Remarks by Commissioner Julie Brill
United States Federal Trade Commission

Conference of Western Attorneys General Annual Meeting
Privacy 3.0 Panel

Santa Fe, New Mexico
July 20, 2010

Good afternoon and thank you for that very kind introduction. "Privacy 3.0" is a good title for this panel. I have spent a lot of time thinking about privacy over the past twenty years. From my perspective, during the last year or so, it seems we have entered a third realm of privacy regulation, the "3.0" stage. What I would like to do this afternoon is spend a little bit of time talking about the different stages of privacy regulation from the perspective of the Federal Trade Commission as well as from the states.

Please note that the things I say this afternoon are my personal views. I am not here representing any of the other Commissioners or the Commission as a whole.

Let us go back and think about the early stages of privacy regulation in the 1990s. "Privacy 1.0," from my perspective, was the "Notice and Choice" Model. We called it the "Fair Information Practices" principles. Although you might not be familiar with that title, everyone is familiar with the underlying principles. During this stage, the FTC and the states looked at privacy issues through a regulatory framework that called for notice, choice, access, and security with respect to information. We evaluated privacy policies that way: privacy policies on the web, practices of companies, and various self-regulatory regimes were all examined through the lens of Fair Information Practices.

The FTC, the states, and many consumer advocates called on Congress to enact these Notice and Choice principles into law. However, Congress did not enact sweeping legislation on these broad principles. But it did enact the Gramm-Leach-Bliley Act, which many of you are familiar with.[1] The GLB Act embodies Notice and Choice principles. Consumers are given a one-time notice. They are required to read it, understand it, and make an intelligent choice that often will last for a long time. It is an interesting model and I am going to have some thoughts and critiques about it in a moment.

Shortly after GLB was enacted, the Federal Trade Commission, as some of you know, switched gears, and moved from "Privacy 1.0" to "Privacy 2.0." It moved from a regulatory framework focused on Fair Information Practices to one focused on principles of harm. The Harm Model was first launched by former FTC Chairman Tim Muris, but it since has been embraced by many people, including in the states. The Harm Model focuses on harmful privacy practices that present risks of physical security or economic injury. As a result, the Federal Trade Commission, and the states, started focusing on

[1] 15 U.S.C. §§6801-6809 (1999).

data security, data breaches, identity theft, and children's online privacy, as well as issues such as spam, spyware, and telemarketing, including the Do Not Call list.

Let me expand a bit on the first two issues, data security and data breaches. During the Privacy 2.0 timeframe, regulators focused on enhancing tools to address data security and data breaches. First and foremost, the states, led by California but followed by many other states, enacted data security laws that required notification to consumers about data breaches. The Federal Trade Commission and other federal regulators adopted the Safeguards Rule under the Gramm-Leach-Bliley Act.[2] The GLB Act focused on financial institutions, and in that context included data security issues.

Within the Privacy 2.0 framework, the FTC started looking at various cases that came to light as a result of the states' breach notification laws. The FTC and the states analyzed the matters under Section 5 of the FTC Act and similar state laws, which prohibit unfair and deceptive acts and practices in commerce. In investigating security breach matters, the FTC asked "Was there deception or unfairness in the way that the companies were notifying consumers about their privacy practices, and in the way that they were implementing their privacy practices?" This analysis fell within the Harm Model—we were looking for harm to consumers—and it employed the FTC Act and the states' unfair or deceptive acts and practices laws to examine privacy issues within that rubric.

There were many enforcement cases brought during this era by the states and the FTC. Cases like *ToySmart*, *BJs*, *ChoicePoint*, *TJX* (parent of TJ Maxx), *LifeLock*, and most recently the Commission's settlement with Twitter fall under Privacy 2.0 and the Harm Model.[3] With respect to the recent *Twitter* case, we asked: "what did Twitter say it was going to do with respect to customers' information, what were its actual practices, and did the deviations between its promises and practices present potential harm to consumers?" This was typical of the type of Harm Model issues we examine in the Privacy 2.0 framework.

