

Unclassified

DSTI/DOC(2007)7



Organisation de Coopération et de Développement Economiques
Organisation for Economic Co-operation and Development

29-Feb-2008

English text only

DIRECTORATE FOR SCIENCE, TECHNOLOGY AND INDUSTRY

DSTI/DOC(2007)7
Unclassified

AT A CROSSROADS: "PERSONHOOD" AND DIGITAL IDENTITY IN THE INFORMATION SOCIETY

STI WORKING PAPER 2007/7

Information and Communication Technologies

JT03241547

Document complet disponible sur OLIS dans son format d'origine
Complete document available on OLIS in its original format

English text only

STI WORKING PAPER SERIES

The Working Paper series of the OECD Directorate for Science, Technology and Industry is designed to make available to a wider readership selected studies prepared by staff in the Directorate or by outside consultants working on OECD projects. The papers included in the series cover a broad range of issues, of both a technical and policy-analytical nature, in the areas of work of the DSTI. The Working Papers are generally available only in their original language – English or French – with a summary in the other.

Comments on the papers are invited, and should be sent to the Directorate for Science, Technology and Industry, OECD, 2 rue André-Pascal, 75775 Paris Cedex 16, France.

The opinions expressed in these papers are the sole responsibility of the author(s) and do not necessarily reflect those of the OECD or of the governments of its member countries.

<http://www.oecd.org/sti/working-papers>

All OECD Working papers on ICT can be found at:

<http://www.oecd.org/sti/ict/reports>

**AT A CROSSROADS: “PERSONHOOD” AND DIGITAL IDENTITY
IN THE INFORMATION SOCIETY¹**

Executive Summary.....	4
Introduction.....	6
Definitions	7
From “Personhood” to Digital Identity.....	9
Data Protection in the IDM-Enabled Ubiquitous Information Environment.....	12
Data Protection and User Control.....	17
Market Demand for User Control.....	22
The Properties of Identity	26
The Properties of Identity and Data Protection.....	27
The Properties of Identity for Policy makers and Software Developers.....	37
Current Conceptions of IDM	39
Decisions and Constraints.....	47
Conclusion.....	49
Annex: OECD <i>Privacy Guidelines</i> (Excerpt).....	i

Tsze-lu said, “The ruler of Wei has been waiting for you, in order with you to administer the government. What will you consider the first thing to be done?”
The Master replied, “What is necessary is to rectify names.”
(Confucius, *Analects* XIII, 3, tr. Legge)

¹ This paper has been prepared by a team of authors, led by Mary Rundle (Managing Editor) and including Bob Blakley, Jeff Broberg, Anthony Nadalin, Dale Olds, Mary Ruddy, Marcelo Thompson Mello Guimarães, and Paul Trevithick, with authorship noted according to sections. The authors are grateful to Wendy Seltzer and Ioanna Tourkohoriti for reviewing sections of the report.

The Managing Editor's work was conducted under the Net Dialogue project (see <http://cyber.law.harvard.edu/home/research/netdialogue>), funded by the Lynde and Harry Bradley Foundation.

EXECUTIVE SUMMARY

In its “Introduction”, the paper sets the scene: Law and technology must be crafted to respect certain “Properties of Identity” in identity management (IDM) in order for the information society to be free and open. Respect for the Properties of Identity is necessary for data protection; data protection is necessary for accountability; and accountability is necessary for trust.

Before advancing arguments, the paper sets out some definitions of terms it uses.

The first substantive part of the paper, “From ‘Personhood’ to Digital Identity”, looks at the issue of “personhood” – or the recognition of a person as having status as a person – in light of two highly influential strands of classical philosophy that influence today’s conceptions of data protection. Despite differences in view over the means, respect for “personhood” is a shared value among countries holding to democracy and an open economy. As IDM systems become more prevalent, data protection can help defend “personhood” and allow people to enjoy greater autonomy by exercising control over their digital identities.

To show some threats that may arise if a sufficiently protective framework for identity information is not in place, the section on “Data Protection in the IDM-Enabled Ubiquitous Information Environment” tells a story. Here the paper looks at emergent information and communication technologies (ICT) and postulates that IDM promises to be a unifying component. With IDM all-pervading, data protection will prove vital.

The paper then addresses “Data Protection and User Control”. Here it suggests that IDM systems must be built with fair information practices in mind.

The section on “Market Demand for User Control” deals with the question of whether the market will support user control in IDM. Trends in demand from individual users and business seem to suggest it will. As a result, organisations will need to transform their thinking and business processes. Among other changes, they will need to: *i*) build appropriate notice, consent, security, and access into business process design, *ii*) limit data collection in transactions, and *iii*) securely dispose of information that is no longer required. Not surprisingly, these are key concepts for data protection.

As the market demands IDM systems that protect data and give control to users, people responsible for designing sound legal and technological systems relating to IDM will need to know that their designs will hold up under pressure. Fundamentally, they need to factor in the way identity behaves. To help them do so, the paper shifts to introduce the Properties of Identity.

The Properties of Identity can serve as a guide for data protection and so help undergird a free and open information society. With this in mind, the section on “The Properties of Identity and Data Protection” explores the adequacy of the OECD’s Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (*Privacy Guidelines*) for IDM.

While the relationship between the Properties of Identity and data protection may be clear, it is people who must bring these ideas to life. The section on “The Properties of Identity for Policy makers and Software Developers” tells how people in government and the IDM industry have distinct roles to play.

Even if the logic of the Properties of Identity and data protection seem obvious, there is still the question of how the identity infrastructure will get from here to there. The paper describes “Current Conceptions” of IDM, shedding light on “core” identity information for use in various IDM contexts, the role of individual user control, identity information that does not need to be commonly used, and the extent to which core identity is compatible with partial identities and pseudonyms. In addition, this section maps out current conceptions of the management of identity information, indicating similarities and differences among IDM approaches – user-centric, service provider/organisation-centric, and network-centric/federated.

To bring discussion back around to immediate issues facing leaders, a section on “Decisions and Constraints” first lists some decisions that must be made in the near term regarding IDM policy and technology. It then calls to mind some of the constraints that set the larger context within which these decisions must be made.

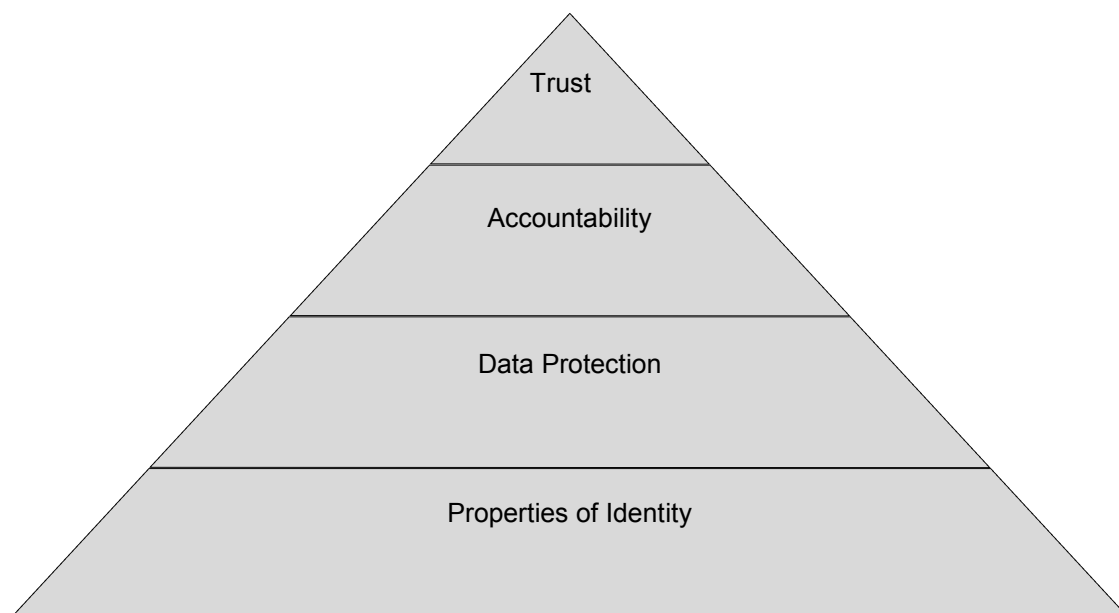
The paper concludes that, given the importance of these issues for the future information society, more investigation is needed into how to address gaps in international data protection in light of the emergent identity infrastructure.

INTRODUCTION

There is a growing sense in the online environment that a free and open society may not be as certain as previously assumed. With a lack of identity controls, society will be susceptible to identity theft, fraud, and the shutting down of businesses and even news media through denial of service attacks. As emergent technologies bring the information society to uncharted territory, even people who see data protection as providing guidance are questioning the adequacy of safeguards conceived years ago.

In 1980 OECD members adopted the *Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, and those *Privacy Guidelines* have remained relevant to this day.² However, as society starts to head into the ubiquitous information environment, a key question is whether those data protection principles need bolstering. In particular, are they capable of protecting data when it is separated from the control of the individual to whom it relates?

In terms of identity management (IDM), unless law and technology are crafted to respect certain “Properties of Identity”, there is no data protection; and if there is no data protection, there is no accountability; and if there is no accountability, there is no trust. The diagram below depicts how these elements build upon each other.



The paper elaborates on these concepts. A common theme is the importance of user control.

² For ease of reference, an excerpt of the Privacy Guidelines is provided in the Annex.

DEFINITIONS

This section sets out some basic definitions of concepts as used in this paper.

As this paper uses the term “person”, it refers to a human being, or a natural person. The paper’s arguments could be adapted to apply to juridical persons (*e.g.* corporations) as well.³

“**Personhood**” is used in the traditional world to mean recognition of an individual or entity as having status as a person. This paper uses the term “personhood” or “**digital personhood**” to discuss recognition of a human being as having status as a person in the electronic realm.

Identity is both a “real-world” concept and a digital artifact; this paper uses the term “**digital identity**” or “**identity**” to refer to what technologists in the field of IDM conceive as “a digital representation of a set of claims made by one party about itself or another data subject.”⁴ As in the real world, a person may have any number of different identities in the electronic world. In the real world identity is considered to entail a rather comprehensive set of “individual characteristics by which a thing or person is recognised or known,”⁵ whereas in the electronic realm an identity can be a very simple subset of identity information (*e.g.* an address). Despite the paper’s discussion of the philosophical concept of personal identity as the “sameness of a same person in different moments in time”, the term “identity” as used in the paper refers to that more limited notion of a set of claims. Digital identity, for the paper, is a “thing”, a man-made thing (an “artifact”) that refers to a person, and that is different from such person.⁶

The term “**partial identity**” is used to refer to subsets of identity information as the “thing” may not be sufficient to identify a person at different moments in time.

The term “**identity attributes**” is sometimes used to refer to the contents of those partial identities or digital identities.

The term “**identifier**” is sometimes used to refer to information that points to a person.⁷

A person acting through digital identities may be familiar to others due to **personas** that he himself develops (with a persona being “the role that one assumes or displays in public or society; one’s public image or personality, as distinguished from the inner self”⁸). In addition, a person acting through digital

³ See, for example, discussion of the referential property in the section on “The Properties of Identity and Data Protection”.

⁴ This definition was developed on the mailing list of the Identity Gang, a group comprising over 2 000 professionals in this field. The definition of “digital identity” as used by the group appears in the Identity Gang’s Lexicon at http://identitygang.org/moin.cgi/Digital_Identity.

⁵ Definition for “identity” at <http://wordnet.princeton.edu/perl/webwn>, as viewed on 3 December 2007.

⁶ Again, “identity” refers to the set of claims itself, to “Joe’s documents”, to “Joe’s ID card” (to the claims represented in them), and not to Joe himself, not to the identity between Joe in T1 and Joe in T2. For this reason, the paper later explains that identity is referential, because the “document”, the “thing”, must refer to a person.

⁷ The distinction between identifier and the person identified breaks down with biometrics, which at once refer to subsets of identity information and (part of) the actual person.

⁸ The American Heritage Dictionary of the English Language, Fourth Edition. Houghton Mifflin Company, 2004. <http://dictionary.reference.com/browse/persona> (accessed: 7 January 2008).

identities may be familiar to others due to **profiles** that others develop about him (with a profile being “a set of data exhibiting the significant features of something and often obtained by multiple tests”⁹).

A data “**subject**” is the person to whom a digital identity refers.

As the terms “persona” and “profile” suggest, identity information can be used by different people to describe a person. Applying the ideas of philosopher Paul Ricoeur: When the data subject himself is initiating new actions through a digital identity, that identity may be referred to as “**ipse identity**”; when others act based on what they know about a person over time, the identity may be referred to as “**idem identity**”.¹⁰

As understood in the European Union, when information in a digital identity relates to an identified or identifiable natural person – meaning “one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity” – they constitute “**personal data**”.¹¹

Of course, IDM and other technologies can allow a person to remain anonymous as data pertaining to him is exchanged, in which case the term “personal data” would not be appropriate. “**Identity information**” serves as a more generic term for data relating to a person, whether identified/identifiable or not.

⁹ Merriam-Webster's Medical Dictionary. Merriam-Webster, Inc.
<http://dictionary.reference.com/browse/profile> (accessed: January 07, 2008).

¹⁰ These conceptions stem from work by Paul Ricoeur (1913-2005). As detailed on the Stanford Encyclopedia of Philosophy: “Following a distinction in Latin between *idem* and *ipse*, Ricoeur holds that the self's *idem*-identity is that which gives the self, among other things, its spatio-temporal sameness. Its *ipse*-identity gives it its unique ability to initiate something new and imputable to himself or herself.” <http://plato.stanford.edu/entries/ricoeur/> (accessed: January 07, 2008).

¹¹ This definition of personal data is contained in Directive 95/46/EC.

FROM “PERSONHOOD” TO DIGITAL IDENTITY¹²

This section explores the relationship between personhood and digital identity. As it considers some philosophical influences that shape today’s conceptions, it underscores the importance of user control in the electronic realm for promoting a culture of accountability and trust.

Classical influences on personhood in identity management

Markings of modern philosophy may be seen in the IDM framings of personhood and digital identity. Both Georg Wilhelm Friedrich Hegel and John Locke set out ideas that have affected society’s willingness to use identity information to refer to a person. Hegel, in his *Phenomenology of Mind*, asserted that “it is only by being acknowledged, or ‘recognised’,” that a person is known to exist.¹³ With IDM, recognition comes from information pointing to a person.

