www.icele.org

Net-ID 2008
Identity, Trust, Privacy and Security in Europe

LIBERTY ALLIANCE PROJECT

Report of the
Privacy Summits

hosted by the
Public Policy Expert Group of the Liberty Alliance
in conjunction with the
ICELE e-Participation Symposium 2008
and the
Net-ID 2008 conference

Thursday February 28th 2008, London
and
Monday March 3rd, Basel

Foreword

This document presents a joint report of the fourth and fifth Privacy Summits, held in conjunction with the ICELE Symposium on e-Participation 2008 and the NetID 2008 conference, in London and Basel respectively. I would like to express my gratitude to Dr Julia Glidden of 21C Consulting and Ms Stefanie Geuhs of Computas for their help in arranging for the Privacy Summits to be held under the auspices of these two events.

“Why only one report for two Summits?”, you may be wondering. The close timing of the two 'host' conferences meant that these two Summits happened in quick succession, only days apart; subjectively, I felt there were more common themes than usual across the two. Also, as the Privacy Summit series evolves and matures, we have found that some of the basic elements of each Summit are consistent – specifically, the same basic models are often used to introduce the session - and then discussion continues in ways which are more specific to each event.

Therefore what I have chosen to do with this report is briefly to re-cap the basic models for identity and privacy, adding two more which have been developed during the series, and then summarise those discussion topics which followed in each of these two Summits.

Some other contextual factors stayed the same for the London and Basel Summits:

- The aim of the Privacy Summits continues to be to bring together a wide range of stakeholders, so as to stimulate a multi-disciplinary discussion of the technical, legal, social and other perspectives on identity and privacy. To that end, the participants included academics, policy-makers, user organisations from the public and commercial sectors, technologists, industry analysts and so on.
- The discussion was held under the 'Chatham House Rule'[1] – so the participants are free to repeat what was discussed, but may not reveal who said what, nor the affiliation of any of the participants.
- This document is an attempt to summarise the topics discussed; it is not an exhaustive record of the meeting, and participants are encouraged to reply with any further notes or comments they would like to add to the output of the Summit.

Anyone wishing to comment on, add to or correct this document should do so by sending an email to robin.wilton@sun.com.

The Summit Reports are published via the Liberty Alliance website, here:

http://www.projectliberty.org/liberty/public_community/privacy_summits

Many thanks for your participation

and your contribution to this ongoing programme.

Dr. Hellmuth Broda - Chairman

Robin Wilton - Moderator

1 The Chatham House Rule

"When a meeting, or part thereof, is held under the Chatham House Rule, participants are free to use the information received, but neither the identity nor the affiliation of the speaker(s), nor that of any other participant, may be revealed" (http://www.chathamhouse.org.uk/index.php?id=14).

Document date: 4th July 2008

Document reference: PS-Lon-Bas-2008

Table of Contents

The Basic Models – updated

As the programme of Privacy Summits has progressed, we have found a number of simple models useful in guiding the discussion. The core set of simple models currently stands at four, each of which helps to introduce a particular aspect of digital identity and privacy.

In this section I will first describe the first two new models (the "Stakeholder" model and the "Ladder" model), which help to explain why this is so often necessary. I will then briefly re-introduce the "Onion" and "Silo" models so that they are at least all present in a single document. For the more detailed background on the "Onion" and "Silo" models, you are invited to refer back to the Brussels Privacy Summit Report[2].

Two New Models

1 - The "Stakeholder" model

The 'stakeholder' model emerged from discussion of 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".

