

# Laws of Identity

Authored by Kim Cameron<sup>1</sup>

Chief Identity Architect, Microsoft Corporation

---

The Internet was built without a way to know who and what you are connecting to. This limits what we can do with it and exposes us to growing dangers. If we do nothing, we will face rapidly proliferating episodes of theft and deception which will cumulatively erode public trust in the Internet.

This paper is about how we can prevent that loss of trust and go forward to give Internet users a deep sense of safety, privacy and certainty about who they are relating to in cyberspace. Nothing could be more essential if new Web-based services and applications are to continue to move beyond "cyber publication" and encompass all kinds of interaction and services. Our approach has been to develop a formal understanding of the dynamics causing digital identity systems to succeed or fail in various contexts, expressed as the Laws of Identity. Taken together, these laws define a unifying identity metasystem that can offer the Internet the identity layer it so obviously requires.

---

<sup>1</sup> The ideas presented here were extensively refined through the Blogosphere in a wide-ranging conversation documented at [www.identityblog.com](http://www.identityblog.com) that crossed many of the conventional fault-lines of the computer industry, and in various private communications. In particular I would like to thank Arun Nanda, Andre Durand, Bill Barnes, Carl Ellison, Caspar Bowden, Craig Burton, Dan Blum, Dave Kearns, Dave Winer, Dick Hardt, Doc Searls, Drummond Reed, Ellen McDermott, Eric Norlin, Esther Dyson, Fen Labalme, Identity Woman Kaliya, JC Cannon, James Kobielus, James Governor, Jamie Lewis, John Shewchuk, Luke Razzell, Marc Canter, Mark Wahl, Martin Taylor, Mike Jones, Phil Becker, Radovan Janocsek, Ravi Pandya, Robert Scoble, Scott C. Lemon, Simon Davies, Stefan Brands, Stuart Kwan and William Heath

# The Laws of Identity

We can now look at the seven essential laws that explain the successes and failures of digital identity systems.

## 1

### **User Control and Consent**

*Technical identity systems must only reveal information identifying a user with the user's consent.*

No one is as pivotal to the success of the identity metasystem as the individual who uses it. The system must first of all appeal by means of convenience and simplicity. But to endure, it must earn the user's trust above all.

Earning this trust requires a holistic commitment. The system must be designed to put the user in control - of what digital identities are used, and what information is released. The system must also protect the user against deception, verifying the identity of any parties who ask for information. Should the user decide to supply identity information, there must be no doubt that it goes to the right place. And the system needs mechanisms to make the user aware of the purposes for which any information is being collected.

The system must inform the user when he or she has selected an identity provider able to track internet behavior.

Further, it must reinforce the sense that the user is in control regardless of context, rather than arbitrarily altering its contract with the user. This means being able to support user consent in enterprise as well as consumer environments. It is essential to retain the paradigm of consent even when refusal might break a company's conditions of employment. This serves both to inform the employee and indemnify the employer.

The Law of User Control and Consent allows for the use of mechanisms whereby the metasystem remembers user decisions, and users may opt to have them applied automatically on subsequent occasions.

## 2

### **Minimal Disclosure for a Constrained Use**

*The solution which discloses the least amount of identifying information and best limits its use is the most stable long term solution.*

We should build systems that employ identifying information on the basis that a breach is always possible. Such a breach represents a risk. To mitigate risk, it is best to acquire information only on a "need to know" basis, and to retain it only on a "need to retain" basis. By following these practices, we can ensure the least possible damage in the event of a breach.

At the same time, the value of identifying information decreases as the amount decreases. A system built with the principles of information minimalism is therefore a less attractive target for identity theft, reducing risk even further.

By limiting use to an explicit scenario (in conjunction with the use policy described in the law of control), the effectiveness of the “need to know” principle in reducing risk is further magnified. There is no longer the possibility of collecting and keeping information “just in case” it might one day be required.

