>> Manas Mohapatra: All right, good morning, everyone. My name is Manas Mohapatra, and I'm an attorney here at the FTC. And I have the pleasure of serving as the moderator for this panel. This panel is going to discuss facial detection technology, which -- it was alluded to earlier this morning -- focuses on using someone's facial characteristics to determine certain general characteristics about them, such as their age range and gender, but isn't focused on identifying who that person actually is. Today, we're going to be hearing about different ways that this technology is being implemented and explore what the privacy and policy concerns regarding such use are and how best to address them. Before we get started in the substantive discussion, I'd like to just start off with administrative details. We're going to begin with each of our speakers here giving a presentation, after which I might chime in with a few questions. And after everyone has presented, we'll have a facilitated discussion exploring some of the issues that have been raised during the presentations. For those of you here in person, just like the last panel, there are question cards in your folders. If you have questions, you can fill them out. And we have FTC staff around the room that will take the question and bring it over to us. For those people watching on the Webcast, we have someone monitoring the E-mail address facefacts@ftc.gov. You can E-mail your questions, and we'll get them included in the panel to the extent that we can. And with that, I'd like to go ahead and introduce our panelists, from closest to me onwards. To my left, we have Beth Givens of Privacy Rights Clearinghouse. To her left is Fred Carter from the Information and Privacy Commissioner's Office of Ontario. We have Andrew Cummins of SceneTap, Jai Haissman of Affective Interfaces, Harley Geiger from the Center for Democracy and Technology, and all the way down the left, Brian Huseman is gonna start us off. And Brian is the senior policy counsel at Intel's Washington, D.C. office. He focuses on a variety of issues, dealing with privacy, marketing, and trade policy. He joined Intel from the FTC, where he most recently served as Chief of Staff and as an adviser to the Chairman on a variety of consumer-protection issues. In addition to being a former FTC attorney, he's also held positions at the criminal division at the Department of Justice and various federal courts. Brian, can you start us off?

>> Brian Huseman: Terrific. Thanks, Manas, for organizing this workshop. And I also heard that we need to thank Maneesha Mithal, the head of the Privacy Division, who, out of her own pocket,

donated the coffee out there. So, I thank both Manas and Maneesha, as well. So, as Manas said, I'm with Intel's Washington, D.C. office. Intel is mostly well known for being a microprocessor company, but we also have a large software business. I'm here to talk today about one of our new software products called Intel's AIM Suite, which stands for Audience Impression Metrics, and this is a facial detection software that is used in digital signs. So, you know, the digital signage industry -- there are millions of digital signs around the world, including the U.S. and Canada, but most of these signs are not smart signs. They don't use facial detection. So, for example, the signs that currently use facial detection are in the very kind of low single percentage of all the digital signs around the world. And this is a very nascent industry. Intel's AIM Suite software just launched in August of this year, so it's been officially on the market for only three months. Let me kind of talk a little bit about what AIM Suite does. I'm going to try to do a demo afterwards, which will kind of explain more and kind of show you behind the scenes. But Intel's AIM Suite uses AVA, or anonymous viewer analytics, to determine your general age category and your gender. So, for age category, we divide it up into four age categories -- 18 and under, which is the youth, the young adult, which is 18 to 35, adult, which is 35 to 59, and then senior adult, which is 60 and above. And advertisers can use this information for two purposes. One, they can do content triggering of the ads, as a viewer is passing by, or they can do aggregation of statistics. So, especially the aggregation part is useful for advertisers because they can know the contents that appeals to these different demographic groups, and they can know where the best place is to put a sign. They know what ads work and what ads don't work. And of course this helps with the advertising dollars. In addition to these age and gender demographic categories, the facial detection software can also tell the advertiser the distance that the viewer is from the sensor and also the amount of viewing time or like how long they focused on the ad. So, did they just go by it in passing, or did they actually kind of pay attention to it and actually read it and look at it? In addition to the advertiser, there are also benefits for the viewer. You do get more relevant information that you might be more interested in. And although the second part, the ability to get real-time discounts or coupons, maybe through Q.R. codes or so forth -- that is coming in the future. Even though these facial detections is a very low percentage of digital signs currently, the coupons and Q.R. codes are even a lower percentage of that. You know, Intel -- we strongly believe in the concept of privacy by design, which was first championed by Dr. Ann Cavoukian from the Ontario, Canada Privacy Commissioner's Office. And we've tried to incorporate the principle of privacy by design into this technology. So, we do facial detection. We do not do facial recognition with Intel's AIM Suite product. So, we only look at a general age category and gender. And no images are recorded. No personal information is collected. Images are not stored. Nothing like that. So, we take privacy very seriously and have tried to build privacy protections into this technology. Here is just one example of what a digital sign might look like. And let's assume that you're passing by this kiosk in the mall. So, you might have a sensor up top that would do the detection. Then there would be a screen with some ads. The Intel software would be in the kiosk. You would also have a processor. That is if you wanted to kind of process and store the data locally on the kiosk. You can also have all the aggregate information collected via the cloud as opposed to on the individual kiosk. Currently we're working with some large-name retailers -- Adidas, Kraft, Proctor & Gamble, Harley-Davidson. I'm just going to give you quickly a couple of examples. This from Adidas -- I mean, I'm gonna put it bluntly -- is really, really cool. I mean, it's really going to change the retail environment. And I wish that we could have brought the whole display here. But how this is going to work is that you might walk into a footlocker, and there would be a huge digital screen from Adidas there. Using facial detection software, it would detect that you're a male of a general age category, and at your eye level, it would display the most popular shoes that those demographics might be interested in. The retailer also could have information about which of those shoes are in stock. You could virtually touch the screen and look at -- you know, kind of you do on Zappos -- you look at all sides and shapes of the shoe. I think it will really provide a lot of benefit to consumers, as far as helping to make their purchasing decisions and making the retail experience better. And then Kraft is another example that AIM Suite is working with. So, in the grocery store, we'll use facial detection to help provide you with ideas for dinners and different menus and stuff that you might be interested in. I want to give these last few slides before I go into the demo. I just wanted to give a few examples of some of the analytics and the charts that we provide to advertisers. This one shows the viewing time -- so, how long the individual -- or how long the people looked at the various ads, based upon the different products. Here, for example, is the total impression -- so ,the total number of viewers at each of these different sites. Here is an aggregate information which shows the gender breakdown for who looked at these various screens during this one month. So, I think we can go into the demo now, so... So, we'll see if this works. Okay. [Up-tempo music plays] [Laughing] Okay. Okay, so, what we're doing -- Now, if you look in the left-hand side here, this is what you would -- Okay, get this closer. Okay, if you look in the left-hand side -- these ads that are being displayed -- that is all that you would see. The rest of this demo is kind of taking you kind of behind the scenes. So, I am going to get in front of this sensor here. And you see that that box means that the sensor has detected a face. And because it is blue, that means that it is a male face. So, how do we do this? First of all, the sensor tries to detect that it is a face. And I need to turn, actually. So, this is just the ad display. What I'm gonna do -- Now the targeting is on. So you will see that, since I'm a male, they are showing a BMW ad to me. What this does is that the sensor first looks to find that it's a face. And it looks for eye sockets -- not the eyes. We don't know kind of the color of their eyes or anything like that, but it looks for eye sockets. Once it's determined that you have two eye sockets, that it's a face, it then looks to other things. It looks to the ears. So, for example, if two ears are showing, there's an 85% chance that that person is a male, as opposed to a female. There are females that have their ears showing, but the software and the algorithms just go down different decision trees. We also look at the nose and the lips and the cheekbones. And so there are mathematical differences between male and female cheekbones, and there are mathematical differences based upon age. So, for example, let me show you that if I put one hand over my eye, that two eye sockets have not been detected ,so it's not going to detect a face. And then once that's removed, there is a face again. We don't track. We don't record information. If I look away and then I look back again, it will register that as two different viewers or two different individuals. And you'll see kind of a rotating down here. It also shows kind of the age range. So, here is the age. So, half time, it detected me as a young adult. Half the time, it detected me as an adult. So, again, I'm 39, so I would fall in the adult category. The young adult is 18-35. So, you know, sometimes the software was a little off and made me look a little younger than I am. So, if there are any female volunteers that want to -- just to see. Do you want to just come up real quickly just so we can show how that works? You just want to stand in front of the sensor here. I'll get out of the way. So, there -- it detects that it's a face. It actually shows that it's a male face. [Laughter] So, there we go. Okay, let me look again. Okay, so for -- it's, you know, the lighting, whatever else. You people have asked about accuracy, so we have -- For gender, we have usually about like kind of a 94% accuracy for age range. We have about a 90% or so accuracy rate. Thank you. So, anyway -- [Laughs] So, that is the -- you know, with the lighting and all that. Thank you for volunteering. [Applause] And so I think, Manas, that's the denouement