2 http://www.projectliberty.org/liberty/public_community/privacy_summits

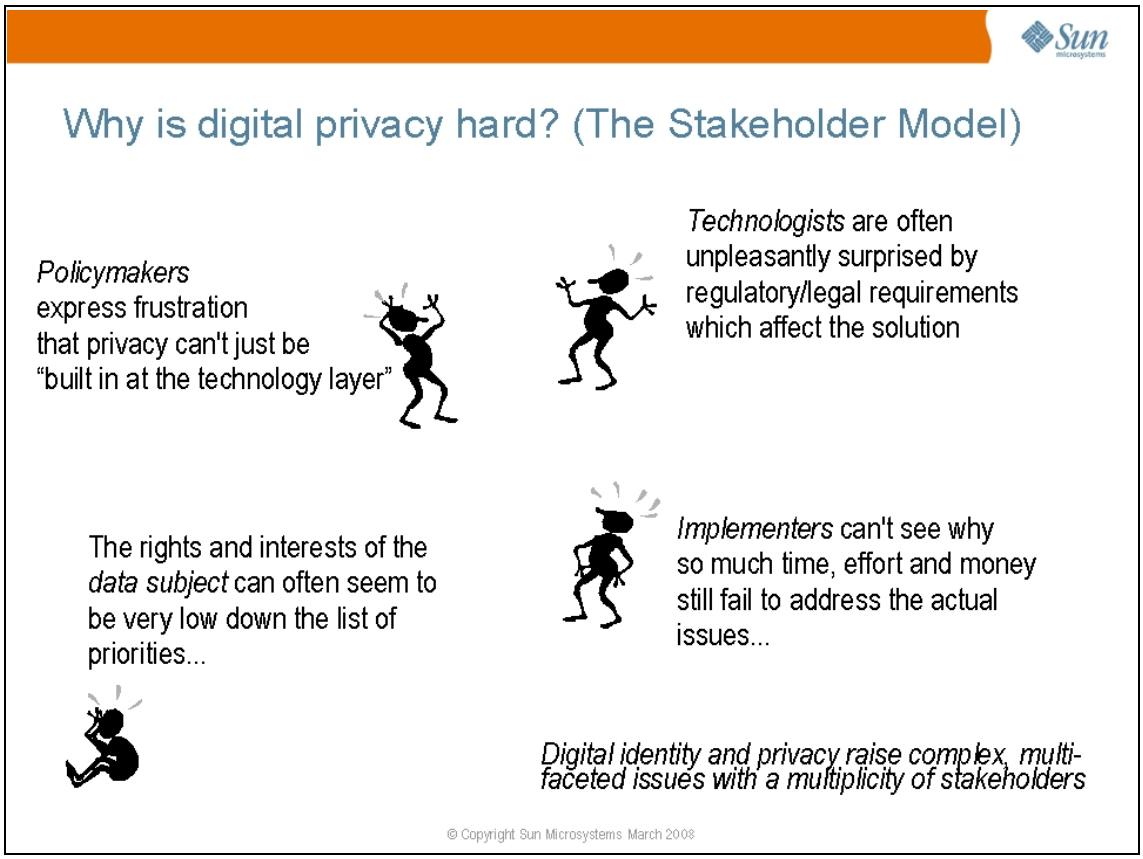

We drew a number of conclusions from these and other comments. The first was simply to confirm a principle which has been stated in the invitations to every Privacy Summit: "The identity and privacy challenges which confront us cannot be solved by a 'single-stakeholder' approach." This is true both of the 'problem analysis' phase (which the Privacy Summits seek to address) and the 'solution development' phase, which is the broader realm of the Liberty Alliance and its members.

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"].

Another important factor which the 'stakeholder' model hints at is this: 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 better termed "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" and "Silo" models, and the notes from the Basel Summit (Appendix B).

The answer to the 'multi-stakeholder' 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, which was in fact suggested by the Brussels Summit, took the form of the second new model, the "Ladder".

2 - 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 means by "trust", "personal information" or "ownership of personal data" may be very different to what another stakeholder means by the same term.

This rapidly emerged as one of the key challenges which the Privacy Summits would have to address, if the format and the programme were to be of any use. The "Onion" and "Silo" models were first steps towards establishing a common conceptual framework which might allow very diverse stakeholders to have a productive discussion despite their different perspectives, assumptions and vocabularies.

Over time, we came to realise that such a conceptual framework was itself just one part of what was needed in order to understand the complex dynamic of a multi-stakeholder discussion. After further thought, we sketched out the "Ladder" model.