In addition to security breaches, there has also been an emphasis on identity theft issues, with attention paid to enhancing tools regulators have with respect to identity theft. The Fair and Accurate Credit Transaction Act (FACTA),[4] for example, allows consumers to obtain free credit reports once a year from each credit reporting agency, to allow consumers to determine whether or not they have been victimized by identity theft. Consumers can look at their credit report and make a determination on their own as to whether suspicious accounts have been opened in their name, or whether other suspicious activity appears in their credit report.

[2] 16 C.F.R. Part 314.

[3] *See, e.g.*, FTC v. Toysmart.com, LLC and Toysmart.com, Inc., No. 00-11341-RGS (D. Mass. 2000); In the Matter of BJ's Wholesale Club, Inc., FTC Dkt. No. C-4148 (2005) (consent order); United States v. ChoicePoint, Inc., No. 1:06-CV-0198-JTC (N.D. Ga. 2006); In the Matter of TJX Cos., FTC Dkt. No. C-4227 (July 2008) (consent order); FTC v. LifeLock, Inc., No. 2:10-CV-00530-NVW (D.Ariz. 2010) In the Matter of Twitter, Inc., FTC File No. 092-3093 (June 2010) (consent order).

[4] Fair and Accurate Credit Transaction Act. Pub. L. 108-159. 117 Stat. 1952. 4 Dec 2003.

The tools and principles of the Harm Model have been very important over the past decade, and have been employed to good use by regulators. But industry has been moving forward in ways that are not necessarily addressed by this Harm Model. There have been substantial developments with respect to the Internet and electronic technology, which have become much more sophisticated in terms of how consumers' information is gathered, retained and used. Very rich ecosystems of data are being created and deployed, paving the way for some very sophisticated forms of advertising.

I would like to talk about one of these forms of advertising, which is known as behavioral advertising. In my view, our two prior privacy models—"Privacy 1.0" and "Privacy 2.0"—do not really address the concerns that are now arising in the realm of behavioral advertising. I'd like to bring you up to date with respect to some of the things that people in Washington are thinking about in terms of how to better address the privacy concerns now arising in this realm.

But first, what is behavioral advertising? All of you receive behavioral advertising. If you use Gmail, if you are on a social network, if you do Google searches, you receive behavioral advertising. It is advertising that is served to you based on things that you have bought, things you have looked at on the Internet, things that you are actually saying in the content of some of your emails, people you are communicating with, where you are located, and what you are reading. Through this type of advertising, you receive very specific ads based upon your behavior on the Internet.

I will give you just a very quick example. Let us say that, as you are sitting here in Santa Fe listening to me, you are looking at a travel website, specifically at some flights to Miami. You do not have to actually buy the tickets to Miami, but you are looking at them. Then maybe you go to an online newspaper website a little bit later on today, and you read an article about basketball. Then you go to yet another website. There, you might receive an advertisement, designed specifically for you to read, about the price of airfare from Santa Fe to Miami. The ad might even mention a package for Heat tickets. So the ad very specifically reflects your behavior on the Web, even though you have not necessarily bought anything. That is the kind of advertising you might receive based upon your online activity. This is what is known as behavioral advertising.

Now, this type of advertising has some very important benefits for consumers, for industry and for the economy as a whole. Consumers receive information about products and services that they are more interested in and more likely to care about, which is good for consumers. We do not want to receive a bunch of ads about products and services that we are not interested in. And businesses are much better able to target their advertising towards consumers who want to see their products. This is the Holy Grail for businesses: how to determine which consumers they ought to be targeting and focusing on. But most importantly, this kind of advertising is what is underwriting the vast bulk of the free information that is available on the Web. It is very important to consumers that this kind of advertising be allowed to continue, because otherwise, much of the information we receive on the Web would no longer be free, and would instead become fee-for-service, which would become very expensive for many consumers.