Locke emphasised how personhood entails a consciousness of being the same identity over time; with this conception, personhood stems from an intelligent, thinking being’s ability to know oneself to be the same thinking being in different contexts. Consciousness acts through a material body, and accountability for choices made attaches to the consciousness. Locke’s ideas implicitly lie behind authentication in IDM systems: IDM systems recognise people through external traits which remain stable over time, but they also authenticate people. The act of authentication represents and depends upon the person’s memory and consciousness of being the same identity over time. The act of claiming an identity in an IDM system and passing the authentication challenge represents an assertion by a person of a Lockean personal (as opposed to bodily) identity, and this authentication in turn creates a voluntary and conscious basis for accountability.

Classical influences on digital identity

In terms of how personhood relates to digital identity – especially personal data – Hegel and Locke have influenced today’s dominant legal theories in two different ways. In Europe, the law reflects a Hegelian sense that the person has a property interest in being able to control personal data. Hegel saw property as allowing an individual to have autonomy over resources,¹⁴ and so property was a feature of personhood: “Not until he has property does the person exist as reason...”¹⁵ In addition, European law reflects a sense that a person experiences freedom as he enjoys property in the context of the community, which the state affords. Freedom, to Hegel, is fully realised only in community with others: “the person must give its freedom an external sphere in order to exist as Idea.”¹⁶ Freedom as experienced through life in community has an overall connectedness in the *Geist* – that is, the mind, or spirit, of the state. In other words, personhood demands control over property for expression, and freedom accompanies that control

¹² This section was written by Mary Rundle, Bob Blakley, and Marcelo Thompson Mello Guimarães.

¹³ Georg Wilhelm Friedrich Hegel, *Phenomenology of Mind*, 1807, §178.

¹⁴ Margaret Jane Radin, Property and “personhood”, 34 STAN. L. REV. 957 (1982); reprinted in part in Property Law 53 (E. Mensch & A. Freeman eds. 1992).

¹⁵ Georg Wilhelm Friedrich Hegel, *Elements of the Philosophy of Right*, 1822, § 41.

¹⁶ Id.

over property when it is enjoyed within community. Hence, European data protection reflects a sense that a person should be able to control data relating to him, and that the state helps him enjoy those rights.

In the United States, another dominant regime affecting the treatment of personal data today,¹⁷ the law reflects a Lockean sense that people reign over their separate private spheres and have defensive liberties against the state. In particular, the Lockean theory of property, normally called the “labour theory,” has influenced the US conception of the relationship between personhood and private property. According to Locke, “every Man has a Property in his own Person”, from which it follows that “[t]he Labour of his Body, and the Work of his hands... are properly his.”¹⁸ This property is in the private sphere and is therefore under the domain of the person, as opposed to the state. Hence, in US law the concern is with preventing state interference with the private sphere; personal data needs to be protected from interference by the state.

The important distinction between personhood and digital identity

Hegel and Locke, and the legal systems influenced by them, place a great deal of importance on recognition of the person. However, depending how identity information is controlled, IDM could threaten to undermine personhood. Stated simply, when individuals are not in effective control of identity information, personhood and the enjoyment of human rights shrink.

To understand why, it is important to remember that human rights attach to the person, as opposed to the profiles that may be built up from identity information about him. While it has always been the case in human history that a person’s reputation and actions have influenced others’ treatment of him, this tendency is magnified in IDM as a person is effectively recognised by the spin-off profiles that begin to accrue for him. In other words, the danger is that what is relevant is no longer *personhood* – the recognition of a person as having status as a person – but rather a *profile* – the recognition of a pattern of past behaviour. Those past actions themselves are not the source from which his human rights derive; rather, the state of being a person gives rise to those rights.

If IDM profiles substitute for the actual person – to the point where recognition is transferred to the profiles rather than to the person behind them – the concept of personhood dissolves. Taking the example of a faulty credit report, if a person has no way of knowing about an error that negatively affects his credit rating, he may be denied a loan to purchase a house. His enjoyment of rights suffers (in this case, the right to know what information is held about him and the ability to correct it so as to buy a home). The problem is caused by the fact that identity information is detached from the person’s control. Such detachment can lead to a diminishing of a person’s participation in society and basic enjoyment of personhood. The ability to control the use of one’s identity information is crucial for reminding others that there is a person behind data and enabling that person to have full status when dealing with others.

To Hegel a person’s freedom increases the more he experiences connection in society; translated to IDM, a person’s freedom increases the more he interacts with others through digital identities that remain under his control. In Lockean terms, actions may be thought of as an extension of the person; acting through digital identities expands a person’s domain. Both philosophical traditions would see greater fullness for the person as he acts through digital identities. Vital to both would be user control over digital identities, which activates personhood and promotes both social connectedness and autonomy.

¹⁷ US conceptions of data protection are echoed in the Asia Pacific Economic Cooperation (APEC) Privacy Framework (2004).

¹⁸ John Locke, *Second Treatise on Government*, 1690, §27.

In addition to effects on the individual, society as a whole suffers when the tie between a person and his data is severed. Data becomes plagued by inaccuracies, people begin to fear that decisions are arbitrary, and a sense of justice wanes. In this climate, people have fewer means to assess the riskiness of relationships and so opt not to interact. User control over identity information can reverse this trend and generate accountability and trust.

IDENTITY LESSONS IN LITERATURE

Novelists and dramatists have driven Locke's points home over and over again: *Dr. Jekyll and Mr. Hyde* presented several persons in one body; *Twelfth Night* gave us multiple persons in multiple indistinguishable bodies; and the theme of one person in one body appearing to be different bodies goes all the way back to Odysseus in the cave of the Cyclops.

These are not just pretty stories; at least two very practical problems arise out of these confusions of identity:

The first is that since a body can change without the (Lockean) person changing, people can have difficulty recognising each other. This is why disguise is effective; it's also why ID card photos are re-taken periodically. The philosophical version of this problem, due to Plutarch, is "the ship of Theseus" - if Theseus sails the Argo around the world, stopping to replace parts of the ship at each port as things wear out, and if eventually every single part has been replaced, is it still the same ship?

The practical problem arises, for example, in the use of biometrics. Faces, fingers, and other body parts change or are even lost over time, undermining the accuracy or even the feasibility of biometric identification. Biometric researchers know this, and they select physical traits as candidates for biometric modalities based on two criteria: stability and distinctiveness.

Stability is the tendency of the trait to change slowly, in a single individual, over time. Distinctiveness is the (low) probability that a particular configuration of a trait (for example a pattern of fingerprint ridges) will be shared by two different individuals. No trait is entirely stable, and therefore biometrics can in principle only establish some level of probability that an individual has been correctly identified - there can be no certainty.

The second problem is that a person can change without the body changing. This is the basis for legal exemptions from accountability; for example, a defense attorney may be able to shield his client from punishment if he can establish that although the defendant committed the crime in question and knew at the time that it was wrong, his mental state at the time of the trial is such that he cannot understand either the wrongness of his act or the reason for which he would be punished.

This problem also causes problems for identification of persons; people forget passwords - and it is this lack of the mental continuity Locke asserted was the basis for identity of a person that makes it impossible to identify them after this happens. When systems try to recover from the effects of users' forgetfulness, they typically do one of two things: ask for other remembered facts (but this may not work) or ask for physical objects like ID cards or birth certificates - but events like the Asian tsunami and Hurricane Katrina show vividly that these things too cannot be depended upon as an infallible mechanism of identification.

DATA PROTECTION IN THE IDM-ENABLED UBIQUITOUS INFORMATION ENVIRONMENT¹⁹

Taking the example of a hypothetical person named Yong Ai Tun, this section projects an image of a potential ubiquitous information environment. It shows several threats that arise if data protection for identity information is not sufficiently in place.

A tale of a ubiquitous information environment

As Yong Ai Tun heads to the subway one morning with her music on, she chooses to be simultaneously connected to dozens of information and communication technology (ICT) systems. Some are Internet based, some are mobile radio based, and some are running on new communications channels, with a mix of public and private networks. In the span of time between when she walks from her home to the subway, she checks inventory at her company and orders more production materials, agrees to purchase and send the flowers and books that an electronic agent recommends for her aunt's birthday, dispatches an avatar to go participate in a spectrum protest, and plays a round of poker.

In seamlessly hooking into these multiple systems at once, Ai Tun controls what information she releases to each channel. She uses different identifiers: Yong Ai Tun in business contexts, AiTun89 in social contexts, and Rainmaker in some other world system. The authentication methods used by these systems have varying levels of security.

Ai Tun engages in these activities through virtual worlds: Social computing has moved towards real-time interactions in 3-dimensional (3-D) spaces, integrating different media.

Accelerated by her generation's demand for easier, faster and overall better person-to-person communication, several types of access devices coexist or have converged. These include handheld always-on devices, embedded IT, and ambient intelligence, which serve as game consoles, TVs, mobile phones, music players, video recorders, office machines, remote controls for appliances, and more. The widespread use of these always-on devices has led to a convergence of real-world and virtual-world identities.

As communication and information volumes have grown, the issue of attention management has become more urgent for Ai Tun. To help her manage her presence and participation levels, the IDM platform makes it easy for her to employ multiple agents (and subagents) to operate on behalf of the various roles she plays. These intelligent software agents enable her to stay appropriately connected and competitive, assisting her in managing relationships, getting the best deals on the coolest products, staving off continuously adaptive spam and advertisements, and discovering the new information necessary to be effective as she acts in various capacities. In addition to intelligent software agents' helping to manage connectivity and participation, these agents help guard privacy; for example, by auditing the treatment of data, they help prevent others from colluding to consolidate information about Ai Tun's various roles.

As Ai Tun emerges from the subway, a sensor in the turnstile gauges her temperature and detects that she is carrying a fever. The sensor transmits this data to the transit card reader, which determines her name based on her radio-frequency identification (RFID)-enabled card that serves for transportation, banking, and other needs. The device then sends data about her fever over a network, which alerts Ai Tun's mobile

¹⁹ This section was written by Mary Rundle, Paul Trevithick, and Mary Ruddy, with contributions from Anthony Nadalin.

phone service provider that it should track Ai Tun. Cell towers triangulate her location as she walks down the street; meanwhile, the phone company gathers her recent location profile and requests another company in its federation to provide the names of all individuals with whom she has been in contact in the last 48 hours. The company sends the data on to a government agency in the healthcare sector. The agency determines that Ai Tun has not been exposed to any contagious diseases, updates Ai Tun's dossier, and sends out a "clear" message to the neighbourhood disease-control intervention points that have been alerted. This health check has been conducted in under 20 seconds, enough time to grant Ai Tun clearance to enter a nearby building, even though she never knew she had to obtain this clearance. The healthcare agency is part of the Program for Efficiency in Government (PEG). As such, it makes the information it has collected available within the PEG federation.

Another PEG participant organisation, the Wealth Resource Observation Network for Governments (WRONG), receives the updated location histories of Ai Tun and the people with whom she has been in recent contact. Matching that data with other records, it determines that Ai Tun was in close proximity to someone who had withdrawn a significant sum of money from a virtual world. If Ai Tun and this person who had withdrawn funds were each carrying devices with the latest version of Bluetooth, it is possible the funds were transferred to her as they passed each other. Ai Tun's name, as transliterated into roman characters, is the same name as one appearing on a suspected terrorist list; therefore, the WRONG analysts send an alert to Ai Tun's banks to notify them to monitor her activity. One of the banks applies an extra degree of care to its accounts due to their large amounts. Upon receiving the alert, the bank freezes Ai Tun's account.

Unfortunately, Ai Tun is currently in the process of trying to purchase medical treatment for her uncle, Yong Kurzweil, who has experienced heart failure. With the freeze on her account, the transaction is blocked, and Ai Tun's uncle is denied treatment and taken off life support. As he dies, a change-of-circumstances bulletin propagates through the communication Grid, and all his accounts are deactivated. Ai Tun, however, does not receive word because she is temporarily disconnected while her status on the watchlist is being investigated. In fact, Ai Tun has become completely unrecognisable to others because her Personal Electronic Recognition System for Online Negotiations (PERSON) has been disabled.

Noticing that her music has stopped, Ai Tun tries to contact her local ombudsman to find out what is happening; but with her Grid access denied, she cannot. Fortunately, her Rights Enforcement Safeguard Club for the Ubicomp Environment (RESCUE) sees that her presence indicator is not working and notifies the ombudsman, who obtains the necessary governmental clearances and checks what triggered the denial. Realising it was a mistake, he authorises a re-authentication prompt for her devices. Thankfully, the glitches are sorted out in time for Yong Kurzweil to be resuscitated; a sample of his DNA is then taken to re-enroll him with the global registry service that reactivates his accounts. Ai Tun is aware only that her music stopped for a moment.

IDM as a critical enabler of the ubiquitous information environment

Many types of ICTs combine to make up the ubiquitous information environment in this story. Mesh networking allows users to form spontaneous communications networks. Grid computing enables devices of low capacity to tap into pools of computing power and databases elsewhere. Feeding information into these communications channels are not just websites and 3-D applications, but also such applications as networked RFID tags, sensors, and location-based services (LBS), which convey information about the user and surroundings. Finally, the Semantic Web processes data about data (metadata) so that relevant information can be digested and presented in a useful form to the user through such applications as search engines and reputation rating systems. Together these technologies connect people and put practical information at their fingertips through cheaper, more efficient, and more tailored services. An identity

infrastructure is a critical enabler of this ubiquitous information environment because it is sufficiently open and can integrate these emerging ICT technologies.

Although hypothetical, the story of Ai Tun conjures up some of the risks heralded in the report on “Dilemmas of Privacy and Surveillance: Challenges of Technological Change”:

Technologies for the collection, storage, transmission and processing of data are developing rapidly. These technological developments promise many benefits: improved means of storing and analysing medical records and health data could lead to improvements in medical care and in management of public health; electronic logging of journey details can promise improved provision of public transport and more logical pricing for road use; and more details of peoples’ everyday behaviour offer the possibility for developing better public policy generally.

