Digital Identity and Privacy

Establishing a common conceptual framework
for discussing digital identity and privacy
across industry and sectoral boundaries

FINAL DRAFT

Robin Wilton
Director of Privacy and Public Policy - Kantara Initiative

July 2010

Foreword

This document is for anyone needing to understand basic concepts of digital identity and privacy, whether for their own interest, or in order to convert them into policy or practice, or so as to discuss them productively with stakeholders from other domains.

The paper owes a great debt to many people, not least the participants in a global series of Privacy Summits run under the aegis of the Liberty Alliance. Several individuals have made key contributions to the development of the conceptual framework which this document describes:

- Hellmuth Broda, former CTO (EMEA) of Sun Microsystems
- Peter Lord, formerly of Oracle Corporation
- Caspar Bowden, Microsoft
- Piotr Cofta, British Telecom
- Lizzie Coles-Kemp, Royal Holloway University of London
- Susan Landau, formerly of Sun Microsystems; author and Distinguished Engineer

My sincere thanks to all of them.

Robin Wilton

Document date: July 2010

Document reference: KI-CCF-v0.9

Table of Contents

Why do we need a "common conceptual framework"?

At its heart, digital identity is both a technical problem and a societal issue: a purely technical solution may not be the appropriate one. Repeated experience tells us that techno-centric solutions to digital identity problems turn out to be inadequate, because the relationship between technical mechanisms and societal issues has not been properly understood.

This paper attempts to tease apart these issues in order to aid future discussions on digital identity solutions and, ultimately, give rise to better privacy outcomes for all of us, no matter which stakeholder category we belong to.

Over the last decade, the identity landscape has spread and evolved at least as fast as any other domain in the information society.

- At the turn of the millennium, the hot topic was how to federate authentication beyond the boundaries of the enterprise... so as to overcome some of the obstacles to multi-party e-commerce models. Collaborative service provision remains a challenge and an opportunity in both the commercial and public sectors;
- Over the following five years or so, technology vendors competed for the 'enterprise identity management' market, building on directory services to create authentication servers, and products for authorization and entitlement management, and provisioning. Organizations still need to manage and make productive use of identity data about their employees, partners, suppliers and customers;
- As the technology continued to evolve and mature, identity-related concerns extended beyond the possible and into the ethical, with a growing (and continuing) focus on data-sharing, data mining, social networking, and the implications of all these for personal privacy in the online world. "The question is no longer 'can we do it?', but, in many cases, 'should we do it?'"[1]
- Technologies such as biometrics, CCTV and geo-location, all in a frontierless digital ecosystem, raise challenging issues of public policy.

Each of these four bullet points embodies a challenge. What the challenges have in common is that no single stakeholder has the solution... and that addressing one part of the problem solves nothing if the other parts go unfixed.

What is needed is a collaborative, multi-stakeholder approach to defining and addressing the issues of digital identity and online privacy.

Unfortunately, that is more simply said than done. As we assembled diverse stakeholder groups for the Privacy Summit programme, we often found that any ten stakeholders have at least a dozen different perspectives on the same problem... and usually several other problems which seem vital to them and irrelevant to the others. Add to that the fact that often, terms and expressions are used to mean different things by different stakeholders, and the problem of even having a productive discussion starts to look intractable.

Our solution was to define a common conceptual framework, using which we could quickly establish some simple reference points for the discussion, so that instead of wasting time re-hashing the basics, the participants could move on to the substantive issues and contribute their specialist expertise.

This document sets out that conceptual framework, in the form of a few simple models.

1 Fulup ar Foll, Oracle Corp. (formerly of Sun Microsystems)

The "Stakeholder" model

The 'stakeholder' model emerged from a very simple, fundamental question: "why is digital privacy perceived as a problem?" - less in the sense of "why bother to do anything about it?" than in the sense of "whose problem is it, and why?".

As we tried to answer this question we observed that in most implementations, few – if any- of the participating stakeholders profess themselves happy with the outcome. Here are examples of some of the comments we heard from different stakeholder groups:

- Policymaker: "I just don't see why privacy has to be so hard. Why can't it be 'baked in' at the technology layer?"
- Technologist: "We seem to devote a lot of time and effort to developing technical solutions to this, only to get 'blind-sided' by some regulatory requirement which is either new, or new to us, or different from one geography to another"
- Implementer: "This project has soaked up enormous resources, and we still don't feel it has addressed the primary business objectives to do with trust and compliance"
- Citizen/consumer: "My privacy preferences seem to be off the bottom of everyone else's priorities".