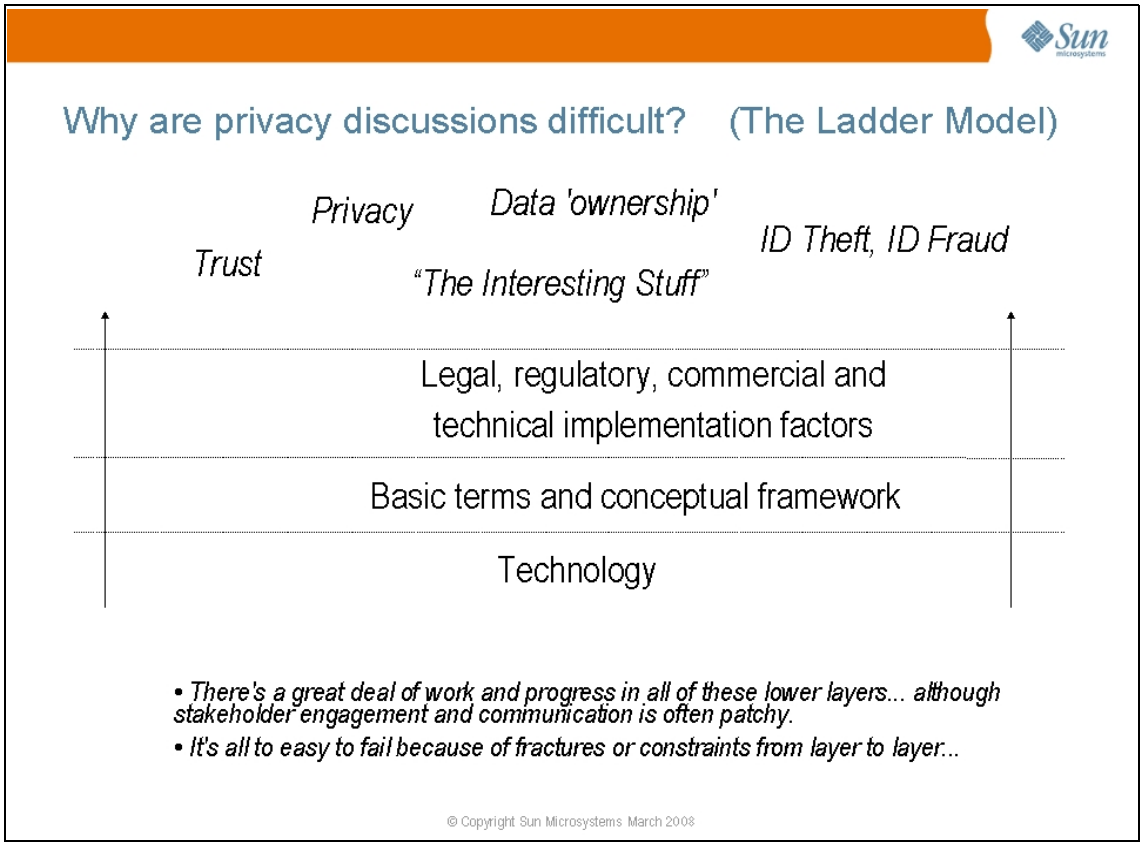

What we found (and this is partly simply a reflection of the 'technical' origins of the Privacy Summit programme) was that it is all too easy for a privacy discussion to start at the bottom, technology layer and work its way upwards from there. There is nothing inherently wrong with this approach, except that by default, it can unwittingly open the way to a number of shortcomings, and ultimately, can mean that the top-level topics (which I refer to as the 'second-order' concepts) are not meaningfully addressed.

The picture we built up 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 between 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' through 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. Let's look at 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, what we think is happening here is this: the 'customer' tends to think of their ultimate requirements in terms of the 'second-order' 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 "second-order 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.

