

>> Christopher Olsen: All right. If everyone will get settled, we're going to start now. I'd like to welcome everyone to the third and final roundtable in our series exploring privacy. It's great to see that we've carried the momentum over from roundtable one through our Berkeley event and now to our final event here in D.C. I need to make a few housekeeping announcements. The first and perhaps most important to at least one individual is we located an iPhone charging in a wall outlet, And it's available at the registration desk up front. There are food and beverages available out in the hallway. There's also a list of other eateries available at the registration desk. Restrooms are located through the lobby. Don't go through the security stands. Go around past the elevators. There's a Wi-Fi code for you to use to get broadband. The code is CABA 010808. There's always a brochure outside that has that code. Anyone who goes out of the building without an FTC badge will have to come back through security, so make sure you build in some time for that. This is perhaps the most exciting announcement. In the case of an emergency, you'll have to evacuate the building, and you'll go outside the building, New Jersey Avenue is just in front. Across the street is Georgetown law school. You go to the right front sidewalk at Georgetown law school. We actually have a rallying point in case of evacuation. We'll have questions today. We'll have people in the audience with question cards. So if you have a question during the event, please raise your hand. Someone will come to you with a question card. I think you have cards in your packages, as well. So you can fill out a question, and people will pick them up and deliver them to the moderators. We also -- For the web audience, we're also accepting questions. You can e-mail them to privacyroundtable@ftc.gov. So, that takes care of the logistic announcements. We're very pleased this morning to have Commissioner Pamela Jones Harbour provide opening remarks, all the way from Barcelona. And we're very pleased that we worked the technology out hopefully so that this will be a seamless process. So, Commissioner Harbour, welcome.

>> Pamela Jones Harbour: Hello. Welcome, Chris. Good morning. And welcome to the third FTC exploring privacy roundtable. Thank you very much, Chris, for your introduction. And let me personally thank all of the talented FTC staff who have worked tirelessly this past year to make these events happen. You heard where I am. Yes, I am in Barcelona, Spain, coming to you by video. A few hours ago, I delivered one of the keynote speeches to the Secure Cloud Alliance 2010 event. But I certainly didn't want to pass up the opportunity to deliver remarks today at the third and final privacy roundtable. And when I spoke back in December, I mentioned that I soon would

be leaving the Commission. This time, I am really serious. I recently announced that I will depart on April 6th, and this will be my final speech, albeit 3,500 miles away. And for the last time, I note that my remarks today are my own and not necessarily those of the Federal Trade Commission or any individual commissioner. I've said it many times before, and I will say it again today.

Protecting consumer privacy is of utmost important. It must be a driving force for businesses in all stages of product and service development. Unfortunately, many of the companies that consumers look to as leaders and that we expect to be leaders still have not taken this message entirely to heart. First, I want to challenge what I see as a dangerous precedent being set by some of the biggest and most influential technology companies when they publicly expose consumer data. And second, I want to challenge companies that are not adequately protecting consumers through SSL technology. At the last roundtable in Berkeley, I discussed the comments of a technology executive who claimed that privacy expectations and norms are changing. More recently, since the Berkeley event, the press has recycled the comments of another prominent tech executive who stated, "If you have something that you don't want someone to know, maybe you shouldn't be doing it in the first place." Speaking for the last time as a regulator, let me be very clear. I could not disagree more with that assertion. Privacy is a fundamental right that people do care about, and I believe that the Commission and my fellow commissioners would share this opinion. The Commission will continue to view privacy as an important value, as reflected in the forms and expectations of consumers until it is proven that consumers feel otherwise about their privacy. The Commission will continue to evaluate consumers' preferences, and armed with these insights, I hope and expect that the Commission will continue to shape the conversation about the intrinsic value of privacy. But make no mistake, the Commission will unfailingly step in to protect consumers where we believe the law has been violated, and that includes violations relating to privacy promises. And I'm going to be even more specific in my admonition to provide some concrete examples for today's discussion. The recent launch of Google Buzz was, quite frankly, irresponsible conduct by a company like Google. I would use that same word to describe the prior rollout of Facebook's new privacy settings, as well as the November 2007 release of Facebook Beacon. But for now, I will focus on the Buzz example. Google is one of the greatest technology leaders of our time. Google consistently tells the public to "just trust us" and have adopted a company motto, "Do No Evil." We have high expectations for Google as a corporate citizen. But for me, based on my observations, I do not believe that privacy -- consumer privacy -- played any

significant role in the release of Buzz. In the rush, perhaps, to compete with Facebook, Foursquare, Twitter, FriendFeed, Loot, and a host of other companies, it appears that Google did not think through the privacy implications of this launch. New technology, such as Buzz, like some of the updated features offered on Facebook, represent a laudable effort to help consumers integrate and make sense of the daily overload of information that bombards them via e-mail, photos, blogs, tweets, news feeds, and the like. And today, consumers tend to have separate online accounts for a variety of services, and often, they maintain multiple profiles to separate their personal and professional uses. Plus, many companies do one thing very well, and accordingly, consumers are then willing to enter relationships with multiple firms. A common characteristic of the most successful Web 2.0 companies is that they thrive on the network effect. That is to say, the greater the number of users or number of inputs, the better the experience, which further enhances the trend towards interacting with multiple data sources. When Buzz was launched, Google described its function as finding relevance in the noise. It is no wonder that seeking to capitalize on network effect, Google decided to build its service by turning to its installed base of approximately 150 million Gmail users. Unfortunately, to my knowledge, none of those users were consulted before Google unilaterally decided how best to use their data. When users created Gmail accounts, they signed up for e-mail services. That is their primary use of Gmail. Several years ago, when Google first introduced Talk, many users were taken aback that their e-mail address book contacts were automatically suggested as Talk contacts. Publicly, there was a backlash, and Google rolled back the Talk offerings. But the company apparently failed to learn from that prior mistake. Buzz was designed as a social network for users, but the net was cast too widely. News reports indicate that the company claims to have tested Buzz extensively with thousands of employees. The problem is, Google employees are in no way representative of the Gmail user base, a combination of young, old, tech-savvy, novice, and so on. The Buzz product manager admitted as much, saying that, "Getting feedback from 20,000 Googlers does not equal Gmail users in the wild." So think about it. When Gmail first emerged, social networking was barely even a reality. When consumers, especially early adopters, created their Gmail accounts, their expectations did not include social networking. In my view, therefore, a reasonable consumer would consider the initial opt-in of Buzz to be a material change in her relationship with Google. Consumers, not companies, should exercise the ultimate decision on whether they want to sign up for these new features that might expose additional data. I am especially concerned that technology companies are learning harmful

lessons from each other's attempts to push the privacy envelope. Of course, providing new features to users and making the user experience more enjoyable are excellent goals. These efforts may win new users while also building additional loyalty in the existing user base, but even the most respected and popular online companies, the ones who claim to respect privacy, continue to launch products where their guiding privacy principle appears to be "throw it up against the wall and see if it sticks, and if not, we can always pull it back." Deeds speak louder than words. And this is turning into a dangerous game of copycat behavior. And unlike a lot of tech products, consumer privacy cannot be run in beta. Once data is shared, control is lost forever. In the extreme, it is only a matter of time before one might imagine the introduction of new features that incorporate, for instance, genomic information or data from public-health records. The privacy stakes will only get higher. And I realize that perhaps companies continue to take a "testing the water" approach to privacy because no regulatory agency has sent a clear message that this behavior is unacceptable. In my opinion, that message may need to change, and I would like to see the Commission take the position of intolerance towards companies that push the privacy envelope, then backtrack and modify their offerings after facing consumer and regulator backlash. In the meantime, however, companies should exercise greater responsibility and be more circumspect before launching game-changing products. Computer algorithms should not be trusted to interpret consumers' privacy expectations. Consumers still have an expectation of privacy. These norms do not change and cannot be assumed away every time a company wants to compete in a new market. We cannot expect a new paradigm where products and services do not offer user choice, materially changing the bargains consumers understood when they first established the relationship. Now, I don't want to be accused of harping only on Google, so let me turn to my second admonition, which is targeted at a large number of prominent firms and which addresses an important issue of data security. I worry that many consumer-facing computing services have significant data-security vulnerabilities, especially services offered in what we call "the cloud." Encryption technology is already built into every popular web browser, but here is an unpleasant truth. Many popular services employ encryption technology and only transmit initial log-in information, such as user names and passwords. All subsequent data is sent in the clear, unencrypted. This problem affects services such as Microsoft Hotmail, Yahoo! Mail, Flickr, Facebook, and MySpace. This practice exposes consumers to significant risks when they connect to popular cloud-based services using public wireless networks in coffee shops, airports, and other public hot spots. Without encryption, user

data is easily intercepted using freely available, off-the-rack hacking tools. And I spoke last fall at the International Conference of Data and Privacy Protection Commissioners in Madrid, and one of the most memorable speakers was a white hat, or ethical hacker, for those who aren't familiar with the term. And during his presentation, this hacker -- ethical hacker -- demonstrated how he easily could break in to a netbook computer in a matter of mere minutes. It was very sobering, indeed. Many users of cloud-computing services lack the basic security protections that users of traditional PC-based software often take for granted. These vulnerabilities are easily preventable. Many web-based services, including online banking and certain online merchants, operate securely over wireless networks. As a notable example, many banks in the financial sector use the industry standard secure socket layer, SSL, encryption protocol to protect their customers' information. These encryption technologies are widely available, yet many service providers choose not to implement these technologies for all data transfers and instead continue to provide a product and services with unsafe default settings. Even though the service providers know about the vulnerabilities and the ease with which they can be exploited, the firms continue to send private customer information over unsecure Internet connection that easily could have been secured. And so, my bottom line is simple. Security needs to be a default in the cloud. Today, I challenge all of the companies that are not yet using SSL by default -- that includes all e-mail providers, all social-networking sites, and any website that transmits consumer data -- step up and protect consumers. Don't do it just some of the time. Make your website secure by default. Let me end by saying that I've been speaking publicly and have been very outspoken on privacy and data-security issues for 6 1/2 years now, and I have continually pushed companies to be leaders on privacy and data security. And I hope my words have resonated with some of you and that commentators and industry representatives will thoughtfully address my concerns. And now that I am leaving the Commission, the voices of two new commissioners will emerge -- Edith Ramirez and Judy Brill. They're both incredibly bright and talented. And I know they will continue to fight on behalf of consumers as I have tried to do all of these years. Let me end by saying it has been my great privilege and pleasure to serve the American public. Thank you. [Applause]

>> Christopher Olsen: Thank you very much, Commissioner Harbour. Now I'd like to welcome bureau director David Vladeck for opening remarks.