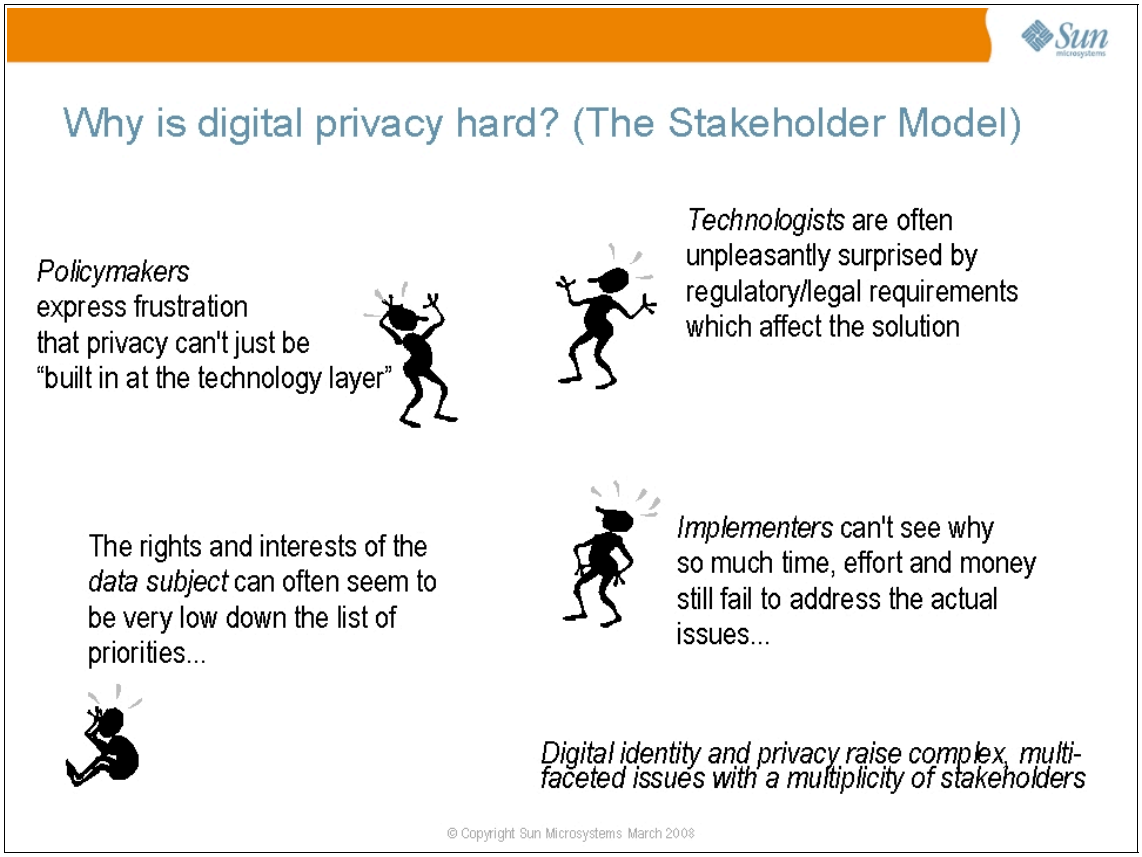

We drew a number of conclusions from these and other comments. The first was simply to confirm a principle which ran through every Privacy Summit, and which is true both of the 'problem analysis' phase and the 'solution development' phase:

"The identity and privacy challenges which confront us cannot be solved by a 'single-stakeholder' approach."

Privacy is a multi-stakeholder problem, and the stakeholders have very diverse requirements and expectations as to the outcome. The diversity of those requirements means that the solutions have to be correspondingly diverse, addressing political, regulatory, personal and commercial requirements with a range of technical and non-technical measures. We will consider this further in the next section [“The 'Ladder' Model”].

Privacy is not secrecy

The Stakeholder Model model hints at another important principle: privacy is not a state, it is a relationship. There's no such thing as “one-party privacy” - or, if there is, it is probably what we would usually call “secrecy”. If I simply never disclose my data, it's hard for me to have meaningful or productive interactions.

Privacy, it seems, is not about keeping everything to myself – it is about making disclosures, but making them in a way which preserves my consent and control, while respecting the needs of both myself and the recipient. As such, privacy is also highly contextual; for more on the importance of the concept of 'context', see the discussion of the “Onion” model, below.

The answer to the 'multi-stakeholder' question seems simple: involve them all in defining what should be done about digital identity, and make sure their various concerns and priorities are taken into account. Problem solved. Except that that was where the Privacy Summits had started from in the first place – born out of a recognition that digital privacy is a multi-stakeholder issue with no single-stakeholder solution. So what was going wrong? The answer took the form of the second model, the “Ladder”.

“Privacy from whom?”

It is also worth noting, at this point, that many of today's privacy issues arise out of the fact that, while online interactions often *appear* to be between two parties, there are frequently third parties involved, whether or not that is immediately obvious to the participants at the time. For example, when two individuals exchange emails, those emails will pass (at least) through the respective ISPs and email services of the two parties.

There may well be rules which govern what the intermediaries are allowed to do with the traffic they handle, but nevertheless, the sender and receiver may have differing assumptions about what each of those intermediaries does with the information as it passes through them – and those assumptions may or may not be well founded. Another example is that of internet search engines: in normal use, it is not generally evident to the end user that their search queries are catalogued by search engine providers, and the resulting data mined for its profiling and marketing value.

