Microsoft

Trustworthy Computing

End to End Trust Update

When we proposed a vision for End to End Trust at the RSA Conference 2008 in April, we committed to a number of actions. We began an internal effort to identify scenarios that would realize the End to End Trust vision in a concrete form, and we began to discuss the vision and its realization with governments, industry partners, customers, industry analysts, and public interest groups. We also solicited feedback on our white paper outlining our vision, both as part of our discussions and on the End to End Trust Web site. The purpose of this communication is to share some of the feedback we have received, share our thoughts on that feedback and to provide an update on one scenario we have been working on. Future communications will share details of other scenarios we are exploring, and our plans for demonstrating our concepts to bring us closer to the End to End Trust vision.

General Feedback on End to End Trust

Most of the feedback we have received on the End to End Trust vision dealt with matters of identity, accountability and privacy, and focused on the "people" layer of the trusted stack (as opposed to hardware, software, and data). Overall, the feedback has been positive, with most individuals recognizing the need for enabling more effective trust decisions based on identity attributes (such as age) and creating greater accountability for online activity. That said, a few key areas emerged:

Anonymity and User Control

First, there is concern about how we protect anonymity (and the values that anonymity supports, such as free speech) in a more highly authenticated Internet. Most have addressed this issue by noting the importance of allowing users to control what they disclose and when, a very important privacy principle (i.e., user control). One commentator noted, for example, that "I imagine this won't be perfect for a long time, but the last things I would want to see from these changes are lost privacies, and loss of control. The ultimate control should remain in the end-user's hands." Similarly, another commentator noted that people have the right to "own and control their identity" and "be anonymous while controlling their identity at the same time." As I noted in the original paper, we believe that End to End Trust must be about empowering users and their agents to make good decisions about what to share and whom to trust.

While recognizing the need for user control, several people also noted the inherent tension between offering control and creating a good user experience. More specifically, many said that for a user to make an informed choice, they need to be presented with relevant data at the right time. At the same time, however, presenting the user with too much data and too many options can degrade the user experience. Thus, identifying the right balance continues to be a challenge.

Access

In one particular area, broadly described as "access" by most people, there was greater divergence of opinion. A small minority of individuals, likening the Internet to the public roads system, suggested that access to the Internet be gated and individuals be required to register to gain access, with such registration information accessible to law enforcement pursuant to appropriate process. Others implicitly warned of the dangers of such an approach, noting that "end users will not accept a net that is regimented and controlled and limits the freedom of personal expression." Of course, many countries have long permitted anonymous calls to be made (e.g., through pay phones or disposable cell phones), anonymous mail to be sent (e.g., delivering mail without a return address), and provided legal protections to anonymity, even though there is some risk of harm. This suggests that societies have concluded that the social and economic benefits of anonymous speech outweigh the risks such anonymity presents.

Practicality of Achieving End to End Trust

Quite properly, several people focused on the practical aspects of reaching the End to End Trust vision. Some noted the need for "certified authenticators" who would validate identities. Others noted the unique role that Internet service providers (ISPs) and access providers can play, from enforcing some of the rules (e.g., filtering unsigned mail) to helping ensure greater accountability (e.g., trace communications in the event a civil authority makes a lawful request for assistance). Others wondered (and worried) about the power granted to these organizations to monitor network traffic and establish behavioral profiles.

People also had feedback on the practicality of making progress toward this vision, which centered on two themes. The first theme is that many think the paper did a really good job of "decomposing" the security problem in a way where progress can be made in certain areas, even if progress is not made in all areas. By way of example, progress can be made on application signing even if other parts of the trusted stack fail to take off. The second theme relates to the enormity of the task in light of the inertia that often exists when change is required. Achieving significant success in establishing the broader End to End Trust vision (as opposed to just a piece of that vision) requires many organizations to come together and sing in harmony. Although many have noted that the focus on the integration of technology, society, politics and economics is the right focus, it was clear that the odds of integrating and aligning those four disciplines across multiple stakeholders internationally seems daunting to many.