>> David Vladeck: Morning. Let me thank -- Let me start out with some thank-yous. First of all, thank you all for coming. We are now down to the hard core, but it's great to see that there's such a good turnout today. Next, I really would like to thank Commissioner Harbour, not just for her thoughtful remarks this morning but for her stalwart leadership within the Commission on privacy matters. We will miss Commissioner Harbour, but we know that her departure from the Federal Trade Commission will not still her voice on privacy matters. I also want to thank our panelists today for sharing their formidable expertise. These roundtables have been greatly enriched by the participation of panelists like the ones today, and we are very grateful for their participation.

Before we get started today, I'd like to highlight four themes that have come up time and time again in the roundtables and end by explaining where we're going with all of this. First, we've discussed extensively the benefits and risks of technology in the privacy context. It's hard to believe that the Netscape browser revolutionized the Internet, opening the way for commercial uses of it just 15 years ago. Since then, geometric increases in the computational capacity and data-transmission speeds and cheaper and cheaper storage of data have had huge implications. These steady innovations have created benefits to consumers thanks in large measure to the flow of information that it makes possible. But these advances have also created new risks for consumers. A few years ago, Tim Berners-Lee cautioned that I.T. professionals must keep in mind that -- and now I'm quoting -- "Data is a precious thing and will last longer than the systems themselves." Well, when data hangs around, odds are it will be useful for some purpose that may not have even been envisioned when the data was collected, and that presents challenges. In addition, the march of technology has blurred and, indeed, threatens to obliterate the distinction between PII and nonpersonal information, especially given the sheer volume of information that is now collected about individuals. Catherine Deneuve, whom I've always admired, once spoke for all of us when she quipped, "I like being famous when it's convenient for me and completely anonymous when it is not." On the web, at least, it is getting harder and harder for individuals to choose anonymity. And technology has enabled companies to surveil people to an unprecedented degree, both online and increasingly offline. Second, we have discussed privacy challenges raised by emerging business models. Business models have changed as quickly as the technology, creating new markets overnight. What did consumers know about cloud computing or even social networking as recently as five years ago? The continual emergence of these new models, too, means that consumers are often presented with unfamiliar or confusing situations, where the nature of the

commercial bargain, in terms of privacy, may not be clear and may be constantly shifting. Not surprisingly, consumers understand little about how their information is handled, whether by companies they share with directly or by companies that work behind the scenes, like data brokers, ad networks, and application providers. Third, although new technologies and business models have raised privacy concerns, they have also been used to innovate to protect privacy. For example, several companies have introduced tools that consumers may use to access the Internet categories they've been placed in and to change how they've been categorized. In nonprofit think tanks, the future of privacy forum, together with marketing communications company WPP, has led an effort to develop and test an icon that would alert consumers how to get more information and how to make choices about how their information is being used for behavioral advertising. This is all to the good. Fourth and finally, there's been little satisfaction with the privacy approaches that have been pursued to date. Privacy policies are not located where consumers can find them. They're too complicated, they're too vague, and too long for consumers to really understand them. While there's widespread agreement that the information processes we use should be transparent, we're still exploring effective ways to disclose what information is being collected and to give consumers a meaningful opportunity to control its use. And, of course, we all know that once information has been shared, there's no way to get the genie back into the bottle. Although we've covered a lot of ground in the first two roundtables, we've left some big questions for today. Our first panel tackles one of the big questions of the Internet. Can we build security and privacy into the Internet after the fact? That is, can we create a secure, authenticated structure on top of the foundation that was built to be trusting and open? Next, we'll tackle health-privacy issues, examining another great puzzle. How do we reconcile individual interest in privacy, particularly about health issues, with society's interest in getting research, epidemiologists, and others the information they need to improve our collective health? Then we'll address the -- Then we'll address the question about sensitive information more broadly. Is there a consensus that particular categories of information are sensitive and deserve heightened protection, or is information about certain kinds of people so sensitive that they should be treated with special care? For instance, information about children. Or is sensitivity simply in the eye of the beholder. Are there policy approaches that would enable people to apply their preferences themselves without the need for some kind of consensus? The final panel will wrap up with a discussion about what we've learned and where we go from here. I expect we'll hear a lot of the same themes and questions come up.

How do we make information practices transparent to consumers, and how do we give consumers appropriate tools to make their preferences known? Also, how do we create incentives for companies to consider privacy before rolling out new business models or new service models? Many people have asked me, "Where do we go from here? Once this roundtable is concluded, what are the FTC's next steps?" Well, I think, to be candid, we're not certain. The first thing we are going to do is we're going to sit back, and we're going to digest everything we've heard. We've made detailed records of the first two panels. We'll do the same with this. We'll need to go back and study them. We will put together our thoughts and recommendations, if any, and we will make those public. We will then solicit your input. We want to be as open and transparent as we can, and we will need your help and your thoughts. So we will have a very public process on this. We've had great, great assistance as we go forward. We look forward to more of that in the future. Before we conclude, I want to say one final word of thanks, this time to the staff that has worked for months to make these roundtables happen. It's really hard to explain just how much time and effort goes into putting these panels together, into doing the research that is discussed at these panels. No one would have ever thought that a president from Chicago would shut down the government because of a little snow. But during D.C.'s recent Snowpocalypse, the entire city was shut down for more than a week. But the roundtable team did not miss a beat. They worked tirelessly through the storm. We're not like the postal service. A little snow, sleet, rain, or 4 feet of snow is not going to stop the FTC. They worked throughout that stretch to put this roundtable together. I greatly appreciate the dedication and care they've shown throughout in making these roundtables a success. So thank you all very much, and thank you for coming. [Applause]

>> Christopher Olsen: Thank you, David. We're going to take a very brief break. I'll ask the panelists for panel one to come up and take your seats. We'll start promptly at 9:15. So if you want to take a couple minutes while panel one gets settled, and then 9:15, we'll begin. Thank you.

>> Loretta Garrison: If folks could please start coming back in and taking your seats, we're about to start. Thanks. Good morning, and welcome, everyone. I'm Loretta Garrison, and this is my co-moderator, Naomi Lefkowitz, and we're going to moderate the first panel for the final roundtable this morning. We're going to open today's final roundtable by stepping back and taking a hard look at the architecture of the Internet. We want to present a challenge to all of those technical folks in

the audience and those of you who are listening in and to our distinguished panelists who have come prepared today with all the answers to the questions we're going to ask them. The panelists today, we're very delighted to introduce to you, are John Clippinger. This is to my immediate left, and we're going all the way down the table. He's the co-director of the Law Lab at Harvard University Berkman Center for Internet & Society. Next to him is Jules Cohen, Director of Trustworthy Computing, from Microsoft. Then Peter Eckersley, who's come all the way in from California. He's a senior staff technologist with the Electronic Frontier Foundation. Next to Peter is Ed Felten, who's the Director for the Center for Information Technology Policy at Princeton University. Next to Ed is Lucy Lynch, Director of Trust and Identity Initiatives from the Internet Society. And then we have Drummond Reed, who's the Executive Director for the Information Card Foundation. And last but definitely not least, Ari Schwartz, who's the vice president and chief operating officer for the Center for Democracy and Technology. I'd like to remind audience members that you can submit questions to the panel by filling out a question card and handing it to FTC staff that will be walking around the room. For those of you watching this panel via the webcast, you can submit your questions by e-mailing them to [privacyroundtable](mailto:privacyroundtable@ftc.gov) -- that's all one word -- at ftc.gov. And to our panelists, if you want to speak at any time, please turn your name tent on end and wait to be recognized. As you know, when the Internet was initially created, it was technically designed to facilitate communications among a number of researchers at various universities around the country and what is now known as DARPA, the Defense Advanced Research Projects Agency. This was a small, known, trusted environment designed strictly to share information as the participants worked on common projects. Since then, we've built on top of that architecture a complex commercial enterprise, a social-networking system, search functionality, none of which was contemplated or even envisioned at the time of the original design. The challenge today to our panelists is to engage in a thought experiment and examine and discuss how you would construct an Internet today to accommodate these various enterprises and what change from that design we can apply to the existing architecture, short of blowing up the Internet. So, Peter, if you can start us off, you can start afresh, how would you design the architecture of the Internet to design all of these activities, to address the privacy and security concerns that we've heard throughout these round table discussions?

>> Peter Eckersley: So, I can't necessarily give you a single answer to that question, but here's the new design. "Let's just go with this instead of the current Internet." But I can say that, if you want to understand the privacy problems we're having on the Internet today, it's helpful to imagine taking a time machine back to the early '90s and looking at where today's Internet came from. And what you would find is that there are a lot of the problems that we're looking at that were essentially side effects or inadvertent design decisions that were made back in the '90s with the intention of just making the web work and making it work better. We could look TCP/IP, which is the basic protocol that most Internet software uses to communicate. And it has this property that the other side can always see the address that you're using to communicate. And then you can look at the web, which is a simple client-server protocol, and say, "Ah, so a web server always sees the addresses of the people who are reading each document." And these sort of things seem inevitable, but perhaps they weren't inevitable, because if you look at the web at the same time people were also using other protocols, like e-mail even and Usenet, where you couldn't necessarily see the other person's address whenever they communicated with you. In fact, back in the '90s, often there was a separation where some sort of federation of machines was talking to each other, and you never got a message with the other person's address traveling the whole way through the chain of communication. And there were special protocols, like finger and ident, that were used to separate policy with respect to privacy from policy with respect to communication. And so I think one way that we could think about re-architecting the web, if we could do things over -- we can't necessarily do that now easily -- would be to say, "Well, does a web server need to see the user's IP address every time they connect, and can we find a different model for that?" There are a bunch of other decisions that were made later in the '90s the way that cookies and JavaScript and other things were added to web browsers that also have serious privacy consequences, and I'm sure other panelists will talk about some of those.

>> Loretta Garrison: John?