>> Manas Mohapatra: Great. Just a few follow-up questions. You had mentioned, in terms of the prevalence of this technology, that this was in the low single percentages. Can you just give some raw numbers? Is that over 100,000 signs, you would think, in the U.S. right now that have it or less?

>> Brian Huseman: So, as far as numbers and kind of market share, we don't have that. But what we, you know, do know is that in 2011 there were probably about one million faces that were detected by our AIM Suite software. So, that's the number that we have.

>> Manas Mohapatra: Great. And just one other question that was alluded to in the earlier panel was that some of these systems can detect ethnicity, and I was just wondering if the AIM Suite can detect ethnicity.

>> Brian Huseman: No, we do not detect ethnicity. We decided we did not even want to try to detect ethnicity. There are mathematical differences between different racial categories -- you know, bone structure and those things. So, the algorithms kind of take those into account, but there is no detection of ethnicity at all.

>> Manas Mohapatra: Okay, great. Thank you very much. Thank you. Our next presenter is Harley Geiger. He is a policy counsel at the Center for Democracy and Technology. His work at CDT is focused on consumer privacy, health information technology, and national security. He's worked extensively on the issue of out-of-home behavioral advertising, serving as a member of the Standards Committee of the Digital Signage Federation, where he led the trade associations initiative to adopt digital signage privacy standards, which cover facial recognition. So, Harley.

>> Harley Geiger: So, I'm Harley Geiger, and I'm policy counsel with the Center for Democracy and Technology. I'd like to thank the Federal Trade Commission for having me here and also to let everyone know that CDT has recently released a paper on facial recognition and privacy. I have copies of it with me, so if anyone is interested, then please just see me after the presentation. So, as you've already heard, facial recognition occurs on a spectrum, with a corresponding level of privacy impact. Facial recognition itself -- the key characteristic there is that it can detect unique facial

