DEPARTMENT OF HOMELAND SECURITY

PRIVACY OFFICE

PUBLIC WORKSHOP CCTV: DEVELOPING PRIVACY BEST PRACTICES

MONDAY, DECEMBER 17, 2007

Hilton Arlington

Gallery Ballroom

950 North Stafford Street

Arlington, VA 22203

PANEL ON TECHNOLOGY PERSPECTIVES

MR. SAND: Thank you very much. My name is Peter Sand. I serve as director of privacy technology in the DHS Privacy Office.

This is the first panel of the workshop. It's talking about the technology itself. What we hope to do in this panel is establish, kind of, the baseline, to establish how the technology works and what it does now and what it can do in the future. So, we're hoping that this first panel tees up the rest of the workshop.

And, as our panelists talk, I'd like to ask you to think about questions you may have about the technology itself, if there are any outstanding questions or things you've been curious about. We're going to have about 15 minutes at the end of our panel to have open questions for the audience. So, if there are things that you think were raised during the technology discussion, or questions you're bringing to this workshop, we hope that you'll ask them during those 15 minutes so that everybody can have a better understanding of what the technology can do now and what it'll do in the future.

We're going to talk about two separate things in this first panel. The first is the technology itself, the hardware, and also the analytics, what you can do with the data once it's generated by the video hardware, and then to talk about some of the issues and challenges.

I'm going to ask the panelists to just quickly introduce themselves, give you an idea of their background, and then we're going to their presentations.

So, Sam?

MR. DOCKNEVICH: Hi, I'm Sam Docknevich. I work for IBM Global Services. I work in our services group, and I lead a practice that focuses on physical security solutions. So, we work with clients in all industries, in all sectors across the U.S., to help them design solutions to enhance physical security for their people, their sites, and their IT infrastructure.

MR. DAVIS: I'm Larry Davis. I'm with the computer science department at the University of Maryland in College Park at the Computer Vision Lab. We develop computer programs and systems for analyzing video for video surveillance.

MR. SAND: Larry?

MR. STRACH: You've got two of 'em. Good morning. My name is Larry Strach. I'm a vice president of engineering for Duos Technologies. We are actually an integrator. We install CCTV systems, mainly using intelligent video.

MR. HOFFMASTER: My name is Randy Hoffmaster. I'm director of engineering for Epsilon Systems Solutions, prime contractor for three DHS contracts to conduct pilot projects.

MS. KING: I'm Jennifer King. I'm a researcher with U.C. Berkeley at the School of Law, and I focus on social issues in technology, such as privacy, as well as human/computer interaction.

MR. SAND: The format for our panel is that we're going to have short presentations, followed by a longer period of discussion on this topic. I will pass the control over to you, Sam.

MR. DOCKNEVICH: Pass the baton. So, I have the daunting task of, in 8 minutes, to make you experts on CCTV, which I don't think is possible, even for me, but I will do my best. So, I'm going to give you an overview of the technology. As Peter said, the first step is to just give us a foundation so we can build on for the rest of the conference.

We look at CCTV as a complex system integration project, because that's exactly what it is. You have devices out at the edge, cameras, and it doesn't necessarily have to be just cameras, it can be audio sensors, badge readers, door readers. There's a variety of other edge devices that you can use to enhance physical security. So, you have cameras at the edge, then you need to do something with that camera input. It has to get back to your user, right? If you look at the left of the slide, we have the cameras; the right of the slide is the user. And the user could be somebody in a command center, it could be a police officer in a car who would like a feed from a school or another site where there might have been an incident. And so, you've got to get the video to that person. A core part of this technology or this solution is networking. The video has to get back. One of the advantages of moving to digital is that

you can run your video over the same network that you run your business applications, your SAP applications. Everything can run over one network.

Now, of course, that introduces some challenges, in and of itself, because now you've got bandwidth issues you've got to contend with, you have quality-of-service issues, you want to make sure you have your video -- your users, on the right -- get their video when they want it with the quality that allows them to do their job, whatever that is.

In the center, you're going to have some type of hardware device. It could be a DVR. A lot of us have DVRs in our houses now. It could be a network video recorder, rather than just a digital video recorder, which then allows people to access it from afar.

Trends nowadays are for the video storage and the application that controls your cameras, encodes the video, sends it to your archive, and allows it to be played back, to run on standard enterprise servers -- Windows-based, Intel-based servers -- more scalable, more expandable, you can run multiple applications on the same group of servers. So, we're, kind of, moving away from the proprietary-device-type of appliance to record video.

And then, of course, you've got storage. Storage is very important, because now you're storing, probably, something that's very critical to you. You want to be able to have it for whatever your retention policy says you should have it, so that video needs to be backed up, you need to consider it like you would any other piece of enterprise -- valuable enterprise information.

In the olden days of analog, co-ax, guns, guards, and dogs video, the only thing people were really worried about was what was above the line, the cameras and the video equipment. Our approach, and what our clients have pushed us to help them with, is the integration. So, think of video surveillance CCTV as another enterprise application that plugs into your enterprise network. The management of the information is very important. The storage of the information is very important. The integration is very important. And also, the support. Now, one of the areas we really work with clients is to help them design the support structure so that this critical piece of enterprise technology is available when the end users want to use it. Whether it's law enforcement or first responders, it has to be available. So, it requires a different level of IT integration, and it really requires IT discipline to ensure that your systems are available when you need them.

So, this always reminds me of what Spiderman said, “With great powers come great responsibility.” Well, we have new technology, which has benefits, right? Digital technology can reduce operating costs. Now you have the ability to send your video with a keystroke rather than going to your rack of analog VCR tapes, and pulling a tape off, and putting it in a box, and FedExing it to somebody or making a copy. But, with that, becomes a lot of challenges, so you have to think through your policies and procedures for your video. How long do you want to retain your video? We're working with several of the police

departments right now where they're trying to align their video retention policy with whatever is in place for their State and local government, to make sure that they have consistency across the different organizations.

Not only are cost factors important -- I don't have it in this presentation, but when you start recording video and storing it for hundreds of cameras, and you want to have high-quality video, you're not talking gigabytes, you're talking terabytes of storage, and that's only for 30 or 60 days. So, you really have to think through, Do I want to look at high-quality video, but maybe I only want to store something that's less quality? You really want to look at other technologies that are going to help you deal with this flood of information. With great power comes great responsibility. Now that you can record and put cameras, and have those cameras, maybe, wirelessly connected back to your central site, you can have more and more and more cameras. Well, does that model really scale? Are you really going to get value from that video, or are you just going to be recording more and more terabytes?

So, when we work with clients, our experience has shown that there's really eight key design points for designing your video, your CCTV solution. First of all, you really have to have a partnership between the physical security team and the IT department and also the users. When we did conferences with customers, years ago, we would find that the audience would be filled with just the security folks, the guns, guards, and dogs guys. Now, IT comes together with them to a conference, because it's an IT-based network application, you've really got to have both teams together to design it.

You really need to build on open and standards-based architecture. The worst thing in the world is to deploy a solution to support your initial deployment of 50 or 100 or whatever cameras, and then find out that it doesn't scale because you've built yourself a proprietary solution. When we work with partners we're always looking for partners whose products have open standards, open APIs, where we can integrate and then customize for our clients. You have to have enough network capacity and performance. Very, very critical. High-quality video, hundreds of cameras chews up a lot of network bandwidth. If you don't have it, there is a delay back to your central command site, and you want that video available when it happens, not 3 or 4 minutes later.

What's more important now, or maybe as important now, is optimizing the infrastructure back end -- your servers and your storage -- to make sure that you have enough performance to enable you to play back the multiple streams of video to multiple users, many of which may not be in your organization if you have a mesh network, and you have the fire and the police and the schools and everyone connected together.

You also need a different level of security, now that everything's digitized. Things can disappear from your archives unless you set up access rights properly. Most of the video management systems allow you to assign access rights to an operator; he can see certain cameras, he can record, or he can pull video from the archive for only certain cameras. Very,

very important that you have the right security in place for the video, so that it doesn't get accessed or transported, other than when you want it to be.

Analytics -- I'm going to talk a little bit about analytics in the next section, but we find that, as you record more video, the question is, How do you get more value out of it? And we're going to talk about some other technologies that will allow you to do more with the video than record it. And you have to support the application to ensure adequate uptime. That's extremely critical, that you have the right policies in place. And establish a data-governance policy and practice for the retention of your video, the security of your video, so you can ensure the privacy and the security, when you need it.

And that ends my 8-minutes -- how did I do? Well, I went 1 minute over. But that's my 8 minutes. Now you know everything you need to know about designing a solution for video surveillance, and I'm going to turn it back to the rest of the panel.

Oop. No, now we do a panel discussion, right? But I'm next.

VOICE: Oh, you have more?

MR. DOCKNEVICH: Right –

VOICE: Oh, I see.

MR. DOCKNEVICH: I have analytics. We were supposed to have a discussion of technology.

VOICE: Okay, I'll reverse those two.

MR. SAND: So, one of the things you mentioned was the network capacity. Are you finding that there actually is the capacity to handle the kind of video that people are asking for?

MR. DOCKNEVICH: I just did a briefing last week for the city of Norfolk, and we had IT, and we had security people, and we had all the end users -- parks-and-recs, and the libraries, and the school system. They were pretty shocked by the amount of network bandwidth that you really need, to transport 30 frames per second, four SIF, which is high-quality video, 100 cameras over their network. They were shocked. They don't have a network that could support that. So, I think, when you're designing your solution, you always start with the user perspective, right? Remember, we're talking technology up front, but really you want your CCTV to provide value to somebody. You're going to do something with it. So, figure out what you want to do, first, and then architect a technology to deliver value to that end user. But, no, nobody really understands until they see the facts and figures about how much bandwidth you really need.