The concept of “least identifying information” should be taken as meaning not only the fewest number of claims, but the information least likely to identify a given individual across multiple contexts. For example, if a scenario requires proof of being a certain age, then it is better to acquire and store the age category rather than the birth date. Date of birth is more likely, in association with other claims, to uniquely identify a subject, and so represents “more identifying information” which should be avoided if it is not needed.

In the same way, unique identifiers that can be reused in other contexts (for example drivers’ license numbers, social security numbers and the like) represent “more identifying information” than unique special-purpose identifiers that do not cross context. In this sense, acquiring and storing a social security number represents a much greater risk than assigning a randomly generated student or employee number.

Numerous identity catastrophes have occurred where this law has been broken. We can also express the Law of Minimal Disclosure this way: aggregation of identifying information also aggregates risk. To minimize risk, minimize aggregation.

### **3**

#### **Justifiable Parties**

*Digital identity systems must be designed so the disclosure of identifying information is limited to parties having a necessary and justifiable place in a given identity relationship.*

The identity system must make its user aware of the party or parties with whom she is interacting while sharing information.

The justification requirements apply both to the subject who is disclosing information and the relying party who depends on it. Our experience with Microsoft’s Passport is instructive in this regard. Internet users saw Passport as a convenient way to gain access to MSN sites, and those sites were happy using Passport – to the tune of over a billion interactions per day. However, it did not make sense to most non-MSN sites for Microsoft to be involved in their customer relationships. Nor were users clamoring for a single Microsoft identity service to be aware of all their Internet activities. As a result, Passport failed in its mission of being an identity system for the Internet.

We will see many more examples of this law going forward. Today some governments are thinking of operating digital identity services. It makes sense (and is clearly justifiable) for people to use government-issued identities when doing business with the government. But it will be a cultural matter whether, for example, citizens agree it is “necessary and justifiable” for government identities to be used in controlling access to a family wiki – or connecting a consumer to her hobby or vice.

The same issues will confront intermediaries building a trust fabric. The law is not intended to suggest limitations of what is possible, but rather to outline the dynamics of which we must be aware.

We know from the law of control and consent that the system must be predictable and "translucent" in order to earn trust. But the user needs to understand *who she is dealing with* for other reasons, as we will see in law six (human integration). In the physical world we are able to judge a situation and decide what we want to disclose about ourselves. This has its analogy in digital justifiable parties.

Every party to disclosure must provide the disclosing party with a policy statement about information use. This policy should govern what happens to disclosed information. One can view this policy as defining "delegated rights" issued by the disclosing party. Any use policy would allow all parties to cooperate with authorities in the case of criminal investigations. But this does not mean the state is party to the identity relationship. Of course, this should be made explicit in the policy under which information is shared.

## 4

### **Directed Identity**

*A universal identity system must support both "omni-directional" identifiers for use by public entities and "unidirectional" identifiers for use by private entities, thus facilitating discovery while preventing unnecessary release of correlation handles.*

Technical identity is always asserted with respect to some other identity or set of identities. To make an analogy with the physical world, we can say identity has direction, not just magnitude. One special "set of identities" is that of *all other identities* (the public). Other important sets exist (for example, the identities in an enterprise, some arbitrary domain, or in a peer group).

Entities that are public can have identifiers that are invariant and well-known. These public identifiers can be thought of as beacons – emitting identity to anyone who shows up. And beacons are "omni directional" (they are willing to reveal their existence to the *set of all other identities*).

A corporate web site with a well-known URL and public key certificate is a good example of such a public entity. There is no advantage - in fact there is a great disadvantage - in changing a public URL. It is fine for every visitor to the site to examine the public key certificate. It is equally acceptable for everyone to know the site is there: its existence is public.

A second example of such a public entity is a publicly visible device like a video projector. The device sits in a conference room in an enterprise. Visitors to the conference room can see the projector and it offers digital services by advertising itself to those who come near it. In the thinking outlined here, it has an omni-directional identity.