features. There are different names for it, different techniques that are used for facial recognition for us. That's the key difference. And you can use facial recognition, of course, to identify an individual across systems or across photos. For facial detection, it may know the face. It may know demographic or emotional state, but, again, no images is saved. There is no unique facial feature that is also saved. From our perspective, this is much less of a privacy problem. Some of the commercial applications you've already heard and you're gonna continue to hear as the day goes by. But it's increasingly being used -- both facial recognition and detection -- in offline and online circumstances -- offline, such as in digital signage targeting ads to individuals, online such as in social networks like Google+ and Facebook, with their "Tag Suggestions" feature. Importantly, facial recognition is coming to smartphones. A lot of the major operating systems right now support facial recognition. So, very soon, it is entirely possible that we'll see millions of consumers casually wielding facial recognition through a variety of apps. So, the privacy interests for facial recognition, naturally, is that it can identify an individual based purely on facial features alone, which sounds simple, but actually it really fundamentally changes the way that we conceive of privacy in public. And it doesn't just identify an individual. With that, you can also pull associated content about the individual. You can search, do an image search on numerous open platforms for photos that the individual appears in. You can pull online content, such as blog posts or possibly travel patterns, shopper profiles, depending on what sort of database you're running it against. For facial detection, much less of a privacy impact. However, from our perspective, the key issue here is transparency. The companies that use facial detection ought to be transparent about the fact that they're doing so. The reason is because consumers have consistently rejected tracking for marketing purposes, even on an anonymous basis. Businesses have a strong interest in being transparent and in giving consumers choices in whether or not they're participating in a facial detection or facial recognition system because not doing so will sensationalize the issue and can easily lead to backlash. The privacy issues with facial detection and recognition are compounded by the fact that there are no laws currently that address facial detection or facial recognition, with the exception -- believe it or not -- of Illinois, at least from a consumer-privacy perspective. However, CDT does not recommend that the Federal Trade Commission endorse facial-recognition-specific privacy legislation. Rather, it would be ineffective. If you were to prohibit facial recognition tracking, then consumers will simply be tracked through innumerable other ways. Instead, we think that facial recognition, through biometric information, needs to be wrapped into consumer privacy legislation that is comprehensive and that covers numerous information categories. However, that legislation will also have an exception -- and properly so -- for publicly available information. And there's a strong argument to make that if you're walking around without a mask, you're making your facial features publicly available. And prohibiting individuals from taking a picture of publicly available information could raise some pretty serious free-expression and First Amendment concerns. So, instead, we think that the baseline consumer privacy legislation should support a safe harbor that is based on codes of conduct for industry that are voluntary, but that should offer, as companies, some tangible incentives such as a reduced form of liability if they are to adhere to them. And any code of conduct like that must be enforceable. So, we also recommend that the Federal Trade Commission and state A.G.s oversee the compliance with those codes of conduct. And this approach is more or less endorsed in the Department of Commerce's green paper on privacy. So, the industry already has a head start on codes of conduct -- the digital signage industry, at least. So, POPAI, the Point of Purchase Association International, and the Digital Signage Federation both have codes of conduct that do cover facial recognition and facial detection. And they did this in the absence of major scandal or significant government pressure, which is really quite different from the way things went with the online behavioral advertising industry. So, that's very encouraging. So, the digital signage privacy standards with the Digital Signage Federation -- they apply -- For full disclosure, as Manas had mentioned, CDT actually wrote the digital signage privacy standards for the Digital Signage Federation. But it includes facial recognition and detection and incorporates a full set of the fair information practices. And we think that it's pretty solid. [Chuckles] So, you don't have to read all that. I know that that's a lot, although the point here -- and this is, again, part of the Digital Signage Federation's privacy standards. But the point here is that it distinguishes between facial recognition and facial detection and has different requirements for each. The distinction is, as you can see, perhaps, whether the information that is collected can be linked to an identity or to an individual's property. Notice in the guidelines, in the digital signage privacy standards were actually a very difficult thing to craft because of the sort of unique circumstances that digital signage uses facial recognition and detection. But we, like the Federal Trade Commission, endorse the concept of layered notice. And so here you see three different layers of notice. First is a privacy policy that is available on Websites, the Website is the owner of the device of the digital sign or the facial recognition camera. We also recommend that the owner of the location in which the device appears have a privacy policy available on the Website because if you walk into a Walmart, for example, you don't know the name of the company that is using the device. You do know that you're in a Walmart. So, second, we believe there should be a notice at the perimeter of the area that is using the facial recognition or facial detection program. Obviously, this only works for an offline circumstance. Third, there should be notice on the device itself -- so, in the case of digital signage, if there's a physical card on the sign or it can appear digitally on the sign as parent of the noticing. The notice should clearly alert consumers that facial detection or facial recognition is ongoing in the area, and it should let them know what that information is being used for. In terms of consent, it's somewhat basic. For facial detection, we think that the privacy risk is low enough that consumers can opt out via notice. So, if there's a sign in the breezeway of a supermarket, a consumer can hopefully see that sign. Hopefully it's prominent enough. And the consumer can then avoid that supermarket if they don't want to engage in a facial detection program. For facial recognition, on the other hand, we think that it must be opt-in. And companies will have different ways of having an opt-in, depending on their business practices. But the opt-in must be informed, and we think that the opt-in should apply only to the facial recognition that's occurring in a particular area. So, if you're opting in to facial recognition at a mall, for example, you should not be opting into a distributed network of facial recognition cameras that are being used by the company that owns the system in the mall. Lastly, accountability. The digital signage privacy standards that we wrote advise companies to have internal checks on compliance and also to have training for employees, but as everybody in the room, I'm sure, knows, self-regulation -- It's very difficult to actually have strong enforcement and strong accountability. And going beyond the digital signage privacy standards, the digital signage privacy standards are just niche. They're just one application that facial recognition and detection are being applied to. And as you're hearing, there's a broad swath of things that facial recognition can be used for, in a variety of different contexts, and obviously the Digital Signage Federation standards or the POPAI standards won't apply to them. So, our solution, as I said, was to codify incentives for companies that use facial detection and recognition into consumer baseline privacy legislation and to enforce the codes via the FTC and the state attorneys general. And that's it. Thank you very much. [Applause]

>> Manas Mohapatra: Thanks very much, Harley. I just had a few follow-up questions. You emphasized the need for transparency. Are consumers being made aware that these signs are out

there? Is it just the fact that not many signs using this currently, or could industry be doing a better job in terms of providing notice?

>> Harley Geiger: So, industry could be doing a better job, but it depends on the company. And some companies I know are compliant with the guidelines, and they do provide that notice. However, I have seen in press reports companies literally declining to point out which signs actually have facial detection or facial recognition and -- quote -- say that they don't want their customers to feel uncomfortable." But my argument is that the secrecy will sensationalize the issue and will actually lead to the consumers feeling much more uncomfortable. So, the industry could be doing a better job, but, you know, some companies are doing a fine job.

>> Manas Mohapatra: All right. And one other follow-up question -- if I understood you right, for the lowest level of privacy impact for technology that is using facial detection, the idea is that you can opt out by not entering the premises of a place that has these signs. Is that right?

>> Harley Geiger: Right.

>> Manas Mohapatra: So, do you have any thoughts about, should this technology not be implemented in particular areas -- for example, you know, healthcare facilities 00 so that people don't have to make that kind of choice?

>> Harley Geiger: Yes, actually. So, healthcare facilities, locker rooms, bathrooms, a lot of places that people must go to and really have little choice in going to. You could argue that a supermarket qualifies, but I think that's going a little far. So, yeah, but we did not cover that in the guidelines.

>> Manas Mohapatra: Great. All right, thank you very much.

>> Harley Geiger: Thank you.

>> Manas Mohapatra: So, our next presenter is going to be Jai Haissman. He is the founder and C.E.O. of Affective Interfaces, who are builders of a motion-sensing technology. He founded

Affective Interfaces to build the emotion layer of the Internet and further a new means of human-computer interaction and user interface. Jai's going to tell us a little bit more about how his company utilizes facial detection technology in their product. Unfortunately, we couldn't get his slides loaded. For the archive versions of the Webcast, hopefully the slides will be up there.

>> Jai Haissman: Thank you. Okay, maybe I can act the slides out. [Laughter] And you guys can guess what I'm doing. We do have some very nice demos in the deck, so I'll post the presentation on our Website and please, if you'd like to have it delivered to you or a notification when we do post it, then please approach me afterwards. Okay, so, my background is as a psychologist. And I became very interested in how to make a scalable technology that could address the mental healthcare needs of the population in a way that an army of therapists would not be able to. So, the issues are at hand are soldiers returning with post-traumatic stress disorder, children that don't have a secure emotional context growing up and the lack of appropriate feedback systems for appropriate emotional learning or emotional regulation. So, when I found out about the potential for this technology, it was very, very compelling, from a clinical perspective. What if you had a real-time feedback system that could tell you how you're feeling and then also could educate you about the emotional expression, through facial expression, of the people that you're encountering? So, I founded Affective Interfaces to answer some of these questions. And the applications are much broader than the clinical and the educational. So, yeah, let's get into it. It looks like my Website is not up there. Okay, well, we're at affectiveinterfaces.com -- all one word -- affectiveinterfaces, and then .com is the Website. I encourage you to come to have a look. All right, so, what does it mean to be the emotion layer of the Internet? We intend to have a client that will sit on devices and be able to track facial expression on demand. And it can be utilized in a variety of ways, but this is where your mobile device or your computer will be able to see how you're feeling moment to moment and then correlate that with whatever it is that you're interacting with. So, say you're playing a video game and your excitement is waning. Then the video game can pivot in real time and start to deliver content that's highly personalized to you and generates more excitement. Or say you're a movie director and you're looking to create a more powerful edit. Split-test your edits. See how your demography, your target market, is responding and then change the edits as you make your decisions. The opportunity with a computer that's able to do this accurately and in real time is very promising for interactive content -- so, where your device is