However, the development of these technologies also has the potential to impact significantly on privacy. How they develop is to a large extent under the control of society. They can be allowed to develop in a way that means personal data are open to the view of others – either centralised spies or local peeping toms. Or, they can be allowed to develop so that personal data are collected and stored in an organised, controlled and secure manner. There is a choice between a “Big Brother” world where individual privacy is almost extinct and a world where the data are kept by individual organisations or services, and kept secret and secure. The development of technology should be monitored and managed so that its potential effects are understood and controlled. The possibility of failures of technologies needs to be explored thoroughly, so that failures can be prepared for and, where possible, prevented.²⁰

Because IDM promises to be a unifying thread for so many emergent technologies, it deserves special attention. In particular, a decisive factor in whether the ubiquitous information environment makes for a healthy or repressive information society is whether the user enjoys control of identity information relating to him. Data protection is vital for this control.

Ai Tun’s story points to some challenges that the OECD’s *Privacy Guidelines* aim to address. Specifically:

- In terms of limits on the collection of personal data, was there an international legal framework by which to judge if the data was obtained by lawful and fair means? Was that data captured with the knowledge or consent of Ai Tun?
- Regarding data quality, was the data that was collected relevant to the purposes for which it was to be used, and was it accurate, complete and kept up-to-date?
- As for purpose specification, were the purposes of data collection specified in advance, and was subsequent use limited to the fulfilment of those purposes?
- With respect to use limitation, was it the case that personal data was not disclosed, made available or otherwise used for purposes other than those specified except *a)* with the consent of the data subject, or *b)* by the authority of law?
- Did each entity handling Ai Tun’s data take reasonable security safeguards?
- Were the practices and policies applying to her data open and clear?

²⁰ The Royal Academy of Engineering, 2007, p. 7.

- Did Ai Tun enjoy an effective right to have access to that data relating to her? Was there a way for her to have timely redress if denied access? Was she able to challenge her profiles and, if the challenges were successful, to have the data erased or corrected?
- Did each of the entities treating Ai Tun's data have a data controller who was accountable for ensuring compliance with these principles?

Tomorrow's enhanced services

It is likely that IDM will become pervasive as the technology helps to solve simple transactional problems that have plagued ICT. Today most people repeatedly fill out billing and address information when making purchases remotely; they often do so with no guarantee that the party on the other end of a transaction is who he purports to be, and they have no binding way to indicate preferences for how they wish their data to be treated, assuming the transaction is legitimate. In the marketplace, IDM will likely grow popular by allowing parties to establish trust and to transfer personal information with ease, according to user preferences. In policymaking circles, IDM will likely be endorsed as a cure for the failing confidence the public has in e-commerce, and as a way for governments to provide services more efficiently.

Entirely new capabilities will rapidly emerge as the user-centric²¹ IDM layer transforms into what is essentially a new paradigm and applications platform. It is impossible to predict what most of these might be, but it is fairly easy to compile a list of potential candidates. An illustrative list is offered here:

One capability will likely entail the **metaverse**. Originally thought to be separate and distinct from "real" life, people now see that new possibilities arise when one attempts to break down the barriers between these worlds and mash them together into what some call an "augmented" reality. To achieve this, the same IDM layer that will be pervasively integrated into traditional systems, websites, and devices people use every day, must also be embedded in almost exactly the same way into the corresponding objects and interfaces in simulated, virtual worlds.

Another capability will involve **central control with distributed data**. A common, consistent IDM layer will afford users a single, *centralised* dashboard (or control panel) for their *distributed* identity information. They will be able to link together their information across distributed external systems and then update and synchronise those aspects of their information over which they are authoritative, at the push of a button on the centralised service. A person will be able to change his address and propagate this change to potentially hundreds of external systems at once.

The concept of **multiple digital personas** (the common parts of each which are linked together for quick update) enables the introduction of more nuanced relationships into online social networks. It is clear that the requirement for allowing a person to express himself differently in different contexts is fundamental to society in many ways. And it will be a requirement of all future IDM solutions as well. Creation and maintenance of multiple simultaneous, digital personas will become a new social norm.

User-centric IDM makes deep **personalisation** of service and dynamic **discovery** of new of services practical. For example, in the e-commerce realm only a few systems currently do a good job of presenting compelling product recommendations. This is partly because the e-commerce "silo" has access by direct observation to only that silo's specific slice of the user's clickstream, choice of navigational path through the site, behaviour, interests, preferences, and so on. Further, for a variety of reasons, silos often have trouble even knowing that this is the same person who showed up before. However, new IDM intelligent

²¹ The section on "Current Conceptions" elaborates on the concept of user centric IDM architectures.

software agents controlled entirely by the user are able to project (in a privacy-enhancing pseudonymous manner wherever practical) rich preference and interest data to external systems that can now “subscribe” to these agent-managed customer data feeds. So, too, the IDM layer makes possible the emergence of dynamic, contextualised (*e.g.* location-based) discovery of new services.

Continuous improvement and adaptation. User-centric IDM creates a virtuous cycle: enriched profile and preference data enables personalisation and discovery of new services, and by implicit and explicit feedback based on the user’s reaction to these personalised/new services, the profile is again enriched. New categories and sub-categories of preference emerge.

At a deeper level, policy makers must understand that the design of IDM systems can determine who has access to what information. By choosing options that promote data protection and foster user control in IDM, policy makers can support the wider goals of democracy and an open market economy.

DATA PROTECTION AND USER CONTROL²²

Here the paper describes how fair information practices can boost user control in IDM systems. Demand for this data protection is likely to grow with, and in turn help shape, the evolution of IDM.

Discussion of user demand for data protection here is based on the following assumptions about the future evolution of IDM:

1. Today's "islands" of identity technology protocols and systems will be bridged by an overarching IDM layer in the identity infrastructure.
2. Every device, platform, and system used by people will plug into this IDM layer.
3. This IDM layer will be based on the understanding that one person wields multiple digital personas and that these personas are contextual.
4. Users will learn and become familiar with a set of new, standardised, everyday user experiences around such "ceremonies" as authentication, release of identity information, selection of digital personas for different situations, and review of user-friendly privacy policies; and
5. Users will enjoy new capabilities, especially those based on delegation of user authority to identity agents that work on their behalf.

Against this backdrop, users are likely to demand IDM tools that allow *a) notice* of other parties' treatment of identity information, *b) an opportunity* for the user to *consent* to or refuse this treatment, *c) an assurance of security* (including privacy), and *d) access* to information on actual practices affecting their data, with an opportunity for redress. The upshot is a market push for fair information practices in identity management.

Notice

IDM systems must allow the user easily to see and understand the way that other parties will treat his data.

Users need to know these data policies in order to make informed choices about whether to release identity information to "relying parties". They must be able to see what information is being demanded from them as well as what the relying party's data retention and handling policies are. Having to read a new, multi-page legal document every time one deals with a new party is tedious. Then, two months later when the policy changes, it becomes necessary to read a new, lengthy explanation. With these obstacles, people tend to ignore the document altogether, trusting that the company is not doing anything inappropriate.

²²

This section was written by Paul Trevithick, Mary Rundle, and Anthony Nadalin.

Similarly, the user needs to be able to see the data treatment policies of entities they trust to securely store and manage personal data – that is, they need to see the policies of “identity providers”. None of these requirements are met by existing, legacy IDM systems.

Adding to pressures are the unique problems spawned by international transactions as laws of different countries allow parties different levels of default access to personal data. For example, purchasing a widget from a vendor in the United States may expose a person to different treatment than using a vendor from France would. The vendors may have the same policies (they may even be the same company), but they are subject to different exposures from the sovereigns. Greater insight into this dynamic may cause certain countries to become business havens because of the way they treat data, just as Switzerland and the Cayman Islands have for bank accounts or the US state of Delaware has for corporate headquarters. Users need to know what the default treatment of their data will be.

It would be helpful to create a standard, checklist of data protection policies so that persons could rapidly assess the way their data will be treated by each entity.

The need for simplicity here suggests some standardisation of terms so that people might rapidly ascertain the notice conditions described above. Standards for simplicity could address the complexity of current legalese and allow people to make market comparisons – thus fostering competition among vendors to offer better levels of data protection.

Consent

In addition to IDM systems’ allowing users to receive notice of other parties’ data practices, users need to be able to express consent for that treatment or decline interactions given those practices. Better still, users ideally should be able to define the conditions and obligations that others must adhere to when dealing with their identity information. As with the area of notice, approaches may emerge to allow people to negotiate contracts for data treatment.

In the new identity infrastructure, it has been suggested that users could choose among “icons” to express preferences, with these icons representing legal agreements that are bound to “machine readable” and “human readable” policies.²³ The vision is for terms for the treatment of identity information to be enforceable. This enforceability would allow the IDM icons approach to go beyond previous efforts of this kind (e.g. the Platform for Privacy Preferences, or P3P).²⁴

Alternatively, there could be a place where people would post their default preferences for the treatment of their personal data; parties dealing with that data could have an affirmative legal duty to consult that posting and abide by those conditions.

²³ Such an approach would emulate that of Creative Commons, which has developed a system for creators of digital content to waive aspects of their intellectual property rights. Creative Commons has become popular around the world by using universally recognisable icons to represent copyright policy, with these icons having computer code, lay-term explanations, and legal statements attached to them to be “machine readable”, “human readable”, and “lawyer readable”. For information on Creative Commons, see: <http://creativecommons.org/about/licenses/how2>; for a sketch of this idea to use icons for data protection, see <http://www.w3.org/2006/07/privacy-ws/papers/21-rundle-data-protection-and-idm-tools/>.

²⁴ If developed, such a system might eventually allow users with similar preferences for data protection to form groups that would then have greater negotiating leverage when dealing with vendors. Of course, the expected increase in delegation of authority from users to their identity agents will add new complexities to handling consent issues when the user is not, at least initially, in the loop.

The hope is that in combination, legal provisions and technologies will allow an individual to have an effective consent right in matters relating to his personal data.

Security (including privacy)

Practitioners developing new and especially user-centric IDM technology are faced with a difficult challenge from a security and system-hardening perspective. Whereas the intent and promise of the new technology is to protect users from malicious attacks (*e.g.* phishing and pharming), the introduction of IDM layer components tends to centralise (or at least correlate) identity-related data flows through a small number of standardised infrastructure components. IDM does not necessarily result in a logical or physical concentration or aggregation of personal data, but it does result in the fact that identity data *flows* are concentrated through new, standardised IDM components installed on the user's computers and devices. It is almost guaranteed that adversaries will now turn their attention to these new common IDM components with intense energy, due to the far greater potential rewards of a successful attack. Today's relatively insecure Internet has many diffuse points of vulnerability. With the introduction of IDM components the situation is inverted: Now a few common components are relied on for many interactions. This is a challenge that experts understand to be serious. Some potential threats to a user's security are higher than without the pervasive IDM layer.

Some key elements of the user's identity could be kept out of the system and embedded on a smart ID card. The card would hold just enough data to unambiguously authenticate the user as well as public and private keys to encrypt and decrypt data stored in the network. This would remove a key link between the private data stored in the IDM and the owner of that data whose identity is stored on the card.

One of the particular challenges from the user's perspective is that almost all measures and new technologies that can be used to reduce vulnerabilities, strengthen authentication, and provide further security measures do so at a cost to the user's convenience. Experience has vividly shown how insensitive users are to chronic, low-level security threats and conversely how much they value ease of use and convenience. Balancing the competing requirements for security vs. convenience will be one of the most difficult challenges for the user.

User-centric IDM has a clear benefit to security due to its introduction of common, consistent user-experience ceremonies for basic identity interactions (*e.g.* authentication) across platforms and devices. The reason is simple: In a world without standard interactions, users tend simply to "click through" dialog boxes without asking questions. They have neither the time nor the patience to examine each and to try to make sense of what is being asked of them. This makes it far easier than it should be for adversaries to fool a user into providing information when he should not do so. Many studies have documented the ineffective nature of visual cues and indicators.

There are also security challenges in the area of privacy. Whether or not user-centric IDM technology yields an overall increase or decrease in end-user privacy depends partly on the design of the identity infrastructure, partly on its implementation, partly on how people actually use the infrastructure, and (lastly yet probably most importantly) partly on legal regimes and social norms outside the technical realm.

The technical qualities that users are implicitly demanding for the privacy aspects of user control include the following:

- i)* **Decentralisation.** Maximal decentralisation of identity information into as many separate data contexts as possible.
- ii)* **Data minimisation.** The common sense notion that only the minimal amount of identity information necessary to support all of the required transactions should be stored.

- iii) **Local identifiers.** In order to prevent privacy-destroying linkage and aggregation of identity information across data contexts (e.g. repositories), each context should wherever possible define its own localised naming scheme (preferably using local pseudonyms) to identify the set of identity information associated with a person and thereby avoid using more global identifiers such as a government tax identity number.
- iv) **Verifiability.** Because relying parties sometimes require that claims made about the user be verifiable, the system must support mechanisms for verification of claims.
- v) **Selective disclosure.** Beyond minimisation of what's stored, in any given exchange of identity information only the information necessary to enable the specific transaction anticipated should be involved. For example if a person's age is required to be stored, it is often the case that merely revealing that the person is over or under certain age thresholds is all that is required in a specific transaction. In some cases the user must also be able to combine selected claims made about them by more than one identity authority into a minimal composite set of claims and be able to present this to a relying party in such a way that the relying party cannot repudiate the original claims. New cryptographic approaches (a new kind of Privacy Enhancing Technology²⁵) are required to meet these (and other closely related) privacy requirements.
- vi) **Composability.** It is best if the user is able to assemble reusable groups of related partial identities into convenient digital "persona" composites that can be used in recurring social, commercial or governmental settings. Without this, the tendency will be for users to rely on a smaller number of less-minimal and more easily correlated digital identities that reduce privacy.
- vii) **Auditability.** The identity infrastructure should be designed to allow audits. Audits can provide accountability, helping to assure parties that the risk of corruption is low and enabling records for legal redress.

Together, these security concerns are sizeable. Still, addressing them is a prerequisite for an identity infrastructure that affords user control to enable accountability and trust.

Access

In terms of access, people want to be able to see what data other parties have on file concerning them, and they want an opportunity to contest those records. International data protection principles seek to provide a means to check that the treatment of personal data is in line with expectations.

Although these principles have shown concern for access, a practical difficulty with the concept is that the burden is on the individual to find out who has his personal data in the first place. Given the number of entities that deal with personal data today, the task of chasing up who has what, and how they are treating it, is so arduous as to render a person's right of access nearly useless. However, IDM tools that are now being conceived could reverse this situation, making it practicable to know who has a person's identity information and how they are handling that data.

