



Achieving a Resilient Cyber Ecosystem: A Way Ahead

Luncheon Keynote Address

Peter Fonash

I am not going to talk about payments. I am only going to talk about cybersecurity in general, and some of our efforts at the Department of Homeland Security (DHS).

First, I am going to talk about our responsibilities within DHS. I come from an organization within DHS called Cybersecurity and Communications. Within the federal government, there is a split role for cybersecurity. Each department on the dot-gov side, on the civil sector side, has a chief information officer (CIO) who is responsible for protecting in networks. The FBI and the Office of Management and Budget (OMB) also have roles. Our role, first of all, is to protect the dot-gov; in addition to the CIO's responsibility, we provide common services across the dot-gov domain. We also work with the intelligence community, law enforcement, as well as commercial partners, like the Financial Services Information Sharing and Analysis Center (FS-ISAC). We work closely with the FS-ISAC. In that role, we provide protection and we have a program called Einstein. You have probably seen that recently in the newspapers. Einstein provides perimeter protection; it is an intrusion prevention system. We have done something called "trusted Internet connections"—an initiative to reduce the number of connections to the Internet from agencies. In general, agencies are being forced down to two connections per agency. Einstein would be placed in line with that connection, and additional perimeter protection also would be in a "trusted Internet connection." That is the second thing we do.

The third thing we do, in terms of programs, is called Continuous Diagnostics and Mitigation, which gives you, at the enterprise level, a set of tools that, if you are familiar with the SANS Top 20, implements about 16 of the SANS Top 20. It does not address mobile security, but it gives you the ability to identify assets, to ascertain the vulnerabilities of those assets

and to do patch priorities. It reports to a dashboard up to OMB what is going on in that federal agency, and how protected they are.

We also run the National Cybersecurity and Communications Integration Center (NCCIC), which is composed of three pieces. The first, which is probably the most well-known, is the U.S. Computer Emergency Readiness Team (US-CERT). US-CERT is responsible first as a watch-and-warning function—watching what is going on in the Internet, and trying to give warnings if there are vulnerabilities detected or particular attacks detected. We also are going to start providing information in automated fashions, for example, reputation information. We are collecting information from many commercial sources on reputations, in other words, reputation of IP addresses, and we are going to be providing that shortly.

The second piece is the Industrial Control Systems Cyber Emergency Response Team (IC-CERT), and the third piece is the National Coordinating Center for Communications (NCC). I was at the NCC, and we were transferred from the Department of Defense (DoD) when DHS was created after 9/11. So, there is a legacy organization within NCC, up and operational, which is the communications ISAC; it also has responsibility for Emergency Support Function 2 under the National Response Framework (a guide to how the nation responds to disasters and emergencies). When there is a natural disaster like a hurricane or cyberdisaster, the different emergency support functions are activated, for example, transportation and health, and we respond and are responsible for managing the reconstruction of communications. Within that activity, some things we did were: during 9/11, we did the communications restoration for Wall Street and we had the responsibility for restoration of communications during Hurricane Katrina.

There also is the Office of Emergency Communications, and there are two priority service programs it runs. One is the Government Emergency Telecommunications Services (GETS), and some of you, I think, have GETS cards. That is for wire lines. And then there is WPS, Wireless Priority Service, which is for your cell phones. If you qualify for those programs, you can get priority communications over wireless and landline. The Federal Reserve has used those services in the past for restoration injection of currency into the marketplace.

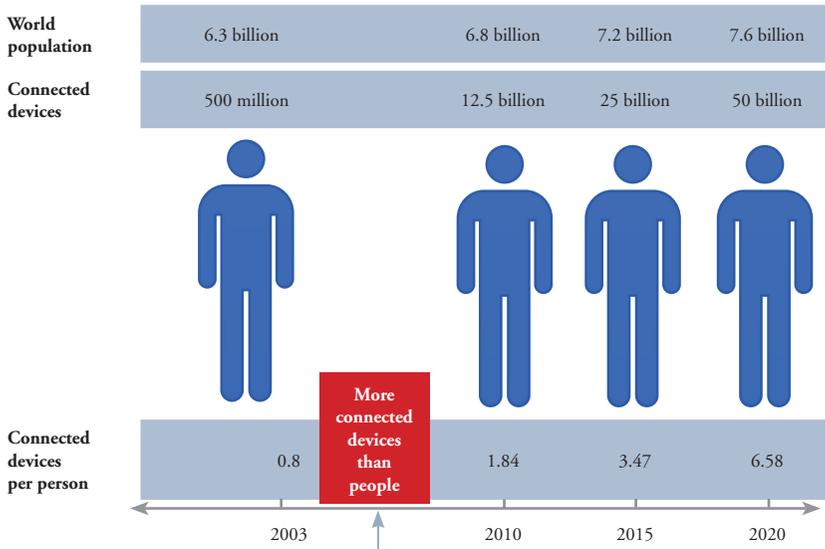
Now I am going to talk about what we call the cyber ecosystem. There are two reasons why I am going to talk about this. The first is that we are all in this cybersecurity problem together. Even though you think you are