MR. SAND: So, do you find that shows up in the initial conversations, or are you the one to raise that issue?

MR. DOCKNEVICH: When we meet with a client, you have to raise these issues up front, or you will be in a position where they'll implement a solution that just doesn't work, or they'll get sticker shock when they price out what they need to buy to support what they think they need to deliver -- 100 cameras, 30 frames per second, high-quality video. If you don't consider these issues up front, you're not going to have a solution that meets your end-user needs, that's affordable, and that's actually deployable. Early design and early planning is key, just like with any other IT project.

MR. SAND: Any other thoughts on the technologies? So, is that an initial show-stopper, and is it possible that during the first conversation, the client could decide, Well, forget it?

MR. DOCKNEVICH: Well, I try to work with our clients -- when I'm with a client -- I first pick a project that provides value to somebody, and show people that we can deploy -- and the city of Norfolk's a good example, because they want to deploy their video surveillance, and they haven't done any video surveillance, any CCTV, and they are looking to deploy a solution to provide value. So, the question is, What are the threats and the vulnerabilities you really want to protect against? Define a solution for that situation, show value to your end users, and then you can then build upon that. So, don't try to do too much in the beginning. Do a phased approach and provide value.

MR. STRACH: I'd like to add something to that. It's been my experience that anytime you deal with an IT department, whether it be in the public or private sector, it's an uphill battle to get them onboard. Usually, they do not like anything on their backbone, clogging their network. Video is a huge, huge bandwidth hog and creates a lot of issues and problems with the IT department. And, to build upon what Sam is saying, it's very important to get them onboard up front, because you could put the most beautiful system in place, go turn it up, and then you find out that you're getting one frame per second. For those who don't know what that is, when we talk about 30 frames per second and one frame per second, if you look -- go back to the old black-and-white movies of -- silent movies, and you see the people walking, that's roughly, I think, eight to ten frames per second. The human eyes see, generally, 25 to 30 frames per second. And if you're trying to do video analytics or forensics or anything along those lines, if you don't have good, fluid motion or good resolution, or the end user's expecting something more fluid than what he's actually getting, then it doesn't matter what you put in, then it's an uphill battle, then, just to go ahead and get the thing right and get everybody on board to accept the end product.

MR. SAND: Just to jump to the other extreme for a second, have you heard people asking for a full integration of video into their existing IT systems, and say, Well, now that we can have video, we want to put it everywhere. If it's running on our network, we can just grab it and feed it into our existing applications?

MR. DOCKNEVICH: You mean, integrating video with other applications.

MR. SAND: Yes, I was just thinking of how you were talking about having the video itself run on the network. You talked about integration, in general. I'm wondering if there are people out there that really want to push the envelope, in terms of how they can use moving pictures in their application or in their environment.

MR. DOCKNEVICH: Well, we really haven't seen that. Where we're seeing more emphasis is to provide a more holistic view of physical security throughout the organization. So, clients are looking, How can I integrate video with access control and door locks and badge readers and their HVAC system, into some command view that will give them one picture for everything that's happening in their site. But how to integrate video into other applications -- business applications, I'm not seeing that yet.

MR. SAND: Just one other question about the network. I'm wondering if you get a sense of what people expect about where the video goes. Like, we're going to set up the camera and turn it on; where does it go? When you talk to your clients, do you have a sense of what their understanding is of where all this information goes once the red light's on?

MR. DOCKNEVICH: Well, that's a good question. The first use of video is just to record an archive so you have it for forensic use after the fact. I was in a hotel the other day, and in their check-in guide, they had a little disclaimer that said, The video we're recording is for forensic postevent analysis only. It's not hooked up to the police department. And if you're mugged on camera, it's not going to send a police officer to rescue you. Pretty interesting. But then, when you find that, once you are archiving video so you can do something with it after the fact, you want to start being more proactive, rather than reactive. So, then you start looking for, How can I integrate the video into a more realtime system so I can respond quicker and allow first responders to get information about the event that they're responding to?

Once you see value, you're going to look for different ways to use it. But, you're right, people haven't thought about these issues and questions, and that's why a lot of the preplanning work is much more important than the technology that you're going to deploy to solve the problem.

MR. SAND: One last question on the hardware. Are we at the point now where people can realistically expect that they can sit in front of their computer, tap keys, and see video from all different places? They can just hit this button and see that room, and that button and see that location? I mean, can the technology itself support that type of TV-show fantasy world?

MR. DOCKNEVICH: Well, if you have cameras, you can see them. If it's IP-based, you can then transport it over your network. Most of the applications are Web-browser user interface, so anyone with a Web-browser connection could then access the video. The most important thing, though, is to ensure that you have the right access rights on your cameras, both for realtime and archive viewing, so only the authorized people can then view the

video. That's really important. And we work with companies to implement the access rights, implement a directory structure of users, just like you would if you're going to give people access to SAP or customer information or critical HR or business information. Think of your video as another valuable piece of corporate information that you have to protect, just like you would the rest of the information you protect.

MS. KING: I have a quick question. Sam, what advice do you give to your public-sector clients when you do an installation and you find they don't have the in-house expertise or the personnel to support it?

MR. DOCKNEVICH: Well, we try to deal with the day-two support issues right up front. If they don't have -- and when you're talking support, are you talking about break/fix support and keeping the systems up and running and –

MS. KING: Yeah, exactly.

MR. DOCKNEVICH: You want to get IT involved, as Larry said, early on. That's one of my key design points, is, establish the partnership between physical security and the IT infrastructure. And you have to think through your day-to-day support plan, just like you would for any other critical application. I mean, IBM offers services for day-to-day support. Clients approach it differently. Some clients want to do it all themselves, so they'll have a Help Desk internal that will take the call, and then they may use the vendors of the particular pieces and parts as their second- and third-line support. We provide that service to our clients. It really depends on the organization. But the key issue is to think through your support before you put the system in place and have something break when you don't want it to break and you can't see an important critical event in your geography.

MR. SAND: All right, let's move on to the analytics. And the issue here is basically, you set up your cameras, it starts to feed video, What can you do with the stuff that you're receiving? What can you do with the data?

MR. DOCKNEVICH: Well, as I said before, again, we have 8 minutes to make you experts on video analytics. So, the basic problem, I think, is pretty easy to understand. People find value, so you want more cameras, and then you want more cameras, but that model doesn't scale very well.

First of all, how many people -- what's your ratio of operators to screens to cameras? You know, how many screens can one person look at? And then, how many cameras can that number of screens support in the background? It just doesn't scale.

So, say you have one operator, and he has three screens, and each operator has 60 cameras. Then you're going to be cycling 60 cameras through two or three screens, so how many seconds of a minute, or minutes of an hour, is that operator actually looking at that screen? So, the other approach is to have a giant wall, where that operator has 20 screens and 60 screens. And that doesn't scale, either.

And it's been shown that effectiveness of operators looking at screens really, even if they're trained and they're professional, becomes very ineffective, and you miss most of everything.

Also, the model doesn't scale, from a cost perspective, right? Who wants a wall of Barco screens that costs a gazillion dollars, or monitors that are very expensive, or the people? Most of our clients want to put their labor budget in actually using their enforcement for catching the bad guys, not watching screens. So, because of the proliferation of cameras and the scalability issues, organizations have been looking for more ways to get value. And that's why video analytics has been such a hot topic. It's probably the hottest topics in CCTV.

While it's the hottest topic, it's also probably the most misunderstood topic, because people think, it's a panacea, which it isn't. It's not a program that looks at people and makes a determination about their behavior, if they're suspicious or not. You know, what's suspicious behavior? I mean, I see some of my colleagues out in the audience, and they look pretty suspicious to me right now, but how would the camera know that? It doesn't.

There have been a lot of installations of analytics which have not reached their goal -- meaning, providing value -- because they've had too many false negatives, they haven't had the right number of positive alerts, and customers have pulled the systems out, because the technology was too immature, and it was too new, and it didn't provide value.

Remember, also, that analytics sit on top of your CCTV solution, right? You have to have cameras and encoders and storage and networks, and then you also take the feed from your cameras and you put it through an engine that does something with that video, ingests the video and provides some type of information to an operator. So, there's an additional expense.

Now, you can imagine, if you're looking at every frame of video and trying to pull data from that video, it's going to require horsepower, right? You're going to need servers, you're going to need software, you're going to need a database to store the information on. So, it doesn't come without cost. And, because of that, many of the first uses of video analytics didn't provide value in proportion to the cost.

We've spent a lot of money, in IBM, developing our own technology, and I'm not here to do a commercial, but our approach has really been to work with customers, on a consulting basis, to find uses for analytics that will be valuable. So, you need to show value. You're going to catch a bad guy, you're going to decrease the number of operators watching screens so you can deploy more people out into the field. The cost of the solution has to equal the value.

So, how do analytics actually work? There are some basic capabilities of most analytic packages, and they really work by looking at movement in a frame. Well, first of all, they only work with fixed cameras. You can't use analytics on a PTZ, pan-tilt-zoom camera, unless it's locked into a fixed position. And it's going to ignore the background, because the

background is meaningless. Background is something that doesn't move. So, every object in the frame that's moving is assigned a value to it -- size – think of the attributes of the object, speed, direction. And, because you can then code the attributes of an object, you can start doing things like, Tell me when an object crosses a line. Tell me when I see this object twice. Tell me when an object that was moving is stationary. So, you have these different capabilities. And, typically, a video analytic system generates metadata, -- metadata is data about the data -- which it then stores in a database, and you can use this data to do two things: send realtime alerts and do postevent forensics. And both of them are very important.