Still, individuals themselves cannot always be privy to what is being done with their data by public and private actors spanning different jurisdictions. The law might generally require consent by a person if his personal data is used by an entity, but it might at the same time authorise use without consent for certain purposes, such as national security. In such cases an individual will be in the dark as to whether his identity information is being used for purposes to which he has not consented. There is arguably an

²⁵

[http://www.oilis.oecd.org/oilis/2001doc.nsf/LinkTo/dsti-iccp-reg\(2001\)1-final](http://www.oilis.oecd.org/oilis/2001doc.nsf/LinkTo/dsti-iccp-reg(2001)1-final).

ombudsman role for democratically accountable officials to play in verifying that citizen data, if shared without consent, receives proper treatment and is safeguarded from subsequent misuse by downstream actors.

Another area with potentially large implications for access is that of national ID cards. A government might issue official identity documentation, or root identities, in order to enable people to exercise digital personhood and enjoy autonomy in the information society. However, the absence of user control over a root identity could lead to a situation where a person would have to “show papers” in order to participate in society. The ultimate risk would be what author George Orwell referred to as “unpersonhood”, namely: the destruction of personhood through denial of access to one’s dossier – or, in this case, the lack of user control over the identity information that is deemed necessary for participation in society.

Implementing fair information practices

The emergence of a ubiquitous and relatively consistent (especially from a user experience point of view) IDM layer provides for the first time the *technical* means to meet these demands for fair information practices. However, it is not yet clear how the IDM development community will achieve the required levels of shared understanding and collaboration to implement the technical solutions. If a legal responsibility existed, and the technical means were within reach, could the providers of IDM systems be understood to have an obligation to offer tools for increased data protection? Might market forces lead them to build in interoperable tools for data protection as a result of demand for user control?

MARKET DEMAND FOR USER CONTROL²⁶

This section looks at the business perspective and predicts a growth in market demand for user control.

The perspective presented here is from the consumer-to-business and business-to-business vantage points. This section starts by spelling out its assumptions about the evolution of today's IDM business challenges landscape. It then examines this projected state in light of business-process transformation. The perspective suggests that the market is likely to favour those IDM solutions that give control of identity information to the user.

As businesses grow increasingly dependent on the Internet, they face challenges such as the rising threats of fraud and identity theft; increasing regulatory compliance requirements; a surge in consumer demand for privacy protections; and the competitive necessity to have dynamic partnerships with other businesses to interconnect their online services. These new market forces will start to fundamentally shift the direction of the IDM market.

Almost all online activities – including sending e-mails, filing tax declarations, managing bank accounts, buying goods, playing games, connecting to a company intranet, and meeting people in a virtual world – require identity information to be given from one party to another. The abundance of different situations and types of identity information suggests the need for a flexible and user-centric IDM infrastructure. This infrastructure must be flexible to support the multitude of identity mechanisms and protocols that exist and are emerging, and the different types of platforms, applications and service-oriented architecture patterns in use.

User-centricity is an emerging concept for IDM. There are some subtle and often overlooked points about user-centric IDM that are mentioned here. For starters, IDM must be user-centric since the end users are at the core of IDM: the infrastructure must empower the end users to execute effective controls over their identity information. These requirements have far reaching consequences, not only for user interactions with the IDM systems, but also for the infrastructure itself and how it must be built.

Some aspects of this trend have been materialising, morphing, and maturing for awhile already. For example, IDM is moving into an execution phase with wide deployment of smart cards, enrollment services, and IDM systems as customers rush to meet mandates and regulations, such as U.S. Homeland Security Presidential Directive 12 (HSPD 12) and the US Real ID Act. At the same time, the management models for IDM are evolving as all entities on a network, whether physical such as devices, or virtual such as policies, need to be identified and represented coherently.

As such, the industry is giving rise to competing frameworks from various vendors as they seek to capture this evolving abstraction of identity. At the same time, identity sensing and resolution capabilities are taking hold as governments and transportation entities struggle with the identity profiling issues associated with terrorism, as supply chains begin to address pedigree and traceability identity and location issues through the use of RFID and GPS, and as financial institutions deal with identity theft through better

²⁶ This section was written by Anthony Nadalin.

security measures. These two evolving forces (IDM and identity sensing/resolution) are creating ripple effects, in turn demanding technology related to information management, information integration, and privacy.

Protecting sensitive personal information is critical, and privacy regulations are on the rise. Although from a technology viewpoint, the priorities may be authorisation and control, what seems to be different and evolving is the notion of equipping the end user with the necessary controls to protect his identity information: Users are informed about what data is requested from them and how their personal data is treated, *e.g.* for what purpose it is used and who can access it. Through this process, users can decide whether to provide their data and to consent to the service provider's data handling policies. Ideally, the service provider employs technical components such as access control systems to enforce the consented policies (*e.g.* ensuring that a user e-mail address is not used for marketing but only for the consented billing purpose).

At the same time, a new wave of individual and enterprise productivity will be sparked due to the integration of people with business processes. Information about people in the enterprise is abundant and growing, both in richness and in volume; while it is currently scattered in many disparate databases, this information will become more integrated. Also, given the trends towards social networking, collaborative computing, and people being a core part of processes, users will demand to be empowered to better manage their identity information, and control access to the same.

Major aspects of the value proposed by various portal products include seamless integration with heterogeneous back-end systems, management of user-specific personalisation data, and improving people's collaboration by simplifying communication processes (*e.g.* through instant messaging, presence awareness, communities, people tagging, team spaces, calendar sharing, team calendars, to name a few). These aspects impart a set of requirements of the identity infrastructure that generates an increasing demand for corresponding IDM features.

In addition, there is tremendous interest in the new Web 2.0 technologies, which not only offer the promise of a richer and more responsive web-user experience, but also are specifically addressing the value proposition of connecting people and amplifying the power of working together. The specific challenges for the identity infrastructure in that area are making people the primary concern in the overall system and accommodating the highly dynamic and self-organising nature of such environments. For example, a typical requirement that portals are currently facing is "community isolation": restricting user visibility based on the set of collaborative communities in which individual users are participating. More concretely, a given user named Bob should be aware of the existence of another user (say Alice) if and only if Bob is member of at least one of Alice's communities. If this is actually the case, Bob should have view access to a dedicated sub-set of the user profile information associated to Alice (*e.g.* her e-mail address, as relevant to their shared community) while other user profile information shall remain hidden from Bob. In this example, if Bob invites Carol to join this community, Alice should immediately become aware of Carol and see (parts of) her profile information. This requires the access-control layer of user-profile information to be highly dynamic because such changes are assumed to occur *ad hoc* and at a relatively high rate.

User-centricity distinguishes itself from other notions of IDM by emphasising that the user (or some agent of the user) – and not some authority – maintains control over "what, where, when, and to whom" a user's identity information is released. Part of this notion enforces user consent which requires that (a) the user's view of any transaction (including subsequent treatment of data) corresponds to the actual transaction and (b) the user agrees to the execution of the transaction. With user consent, the user has the ability to opt in or out of the release of information. For example, before a user logs into a banking website, he is told that he must prove his name and birth date, and only when the user agrees to this transaction and proceeds is his data released. In user-centric IDM, the user may choose from many identity providers and

also move his information between them. Two important components are mechanisms to protect a user's privacy and anonymity, and yet simultaneously hold a user accountable if he misbehaves.

Any enterprise wishing to adopt user-centric IDM will require some corresponding changes in its business processes, but perhaps more challenging and more importantly, it will need to change its business thinking and culture.

Probably the most fundamental paradigm shift for organisations is to move from believing that they own the personal information of their clients to believing that they are really stewards and custodians and that the individual is still the ultimate owner of their own information. This in itself does not necessarily lead to specific business process changes, but forces organisations to consider the needs and desires of the individual. This will often translate into fair information practices – providing the individual with explanations (notice) regarding use of his data, seeking his consent, giving him access to that data, and ensuring its security. It also leads naturally to thinking about the diverse needs and attitudes of the client base and therefore to the range of options that should be offered regarding information handling. Instead of companies dictating to clients what is needed, design emerges from a more collaborative negotiation process.

At present, most organisations view every client contact as an opportunity to begin building an ongoing relationship with the client. This relationship may lead to more opportunities to do business with the client or to build client satisfaction and loyalty. Consequently, the company seeks to gather information from an individual the first time he requests a service, with a view to building an ongoing relationship. This orientation may lead a company to gather information that is not strictly required for the transaction, and it may prevent the company from deleting information once the transaction is completed. A shift would not mean that organisations could not build client relationships; it would just mean that they would have to do so through explicit relationship-building transactions to which the individual would consent. Organisations must come to see that the personal information of their clients is not only an asset, but also a potential liability, *e.g.* a source of law suits over the failure adequately to protect such data, particularly in the absence of a client driven/consented reason for having it. As regulatory controls over personal information increase, the amount of liability associated with data collection will also force companies to re-evaluate their data gathering and retention requirements.

Despite the human tendency to want to know the identity of the individual being served, for many situations this may not be necessary and may not be desired by the individual. To process transactions with little or no identifying information will often mean reliance on a third party assertion or assurance on behalf of the individual. This will require an enterprise not only to be confident in the technical trust assurances (*e.g.* digital certificates) provided, but also to develop new business and operational relationships with those third parties. This may include regular assurances/audits of third parties and co-operation in trouble-shooting and investigations.

A corollary to this is that the individual client also has to become a trusted and competent player in the IDM scenario. No matter how well designed the data protection solution is from a technical perspective, part of it will be running in the client's environment, and there will be a need for some minimal competence on the part of the user. This implies technology vendors' developing robust support processes to assist users in setting up and maintaining their environments, and explaining the benefits of the new metaphors, and the risks associated with the proliferation of identity information.

User-centric IDM can also carry some unique requirements that in turn require specific processes to support them. A case in point is designs where transactions are routinely anonymous but where, under certain circumstances (investigations, etc.), identifiers need to be re-attached. Clearly a specific process with all of the appropriate approvals and checks and balances is required to support this scenario.

Given these points, it is clear that organisations need to transform their thinking and business processes to:

- i)* Build appropriate notice, consent, security, and access into business process design.
- ii)* Limit data collection for each transaction to what is strictly required for the transaction.
- iii)* Securely dispose of information that is no longer required once the transaction is completed.
- iv)* Limit the amount of personal identifying information strictly to that which is required for the transaction.
- v)* Develop contractual definitions of obligations with the third parties that will be used to provide trusted assurances or assertions.
- vi)* Develop better support processes for clients; and
- vii)* Develop specific processes to support any unique design features (*e.g.* re-identification).

Although these business transformation requirements will require some sustained effort to achieve, they will likely have collateral, transformative influence on organisations in ways that are very beneficial. Organisations that adopt user-centric IDM and the business transformational changes that go with it will be moving a big step closer to their clients. They will become organisations that always put a premium on the client perspective when designing new services and processes. As a result, it will be natural for them to be rewarded with increased growth and customer loyalty.

Even as the market drives business to adapt products, people responsible for designing sound legal and technological systems for IDM need to know that their designs will hold up under pressure. Fundamentally, they need to factor in the way identity behaves – that is, they need to factor in the Properties of Identity. These Properties of Identity are introduced in the text box below.

THE PROPERTIES OF IDENTITY*

Identity behaves according to a number of observable properties, as follows:

1. **Identity is social.** Humans are naturally social. To engage in social interactions (including commerce) people need something that persists and that can be used as a basis for recognition of others – an “identity”.
2. **Identity is subjective.** Different people have different experiences with the same individual and therefore attribute different characteristics to that individual; that is, they will construct different identities for him.
3. **Identity is valuable.** By building a history of a person’s past actions, exchange of identity information creates social capital and enables transactions that wouldn’t be possible without identity. In other words, identity lends predictability to afford a comfortable level of confidence for people making decisions.
4. **Identity is referential.** An identity is not a person; it is only a reference to a person. Even if a person develops spin-off personas so that other people know him through those various digital identities, and even if others create profiles of a person, ultimately the collection of characteristics that signal who a person is need to point back to that person.
5. **Identity is composite.** Some information about a person arises from the person himself; he volunteers it. But other information about him is developed by others without his involvement.
6. **Identity is consequential.** Because identity tells of a person’s past actions, the decision to exchange identity information carries consequences: Disclosure of identity information in a certain context can cause harm; failure to disclose identity information in another context can create risk.
7. **Identity is dynamic.** Identity information is always changing; any particular identity dossier might be inaccurate at any given moment.
8. **Identity is contextual.** People have different identities that they may wish to keep entirely separate. Information can be harmful in the wrong context, or it can simply be irrelevant. Keeping identities separate allows a person to have more autonomy.
9. **Identity is equivocal.** The process of identification is inherently error-prone.

* *The Properties of Identity were articulated by Bob Blakley, Jeff Broberg, Anthony Nadalin, Dale Olds, Mary Ruddy, Mary Rundle, and Paul Trevithick.*



Content in this text box is licensed under a Creative Commons Attribution 3.0 License.

THE PROPERTIES OF IDENTITY AND DATA PROTECTION²⁷

As policy makers consider data protection and IDM, a useful point of inquiry is whether the OECD's Privacy Guidelines accommodate the Properties of Identity. If they do, they should in themselves be adequate to guide data protection policy for IDM; if they do not, the gaps may signal areas for policy maker attention.

The paper here takes up the Properties of Identity in turn to highlight relevant data protection principles found in the OECD's *Privacy Guidelines*. For each property, the paper *i)* restates the property, *ii)* describes how that property is accommodated by the *Privacy Guidelines*, and *iii)* suggests how the *Privacy Guidelines* might be augmented to address the property more fully.

The *Privacy Guidelines* are used as a benchmark here because they serve as the statement of data protection principles that has been agreed among the widest representation of countries to date. For ease of reference, an excerpt of the *Privacy Guidelines* is reproduced in the Annex to this document.

1. Identity is social

i) Restatement of the property

Humans are naturally social, and to engage in social interactions requires that people be able to connect the past to the present, and the present to the future. People need, in other words, something that persists and that can be used as a basis for recognition of persons – an “identity”.