responding, very much personalizing to you. Not feeling so great today? Then it delivers a new iTunes playlist, very different from if you're feeling frustrated or very happy. So, these are the directions that we're moving with our technology. Effectively, it allows us to capture data anywhere and then using just simple off-the-shelf Webcams -- no specialized equipment. And most of the processing can be done locally with our proprietary algorithms. We've done work with Anheuser-Busch InBev and with Proctor & Gamble on market research studies, and we're working with some digital media companies for building this new interactive type of content. So, you might be wondering if you're able to accurately assess emotion, that's very private information. You know, when we have face-to-face contact, then we're able to do that more or less accurately, depending on our own sensitivity and empathy skills. But if we've got devices that are using this, then how can we protect our privacy interests? How can we feel like our own private emotional world is taken care of? So, our policy around that is really about transparency and informed consent and then reciprocity, providing some sort of value -- if we're realizing value -- from the exchange. So, every user knows when the technology is turned on, and they have to authorize turning it on in order for it to be working. And we feel that that's important, from a brand perspective. So, our perspective is really from industry. How is it that we create a context for our technology to be welcomed and perceived as valuable, as opposed to as a privacy threat or as something that's intrusive? People have, more or less, comfort with the technology just based on your own perspective. But creating that basic trust around brand, we think, is very important and also very important for the industry. So, we really encourage those that are involved with these behavioral targeting metrics to shape policy or company policy around that very specifically. Okay, so, let's talk about just why is emotion useful and relevant. Emotion is a fundamental driver to human behavior. So, we're talking about the primers for what we decide and how we associate values to things. We think about our decisions, and how we preference those decisions is largely informed by some kind of emotional tone. "I like this. I don't like this. That makes me nervous. That's really exciting." So, to capture this kind of data involves -- or provides an extremely useful feedback system for enterprise to understand how consumers are responding to their content, to their messaging, through content development. For the user, it provides a very compelling feedback loop for understanding how people -- "How am I doing today, and how's that trending over time?" For government, it's very useful potentially as a truthfulness assessment tool that is likely more accurate than any other existing technologies. There's a lot of work on this by Dr. Paul

Ekman, whose 40 years of research has established that these facial expressions, the microexpressions that we are constantly communicating with, are universal across cultures and are highly reliable indicator of emotional status. So, let's talk a little bit about the alternatives to facial expression for understanding emotion. We can use neurofeedback systems that are these dry-sensor headsets, and they track our brain waves and give us a little bit of information about how we're feeling, mostly on the level of excitement and arousal and cognitive processing. But it can't really get down into the emotional brain center because there's too much brain tissue in the way. So, we use behavioral indicators which are hard wired to the emotional brain. The face is displaying how we're feeling in real time. And unless we're trained to mask it, it's a constant broadcast of the inter-emotional state. So, that's a very high utility in understanding customer response, but then also for generating self-awareness. So, if you can see a trend line of your happiness or frustration, that's very important. The other applications for the interactive media space we think are extremely promising, this real-time gaming content and so on. We're announcing a large data set, which will allow us to do big data calculations on distributed emotion. So, that means if we know how people are feeling based on the area in which they live because they're using our technology, we'll scrub that of identifying information so that it's not personal to that individual. And then we can represent it as regional hot spots for emotion. And we think that'll be extremely useful for predictive health metrics and correlating heart disease to anxiety or depression and then also seeing how markets respond to local changing emotion states. How much time do I have? Okay. All right, so, yeah. I just want to underscore that I think the really important thing we want to bring to this is that we're shaping a culture around transparency and around, you know, informed consent and reciprocity. Thank you very much. [Applause]

>> Manas Mohapatra: I just had a few follow-up questions. You know, you had mentioned the capturing of data and the amassing of this data set. Are the images that you are capturing -- are those retained? And if so, then who gets access to those?

>> Jai Haissman: Yes, and it depends on the application. So, in the case of a user where they would like to keep a profile and then have that accessible to them, then we will retain the identifying information. For applications where it's part of a larger data set, then we'll likely scrub it of any kind of identifying information. We've been approached by companies that would like to

have these, I guess, location-based facial tracking for immersive labs or AIM-type applications. And we would not use any kind of identifying information. So, we always have informed consent about our policy and what we're going to do with the data.

>> Manas Mohapatra: All right. I guess the question is, in terms of identifying information, I think there's a question as to whether or not an image of your face is, by its very nature, identifying, and so if that image is actually being retained or shared with anybody.

>> Jai Haissman: Well, the other thing is that we're not doing facial recognition, so there's no identity in the acquisition of the data and the analysis of the data. But we might ask a person for their E-mail contact information. But we're interested in making meaning of facial expressions rather than trying to identify the individual. So, that's fundamentally first.

>> Manas Mohapatra: Great. That's very helpful. Thank you very much. Great. Next we have Andrew Cummins, who's the Chief Strategy Officer of SceneTap. Andrew is a strategy expert in the technology and defense markets and previously worked in multiple strategic development roles at Boeing Defense Space and Security. Andrew and the SceneTap team leverage facial detection technologies and apply them to both data analytics and the social-media industries. Thank you.

>> Andrew Cummins: All right. Thank you, Manas. I appreciate it. I hope everyone's doing well this morning. It's a pleasure to be here. I want to thank the FTC for putting on the workshop. As Manas said, I'm Andrew Cummins. I am the Chief Strategy Officer at SceneTap. SceneTap is a data analytics company. We also have a social-platform component that's involved. What SceneTap is really doing is we are applying -- we're an example of how you can apply facial detection technology in a unique manner to really produce a positive outcome, commercial outcome, on many different industries. And we are just one example. So, over the next few slides, I just want to talk about a few different areas. First I want to talk about what SceneTap really is, who we are, what we do, and why our customers are so excited about the services that we provide them through this technology. I want to talk about facial detection technology and how we use it specifically. And last, we'll have a discussion on our proactive approach to privacy and how we guard individual privacy. So, let's see. So, first of all, I want to start with a foundational