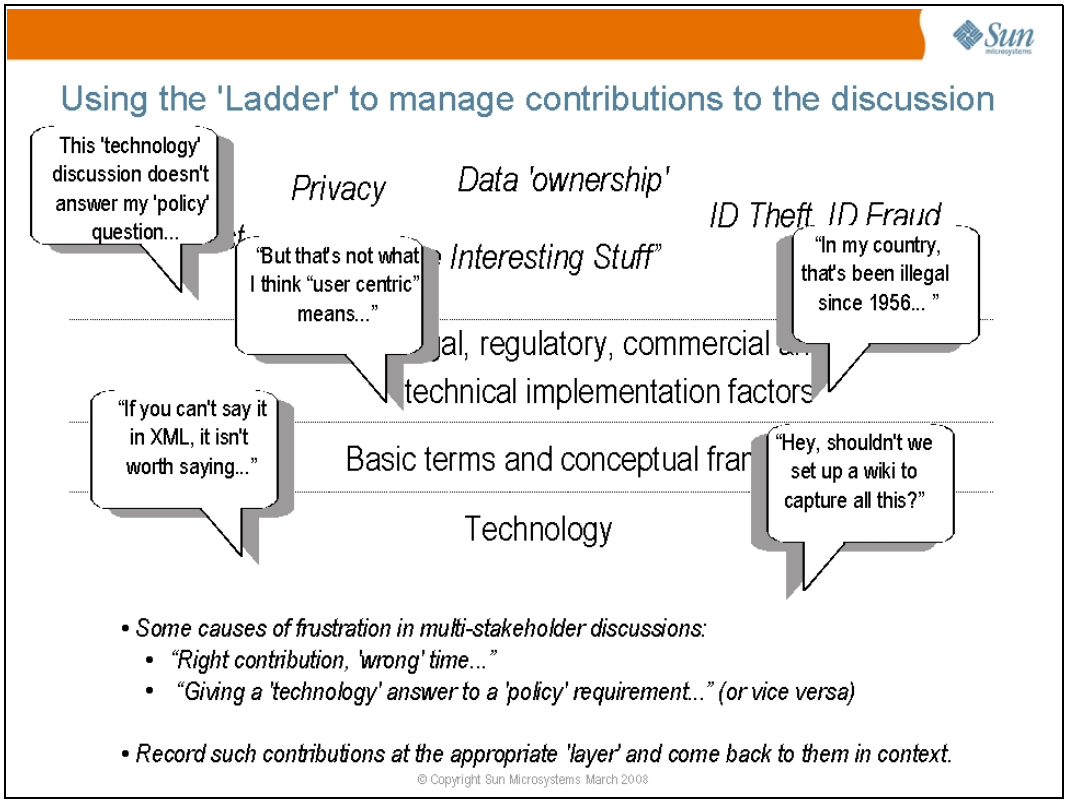

This helped put into context the models which had already arisen out of the previous Summits – for instance, we could now clearly position the "Onion" and "Silo" models as tools for establishing that common terminology and conceptual framework.

Existing Models

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.

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. 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. 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

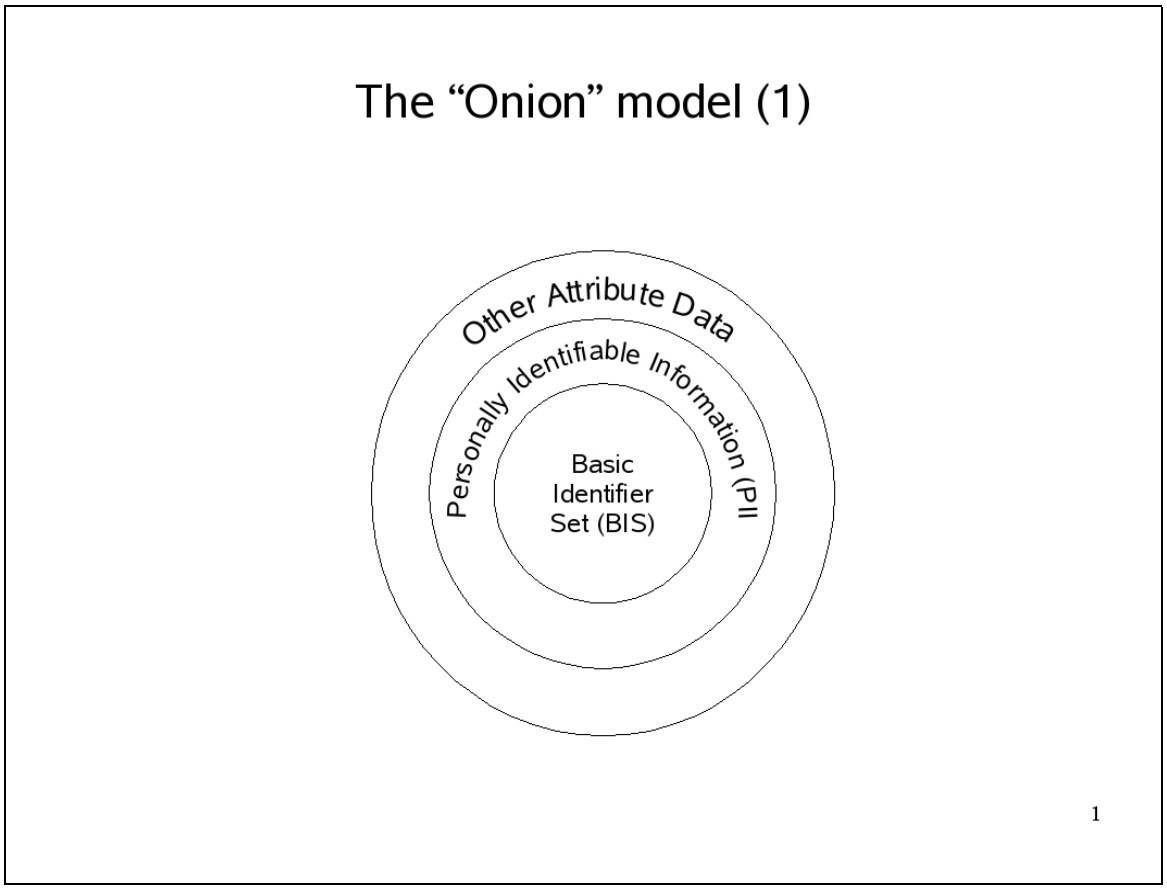

The first diagram, above, illustrates the basic principle of layers of identity data.