On the other hand, a consumer visiting a corporate web site is able to use the identity beacon of that site to decide whether she wants to establish a relationship with it. Her system can then set up a "unidirectional" identity relation with the site by selecting an identifier for use with that site and no other. A unidirectional identity relation with a different site would involve fabricating a completely unrelated identifier. Because of this, there is *no correlation handle emitted* that can be shared between sites to assemble profile activities and preferences into super-dossiers.

When a computer user enters a conference room equipped with the projector described above, its omni-directional identity beacon could be utilized to decide (as per the law of control) whether she wants to interact with it. If she does, a short-lived unidirectional identity relation could be established between the computer and the projector - providing a secure connection while divulging the least possible identifying information in accordance with the law of minimal disclosure.

Bluetooth and other wireless technologies have not so far conformed to the fourth law. They use public beacons for private entities. This explains the consumer backlash innovators in these areas are currently wrestling with.

Public key certificates have the same problem when used to identify individuals in contexts where privacy is an issue. It may be more than coincidental that certificates have so far been widely used when in conformance with this law (i.e. in identifying public web sites) and generally ignored when it comes to identifying private individuals.

Another example involves the proposed usage of RFID technology in passports and student tracking applications. RFID devices currently emit an omni-directional public beacon. This is not appropriate for use by private individuals.

Passport readers are public devices and therefore should employ an omni-directional beacon. But passports should only respond to trusted readers. They should not be emitting signals to any eavesdropper which identify their bearers and peg them as nationals of a given country. Examples have been given of unmanned devices which could be detonated by these beacons. In California we are already seeing the first legislative measures being taken to correct abuse of identity directionality. It shows a failure of vision among technologists that legislators understand these issues before we do.

## 5

### **Pluralism of Operators and Technologies:**

*A universal identity system must channel and enable the inter-working of multiple identity technologies run by multiple identity providers.*

It would be nice if there were one way to express identity. But the numerous contexts in which identity is required won't allow it.

One reason there will never be a single, centralized monolithic system (the opposite of a metasystem) is because the characteristics that would make any system ideal in one context will disqualify it in another.

It makes sense to employ a government issued digital identity when interacting with government services (a single overall identity neither implies nor prevents correlation of identifiers between individual government departments).

But in many cultures, employers and employees would not feel comfortable using government identifiers to log in at work. A government identifier might be used to convey taxation information; it might even be required when a person is first offered employment. But the context of employment is sufficiently autonomous that it warrants its own identity, free from daily observation via a government-run technology.

Customers and individuals browsing the web meanwhile will in many cases want higher levels of privacy than is likely to be provided by any employer.

So when it comes to digital identity, it is not only a matter of having identity providers run by *different parties* (including individuals themselves), but of having identity systems that offer *different (and potentially contradictory) features*.

A universal system must embrace differentiation, while recognizing that each of us is simultaneously - in different contexts - a citizen, an employee, a customer, a virtual persona. This demonstrates, from yet another angle, that different identity systems must exist in a *metasystem*. It implies we need a simple encapsulating protocol (a way of agreeing on and transporting things). We also need a way to surface information through a unified user experience that allows individuals and organizations to select appropriate identity providers and features as they go about their daily activities.

The universal identity metasystem must not be another monolith. It must be polycentric (federation implies this) and also polymorphic (existing in different forms). This will allow the *identity ecology* to emerge, evolve and self-organize.

Systems like RSS and HTML are powerful because they vehicle any content. We need to see that identity itself will have several - perhaps many - contents, and yet can be expressed in a metasystem.

## 6

### **Human Integration:**

*The universal identity metasystem must define the human user to be a component of the distributed system integrated through unambiguous human-machine communication mechanisms offering protection against identity attacks.*

We have done a pretty good job of securing the channel between web servers and browsers through the use of cryptography – a channel that might extend for thousands of miles. But we have failed to adequately protect the two or three foot channel between the browser's display and the brain of the human who uses it. This immeasurably shorter channel is the one under attack from phishers and pharmers.