conversation on really just what SceneTap is in our business so you can have a better understanding of, really, what we do. So, first, SceneTap is really leveraging innovative facial detection technology and we're combining that with our own proprietary systems and we're producing positive benefits in the nightlife industry. So, when I say nightlife, I mean bars, I mean nightclubs, I mean lounges, those types of venues. Since we are at the Face Facts workshop, I decided it would be fun to talk about some SceneTap facts. First of all, SceneTap, I want to mention, is a consumer of this technology. We do not develop this technology, nor do we produce this technology. So, again, it's really just -- We are an example of how you can use this technology in a creative, positive way. Secondly, we use Intel's AIM Suite of technologies. So, Brian was able to give you a great overview of how that works. So, what he showed is essentially what's going on behind the scenes. So, the safeguards that he talked about that's integrated into that technology also applies to us. We'll talk about the next couple bullets a little bit more on the next slides. First of all, what we do -- We're in the business to help these venues really optimize their business, okay? That's really our main goal and our value proposition. Alternatively, we also help the public plan their evening more effectively. We're able to give them information that they've never had before. And then last I just want to have a conversation on, you know, really how we are working hard to ensure privacy, our proactive approach, and we'll explore this later. So, to talk a little bit about our business approach, to give a little bit more background -- we're built on what I like to call a tri-benefit model. That refers to the three customer sets that we serve. Really, our main value proposition is targeted to venue operators. Like I said, we are in the business of helping them optimize their operations, okay? So, what we do is we are able to -- Using our technology system, we are able to essentially track information of patrons as they walk in and out of these venues. This is not facial recognition technology. We're not capturing any unique individuals. What we're doing is we're taking the detection, grouping people into buckets -- gender and age and then obviously the count -- and being able to use that data to show them trends and help them optimize based on trends. So, that's the consulting services that you see up there. We use the data to help them optimize. Now, for the venue patrons, what we do is we take this data and we send snapshots to the public via our Website and our smartphone applications. What this then does is allow the public to be able to see what's going on in real time in a certain venue. So, you can see, for example -- and I'll show this on the next slide -- how many people are there or -- excuse me -- how full it is, really -- the average age, and the gender split. So, it's just basic, aggregated, demographic data. Again, it tells you nothing who individually -- you know, Andrew Cummins. It wouldn't say Andrew Cummins is in "X" venue right now. And then lastly we have relationships with third parties, and really, again, that's just consulting services again using the data and the technology that we have, helping them optimize their business. These are businesses that are, you know, in related and adjacent industries to the nightlife industry. So, you're probably asking yourself now, "How does this technology work?" How do we collect the information? So, really, it's our technology. It's our unique technology system that differentiates us. And so how it works, how we capture the data, is it's -- it's essentially a combination of sensors that are inside the venues, facial detection technology, the AIM Suite technology, people-counting technology. And around this whole thing, we have a proprietary system which is wrapped around the hardware and the software. So, again, what we're doing is we're capturing data in a real-time manner. So, what you'll see on the right-hand side -- On the upper part, you see a graph of what a potential patron, the public, could see on our Website of a venue. So, you can see it says right now -- you know, at that moment, it was 34% full. The average age was 30 of both females and males, and you have a 62% male ratio, 38% female. So, again, aggregated, general basic groups of information. It's just another way to help, you know, make a decision as to what you want to do for that evening. Now, really where the value comes in and where our value proposition is targeted is at the venues. You'll see in the bottom part, this is an example of a graph that we could produce for venues. And again it's this demographic information. We can show trends. We can see what's happening, you know, performance on the last four Thursdays. We can overlay this data with sporting events, weather, if there's events in town, concerts, other types of events. So, whatever it might be. So, they can really get a feel for how their traffic is driven based on the macro environment around them. So, then lastly I really want to end the discussion with the reason why we're all here, with the privacy discussion here. SceneTap really, truly is committed to privacy. As I discussed earlier, SceneTap really does have a proactive approach to ensure individual privacy. We've been operating to the best of our ability within the current frameworks and guidelines that exist. So, just to go over some of our, you know, beyond the AIM Intel Suite that we discussed earlier, to go into some of our own protections that we've built in to ensure individual privacy. First, this is anonymous tracking, okay? We do not track unique individuals, nor do we care about unique individuals at all. This is only grouping people into demographic buckets. Next, it's non-individualized data. We're not sending unique, individualized data points to the public. We're not giving that to the venues. This is all aggregated data when it's lumped into the buckets, okay? Let's see. So, we do -- All the data is secured. We have a closed-,circuit password-protected video feed in each of our venues. So no one within the venue can access any of the data. This is all behind, again, this encrypted, protected system. There are no recorded videos, as Brian was mentioning. There's no recorded video streams. There's no stored images. The data is logged, and then the video itself is actually destroyed on the fly. So, really, the sensors are truly just an eye for the software, and that's all it is. And then lastly, I wanted to talk about transparency because I know that's big issue in this area. SceneTap actually protects -- or, excuse me, you know, works in the area of transparency in two different ways. So, first we do have decals to provide consumer notice. So, in each one of our venues, there is actually a SceneTap decal that's in the front window, right by the entrance, okay? So, that essentially gives consumers the opt-out choice that Harley was just mentioning. So, if you can see the SceneTap logo, if you don't want to walk in, that's completely your choice. Secondly, we also have a privacy policy on our Website and on the application. It describes how we collect the data, why we collect the data, what we do with the data. It's available to the public. So, those are really two layers of which we provide this transparency to give consumers the knowledge of everything that we're doing. So lastly, I just want to thank everyone for your attention. I hope you now have a better understanding of SceneTap, how we're using the technology in a unique, creative manner to really leverage this innovative technology to produce positive benefits in different markets and for society. If you have more information, you can please visit us at scenetap.com. If you want to become a user, we'd also appreciate that, as well. And I appreciate your time. [Applause]

>> Manas Mohapatra: Thank you very much. I was just wondering -- in terms of the decal that you were mentioning, is there anything besides your logo that appears on it -- just the fact that there's a camera or the Website address? Or is it just the logo?

>> Andrew Cummins: It's not just the logo. It provides the Website address, and it also provides pertinent data to the company and kind of what we're doing.

>> Manas Mohapatra: Okay. And in terms of your privacy policy, do you list all the locations that you're operating in?

>> Andrew Cummins: The locations are not, but essentially all of our locations are listed as soon as you enter the front of the Website or the mobile application. They're all listed there, so...

>> Manas Mohapatra: And how many locations are taking advantage of this right now?

>> Andrew Cummins: So, in Chicago, which has been our beta market and our launch market, we've been up and running for about 6 months, and 50-plus venues in Chicago. And actually we are launching the market of Austin, Texas, tomorrow, which we're very excited about that. So, we've got about another 25 to 30 on that market, and hopefully we'll be expanding beyond that.

>> Manas Mohapatra: Great. Thank you very much.

>> Andrew Cummins: Thank you. [Applause]

>> Manas Mohapatra: So, next we have Fred Carter, who, since 2004 has served the Ontario Information and Privacy Commissioner as Senior Policy and Technology Adviser. His primary responsibilities involve providing strategic research, information, and advisory services to IPC commissioners, management, and staff on a wide range of technology and privacy policy issues. Thank you.