>> John Clippinger: I think Vint Cerf was asked that question. Said, "If you had a chance to do it over again, what would you do?" And he said it was missing an authentication layer. And by that - - and this is something that we've been very interested in is, okay, how do you know who you're dealing with, and how do you start to develop -- And I think we're going to have to think about -- I

don't think about blowing up and starting new, but how do you build a new layer, or how do you build something on top that allows you to have a principled way of knowing who you're dealing with and having -- sort of creating a kind of consequence for behavior, how people treat and disclose information? And I think what we've been talking about here with the traditional privacy format has been, how do you sequester information? I think the issue is going to be how you control it and who has access to it and how do you enforce certain contracts or conventions of access through information. You can't sequester it. That point was made earlier. Personally identifying information can be constructed from non-identifying information. I think there's a need to create a new kind of governance regime, a new kind of -- You have to approach it in a systemic way, not just in a piecemeal way. And my view is that it's very important the locus of control really has to be on the user. You have to have a user-centric, interoperable system that allows people to control information about themselves and have a chain of trust that can be traced back to the individual. It's not to say that people are going to make all those decisions, but architecturally, I think that's a critical consideration. So, going forward, I think we have to think about not just little piecemeal-type fixes but a very systemic way of thinking about it that uses a variety of methods all the way from new kind of encryption technologies to contracts, to what kind of business models are to be used. What are the incentives? What kind of incentives do different companies, employers, and identity providers have that are aligned to take a race to the top rather than a race to the bottom? So, my admonition is that we're moving from a technology to a social era. And in doing that, we're making very profound decisions about how people are going to participate and be protected in our society. So we're building new kinds of institutions that have far-reaching implications.

>> Loretta Garrison: Well, one of the critical aspects of any redesign is going to be the usability and, in a sense, the invisibility of the change to the consumer. We're going to talk a lot about this later on, but certainly one of the ideas behind looking at the basic architecture is whether or not something technically can be built in or designed that would change the default so that it would make it easier for consumers to understand what's going on and to make informed decisions.

>> John Clippinger: I couldn't agree more. I think that one of the things we're looking at experimenting with is how people can see their information being used -- who sees whom, who

sees what. And you might have red, yellow, green information, and when green information goes red, what causes that? A piece of information is somewhere where it shouldn't be. It says, "I'm here. What am I doing here?" Have audit trails. I think we have to move away from sort of complex, inscrutable legal agreements to where people can have an intuitive understanding. Expectation of privacy has to be reflected in the experience and certain norms that are adopted and relied upon. And this is new territory. But you're starting to see some very interesting designs. But on top of that, I think you have to have the audit mechanisms. You have to have some kind of independent party holding others accountable for that. We'll talk about that later.

>> Loretta Garrison: Okay. Ed, did you have a comment?

>> Edward Felten: Sure. Certainly it's important to give users some kind of visibility and control over their information -- where it goes, how it's used, and so on. But this is much harder to do in practice than you might expect. Today, users, in principle, have a certain amount of control over things like the privacy settings in their browsers. But in practice, they really don't, because the mechanisms that are used either involve asking the user millions of questions in pop-up boxes that the users quickly learn to click away, or some kind of very detailed browser privacy preferences dialogue that hardly any users even open, let alone understand. The real challenge is how to let ordinary users have effective control and real autonomy in this area without having to invest a huge amount of effort or learn a lot about how the technology works. And I think that requires some -- that's going to require some really clever advances in the basic models of user interaction online. I don't think we can add this on with patches. I think we really need to think, "How does the user interact with the technology on a minute-to-minute basis?" And we need to build the technology where the user is revealing to the technology through the things they already want to do, what they want.

>> Loretta Garrison: Lucy.

>> Lucy Lynch: I want to drop back just a little, to the question about actually re-architecting the Internet and about whether or not you need to blow it up to change it. Because there are two different conversations going on here. There's a conversation about the Internet -- the entire

Internet, the network that communicates -- and there's a conversation here about the web, which is most end-users' experience, and what happens above that layer. And referencing John's comment, authentication needs to be built in actually at that network layer. Users aren't the only ones who communicate on the network. Entities communicate. Machines communicate to one another, and it's essential that they be able to continue to communicate and to identify end nodes. The benefits of the Internet come from the distributed, decentralized, hierarchical model that allows any entity to communicate with another one. You don't want to break that. There are technologies being designed, and some of them are privacy-aware, like Geopriv, that are built in at the network layer. And you need always to think in a distinction above and below the web when you're talking about this. There are privacy concerns below the web, as well.

>> Loretta Garrison: Drummond.

>> Drummond Reed: I wanted to find an example of how you can achieve privacy by design but also how hard it is. Again, I'm here as Executive Director of the Information Card Foundation, and I want to point to information cards as a technology that has -- is about six or seven years worth of work in its development to try to address exactly what you brought up, which is the usability of privacy. A good part of what we're doing right now is educating audiences about, if you want to give end-users a very easy way to authenticate to websites, to share information about themselves or from third parties about themselves while also, at the same time, protecting their privacy, it's a difficult job. I'll give a very specific example. With Information Cards, as an authentication technology, as a way to sign into websites, the end-user experience is simply one of picking a card out of a wallet. There's no typing of user names or passwords. And yet the underlying technology will automatically -- you can pick one card and use to sign in to a hundred different websites. The underlying software will give a different private personal identifier to each of those 100 websites. They will not be able to correlate the information, the login experience. That's because it's carefully designed to do that. The user doesn't have to understand anything about that. It's a simpler experience than login today, but the technology's been designed to ensure that no correlatable identifier is being shared across all those sites. That's an example of the kind of approach that you have to take if you're going to address what John talked about -- privacy at this relationship layer that, I believe, is evolving and that we're going to need to address this issue.

>> Loretta Garrison: Jules.

>> Jules Cohen: I wanted to echo a couple of points that have been made. It's not just about some of the privacy questions. I think that it's up-leveling to this interesting question of, with respect to the hardware that I'm using on the Internet, the software that runs on top of that hardware, and the people that use these other two things, how do I make trust decisions? And those trust decisions are broader than privacy. Should I trust this piece of hardware? Where does it come from? Should I trust this operating system, this application? Where do they come from? Who are the people behind them? And then, should I trust the people that are using them? And sometimes those are privacy decisions and where the data that flows through the system are. And sometimes it's security decisions, and, "What tools does the user have to actually make those trust decisions and what information do they have on hand and what cues in the user experience are they provided with?" are some of the hard questions that we're grappling with.

>> Loretta Garrison: I wanted to ask all of you about some of the basic premises on which the Internet was developed or the web. You have a system of networks. It's peering, which is -- In a sense, it's, in my mind, it's related to the federal highway system, the superhighways. Then you've got the state roads. Then you've got the county roads. Then you've got private roads. Each of these can be built either in an organized way or, in the private sense, if two companies want to share information, they can simply create that network or that connection and do it. Does this basic autonomy of the Internet design create barriers or difficulties to addressing -- in a structural way -- to addressing the privacy and security issues that we have? Peter?

>> Peter Eckersley: I want to say, no, actually. I don't think there's a problem with peering and the way that the Internet is a hollow of little networks that are stitched together in a patchwork quilt. I think that works pretty well, and I think if the protocols that those networks are talking to each other solve your privacy problems at some layer, then that's going to work well. I think the dynamic that is problematic is one where no one really owns the whole privacy problem. In the example I told before, where the web creates this privacy problem by showing the reader's address to the server every time and it didn't have to be that way. E-mail didn't have that property. That

was, like, a low-level consequence of one protocol, and it was a privacy problem that wasn't solved there. And the privacy problem got kind of kicked upstairs. Someone said, "Hey, we won't deal with this in HTTP, but if people want to solve this privacy problem, then they can go and invent a separate proxy protocol to hide their I.P. addresses. And the problem is that, when you kick these things upstairs, suddenly only 10% or 5% or 1% or less of people actually get the solution. And so I think the problem isn't with the way that the networks are stitched together. It's with making sure that someone is designing privacy, and they're answerable to the users, ultimately, when they say, "Hey, why was I tracked by this person?" You need one kind of place that you can go to and, say, fix this protocol until I get the privacy properties that I need from it.

>> Loretta Garrison: Ari, I saw you nodding your head. Did you have a comment?

>> Ari Schwartz: I mean, I strongly agree with what Peter just said. There's -- I think there's a tendency to be concerned with how information is passed back and forth across the Internet because of the way that some of the original design went in and also because of the way that some network operators have been talking about discriminating against certain kinds of content. But if we didn't have this discussion about going in and looking into the content of the packets, then we would have less concern about that information being passed. If we can build a more secure system that respects privacy in the protocols itself, then those concerns are addressed. And it has nothing to do with the -- You have less of the worries of the discrimination of packets. As long as we keep that basic end-to-end principle, we shouldn't have a problem of the way -- of the structure of different kinds of entities and different kinds of peering agreements.

>> Loretta Garrison: John?

>> John Clippinger: One caution, I would say, I think one has to look at the emerging business environment and where information's going to create value. And so, there will be business models based upon aggregating information and making it available. It's sort of the next generation of Google. So there are going to be very strong forces in the market to test the limits of those protocols and to reinterpret. And so I think it's very important not only to have the correct business incentives but have the correct audit mechanisms. Because we're really talking at another level that

has never existed before. And recognizing that there's great wealth and opportunity and things that could happen when you use this information effectively, and you can fight disease, and a whole number of things can be done, so you don't want to sequester it. But at the same time, you want to have a set of rules -- rules of the road -- that are governance principles that are enforced quickly, transparently, and effectively and also grow with the technology. Otherwise, it will get co-opted.

>> Loretta Garrison: Okay, I want to turn to the I.P. address protocol, because as you've discussed, an addressing system is fundamental to sending data packets over the Internet. Drummond, are there technical limits to masking this information to avoid tracking? Or are there other ways to address the tracking issue?

>> Drummond Reed: And it's a very complex question. Another hat that I wear is I'm co-chair of a technical committee at an Internet standards body called OASIS, and that technical protocol is called XRI. It's for a new type of identifier for the Internet. And an easy way to explain where that fits is that, if the plumbing layer that we were talking about between the hardware is using I.P. addresses to communicate, and this next layer of the web is using URLs, you're connecting between browsers and servers for pages, XRIs are designed for this relationship layer. XRIs are really designed to identify people, organizations, concepts and to have communications directly at that layer. Imagine that you can actually have a messaging relationship where you're not communicating necessarily between I.P. addresses or between e-mail addresses but person to person and the communication is actually able to route itself to the right device, depending -- Am I trying to send John a short message about "I'm five minutes late for this meeting," or am I trying to send Ari a PDF file for something? It's a matter of being intelligent about the choice. That kind of communications routing and the associated rules -- for instance, the privacy or security -- that can be apply to the message, that's the kind of thing that that layer can address. It's one approach. I believe there are issues. The transition from IPv4 to IPv6 has introduced both new capabilities and new vulnerabilities at the layer of I.P. addressing. URLs have their own set of issues. We're trying to address some of those at the XRI layer. It's one way that we can help address those things.