“Social networking” services, also, tend to mask users from the fact that they (the users) are not conducting a one-to-one conversation with their online buddies, but are doing so in the presence of an all-seeing third party with specific commercial interests.

Thus, defining privacy as a multi-stakeholder issue is not just a matter of acknowledging that policy-makers, enterprises, technologists, individuals etc. all have their own perspective on the question. It is also a matter of clarifying exactly who the participants in a given online service really are, and reconciling their interests with the privacy of the service user.

Unfortunately, as the Stakeholder Model suggests, the service user's interests are often very low on the list of the other stakeholders' priorities, and are often outweighed, in cost-benefit terms, by financial imperatives. As an example, consider the UK's Revenue and Customs data breach[2], in which tens of millions of taxpayers' personal details were exposed, in part because the cost of

2 UK HMRC data breach, 2007
http://news.bbc.co.uk/1/hi/uk_politics/7103566.stm

simply copying them all was a few thousand pounds lower than the cost of paying for a selective database extract of just a few hundred records (which was all that was needed); the resulting records were then written, unencrypted, to CDs which in turn were apparently lost at some stage in the mailing process. Successive failures of cost-benefit analysis, risk mitigation, operational processes and auditability resulted in the (to date) untraceable loss of sensitive personal information relating to what will, through the coming decades, be an entire generation of economically active citizens and taxpayers.

The "Ladder" model

As mentioned above, in the diverse, multi-stakeholder environment, privacy requirements are both varied and expressed in many ways. What one stakeholder understands by "trust", "personal information" or "ownership of data" may be very different to what another stakeholder means by the same term.

The "Onion" and "Ladder" models were first steps towards establishing a common conceptual framework which might allow diverse stakeholders to have a productive discussion despite their different perspectives, assumptions and vocabularies.

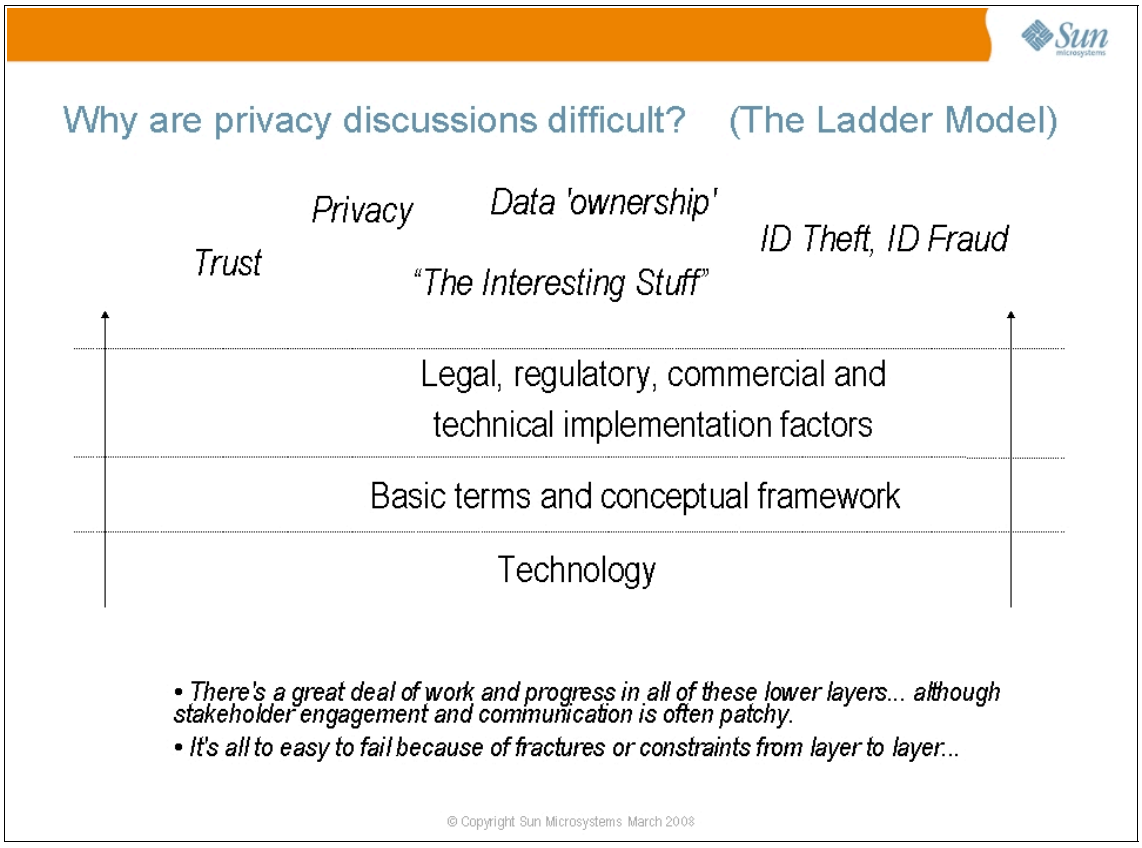

It is common (and easy) for a privacy discussion to start at the bottom, technology layer and work its way upwards from there. Unfortunately this can unwittingly open the way to a number of shortcomings, and ultimately, can mean that the top-level topics are not meaningfully addressed.