secure, you have to make sure your supply chain is secure. You have to make sure your partners are secure because if you look at the Target intrusion, for example, it was not Target but instead one of its vendors that was actually intruded. And there are many, many cases in which the actual organization was not the one that was actually invaded, but it was through another mechanism. So, we are all in this together. It is an ecosystem, and we need to raise the overall security of the ecosystem. The second reason is that in addition to protecting dot-gov and critical infrastructures, we try to protect the general public and, in general, cybersecurity services in the United States. What we are trying to do with the initiative is to raise the efficiency and effectiveness of cybersecurity for the whole country.

I am going to try to go through where we are and why we should be concerned about doing things better. I hope everybody has heard about the Internet of Things (IoT). The point of this is that we have problems today in effectively providing security for controlled enterprises. Where we are going with the IoT, there are going to be all these devices—cars, refrigerators, home heating systems—that currently are under no one's security control. The number of devices is going to be in the billions, actually 50 billion (Figure 1). The figure shows we are really at a curve in terms of the use of the IoT: we have dramatically increased the use of it, and it is under nobody's security control. You are going to see auto manufacturers do things about IoT and address the safety of their cars. That is going to be a real problem as we get to auto-driving and things like that. You also actually can be attacked by your refrigerator some night when you go down for a snack. So just be aware. Attacks are continuously expanding. It seems like we get a new attack every week. Basically, the data breaches are increasing both in numbers and in scope.

In terms of how we are doing on cybersecurity, how we are protecting ourselves, there was a survey that said budgets in 43 percent of organizations are going to be flat from 2014 to 2015, so there is no additional money. Five percent are actually going to cut their cybersecurity budget. And 53 percent said they do not have enough people to do the job. We get into this efficiency issue.

It is interesting to note that based on the numbers from US-CERT, which are a couple years old, we had more than 160,000 reported incidents a year, and those are just the ones reported to us. There were far more incidents going on than the ones reported to us. Chart 1, taken from the

*Figure 1***The Internet of Things was ‘Born’ Between 2008-09**

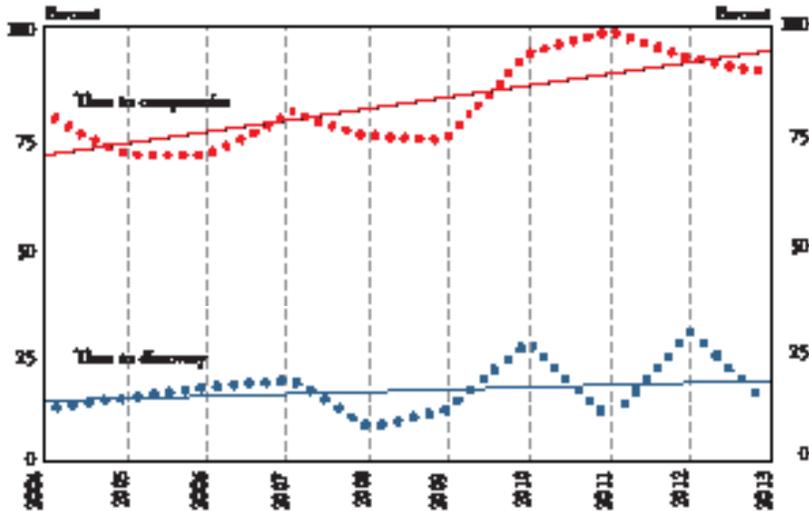
Source: Cisco IBSG, April 2011.

Verizon 2014 Data Breach Investigations Report, shows that in 2004, about 20 percent of the time we were able to detect intrusions in a day or less. And the bad guys were able to get in and attack about 70 percent of the time. There was a gap of 50 percent in how effective they were versus how effective we were. In 2014, that gap had grown. We were about 25 percent effective, and the invaders/intruders were about 90 percent effective. They have gotten much more effective and much more efficient than we have. They were better before, and they are much better now.