- The core BIS is the minimum set of attributes accepted as being sufficient to establish the uniqueness of a given individual.
- 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).

Annotated "Onion" model

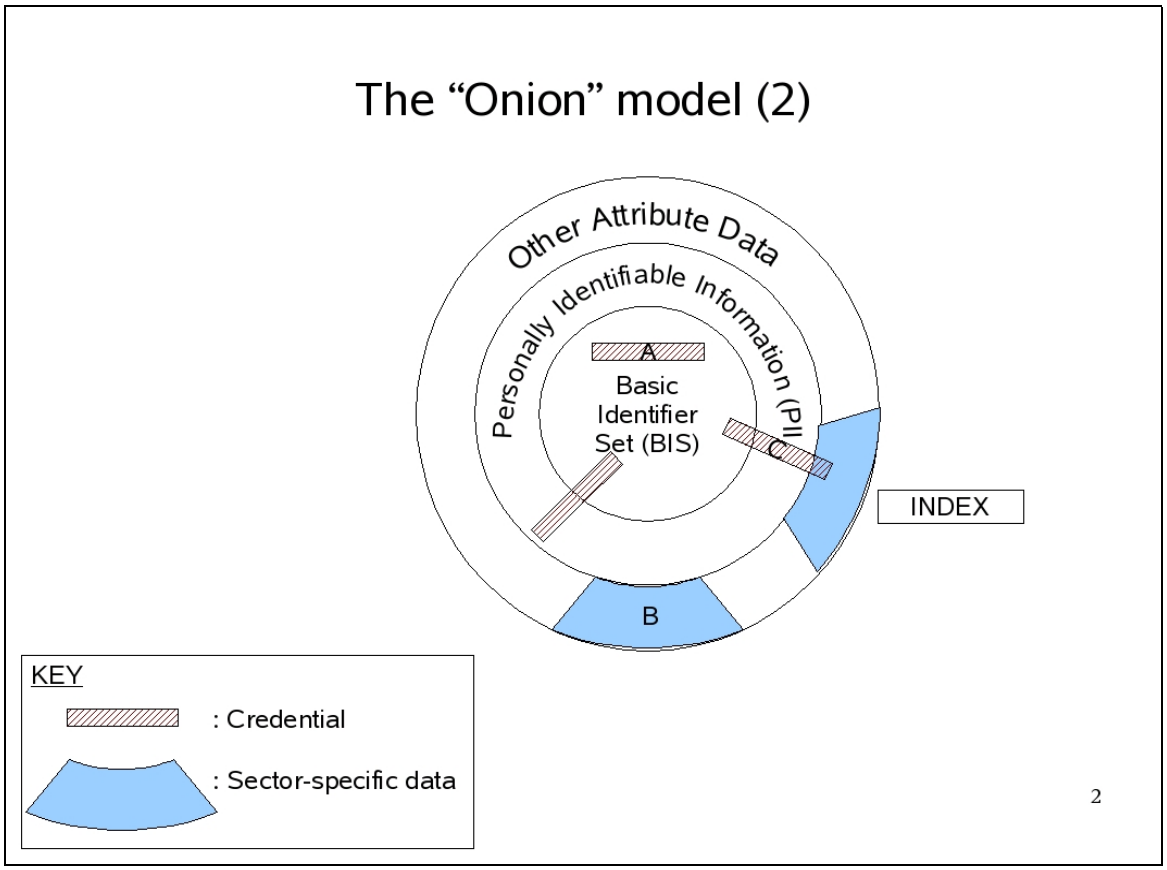

This version of the diagram is annotated to illustrate some further useful points.

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).

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. This is not strictly true in practice, but that is a problem which is beyond the scope of this report 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. 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 governments 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.

Authentication vs. Authorisation

I am grateful to David Chadwick (University of Kent) for clarifying a useful distinction here 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 someone's BIS is also, strictly speaking, an assertion of attributes.

The “Silo” model

As noted in the summary document from the Berlin meeting, the Austrian ID Cards implementation can be used to illustrate a number of useful concepts. A concise description of the Austrian scheme can be found here[4] on the European Commission's IDABC website, and a web search using the argument “Austrian Citizen Card” will return a wide range of further documentation.

In the Austrian government example as described, a single state-issued identifier is used as the basis for multiple sector-specific identifiers – which can only be correlated by the Privacy Commissioner. The sector-specific identifiers are generated using one-way cryptographic functions, so that the sector-specific identifier can easily be derived from the original identifier on the citizen card, but conversely, the original identifier cannot be derived from the sector-specific identifiers. This gives rise to an architecture as illustrated in the diagrams below.