What we observed was this: that as you work your way up the conceptual 'ladder', it is very common for the discussion to narrow the further you climb, but very rare for it to broaden out. As a result, by the time you reach the top layer you may find you have nothing useful to say about important aspects of the solution, such as data ownership, trust, privacy and so on.

The Ladder model suggests two ways in which to solve these problems. First, it helps greatly if the stakeholders are aware that this is what is happening. In the course of a discussion, it is immensely useful to be able to 'situate' a given comment or requirement at the appropriate layer of the model. For instance, to be able to say "*x* is a comment about the technology layer, but *y* is an expression of a legal/regulatory requirement...". This helps to ensure that the interests of the various stakeholders are made explicit – and, if necessary, that part of the discussion can be 'put on hold' until a given

stakeholder group can be appropriately involved.

Second, we found that if the different stakeholders can be given a consistent, mutually-understood terminology and conceptual framework, it's much easier to identify, correct or pre-empt confusions among the stakeholders and between the different layers of the ladder.

The “Ladder” model does not address privacy issues as such – but it helps stakeholders understand and navigate a discussion process in which there is otherwise enormous scope for confusion, misunderstanding, and 'circling round the same arguments' almost indefinitely. Since we formulated the model, we have been able to test and validate it in subsequent discussions, and found it to be both useful and robust. Here are a couple of illustrative (hypothetical) examples:

1 - “Narrowing versus widening”

Looking at privacy from the technology perspective, here's what can often happen: an organisation wants to do something about secure authentication and privacy, so it asks for technology suggestions. One of the technologists asked is a smart-card specialist, and her instinct is to make the case that smart-card technology can solve the problem - say, through some combination of strong authentication and digital signing of assertions. Her understanding of the technical elements is flawless, but her conceptual framework is constrained by the technology with which she is most familiar. As the saying has it: “if the only tool you have is a hammer, every problem starts to look like a nail”.

As the project progresses 'up the ladder', it is found that for the core technology to address the digital signature aspects, further work has to be done on key management, certification, and protocols to provide “proof of origin” and “proof of receipt”. The technical scope of the project starts to creep.

Then perhaps it's discovered that because of different regulatory environments in different countries, digital signing will offer acceptable proof in some but not in others.

Then, at the top of the ladder, it may turn out that the resulting solution doesn't conform with the privacy laws in one country unless the organisation is able to retrieve, audit and delete users' records at the request of the data subject – so a further layer of administrative and process functions has to be added.

In simple terms, any given 'customer' tends to think of their ultimate requirements in terms of the high-level concepts at the top of the ladder. Other stakeholders (such as, in this case, the technologist) see the requirements from a different perspective, and through the constraints and filters of their own conceptual framework. It seems to be much easier for the resulting work to “narrow” as it progresses up the rungs of the ladder, and much more difficult for it to “widen”. Consequently, it may well end up addressing only a subset of what the customer sees as their requirements, or possibly miss the mark altogether.

Note that this in no way invalidates the input or the subject-matter expertise of the technologist (in this example); it just illustrates the importance of recognising that multiple stakeholder perspectives, requirements and contributions all need to combine in defining and fixing the problem.

2 – Managing contributions to the discussion

Particularly in these multi-stakeholder round-table discussions, we found it was all too easy for the discussion to be 'diverted' (albeit with the best of intentions), simply because of the wide range of perspectives stakeholders will feel the need to contribute.

The 'Ladder' helps with this, because it allows such contributions (which may well be valid and important) to be noted, positioned accordingly on the ladder, and dealt with when the time is right. This not only reassures contributors that their perspective will be taken into account, it also helps deal with the tension which can arise when one stakeholder feels that their "high-level concept" is being unfairly reduced to a "technology" discussion... or, conversely, when the technologists get frustrated because the conversation is revolving around policy and conceptual matters rather than tangible technical bits and pieces.