>> Loretta Garrison: Yeah, the IPv6 issue, we had discussed among us. And although there had been some doubt that this would create any issues with I.P. addresses being collected and linked to

information, there was a recent article about a company, ClearSight Interactive, which has acquired something like 100 million I.P. addresses, and of those were actually able to link e-mail address, postal addresses, names, and other registration information to actual individuals, and they're going to initiate targeting based on that. And, of course, with IPv6, you're going to have static I.P. addresses increasingly assigned to individual PDAs. So you will have this direct linkage. This is actually exacerbating the problem of linking all of these bits of data, this sticky data, to individuals. John or Peter, Drummond, any of you want to take that up? Drummond? Or Lucy. Okay.

>> Lucy Lynch: I think there are a number of problems involved in that convergence that you're talking about, and it's exacerbated, actually, by introducing identity-management technologies because there's a set of passive data that's collected that are the system identifiers. They give you one profile. But as people volunteer personal information in conjunction with that data, that's where -- That's where that identifiability comes. Because with a few exceptions, like WHOIS data, which is directly tied to the ownership of an I.P. address, the way they're building that conjunction is by taking volunteered data and system data and conjoining them. That is not a problem with the design of the network. That's a problem with understanding what data you volunteer and how it gets used in conjunction with the other data that's available. So there's a user-education issue, and there's a compliance issue there. You need to gain consent in order to use that data in that way.

>> Loretta Garrison: But, Lucy, if you go back to the earlier point that the I.P. addresses were not necessarily intended to be identified but now they are and now they're linked, is that also not a structural problem, a design problem, as well?

>> Lucy Lynch: No.

>> Loretta Garrison: Okay.

>> Lucy Lynch: No. Users want service. You need to be able to deliver to their end node. Trust me -- users want service. Whether or not they should be exposed because they get service is not the problem, but you need that identifier in order to deliver service.

>> Loretta Garrison: Ed.

>> Edward Felten: Sure. Talking about I.P. addresses as a way of tracking or linking activity really, I think, puts the focus on part of the problem around tracking and linking, which is that there are so many different technological ways that sites or different parties can track and link what people are doing. And if I, as a user, want to avoid being tracked or linked, I need to have a strategy for dealing with all of those different tracking methods. There's a very large perimeter that I have to defend, technologically, to maintain my anonymity when I want to. And if we're going to make progress to give users more control, we have to reduce the size of that perimeter, either through technical or other means.

>> Loretta Garrison: Peter.

>> Peter Eckersley: Look, I think that's absolutely correct. I mean, I was going to make a point just about IPv6, which is -- this is an interesting story. If you compare IPv6 to IPv4 -- We use IPv4 today, but people are hoping that, one day, the Internet will use IPv6.

>> Lucy Lynch: I use IPv6.

>> Peter Eckersley: And a few people do. But there's a bit of a switching problem because you don't get much from using IPv6 until almost everyone uses it. So it's this hard bump for the Internet to get over. If you look at IPv6, if everyone implemented it naively, it would be a privacy disaster, in the sense that the specs tend to publish your MAC address in public view to the whole wide world, so, in fact, there's almost nothing you can do by default to avoid being instantly identified as soon as you get onto the Internet. And so that's kind of a bad thing. But then the number of addresses that you get from IPv6 has a much larger space than the mere 4 billion addresses in the current Internet. Those addresses, perhaps if we shuffled them the right way, that would actually give us the opportunity to make I.P. addresses less trackable because you could give people a new one every single time they popped onto the network, and then you wouldn't have a problem with tracking by I.P. address. Now, some people would say, "Oh, no. That's not good, because it means that people can't run their own little servers on their own machines that have a

persistent address for those," and maybe that's a problem that we can solve by some other intermediate ways. There is a way to look up an address for a transient server and get the different shuffled I.P. address every time it changes. So I think there are these consequences that come from these technical protocols, but Ed's point still stands that, if we want to talk about privacy, we need to not talk about just one of those things. We need to deal with this bewildering mass of different tracking mechanisms all at once, unfortunately.

>> Loretta Garrison: Jules, I want to turn to you because when question had a discussion before, you said that in the work you're doing, as you build new applications for the web, you look at the user experience in offline in order to design for online. I want to look at the addressing issue and the sending of information in a very oversimplistic way. If I mail a letter to you, the post office delivers it. They don't record that I sent it to you on a certain date, and they don't open it and read it. So that's an offline experience. But certainly, that's not the case online. Do you have any thoughts about that or want to talk a little bit about the way in which you are mapping online -- or, offline to online?

>> Jules Cohen: The context in which I think that it's really helpful to think about the relationship between the online world and the offline world is more in the identity-management space. And I'd be happy to talk a little bit about that now. So, just for a little bit of context, I think we've figured out a lot of the identity-management problems in the offline world reasonably well. I mean, we have methods that have grown up over generations, decades to figure out, in a particular context, if you want to prove who you are at a given level of assurance, you can do so. You carry around a wallet. As Drummond said, it has maybe a driver's license, maybe a student I.D., ATM card, Starbucks card, corporate I.D. And they all provide different information about yourself and about who you are in the real world. And depending on the context, you might choose, "Oh, I'm going to show my driver's license because I'm at TSA. Oh, I'm going to show my student I.D. because I want a discount at a museum. Oh, I want to use my ATM card to get cash. Oh, I want to use my corporate I.D. because I want to opt-in to some kind of service." And they're all used in different ways to access different services in the real world. And that model works pretty well. But on the Internet, as John was saying, we don't have that functioning interoperable identity layer. It doesn't exist. And as a result, we have what's essentially a rather kludgy method of using user names and

passwords and shared secrets. I think we all know the challenges with those -- with phishing and with identity theft and the like. When the thing that you use to prove your identity is something that anybody can type in on any computer anywhere in the world to access the kinds of rich information we've been talking about -- and not just the kinds of information that have been discussed here but also bank-account information, healthcare information, life-and-death information, the really hard stuff -- there's a challenge there. And so one of the questions is, we know we have this working in our operable model that works at a reasonably high level of fidelity in the offline world and provides reasonably good privacy protections in a bunch of contexts. How do we take what we have in the real world and move it over into the online world? And one of the things, I think, that you figure is that, in the offline world, in the real world, there are moments when trust is created. When I go to the DMV and show my utility bill and my Social Security card and relate it, there's this trust that's created over the counter there with a human. There's an in-person proofing moment. And at that point in time, that trust is bound into a plastic card, and they hand it to me, and then I can go reuse that trust offline to get services. The same thing happens when I register for school. The same thing happens when I become an employee. The same thing happens in a bunch of contexts. And in those contexts, there is a relatively high bar that I cross offline, and trust is created. And one of the challenges that we have online is that there are no similar in-person proofing experiences. It's pretty hard to get that level of trust to be created online because you don't have a human making a trust decision at the outset. So one of the things I think we need to do -- and we can talk about this more over the course of the panel -- is figure out ways to reuse pieces of trust that exist offline in online contexts at a reasonably high level of security and figure out ways to use those to make good privacy decisions about what happens with the subsequent data.

>> Loretta Garrison: Ari?

>> Ari Schwartz: Let Ed go first, 'cause he's going to make one of the two points I was going to make, but he's willing to empty out his wallet to do it, which I was not.

>> Loretta Garrison: Okay, Ed.

>> Edward Felten: Thanks. So, I think we need to be careful about the analogy to the real-world plastic cards. Here are the plastic carts that were in my wallet right now. All of these people know who I am. They know my name and address, and they could trivially link back to my identity and link their records together. There's a library card, frequent flyer, my work I.D., credit cards, driver's license. All these people know who I am. They can link my activities together. So I don't have great privacy protection there, at least as a technical matter.

>> Ari Schwartz: My point was exactly the same one, which is, I think we can do it better online than we do it offline, and that should be the goal. The goal shouldn't be to do it exactly the way we do it online. I think we can learn from the way that we do it offline, to help to try and figure out kind of a process to go about doing identity online, but the goal should be -- As Drummond was saying before, how can we de-link information to solve the problem that Ed was talking about. How do we build transparency in enough that people can see what information's held about them and make changes to it if it's wrong and it's something that's used to make decisions about them in the ordinary course of business? Those are things that you can do online that you can't do with systems that were designed in the world of file cabinets.

>> Loretta Garrison: Before we migrate, clearly, into identity-management issues, can we wrap up with just a couple of issues related to going back to the architecture and whether there are some structural changes? Whether they're big changes -- I haven't heard any big changes. But maybe there are some small ones that we can consider, which still would be important. Peter, you've talked about some which you call "low-hanging fruit."

>> Peter Eckersley: Absolutely. So, I think -- Maybe this is a terrible analogy, but if we're going to talk about low-hanging fruit, perhaps privacy is like we're trying to make a fruit salad, and in order to be a tasty fruit salad, it's got to have everything. There is low-hanging fruit. And one point I really want to emphasize is Commission Harbour's call for SSL encryption. I think that's a tremendous idea. It's really low-hanging fruit. It's a protocol that we have that's already developed. It's already widely in use. And, in fact, it actually addresses the question you were asking just before about, well, the post office doesn't open our mail. Using SSL prevents the network from opening your mail. And it's a great idea. It protects against hacking. It protects against all sorts of

privacy problems. Not all of them, but it's low-hanging fruit. Let's get it. Let's put it in our fruit salad. There are harder things that I think we should try to do. Ed mentioned before the fact that browser user controls -- currently, you need to be an expert, frankly, and very patient -- both an expert and very patient -- in order to get anywhere with the browser privacy controls. A question I wanted to ask him was, could we do better blacklists, with something like the Adblock Plus model, where you have a list of the bad things that you need to block? And that's socially constructed. It's an institution. We could crowdsource it. We could have everyone sitting down and studying the web and saying, "Wait, here's a new tracking company that has no relationship with the people they're tracking. Let's just block them." Could we do that? Would that be a feasible model? So that's a harder fruit to get, but maybe we could get it. And then maybe there are other really important ingredients that tie into the next subject we're going to talk about, and this is a question for all the identity-management people -- can we get anywhere with identity-management systems that give you throwaway identities that are nonetheless trustworthy? Like, I know that the cryptography is there. There are these fancy protocols called zero-knowledge proofs that, in principle, allow you to show up at a website and say, "Hey, I'm not going to say who I am, but I can prove that I'm a person in good standing." Is this a solved problem? Are we close to solving this problem? It's not exactly my field, so I'd actually love to know the answer.

>> Loretta Garrison: Okay. And for John, we have a question for you from the audience. You said earlier in our conversation that we're starting to see interesting designs for giving consumers more control over their data flows. Can you briefly describe some of those?