>> Fred Carter: Thank you. Thank you, Manas. I'm really grateful for the invitation and the opportunity to talk to you here today. I'm here on behalf of Dr. Ann Cavoukian, who is the information and privacy commissioner of Ontario, Canada. It's our largest province. Dr. Cavoukian has been working in the privacy business for about 25 years, and as a commissioner, this is her third term. She's probably the longest-serving privacy commissioner in the world. Our office is sort of broadly similar to the FTC. We're independent. We carry out investigations. We can issue orders in some cases. We're an ombuds-type organization. We oversee three laws -- access to information, as well -- over the public sector and the healthcare sectors in Ontario. We have a number of functions. What is most pertinent to here is the research and education mandate. Our office is well-known around the world for its proactive views on emerging technologies and their impacts. And in that context, we've been very active. We've done a lot of work in biometrics.

I won't go through this list in great detail, just to say that, you know, about two decades ago, we, you know, embedded in statute requirements for use of biometrics in Ontario, way ahead of anyone else. We've issued numerous guidance and discussion papers in 1999 on policing in biometrics. We've taken positions on public-policy issues, national I.D.s and sponsored international resolutions on biometrics. And I was a member until very recently the International Biometric Advisory Council of the European Biometrics Forum. I'm not gonna read this list. It's just a list of publications. I just want to give you an indication that we've really ramped up our activities in biometrics in the last four years with the acquisition or the hiring of world-class staff who are able to actually go under the hood and really examine biometric technologies and write good guidance papers and technical papers. We help out other data-protection authorities in evaluating various technologies and coming to their conclusions. All these publications are available on our website. Okay, so, about a year ago, we were approached by a small start-up firm in Toronto called CognoVision. They had a technology that involved face detection, and they wanted to do it right. They wanted to build privacy in from the start, and they wanted to apply privacy-by-design principles. We had a look at it, and we thought, "This is pretty good. Let's write a paper. Let's write two papers together." This is the paper that was out in the hallway. It's gone now, but you can get it online. This was the short version. And essentially, the long version may not get written because we're so impressed with Harley Geiger's overview of the industry that that's basically, you know, pretty good. So I'll just focus on this paper. And CognoVision got bought out by Intel, so you really actually already know about this technology, because Intel has described this as the AIM Suite. And we were really impressed, and we applied privacy-by-design principles to it. And we were impressed for a number of reasons -- the proactive commitment. I mean, if you're gonna say that your video analytics is anonymous, we're gonna hold you to it. It really had better be anonymous. You know, really important is that nothing is retained on the devices. There's no image or a template retained or transmitted. There is, as a result, no possibility of secondary uses of it. That's really good security, and it's really good data minimization. In Canada, your face is personal information. And as was pointed out, if you go away and come back, you're recorded as a separate individual. There's just no linkage between the impression. No identification -- another really important threshold. Bright red line, as it were, is that identification does not take place, at least in this technology that we've seen. No matching or linkage. It's always possible in the future, and I'll get to that in a second, but that's a really important criterion. So, we know, of course, that any claim of anonymous, you know, maybe shouldn't be taken at, well, face value, but re-identification is always possible, so it's always gonna be important to look at how the technology is actually deployed in the field rather than shipped from the vendor. Appropriate signage has already been discussed. it's worth reflecting that this is remote technology, and you are detected whether you know it or not. And this has very important implications for applying Fair Information Practices. You don't know it's happening. It's taken at a distance. You have no knowledge. You aren't able to give consent. So signage is very important, just as is building in the privacy safeguards by default. It really should be not just a matter of policy, but built into the technology that your identification isn't gonna occur. And this seems to be the case with the AIM Suite that we looked at earlier this year, so we're really bullish on it. And we really appreciate the work of CDT and Privacy Rights Clearinghouse and the work of the Digital Signage Federation, POPAI. It's very heartening to see these proactive adoption of principles that apply to the industry. I think what's gonna be needed is some verification matters. We don't know whether it's being applied or what's happening behind the scenes. It may not be the case. It's very easy to change a few settings, and suddenly it's a different system. I was asked to mention these types of applications in Canada that we're aware of, and there certainly are a number of others, small start-ups that use infrared that detect the presence of eyeballs. So it's not as sophisticated as Intel's solution, but it's certainly able to deliver metrics about who is looking at it. Being deployed in Tim Hortons to see who's looking at the screens. They don't know who's looking in, but there you go. We have done a number of other -- very briefly worked with other organizations. Notably, the Toronto Transit Commission has deployed video-surveillance cameras. About three years ago, we carried out an investigation and asked them to do audits and more importantly asked them to do a pilot project of technology that is also homegrown in the Toronto area that can defect the faces in real time of video feeds and essentially apply cryptographic techniques and erase them or scrub them unless the feed is needed for forensic purposes if there's an incident, and then it can be decrypted. So that's a really interesting technology. It isn't quite ready for prime time, but we are aware of other firms that provide this type of service of scrubbing the identities of faces from videos, perhaps for evidence, court, you remove people who aren't subjects of interest, and so on and so forth. In this case, it's a paper that we worked on with the Ontario Lottery and Gaming Corporation there. It's a monopoly in Ontario. They manage our racetracks, our casinos, our gaming facilities, and they have a lot of patrons, And they have about 12,000 people who have voluntarily submitted themselves to be excluded from their gaming facilities because they've got a problem, and they voluntarily enrolled. And the biometric -- sorry, the system in place here, very high-level, is that the challenge is not to identify in any way the hundreds of thousands of patrons that come in out of these facilities, but only to detect them, but only to be able to identify in the most privacy-protective way the people who have voluntarily enrolled, and we're really pleased with the early results. It's very, very encouraging, and this paper describes it. Again, I brought some copies, but they're gone, but it's all available on the Website. So, when you really boil these concerns down, they come down to four what I call meta Fair Information Practice principles -- safeguards, data minimization. That's limiting purposes, limiting collection, limiting retention. We want to see these sorts of things happen. Then, on the other side, user participation -- notice, consent, access, redress. These all speak to the ability of the individual to be a participant in the data life cycle. And then accountability, which is not just to the individual, but to regulators, to other business partners and so on, so forth. We want to see demonstrable adherence to standards as they are evolving, and they seem to be evolving quite nicely. This is the last slide. This is the privacy-by-design principles. You'll see that they comprise seven sort of principles. They map very nicely to the meta Fair Information Practices, with three new additions. Essentially, I encourage you to think of these privacy-by-design principles as sort of a robust implementation of the Fair Information Practices. Set really clear leadership roles. Have verifiable, systematic methods to build privacy into your system, and show the results. Show the results. We want to see metrics that actually demonstrate that you're actually adhering. And with that, I think I'm out of time.