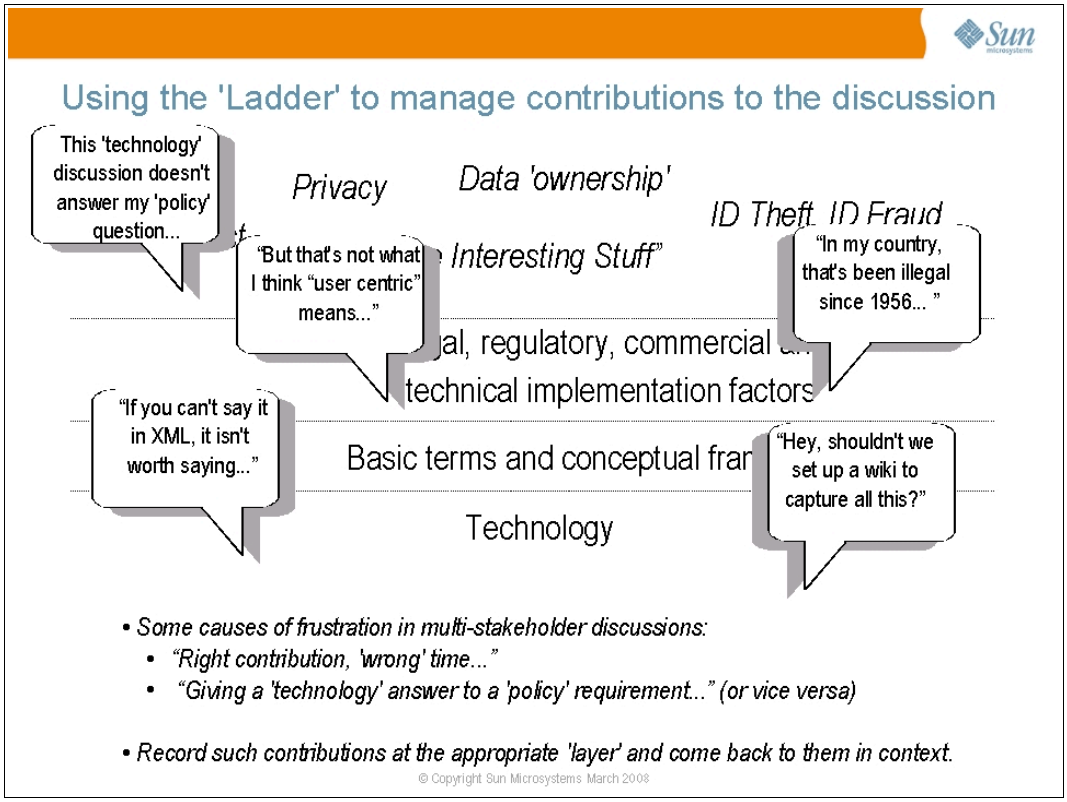

This also helps show the role the "Onion" model (below) plays as a tool for establishing that common terminology and conceptual framework.

The "Onion" model

Even among experts (and perhaps especially among experts), the phrase "identity data" can mean a wide range of things. For instance, how do identity, identifiers or credentials, and personal attributes relate to one another to constitute a 'digital identity'? If digital identities are composed of multiple elements, do those elements or types of element need to be managed differently? Is my digital identity made up of all the digital facts that are associated with me, or is it important to be able to segregate some facts from others (and if so, how and why)?

One principle we noted was that there seems to be a strong theme of 'uniqueness' about digital identity. This has both philosophical and practical roots. Philosophically, Leibniz formulated (in the late 1600s) the principle of the 'Identity of Indiscernibles', which states that if two things have exactly the same set of properties then they are one and the same thing – they are identical. A relation of 'identity' obtains between them. Practically, we can determine that two things (or people) are not identical by looking for some attribute that they do not have in common – in other words, we look to prove identity through uniqueness.

When we apply this to digital identity, it reveals that what digital authentication credentials seek to do is establish the relationship of identity between "the person to whom the credential was issued at some point in the past" and "the person presenting the credential now".

To do this, many national identity schemes are based on a so-called "Basic Identifier Set" (BIS). These are the small set of data attributes which are generally considered, in such cases, sufficient to establish the uniqueness of a given individual within a given population, group or community. A classic example of a BIS is:

- Name (Given name, family name)
- Date of Birth
- Place of Birth
- Gender

Examples have been given[3] of cases in which any or all of these attributes might not be immutable, but for practical purposes they form the core of most large-scale identity schemes such as passports, identity cards and so on. The credential (the passport or ID card) encapsulates key items of personal data in a form which is hard to tamper with and easy to check reliably. At least, that's the principle.

However, identity-related data clearly also encompasses a much wider range of data than just the BIS. The model we derived was a layered one, in which the BIS is the centre, surrounded by other Personally Identifiable Information (PII), which in turn is surrounded by other attributes and historical data relating to an individual. This is illustrated in the diagrams below.

3 Gillian Ormiston, OECD Workshop, Trondheim May 2007: http://www.oecd.org/document/41/0,2340,en_2649_34255_38327849_1_1_1_1,00.html