Article 6 of the Universal Declaration of Human Rights states: “Everyone has the right to recognition everywhere as a person before the law.” In essence, the right to recognition as a person is foundational to a person's enjoying all other rights. As Hannah Arendt explains, the right to have rights is a “pre-legal premise, a ‘proto-right’, in which it is left open, what a human may be, who a human may be, and which rights may be granted to him aside from this unique one of belonging to humanity and of formulating his rights correspondingly”. Fundamentally, this right to have rights is the right of every individual to belong to humanity.²⁸

As Cospedal García notes with respect to the digital environment, if every human is a person, it follows that “the accreditation of personal identity is a necessity of the individual in his public and private relations, which is translated into the exigency of having available a reliable means for its perception, without ambivalences, in the real or physical world and in the virtual.”²⁹

Indeed, the right to effective digital personhood will arguably be the most fundamental right in the future information society, as it will determine the possibility of a person to enjoy all other rights – including civil, political, economic, social and cultural rights.

²⁷ This section was written by Mary Rundle, with contributions from Bob Blakley, Marcelo Thompson Mello Guimarães, and Dale Olds.

²⁸ Hannah Arendt, *The Origins of Totalitarianism*, Meridian Books, 1967.

²⁹ García, Ma. Dolores de Cospedal, “Utilización de la Firma Electrónica en la Administración Española IV: Identidad y Firma Digital. El DNI Electrónico” in edit. Ministerio de Economía, Administración Electrónica y Procedimiento Administrativo (España: Ministerio de Economía, 2004) at 189.

ii) How the property is accommodated by the *Privacy Guidelines*

Because identity is so foundational to society and the individuals living in it, government arguably has a role to play in promoting a system whereby a person enjoys personhood – that he is recognised as having status as a person. In the information society, this role of government translates into ensuring that people can enjoy personhood and the rights and responsibilities that come with it, and that they can be recognised as having status as a person through their various digital identities that rest under their control. Designed to give individuals better assurance of how their data will be treated in cross-border dealings, the *Privacy Guidelines* implicitly recognise the right of a person to enjoy personhood in digital interactions.

iii) How the *Privacy Guidelines* might be augmented to address the property more fully

In light of the importance of personhood for the social property of identity, an area to be strengthened in the international system is the bond between data protection and the right of a person to “recognition everywhere as a person before the law.” User control over personal data is essential for this recognition in IDM.

The ability of a person to wield control over persistent identity information can help ensure his recognition over time. However, the *Privacy Guidelines* do not endow a person with a right to the continuation of identity.

The *Privacy Guidelines* have a strong focus on protecting a person’s data against inappropriate treatment by other actors; however, they place the individual in a rather passive role and so fail to provide him with the proactive right to use his own identity information as he sees fit. The law may need to lend its support to emergent IDM tools so that the user will by default have a right to make use of his personal data.

2. Identity is subjective

i) Restatement of the property

Different people have different experiences with the same individual and therefore attribute different characteristics to that individual; that is, they will construct different identities for him.

The identities others construct for a person form the peg on which his reputation is hung. For example, a business partner’s view of a person’s reputation includes his view of her identity and his interpretation of the individual attributes of her identity; he may view her identity in terms of her credit rating, her reluctance to participate in lawsuits, her history of being late, or the college she attended. Her own view of what comprises her identity may be different. People can therefore disagree about an identity, even as the information is accurate, complete, and current.

ii) How the property is accommodated by the *Privacy Guidelines*

Generally speaking, as governments and private parties use people’s personal data to provide services, democratic societies would want IDM systems to permit the people whose data is in question to know how their data is being used to make representations about them. The *Privacy Guidelines* foresee this need and call for transparency via the “Purpose Specification” and “Openness” principles.

iii) How the *Privacy Guidelines* might be augmented to address the property more fully

Regarding disputes, the *Privacy Guidelines* provide for challenges in the “Individual Participation” principle – but the idea there is to ensure conformity with other data protection provisions concerning

notice, consent, security, and access. The matter of subjectivity in IDM, however, is different: Even if all the data protection principles are followed, people can still disagree about an identity, and the disagreement might be irreconcilable if assessments are based on value judgments rather than objective facts. As a consequence, individuals or organisations who rely on third-party reports (as opposed to firsthand experience) about a person's identity need procedures for investigation and resolution of disputes concerning the information provided.

As suggested previously, there will be instances when governments need to use people's personal data without their knowledge. The *Privacy Guidelines* envision such activity in the "Use Limitation" principle if the law authorises it. However, something that could be strengthened for the sake of public confidence is the ability of people to know that any such use is in fact in line with what law provides. For example, the international system could guarantee access for independent ombudsmen who are democratically accountable at the local level, authorising them to check that any use of personal data not specified in advance strictly complies with law.

3. Identity is valuable

i) Restatement of the property

By building a history of a person's past actions, exchange of identity information creates social capital and enables transactions that would not be possible without identity. In other words, identity lends predictability to afford a comfortable level of confidence for people making decisions.

ii) How the property is accommodated by the *Privacy Guidelines*

For identity information to be valuable, it needs to be relevant for the purposes for which it is used, and it needs to be accurate, complete, and current as necessary for those purposes. The "Data Quality" principle spelled out in the *Privacy Guidelines* captures exactly this notion. The "Accountability" principle then helps to ensure that the source of information can be held to account for its accuracy, completeness, and freshness, thus allowing others to trust it.

iii) How the *Privacy Guidelines* might be augmented to address the property more fully

Reputation-rating technologies will prove increasingly important as the value-generating nature of identity is understood. Technologies such as those allowing portability of reputation across systems will help unlock this value-generating potential. An improvement on the *Privacy Guidelines* could be to update them to recognise the individual's right to bring a reputation with him into new environments, thus unleashing more of identity's value.

Similarly, interoperability in IDM systems will be important for the value-generating potential of identity information to be realised. Though admirable, data protection principles are arguably limited in providing a person with notice, consent, security, and access only. Rather than taking a defensive stance against the abuse of identity information, the principles could be improved to recognise positive rights of individuals to make use of their identity information. Interoperability would facilitate the exercise of such rights because it would enable people to make use of their data across different systems and not suffer from lock-in. The international system could encourage interoperability through competition policy or legal mandates.

For many contexts, it would make sense for identity systems to allow for collectively controlled identities. For example, many cultures emphasise group identity, and people everywhere are increasingly using ICT for group activities. Such situations conjure up notions of "juristic" or "juridical" persons – that is, entities (e.g. corporations) that enjoy legal personality with rights and responsibilities. The law

recognises juridical persons in large part because they are connected to natural persons who may be held to account. New literature is developing that fleshes out ideas for allowing the information society to reap additional rewards by allowing jointly controlled digital identities as well as “limited liability personas”.³⁰ Development of law in this regard would allow additional value to accrue through, *e.g.* the transfer of rights to control identity information.

An assurance that the *Privacy Guidelines* were operative might persuade a person to consent to the use of his data by others; similarly, if other parties knew that the *Privacy Guidelines* would be enforced, they would be more inclined to give notice to and seek consent from people whose data they were using. IDM tools may emerge to address enforcement problems of the *Privacy Guidelines*; as such, they could open up new markets for the exchange of identity information.

While the *Privacy Guidelines* envision user consent in such provisions as the “Collection Limitation” and “Use Limitation” principles, part of the challenge for effective data protection is to raise consciousness about the value of personal data, and to work toward informed user consent based on an appropriate appreciation of the value of identity information.

4. Identity is referential

i) Restatement of the property

An identity is not a person; it is only a reference to a person.

Since an identity relates to a person, a tie must be maintained between the identity and the person. The reference is necessary because rights and responsibilities attach to the person rather than his identity information. A person may develop spin-off personas so that other people know him through those various digital identities, and others may create profiles of a person; nonetheless, ultimately the collection of characteristics that signal who a person is need to point back to that person. This tie allows a person to enjoy user control over identity information; it also engenders accountability and trust among parties and society generally.

As suggested in the description of the valuable property of identity, law could develop to unleash additional value in IDM through, *e.g.* allowing the transfer of rights to control identity information. New uses of identity information underscore the need for a referential tie back to the person or people who control identity information to ensure accountability and engender trust.

ii) How the property is accommodated by the *Privacy Guidelines*

Currently a person has little knowledge of who is exercising control over his data and what they are doing with it; as such, an individual does not have an effective right “to have data erased, rectified, completed or amended” when successfully challenging data in others’ care, as provided in the “Individual Participation” principle of the *Privacy Guidelines*. The “Collection Limitation” and “Use Limitation” principles in the *Privacy Guidelines* anticipate the need to prevent the build-up of profiles detached from a person’s control. By strengthening a person’s control over his personal data, these principles can all support personhood in IDM.

³⁰ See, *e.g.* “The Limited Liability Persona,” entry posted by Bob Blakley on the Burton Group’s Identity and Privacy Strategies Blog, 17 November 2006, at: http://identityblog.burtongroup.com/bgidps/2006/11/the_limited_lia.html.

iii) How the *Privacy Guidelines* might be augmented to address the property more fully

To say that identity information should not be detached from the control of the person to whom it relates, is not to say others should necessarily be able to associate his data with him. Privacy enhancing technologies (PETs) in IDM can afford user control in a way that enables a person to minimise the extent of identity information that is disclosed and to prevent linking among transactions. By reducing the build-up of profiles, PETs help to keep the person, rather than those profiles, as the locus of recognition and accountability in the information society. In terms of the *Privacy Guidelines*, PETS can be seen as bolstering the “Collection Limitation” principle.

Again, it is persons who have rights and responsibilities; although identity information stems from persons, the identity information itself has no rights and responsibilities. Data protection is at its heart geared toward helping people rather than personal data. Given how important it is for identity information to continue to refer to the person to whom it relates, it is arguable that a person should have an inalienable right to control his personal data. The *Privacy Guidelines* do not go this far.

5. Identity is composite

i) Restatement of the property

Some information about a person arises from the person himself; he volunteers it. But much information about him is developed by other actors without his involvement.

An actor developing information about a person is the only authoritative source for it; as such, it makes sense for that actor to have certain rights and obligations regarding that information. For example, a credit agency examines a person’s credit record and creates a credit score. This credit score is a business asset of the credit agency, and the algorithm used to generate it is a trade secret or an item of protected intellectual property; the credit agency is the only authoritative source for the person’s score, and that agency may impose restrictions on the data’s use.

Because other actors have legally-recognised interests in the information they generate about a person, that person cannot have an absolute right to control it independent of other influences. Even if a person’s right to control his personal data were considered inalienable, the legitimate rights and interests of all parties must be respected when designing systems which use identity information. Similarly, other actors have obligations with respect to the data they generate about a person.

As noted above in the Definitions section, identity information can be used by different people to describe a person. Applying the ideas of philosopher Paul Ricoeur: When the data subject himself is initiating new actions through a digital identity, that identity may be referred to as “*ipse identity*”; when others act based on what they know about a person over time (for example, through others’ accounts of that person’s reliability), the identity may be referred to as “*idem identity*”.³¹

ii) How the property is accommodated by the *Privacy Guidelines*

The “Data Quality” principle applies to a person’s identity information even as others contribute to that information. Thus, those other actors may be seen as having a responsibility to ensure that the

³¹ These conceptions stem from work by Paul Ricoeur (1913-2005). As detailed on the Stanford Encyclopedia of Philosophy: “Following a distinction in Latin between *idem* and *ipse*, Ricoeur holds that the self’s *idem*-identity is that which gives the self, among other things, its spatio-temporal sameness. Its *ipse*-identity gives it its unique ability to initiate something new and imputable to himself or herself.” <http://plato.stanford.edu/entries/ricoeur/> (accessed: 7 January 2008).

information they contribute is “relevant to the purposes” for which it is used, and that it is “accurate, complete, and kept up-to-date.” The “Accountability” principle supports this requirement by ensuring that the source of information can be held to account for its accuracy, completeness, and freshness, thus allowing others to trust it.

iii) How the *Privacy Guidelines* might be augmented to address the property more fully

Some people have suggested that, rather than requiring that a mistake be made before a person has the right to have information removed from another party’s system (as provided in the “Individual Participation” principle), the law could provide that a person automatically has this right, as a sort of statement of the inalienability of personal data. Others would argue that the principles in the *Privacy Guidelines* together ensure fair information practices and that if those principles are followed, there will be no usage of a person’s data without his prior consent. The difficulty with both views is that they fail to make room for the composite property of identity. To bridge this gap, data protection could provide for inalienability of personal data and at the same time allow for others to use that information through arrangements such as fixed-length licensing schemes for the use of personal data, rights of redemption for licensed data, and periodic expiry or wipe-clean dates. Backed by the assurances of the *Privacy Guidelines*, such arrangements could ensure that a person retained control of his personal data, while at the same time allow other entities in society to build on that data.

While data protection principles provide a person with assurances of notice, consent, security, and access as regards his personal data – and thus carry obligations for other parties treating that data – the *Privacy Guidelines* are rather quiet as to the rights of actors who add value by contributing information about a person. They could be adapted to include this notion.

6. Identity is consequential

i) Restatement of the property

Because identity tells of a person’s past actions, the decision to exchange identity information carries consequences: Disclosure of identity information in a certain context can cause harm; failure to disclose identity information in another context can create risk.

Disclosure of identity information in an improper context (ethnicity in housing applications, sexual orientation in employment applications, etc.) can cause harm. At the same time, failure to disclose identity information in a certain context (credit history in a credit application context, criminal history in a primary education context, etc.) can create risk for individuals or organisations relying on identity information.

Therefore, it is important to avoid disproportionate bias in favour of the interests of any party: If socially stable structures are to be created, risks to all parties, not just the person whose data is in question, must be considered when rules are designed. Because identity allocates risk, accountability (and traceability of actions to identified individuals) is necessary in cases where risks created are large.