From a realtime alert capability, what do people use analytics for? One of the most common is an object left behind. Typical scenario, you have an outdoor plaza, and you have people milling around. Wouldn't it be nice to know if someone left their backpack or their suitcase? Could be filled with schoolbooks, it could be filled with something else. But an analytic can send that type of an alert, because it'll then categorize the person as a moving object that has a backpack or a suitcase as another object. You can set up an alert to tell you when that object has been stationary for more than 60 seconds.

The key to this, of course, is camera placement. We've worked with several outdoor projects with law enforcement, and you find that one of the key characteristics of whether or not your analytic is actually going to alert you when you have a package that's left behind is number of cameras, field of view -- you can't have one camera on a building at 20 stories high and covering a plaza where you can't even read somebody's face. So, again, realtime alerts can provide that type of capability of noticing when the same car went around your building three times. An analytic can provide that.

For monitoring operations in a port, in sending alerts when people are in areas where they shouldn't be, you can use an analytic to provide that type of realtime alert.

You can also use an analytic to provide postevent analysis. A good example is, you have an incident, and you know the perpetrator, the bad guy, is fleeing in a yellow vehicle. Well, you have cameras, and you have all the video, how do you go find all the yellow vehicles? Well, you can look at all your hours and hours of used-to-be-tape, and now it's digital files. – Remember that an analytic has already coded every moving object, so it knows what's yellow and what's been moving, and so, you can go to your keyboard and do a SQL-like query, and it'll then pull up the video and show you all the yellow objects. And then you can define the size of the object to differentiate people and cars and buses. So very powerful for postevent analysis.

Customers are using it in public safety; they're also using in retail. A lot of different uses. The key is to make sure that the use of the video, the value you're going to get out of it, is equal to or greater than the cost of deploying the solution. So, again, it requires a lot of

consulting up front. Rather than focus on the technology, focus on the problem you're trying to solve.

And, lastly, video analytics can also be used to preserve privacy. We talked a little bit about using access rights, or assigning access rights to live camera views and archive camera views. That helps preserve privacy, because you're limiting the users who can actually look at the live cameras. You can also use analytics to preserve privacy. This is a good example. We've used analytics to mask out the faces of the people in this screen. That's a very typical use, especially if you're going to court -- police department's going to court, and they have a piece of video, and they want the suspect, but they don't want all the other faces to be seen. And you can also use the video to define and build a hierarchy of what video is then available to the end user, so you can then lock in and ensure privacy -- preserve privacy.

Okay, my 8 minutes is up. Well, I did 1 minute over on that one, too, but you are now experts on video analytics, and we'll throw it back to the panel.

MR. SAND: One of the things that struck me in your talk was the difference between the hardware itself and the analytics, in terms of when you which part you want to use to decide what you want to do. It just struck me that one of the first questions you would ask in an IT project is, What do you want to know? What is it that you want out of this? What is it that you want to know? And it sounds like a lot of that, in the CCTV world, is actually about the analytics. So, if you end up picking out a really cool camera, that can pan and tilt and zoom and turn upside- down and go infrared, and all that kind of stuff, that may actually prevent you from being able to analyze the resulting video. Is that something that you find yourself talking a lot about to your clients?

MR. DOCKNEVICH: Well, absolutely. You're going to think that's the only thing I talk about is planning and upfront design work, but it's really critical. Where clients have succeeded with their CCTV implementations is when they have a vision. And if they don't, we like to help them build their vision, to get the users together and decide, What do you want to do today, tomorrow, the next 3 years, the next 5 years? Where do you see this evolving? And then, you build the technology solution to provide some value up front with your first-phase deployment, without allowing you to back yourselves into proprietary or buying technology that's not actually going to deliver.

In the example that you used, Peter, you may have very expensive pant-tilt-zoom cameras and all the great capabilities. For analytics, you might want to put less-expensive fixed cameras. But you can add them later. That's the beauty of a scalable system, you can add devices to it with relatively little configuration. Back in the old days, you had to send people up with bucket-trucks and configure cameras up on poles. And now that everything is IP-based, your operators, wherever they are on the network, can configure all your end devices.

MR. HOFFMASTER: I have a question for you, Sam. Of all the analytic packages -- there are objects left behind, objects moved, loitering, trip wires -- which ones, in your experience, do you think are the more reliable?

MR. DOCKNEVICH: And reliability is determined how?

MR. HOFFMASTER: Well, with proper programming, proper setup, which ones have the most reduced false-alarm rates?

MR. DOCKNEVICH: Well, that's a hard question to answer.

MR. HOFFMASTER: I realize –

MR. DOCKNEVICH: But, our experience has shown that typical when you deploy an analytic solution, there's definitely a learning curve. So, you can draw a graph. And so, you deploy it -- and we've done many pilots with public safety and retail -- and you deploy it, day one, and you think you've decided on the business scenario that you want to use your analytic to provide information about, and then you put it in play, and then you find out you either get no alerts, or you get a thousand alerts. And then you tune it.

Typically, we're not deploying individual capabilities. Customers are taking the capabilities and putting them into a solution that solves a specific problem. I mean, one problem we were working with a law enforcement agency is, they wanted to understand what cars went around their building more than two times, because that could be an indication that somebody was going to do an evil deed. So, then we used some of the basic capabilities that we've discussed, and packaged them, and put them together into a solution. But then there's a learning curve. So, you put it in play, and you tune it. And if you do that, you will get to either the level that you're getting the right number of alerts, the least number of false positives, and you'll get some value out of it.

It reminds me of a story, back when I was in systems management consulting, and you're installing a systems management tool to monitor your network and your servers, and the IT manager says, I want every alert from your tool. I said, You're sure you want every alert? And, you know, the first day they got 57,000 alerts, and he goes, Well, maybe I don't want every alert. Maybe it's not important enough to know when the temperature goes up a tenth of a degree, but maybe tell me when it goes up 5 degrees. So, then you start tuning that until you get to the level where you get -- you're getting value from it.

MR. HOFFMASTER: Thank you.

MR. SAND: Thank you very much. Larry?

MR. DAVIS: Thank you.

MR. DAVIS: All right. I'm going to talk some more about video analytics and computer vision technology for video surveillance. Sam did a very good job of giving you an overview

of what current capabilities are of video analytic systems. So, I'll go through them very briefly here again. They're probably very effective at intrusion detection for perimeter and facility security. The use of all these techniques depends a lot upon how complicated the situation is in which they're deployed, so it's very hard to get video analytics to work in very crowded situations. It's one thing to be able to, for example, do change detection to see that somebody left a package behind in an area where there are very few people and the package is in plain view, but it's very difficult to do under normal operating conditions in an airport or a train station, where there are people moving all over the place and the package might have been left behind by somebody putting it underneath a bench or putting it into a trash can, you know, which, maybe, is even a more dangerous situation. Sam was talking about tuning, essentially, the detection thresholds controlling the false-alarm rates from all of these video analytic procedures. It's a very, very difficult thing to do, because the range of operating conditions is really fairly limited.

We talked about detection of anomalous or suspicious behavior. Again, you might be monitoring for people who drive erratically -- well, if they leave the road, they're probably driving erratically -- or people entering through doors that you've designated as exits- only. These are things that you can do. And person authentification -- of course, this is the source of greatest concern when you talk about privacy concerns –you know, face or iris recognition, or other self-identifying information -- the license plate on a car that people are driving -- so, there are very, very good systems that detect and read license plates, and these are in place in many systems around the world.

All right? All right.

So, how do they do it? Well, I think that Sam, once again, gave you a pretty good overview of what the basic high-level technical approach is to video analytics, but there are really two parts of it. One is motion detection, and the other has to do with very specialized operations for analyzing people and vehicles, which are applicable whether or not you have a camera that's moving or a police car moving through the city. All right? So, there are very powerful techniques that haven't yet been integrated into commercially available video analytic systems, although some of them are in your digital cameras, being able to, for example, find all the faces in an arbitrary image, without the camera having to be stationary. This, in fact, was a camera – an image taken be a stationary camera, but it works just as well for an arbitrary frame from any video system. And if you've bought a digital camera in the last 6 months, then it probably has capabilities in it for red-eye reduction that uses this type of human face-recognition software, or even focuses on faces for family pictures automatically.

All right. What are the types of future applications that people are working on, mostly in laboratories like mine, around the world? Well, most video analytics are limited to analysis within a single camera. And, when you have to analyze some type of activity that occurs

over a much larger geographic scale, it becomes very, very difficult to do this with the tools that are available today.

So, for example, being able to track a person or a vehicle through a network of hundreds or thousands of cameras is way beyond the state-of-the-art. There's a tremendous amount interest for this in DHS, in order to track people, who are designated to be suspicious, through airports. All right? But it's a very hard problem.

Video forensics -- again, some of it is limited to analysis within a single camera. But you might remember the London bombing; when they try to backtrack the people who were involved in that bombing, through thousands of cameras, and there were essentially no tools of any sort. It was a tremendous amount of human labor, to go back into that video collection and -- from the incident -- track people back to their origins. And, even in a retail establishment, we have a relatively small area to monitor. You still might have an order of 100 cameras; and, after a shoplifting incident, you have to create this video forensic trail, and there are no analytic tools to really help you do that and find the group of people who were involved in some theft. All right?