Of course, while dialogue is very important, people also want to see real action. In that regard, we are, in fact, addressing one of the scenarios discussed in the paper, detailed immediately below.

In-Person Proofing Scenario

In the End to End Trust white paper we released in April, we talked about the concept of in-person-proofing (IPP) as the basis for authenticating identity attributes on the Internet. One potential IPP scenario we briefly discussed in the paper was schools identifying their students and issuing them digital credentials so that they can visit hosted electronic playgrounds where all the other members are children of similar age. With increasing concerns about protecting children online, age verification solutions are being explored by governments, child safety

advocates, academics and technology leaders as one way to maintain a separation between the interactive online experiences of children and those of adults.

Over the past few months, we have made solid progress in building out this idea further and have developed a conceptual approach to this scenario, which can be found at http://www.microsoft.com/endtoendtrust. The approach is based on using digital information card technology (Windows CardSpace is one example, the open source Higgins project is another) as the infrastructure to support the creation of digital identity cards. Children can use these digital ID cards to present verified age claims online, supported by offline in-person proofing events at schools or other appropriate institutions. We have submitted this concept to the Internet Safety Technical Task Force (ISTTF), led by Harvard University's Berkman Center for Internet and Society, as part of our work with other leading technology businesses and organizations to identify effective tools and technologies to create a safer Internet environment for children. We have also submitted this proposal to the European Commission, which has launched a public consultation on age verification. Both groups will evaluate and report on proposed solutions in the next few months. We commend both Harvard's Berkman Center and the European Commission for taking the lead on exploring age verification solutions.

The framework we have outlined is a largely technical solution to the age verification challenge, but we believe that the nontechnical aspects of the problem will be as difficult to solve as the technical ones, if not more so. To make this vision a reality, child development experts, governments and industry must work together to address a variety of additional challenges and questions.

Beyond the use for age verification specifically, if IPP is going to become the foundation for authenticating identity attributes on the Internet, there are many questions that need to be answered, many of which we heard in the feedback on the End to End Trust paper. For example:

- Can we leverage current IPP systems and capabilities, such as post offices that proof for passports, schools that proof for attendance, motor vehicle departments that issue identity cards and drivers' licenses, and banks that proof to know their customers?
- What type of organizations could and should incubate this effort and provide proof of concept implementations?
- What organizations are logical places to first "consume" IPP-based identity attributes? Is it the private sector, where social-networking sites may be looking for effective ways to protect children? Is it in the public sector where robust identity is and should be required for particular governmental interactions (e.g., claiming monetary benefits from the government)? Or both?
- How do regulatory and economic risks affect the willingness of organizations to enter the IPP and identity authentication market? Is it necessary to construct a more favorable regulatory environment or provide other safeguards (e.g., limiting consumer fraud losses to $50 helped promote the safety of credit cards)?
- How do we address failures, such as the inappropriate granting of, or the loss of, an identification credential?

We are hopeful that initiatives like those of the Berkman Center and the European Commission will help to foster these important discussions.

Next Steps to Make End to End Trust a Reality

The challenges with aligning social, economic and political requirements with IT capabilities are critical and real, but as discussed in my paper, we are optimistic that it is possible to address these challenges collectively. In light of the paper and the feedback on it, we are thinking more about (1) the appropriate ways to engage all of these disparate interests, (2) how to more clearly identify which problems need to be solved in which spheres, and (3) how to catalyze the right actions by the right parties at the right times. We will also factor this feedback in as we continue to build out scenarios for making End to End Trust a reality in the future.

In sum, there was both support for and concerns expressed about the concepts in the End to End Trust paper: this is the hallmark of a healthy debate about the future of the Internet. Thank you all for your comments and focus on these issues. For those of you who haven't yet read the paper, or joined the dialogue via our forums or otherwise, I encourage you to do so by visiting http://www.microsoft.com/endtoendtrust. Over the next few months we will be updating you on the results of this continuing dialogue and our next steps.

Microsoft Corp. • One Microsoft Way • Redmond, WA 98052-6399 • USA