The challenges here, and why we need to get more effective, are the following. First, the security analysts that we have, every organization has, have incomplete knowledge of their individual organization as well as what is going on in the Internet in general. Second, adversaries are getting better and faster than we are. Our ability to detect and respond to intrusions is way too slow. There are some charts that show the average detection is 205 days, and the bad guys are getting in and out in a few days. That is a real problem. There is enormous growth in the scope of the potential cybersecurity intrusions because of the IoT. Third, trust among organizations is not sufficient to automatically share defensive courses of action; we do not share information. There are legal reasons why we do not, but there also

Chart 1

Percent of Breaches Where Time to Compromise (Red)/Time to Discovery (Blue) was Days or Less



Source: Verizon 2014 Data Breach Investigations Report.

are trust reasons—do I trust that you will protect that information, will you use it appropriately? People also are afraid they will give away some competitive advantage if they provide this information. Fourth, there is no resilient infrastructure that can support assured communications. What I mean by that is yes, we have those priority service programs, but right now we are moving from circuit-switch technology to next-generation technology, VoIP-type technology or IP-based technology, and we are not going to have the programs in place for about three or four years to provide those next-generation capabilities. So, the communications infrastructure is vulnerable to attack. I think 2017 is where we plan to start having operational capability. However, until then there is going to be a gap.

What I am going to propose is that we need to improve the effectiveness of cybersecurity. We need to make the analysts more productive. We need the ability to reduce the time to detect and respond from months to days or minutes. We also need to have much more innovation than we currently have in terms of the insertion of the innovation. There are a lot of innovations going on in the research community. I probably have six different companies come to see me every week—some of you probably have the same thing—telling me about new technology. But actually getting it out,

using it and putting it into an existing system is yet another challenge. We need to be able to better manage that process of innovation insertion. We do not manage our risks very well because many times we treat all data as equal, but all data are not equal. We need to move away from that model; we need to move to a risk framework.

How do we propose to do this? We feel there are solutions we can provide if we can get industry consensus on these things. For example, we need to get interoperability, automation, trust and information sharing. If we get those things, we will have much more effective and efficient cybersecurity than we have today.

What I mean by interoperability is that the tools we have today mostly are not integrated. Our analysts get data from different sources in different formats from different tools. We have to integrate those. Analysts are spending too much time manually changing the data or interpreting why these data look different than those data, even though the data are the same. That is why we have a manpower shortage; a lot of time is spent on rote efforts as opposed to analysis. If we can get to interoperability of tools, all with common semantics, understanding and syntax of data, then the tools can seamlessly provide data to the analyst. The analyst then will have a common understanding of what that information means as well as the tools.

Once you have interoperability, you can go much more to automation. We want to get to automated courses of action. For common events and common occurrences, we want to be able to detect something and then respond to something in an automated fashion. We do not want the analyst involved. We want analysts to be addressing the hard problems: we want to move the analysts away from being just involved in the rote activities to where they are actually being analysts and actually seeing unusual things. We also want to move to machine learning, so that the machines understand things better, see things and learn the analyst's intervention. After the intervention, a similar intervention is no longer needed because that would now be part of the machine learning of that environment. The machine learning will then allow the machines to take that automated course of action.

As for trust, go back to the idea that the analyst only has a partial understanding of what is going on in the rest of the Internet. We have to get to the point where we do much more information sharing, and to do that we need to have trust in partnerships so that people are willing to share

information. But we also have to resolve the technical issues of authentication mechanisms. Even if you have authentication mechanisms, if you do not have the trust, you cannot share information.

Once you have enabled interoperability, automation and trust, then you can really get into information sharing. That information sharing basically will be in the physical side of DHS. We want to use the motto “see something, say something” for cybersecurity. In other words, if you see something, we want you to report that to the rest of the community so they can take action on it and patch that vulnerability so that potentially they do not even get attacked.

Where are we today in interoperability and where do we need to go? There is something called orchestration, applications that turn tools into tool sets. The orchestrators basically manage—orchestrate—the activities of the suite of tools. They have to develop configuration files and things like that so as to get a set of tools to work together. You have to spend significant efforts in getting the orchestrators. Every time you bring in a new orchestrator, you have to redo that work. Where we want to get to in interoperability is that we have this common data model, common application programming interfaces (APIs), the tools just plug and play, and so the orchestration is automatic. We are going to talk about tools that do sensing, sense making, decision making, and an action. You want to have a set of tools that do these; you want to have a tool that senses an intrusion, then a tool that makes sense of that intrusion, then a tool that makes a decision on how to block it, and then tools that implement those decisions. That is where we want to go.

Where we are today in terms of future automation, again, we are at the orchestration level, but we want to get to automated response. This area is very controversial to many people who have concerns about unintended consequences. The National Academy of Sciences, and just about everybody else, has told me that is an issue. For example, if I detect something, I direct my firewall to do something and that firewall starts blocking normal corporate email. The unintended consequence is that normal business email is now being blocked. I did not think that was going to happen. It is an unintended consequence of an automated action. We need to get to the point where we have a much better understanding of what automation means and what are the consequences of that automation. We also have to have mechanisms to allow us to reverse automated actions, so we

can remove them very quickly once we see unintended consequences. We talk about getting the human on the loop as opposed to the human in the loop. Right now, the analyst is in the loop so that the human gets involved in making the decisions. We want to get to the point where the human is on the loop observing what is going on. That is where we want to move to.