And then there are really hard problems that aren't going to make it out of the research lab for at least a decade. So, this generation notion of being able to detect that somebody is suspicious or nervous, for airport security, immigration screening, crime detection, things of this sort. And probably the holy grail of video analytics is to be able to stop, ahead of time, some terrorist attack. But this involves a very, very difficult problem in analyzing extremely large networks of video cameras to find lots of people and places and vehicles involved in dangerous activities.

All right. So, I want to talk a little bit about privacy protection and visual surveillance. How can we protect the privacy of people as camera networks proliferate? Well, there -- there's a lot of self- identifying information when you take images of people, the things that they own, and the vehicles that transport them. And some ways of protecting privacy involves actually destroying information so that it can't be recovered in the video. So, Sam was talking about, just for purposes of showing a video in court, blurring out faces. But, suppose you actually want to ensure privacy at a much more fundamental level, at the collection level, to make sure that even your operators can't use these videos to track people that they're interested in, aside from people who were involved in some kind of criminal activity? All right? So, the simplest thing you can do is just blur, as this image shows, and Sam shows, a face or a license plate -- right? -- or some other biometric. This is destructive; you can't recover the face after the image has been blurred. So, if you actually saw some action, and you wanted to determine who the people who were involved in it, you can't go back and unblur the image. This is not completely true; I should tell you that even if you had a video, and you blurred every frame, there are actually techniques that you can use, just like CSI, to recover a high-resolution accurate image of the face. So, it's only superficially that you think that this might

preserve people's privacy; in fact, if a clever-enough digital technician got ahold of that video, they could, in principle, recover a high-resolution version of the face. All right? So, probably this solution to providing what's called everlasting security, in the context of video surveillance, is to use the same techniques that we use to provide security when we perform financial transactions over the Web. Right? And that's to use the same type of cryptographic or encryption techniques that we use to send our credit card numbers and other personal information, self- identifying information, in the context of a financial or personal transaction. And the fact of the matter is if you could find the faces, like we did in that amusing photograph of that crowd with our digital camera technology, then we can take all those pixels that make up the face, and we can treat them just like a credit-card number, and apply the same methods that would allow us to create a visually-unrecognizable face, but a face that's perfectly recoverable if you're a person that has the secret key, so to speak. And this would give us as good as security you have in our everyday financial transactions.

But this isn't actually true, either. And that's because of some of the points that Sam made, also, in that all of these detection algorithms make mistakes. They have what are called false dismissals; there are faces that they will miss. When a face is missed, then it won't be encoded using these techniques, and it will be available for everybody to see. And, in the long run, it's also important to keep in mind that there are many other weaker biometrics that people are investigating in the research community to see if they can be used to identify people when you can't see their face; this is wholly true of what's called identifying people at a distance, where there simply aren't enough pixels on the face of a person to identify them, but there's lots of other information that can be used to identify you, including your body size and shape, the way that you walk, your hair, the type of clothing that you wear, that can be recovered in these situations, and combined with other information to identify you, even if your face is obliterated by one method or another.

For those of you who are really concerned about privacy of individuals in large-scale video surveillance systems, these are some factors that I think you need to be take into account and think about.

MR. SAND: I have just one question. Do you find that most of the research -- or that the questions that people have are focused more on the analytical side, or are they focused more on the realtime-alert side?

MR. DAVIS: I think that they focus most on the analytic side. So, a lot of concern with video forensic and being able to go into large video collections and create very accurate descriptions of people and packages and vehicles so that they can be used to identify perpetrators of events.

MR. SAND: After those events –

MR. DAVIS: After those events have occurred. But there's also a significant amount of work on realtime online. Video surveillance, what you would call triggering events -- so, I'd say, probably about 50-50.

MR. SAND: Thank you very much. Larry?

MR. STRACH: Thank you, Peter.

My presentation this morning is going to cover in two areas, and these are -- come from the perspective of some recent installations my company has done. First, we'll talk a little bit about the equipment. Many of the items under equipment have already been discussed, so I'll briefly just go over them. And, secondly, is architectural. More and more, it's becoming apparent that we have to be good CCTV neighbors, and certain things have to be taken into account, other than the technical aspects, to go ahead and put a good project together. And I'll show you some examples of that.

First, on their equipment and camera placement, we already talked about intelligent video, yes or no. That's one of the first things you have to decide, because that greatly determines what you can and cannot do with cameras, and where to put them.

Second item is, Can you see what you want to see? We've already talked about that. You've got to make sure that everybody who uses the camera systems has a different perspective of what they want out of them, and you've got to be able to go ahead and determined and decide and satisfy all those needs. So, just putting a camera up may or may not be good for one person or one department, but it may be required for another.

Also, you need to take into account seasonal changes. You put a camera up in the fall; come spring, a tree has grown. What do you do now? Or a building was put up where we installed a camera, blocking our field of view. So, those are the types of things you have to take into consideration.

And, lastly, you've got to make sure you could do what you want to do. And, again, that goes back to defining what it is you want to do with your system.

One of the biggest things of putting in a CCTV system is the infrastructure. What is existing or not existing? Whether you use poles to hang the cameras, and the type of poles you're going to use. Some people like telephone poles, some people like I-beams, some people like attaching the camera to buildings. You've got make sure that you can go ahead and put those devices in, whether it be a pole or an I-beam.

Also, if you want to attach your cameras to buildings, you've got to make sure it's okay with the building owner or the operator, and there are no covenants or regulations preventing you from doing that.

As part of the infrastructure, you've got to get all the information from point A to point B. Generally, you need conduit and cables to do that. You can go wireless but, generally, you

do need some type of wire. And, along those lines, Is there existing conduit or existing wires that can be used? If not, do you have to put in new ones, and what are the challenges associated with putting in new conduit and new cables?

The IT infrastructure. The issues here are the three ways of getting information from point A to point B. You could do it over copper, you could do it over fiberoptic cable, or you could go wireless. Copper has its challenges, in that sometimes it's difficult to get sufficient bandwidth, even though lately there have been great advances in technology, allows you to go ahead and push video over existing copper wires, whether it be telephone or even abandoned power cable, that allow you to go ahead and push video over a long distance. Fiberoptic cable is probably the best way to go, in terms of transmitting information from point A to point B. It's almost endless in bandwidth, it's immune to interferences from either lightning or surges, things of that nature; however, it is expensive. At first wireless, was very, very popular, because you didn't have to go through the expense of putting in a copper or fiberoptic infrastructure, but, unfortunately, you're limited to a certain bandwidth -- frequency spectrum, and more and more, those spectrums are getting populated and crowded as more people are starting to use them, and you run into interference issues.

And, lastly, power. You need power everywhere you go, and you need to go ahead and determine whether you could tie into existing power infrastructure with existing grid, or can you use the end-user's power, or do you need solar, or do you need backup generators, things of that nature.

One of the most important things with cameras is lighting. You can only see what you can see. And there are three issues involved with lighting. One is whether you're determined to use visible lighting, which is lighting from the sun or lights that are currently there; infrared lighting for low-light conditions. At night, it would be best to put up lights to go ahead and light up the area. Lighting up the area does two things: one, it acts as a deterrent; and, it gives you good camera views. However, nobody wants lights blaring in their bedroom or their house or in their office, so sometimes you have to go to infrared lighting. But infrared lighting does have some drawbacks, especially in terms of intelligent video. These types of things have to be taken into consideration.

Lastly, what I mean by existing is, you could have a camera, you could put it in place, a light comes on at night, and it shines right into the camera. That light could be from streetlights, it could be from cars, it could be from anything. And what you get is what they call blooming, where this big white light all of a sudden shines into your camera and blinds you. So, that's another thing you have to take into consideration.

Shadows -- shadows from cars, parked cars, clouds -- could have a big impact on what you can and cannot see. A shadow from a building. You have a field of view of your camera that's half in the light, half in the dark -- makes it very difficult for the camera to settle in on the optimum settings.

And, lastly, is weather conditions. What happens down South, where it's generally warm, versus up North, where it's snow and ice and fog conditions? All those have to be taken into consideration.

Next part is architectural. I said, before, about being a good neighbor. I have, up here, two slides of a project we are currently under construction here in D.C., and it goes along a railroad line -- and what it shows is two views. You can see, on your right, the VRE; on the left is just a railroad trestle, with no trains on it, but you have a view of the Jefferson Memorial and a view of the Capitol. Our original design had aerial cables and poles going in these areas. After meeting with certain local authorities, it became clear that, you know, this is not good neighborhood practices, not something we should do. And they were right, we should not go ahead and hang up aerial cables or poles and block these views. We need to go ahead and go underground or go in conduit in the railroad infrastructure, to hide those cables. This was not something that was contemplated up front, but this is now something that we take viewshed issues into design up front. Is there anything that's going to block the view?

Another recent job we did was a tunnel security project located in New York. And this particular tunnel is a historic landmark. They wanted to secure this tunnel with cameras and intelligent video, which required installing the cameras onto the tunnel face. Couldn't do it. It was a historic landmark. So, after meeting with the people responsible for this area, we came to a compromise, and what we did is, everywhere we put bolts or anything we needed to install to the tunnel face, we did it into the mortar in between the bricks. Mortar is easy to replace; those bricks, you can't. So, that required special mounting brackets and things of that nature. Another design consideration.

Getting in the spirit of being good neighbors, all the intelligent video, all the cameras, all the sensors, everything comes back, generally, to what we call a head end. This head end, in times past, used to be, basically, a shack. Not anymore. We have to look at making it look like a good building, maybe a nondescript. The building you see here is exposed aggregate face, it looks like a communication building you see anywhere else along the highways or railroad tracks or anything like that.