>> John Clippinger: Actually, this builds on an earlier point with the notion of I-Cards. We were involved in developing something called Project Higgins, and the analogy was to having different kinds of cards, but the difference is -- and I think we can do it better in the online world -- is that only the end-user, the person, knows they can link them. And you can have a different card generate an identifier, an ephemeral identifier. And I think we're going to move to a point where you're going to have authenticated anonymity, and you need to separate -- this is my view -- separate out sort of the physical person from the virtual, authenticated person. Because the real consequential damages are done to the individual. And they have life consequences when the abuses happen. So they take the DNA information. But you also have a social contract, in the

sense that information is jointly created about you. You have medical treatment, you get FICO scores, things like that. So you can't disassociate. But there may be mechanisms that allow us to have the cake and eat it, too. And I think this is new ground. There's new thinking on this. As your knowledge proves, I find it fascinating and very promising. And we worked with Microsoft with the Windows company Credentica that they acquired. And that technology's coming on board that, I think, could have an amazing impact to set up.

>> Loretta Garrison: Ed?

>> Edward Felten: This idea of authenticated anonymity is actually something that today's password system gives us, when it works, when you have secure passwords and so on. That is, I can set up a user account and password on one site, a different user account, different password on another site, and they're inherently unlinkable if I choose those well. But the problem is that there are so many other ways that those sites can link together -- the fact that, yes, this really is the same person. And once they have connected those dots, it doesn't matter how I authenticate myself to the sites. Again, there's this perimeter you have to defend, because if a link is made ever between my activities on the two sites, then there's no one doing that.

>> Loretta Garrison: Jules.

>> Jules Cohen: I just wanted to make a comment about the zero-knowledge proof -- I'll make two comments. One is that we're on the cusp -- Sorry. I'll make two comments. One is that we're on the cusp of a very sort of broad and deep conversation on identity management. And we could do it right, and we could do it in a way that doesn't exacerbate all the kinds of problems we're talking about. And the zero-knowledge proofs are a way to do that -- to allow unlinkability, to allow a number of properties unlinkability, untracability -- a number of properties that can improve the situation and, at the very least, don't make it worse. But it certainly doesn't address the plumbing-layer issues. I would just note that we released -- a couple weeks ago, at RSA, we released the foundational pieces of the zero-knowledge technology. It's called U-Proof, under the open-specification promise. So developers can go build on top of that freely, and we're hoping to see a significant uptake of the use of that technology.

>> Loretta Garrison: Okay. Oh, Drummond? Sorry.

>> Drummond Reed: Before we leave the technology layer, I want to build on what Jules just said. If folks are not clear when this term is used -- zero-knowledge-proof technology -- I want to make it very clear. Imagine you have an information card that is able to prove -- that actually has your birth date on it, okay? If you share that information with a site, it actually doesn't take much more than your birth date and maybe your zip code, one or two other pieces -- Even if you're not sharing a linkable identifier, you're going to be -- they're able to correlate you, or they're going to be able to link you. And this is just one example of the many ways that can be done. With zero-knowledge-proof technology, that information can be there, and when it's shared with a site, technologically, the site can prove that you are over a certain age but not get your birth date, okay? And it is a significant step forward. It's been widely vetted. Microsoft's acquisition of Credentica -- now their release of U-Proof at RSA, I think, is a major step forward. Information cards were designed to carry any type of token, including these new U-Proof tokens. So this is something we hope to see coming into use fairly quickly now. It's been theoretical for quite a while, but now it's a real thing, and what it could mean for privacy or authenticated anonymity, as John puts it, is, I think, significant. I want to say one other thing on the technology layer, before we move up. The other thing that I think is happening, and I'm putting on another hat, which is the XDI Technical Committee at OASIS. XDI is a protocol based on XRIs, and one of the key things it does is bind data and policy. It is a way of, whenever you share information, if you're able, on the part of the person sharing the information, to say, "This is the policy bound with it," okay? "This is the terms under which I'm sharing the data." In XDI, we call that the concept of the link contract. If you're able to do that, it introduces a new paradigm for how that information, that data and its use, can be respected throughout that life cycle, throughout that chain of trust, as John was talking about. Now, actually observing those policies is not something necessarily technology can enforce. But doing the binding, having a cryptographic way that that binding can be observed, is something that technology can do, so it's sort of how the two pieces can work together.

>> Loretta Garrison: Well, we want to talk a little bit more about the technology policy together and also bring in enforcement. We'll do that after we do more discussion on identity management.

But there's a question that's also come from the audience, Peter, I think this is to you. "Isn't the focus on SSL and the allegedly new problem of 'in the clear' traffic to the cloud really an old problem? How is this any different from truly ancient e-mail transport protocols like SMTP or POP, that involve similarly unencrypted traffic?"

>> Peter Eckersley: It's true that we've had a long struggle to move from plain-text Internet, where we all used Telnet and unencrypted SMTP and POP to an encrypted Internet where, unfortunately, as the network grew, it just became less true that could you trust the network never to listen to your user names and passwords or the content of your communications and do things that you didn't want done with those communications. So all of those examples of protocols are protocols that we're trying to encrypt. SMTP, you know, ideally, should go over SSL. POP definitely goes over SSL these days, really. And if you're not sending it over SSL, you're doing something wrong, and so I think the same lesson applies to the web. We've got the SSL protocol. It has its flaws. Those are fixable, I think. It'll take some work to get rid of the flaws. And right now, flawed SSL is a million times better than a plain-text Internet.

>> Drummond Reed: I want to know one thing. On Commissioner Harbour's list this morning, when she was talking about the e-mail provider, she didn't mention Gmail, and that's because after the China incident, Gmail switched over to use SSL by default, all Gmail connections.

>> Peter Eckersley: Yes. Many, many congratulations to Google for doing that. They showed that it was -- I mean, there was an argument, until Google did it, that it was too expensive for a huge cloud provider to encrypt everyone's e-mail communications. And Google has demonstrated, actually, it's not that expensive anymore. Computers have gotten fast enough that we can encrypt everyone's e-mail and still have it as a free service.

>> Naomi Lefkovitz: I mean, one question. And if you look at the -- even at the postal mail, you know, if the post office sees white powder leaking out of an envelope, they're going to open it, right? And then they're going to do everything they can to try to track it down. So is there some role for law enforcement and tracking of data, and is there some way to make those two compatible?

>> Peter Eckersley: I mean, I would love to be able to say that, just by turning on SSL, law enforcement is completely disempowered and needs to go get lots of warrants in order to access things. Realistically, that e-mail is still stored on the cloud provider's servers, and, frankly, a lot of the time, I think it would be not that hard for law enforcement to find a due-process way to access e-mail. The people who are really being locked out here are authoritarian regimes that don't have a legal process way to access that cloud provider. I mean, I think Iran was very unhappy about Gmail turning on encryption because it meant that they couldn't eavesdrop on their citizens anymore because the Iranian government couldn't go to Google and use legal process to obtain that e-mail.

>> Naomi Lefkowitz: Okay, well, let's move up a layer. Let's talk about browser controls. So, are there any technical changes that could be developed or implemented to address some of the privacy issues that we're talking about? And how easy or difficult would they be to implement, and how usable are they for consumers? Ed, do you want to start us off?

>> Edward Felten: Sure. This is an area where, I think, all the major browser vendors are trying to find ways to innovate, to give better technical controls, to give users more effective control over when they can be tracked and linked and what information gets provided. Historically, browsers have just promiscuously provided all kinds of information about the user, information which Peter and some of his colleagues and others have shown is often sufficient to uniquely identify a user. And that doesn't have to happen. It's not technically necessary, but it's a matter of really careful engineering in designing the browser to make sure that you're not inadvertently giving information that's useful. It's a matter of thinking about what information really needs to be released ever. It's also important to give users more control over what information gets released and to which sites. There needs to be a lot of change, I think, inside the plumbing of the browser. And then, to the extent you're giving users control and choices, you need to think really hard about how to present those choices to them in a way that's better than than we've historically presented privacy choices to users.

>> Naomi Lefkowitz: Is there any work going on in that?

>> Edward Felten: Well, there's a lot. I mentioned the browser vendors. This work that we're doing in our lab at Princeton, as well, to try to look at browser architecture and try to figure out how to let users compartment the information that's given to different sites and give users control over when sites can connect what they do on one site to what they do on another. And that means engineering the browser so that it keeps track of where information came from and so that it's careful about which information is given to who. And that's an issue that we're working on and also some folks involved with browser vendors, as well. So I'm hopeful that we'll make progress in this area. But it's a constant arms race, if you will, between people who are trying to find new ways to track and identify users and those of us who are trying to establish technological control over those. That perimeter that I talked about before seems to be getting larger, and there are people out there who are working to make it larger.

>> Naomi Lefkowitz: Okay. Anybody else?

>> Ari Schwartz: Go ahead.

>> Lucy Lynch: There's a little bit of an elephant in the room here, which is the user experiences, user driven. And in many cases, the user will do what is convenient and what delivers to them the experience that they've learned to expect. So, in many of the cases that we're talking about, we're talking about the user making an intervention, the user making a decision, the user making a choice, and in many cases, people will make that choice once, so it's good to get the defaults right. In some cases, people are willing to make that choice for a trigger event that they have to be notified about. So getting that balance right -- because in many cases, the browsers are promiscuously sharing information so that your experience is a positive one. You get the right plug-ins, you get the right whatever without the user having to actively manage their sessions all the time. And getting that balance right is one of the big difficulties here.

>> Naomi Lefkowitz: Sure. Go ahead.

>> Edward Felten: If I can just jump in briefly. What Lucy said is absolutely right, that -- And this is one of the reasons this problem is really hard. You need to give the users the experience, the benefits that they want from using the Net, and you need to do it in a way that is realistic about how much decision-making they want to do and how...

>> Lucy Lynch: And how often.

>> Edward Felten: ...how well-equipped users are to actually make those decisions. It would be easy if we didn't have these problems to deal with.

>> Peter Eckersley: So, I kind of made a point about this before, but I want to try and make it more clearly. This problem of having too many choices that are crucial. Essentially, you can imagine the innards of these browser settings as being a gigantic switch box with "Allow this site to send this bit of information to this other place," "Allow this thing over here to talk to that." And then you only -- And Ed pointed out. You only need to get this wrong once. You only need to allow Facebook and Amazon to link your accounts together once, and suddenly, forever, even if you chose a different user name and tried to keep those things separate, they're now associated in those firms' databases. And so the question is, "Okay, can we realistically expect human beings to be in there, in the switch box in their browser, saying, 'Yes, this site can talk to this one. No, this one can't.'" I think the answer is no, especially if we have any notion of what reasonable user ability looks like. And so the idea that I was talking about before, when I talk about crowdsourcing these things, you're saying, well, for a lot of us, the answers to these switch-box questions will be the same. It's going to be a complicated pattern of YESes and NOs about which things you want to allow to talk to which other ones. But let's try to solve this problem collectively. Let's all get together in some technical process and say, "Okay, let's try to answer the switch-box questions." And I don't know if this approach will work, but I think its the best one that we've got to try at the moment.