But there are some serious privacy concerns related to this behavioral advertising activity. I just gave you an example of a relatively benign Web search or series of Web searches that led to a fairly innocent piece of advertising served up to a consumer. But you can imagine a much more invasive type of advertising, based on a much more personal or sensitive Web search. Say, for example, if I or someone in my family had a medical condition that was serious or very confidential, and I searched for information about it on the Internet. Through behavioral advertising, I could be served up advertising based upon that medical condition, which I might consider to be very intrusive. These are some of the very sensitive kinds of issues that people in Washington are thinking about right now in terms of how to deal with privacy in this new realm of behavioral advertising.

So why do our two prior privacy frameworks not work adequately in this realm? First, the traditional Notice and Choice Model places a lot of the burden on consumers. It requires us, as consumers, to read and understand very complex notices and make choices based on them, and we have to make our choices without really understanding all the ways in which our information might be used in the future. These privacy notices have really morphed from informative documents designed to tell consumers how their information is collected, used and stored, into legal documents designed to protect companies from liability. I would contend that companies issuing these legalistic notices are not behaving deviously; rather, they are simply responding to the design of the regulatory framework contained in GLB and the underlying Fair Information principles. When you consider that a great deal of behavioral advertising is now being served on mobile phones or mobile devices, you quickly realize that there is no way to effectively communicate what is happening to consumers' information and provide consumers with the ability to make a reasonable choice under the traditional Notice and Choice Model.

So what about the Harm Model? This model looks at whether there is some kind of tangible harm to consumers—economic, physical, etc. But there are some issues that the Harm Model does not adequately address. It does not really address the exposure of sensitive information through behavioral advertising targeted to children. Nor does it adequately address exposure of medical information. And more fundamentally, the Harm Model is a reactive model.[5] It stipulates that once we determine that there has been harm, we will try to recompense those who have been harmed. The Harm Model does not take a forward-looking approach, one where we would ask "How can we set up an architecture or a system that tries, in the first instance, to avoid as much harm as possible?" A framework focused on that question would be proactive, rather than reactive.

The debate underway in Washington is very much about what we can do to try to create a better model for dealing with privacy in this "3.0" realm. There has been a lot of discussion about this at the Federal Trade Commission, in Washington, and in Silicon Valley. I do not think there is any consensus about the precise shape that our privacy

[5] Daniel J. Solove, *Identity Theft, Privacy, and the Architecture of Vulnerability*, 54 Hastings L.J. 1227, 1232-45 (2003).

model ought to take in this “3.0” realm, but I do think there is some consensus emerging about the issues that ought to be addressed. Let me outline some of those emerging consensus points.

There is an emerging consensus that consumers do not understand the extent to which companies are collecting, storing, and using their information. Consumers have no idea when they are told that a company may share information with its affiliates that that could mean literally hundreds of other companies. There also seems to be a fair amount of agreement that the distinction between Personally Identifiable Information and non-Personally Identifiable Information has become blurry. It is no longer possible to be truly sure that when information is said to be “de-identified” that it will remain “de-identified.” Some of my co-panelists have done a great deal of work in this area, so I’ll let them describe the state of the art in re-identifying formerly de-identified information. And finally, with regard to notices to consumers, it is widely recognized that the static, one-time notice and choice model is just not working.

Now for my personal wish list for a new Privacy 3.0 framework. I would like to see industry develop “just in time” notices, so consumers are continually advised about how their information will be used. Icons and other forms of communicating these practices should be made simple and clear. I would like to see development of a universal icon, along with universal placement recommendations, so that consumers can really understand what the icon means and become familiar with using it. I also think notices should focus on unexpected uses. Right now, notices spend a lot of space telling consumer things like: “We are going to give your information to the company that is going to send you the product that you are buying.” But consumers already expect that kind of information sharing. They know their name and address has to be given to the company that is going to mail the product. That is an expected use. I think notices really need to focus on unexpected uses, first and foremost.

Of course, there are many other ways that privacy practices can be improved in this new age of behavioral advertising. I expect the Commission will issue recommendations for a new privacy framework that can address those issues. Needless to say, this issue is very dynamic, and important. You can expect more attention to this issue by the Commission, as well as Congress, in the months to come.