Lastly is color scheme. Believe it or not, when we hang equipment up there, it does make an impact, in terms of the surrounding area, what color it is. What you see here are are two camera clusters. The one on your right is painted black, the one on your left is painted white -- and junction boxes associated with them. This is within the same project, just a few hundred feet apart, but it was determined that it would be better, from an architectural standpoint, to have that equipment painted black. It's not just the cost of painting the equipment black, though, that you have to take into account. By painting them black, it's prone to higher heat, especially during the summer, so we had to speak with the vendors, get

special design-approval changes to go ahead and compensate for the added heat increase. But, overall, it does look pretty good, so I'm glad we did that.

And that's about it. So, in summary, what I'd like to say is that, not only do you have to look at equipment -- at technical issues -- you also have to look at architectural issues. Be very flexible, because everybody has a different perspective and need of what they want out of their camera system. Try to match all those needs. Thank you.

MR. SAND: I just have one quick question. As you were talking, I was imagining the person buying the system, sitting at his desk getting a big red button that's going to turn everything on, and just sitting there waiting for the "go" sign, and people regularly coming up to him and saying, "Well, you know, there's a tree, there's a shadow, it's too bright, somebody turned on their light and we can't turn it on" – I mean, all of these very practical mechanical things that would just throw the whole project off. I would imagine that a client would just be totally shocked by how much of this stuff has nothing to do with the technology at all.

MR. STRACH: Yeah, that's right, that's the big issue. And you're absolutely right, Peter, this is not a mature technology, even though it is maturing. And a lot of people are continuing to learn. We don't have a lot of standards and regulations in place, things of that nature. So, everybody has a different perspective of what to expect. And, unfortunately, sometimes it's just OJT, on-the-job training, things you brought it up, we had to learn the hard way. Our company is now taking upon itself to go ahead and develop programs that allow us to go ahead and place cameras at certain height, at certain angles, and we build a 3D model within the area that it's looking, so we could get up front, as much as possible, what we could expect. And the 3D modeling nowadays is just phenomenal, because you could adjust the lighting schemes over the course of the year, so you could see what the effect of light over the course of the year is doing; the field if views, if there's a highway overpass where big trucks go by and cast shadows -- and you know what the problem with shadows is, it's not so much if a shadow shows up and stays, it's the intermittent shadow of car coming by, shadow/sun, shadow/sun, shadow/sun. All those things cause problems. So, you have do computer modeling up front, to try to get most of the errors out of the way. However, I could assure you that every time we start up a system, we've got to adjust this camera, we've got to move it here; we have to put a shroud over there; we may have to add a camera there, because this particular field of view just isn't working. And you have to build that into your design, recognize that you're going to go in after the fact and make some adjustments.

MR. DOCKNEVICH: Peter, I'd just like to add a comment. When we work with clients, we tell them, Consider your CCTV system like buying SAP or any other business application. You can't just buy this package and put it in and have it work, because you don't do that with any other application; there is an extensive customization period that happens after you do the requirements definition up front, right? The first thing you do with a business package, here's the package, and, What do I really want to do? And then I customize it. So, as Larry

was saying, you can do all the preplanning and use computer models, and then you still have issues that come up that were unexpected. And that's the tuning phase that every project has, to make it very successful.

MR. STRACH: And, again, to build upon what Sam said, it is not plug-and-play. And technology is changing so much. What we thought would be good for a proposal in a project 6 months ago, may not be there by the time we get to project execution. The camera technology is changing so quickly. I remember, a couple of years ago, the -- what they call the lux level for a camera, which is how much the camera can see in dark situations, has improved greatly. What used to require, for sure, an IR light or a visible light nowadays may not need it. So, take that into account, too, as well, that -- you know, projects -- the gestation period for projects could be quite long, could be a few months, it could be several years, and so, think about that, the -- how technology changes every few months. When you go to your final design, final installation, you might want to take a fresh look, step back and see what's out there.

May I add just one thing, a lesson learned really recent on that. It has to do with infrastructure. We were going to go ahead and add some additional cameras to a project, and we were going to run dedicated wires for those cameras. A new technology came out which allowed us to go ahead and build on a single fiberoptic cable - - what they call daisy chain -- we could add many different equipment devices to that single fiberoptic cable. And we just found that out, literally, two weeks before we were about ready to install that camera, but it saved us thousands. And that's something that we could all learn.

MR. SAND: Randy?

MR. HOFFMASTER: My name is Randy Hoffmaster. I'm director of engineering for Epsilon Systems. And I'm going to talk from the approach of a prime contractor executing DHS-level contracts. And when I talk about DHS level, I simply mean rather complex, large-scale surveillance systems, not just one municipality. And I'd very much like to leverage off of something that Sam said up front, and that is, you can't do too much planning, not only in the system design, but in the actual implementation into the field.

So, one of the most important factors that's been our experience in recent contracts is -- well, actually, three contracts -- is the ability to build teams. And what I mean by that is, teams that are not under contract, but they have an influence over the outcome of your project, whether it's by schedule, by configuration. For instance, the power companies seem like that would be a pretty simple thing to deal with. We've had cameras up, a project nearly complete, have been waiting on two out of three power drops for closer to a year. Those are long-lead things that you may or may not have into the end of the plan.

Municipalities, there's been some talk about some historical and architectural perspective. Keep in mind that there's also a planning commission, and that the day that you do your site

survey and you start to plan for a project, you should put on the glasses of the municipality, who has a 5-year plan, or you may be putting up a rather industrious-looking security system that is going to wind up in the middle of a park, and then you have a big setback.

Departments of transportation, not too big of a deal. We've had some right-of-way issues we've had to attach to some of their structures. But, again, it's one of the things you've got to plan for.

The owners of the critical national infrastructure -- and what I mean by that is, the actual facility for which the surveillance system is installed -- they've typically already working in a rather dangerous atmosphere. They often have their own security in place. And if they don't, they at least have their own safety personnel certifications for workers. The cost associated with that can be significant. You may have to pay for that track protection, you may have to pay for those certifications. Just something as simple as using lifts on property can be four to five times more expensive on some sites, because you may not be qualified to use it, because only one vendor is able to do that on particular sites.

There are some union issues, only from the perspective of cost management. There have been times that we have costed out work, you get to some facilities that are unionized, and the union may claim the labor. That's okay, except that you lose control over your budget, you lose control over the schedule, and you just have to be willing to go into that with your eyes open.

CONOPS, I don't have an answer for CONOPS. I simply need to put it on people's screens. Security, typical security, is a raw-data provider. What you do with that raw data is difficult, -- extremely difficult in some cases -- and a lot of that has to do with how many people are in the chain of command. Is it for one municipality? Is it just for the city of Norfolk? Is it just for one particular municipality, or is it going to be daisy-chained to where it's going to be escalated from a railroad to local law enforcement to State law enforcement to Federal law enforcement?

Again, I don't have an answer for that, but, as surveillance systems become more prevalent, I believe that that's something that needs to be addressed more and more.

Adjacent landowners -- that's simply a lot of what Larry was talking about, with making -- there's a lot of influence that the adjacent property owners may have over the configuration and deployment of a system, both from a privacy point of view and from an aesthetic point of view.

The pre-award assumptions that drive cost. What we've found is that, oftentimes, the receiver of a system underestimates how intense the server solutions are for surveillance systems. They may underestimate how much server space they have. They may underestimate how much HVAC is required. All of those are very costly upgrades if they're

needed. We've installed systems between 2 and 22 servers. You can imagine how much heat is put out of several racks of servers.

Power and vital power. Another thing I'd like to leverage off of, that we've been talking about, is to look at what you're trying to solve, and scale the solution to the problem. Don't just assume that you need 30 frames per second, don't assume that you just need vital power. But, what you do need to look at is most of these system provide vital data. You don't want to put it on an unmanaged line. You don't want to put it on an Internet connection that doesn't have guaranteed bandwidth or that doesn't have priority service if the line goes down. If you require vital backup power, you have to acknowledge that computer systems require a very quick transfer of power. Determine whether your system needs it, or whether a part of your system needs it. It's a very important factor in the costing out of transitioning between power and vital power.

The backhaul communications, part of that I had talked about, was the Internet connectivity of managed service versus an unmanaged service. You can transmit video across a DSL. It may be unmanaged. If you're looking for vital information, and you are not guaranteed a bandwidth, you're not guaranteed service, that may not be a good plan. We've even looked at some redundant measures, where we would use a T1, with a managed DSL as a backup in case we lost the T1.

LAN versus standalone. Primarily, I look at this as riding on most LANs is a difficult thing to do. From an execution point of view, from an installation point of view, from an integration point of view, a standalone system is easier. Now, there are complications that go along with that, in that you have to have a separate backhaul. But I don't want to belabor this point; a lot's been talked about -- bandwidth, storage, and, not so much video compression, but if you go and ride on someone else's LAN, the amount of storage require usually makes the normal storage for that LAN pale.

Web-cam mania. I term that as -- everybody thinks that they can see video from home. There are huge factors that drive that. If you can look at your college campus, your alma mater, and it's trickling data, and it doesn't have very good resolution, you can't push this data anywhere across the Internet. Frame rates have been discussed. Resolution. And now, with megapixel cameras, that really escalates things; it goes from an analog, that might have 3-/400,000 pixels, to a megapixel camera that carries, maybe, 3.1 million pixels. Encryption also slows down transmission.