To say that identity is consequential is to say that the operation of an identity system itself may carry risks. Thus, there needs to be attention to ensuring that there is *no release of data that shouldn’t be released*. In other words, this property is concerned with inner-system workings. (This concern differs from concern over data release in the valuable property of identity: with identity being valuable, the concern about data release is outside the system and focuses on the active usage of data through a user’s purposeful release of his data to increase functionality and reduce risk.)

ii) How the property is accommodated by the *Privacy Guidelines*

Fair information practices generally, and the *Privacy Guidelines* specifically, may be viewed as a testament to the fact that identity is consequential: They are geared towards giving a person notice, consent, security, and access with respect to his personal data. In IDM, it is technically possible for parties to provide the type of notice and negotiate the kind of consent envisioned in the principles of “Collection Limitation”, “Purpose Specification”, and “Use Limitation”. Similarly, IDM systems can include mechanisms to provide users with security as spelled out in the principles of “Data Quality”, “Security Safeguards”, and “Accountability”. In terms of access, IDM tools can enable a person to know how data is treated and contest that treatment, as foreseen by the principles of “Openness” and “Individual Participation”.

iii) How the *Privacy Guidelines* might be augmented to address the property more fully

A key question is whether law will require IDM systems to provide this technical support for fair information practices, or whether governments will leave it to the market to manage risk.

Another question is whether the same standards will apply to private and public actors in managing the risk associated with identity information.

7. Identity is dynamic

i) Restatement of the property

Identity information is always changing; any particular identity dossier might be inaccurate at any given moment.

People move; their names change, their ages change, their employer changes, their health status changes. Identity information is always changing; any particular identity dossier might be inaccurate at any given moment.

ii) How the property is accommodated by the *Privacy Guidelines*

An individual must have a right to access his data to determine accuracy and to correct errors and out-of-date entries. The “Data Quality” and “Individual Participation” principles in the *Privacy Guidelines* anticipate the need for such quality and access.

iii) How the *Privacy Guidelines* might be augmented to address the property more fully

The data protection principles do not attend to *others'* interests in having a person's identity information be accurate and up-to-date. IDM tools of the future could enable automated checks for accuracy, completeness, and freshness, along with corrections. Future data protection principles could factor in such capabilities.

Certain types of identity information should be stored as a snapshot in time for later verification in the event of a subsequent dispute or investigation. To safeguard the interests of all parties, data protection might delve into these requirements.

8. Identity is contextual

i) Restatement of the property

People have different identities that they may wish to keep entirely separate. Information can be harmful in the wrong context, or it can simply be irrelevant. Keeping identities separate allows a person to have more autonomy.

An identity attribute that is relevant in one context (say, the fact that someone has billions in assets when applying for a loan) perhaps should not be mentioned in another context (say, when that same person is drafted for the army).

The fact that people put on different faces is not just normal but good – because the combination of riskiness and contextuality requires it, and because people sometimes simply enjoy being able to act in different capacities. So, for example, a policeman may project resoluteness on the job but be a soft daddy at home.

A pseudonym can become a legitimate and predictable identity in itself, much as James Bond behaves similarly regardless of which actor portrays him. From the point of view of relying parties, however, anonymous interactions, though, must be entered into with care because there is little background information about the identity of counterparties to calibrate behaviour.

To say that identity is contextual is not to say that digital identities should be “siloeed” the way they have been in early IDM systems, where separate contexts have been maintained by preventing the transfer of identity information beyond limited domains. Rather, what is important here is that the exchange of identity information should not allow a linking of information from different transactions, unless a user specifically wishes to do so.

ii) How the property is accommodated by the *Privacy Guidelines*

The “Use Limitation” principle calls for consent by a person if his personal data is used, unless the law authorises use without consent.

Backing up Use Limitation, the “Openness” and “Individual Participation” principles afford means for checking that treatment of personal data is appropriate.

As noted already in discussion of the referential property, PETs in IDM can provide user control in a way that enables a person to minimise the extent of identity information that is disclosed and to prevent linking among transactions to keep contexts clean. By reducing the build-up of profiles, PETs help to keep the person, rather than those profiles, as the locus of recognition and accountability in the information society. In terms of the *Privacy Guidelines*, PETs can be seen as giving effect to the “Collection Limitation” principle.

The “Openness” and “Individual Participation” principles can be viewed as providing ways to check that unwanted linkages are not made between different identities.

iii) How the *Privacy Guidelines* might be augmented to address the property more fully

If a person could rely on the principles of “Collection Limitation”, “Purpose Specification”, and “Use Limitation” in IDM, he could choose the contexts in which he let his information be used and thereby have a consent role in any linking that occurred. Of course, to exercise this capability he might need the

assistance of electronic agents or recommendations from groups he trusted. The challenge here is to choose the right policy for encouraging IDM systems that will give effect to these principles.

Regarding the perceived threat to others posed by people acting anonymously, many cryptographic schemes can maintain user anonymity in IDM while making it possible to detect and trace fraud.³² The *Privacy Guidelines* could be revised to recognise people's right to remain anonymous or pseudonymous so as to protect users' rights proactively to use their different identities.

9. Identity is equivocal

i) Restatement of the property

The process of identification is inherently error-prone.

People can have similar names and histories, and they can look alike. Secrets used to authenticate individuals can be lost or stolen. Malicious people can impersonate others and "steal" their identities. Technical identification systems always generate some number of false positives and false negatives.

Given the tendency for errors in the identification process, identification should be subject to appeal and correction, whether it is a data subject or another party who is hurt by a mistake.

ii) How the property is accommodated by the *Privacy Guidelines*

The principles of "Data Quality", "Security Safeguards", and "Accountability" aim to provide appropriate incentives and reduce the probability of errors. Meanwhile, the "Openness" and "Individual Participation" principles in the *Privacy Guidelines* are geared toward promoting transparency and affording opportunities to redress grievances.

iii) How the *Privacy Guidelines* might be augmented to address the property more fully

To protect the interests of parties relying on a person's identity information, individuals could be asked to verify, on an ongoing basis, the accuracy, completeness, and freshness of data they have chosen to release. Here again, as the technology is available, the law could require this verification, or government could take a *laissez faire* approach.

Knowing what has taken place is key to resolving disputes. The OECD could initiate work on defining auditability standards for IDM tools. More generally, the OECD could even aim to establish a "gold standard" for IDM tools that conform to the *Privacy Guidelines*.

All in all, there are serious hurdles for implementing effective notice, consent, security, and access in IDM systems. If resources are to be allocated to trying to meet demand for these fair information practices, decision-makers would be wise to factor in the Properties of Identity.

³² For an example of how privacy and auditability can simultaneously be designed into a system, see e.g. Choi, J.Y., Jakobsson, M., and Wetzel, S., "Balancing Auditability and Privacy in Vehicular Networks," in Proceedings of the 1st ACM International Workshop on Quality of Service & Security in Wireless and Mobile Networks (Montreal, Quebec, Canada, October 2005), Q2SWinet '05, ACM Press, New York, NY, 79-87, at <http://www.cs.stevens.edu/~swetzel/publications/balancing.pdf>.

TABLE: ASPECTS OF IDENTITY PROGRAMMES THAT DEMONSTRATE THE PROPERTIES OF IDENTITY

One immediate area for which the Properties of Identity are relevant is IDM in national identity cards. This table highlights how various plans to use IDM in identity documents demonstrate the Properties of Identity.

<i>IDENTITY IS...</i>	
Social	Costa Rica is considering a constitutional amendment providing that every person has a right to have, or not to have, what is referred to as “virtual personality”. Meanwhile, the Spanish Law on Citizens’ Access to Public Services expressly frames as one of its objectives to “facilitate the exercise of rights and fulfilment of duties in electronic means” and to recognise the right of every citizen to acquire the necessary means of electronic identification — <i>i.e.</i> linking legal personhood and a right to electronic means of identification.
Subjective	Governments require people to use machine-readable travel documents (MRTDs) to cross borders; presumably most of the identity information matched with a given passport has not been compiled by the individual in question.
Valuable	As noted immediately above, governments rely on MRTDs and profiles to see if people should be permitted to cross borders. With the efficiency and security these identity cards afford, large numbers of people can cross borders with relative ease every day.
Referential	The US Homeland Security Presidential Directive 12 (HSPD12) of 2004 calls for standards for secure and reliable forms of identification for government employees and contractors to prevent terrorists from gaining access to federal facilities. “Identity proofing” processes use biometric technology, with personal data stored on a smart card. Information is to be decentralised, under the control of the individual.
Composite	Germany has its citizens use e-Health cards to administer the healthcare system. While the citizen is understood to have certain rights regarding data relating to him, it is accepted that the government is allowed to use that data for certain specified purposes.
Consequential	A recent Resolution of the Mercosur/Mercosul (including Argentina, Brazil, Paraguay, and Uruguay) Common Market Group approved common rules for the recognition of electronic documents, electronic signatures and advanced electronic signatures within the trade bloc. There are still very different levels of technological development among the participating countries’ infrastructures. To offset risk, the Resolution establishes strict liability and minimum data protection standards.
Dynamic	The United Kingdom’s Department of Work and Pensions (DWP) is working on a “change of circumstances” mechanism to give citizens a single point of contact anytime they need to change information; the Identity and Passport Service is to base the national identity card on the DWP’s database.
Contextual	The Belgian <i>Carte d’Identité Électronique</i> , with several million cards already issued, keeps identity information separate. Besides having the citizen’s picture, national registration number, and other data printed and stored on it, the card contains a pair of keys and two certificates – one for identification, the other for signature. With the Austrian identity-card system, each citizen is assigned “Sector Specific Personal Identifiers.” For example, according to this scheme, a person can choose whether to present his medical identity or his education identity.
Equivocal	In 2007 the People’s Bank of China and China’s Ministry of Public Security (MPS) jointly launched the Online Verification of Citizens’ Identity Information initiative for banking. A resident permit may be used to contest an identification.

Of course, while a particular plan may be listed as an example of a given property, that plan would need to encompass all of the properties together to be consistent with the Properties of Identity as a whole.

THE PROPERTIES OF IDENTITY FOR POLICY MAKERS AND SOFTWARE DEVELOPERS³³

*The Properties of Identity can serve as a helpful guide as governments seek to set and administer policy, and as software developers design for an identity infrastructure. Again, the Properties of Identity serve not as instructions for the way things **should be**, but rather simply as a factual observation of the way things are, and, as such, can serve as very basic predictive measures for the soundness of legal and technical choices relating to IDM.*

Properties of identity for policy makers

As policy makers work with technologists to consider approaches for the identity infrastructure, they would do well, consistent with the values of democracy and an open economy, to use the Properties of Identity as a checklist to test the viability of policy choices at an early stage.

The role of policy makers is important for IDM development because these people develop the framework for bringing law and computer code into alignment with policy goals and can shortcut the time-consuming trial and error of the competitive marketplace.

It is helpful to note that there are two groups of policy makers here: *i*) those who set and administer policy at the international level, and *ii*) those handling policy at the national or local level. Both have a critical role to play in ensuring that the identity infrastructure supports the values of democracy and an open economy for a free and open society.

International policy makers dealing with IDM increasingly are using the settings of intergovernmental organisations to develop legal and technical standards with broad reach. Organisations focusing on IDM include the International Telecommunication Union, the International Organization for Standardisation, the International Civil Aviation Organization, and the OECD, among others. By using the Properties of Identity as markers for sound policy, the international system can encourage the achievement of shared objectives such as the bolstering of data protection and the creation of an identity infrastructure that fosters accountability and trust through user control.

Local policy makers can meanwhile serve on behalf of their publics in a more direct way. For example, when information sharing occurs among governments and involves the participation of private actors across jurisdictions, citizens need to know that their data is in good care. Although the individuals themselves cannot always be privy to what is being done with their data (*e.g.* if authorities are screening for money laundering), democracy demands that there nonetheless be accountability in the system. Hence, there is an ombudsman role for local officials to play in verifying that citizen data receives proper treatment and is safeguarded from misuse if shared with actors in other jurisdictions. Another example is the role local officials have in checking that their constituencies do not suffer economic discrimination, for example if a power attempted to block one market's citizens from participating in trade. While individual citizens might not be able to see patterns of discrimination, democratically accountable governing bodies could look at patterns and defend citizen interests in an open economy.

If policy makers see themselves as having a democratic responsibility to build an identity infrastructure that allows users to control their identity information, IDM policy choices will more likely promote accountability and trust – and thus enable a free and open information society. The Properties of Identity can serve as guideposts as decision-makers craft policy to this end.

³³

This section was written by Mary Rundle, Mary Ruddy, and Marcelo Thompson Mello Guimarães.

Properties of identity for software developers

As engineers develop technology for the identity infrastructure, they would do well, given market demand, to use the Properties of Identity as a checklist to test the viability of their designs at an early stage – with the Properties of Identity serving not as instructions for the way things *should be*, but rather simply as a factual observation of the way things *are*.³⁴

The role of engineers responsible for architecting and coding software systems matters because, as Lawrence Lessig has articulated,³⁵ the code that runs society's computer systems acts as a form of regulation – but this regulation is developed without public input or oversight.

There are two classes of developers to note here: *i*) application developers creating new systems, who should be encouraged to build on IDM components that are in line with the Properties of Identity, and *ii*) the vastly smaller class of developers capable of contributing to the IDM layer itself. This latter group must be made aware of the emerging opportunities to contribute to open source projects that will bring a Properties of Identity-based IDM infrastructure into fruition. Development projects in which they can participate are sponsored by PRIME, Eclipse, the Liberty Alliance, and other organisations.

As the Internet and other enterprise-crossing applications proliferate, and as identity protocols continue to be defined and redefined, there is a growing need to make it easier for developers to support and integrate with different identity protocols without having to recreate and test a new integration method every time. By using existing, vetted identity frameworks and modules, they can save time and reduce the chance of introducing bugs.³⁶ When it comes to applications, it makes sense to use a standard framework to bake in IDM interoperability. This is especially true now that many applications are networked.

A number of trends affect a software engineer's perspective today, namely: Internet and enterprise identity are converging; security risks have become higher profile (even "soccer moms" worry about identity theft); open source technology, vetted by many eyeballs and reusable, has become accepted; there has been a proliferation of identity protocols and it seems likely that this aspect of the infrastructure will stay heterogeneous (*i.e.* there will be many identity protocols in production for a long time); increasingly, applications and services cross multiple contexts, and more systems and applications are networked; and there is a growing awareness that developers need to make Internet identity easier to implement and use – that a different approach is needed. (Some call this goal the identity layer "big bang".) The increase in software virtualisation and the rise of Service Oriented Architectures (SOA) has created a situation where almost every new software application needs to include IDM capabilities, and the traditional IDM tools and techniques are not flexible enough to meet the new demands.

While the market will naturally favour modules that are aligned with the Properties of Identity, waste could nonetheless be avoided by raising awareness about them. Developer kits can facilitate education and implementation of systems that support the properties.