>> Manas Mohapatra: All right. Thank you very much. In the interest of time, so, we let you guys go at a reasonable time for lunch. I'll skip any other questions and go to Beth Givens, who's our next presenter. And Beth is the founder and director of the Privacy Rights Clearinghouse, which was established in 1992. And it's a nonprofit consumer-information and advocacy program located in San Diego, California. Beth has participated in many public-policy task forces at the state and federal levels. And prior to her work as a consumer advocate, she was a librarian that specialized in resource sharing. She's also a member of the International Association of Privacy Professionals. Thank you, Beth.

>> Beth Givens: Thank you very much, and thank you, Federal Trade commission for convening this event. Just a couple of words about the Privacy Rights Clearinghouse. Two-part mission. One is consumer education and what I call our Dear Abby role, taking people's questions and complaints and attempting to troubleshoot them as best we can. I'll mention a little bit more about that in my presentation. Secondly, advocacy. We focus mostly on the California legislature, because as you probably know, California has been in terms of privacy a trend-setting state, and a lot of measures that go through the California legislature do end up in other states and even making it into Congressional legislation. Let's see here. Okay. There we go. Thank you. Okay. Fred talked about the FIPs, Harley talked about FIPs. I don't want to dwell on this, but the Fair Information Principles or the Fair Information Practice Principles or the Principles of Fair Information Practices, all variations of FIPs, are the building blocks of privacy public policy and laws and regulations, but also organizational policies. I use the full FIPs, all eight of them, when, for example, evaluating a company's privacy policy. They are not new. They date all the way back to the early '70s, and there are several variations of them. The OECD FIPs I would call, you know, all eight, the set of them, constitutes full FIPs. Also, the Department of Homeland Security, modernizes a bit, 2008, gets into -- excuse me here -- gets into adding the word "auditing" to "accountability," which is something that Fred brought up. Very important to have that component. And that's all I will say about those, because I think most people here are quite familiar with the FIPs and their importance in terms of guiding privacy policy and, also, giving individuals knowledge of what they can expect or not expect in good privacy policies. I'll just talk about the two standards, which have also been mentioned, the Digital Signage Federation and POPAI, the -- let's see -- Point of Purchase Advertising International. I have been asked to compare two privacy standards. The first one, issued by the World Privacy Forum in late February of 2010, is based on the full FIPs, and it's in a report called "The One-Way-Mirror Society." And it's really the first I think sound of alarm, call of alarm as to the significant potential for privacy invasion and abuse with digital signage, especially when we move from detection over to recognition. The Digital Signage Privacy Principles were signed on to by seven additional groups, including ours, and they incorporate, as I say, the full FIPs. They call for real-time availability of privacy policy, readable labels that clearly disclose purposes. And I want to say a few things about that, because notice is important, but I've heard really just text, and I think English language. We also have to think in terms of ADA, but also languages other than English, and, of course, other capabilities that individuals have. So, the other thing I should

say about notice is what I call the euphemism factor, and that is notice that skirts the issue and kind of makes it a happy face, a feel-good. Video surveillance -- you may have seen these for your safety and to keep our shelves stocked, that sort of thing. That's what I call the euphemism factor. Excuse me just a second. Yeah. Very important that signage not be in areas where there might be any kind of HIPAA implementation -- supermarkets with pharmacies, for example, change areas, locker rooms, areas where children play. And it's also important, I think a very, very key principle is deletion. And in the World Privacy Forum principles, there is mandatory deletion after 14 days of collection, immediate deletion for children under 13. And that's an important distinction I think with the Digital Signage Federation's principles. Also, in the World Privacy Forum principles, only the subject is able to view signage. I haven't really heard any discussion on that. That might be interesting. And then, of course, accountability. A key FIP, who is responsible for compliance. In the interest of time, I have skipped over the particulars. And I'll just go talking about the -- let's see here. Sorry. It's hard to see, actually. Oh, here we are. One of the key challenges for privacy protection is something that Fred mentioned, and that is that it's essentially invisible to the individual. Unless there is robust notice in more than just the entryway, the individuals are really not going to know what is happening in terms of their privacy. Sorry about this. So -- Gee, I'm sorry. It's hard for me to see here. Let me get myself straight. It's also difficult for the individual to detect when there's actually noncompliance occurring. That's one of the themes I think of privacy that we've observed over the years is that individuals don't really know in many cases that their privacy has been abused or violated. We gets individuals coming to us Saying, You know, "I'm not getting a job, and I'm because I'm not renting that apartment, I'm wondering if because." And I think that enters into the situation here, very much so, with digital signage, the invisibility factor -- something that Fred also talked about. One of the key points of criticism that I have of the Digital Signage Federation's principles or discussion is -- And I'll read it. It's the notion that if you are uncomfortable with digital signage, just don't shop in that mall or that store. Don't go into that supermarket. I think that's really unfair and unrealistic. Also, it requires individuals to lose out on opportunities they might otherwise have. Let me read to you that statement from the document that does trouble me. "Notifying consumers that a particular signage unit collects information gives consumers the opportunity to avoid that signage unit." Again, I think that that notion needs to be questioned, and I do find it quite troubling. I think what also concerns me the most, retention, of course, is very important, and I think that the World Privacy Forum principles deal with that very well in terms of calling for deletion. Another area, of course, is dealing with children and deletion and just being a very straightforward statement of policy with hard edges. Not a great deal of wiggle room. And, of course, that's the nature of the beast. On the one hand, you've got a consumer-centric statement at the World Privacy Forum, and the other is an industry-centric statement, the Digital Signage Federation, and it's kind of the nature of the beast of both. But I would say that the lack of specificity in the DSF principles is an issue of concern to me. Finally, function creep. When you are successful in your facial-detection application -- And let's just say maybe visitors ship to the bar. Goes up "X" percent, or sales go up "X" percent. I think the temptation is a very real temptation to move into the next step, which is facial recognition. And one of the key principles, of course, there is opt-in versus opt-out. Individuals can opt out of going into a location where digital signage is used, yes, but again, they're missing out on opportunities themselves. Okay. I was asked to talk about the shortcomings of self-regulation. And I think that's probably fairly self-explanatory, but I will say from the point of view of an organization that works directly with individuals is that individuals are really ill-equipped to identify and report noncompliance. It's very difficult. One reason is that privacy of uses are often invisible. And, by the way, we have technologists to thank for making the invisible visible, and just I wanted to do a shout-out to those technologists who have worked so hard to make these things public. Another problem with self-regulation, lack of benchmarks with which to evaluate and a lack of auditing mechanism for the most part. And let me just skip to an important consideration. I am not aware of many, if any, voluntary code processes that have included the participation of consumers or consumer representatives. I think insufficient public awareness, which I will close on, is also key. And that's one of the things that we do as individuals is work, try to describe to them, "Okay, if you have a privacy complaint on such and such an issue," we lay out the regulatory landscape for them. In this case, you would go to the Federal Trade Commission with your complaint. In this case, it would be Department of Health and Human Services. In this case, the new CFPB. That's a very complicated landscape for individuals, and it's one of the key things that we do, but it's also difficult in terms of consumer education and public awareness. I can be reached on our website. And let me just close by again thanking you, Federal Trade Commission, and appreciate this opportunity. [Applause]

>> Manas Mohapatra: Thank you very much, Beth. I'm sorry. We had a lot to get through today, and I was hoping for more opportunity for facilitated discussion. And in light of the lack of time, I was just wondering if people on the panel had, you know, closing thoughts, if there are things that they'd like to respond to that have been spoken. If we just want to start with Brian, if you do, if you want to chime in or move forward. Harley?