Basic "Onion" model

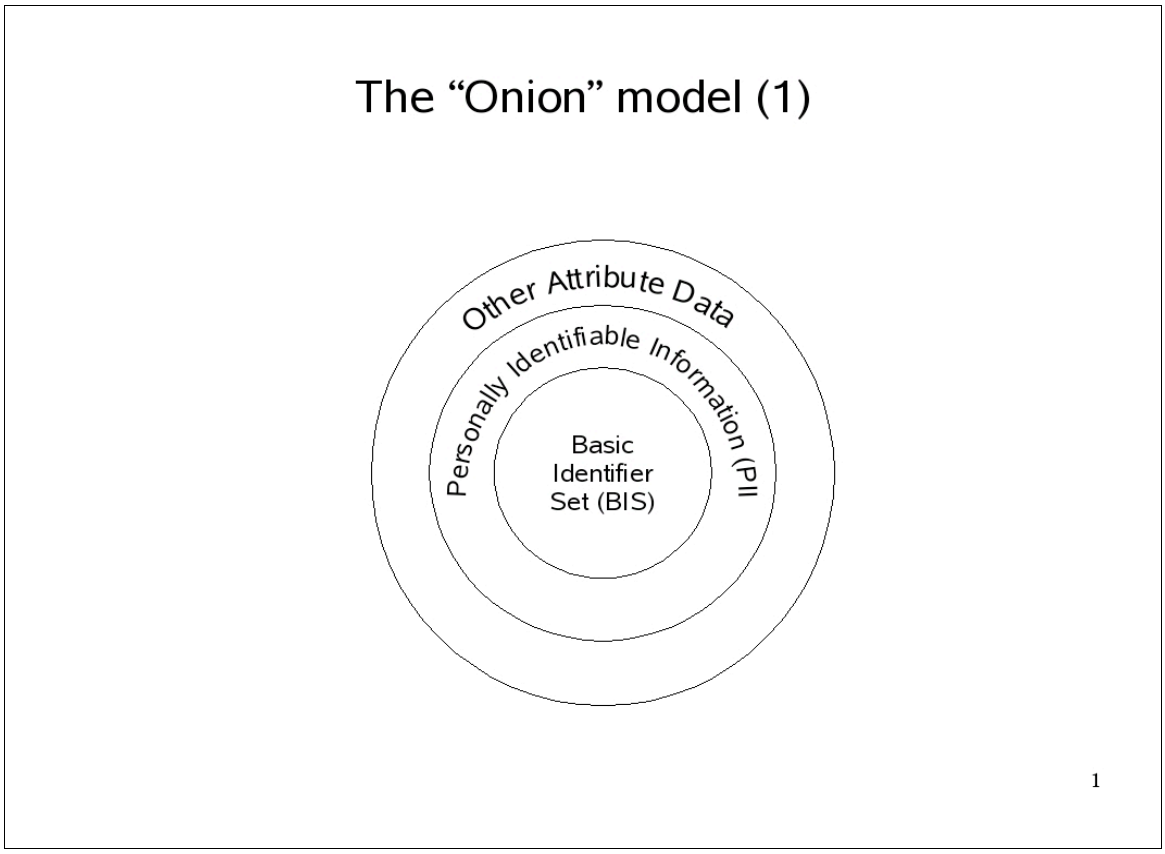

The first diagram, above, illustrates the basic principle of layers of identity data.

- The core BIS is the minimum set of attributes accepted as being sufficient to establish the uniqueness of a given individual within a given population.
- The next ring consists of the kind of personally identifiable data which might not meet the BIS criteria, but is probably covered by national data protection laws or their equivalent. An example of this might be "current address".
- The outer ring consists of other attribute data associated with the individual, such as transaction histories. It also includes sector-specific data such as blood type, which might not in itself identify an individual, but is clearly useless unless correctly attributed to the right person.

It is interesting to note that one effect of increased computerisation (and increased computing power) is that data in the 'other attribute' category which might not previously have been sufficient to identify an individual might now be sufficient to do so. For instance, many web servers accumulate data about the browsing behaviour patterns of users; over time, interaction with a given website would allow the website owner to say, with reasonable certainty, whether a given user is the same one as visited the same website the previous day from the same IP address (as opposed to, say, a different member of the same household).

Some of the privacy implications of this are discussed in a Code of Practice for Personal Information Online (July 2010), published by the UK Information Commissioner's Office[4].

4 http://www.ico.gov.uk/upload/documents/library/data_protection/detailed_specialist_guides/personal_information_online_cop.pdf