The first diagram allows us to relate this multi-sector view to the 'onion' model described above. Each silo will usually contain a sector-specific set of information about the individual, corresponding to a 'wedge of the onion' – that is, some or all of the Basic Identifier Set (BIS), some Personally Identifiable Information, and other data such as transaction history, entitlements and so on.

Note that this architectural model applies equally to public- and commercial-sector systems, and to intra-organisational as well as inter-organisational systems. Within a single organisation, this model illustrates a classic 'application silo' set-up. Between organisations, it illustrates a typical 'distributed identity' set-up. Each organisation (or application) stores the information it needs, regardless of whether this results in duplication. As far as the user is concerned, the applications appear disjunct.

Basic silo model

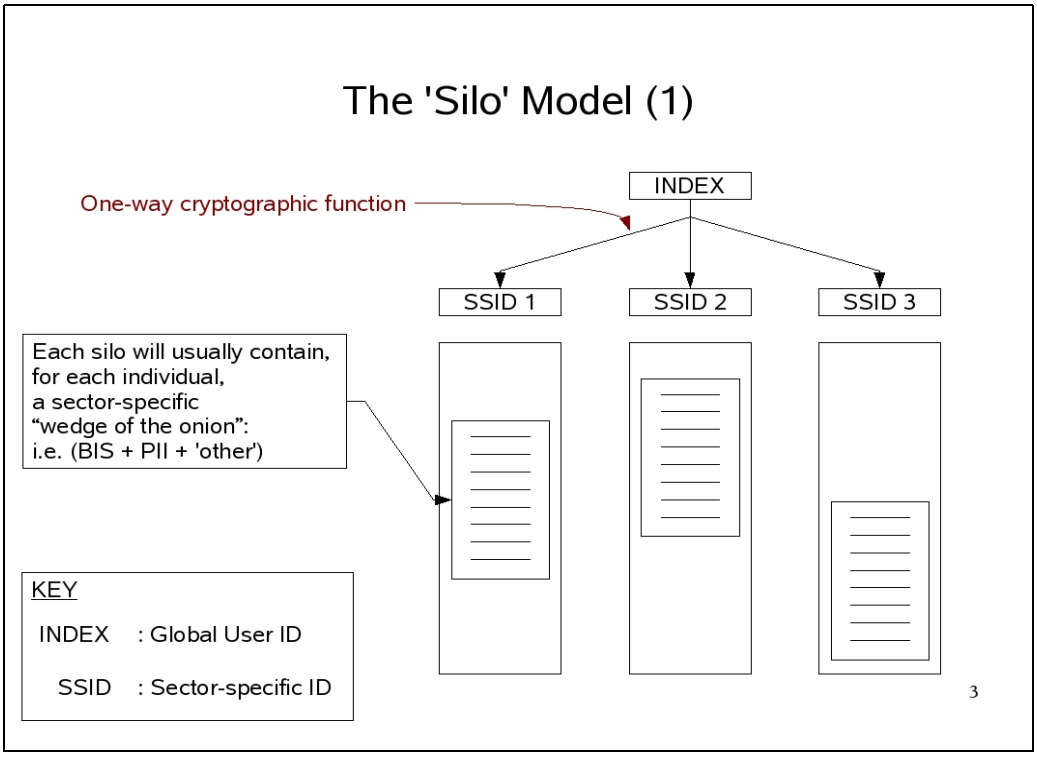

4 http://ec.europa.eu/idabc/en/document/4486/5584

The second silo diagram illustrates some further points about identity in distributed applications. First, note that each silo can be regarded as a 'context', within which the user discloses some information to the service provider. For two silos to be federated, the appropriate technical and non-technical measures need to be in place to establish the context within which the user will disclose information to both service providers.

The Austrian example uses specific technical means to enforce a particular relationship between the INDEX and the SSIDs. In other cases, those relationships may be defined and enforced through policy measures or through different federation technologies. Similarly, data exchange between contexts may be controlled through technical or non-technical measures.

The second diagram also illustrates the principle that having a technically-enforced one-way relationship between the INDEX and the SSIDs does not, in itself, guarantee that data between contexts cannot be linked and attributed to the same user. For example, anyone able to search each database looking for a given BIS, or an attribute such as Postal Code, would quickly be able to find the sector-specific records relating to a given individual, even if they did not know the INDEX or SSID for that individual.