>> Naomi Lefkowitz: Well, just I'm a very practical person. And so how does that work? Who's going to sort of start the crowdsourcing?

>> Peter Eckersley: Well, the precedent that exists right now is this plug-in Adblock Plus and some other similar ones where the model's fairly simple. It's a list of things that your browser isn't allowed to load. Some of those things are scripts, some of them are images -- maybe the one-by-one transparent GIFs that are solely there to track you to make your browser go and fetch this tiny, invisible image from a web server just so that the server can see that you are where you were coming from to fetch that image. And people try to compile lists of these things and say, "Whenever your browser encounters a reference to one of these objects, just don't fetch it." And so that's a reasonable first approximation. Now the way -- Most of the way that Adblock Plus does this is by getting human beings to compile lists, and usually they target advertising rather than tracking. But I think the same model is equally applicable to the tracking stuff, which is probably of more policy concern. And then the question is, "Can we compile a good enough list that's long enough and comprehensive enough and has enough people looking at it and working on it that it gives people a solid percentage level of protection?"

>> Naomi Lefkowitz: Just going to let Drummond and then John.

>> Drummond Reed: I was just once again going to make this point about the difficulty of privacy by design. So, I think we all know that, as egregious as cookies can be for tracking, if we put the choice in front of the majority of consumers today -- "You can stop that problem. Just turn off cookies" -- how many would do it? Even if you made it one big red button right in front of them, their web experience would suffer to the point where a tiny fraction would take that choice. So if we're going to solve that problem, I would submit it has to be -- there has to be a more overarching solution to doing it. And Peter's got a good point in talking about, technologically, there are some ways that I would say are actually reflections in policy. Another structural way that I think we'll talk about as we get through identity management is this emerging paradigm of trust frameworks. And I would submit they are going to be a powerful tool for being able to approach that.

>> Naomi Lefkowitz: John.

>> John Clippinger: I was at a conference, South by Southwest, and Danah Boyd, who is an ethnographer of sort of the web and web behaviors, gave a very excellent talk on privacy. And I

think she came up with a very different perspective on how to look at -- rather than doing toggles and choices and intentions. And she was talking about different kinds of publics and expectations that people have in different kinds of publics and things that are behavioral -- what you call articulated. And people are very good at social signals. I mean, we have -- The bigger part of our brain is dedicated to interpreting complex social signals. The question is, are those signals there in the environment that people can build upon -- intuitively build upon? And they're constantly building new norms and inventing ways in which they create their own little publics. And so I do not think it's going to be -- I think we're just on the edge of understanding this. So I don't think it's going to be a complex set of toggles and switches. Yes, you have to have some kind of fundamental understanding, but I think that you're going to have to build on the sort of dynamic intuitions that people have and rely upon sort of cohort norms that people have. And you have to have some enforcement. You have to have some consequence of violation for bad actors. I mean, there's so much work now that's being done in behavioral economics and trust and how it works and how implicit trust mechanisms are created and enforced and how people develop contracts and conventions among themselves in an emergent way that I think this is going to require a different way of thinking about it. And I think to prematurely rely upon techniques that have not lasted -- it worked in the past -- project into the future -- will not be particularly beneficial.

>> Naomi Lefkowitz: One final point, Ari, before we turn to --

>> Ari Schwartz: Yeah, I just wanted to push back a little bit on what Drummond was saying about privacy by design is hard. I do think that privacy by design is hard if you haven't thought about it in the protocol at the beginning, and cookies is the perfect example of that, right? I mean, there was basically no thought to privacy when cookies were created, and now we have to deal with the consequence of that and create this whole complicated set of controls around it and rules about how they're used, et cetera. If we had tried to build user controls in the beginning, it would have been much easier than what we have today. So I think there's generally a viewpoint among some technologies, particularly among companies, that says, "We'll put the technology out there, and we'll figure out some of the -- We'll do rapid prototyping and figure out how to address those privacy and security problems down the road." It turns out that privacy and security in particular

are much more difficult to build in after the thing's been created. If we had thought about it at the beginning, we could have addressed it much more easily.

>> Naomi Lefkovitz: Okay, so we've been already talking a little bit about identity management, but let's jump in a little further, because we've already heard that many people have said that the Internet doesn't have an authentication layer built in. So, often people talk about identity management as a means of solving that problem. But let's make sure we're all on the same page. Lucy, do you want to give us a little nutshell of what identity management means, and what do people mean when they talk about federated identity management?

>> Lucy Lynch: Well, I think the first thing to recognize is that people are generally talking about that experience at a very high level in the network. They're actually talking, in general, about a web experience. Although, there is some work going on in federation below the web for services like data sharing. But in general, they're talking about the end-user facing experience and a relationship between a service provider and end-user and somebody who is your trusted relay among those relationships. And it's really -- If you think of privacy as not secrecy but as information sharing with context in a context, federated identity is really about allowing a user to select the pieces of information that you need to share to accomplish the service through a proxy so that only those details are shared. But it means that the user needs to move trust from themselves to the proxy. Or they need to be very active in acting as their own proxy with these zero-proof tokens. And that, in a nutshell, is really what you're talking about, is empowering both the user and the service provider to have a successful transaction without having to do a high degree of information sharing.

>> Naomi Lefkovitz: So, what will we get if we could build a good identity-management system? What's it going to do for us online, and what won't it do? And will it even necessarily increase privacy? Jules, do you want to start us off?

>> Jules Cohen: Yeah. Thanks. This is another question where I think it's helpful to look back to the real world. You know, in the real world, we have this way of proving our identity in various situations, and when you ask, "What will get online?" I think we'll get the same kinds of benefits

online, if done correctly, that we get in the real world. We'll have the ability to demonstrate who we are to a service provider. The service provider can make a trust decision about us, be that Microsoft or be that the federal government or the state government or whomever -- whomever you're getting the service from. And if you're -- If the token that you pass is the correct level of assurance and they deem it trustworthy, you get a service back. I want to just follow up on what Ari said, with respect to how we can potentially do it better online, with an example. So, in the real world, if I walk in to a museum and I want to prove that I'm a student at an accredited university, I pull out my student I.D., and they can see my name, they can see the name of the university, they can probably see when it expires and when I graduate and some other pieces of information. But really, they need a far smaller amount of information to determine whether or not I'm actually a student. They just really need something that says "bearer is a student." Doesn't matter whether I go to Harvard or Princeton or some other school. Those pieces of information aren't necessary. On the Internet, we can do that kind of redaction. We can share a token with somebody. A user can choose to share a token that says, "I am a student," redact the pieces of information that are unnecessary, and the relying party, the service operator, can make a trust decision based on that. Same example -- we talked about it a couple times -- when you go into a bar and order a drink in this country, they want to know, is it a valid I.D.? Are you over the age of 21? And that's about it. Does the picture match? They get a lot more information. And in the real world, it's really difficult to stick your thumb over all those extra fields and redact them. But on the Internet, that kind of redaction is possible. So, what can we get out of identity management? We can get those same kind of trust moments that we get in the real world, but we can get them with a higher level of privacy through redaction and through the kind of unlinkability that's possible that prevents the concern about linking all those cards in Ed's wallet.

>> Ari Schwartz: I think one important thing to note is that, in this space, privacy and security are very much aligned. I mean, if -- One of the problems that we have with security online is that too many people have information that they shouldn't have, right? And that's both a security problem and a privacy problem. If we can get the right people the right information at the right time to authenticate, to use services, et cetera, that will help both privacy and security. But that means looking at the whole environment of the Internet, and most companies, I think, are going to be somewhat -- are going to push back. I'm skeptical that companies -- what's called the relying party

in a lot of these trust frameworks -- that the companies are going to be -- are going to be enthusiastic about the idea of getting less information, even though it means that the whole system is more secure.

>> Naomi Lefkovitz: Ed, and then we'll go down the line.

>> Edward Felten: I want to amplify that. I think it would be nice. I, as a consumer, I guess, would like to live in a world where sites only ask for the information they needed to provide a service. But that's not the common practice today. If I want to get an account on newyorktimes.com, for example, to read the newspaper, they ask for my gender, they ask for my year of birth, my zip code, my job title. And, of course, I could lie. But they do ask for that information, and it's obvious why they're asking me for it. It's not to provide the service. They don't need to know my gender to provide the service better. You know, it's a transaction that they offered to me. They're up front about what they're doing. So I don't think they're cheating me, but nonetheless, I think the business practice is that sites ask for more information than they need to provide the service and that people give it. So you can have a better authentication mechanism for logging in to that site, but still, ultimately, they will ask me for the information as a condition of using the site.

>> Naomi Lefkovitz: John?

>> John Clippinger: I'd like to challenge that a little bit. Because this is something that we worked with and talked to a number of companies about. And that meant major financial-services companies, large retailers, a variety of parties. 'Cause the assumption would be, if they can get as much information about you as they want, they're going to do that. And so we had a working group that brought a number of these people. Now, I was very surprised, in a sense, that a lot of them do not want to have the liability of having the personal identifying information. And they would like to have a trusted relationship where you would give them a lot of information, that, in fact, they could really know what your preferences were. And so, if there was a vehicle, a means by which they could get the information they really need -- They don't have to know exactly where I live. They don't have to know, in many cases, the physical me. They maybe have what I would -- The

virtual me to have certain sets of attributes, so it is very valuable to them. They all carry on transactions. And have a sort of social contract, or an economic contract, around that that's enforceable. I think it plays both ways. I think there will be a change, and I was very surprised to see that there's this shift that's taking place. Some people call it the big flip. But I think you'll see that.

>> Naomi Lefkowitz: well, I have a hard enough time negotiating with my phone company to get a better rate. So, that sounds good, but how is that going to work on an individual basis? People are going to trade more attributes for something. Better discounts? I mean, are they really going to do that?