If software engineers see themselves as having a vested interest in building an identity infrastructure that runs consistent with the Properties of Identity, the code behind the information society will be much more likely to put control of identity information in the hands of the user – and thereby support a free and open society.

³⁴ The Properties of Identity presented in this paper describe observable qualities of the way identity behaves, or that are. Kim Cameron has outlined "The Seven Laws of Identity", which spell out characteristics of an ideal IDM system; those characteristics of the way a system should be, may be viewed as a prescription for code design that implicitly factors in the Properties of Identity.

³⁵ Lawrence Lessig, *Code and Other Laws of Cyberspace*, Basic Books, New Ed. Edition, 2000.

³⁶ This is especially important in applications where security and privacy are a concern.

CURRENT CONCEPTIONS OF IDM³⁷

Even if it appears the market and democratic governments should favour user control, there is still the question of how the identity infrastructure will get from here to there. In this section the paper describes “Current Conceptions of IDM” to provide a practical grounding in what might constitute user control in IDM arrangements.

Core identity?

In discussions of digital identity one often sees “identity” defined as “a collection of attributes” or “a collection of claims” or “partial identity”. The question naturally arises, “attributes of what?” or “claims about what?” or “partial identity of what?” These questions have no clear-cut answers. The answer “a subject” is too vague; the answer “a person” leads to arguments about how one distinguishes a natural person from a legal/juridical person, and perhaps even from a non-person; there may also be arguments about whether multiple persons can inhabit the same body, either at the same time or sequentially. And the answer “a human body” creates confusion about matters of intent and continuity of memory, which are important when making decisions about reward and punishment.

In most practical cases it simply does not matter whether persons have immutable core identities; it is usually sufficient to answer an easier question: “Is this person the same person who did ‘X’ in the past?” While the (possible) lack of a core identity can create a small amount of doubt about whether our suspect previously did ‘X’, the evidence for or against the identification of the suspect is normally strong enough to justify confidence in the identification.

This paper takes no position on the question of whether persons have core identities, but it does take the position that core identity is not observable by parties other than the subject himself – so identification systems need to operate on the basis of recognising attributes or establishing the truth of claims.

Looking at it from the other direction, however, the paper takes a strong stance that the concept of personhood must be bolstered by treating personal data as inalienable. Without this tie, a person loses autonomy and society cannot flourish as free and open. The paper acknowledges that some policies that are focused on core identities may have as their aim to support personhood.

Identifiers, attributes and claims

Identity is not a phenomenon of the digital age; it is an ancient and fundamental human trait. Humans believe innately that they have identities, and they also perceive other humans as having identities. Ethnicity, gender, family ties, membership in cliques, societies, and other associations, eye and hair colour, primary language, and many other human traits and behaviours are constituents of human identity. Human conceptions of identity are subjective; a woman living in Detroit whose great-grandparents are six Germans, one Japanese, and one Irishman may self-identify as Irish-American because she inherited the family name “O’Brien” from her father’s father’s father and she has an aesthetic fondness for Irish music and literature. But the census will call her “Caucasian”, and a medical research study might identify her ancestry as “German”. People on a street in Detroit may recognise her as “Asian”, while people on a street in Caracas might see her as “American”.

³⁷

This section was written by Bob Blakley.

Digital systems are not particularly adroit at handling subjective concepts like identity. When today's digital systems translate human identity into digital form, they tend to simplify it and squeeze out ambiguity and contradiction. This simplifies processing of identity information, but it also leads to situations which confuse or upset the humans whose identities are being digitally manipulated.

This section defines a digital identity as the combination of two elements: an identifier and a collection of claims.

An identifier is simply a name – it can be a name which is comprehensible to a human (for example, “Bob” or “Alice”) or a name which is comprehensible to a computer system (for example, “515-99-7777” or “0xfa102b66”). A digital identity's identifier refers to the identity's collection of claims.

Claims are statements about the subject of the identity (for example, “The subject is 5'8” tall,” or “The subject was born on 15 August 1947,” or “The subject holds account number 1234567890 at Last Local Bank,” or “The pattern of ridges on the subject's right index finger look like this picture.”)

Some claims are statements about the subject's behaviour or possessions (“The subject holds account number 1234567890 at Last Local Bank.” or “The subject prefers Vodka to Beer.”) Other claims describe attributes of the subject (“The subject is 5'8” tall.” or “The pattern of ridges on the subject's right index finger look like this picture.”) The difference between claims about attributes and other types of claims is that, at least in principle, claims about whether a subject has a particular attribute can be verified or falsified by examining the subject, whereas verifying other types of claims may require examining historical records, questioning witnesses, or doing other research for which the physical presence of the subject is neither required nor helpful.

Producers and consumers of identity

Identity claims are asserted by a wide variety of persons and organisations. Governments assert claims about citizenship, qualification for regulated activities (such as driving or membership in regulated professions), criminal history, and other attributes. Businesses assert claims about employment status, expertise, and authority to perform certain business-related transactions. Religious organisations assert claims about faith affiliation and marital status. Credit agencies assert claims about payment history and loan repayment risk. Doctors assert claims about health status.

Identity claims are also consumed by a wide variety of persons and organisations for a wide variety of purposes. Insurance companies consume claims about health status. Banks consume claims about credit history. Employers consume claims about work and criminal histories. Governments consume claims about citizenship.

Authoritative sources for attributes and claims

The claim “The subject's credit score is 605” is asserted by a credit agency; the same claim would not be treated as reliable if it were asserted by the subject's priest. The claim “The subject is a non-smoker” may be asserted by the subject himself; this claim would not be considered reliable if it were asserted by the subject's banker. The claim “The subject's name is ‘Bob’” is typically asserted by the subject's parents, or by a civic authority in the municipality of the subject's birth; this claim would not be considered reliable if it were asserted by an online merchant from whom Bob had bought books.

It is often important to accept identity claims only from authoritative sources; the department of motor vehicles would not be authoritative for a claim about the subject's credit score, but it would be authoritative for claims about the subject's qualification to operate an automobile.

Ownership and control of identity information

Who owns a particular identity claim may be complicated. Some legal systems grant subjects the right to control certain identity information; other legal systems do not grant this right. Many identity claims and attributes have complex usage rules; a subject's medical record information can usually be used by medical professionals for many treatment purposes without the subject's permission, but may not be disclosed for purposes other than treatment. A photograph of a subject's face is considered by many legal systems to be "owned" by the photographer, but there are complicated rules for the circumstances in which the photographer may use the photograph - and these rules differ from country to country. Police surveillance photographs are governed by very different rules than those regulating photographs taken by artists and journalists, and these rules in turn are very different from those regulating photographs taken for the purpose of advertising a product.

Simplistic conceptions ("information is free", or "subjects own their identity information") are simply wrong; the real-world complexity of restrictions on and entitlements to the use of identity information have to be taken into account both when formulating public policy which governs identity information and designing technical systems which process that information.

Identity and privacy

Identity is inextricably related to privacy. Some identity claims (infectious disease status or criminal history) are inherently private; other claims (religious affiliation or trade union membership) may be private in some contexts but not in others; some attributes (a woman's facial appearance) may be private in some societies but not in others.

Privacy is respected when identity claims are revealed only in accordance with fair information practices; these practices have been designed to ensure that subjects' dignity is respected by disclosing information only in situations where disclosure does not harm the dignity of the subject (unless some compelling societal interest compels disclosure).

Pseudonyms

Many transactions in which identity claims are important do not depend in any important way on knowing the actual name of the subject; in many commercial transactions, for example, what is important is to know who will guarantee that the merchant gets paid, and to what address the merchandise is to be shipped. A credit card number and address suffice for these transactions – the subject's "correct" name is not really required.

In a theme park, it is sometimes important for safety reasons to know that riders are physically big enough to withstand the forces generated by an extreme ride. In these cases the subject's height and weight are important, but the subject's name is irrelevant.

In transactions like these, subjects are often permitted to use pseudonyms, or to use no name at all (that is, to be anonymous).

The use of pseudonyms coupled with reliably asserted and situationally appropriate identity claims can help preserve subjects' privacy in many kinds of transactions.

Identity and authentication

Digital identity arose from the need to allow users to share computer systems. When multiple users share a system, it is necessary to give each user a place for his or her own data, to ensure that users

consume resources fairly, and to protect users against interference by other users. Early timesharing systems implemented “user accounts” to solve these problems; each user received an account with a resource quota and a dedicated area of storage which could be used to hold programs and data. To prevent users from accessing each other’s accounts, each account was given an “identifier”, and access to the account was protected through the use of a secret “password”.

With the emergence of messaging and electronic mail systems, it became necessary for users to designate other users as recipients when sending messages; to meet this need, messaging system users began to think of account identifiers as “usernames”.

Early messaging systems allowed users to communicate with other users who had accounts on the same machine; as networking became more ubiquitous, however, the requirement to communicate with non-local users became more and more important. This created a problem: how to identify users in a globally unique way, so that a message could be sent from any user on any system to any other user on any other system.

Second-party identity systems

Early solutions to this problem focused on a primary relationship between a user and his computer system host – often an employer. Users were issued accounts on a host system and given passwords to allow them to “authenticate” their rights to use those accounts. Unique names were built by specifying the host name and then the account name (or “userID”) within that host’s naming “domain”. This domain-based naming structure was eventually given a standard syntax by IETF RFC 822, and became the standard naming convention for both e-mail addressing and website location. Thus, John Doe became “jdoe@bigorganization.com” and the web page for the sales department at John’s company became “www.bigorganization.com/sales”.

John’s account is maintained, in this model, by the organisation which owns his domain. This is typically the organisation providing services (for example, e-mail and application access) to John. Because the service provider is managing John’s account, the model is called a “second-party” identity system. John is the “first party”; he’s the one who has the account. The service provider is the “second party”; it offers services, and it also creates and maintains the account which gives John access to those services. Maintaining John’s account requires the second-party organisation to manage John’s password, to allow him to change his password periodically to guard against various kinds of attacks, and to allow him to reset his password when he forgets it. If the services offered by the domain are customisable, the second-party service provider will also have to maintain “profile” information (for example: home address, age, telephone number, and so on) on the basis of which services are customised. This profile information essentially constitutes the organisation’s view of John’s “identity”, in the same way that John’s account identifier constitutes the organisation’s view of John’s “name”.

When John wants to start using his account, he needs to authenticate himself. He does this by contacting the system and requesting an “authentication dialog”. John initiates the dialog by naming the account he would like to claim; the system then challenges him for a password or other authentication data, runs some checks to ensure that it has received the correct response, and (assuming John has responded correctly to the challenge) grants John access to his account.

Issues with second-party identity systems

This is convenient as long as John has relatively few accounts. As soon as John’s use of computing resources starts to expand, however, two problems emerge:

- i) John has to create accounts on many different systems. There is no guarantee that John can get the same account name on each system he uses, so he may have many different “usernames”. He should not, for security reasons, use the same password for all of the systems on which he has accounts; therefore he will have many different passwords. And since the various domains which host accounts for John do not communicate with one another, John has to enter profile information many times in many different systems. This proliferation of account names, passwords, and profile data is inconvenient for John.
- ii) If two different organisations provide service to John and want to co-ordinate their service delivery to him, they need to co-ordinate their separate views of John. Since John has an account with each organisation, there is a good chance the two organisations have different “names” for John. And since each organisation maintains a separate profile for John, they may have different – and even conflicting – information about him. Finally, the two organisations are in some sense wasting resources, because they are maintaining two separate profiles with (mostly) the same information.

Trust characteristics of second-party identity systems

In a second-party identity system, the user relies upon his service provider for account management and profile management. As a result, the user is often faced with a take-it-or-leave-it proposition: accept the service provider’s security and privacy policies and practices, or forego using the services offered by the provider. If disagreements about the use of personal information arise between the user and the service provider, the user is at a disadvantage because the service provider is in possession of the user’s information, and the user has no independent advocate to whom to appeal information-use decisions. This asymmetry has given rise to regulations in various jurisdictions; the European Union created the position of “data protection commissioner” to defend users against abuses of personal information by second-party identity providers.

Federation

As users turned increasingly to e-commerce and mobile work, the second-party identity issues became more and more acute. Enterprises and service providers searched for ways to decrease the burden imposed on users by the need to manage multiple accounts with many different organisations. Two technologies emerged from this search: single sign-on and federation.

Single sign-on reduces the number of times a user has to remember and use a password. Single sign-on does not actually reduce the number of logon events; instead, it uses client-side technology to automate logons and hide them from the user, while still protecting the security of user passwords and account information.

Federation reduces the number of accounts a user has to create. It does this by building trust relationships between service providers, which in turn allows one service provider to rely on another service provider’s authentication of the users.

In a federated environment, John authenticates himself to his primary domain (perhaps his employer – bigorganization.com), and then his primary domain authenticates him to all the other domains in the federation. This way, John sets up only one account (the primary domain account) and logs on only once (to the primary domain). All other authentication is handled in the background.

This is great for John, because he only has to remember one account name and one password. But it is also great for the organisations who are members of the federation, because most of them do not need to create and maintain an account for John in order to offer him services. By relying on John’s primary

domain to authenticate him, the other members of the federation avoid the cost of managing John's account.

Issues with federation

Federation addresses both of the second-party identity system issues enumerated above, but it gives rise to a new issue. Federation works well as long as users get their services from a small, well-defined collection of domains whose operators have business relationships with one another. It works less well for users who need to access services from large numbers of domains that do not have pre-existing relationships with one another. In these environments, if John needs to access a service provided by a domain with which his primary domain has no relationship, he must go directly to the service provider domain and establish an account there. His account is not portable; he is at the mercy of the organisation which controls his primary domain. If that organisation chooses not to establish a federation relationship with one of John's preferred service providers, John cannot use his federated account to access that service provider.

Trust characteristics of federated identity systems

Federated identity systems still have the same basic trust characteristics of second-party identity systems, with the user having to rely upon his primary domain provider for account management and profile management, and, as a result, being faced with a take-it-or-leave-it proposition: accept the primary domain provider's security and privacy policies and practices and the primary domain provider's choice of partner service providers, or forego using the services offered by the provider and all the other members of the federation.