Annotated "Onion" model

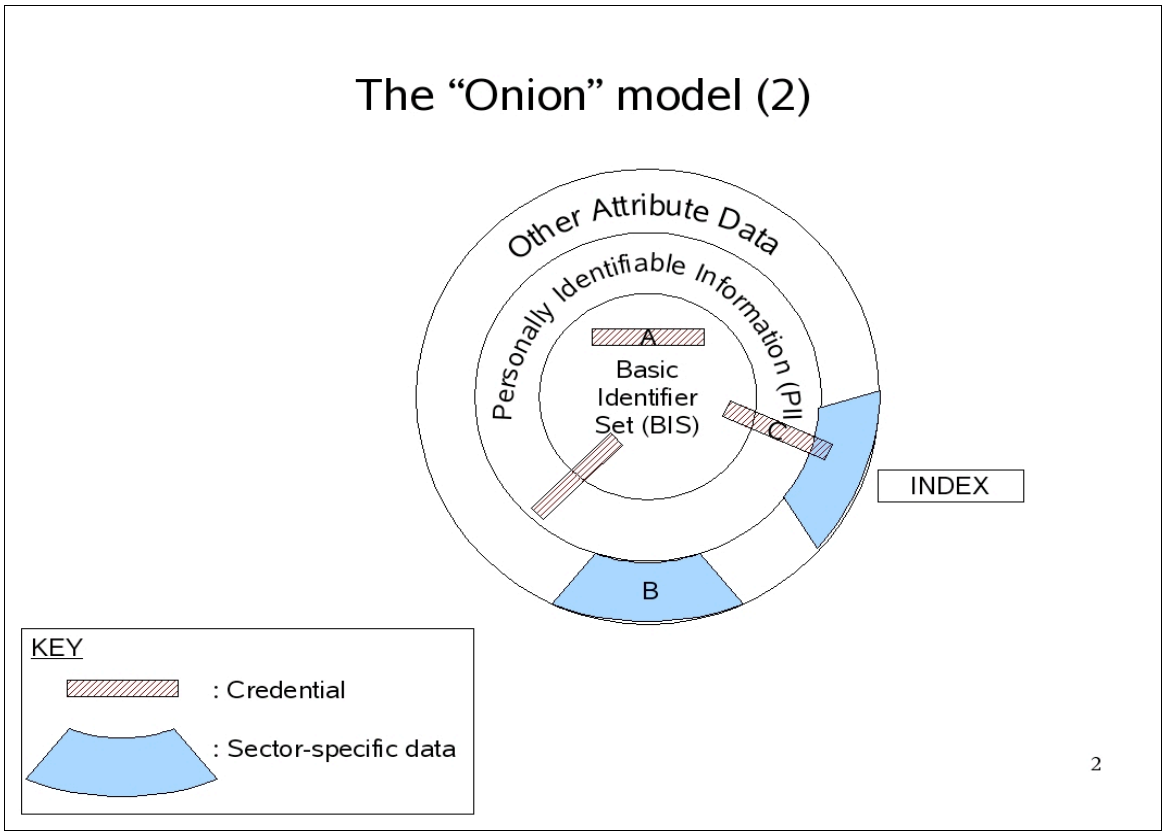

This version of the diagram is annotated to illustrate some further discussion points.

- **A** represents a credential which contains only those data items about the user which form part of the BIS.
- By contrast, **C** illustrates a credential which contains some items from all three rings. An example of such a credential might be a driving licence, which could contain the following:
    - BIS items such as name, date of birth, gender;
    - PII items such as current address;
    - Other attribute data such as 'entitlement to drive heavy goods vehicle'.
- As **B** suggests, the 'onion' will often tend to be divided into sector-specific wedges, some of which may rely on their own sector-specific credentials (such as the driving licence). This also reflects the importance, for privacy purposes, of what Helen Nissenbaum[5] refers to as "contextual integrity." That is: if I disclose personal data in one context (for instance, I tell my doctor my symptoms in order to get a diagnosis and treatment) I do not expect the same data to appear in an unrelated context (for instance, in a letter to the local newspaper). Violations of privacy often take the form of breaches of contextual integrity.

Even in the few years since the "Onion" model was first formulated, data mining and behavioural advertising techniques have grown in use and sophistication, giving rise to increasing concern about the privacy risk of data from the 'outer layers' of the Onion.

5 Helen Nissenbaum - "Privacy as Contextual Integrity" (pdf document)
http://www.nyu.edu/projects/nissenbaum/papers/washingtonlawreview.pdf

Three special cases

There are two other kinds of data which are relevant to the question of identity, but which are often not explicitly considered. These are:

- 'Index' values
- Shared secrets

Index values

Whenever sector-specific data about an individual is stored (for instance, by tax authorities, driver/vehicle licensing agencies and so on), there is almost inevitably an index value which is used to identify each unique record in that store. The index value may or may not appear on a sector-specific credential issued by that agency.

The importance of this point is that such indices can be over-exposed and over-used, and this can undermine the integrity of the identity data in question. An example of this is the US Social Security Number. This fulfils the role of an index to each citizen's social security records, but over time (despite laws to the contrary) has come to be used as a credential. As a result, there is now widespread inappropriate reliance on Social Security Numbers, and their utility as an identifier is greatly compromised.

Where an index exists, it is important that it be appropriately managed (and if necessary, subjected to quite different management disciplines from, say, the credentials associated with it). An example of this is the Norwegian government's policy for national identity numbers. These are, by default, not to be revealed – and applications wishing to use the national identity number as a means of indexing an individual's sector-specific records may only do so with specific legal permission.

Shared secrets

Two types of shared secret are relevant to us here. The first type is the shared secret which 'binds' a user to a given credential. A simple example is the PIN used to link a payment card with its holder. There is an assumption (written into the card's terms of use) that if the correct PIN is used, the user must either be the cardholder or someone to whom the cardholder has disclosed the PIN. Another example is the password associated with a user-ID. Both these examples illustrate shared secrets which it is possible for the legitimate user to pass to a third party.

One advantage often claimed for biometrics is that they represent a value which the legitimate user cannot, in most circumstances, pass on to someone else. Depending on the biometric in question, this is not necessarily true in practice, but that is a problem which is beyond the scope of this document and will not be considered here for the time being.