>> John Clippinger: Well, this is exactly is alluded to earlier As information starts to become valuable, then you get into this asymmetry at the bargaining position. And that has to do then -- And then there has to be a regulatory governance framework. But I think the enlightened -- from the enlightened party -- what they want to do is have a trusted relationship with an aggregate of their customer base, that they get the information they need in order to produce a better product. And if there's the attempt of a provider, a service provider, to coerce or trick and trap, and this is what I fear, then I think you're going to move into a very adverse environment. But I don't think -- There doesn't necessarily have to be that incentive. I think they can have a better business opportunity by not doing that. And I think the trust of exchange is going to be to key to building a brand.

>> Naomi Lefkowitz: Did you want to answer that, Drummond, or you want to talk about trust frameworks?

>> Drummond Reed: Both.

>> Naomi Lefkowitz: Okay. Go ahead. Well, first I want to agree with John. One of the reasons I am such an advocate of trust frameworks is I think that the building -- The identity management is just a first piece of this relationship layer that we're talk about. And in some ways, it's one half of the coin or one side of the coin. And trust frameworks are emerging as the other half of the coin.

And the power is, I believe -- to get into what Ari's -- a very good point, which is, how can you actually align the consumer's incentives for privacy and protection of their information and the business' incentives to get the information they need to best deliver the service. Really, that's what both of them are incented to do. The optimist in me says the trust frameworks are a means to do that. And let me make it clear what we're talking about when we're talking about a trust framework. As Lucy was explaining, the concept of identity management, as we're talking on an Internet scale, is that, from a consumer's perspective, you're able to use -- It's identities of service. You're using a specialized service provider to encapsulate this. I don't need a service provider to go up and show credentials from my wallet in a physical store or in a bar or something like that. But on the Internet, I'm not physically there. I can, with technology like information cards, put this service right here locally on my machine, on my laptop, or on my iPhone. But immediately, you'll start to see one problem, which is, "Well, if it's tied to that physical machine, what if I'm using a different machine? What if I'm over at a friend's house? What if I lose one of those things?" So, you're going to see that migrating into the cloud. You're going to see these credentials. Well, You're starting to suggest something that's very similar to what we have in the financial system, which is to carry out financial transactions. All the time, we use banks. And banks are trusted to be in that position of our intermediary was sharing that information. Of course, it's a highly regulated environment, for many reasons. And what we're potentially evolving here with identity management on the Internet is a class of service provider, referred to as an identity service provider or sometimes identity provider, which scares me a little bit because it's not really providing your identity, right?

>> Lucy Lynch: Broker.

>> Ari Schwartz: Yeah, a broker. And now the concept of trust frameworks is, in that context, fairly simple to understand, which is, "Now you've got a service provider -- an identity service provider -- that's in the position of serving as your proxy or your intermediary, sharing your information selectively with sites that are called "relying parties." They're relying on the information that's being shared by the identity provider. But think of it this way -- you've also got, potentially, an advocate. You've got a service provider who's in the business of helping you protect your information when it's being shared. So, to some extent, this starts to address the asymmetry

that John was talking about. When you go, as an individual, to a site and you're looking at a privacy policy and you have none of the technical legal capability to look at it, that's really an asymmetrical relationship. When your identity provider, which may be -- I mean, there are many examples of companies that are already looking at that business -- PayPal, Google, Yahoo!, AOL. They are in a position to say, "Okay, there are norms for these things. We can tell you what sites we found to have the policies that are favorable to you and which ones don't. So, the emergent idea of a trust framework is that there's a set of policymakers that say, "Okay. Four of these exchanges of identity credentials, there are certain requirements that identity providers need to meet, and there are certain requirements that the relying parties need to meet. We're going to specify those in a trust framework, and then there's going to be a way to actually certify that the identity providers meet what are called levels of assurance in the credentials they're issuing and that the relying parties meet levels of protection. And the best example is the one that's already been implemented by the U.S. government, a group called ICAM, and it's called the trust framework provider for adoption program. And they've actually outlined a process where the federal government using trust framework providers. And just two weeks ago, the first two trust-framework providers, for members of the American public can actually start to use commercial identity providers, to provide OpenID or information-card credentials. U.S. government websites. And those trust frameworks, developed by ICAM and the GSA, have a set of privacy requirements in them that both the identity providers and the relying party sites need to meet. That capability of a trust framework to set up that overarching set of agreements, and not just the security but the privacy norms, is, I think, a potent new tool, and it's not just for governmental interactions. It can be for any type of interactions.

>> Naomi Lefkovitz: Well, let me ask one thing. So, this seems to introduce, like, a whole new privacy issue. I mean, if you look offline, and take you your driver's license and use it at the bank or the airport or the bar, it's not as if the DMV knows where you've used that driver's license. But now it seems as if the identity provider is going to know all the places that you are using that identity. Is that of concern? Should we be concerned about that?

>> Drummond Reed: I imagine you're going to get about six different responses, so I'll let others go first.

>> Naomi Lefkovitz: Okay.

>> Ari Schwartz: I would say, yeah, we should be concerned, but there are things we could do to address that concern. And the optimist in me sees a system that you could set up here that's not too farfetched, that could address those concerns where there are rules about what identity providers can do with that information. You have the unlinkability between the two different sets of credentials that each of these have, so they use it in two different places. They can't compare it. So it actually provides stronger privacy and stronger security while still authenticating at a better level than what we have today. So, I think that the possibilities are out there. And to have a federated system where you could have more than one set of tools to use in different places. So I think there is a path out there, but it means setting up the right kind of rules and putting them into place and giving the right incentives. And I think one of the things I thought was really interesting from the FCC broadband plan was they took this issue head on. They supported the idea of federated identity -- exactly as Drummond laid it out just there -- but also suggested that there needs to be -- and they actually had a separate, pull-out box on this -- the idea of an FDIC-like entity to oversee -- which is a private entity backed by the government to provide insurance in this space. So it's sort of serving as the super-trust framework and setting up the rules by which, if you follow these rules as the trust framework, you get safe harbor, and you can -- in in space -- but you have to follow the basic rules. Driving people to -- The real benefit is you have the liability protection and the insurance backing of the government. And then, also, individuals can feel more trust in it, too, because of the data portability issues and the fact that they know there's some government backing behind this process. So, it's an interesting model. Not saying that it's perfect. I mean, we still have to come up with the right rules. But an interesting model to look at. I think it sort of adds a new dimension to this discussion a little bit. And that's just brand-new somebody?

>> Naomi Lefkovitz: Peter?

>> Peter Eckersley: Sure. I think there are different ways that this could go, and the fundamental question is, you're have an identity broker who's central party that controls where your information goes. We really want to know, just who does this identity broker work for? And I think there's a

spectrum of answers to that. There's the worst possible answer, which is that the identity broker is like a data breaker. That is, they work for relying parties. They work for the people who want to know who more about you, and they're in the business of giving your data to other people for a fee. That's the worst answer. Maybe there's an answer that's in the mill, which is maybe they're a bit like a bank. Do banks work for us? I think that's probably a controversial question. There are probably some situations in which they really do work for their customers and other situations in which the bank works for itself. And so that's a murky case. And talking about a data FDSE is bringing that image up for me. And then probably the best possible case would be if this organization really works for me, then -- and there's a way that they can really demonstrate that they wouldn't work for me. Maybe that's a good thing. Maybe that allows a bunch of trust relationships and allows me to have some competent defending my identity online and telling companies that want my details before giving me an account. You know, "Here are some fake ones," or, "No, sorry. We're going to negotiate as a group for our customers and refused to just close things that you don't need to know. So if we can get to that world, maybe this is a good avenue to go down, but we really got to make sure that we're going down the avenue where these organizations worked for their users.

>> Naomi Lefkovitz: Jules.

>> Jules Cohen: Couple of observations. One, just to follow up on the question you raised about whether the identity provider can see all the places that you go. I think that, again, looking at the offline example, I think it's instructive to think about some of the privacy principles that are inherent in the offline world and thinking about whether they can move over. So, for example, to your point, in the offline world, the person who issued my student I.D. or the person that issued my driver's license can't see where I used it. They're also issued in a decentralized manner, so I confuse which of those ideas in the offline world I use. Which is a huge thing that we need to move over into the online world. I don't always want to be using the same one. And then there's unlinkability. If I used my student I.D. here and I use it here and I use it here, or my driver's license, to Ed's point, yeah. It's possible to link them in the real world. But there's a fair bit of friction there. Somebody has to swipe it or photocopy it or write down the contents. And a lot of those transactions in the real world are ephemeral. Somebody looks at it, makes a trust decision,

moves on, and there's no record kept. To the extent that we can try and view the online world with that principle, as well -- that's great. I would note that there are some I.D.s in the real world where you show them and there is a record kept. ATM cards, store value cards, those sorts of things. They have a slightly different set of principles. The other thing I would note is, just around this conversation about the various parties involved in the transaction. So, you have the identity provider, you have the user, and you have the relying party. At some level, those are the core three. And one of the interesting challenges I think that we're noting here is that they all need to be accountable, and it's really a very shared thing. And if one party is less accountable than the others, then you have an imbalance, and it will work a lot less well.

>> Peter Eckersley: Jules, can I quickly ask you a question?

>> Jules Cohen: Yeah.

>> Peter Eckersley: When you talk about unlinkability, Ed made a good point before about the fact that we can make up user names and passwords for site, and those function a bit like throwaway identities, from an identity-management system, but they can be linked together by all the usual web tracking mechanisms, like cookies and third-party requests. Do these unlinkability properties that these frameworks purport to have actually, really protect us against that?

>> Jules Cohen: So, I think the question is, to what extent can the protections that we're talking about building in at the application layer, the identity layer, actually help us in the plumbing layer. And that's a question of whether those technologies can be boarded down. And I guess the point I made earlier, which is that the nascent identity layer could exacerbate the problem, make it a lot worse, or we can build in protections in that layer, not make the problem worse, and then there can be a different set of conversations about the plumbing.

>> Naomi Lefkovitz: I think Lucy's has been waiting a while.

>> Lucy Lynch: We've actually moved on.

>> Naomi Lefkowitz: All right. Ed, you wanted to answer them?

>> Edward Felten: Sure. I actually wanted to follow up on what Jules just said. Because, from my standpoint, I think there's a lot to like about systems that help you automate the management of your identity and login in an unlinkable way, and so on. And those can all be improvements. But from a privacy standpoint, I think it's a useful goal if those technologies to strive for making things no worse than they are today. Providing the level of unlinkability that you can get with passwords, for example. But, fundamentally, I think they don't solve one of the big privacy concerns that a lot of users have, and that is the kind of market negotiation that goes on in which a user reveals some private information in exchange for getting a service. You can build technologies that make that negotiation more efficient, that makes the result more precise, that improve enforcement of violations, but fundamentally, the negotiation is still going on, and there is a trap in private information that traffic in private information that is inherent in the way that the market operates. And I think you have to step outside of technology redesign if you want to change that. And you're opening a real Pandora's box if you do.