No wonder. What identities is the user dealing with as she navigates the web? How understandably is identity information conveyed to her? Do our digital identity systems interface with users in ways that objective studies have shown to work? Identity information currently takes the form of certificates. Do studies show certificates are meaningful to users?

What exactly are we doing? Whatever it is, we've got to do it better: the identity system must extend to and integrate the human user.

Carl Ellison and his colleagues have coined the term 'ceremony' to describe interactions that span a mixed network of human and cybernetic system components – the full channel from web server to human brain. A ceremony goes beyond cyber protocols to ensure the integrity of communication with the user.

This concept calls for profoundly changing the user's experience so it becomes predictable and unambiguous enough to allow for informed decisions.

Since the identity system has to work on all platforms, it must be safe on all platforms. The properties that lead to its safety can't be based on obscurity or the fact that the underlying platform or software is unknown or has a small adoption.

One example is United Airlines' Channel 9. It carries a live conversation between the cockpit of one's plane and air traffic control. The conversation on this channel is very important, technical and focused. Participants don't "chat" - all parties know precisely what to expect from the tower and the airplane. As a result, even though there is a lot of radio noise and static, it is easy for the pilot and controller to pick out the exact content of the communication. When things go wrong, the broken predictability of the channel marks the urgency of the situation and draws upon every human faculty to understand and respond to the danger. The limited semiotics of the channel mean there is very high reliability in communications.

We require the same kind of bounded and highly predictable ceremony for the exchange of identity information. A ceremony is not a "whatever feels good" sort of thing. It is predetermined.

But isn't this limitation of possibilities at odds with our ideas about computing? Haven't many advances in computing come about through ambiguity and unintended consequences which would be ruled out in the austere light of ceremony?

These are valid questions. But we definitely don't want unintended consequences when figuring out who we are talking to or what personal identification information to reveal. The question is how to achieve *very high levels of reliability* in the communication between the system and its human users. In large part, this can be measured objectively through user testing.

# 7

## Consistent Experience Across Contexts

The unifying identity metasystem must guarantee its users a simple, consistent experience while enabling separation of contexts through multiple operators and technologies. Let's project ourselves into a future where we have a number of contextual identity choices. For example:

- browsing: a self-asserted identity for exploring the web (giving away no real data)
- personal: a self-asserted identity for sites with which I want an ongoing but private relationship (including my name and a long-term email address)
- community: a public identity for collaborating with others
- professional: a public identity for collaborating issued by my employer
- credit card: an identity issued by my financial institution
- citizen: an identity issued by my government

We can expect that different individuals will have different combinations of these digital identities, as well as others.

To make this possible, we must "thingify" digital identities – make them into "things" the user can see on the desktop, add and delete, select and share. How usable would today's computers be had we not invented icons and lists that consistently represent folders and documents? We must do the same with digital identities.

What type of digital identity is acceptable in a given context? The properties of potential candidates will be specified by the web service from which a user wants to obtain a service. Matching thingified digital identities can then be displayed to the user, who can select between them and use them to understand what information is being requested. This allows the user to control what is released.

Different relying parties will require different kinds of digital identities. And two things are clear:

- A single relying party will often want to accept more than one kind of identity; and
- A user will want to understand his or her options and select the best identity for the context

Putting all the laws together, we can see that the request, selection, and proffering of identity information must be done such that the channel between the parties is safe. The user experience must also prevent ambiguity in the user's consent, and understanding of the parties involved and their proposed uses. These options need to be consistent and clear. Consistency across contexts is required for this to be done in a way that communicates unambiguously with the human system components.

As users, we need to see our various identities as part of an integrated world which none the less respects our need for independent contexts.

# Conclusion

Those of us who work on or with identity systems need to obey the *Laws of Identity*. Otherwise, we create a wake of reinforcing side-effects that eventually undermine all resulting technology. The result is similar to what would happen if civil engineers were to flaunt the law of gravity. By following them we can build a unifying identity metasystem that is universally accepted and enduring.

---