In all these cases, there is a basic design principle to do with storage of the shared secret. It is good practice for shared secrets not to be stored in readable form; this does not necessarily mean they cannot be compared with the value presented by the user. For instance, one approach is to feed passwords through a one-way hash before storing them. When a user enters their password to authenticate, the entered value can be hashed using the same algorithm, and the result compared with the stored value. Similar hashes can sometimes be used in the case of biometric data.

I do not wish to understate the complexity of such an approach – and the state of the art in this area continues to evolve rapidly. A 2002 report by the US National Research Council[6] examines some of the problems to be overcome – such as the question of how to achieve reliably consistent binary hash values of what are, after all, rather 'analogue' biological characteristics (such as iris structure, fingerprints or vein patterns). More recently, a UK technology firm, Touch2ID[7], appears to have implemented a robust algorithm for this, at least where applied to fingerprint biometrics.

6 http://www.nap.edu/catalog.php?record_id=10346 "IDs – Not That Easy" - Kent and Millett (Eds.)
7 http://www.touch2id.co.uk/ Touch2ID – database-less proof of age with biometric authentication

Depending on the threat model, such hashing may be supplemented with encryption for further security – though at the cost of an increased burden of complexity and key management.

The other kind of shared secret which bears special consideration is “password recovery” data. That is, the pieces of information a user lodges with a service provider so as to have a fallback authentication mechanism in case they forget or lose their password. Clearly, if an attacker has access to this password recovery data, it is possible for them to impersonate the real user and lock them out of their account. However, anecdotal evidence suggests that password recovery data is often much more weakly protected than password data.

“To point or to store... that is the question”

A further note about credentials is that, for privacy reasons, some take the view that the closer a credential stays to the centre of this onion model, the better: that is, credentials should serve to identify the individual, but not necessarily be loaded with attributes and other personal data.

By analogy, imagine that, in order to establish your entitlement to buy alcohol in a bar, you show your driving licence. The bar staff only need to know that you are over the required age – but the credential might also reveal to them your date of birth, place of birth, current address, driver/licence index number, which types of vehicle you are entitled to drive, and possibly any endorsements you have.

In the online environment, where all that is needed is a pointer to the authoritative source of that information, it seems a sound principle that credentials should gravitate towards the centre of the onion, and point to, rather than hold, PII and attribute data.

By implication, this means that the more centralised a repository is in its design, the more attractive it is for the system to focus on 'proof of uniqueness' as opposed to broader sets of PII and sector-specific data. As the subsequent models will show, this is not a guarantee of 'unlinkability' (if that is an objective of the design), but may contribute towards it.

I am grateful to David Chadwick (University of Kent) for clarifying a useful distinction between authentication credentials and authorisation credentials. What I describe above – as credentials which tend to gravitate towards the centre of the onion – are authentication credentials. They are useful for establishing your uniqueness, and serve to establish your entitlement to something based either on simply 'who you are', or on other attributes to which the authentication credential reliably links you.

However one can also define authorisation credentials, which are those which gravitate towards the outer layer of the onion; these are credentials which assert something about me without necessarily revealing my identity. In the academic community, a common example is “*x* is a member of this institution” - where it is not necessary to identify *x* in order to approve their access... just to establish that they are a member of a qualifying institution. Similarly, “is over 18” and “has blood type O” are not enough to uniquely identify me, but may well be sufficient to establish my entitlement to something.

Authorisation credentials are often also referred to as 'attribute assertions' – though, of course, an assertion of data from someone's Basic Identifier Set (BIS) is also, strictly speaking, an assertion of attributes.

Conclusions

The landscape of digital identity is shaped by many things: economic imperatives, political goals, cultural attitudes, technical possibilities, and the *de facto* results of individuals' behaviour online.

As I hope this paper has helped to explain, digital identity and privacy are societal issues as much as (and probably more than) technological ones. On the one hand, current technologies give those societal issues new ways of emerging; on the other, they present us with a constantly-evolving range of tools to work with.

I also hope I have illustrated the principle I stated at the outset: that digital identity and privacy are multi-stakeholder phenomena, and no single stakeholder has all the answers.

Add to that the speed with which online activity evolves and adapts, and it becomes clear that the way in which diverse stakeholders discuss their shared problem has to mature, and it has to do so fast. To put it bluntly, we cannot afford to re-define the core concepts from scratch every time we convene to define and solve problems of online identity and privacy.

There is a pressing need for all of us, whether from a background of politics, commerce, academia, law, civil society or elsewhere, to attack these problems with a sound basis of common concepts and shared vocabulary, and that is what I hope this document will help to establish.

In addition to the individuals credited at the beginning of the document, it is also right to acknowledge the role played by a number of organisations in helping it gestate:

- This report appears under the imprint of the Kantara Initiative, and I am both grateful for and proud of that organisation's ability to assemble the diversity of participation which makes this kind of framework document possible;
- The Liberty Alliance, especially the Public Policy Expert Group which ran the programme of Privacy Summit meetings from which much of this material came;
- The Net-ID conference series, which has co-hosted several of those Privacy Summit meetings;
- The Internet Society, which has helped extend and broaden our work with the legal community.