Soil disturbance. Another cost driver that we've experienced is that we can not touch soil. In the planning of installations, then, you have minimize how many foundations you make, how many poles you put in the ground, how much fence you put in the ground, how much underground cable you put in the ground. And you have to have a partner in infrastructure, a -- the facility, who's willing to take care of the expense of actually remediating that soil.

Physical vulnerabilities. With the rollout of security systems now, it's somewhat of a legacy system, a legacy approach. The mindset that having cameras deters crime is, I think, true. The idea that you break the world into typical crime and terrorist crime is what I'm getting at.

Typical crime, if you hang a camera, there's a deterrence that goes along with that. People aren't too interested in learning how to defeat that. But, the terrorist level, you're dealing with a very cunning, calculating, educated, sophisticated threat, and, with that, if you make your system entirely overt, it's -- you can see the camera, you can see the transmission line, you can follow the transmission line to where it goes. Perhaps in the future, in -- just from a technology point of view, not from a privacy point of view -- fundamentally, I believe that the systems, in the future, have to be less overt and more covert. Overt enough to become a deterrent, a visual deterrent; covert enough to where it can't be easily understood and defeated.

Engineering, configuration control, and life- cycle management. Again, in the evolution of things, in order to know -- in order to troubleshoot a system for its life cycle, you need to know what the system looks like at any given point of time. It's going to change. Things are going to change from the time it was designed until it was rolled out and while it's being maintained. And there need to be standard operating procedures that are in place in order to do that.

Drawing management's standard operating procedures from initial design, the review and approval authorities, all the vested parties -- when you actually go to deploy, there's going to be deviations. There needs to be a system set up to approve the deviations as it's being deployed, and then a system for incorporating those deviations, and then, as the system is revised, to revise the drawings, accordingly. The same is true with material management and documentation control.

I'd like to thank you very much.

MR. SAND: Just one quick observation; maybe it's a question. But, as you were talking, I was imagining all the things that you'd have to take into account when you decide you want to do a video installation. And it sounds like a person may see the cameras first, and think, Well, I want that great camera that can do all those great things, but it sounded like, from what you were saying, that the camera is actually, maybe, the last thing you should be thinking about, because that's the quickest thing that you could install. And if you buy the cameras on day one, and it takes us a year to get permits and all these things, by the time you're ready to actually install the camera, it's going to be out of date, and the new features will be available, or maybe a better camera or better, you know, transmission stuff, would be available.

MR. HOFFMASTER: Absolutely. Technology is changing very quickly. And, when you look at the cost to do maintenance, if you get your equipment early, you're already cutting

into the time that the manufacturer's warranty period starts -- it starts clicking the day you buy the camera. So, if you buy it, and you shelve it, you're actually losing money just by putting it on the shelf and not deploying it right away.

MR. DOCKNEVICH: You know, one other, just, quick comment on that is, I met with a client, and they actually called their CCTV project The Camera Project. And I said, Wow, can we get rid of this? It's really not about the cameras. Why don't we call it, 'How-can-I-make-the-city-a-better-place-to-live-in Project?' and then maybe CCTV-and-cameras is a solution or a way to improve the city. The Camera Project or a focus on cameras is just all wrong.

MS. KING: I'm Jen King. And I do have to start off by noting that, about a week ago, my wallet was stolen, and the person who stole my wallet promptly went out and used my credit cards at the Bay Area version of the Metro to buy transit tickets. And, when I made this report to the police officer, she kindly offered to have the transit agency pull their video surveillance footage to try to identify the person who used my credit cards. So, it will remain to be seen whether or not I actually benefit personally from surveillance technology.

I don't have a lot of time, so let me just go quickly. I'm going to talk about a few things, some things you've heard a bit today: information overload -- basically, the amount of information you collect with a system; information management and retrieval, what you actually do with that information after you collect it; the personnel and expertise you need, to actually use the system effectively; and a few notes on forensic systems versus live-monitored systems.

And a few notes. I typically research privacy and technology issues, but I'm not going to be talking about privacy in this discussion. Instead, I'm putting on my hat as someone trained in information science, as well as having built software in the past.

And one thing I'd actually like to point out, especially for the law enforcement folks in the room; I think video surveillance, in a lot of ways, is a question of resource efficiency and how efficiently you're able to deploy your resources. And so, I'd just note two things you might want to keep in mind, which is, basically, how efficiently and how effectively does video surveillance allow you to actually deploy your resources? One question that often comes up is, Is it more effective to have more cops on the street versus using these tools? And another thing, which I know we probably won't be delving in deeply at this workshop, but, generally, How effective is it in fighting crime?

Information overload. How are you actually going to manage the data that you collect through video surveillance? Some questions that you might want to consider are actually understanding the context of the information you're collecting. So, being able to go through and determine the goals of your system, versus what's just noise. What are the true safety threats that you're looking for, versus what might be overreaching and considered a privacy violation? And how to avoid issues of bias, especially if you're using humans to sort through your data. So, questions of bias, racial profiling, et cetera. There's a huge question around

this with information coordination. So, suddenly you have this entirely new stream of information coming into your department, and you may also have other sources, such as live dispatch and such, How do you make all these things play nicely together? Suddenly, you have a lot of different inputs.

And, actually, one I wanted to throw on here, just because -- anybody in the audience who's had experience with this -- one thing I've heard mentioned several times now is, First we're putting in the video system, and then the holy grail is when we can pipe that video into patrol cars. And as someone who studies human/computer interaction, I'm extremely skeptical of this notion that, cops who are already dealing with taking calls and driving, can multitask to the point where actually piping in a live video feed is helpful. So, if there's anybody in the audience who has questions or thoughts about this, please do come talk to me, because I'm curious in hearing more about it.

So, information management and retrieval issues, you need to be thinking about long- term storage and archive management. Recently, I went and viewed a system in California for about a medium- sized suburban city, and they have what looked to be a fairly sophisticated software package managing their cameras, but the one thing I noticed was that it was built on top of just a Windows file system, and, you know, they only had, maybe, 15, 20 different files stored at that point, because it was a fairly new system. But when you go into the thousands, using Windows Explorer as your tool for managing your archive is not a very efficient solution. The officer actually -- since there were, like I said, about 15 to 20 files in there, he tried to find us a specific one to show us, and even that took him about 10 minutes. So, how you organize your archives will have a huge effect on how usable your system is for your employees.

Also, video search, we've talked a little bit about this earlier -- the ability to actually search through video itself is something that's coming, but not really there yet. How are you actually going to isolate critical segments of your video, especially for a postevent investigations, and find that sort of thing? And how you manage your archives really will have a huge effect on other issues, such as the general oversight of the system, how you're able to audit it, how effective it is, and actual workflows, as well, if you think about what an investigator has to do, to go through footage and make that as part of their daily job; you don't want it to become an extremely laborious task; otherwise, you lose any gains that you have from implementing this sort of thing.

Okay, so, personnel and expertise, who is actually going to manage your system? One of the things I've seen so far is that smaller jurisdictions literally will take an officer on the force and say, Hey, you're our new video surveillance manager. Have fun. And these are usually officers that have no IT expertise, no training at all; and the few I've talked to have done really heroic jobs in learning how to navigate this sort of thing, but this is obviously not why they joined the force, and they didn't have the expertise to do this sort of thing. And, in

general, these are not desktop support folks, so you need to think about, who's actually going to, on a day-to-day basis, care and feed for the software. It's enterprise software, and so, it requires a lot of care and feeding, essentially. What this might mean is that, if you're a smaller jurisdiction, you may need to sign a contract with a vendor to actually provide you with support after the fact.

Who's actually going to train your investigators? A lot of folks think, Okay, with the system installed, we do a training, and we're done, but that's not usually how these things work. You generally have to always have someone available who will provide you with training and be able to train, after the fact, new folks, you know, relearning, that sort of thing.

And then, one big question is, How will you know if your system is actually working? You know, you kind of have a duty, in a sense, to not only understand if you're spending money in a way that's efficient for the department, but also, you're spending public money, so it behooves you to find out if the system is actually meeting the goals for which you've chosen it for. So, you need to put some thought into how you're actually going to track this. What type of statistics are you going to collect? You know, how are you going to manage tracking requests? How will you know if you've actually used them successfully in prosecutions, that sort of thing? So, things to think about.

Forensic systems versus live monitored systems. A forensic system is one where it's usually just a fixed installation of cameras; there's no people behind the scenes actually directing the input. What happens here is that what you capture is largely accidental; it just happens to be whatever is in the frame of the camera, wherever you put them. And so, you need to consider, basically, what effect this will have on your goals.

One of the things to think about is where you're actually putting your cameras. And, largely based on the characteristics of the area, it may or may not have a deterrent effect. If you put them in a hotspot area with a lot of crime, you need to consider whether or not you think it's feasible that crime will move around the corner. Some things to think about in that sense are, especially if you're in a gang territory, gang areas are very tightly controlled, and there's a lot of fight over territory. And so, if you putting a camera in an area infested by gangs may actually benefit, and may actually have the effect of deterring crime, or it may escalate a gang war over fights when people try to move the crime around the corner. So, there's a lot of, like, minute details you need to think about.

Live monitor systems, those are systems that use human operators. We don't seem to have a lot of those in the U.S., compared to the U.K., at this point, but that might be changing.

Presumably, in that case, you're going to deploy resources based on what the people behind the cameras actually see, so there's a bit of a feedback cycle. That's something that's usually lacking in a forensic system, where you're not necessarily deploying anybody based on what's being caught by the cameras; it's used after an investigation, to do research.