Thus, if 'unlinkability' is a requirement, it must be enforced through good data custody practices as much as through any technical means at the SSID level.

If the user is to trust a system such as this, the system must contain adequate (technical and non-technical) measures to deal with context, user consent, and controls over the exchange of data. These measures must be able to cope with different credentials, different levels of trust, and different control mechanisms between, for example, public sector and commercial sector contexts.

Annotated silo model

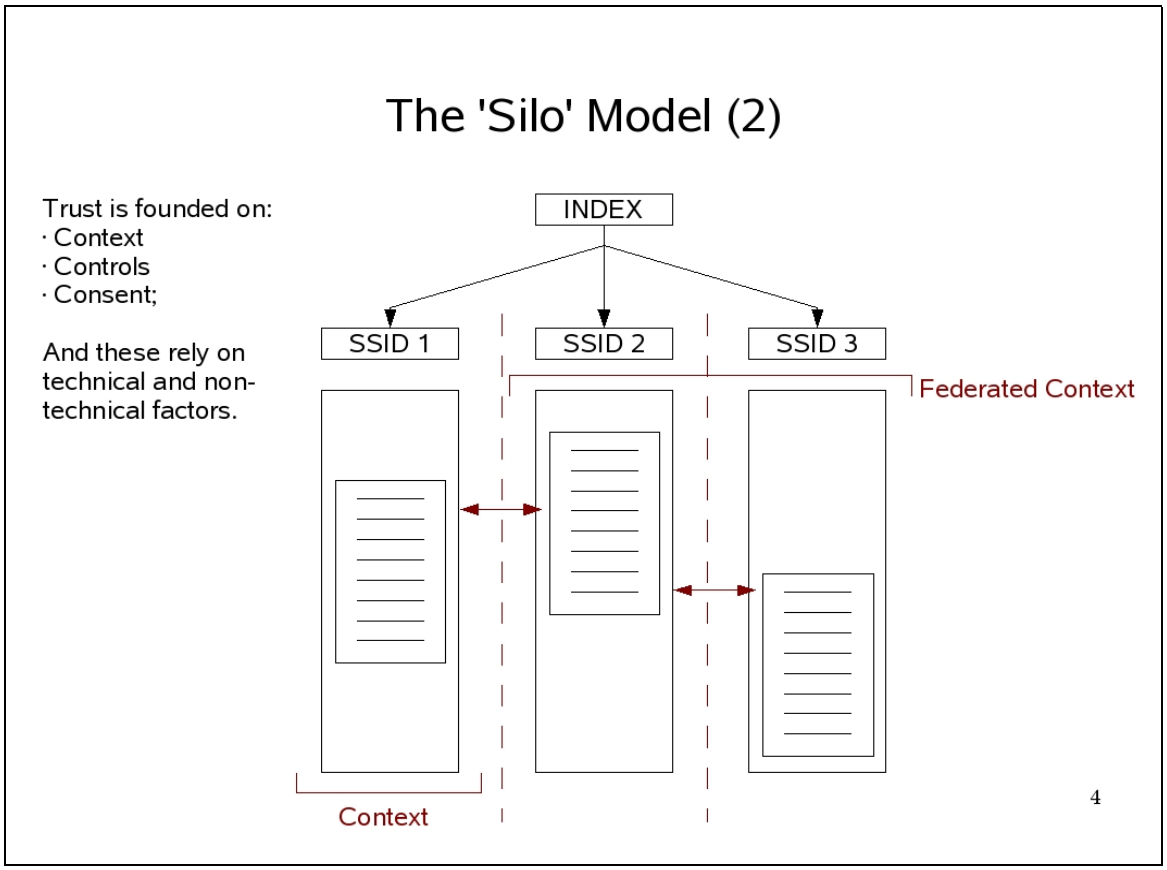

Appendix A - Summary of themes discussed at the London Summit

For completeness, and to serve as a 'memory-jogger' for those who took part, this appendix records – in 'raw' form - the notes captured on flip-chart in the course of the discussion.

Data Exchange/Sharing:

1 - Within the commercial sector
2 - Within the public sector
3 - Across sectoral boundaries

An understanding of identity in the sectoral ('onion') and contextual ('silo') senses was felt to be a basic prerequisite of any discussion of solutions in these areas. Data sharing across sectoral boundaries raises the additional possibility of differences in the legal and regulatory requirements, and in the areas of liability and recourse should something go wrong for any of the parties involved.