In terms of trust, we have a lot of partnerships with the ISACs, and now there are going to be ISAOs, Information Sharing and Analysis Organizations. We are putting out a grant on ISAOs and we will be bringing out best practices through the ISAO Standards Organization so that information sharing can be done in an organized manner. The financial sector, by the way, I think, currently has probably the best information-sharing organization. The energy sector has a very good one as well, but you guys are clearly one of the leaders in that. We want to get to the point where we automatically trust those organizations. What I mean by automatically is I get information, and then I take that information and act upon that information. We are not there yet, but that is where we need to get to.

In terms of information, the right data will arrive in time to take that automated action. So see something, say something; you send that information out, and automated action is taken. That is where we want to get to. Everybody has a common understanding of what is going on.

So, then future communications. Right now we are transitioning from a circuit-switch technology to an IP-based technology. There will be some delay in capabilities for a while, but we need to have resilient communications because the assumption has always been that during a cyberattack you have communications and your security operation center is able to direct the response and recovery. What if they take out the security operation center, take out the communications? So, you need to address that too.

How are we going to do this? We, the government, are going to facilitate our ideas, but we want industry to lead. I am going to make the pitch that we are going to work with the IT industry on this, but we also want the customers of the IT industry—the banks and we are trying to get the healthcare industry as well—to say they want this because we believe we need to go there, but it is going to be market driven. There are reasons why the IT industry does not want to go this way, because right now they can sell proprietary solutions, and they make more money on proprietary solutions than open-based solutions. If we go to open-based solutions in

the very, very competitive IT industry, it is a market share issue. But we feel that the customers want this. I have talked to several banks and they seem to think this is a good idea in terms of where to go. We sent a request for information in January, and 58 companies gave us comments. We also had a roundtable with a much smaller group of industry. Banks were represented as well as the IT industry. It seems like the banks were in support of this. Even the IT industry was expressing interest. I think the IT industry is starting to see security as a service as opposed to providing a tool set. And I think what you are going to see is that as we go to security as a service, they are going to be much more open to having open systems, no pun intended.

So, we want to get to the point where we go from months to minutes and milliseconds in terms of our response capabilities. Part of the overall architecture, as we see it from the DHS perspective, is your example enterprise security system, which could be on the enterprise or in the cloud. You can virtualize the system into the cloud. You have sensing tools, sense-making tools, decision-making tools and acting tools, and they are managed by that management orchestration, and then there is a common database there too. The enterprise security system does boundary protection, infrastructure protection, host protection, endpoint protection. So, within that, there is a lot of information being shared in real time. It also provides information out to other partners, as well as to what we call the cyber weather map.

Our Deputy Under Secretary Phyllis Schneck talks about the cyber weather map. The idea is that we want to model ourselves like the National Oceanic and Atmospheric Administration (NOAA). NOAA collects a lot of information from a lot of different sensors across the country, and then has a model and runs forecasts. So, we are collecting information from the dot-gov domain, we also are getting information from the intel community and law enforcement and we are buying commercial information about what is going on in the Internet. We are starting to combine that information. We are not where we want to be, but we are collecting all this information, and then we will do analytics on that information. We are going to provide it to the enterprises, and we are also going to provide it visually. In the first part of what I call integrated adaptive cyberdefense, which, I should say, is a concept that we have been working on and partnering with the National Security Agency (NSA), there are three pieces; the enterprise piece, the weather map piece and what we call the AIS (Automated Information Sharing) piece, or the infrastructure that shares that information.

We are working this concept with NSA and we are demonstrating the concept in an integration lab at Johns Hopkins University's Applied Physics Lab. We are talking to different partners and doing pilots of this technology. We want to shift it to where we are actually getting faster than the attacker. We have done demonstrations and automations of this: in the laboratory/operational environment as part of the Applied Physics Lab we have been able to detect and self-defend attacks in less than a minute in the best case, and eight minutes in the worst case. In terms of sharing Structured Threat Information Expression (STIX) indicators, which are a threat sharing mechanism protocol, we have been able to share that information in less than two minutes in the best case, and nine in the worst case. We have constructed pseudo communities of interest, and we have been able to share that in less than a minute in the best case, and 45 minutes in the worst case just because of the architecture.

That is where we want to go. When we have looked at the effectiveness of this, and again this is in a laboratory environment with some operational capabilities, we have dramatically increased the productivity if you start multiplying those factors by that much.