so in order to keep the different spheres of their life separate.

I use the term 'persona' to refer to that subset of personal data which it is appropriate to reveal in a given context. There are other related terms. The European FIDIS project refers to 'partial identities', for instance. My conjecture is that most people are at least vaguely aware of presenting more than one real-world 'persona' – although if they are radically different, the inconsistency potentially ranges from the merely odd to the clinically pathological.

When we go online, though, our ability to visualise, understand and manage multiple personas is hampered by the lack of 'cues' we get from the online world. It is not obvious what information a service provider or website can 'see' about us, and how that compares with other sites and other contexts.

In the early 2000s, the term 'network identity' was used to refer to the disparate sets of personal data which all relate to a given individual, but which have accumulated on different service-provider and other sites around the internet. It is very hard for an individual to form any kind of coherent picture of this distributed data, let alone manage it. There is no sense of an over-arching system of 'persona management'. That, in turn, means it is increasingly hard for the individual to ensure that each service provider knows only those things appropriate to the context, and that personal data is not being inadvertently revealed beyond the intended context.

In other words, what appears to be a flaw in the pejorative sense of 'silos' of data might be a benefit in terms of ensuring that personal data stays 'in context'. If we accept that data-sharing is increasingly easy, technically, and that there is increasing incentive on the service-provider's part to share data with other service-providers, then the question becomes one of what technical and non-technical controls can be put in place to ensure that context is respected, and that user consent is properly applied to such transfers of data between contexts.

Thus, we come back to the basic elements identified using the 'silo' model, and to the factors of context, controls and consent described in the introduction to the section on Trust and Privacy.

8 Conclusions

What we discovered as we explored the concepts of identity and privacy through a series of round-table meetings was that each participant in such a discussion was likely to have very different ideas about the meaning of basic concepts such as identity, trust, privacy, personal information, and so on. This was so whether the participants shared a common perspective (as, say, technologists or privacy advocates) or whether they approached the subject from radically different directions (policy-makers, academic researchers, *etc.*).

At first this seemed an almost insurmountable obstacle. What is more, unless we could arrive at some common understanding of the basic elements (What are credentials? What does 'proof of identity' mean?) there would be no prospect of a meaningful, cross-disciplinary conversation about the 'higher-order' issues such as trust and privacy.

However, the discussions allowed us to develop simple models such as those described in this paper (the 'onion', the 'silo') and others, which describe the basics in a simple way and thus pave the way for the more abstract conceptual discussion.

It is always tempting, at this stage in the evolution of a new technology, to assume that it will never amount to anything unless every single adopter becomes an expert. No doubt the invention of the motor car was greeted with derision by those who predicted that it would be unusable by anyone except a highly-trained (and slightly lunatic) mechanical genius. Twenty years ago it was probably still the case that opening the bonnet of the average car would reveal a standard set of engine components which a competent amateur could understand conceptually, and conceivably maintain. These days that really is not the case. The vast majority of drivers have no idea how electronic fuel injection works, nor how an electronic engine management system does what it does... let alone tiptronic gearboxes and hybrid dual-fuel propulsion systems. And yet they are competent users of such complex systems – partly because the complexity of the underlying system is well hidden, and the user is presented with familiar metaphors for what is actually going on.

And so I suspect that over the coming years, the complexity of identity management infrastructures will increase. Despite that, the vast majority of users will not be identity management experts: they will rely on metaphors which express what is being done with their personal data, their online personas and their privacy.

However this does imply a thorough understanding on the part of the designers of those metaphors and the infrastructures which underpin them. There is no single stakeholder or community which can provide those designers with all the input they need. Nor is there a single stakeholder who can give all the necessary input to those who implement, administer, audit and regulate the resulting systems.

What is needed is a long-term and productive engagement between multiple stakeholder groups, able to converse in mutually understood terms about common concepts. It is my earnest hope that papers such as this, and dialogues such as those which gave rise to it, will play a part in making that possible.

8.1 Recommended reading

The discussions on which the paper is based were held under the Chatham House Rule, which encourages free and open discussion on the basis that any resulting material cannot be attributed to the organisation or individual who took part.

That said, there are some materials which I recommend to anyone seeking to read further in this area.

Bibliography

Adams, J. and Birch, D. (Eds.) (2007) *The Digital Identity Reader*, London: The Mastodon Press.

Bellovin, S., *et al.* (2008) 'Risking communications security: potential hazards of the protect America act', *IEEE Security and Privacy*, http://www.crypto.com/papers/paa-ieee.pdf.

Birch, D. (Ed.) (2007) *Digital Identity Management: Perspectives on the Technological, Business and Social Implications*, Gower, Aldershot.

Diffie, W. and Landau, S. (2007) *Privacy on the Line*, Cambridge, MA: MIT Press.

Jenkins, R. (2004) *Social Identity*, London: Routledge.

Sen, A. (2007) *Identity and Violence: The Illusion of Destiny*, London: Penguin.

Wilton, R. (2007) *Liberty Alliance Privacy Summits 2007*, Berlin, Brussels, http://www.projectliberty.org/liberty/public_community/privacy_summits.