>> Naomi Lefkowitz: Drummond?

>> Drummond Reed: And I agree very much with what Ed said, and I want to point out that is sort of the fundamental purpose of trust framework is to say, "Okay, so we have this selection of technologies --" and I want to make a couple points here. First of all, no one's proposing a single, overarching trust framework for the whole Internet. In fact, the model of the two -- and Lucy and I are sitting side by side -- the two trust framework providers that were announced two weeks ago at RSA are Kintera and a new organization called the Open Identity Exchange.

>> Lucy Lynch: They're both only U.S. gov. This is not global.

>> Drummond Reed: Right.

>> Naomi Lefkowitz: Can you talk into your mike?

>> Lucy Lynch: I should point out here that these trust frameworks are U.S.-centric, and we're talking about a global technology and global policies. And at some point, interfederation will require that you, as a U.S. citizen, interoperate somewhere else in the world. So this is the kernel of the beginning of a solution to a problem which is much larger. There are some European programs like STORK already looking that the problem, as well.

>> Drummond Reed: Right. So, the key point of the Open Identity Exchange approach -- and there are two white papers at openidentityexchange.org that I highly recommend on this to address Lucy's very issue, which the OIX approach was to say, "Okay, the U.S. government has really proposed the first trust framework provider program," coming from -- it is a collaboration between the OpenID Foundation and the Information Card Foundation to help create OIX. The approach was there to say, "There are going to be many more trust frameworks. There are going to be trust frameworks." One semi-public, nongovernmental organization that's looking at a trust framework right now is PBS. It's saying, "We could use a trust framework in public media to help connect all of our member stations -- basically, a federation of all the member stations -- and also the websites that service many PBS shows. We'd like to have folks that have credentials from a PBS station as a member or supporter be able to go to shows -- the sites for those shows and be able to exchange information under a certain trust framework, a certain expectation. That's an example that's not governmental. There are a number of other examples given there. So, the point there really is that those contexts, as John putting it, can be represented by a trust framework. And to address the usability issue, from the end-user's standpoint, it's really sort of like making a decision, like a financial decision. "Do I trust this merchant?" "Do I trust this network that I'm working with, in particular, their credit card?" If we can do that, then we can get to a set of policies that are bound to the technology being used such that you're able to establish norms for good behavior at a fairly broad basis. It is a new way of essentially binding technology and policy that I think has the usability characteristics that we need.

>> Naomi Lefkowitz: So, let me ask. So, we started this session by talking about what we should have done in the beginning when we were developing the Internet if we had thought clearly enough about it. Now here we are on the cusp of a whole new system, and we're talking about, you know, how it could be set up to do it right, in terms of getting the benefits, as well as maintaining privacy.

So, how are we going to get to that "could"? Is that going to be -- Do we need regulation? New regulation, we have old regulations that work? Is this going to be self-regulatory? Standard-setting? How are we going to get there?

>> Edward Felten: How much time? [Laughter]

>> Jules Cohen: That's right.

>> Naomi Lefkowitz: John?

>> John Clippinger: I think we have to acknowledge, at least from my vantage point, is that we're designing a new kind of ecosystem, and they're precedent to the things, analogies we can build upon. But I think it's very different, and I think there's going to be evolving process. I think that the proposal that came out of the FCC and FDIC kind of model where you have regulatory as a backdrop to allowing the private sector to do things, I think, is an interesting way of approaching it. I think it's going to evolve over periods of time. I think that there's going to be a lot of learning, exploration. When Drummond was talking about trust frameworks, I think there are going to be lots of inventions in that, over time. One of the things that I want to mention is that we're looking a lot at mobile data, which can be highly identified. I personally think there has to be something like a personal vault that has stewardship with fiduciary responsibilities. 'Cause I think this information's getting very, very valuable. And, I think, as stakeholders come in and see these are the new kind of banks, these are the new kind of business models, there's going to be a lot of pressure. And you don't want to have some of the problems. We had the financial services industry spill over here. That is a very big concern that I have.

>> Ari Schwartz: As I said earlier, I have some concerns about how much the relying parties -- the companies that would be using these services -- would actively promote the idea of these services knowing that they may get less information and they may have to collect information separately. I do think that there are ways, if we can come up with incentives to get them involved in these systems, that there are ways to give users real control, to make this user-centric, as John and others have said on the panel. And to make that happen, one of the things that we look at is -- Sort of the

background to this is to look at what we have in the world today, following off of Jules, and looking at the Fair Credit Reporting Act is sort of a model in this space. We had a pretty detailed filing for this roundtable on that topic, but the basic idea is that we usually think of the Fair Credit Reporting Act as covering credit, insurance, and employment as sort of the main areas. But if you look at the law itself, it's actually much broader than that, in terms of really using background and repetitive information for eligibility for a covered need under the act, including one very broad area called "business need," which, if you look at the FTC commentary, is a whole bunch of areas of eligibility, including landlord/rental issues and even dating services are directly mentioned as examples. And I think what we're seeing in this space when you're talking about people getting access to services or getting access to sites on the web, we're starting to see a little bit of a blur between what the traditional split that the FTC used to have in this area where they said, "Well, if it's about eligibility, that's covered, and if it's about identity, that's not covered under FCRA." We're starting to see that now. Now you're going to start using these authentication services for eligibility to get onto a site, to make decisions for an individual. So we're starting to see that blurring there. I think that, in some cases today, we would have an argument that the FCRA applies, and we make that case in the paper. Some cases, there's a big gray area where some cases are fun. Some cases clearly would be out of the FCRA. But I think that the identity providers and the trust frameworks can learn from what the protections that are in FCRA and build a model around that, build contractual models, whether we put it in our work prior to what the FCC was talking about in the FDIC case, which we'd have to look into a little bit more. But in terms of the trust frameworks, building a three-party contracts between the user, the relying party, and the identity provider to get -- that include levels of protection around this area.

>> Naomi Lefkovitz: Peter?

>> Peter Eckersley: If I take the question is, "How do we get from the Internet of today, where, frankly, most of us have no privacy, to an Internet of tomorrow or the Internet we might have built back in the '90s if we'd had that crystal ball where privacy is there by default and most of us have it. I think the first thing to remember is that there's no guarantee, but we're going to get to that issue. It's actually going to be really hard. The costs of not getting to an Internet that protects privacy are actually very high. There are a lot of bad things that happen when you don't have privacy. But if

we want to get to it, we're going to have to try really hard, and that's going to involve throwing lots of kitchen sinks at this problem. We're going to have to have engineers working on fixing these protocols at the low level, but we're also going to have to have regulatory agencies looking over their shoulder and saying, "Are you doing a good enough job yet? We're going to have to turn around to the browser manufacturers and say, "You guys need to fix the cookie settings and the third-party content settings, and all of these things and all of these things, and we're going to need to have other nontechnical institutions making sure that happens. And then we may also need better privacy rules, as well. It may be that we'll get a third of the way by technical innovation and another third of way by implementing better privacy rules and then the last third of the way by magic and levitation. I don't know. [Laughter]

>> Lucy Lynch: It's a mystery.

>> Naomi Lefkowitz: Wave a wand down there.

>> Drummond Reed: No, I just wanted -- In terms of specific areas of focus, again, with this emergent -- As you put it, if the emergence of identity management trust framework is giving us a new tool at this relationship layer, then I do want to point out one specific area, which was, in the expectation, the initial way the trust frameworks were envisioned, for instance by ICAM, it was the concept of specifying levels of assurance from the identity provider. And what that means is, if you -- If one site needs only a low level of assurance that you are who you are -- a good example is a national park site, the ones taking campsite registration. They don't need to know -- They don't need to deeply proof your identity. They just got to make sure, if you're coming back to the site to change that reservation, you're the same person. Right, that's called level of assurance one. However, if you want to go look at your tax records or health records, you're have to be up at at least level of assurance three. And if you're talking about government employees or defense contractors, that's level of assurance four. Well, that's what -- These four levels of assurance were defined by NIS, and it's very well-established concept on the side of, "What's the level of confidence you have in the information coming from an identity provider?" When we started to look at this and say, "Okay, if trust frameworks are going to now be a tool for establishing policy for identity management, an Internet layer, there needs to be the corresponding concept on the

relying party's side. And it was Mary Rundell at Microsoft and co-author of one of those two papers that I pointed out. Said we should have that corresponding thing. Let's call it "levels of protection." And that's the levels of protection to which the relying party sites win the information shared with them, whether it's non-correlatable identifiers at level 1, and they actually have to say, "You know, our policy is we're taking a non-correlatable identifier. We're not going to try and correlate it. You're giving us other information that is correlatable, but at level-one protection, we're saying we're not going to correlate it, okay? On up to higher levels of protection. It solves the problem, "A," of making it understandable to consumers, and, "B," it establishes again these norms which, as Ari was saying, if they get established in trust frameworks for which there's societal pressure, if not regulatory pressure, to adopt, that's again a tool that could solve this problem on a broader basis.

>> Naomi Lefkowitz: Ed.

>> Edward Felten: I want to agree with what Peter said about the technical opportunities. But I want to add two things to that. One is that although I have high hopes for what we can do technically, and I'm certain we can do a better job in designing the technology to give users more effective control, some of the underlying technical problems are fundamentally hard. The technical issues here are things that we're going to be wrestling with for the longer term. Number two, I think there's an important role for self-regulation. I think there are important areas in which we basically know -- have some idea of where the line between responsible corporate behavior and bad actors would lie. But in some of these cases, there really is not a well-established line that is agreed upon. And I think, in a lot of areas, it's a matter of getting people together and agreeing on some brighter line that responsible companies can agree not to cross and then trying to generate pressure, through all the means available, including pushback from users and help from technology, to try to give companies an incentive to stay on the right side of that line.

>> Loretta Garrison: Well, I think the answer to today's question is that this is really hard. There are some things that we can do. We could have secured URLs, anonymous browsing. I look at browser controls, and I gather the browser companies are doing that. Deal with cookie settings. Look at things like the Adblocker Plus, as Peter called it, crowdsourcing. And identity

management, which addresses a part of what you do when you have to have transactions on the site. Of course, wrapped up in all of this are usability issues. There are also corporate governance issues and enforcement issues, so it's a very complicated topic, and I think that our panelists have done a marvelous job today of introducing us to these complexities and making the information very accessible. Thank you all so much. [Applause] We'd like to start promptly at 11:00 for panel 2, please. Thank you very much.