And so, in a live monitor system it's much more of a live feedback cycle. You see something going on, you'll most likely deploy officers to the scene. And so, again, you're going to think about, you know, what's going to happen in the area where I put the cameras? Is it actually going to deter crime? Is it going to displace crime? Will people find that type of setting to be more big-brotherish or will it have the desired effect?

Okay. So, just a few last words about live monitored systems. And we've heard a little bit of this already today. Monitoring is an extremely active activity. This is not a job for couch potatoes. This is not about just watching TV. It requires a lot of focus and a lot of stamina. So, there is some research out there that suggests the limits of attention at anywhere between 30 minutes and 2 hours before you actually need a break. So, this obviously has a lot of impact on how you staff a live monitored system. You can't set one person there and expect them to work an 8-hour shift straight through and actually be effective in noticing what's going on, on the screen.

There's also been some points brought up about visual attention. So, how many screens can one person actually focus on at any one time? And what is the optimal ratio of how many feeds you're looking at for a single person?

This also brings up more information-overload issues. So, if you're imagining you have an officer trying to watch the cameras, but they're also dealing with live dispatch, and they have information coming in to them, and they're trying to relay information out, there is a lot of issues, how much information can one person deal with, and how are you going to, basically, effectively manage those issues?

And then, finally, some qualifications and training issues. Like I said, I've seen several installations now where they basically choose a single officer to manage the whole project. And, in one case, he was going to be the person sitting behind the camera; or they were going to use officers who, for whatever reason, couldn't be deployed in the field, and have them sit there. But you need to, you know, ask yourself, Are those the people you want watching your camera feeds, or do you want dispatchers, for instance, or do you want to find civilians that you train? How much training do you actually need? You need to think about, not only just how to use the software and the system, but, actually, how to use video, in and of itself, as a tool for public safety. And then, again, those issues of racial bias and profiling and so on, so forth, that inevitably come up when you're using human beings to deal with this sort of stuff.

Thanks so very much.

MR. SAND: Okay, we have about 15 minutes left of our panel. If you have any questions, there's a microphone in the middle of the room. Please come over and talk to us.

In the meantime, we'll ask ourselves questions.

One of the things that got me about your talk was that you could go through all of the issues that we talked about, and you finally have your cameras installed, and they're appropriately fixed so that you can do all the analytics, then the guys you're trying to watch just go around the corner, and you're stuck. I mean, there are so many kind of ironic and unforeseen issues that have nothing to do with the experience that I would imagine people expect when they think about CCTV, based on TVs and movies and that sort of thing, or even what an organization may expect, in terms of the deliverable. I can't imagine how those initial conversations must go. I mean, it must seem almost bizarre that you talk about everything other than exactly the one thing that your client expects to have a conversation about. I mean, it's amazing to me.

MS. KING: I think, also, it depends on precisely what your goal is. So, if you're trying to reduce gun violence, that's a very different goal than if you're putting cameras in a tourist area, and you're there to move around the corner. In some places, you might want to deter people from being in an particular area, but you want them to move crime indoors, for example. If it's a tourist area, one of the goals might actually be to push crime out of the main tourist areas, so that the tourists feel safer -- and then push it to another neighborhood. So, it depends.

MR. SAND: So, a CCTV solution could be, buy a bunch of old cameras that don't work, put them in very visible places, don't hook anything up, and just move the people –

MS. KING: Possibly, although –

VOICE: Just the housing.

MS. KING: Hopefully, the legal panel will discuss issues, if there are any, with putting up cameras and if you're on the hook once you put up the camera, from a public-safety standpoint, of actually responding, or being able to deter crime, you know, if there's a camera there. You could have somebody who is victimized in front of a fake camera, and, you know, basically, said, Well, I was expecting, you know, that the police actually had footage here to help me out after the fact, and they don't. Does that create any legal entanglements for you?

MR. SAND: Or they have too much footage, and nobody can find any of it.

MS. KING: Yes, that would also be a problem.

MR. DOCKNEVICH: Right. Like we said before, the first step is education, requirements, and design, then pick out the technology.

MR. SAND: Okay, we have our first question from the audience.

VOICE: Hello? I've heard, anecdotally, that the gambling industry is far ahead of everybody else in video analytics. Can anybody here say whether that's true or false?

MR. STRACH: They can afford it.

[Laughter.]

MR. STRACH: And you're absolutely right, the things they're doing with video analytics is quite impressive, and there's a real payback. And let me build on the word payback. When you're a private it's difficult to justify a return on investment. Usually, private companies put them in place, maybe to reduce risk, help with their insurance rates, deter theft, a knee-jerk reaction to a public relations incident, to be perfectly frank. But the gambling industry, there's a straightforward return on investment. For every dollar they put in, they know how many dollars they hope to recoup from that, and they're willing to go ahead and spend that kind of money.

MR. DOCKNEVICH: Yeah. Absolutely. There's also regulatory issues on where they have to have cameras and how long they need to store their video, but -- we've said before, the key issue with -- whether it's CCTV or analytics -- is, Does your business case support it. And, as Larry said, they have a great business case.

MR. HOFFMASTER: One other thing about the solution, though, is the analytics is all about understanding the background. The problem is always located right where these four people sit around the table, it's -- and the problem is typical -- it's much easier to come up with an analytic solution for that, versus looking down a fence row and trying to imagine every possible threat that could come across the camera.

MR. BLITZ: I'm Marc Jonathan Blitz. I'm a professor at Oklahoma City University School of Law, and I'll be speaking tomorrow on a panel. Two quick questions about PTZ cameras. You said that the analytics -- or, at any rate, most of the analytics don't work with PTZ cameras. Does that include the privacy-related analytics -- for example, the face masking; does that mean that you -- for example, that you can't use those privacy protection mechanisms in PTZ cameras?

And the other quick question, so you can answer it, as you will, is whether or not there's been any progress at replicating the benefits of PTZ cameras in video forensics. So, if you'd like to zoom in on something when you're watching archive video, how effectively can one do that, and are there improvements in that technology?

MR. DAVIS: Well, I think it's not completely true that these video analytics are restricted to stationary cameras. There are a number of applications, especially in border security, where you use PTZ cameras -- or at marine facilities -- and the motion of the camera is relatively slow, so you can build up the type of models that you use in a stationary camera, use the same type of analytics on a slowly moving pan-tilt-and-zoom camera. The problem is, if you want to move very quickly from one viewpoint to another, there's a certain lag time before the analytics can kick in, because you have to build the same types of models that you would have built for the stationary camera. It takes a little time to do that.

But these so-called privacy-preserving technologies, like finding faces, they're probably more accurate with methods that don't rely on a stationary camera, in some sense, that just directly use models for what faces look like to find them in any frame of video, whether it's from a stationary or nonstationary camera. So, I don't think, in that regard, there's anything that restricts those types of technologies from being applied to arbitrary cameras, because, basically, the best methods work on a single frame of video and don't rely on motion or anything of that sort.

MR. HOFFMASTER: With regard to zooming after video's been recorded, the only way you could do that would be if you had a megapixel image and you were digitally zooming on it, because you can't -- you can't optically zoom on something after it's been recorded. And with forensics and PTZs, those typically have to be on a very fixed path, the analytics has to understand what the background should look like, and it has to know where the camera is pointing in order to compare what it should look like to what it does look like. So, sometimes they have predefined stops, sometimes, like Larry said, it's a very slow panning.

And I would say it has to do with what size of a background you're looking at, whether you're looking at a room like this, or whether you're looking at a desert situation, where you're literally looking over kilometers of area, and looking for changes to a landscape that are very, very few changes, other than the typical shadow and even the weather isn't that significant in a desert kind of situation.

MR. DOCKNEVICH: Yeah. And just one other comment on your second question, about panning and everyone has a digital camera -- right? -- and it has so-many-times magnification optically, and then you can digitally magnify it, but you lose clarity. So, again, it's important, if you want to do that after the fact, to have the right megapixel cameras in place in the beginning, so then you have more pixels, so, when you blow the image up digitally, you can still see what you want to see from your frame.

The other issue to consider is, when you are looking at IP cameras, whether they're interlaced or progressive scan. Everyone understands, interlace is like your old TVs -- right? -- where they have odd and even rows, and they paint the odd rows, then they paint the even rows. Now, the problem is, even at the speed of light, it takes time to paint those, so you're going to have a little interference. So, if you're looking at a video of a camera interlaced, and you're zooming in, you might not see anything, because it'll be blurred. High- quality IP cameras are progressive, which means they paint the entire picture at the same time. So, again, thinking through what you -- if you want that capability, then you've got to get the technology to deliver that result to your user.

MR. HOFFMASTER: And, Larry, you made me think of one other thing, and that was that if you go to a megapixel camera, you're probably already dealing with terabytes of memory -- when you go to a megapixel, you're multiplying that by, maybe, sixfold.

MR. STRACH: And just one other thing. But if you have video analytics installed on your system, if you have an alarm, you could have your cameras pan- tilt-zoom into that alarm, and get images that way. But that's only if intelligent video is active.

MR. DOCKNEVICH: Right. I mean, that's a typical -- go to your average retailer -- okay? -- and you'll see, look up in the ceiling, and you tell by the size of the dome. Look for the little dome, so it's their inexpensive, cheapo, fixed wide-angle cameras, and there'll be a million of them. And then look for the bigger domes. Those are the expensive PTZ cameras, and there are going to be few of those. So, you use the cheaper technology, deploy more of it; and then you have the special technology that gives you the extra capabilities, but you don't employ that everywhere.