>> Harley Geiger: So, Beth, I wanted to respond to a couple of the things that you mentioned in your presentation. You had said that the digital-signage privacy guidelines, which I actually had written, lack specificity. But I actually think that they're very extremely detailed. I have them with me, and I invite anyone to read them. But it covers specific categories of synonymous data and direct identifiers that are not covered in other codes, including the POPAI code. It gives specific notice language. I also would like to hear more on how your troubles about having an opt-out notice for facial detection. So, that was a very difficult thing to sort of implement, but we had three layers of notice, and it's difficult to think of another way to notify consumers and have them avoid facial detection without simply outright banning facial detection in a given area where it might be otherwise appropriate. I mean, obviously, a locker room, HIPAA-covered entity, might not want to have facial detection, but a place like a mall, without banning it, how, if you don't have a privacy policy, a notice at the perimeter and a notice at the device, are you to implement and opt out? And, lastly, still on notice -- in fact, the digital-signage privacy guidelines are quite protective when it comes to notice. The POPAI guidelines, unless they've been changed, I believe they just have one notice covering an entire establishment, whereas for our guidelines, we have one, and each device, it's actually doing the collection.

>> Beth Givens: Yes. Facial detection, I guess it boils down to the creepiness factor. Facial detection I actually think would be offensive to a considerable number of people, and I think the notice Is actually incredibly important. I was going to suggest, and I will now, that in addition to the text notice -- By the way, again, other languages, ADA --

>> Harley Geiger: It is ADA-compliant. The guidelines are ADA-compliant.

>> Beth Givens: Yeah, yeah. But I was thinking it might be useful to have a QR code, as well, so the -- I think, what, those can accommodate, what, 4,200 or so characters? That would be like an eight-page privacy policy. For those with smartphones, could learn even more. And I think with digital signage, it's enough of a new thing for many individuals that they might benefit from learning more and just capturing it on their smartphones with a QR code.

>> Harley Geiger: So that would ease your concern about the opt-out notice, to have a QR code on the notice?

>> Beth Givens: Actually, it doesn't ease my concern about, you know, just don't go into those establishments if you are uncomfortable with facial detection or facial recognition. That I do think is kind of a cop-out. And one of the things I wanted to say -- and, again, I'll say it now -- brilliant minds have gone into the development of these technologies, both on the science and technology side and on the business side. I would love to see those brilliant minds working towards a creative and effective notice mechanism for digital signage. I think that would be very, very constructive.

>> Manas Mohapatra: Just one follow-up question, just related to the facial detection. I'm just wondering what the panelists thought about in terms of should there be some prohibited uses of facial detection -- for example, differential pricing based on somebody's age or their gender or, you know, to the extent that ethnicity is possible? Is that something that, you know, should come out? Is that something that should be prohibited? Is that something that industry should just be aware of? I'm just wondering if people have thoughts on that?

>> Harley Geiger: Well, yeah. We would have major concerns about it, and, you know, we would have to gather more information and talk to the company. But on its face, that sounds like something that we would want prohibited, sure.

>> Brian Huseman: I mean, and that's an issue that's -- I mean, not just with facial detection, with Internet purchasing, I mean, with a bunch of different technologies.

>> Beth Givens: I'm very interested, Jai, in your presentation, and I do see the potential for detecting, you know, micro-expressions being used in a manipulative way, and, you know, the privacy policy, and, you know, the usage application. The boundaries around the usage of detection of micro-expressions I think are very, very important. I think we have a panel of, you know, best practices and good, responsible actors here, but I can see these technologies being used, especially the micro-expression detection, in ways that are manipulative.

>> Jai Haissman: Yeah, that's a real concern that we addressed very early on in our formation. Our approach is to support the use of our technology that's geared towards understanding rather than persuasion, creating an experience, but not to deliver a message. So we're selective in who we license the technology to and for what applications towards that purpose, and I'm appreciative of really good guidelines in order to help us shape our policy. I expect that this will be an ongoing dialogue, both with our customers and both our end users and with our customers, and we're very much interested in the blue-sky applications for this technology. But there's many gray areas that we're also sensitive to, such as if we can more accurately assess hostile intent in a security line at an airport, then arguably that has a high value. But then are people uncomfortable in being assessed for their emotions while in an airport line? Of course. And so weighing this cost is tricky, and, You know, we'll continuously seek council on trying to stay on the light side of these applications.

>> Manas Mohapatra: Fred, if you --

>> Fred Carter: I just want to pick up on the point that you made and just earlier about function creep. You know, as good as facial recognition is becoming, I don't -- and understanding that facial detection is an essential step towards facial recognition, I think the real risk is not so much that you're gonna be identified through your face, but you'll be identified by the collection of other information about you. It could be your cellphone. It could be an RFID loyalty card. It could be other behavioral aspects, you know -- your gait or, you know, your license plate that you parked in the casino car. I mean, all of these things are brought together. It's really the linkage of who you are to other information about you that allows you to be identified, where the privacy issues become really engaged rather than facial recognition.

>> Jai Haissman: And that's, again, why CDT recommends not just having facial-recognition specific legislation. All of the ways that Fred just mentioned that you can be tracked and identified will still be viable if you just have facial-recognition specific legislation, so it needs to be wrapped into consumer-baseline privacy legislation that covers all of those categories for it to be effective at all.

>> Brian Huseman: Yeah. I mean, Intel agrees with that. I mean, we've long supported a comprehensive federal privacy legislation based upon the Fair Information Practices, and we think that that would go to address a lot of these kind of function creep and other issues we're talking about.

>> Manas Mohapatra: All right. Well, there's obviously a number of issues that this panel has raised, and unfortunately, we don't have time to address them all, but thank you all very much. Appreciate it. For the audience, we will be back at 1:15 with Commissioner Julie Brill with remarks, and we have a great second half of the day. So thank you all very much. [Applause] I got to check you guys out.

>> Male Speaker: Exactly.

>> Beth Givens: Well, yeah.