"Citizen" Context

We noted the importance of the following factors when establishing systems for identity and privacy relating to individuals acting in their capacity as 'citizens'... that is, generally cases where the user may have little or no option but to make use of the service in question (for instance, dependence on benefits or medical treatment, cases of legal compliance, crossing frontiers and so on). These cases may often differ from 'consumer' use cases in important aspects – for instance, the user may not have the option to choose between service providers, or to choose not to disclose the requested data. Also, as noted above, there may be differences in terms of compliance, liability and recourse.

- Intelligibility: the view was expressed that this may need to go beyond the provisions of a 'fair processing notice', in that it needs to be clear to the citizen, over time, what the totality of their consents amounts to, and how the cumulative effect of disclosures may be expected to affect them.

- Consent and control: the view was that this needs to apply not just to the citizen's data itself, but also to related meta-data, such as information about other parties' requests for access to the subject's data.

- Governance: in case of exceptions to the principle above (consent and control), the view was that it becomes all the more important for citizens to have well-founded confidence that their data is being appropriately safeguarded. For instance, it may not be appropriate for a data subject to know that there has been a law enforcement access to their personal data – but the citizen nevertheless has a right to expect that the data is legally and appropriately processed.

- Intention: this was proposed as a somewhat broader concept than “purpose of use”, in the sense that, for instance, a system might collect data for which a purpose of use must be declared, but that purpose of use might relate to a broader intention on the part of the data controller. Again, particularly in the 'citizen' context, it was felt important that there should be transparency where such purposes are concerned.

- Expectations of data separation (as tends to happen by default in a 'siloed' model)

– Is there an argument for 'protecting people's privacy despite them'?
  - “Rules + Context = Privacy Value”
  - What's the appropriate response to a naive privacy culture and risk assessment?

Privacy, Consent and User Centricity

Consent is contextual: Give, Revoke, See

1 - Visibility (cumulative 'cloud' snapshots)
- Inline, vs. 'big picture'
2 - Manageability
3 - Proxy use-cases, 'brokerage'

"Beyond First Disclosure"

- DRM: differing attitudes with respect to 'protected published media content' vs. ' my P.I.I.'
- Where's the Policy enforcement point?
- Insider/Human attack

The "Ladder" Model:

- The Interesting Stuff
- Implementation Factors
- Basic Terms and Conceptual Framework
- Technology

A second "Ladder" model - for data ownership:

Social/Cultural

Technological

Contractual (market)

Constitutional/Legislative

"Self-evident truth"

Another Way of Putting it all in Context:

Human Rights Act Sect.8.1

Infringements (of the right to privacy) must be:
- Predictable
- Proportionate
- 'Necessary' (i.e. the least privacy-infringing alternative)

- 'pressing social need'
- law enforcement
- national security

Relevance of Human Rights legislation:

"The rule of law embodies the basic principles of equal treatment of all people before the law, fairness, and a guarantee of basic human rights. A predictable and proportionate legal system with fair, transparent, and effective judicial institutions is essential to the protection of both citizens and commerce against any arbitrary use of state authority and unlawful acts of both organisations and individuals" (Report from a Review of the Regulatory Framework for Legal Services)

Illustrative hypothetical use-case - Social Care and data sharing:

1 - Data Subjects' powerful motivations for getting out of such a system;
2 - Risk of insider abuse;
3 - Audit/monitoring ('internal surveillance')
4 - User-centric control? Vital importance of user's feeling of 'informational self-determination'
Semantic interoperability issues

Risk Assessment based on population statistics:

- # of likely cases (i.e. risk to mitigate)
- # of authorised users
- %age of probable insider attacks

- cost of enforcement
- cost of counter-measure/mitigation
- threat model
- appropriateness of mitigation

Appendix B - Summary of themes discussed at the Basel Summit

For completeness, and to serve as a 'memory-jogger' for those who took part, this appendix records – in 'raw' form - the notes captured on flip-chart in the course of the discussion.

The Digital Footprint

Is made of many parts... of varying sensitivity

Raises questions of:

- Ephemerization
- Time-sensitive disclosures/contexts
- Policy enforcement points and their placement

(Is not the same as the digital sole... ;^)

Diverse Contextual Factors

- Context
- Sector
- Purpose (of collection, of use)
- Scope of usage (e.g. employer, tax, national, international...)
- Time (consent may change over time...)
- Mitigation of Risk
- Quality of Service
- Consent (or otherwise)

Can also be summarised as "what, to whom, how, why, when, where..."

- “Ownership” is another contextual factor
- Yet another may be based on 'rights' to have the PII used in a particular way (or not).

E.g. Use of patient data -

- For treatment
- For a case study
- For a clinical trial