MR. STEVENSON: Boyd Stevenson, American Trucking Associations. I've got two questions that, unfortunately, are not related, like the professor's were.

The first was, as you were talking about the quickly-changing technology of the cameras and how it's -- often makes sense, after you've gone through your design and layout phases, to buy the cameras at the end. After you've gone through that and installed the system, how long are we talking until the cameras you have purchased for your system are obsolete and need to be upgraded to the newest technology?

My other question concerns privacy issues. If you're a private employer, even a public employer, talking about bringing cameras onto your worksite, do you see a lot of pushback from employees in the workplace, especially unionized employees?

MR. DOCKNEVICH: Well, I can address the first part of the question. Quickly-changing technology makes your cameras obsolete, I don't think cameras ever get obsolete. And we still deal with clients who have analog cameras, and they say, Gee, I bought these cameras, they still work -- I mean, I'm getting what I expected to get, which is the same -- it's not like they degrade over time. And so, then you use a converter, convert it to digital, and you use those same cameras.

One of the advantages of IP-based network solutions is, you can just unplug and replug into your network without extensive recabling. So, don't look at it as, The technology becomes obsolete and ceases to be functional. You just have to design, with open standards, open architecture, so, as you get new things in place, then you can plug them into the same network that you have today.

Larry was using an example of new technology with fiberoptics that just made it easier to daisy- chain cameras, and that happens. And, again, if you have an open architecture, standards-based solution, you're going to be ale to do that.

MR. STRACH: Your second question regarding pushback from putting cameras with either unions or workers, yeah, there's quite a bit of that. There's no question about it. But I'm a supplier, so I never really got into those discussions between the end user and their labor

force. But I have been told that there is a fair amount of pushback, a lot. But it also depends how you frame it and sell it to the workforce, too, as well, whether it's a security issue, and not a big brother watching over you, and it won't be used against you, unless, of course, you do something illegal. But it's more of a buying/selling type of approach that they're taking to go ahead and get the buy-in from the workers.

MR. HOFFMASTER: We had one incident of the labor defining -- a labor union defining what could and could not be on video. So, in a chemical facility, we were restricted to just perimeter control; we could not look inside the perimeter. Now, will that change over time? It may change over time, as threats change.

MR. SAND: Last question.

MR. D'ANGELES: Yes. My first one's to Jennifer. It's not really a question, but I'm Dave D'Angeles, with the Department of Homeland Security. First 3 years of Department of Homeland Security, I spent a lot of my time with LAPD out in Los Angeles. And you made the comment about not having the cameras capable for monitoring in the police cars, that you didn't see much value out of it. One of the areas -- and I'd like to discuss with you more -- that we saw, where it absolutely was one of the best tools, was for the incident command and also for potential surveillance, and this was where the cop on the beat could look at his camera in his car -- when there was a possible surveillance, they could point a camera to that, and also, when there was an incident, they didn't have to meet with the incident commander, they actually could go on the screen, and they'd view the cameras and go, Here's where we're staging. But I'd like to talk to you more about that.

My other question is to the group. We touched on it lightly, but every time there's a security convention, you'll see a million people selling intelligent video. And there's all guys running out of trees and everything like that, and they make it sound like you just plug it in and it's automatic. Some of them are honest, and they tell you it's different. But there's no background that is static, and that's one of the problems. And the other thing is, I'd like an explanation -- I work this, but how much time it takes to actually train cameras in the clusters once you get the intelligent video software, just some of the remarks on that.

Thanks.

MR. DOCKNEVICH: Well, you've raise some very valuable points, very important points. I think everyone familiar with the Gartner Hype Cycle, you know, you get the hype in the industry, and everyone's, like, so excited about it, and everyone buys it, and then they go, Oh, golly gee, you know, it really doesn't work like the vendors say. Well, that's, kind of, where we are with analytics.

And now we're moving beyond that by, (a) finding real-use cases where analytics can work, (b) continually improving the technology and the solution so it does what the user wants it to do, and then deploying it in a pilot mode and learning how to tune that analytic and the

case that you wrap around it to provide the benefit. That's why when you're looking at different analytic packages, they do all work differently. Some of them, you go their Web sites, they say, We do 14 things. Well, great. Well, you know, our approach has been more open -- our application is a J2E services-oriented architecture, open application that we can then customize, because every use case is different. So, I think we're past the hype, we're kind of on what downside, and now we're actually seeing pilot deployments providing real value to law enforcement and to end users, where people are saying, Wow, this actually works now. So, we're going to start seeing that, and we're moving up the other side of the curve.

MR. SAND: Could we have the last question, and then the last question?

[Laughter.]

MR. SAND: And I ask the questioners to be quick. We're going to be around during the break and the rest of the conference.

MR. YANG: Good morning. My name is Press Yang, from FEMA IT. This question is primarily directed to Mr. Hoffmaster and Mr. Sand. I'm assuming that when we talk about the municipalities and the law enforcement and public safety agencies, I'm assuming these projects are funded, in large part, by Federal dollars designated for homeland security. If that's the case, are there any types of standards and any thought of, maybe, connecting these different networks together?

MR. HOFFMASTER: For the DHS projects, there is a gold-disk standard. And it's probably one of the reasons that I'm inclined to look at those security systems to be standalone and not integrated with a LAN on a facility. It's -- and, plus, you have to get the DHS IT gold disk to work on a NMCI or some other standard.

Plus, the more security that's put on a system is all part of the planning. Scale the security -- IT security -- along with what you're trying to preserve. And so, there is a gold-disk standard. So far, on the projects that we've deployed, it is not incorporated onto a corporate LAN, it's a standalone system. And I should say that there is actually a very -- speaking of cost controls -- with a 30-day password refresh -- you may have 200 users -- it becomes very time-consuming, manpowerwise, to keep those passwords refreshed.

MS. MULLIGAN: I have two. They are quick questions. One is given that the analytics on the back end are really not there yet, to what extent are you using to use lower-level sensors instead of altogether reliance on video, which clearly has real limitations? One of you could, particularly Mr. Hoffmaster, I think, about some of the integrated projects, and perhaps some others?

Sure. Dierdre Mulligan, UC-Berkeley.

And then, second, sabotage. Some folk -- you mentioned, in particular, or you've thought that we need to be going to a more covert system, and it seemed that you were particularly concerned about people sabotaging systems, and I wanted to hear if there are any stories about sabotage.

MR. HOFFMASTER: I don't know of a particular story of sabotage, at least not -- other than, maybe, in a war zone. But all I'm saying is that the potential for sabotage becomes very plain; if you can see it, you can defeat it. There was -- what was the earlier part of your question? It had to do with –

MS. MULLIGAN: Yeah, well, given that everybody has said the analytics on the back end –

MR. HOFFMASTER: Oh, the sensors.

MS. MULLIGAN: -- are really not there. And, I work on sensor net technologies –

MR. HOFFMASTER: Right.

MS. MULLIGAN: -- as well as video surveillance. And to what extent are you seeing integrated projects?

MR. HOFFMASTER: I think that's an excellent observation. I think that sensors are used for layered detection to reduce the amount of false alarms is a very smart thing to do. And it doesn't push the kind of data that video analytics pushes, so it becomes much more feasible to use seismic pressure, acoustic sensors, even IR sensors, though they are pretty -- what I'll call a dumb sensor -- right? they're not very sophisticated. If you layer them with a number of things that show that I have motion from an IR, I have pressure from something nearby, somebody has touched the fence, because I now have a wire sensor on the fence, that if you layer those correctly, I think that you can greatly reduce the amount of backhaul that's required and the amount of video that you need.

MR. DOCKNEVICH: Yeah. Think of it as, you have all of these tools in your toolkit, so pick the tool that does what you need to get out of the solution at the least amount of cost. I mean, I think there is always a rush to video and a rush to analytics, when there's other solutions that'll work even better.

We were working with a school system, where they had those $40,000 waxers get stolen out of their supply closet, and, Oh, we're going to put cameras in and have video analytics. You know, it's an RFID tagging solution might be much better, cost-effective solution.

Another great example is working with petrochemical, where they want to monitor thousands of miles of pipeline. Do you really want to have cameras, you know, every 50 feet for a couple of thousand miles? Obviously not. So, you use, you know, sensors that smell for the petrochemical molecules, and they can do that job much better.

So, yeah, that's part of this initial, What problem do I have? How do I want to solve it? And what are the different options, technologically, to get me there? That's a great question, though.

MR. HOFFMASTER: Yeah. And also know that there are analytics on sensors, now, similar to a submarine sonar kind of solution, where you have it's analogous to the video. The video has to understand what it sees, and then it has to digest what is normal and what's abnormal. You can do the same thing with noise. You can filter noise. You can take the normal, everyday noise that a sensor sees -- or hears –and you can filter that our and look for the anomalies. That's a very good question.

MS. KING: Yes, a number of of jurisdictions are using gunshot detection in their networks to help with deploying resources and cameras, because their focus has been on violent crime and solving violent-crime issues, and that seems to be an interesting way of approaching the problem. And also, possibly, is to preserve privacy, because you're looking at very specific events, and not just taking everything in.

MR. SAND: Well, I want to thank you for your time and attention. I hope this created a foundation for the rest of the workshop. And I personally want to thank the panelists for putting in their time and their effort into this.

Thank you.

[Applause.]

MS. LEVIN: Try and return in 15 minutes we'll resume, at five of.

Thank you.

[Recess.]

MS. LEVIN: Thank you all. We're going to get started with the international panel now.