Again, as is the case with second-party identity systems, if disagreements about the use of personal information arise between the user and the primary domain provider, the user is at a disadvantage because the primary domain provider is in possession of the user's information, and the user has no independent advocate to whom to appeal information use decisions. Furthermore, John's primary domain will typically pass information about him to federation partners; the rules for sharing information among federation partners are typically defined by business partner agreements between the partners (though the terms of these agreements may of course be influenced by laws and regulations). Users usually have no input into these business-partner agreements, and may not even be aware of their existence.

“User-centric” identity systems

Single sign-on and federation technologies address the user convenience and provider cost issues created by second-party identity systems, but they do not address the asymmetry of the relationship between users and service providers which is created by the requirement for a user to trust a service provider with his personal information in order to receive services.

“User-centric” identity systems were invented to give users more control of their personal information by allowing them to choose identity providers independently of service providers.

The goal of a user-centric identity system is to enable the creation of identity providers who operate in the user's interest rather than in the interest of the service provider. To support this goal, user-centric identity systems incorporate three components:

- i)* “Identity providers” store user account and profile information and authenticate users.
- ii)* “Relying Parties” enable service providers to accept “claims” about users from identity providers. These claims take the place of an authentication dialog with the user. So, for example, instead of

logging on directly to bigorganization.com with his jdoe@bigorganization.com identity, John logs on to identityprovider.org, and identityprovider.org sends bigorganization.com's relying party a claim which says "the sender of this claim is jdoe@identityprovider.org".

- iii) "Identity selectors" allow users to choose which identity provider to use with (and what information to disclose to) a particular service provider.

There are two key differences between user-centric identity systems and federated identity systems:

- i) In the user-centric system, John has more protection against attempts by service providers to coerce him to accept adverse security and privacy terms. The user-centric system allows John to select an identity provider on the basis of its security and privacy policies and practices, and a service provider on the basis of the quality of the services it offers. The service provider may still attempt to coerce the identity provider to adopt terms which are disadvantageous to John, but the identity provider has more leverage in negotiating with the service provider than John does, because it represents many users. In federated (and second-party) systems, John has to choose between using a provider's services and accepting its security and privacy terms, or declining the security and privacy terms but foregoing the services.
- ii) In the user-centric system, John can use his credentials with a wider variety of service providers than in a federated system. In a federated system, service providers must establish trust relationships with each other before John can use a credential from one service provider to access services offered by a different service provider; if John wishes to obtain services from two organisations that compete with each other, they may be unwilling to establish a trust relationship, and John will simply have to establish separate accounts with each service provider. On the other hand, in a user-centric system, both competitors may be willing to establish a trust relationship with John's chosen identity provider, since it is a neutral third party and not closely allied with either service provider.

These differences between federations and user-centric identity systems arise because the user-centric system gives the user a place at the bargaining table, while in the federation only relying parties have seats at the table. In a federation, several second-party relying parties reach an agreement to accept a designated set of credentials (usually each member of the federation accepts all other members' credentials, or all members accept credentials from a single designated authentication provider), whereas in a user-centric identity system, each relying party reaches an agreement with each user on what credential it will accept for that user.

First-party configurations

User-centric identity systems can be used in at least two modes: first-party and third-party. A first-party mode is one in which the user implements and operates his or her own identity provider. The user is thus in complete control of the identity claims stored by the identity provider, and of which identity claims are provided to which relying parties.

Issues with first-party configurations

A first-party system has a fundamental weakness, however: The relying party in a transaction has to depend upon the user to make accurate claims about himself. This reliance, of course, provides the user with considerable scope for fraud and misrepresentation. It is unlikely that relying parties in user-centric identity systems will rely on first-party identity providers for high-value transactions, because of the fraud risk created by the absence of authoritative third-party sources for claims.

Third-party configurations

A user-centric identity system operated in third-party mode avoids most of the disadvantages of second-party systems, federations, and first-party user-centric systems.

A user-centric identity system in third-party mode has the following advantages over a second-party identity system:

- It does not force the user to rely upon his transaction partner to manage and protect his identity information.
- It does not force the user to create a new credential for every transaction partner.

A user-centric identity system in third-party mode has the following advantages over a federated identity system:

- It allows the user to negotiate acceptance of his credentials with each transaction partner, even when transaction partners have no relationship with one another.
- It allows the user to select an identity provider which is independent in the sense that it has no relationship with any of his transaction partners (this makes it more likely that the identity provider will act in the user's interests rather than in the interests of the user's transaction partners).
- It allows the user to select which identity claims will be disclosed to each transaction partner.

A user-centric identity system in third-party mode has the following advantage over a user-centric identity system in first-party mode:

- It does not require transaction partners to risk fraud by relying upon claims the user makes about himself.

Issues with third-party configurations

There are still some issues with user-centric identity systems operated in third-party mode. In particular, identity claims are still transmitted (albeit under user control) from identity providers to relying parties. Relying parties can still compromise user privacy by disclosing identity claim information they receive about users; they can also aggregate information about a user over time if they receive different claims in support of a variety of transactions with the same user over time.

DECISIONS AND CONSTRAINTS³⁸

This section lists some of the concrete decisions that must be made in the near term regarding IDM policy and technology, and then calls to mind some of the constraints that set the larger context within which these decisions must be made.

Decisions

As identity becomes increasingly critical for activities both in the digital world and in a security-conscious physical world over the next few years, a number of important decisions will have to be made. The compilation below (numbered for ease of reference) serves simply as an illustrative list:

- i)* What identity information must individuals disclose, to whom, and under what circumstances?
- ii)* Conversely, what information may individuals withhold, and what questions are they entitled to refuse to answer about their identity information?
- iii)* What rights should be granted to individuals with respect to disclosure and use of their identity information?
- iv)* Conversely, what restrictions should be placed upon the behaviour of recipients of identity information?
- v)* What activities may individuals engage in anonymously or under pseudonyms?
- vi)* What types of investigations of individuals' identity information are permissible, for what purposes, and by what kinds of organisations or individuals?
- vii)* When reputation and profiles are developed and used for national security purposes, can the subjects of these assessments, or ombudsmen acting on their behalf, have an effective legal right and appropriate technical means for checking the assessments' accuracy?
- viii)* What due process should be available to individuals who assert that an identity claim made about them is false?
- ix)* What types of identity services should be available to individuals, what protections must they offer to their customers, and how should they be regulated?
- x)* What role should government play in identifying individuals, what roles should be allocated to private-sector organisations, and what roles should be reserved to individuals themselves?
- xi)* What penalties should be established for misuses of identity information by governments, private-sector entities, and individuals?
- xii)* What organisations will be considered authoritative for identity claims, and for which claims will each organisation be considered authoritative?
- xiii)* What technical measures will be considered strong enough for establishment of identity online?
What technical measures will be considered strong enough for establishment of identity in the physical world?
- xiv)* What technical measures will be considered strong enough for the protection of sensitive identity information?

³⁸

This section was written by Bob Blakley.

Constraints

Technology and law can solve many problems, but identity problems are complex and subtle, and there will be limits to the effectiveness of technical and social identity systems. Decision makers will have to factor in the following types of constraints (lettered for ease of reference) when designing identity systems:

- A.* Many parties (governments, media, private businesses, individuals, and society generally) have conflicting rights in identity information about individuals. These rights (including press freedom, publicity, and privacy) must be balanced carefully in an environment in which information circulates globally at the speed of light.
- B.* The process of identification is inherently uncertain; technical, social, and legal systems must acknowledge this uncertainty and deal with it in a fair and humane way.
- C.* Natural disasters, poverty, war, and other misfortunes will ensure that there will always be undocumented individuals. These individuals' dignity and human rights must be respected despite the impossibility of establishing their identities in the usual way.
- D.* Goals of ascertaining the identity of individuals in some situations must not be allowed to preclude operating anonymously or pseudonymously in others, as there are valid activities, such as voting, where anonymity or pseudonymity is crucial to the integrity of the process.
- E.* Disclosure of information is a one-way process; sensitive information cannot be "un-disclosed". Technical and social processes can provide individuals with only limited protection against consequences of accidental or inappropriate disclosure of sensitive identity information.
- F.* Individuals establish and maintain multiple identities; this is normal, appropriate, and even necessary in some cases. This fact must be taken into account and accommodated by social and technical identity systems.
- G.* Questions of identity will always create conflict, because identity is consequential and valuable. Social and technical identity systems must provide fair processes for resolving disputes about identity claims.
- H.* Any organisation that maintains an extensive collection of identity information about individuals has considerable power, which it will be reticent to give up. That power must be balanced with appropriate accountability.

These constraints should be taken into account when designing and implementing IDM policies. If those policies factor in the Properties of Identity, they will be much less likely to come under strain once implemented.

CONCLUSION

Personhood in the modern era demands that individuals have the ability to create, manage, and protect identities. Each modern personal identity exists simultaneously in the physical world and in the electronic world, and things that happen in one world can spill over into the other. For example, theft of online banking information can result in a loss of real money in the physical world.

Managing and protecting identity requires different things at different times; it sometimes requires keeping identity information secret, but at other times it involves not only revealing personal information but also providing evidence to support the accuracy of the information. It sometimes requires acting under one's real name in order to enjoy the benefits of an established reputation, while at other times it benefits from anonymity to avoid degrading that reputation or to avoid repercussions for expressing unpopular views. At times, people may want to link aspects of multiple personas; at other times the same people may wish to keep identities distinct.

The management of identity takes place under rules negotiated socially and legally. The details of these rules deeply influence individuals' autonomy and the quality of their lives. Today's rules for management of identity have not kept pace with changes in identity technology, surveillance and data mining technology, and electronic crime, and will need to be updated to deal with the new challenges.

The protection of identity requires both legal and technical mechanisms, some of which are not yet in place.

If the legal and technical challenges associated with identity are not successfully addressed, identity-related fraud will escalate, privacy will erode, the quality of social interactions both online and in the physical world will deteriorate, and social trust in government and technology will wane. And solving identity challenges is a daunting task.

A successful legal and technical infrastructure for identity cannot be centred on any single party; it must balance the interests of individuals, businesses, governments, and whole societies. The infrastructure must live comfortably within local, national, and international laws and treaties; it must also sit comfortably within the fabrics of the individual human communities in which stakeholders and technologists will determine what identity management, data protection, and privacy-enhancing technologies are feasible and cost-effective. Governments will have to decide what data protection measures for IDM are just and socially beneficial. And individuals will have to decide how they will use and protect their identities in an increasingly complicated, connected, and public world.

Due to the global footprint of present-day social, commercial, and political life, there is a need for policies and IDM infrastructure that can span national boundaries and at the same time encourage user control.

The future identity infrastructure will not be simple, and attempts to oversimplify it, particularly in pursuit of easy implementation of technical components, should be resisted. This paper has presented a set of decisions which will have to be made in designing both technical infrastructure and policy for identity. It has also provided a set of inherent Properties of Identity – which any successful identity infrastructure

will have to acknowledge and accommodate. And finally, this paper has provided a set of constraints which limit the range of approaches that can successfully be pursued.

Indeed, the Properties of Identity can serve as a compass for policy makers and technologists. Because identity is social, policy makers will realise that identity is within their area of expertise, and will not defer unduly to technologists when making decisions about identity policy. Because identity is subjective, consequential and valuable, policy makers will see a need to balance their constituents' interests. Because identity is equivocal, policy makers will resist suggestions that identity is a "problem" which has a perfect technical "solution" – identity disputes are social disputes, not technical glitches, and the processes policy makers will design to resolve these disputes will be social and legal processes and will have the usual provisions for fairness and due process. And because identity is contextual, policy makers will recognise a need to identify all the important use cases before designing laws, policies, and systems for handling identity.

Because identity is dynamic, technologists will build identity systems which allow for correction, revision, and verification of identity information. Because identity is subjective, technologists will build identity systems which can handle disagreements about identity – in some cases resolving the disagreement and in other cases recording it without a resolution. Because identity is equivocal, technologists will build probabilistic identification systems rather than deterministic ones and will not make overly strong claims about the level of certainty of identification which a system is able to achieve. Because identity is consequential, technologists will build systems with appropriately robust protections for security and privacy. And because identity is contextual, technologists will build systems which enable individuals to maintain different identities and choose which identity to use in each context they encounter.

The paper has argued that a free and open social order requires trust, and trust hinges on accountability; accountability, in turn, hinges on user control; user control, in turn, hinges on data protection; and data protection hinges on respect for Properties of Identity in law and technology. The Properties of Identity, if followed, will allow for robust personhood and digital identity going forward.

Given the importance of these issues for the future information society, more investigation is needed into where gaps exist in current international data protection arrangements and how these will relate to an identity infrastructure.

ANNEX: OECD *PRIVACY GUIDELINES* (EXCERPT)

OECD's 1980 Guidelines on the Protection of Privacy and Transborder Flows of Personal Data Part Two – Basic Principles of National Application (*complete excerpt*)

Collection Limitation Principle

7. There should be limits to the collection of personal data and any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject.

Data Quality Principle

8. Personal data should be relevant to the purposes for which they are to be used, and, to the extent necessary for those purposes, should be accurate, complete and kept up-to-date.

Purpose Specification Principle

9. The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfilment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose.

Use Limitation Principle

10. Personal data should not be disclosed, made available or otherwise used for purposes other than those specified in accordance with Paragraph 9 except:

- a) with the consent of the data subject; or
- b) by the authority of law.

Security Safeguards Principle

11. Personal data should be protected by reasonable security safeguards against such risks as loss or unauthorised access, destruction, use, modification or disclosure of data.

Openness Principle

12. There should be a general policy of openness about developments, practices and policies with respect to personal data. Means should be readily available of establishing the existence and nature of personal data, and the main purposes of their use, as well as the identity and usual residence of the data controller.

Individual Participation Principle

13. An individual should have the right:

- a) to obtain from a data controller, or otherwise, confirmation of whether or not the data controller has data relating to him;
- b) to have communicated to him, data relating to him
 - o within a reasonable time;
 - o at a charge, if any, that is not excessive;
 - o in a reasonable manner; and
 - o in a form that is readily intelligible to him;
- c) to be given reasons if a request made under subparagraphs(a) and (b) is denied, and to be able to challenge such denial; and
- d) to challenge data relating to him and, if the challenge is successful to have the data erased, rectified, completed or amended.

Accountability Principle

14. A data controller should be accountable for complying with measures which give effect to the principles stated above.