UNITED STATES OF AMERICA

COMMODITY FUTURES TRADING COMMISSION

STAFF ROUNDTABLE ON

CYBERSECURITY AND SYSTEM SAFEGUARDS TESTING

Washington, D.C.

Tuesday, March 18, 2015

PARTICIPANTS:

VINCENT McGONAGLE
CFTC

TIMOTHY MASSAD
CFTC

J. CHRISTOPHER GIANCARLO
CFTC

PHYLLIS DIETZ
CFTC

SUSAN STEWART
CFTC

JAMES ORTLIEB
CFTC

ROBERT WASSERMAN
CFTC

DAVID TAYLOR
CFTC

MICHAEL DANIEL
White House

WILLIAM NELSON
FS-ISAC

BRIAN PERETTI
FBIIC

MARK CLANCY
DTCC

LEO TADDEO
FBI

GERARD BRADY
MORGAN STANLEY

PARTICIPANTS (CONT'D):

STEVEN CHABINSKY
CrowdStrike

MURRAY KENYON
NSA

DAVID GARLAND
CME Group

GREG GIST
CitiGroup

CHRISTOPHER KINNAHAN
FBIIC

DAVID LaFALCE
DTCC

RANDY SABBAGH
Schwab Technology

JOHN RAPA
Tellefsen & Co.

KEVIN GREENFIELD
Office of the Comptroller of the Currency

DAVE EVANS
Bank of England

JERRY PERULLO
ICE

THOMAS MILLAR
US-CERT/DHS

RONALD ROSS
NIST

RYAN LIBEL
CME Group

PARTICIPANTS (CONT'D):

ANN BARRON-DiCAMILLO
US-CERT/DHS

* * * * *

P R O C E E D I N G S

(9:08 a.m.)

MR. McGONAGLE: Good morning, everyone. Welcome to the Staff Roundtable hosted by the Divisions of Market Oversight and Clearing and Risk. The topic for today is cybersecurity and system safeguards testing; and we have some introductory remarks from our Chairman.

MR. MASSAD: Well, good morning, everyone. Thank you for being here; thank you, Vince. I think we all know that cybersecurity is the most important single issue facing our markets today in terms of market integrity and financial stability. The need to strengthen the security and resilience of our financial markets against cyber attacks is clear. And the examples of cyber attacks unfortunately are all too frequent and familiar, whether it's JP Morgan or Home Depot, Target, Sony, both within the financial sector and outside. Some of our nation's exchanges have been hit or suffered other technological problems that have caused outages or raised concerns. And

because of the interconnectedness of financial institutions and markets, an attack in one place can obviously have significant repercussions throughout the system. And I guess what's most concerning to many of us is that, while we know some of these attacks are motivated by people whose aim is commercial profit, some are clearly motivated with the aim of simply to disrupt or to even shut down services.

Now, we at the CFTC have responded in a number of ways. We have incorporated cybersecurity standards into our regulations, our core principles now include them, we've required clearing houses and exchanges to maintain system safeguards and risk management programs, to notify us promptly of incidents, to have recovery procedures in place. And we've also made this a priority in our examinations. But, you know, the responsibility for cybersecurity obviously rests with private institutions. As a government agency, we can set standards, we can engage in examinations, but it is up to the private

institutions that run critical financial infrastructure to do the daily comprehensive work that's required. And that's especially true when it comes to testing. Testing that some would say only works when the institution fails, meaning when it is pushed to the point that you truly identify weaknesses or a penetration occurs so that then you can remedy a problem.

And that brings us to today's Round Table discussion. So we are seeking industry and government views on cybersecurity matters, but in particular, on systems testing. The staff is interested in the panelists' thoughts on what constitutes effective and adequate risk analysis in testing by exchanges and clearing houses in particular. And we also want to hear thoughts on what should our role be in promoting testing. Can the agency contribute to cyber readiness by establishing more detailed standards for systems testing? And how do we make sure those standards truly add value to cyber readiness and not simply more work for IT specialists?

So we're delighted to have you here. I want to thank the panelists, in particular, for contributing their time and expertise. I want to thank our staff for all their hard work in putting this together, and I look forward to today's discussion. And let me -- I think Commissioner Bowen is -- did she want to say -- she stepped out but I know Commissioner Giancarlo wanted to say a few words.

MR. GIANCARLO: Thank you, Chairman. Good morning, everyone. Today's Round Table is timely and critically important. Working to make U.S. derivative markets more resilient to cyber attacks is essential to the mission and oversight of the CFTC. And I commend the leadership on this issue by each of my three fellow Commissioners, starting with Commissioner Wetjen, for drawing attention to the issue during his tenure as Acting Chairman, and to Commissioner Bowen in her work in establishing the Market Risk Advisory Committee that has identified this issue as a key part of its mandate. And to you, Chairman Massad, for

making cybersecurity a Commission priority.

I'm interested today to hear from numerous experts on the panels and I thank them for their preparation and their participation, and I thank the staff as well for putting together a terrific panel today. I hope to learn about the range and nature of cyber threats, from cyber crime and vandalism, to terrorism and outright cyber warfare against U.S. and global capital markets. I'm interested to hear about the latest defensive tactics and emerging best practices for market participants in this rapidly evolving and morphing area. And I'm interested to explore how we best balance effective cybersecurity of execution venues and clearing houses without sacrificing marketplace vibrancy and fair access to trade execution and clearing.

And I apologize in advance that during the course of the day I may need to step out to take care of some business, but I will try to be here for a good portion of the day for this very important program.

I thank you all.

MR. McGONAGLE: Thank you, Commissioner. I'll turn it over to Bob for the first panel.

MR. WASSERMAN: So first I'd like to thank Chairman Massad and Commissioner Giancarlo for those remarks. I'd also like to thank everyone for coming today, in particular our panelists. We have, I think, an extremely talented group of panelists here and I expect that today's discussions will be of considerable assistance to the staff as we work to develop proposals to strengthen our rules regarding testing to protect our regulated infrastructures against cyber threats.

I'd like to start with some very important administrative announcements. First, as a public service, we have Wi-Fi available. Instructions are available on the written agendas that are on the table near the door as you came in. We will, during the course of today's proceedings, be taking written questions from the audience in this room and we will endeavor to

insert some of those questions toward the end of each panel. You should find a note card on your chair, and there are additional note cards on the table near the door. If you will please write your questions down as legibly as possible and pass the card down the row to one of my colleagues who will periodically be coming to pick those up. Restrooms are outside this room to your right as you leave, and then at the end of the space to your left. We have some limited quantities of coffee and tea in the back as well as water.

Panelists, if you could please press the button to activate your microphone when you speak. This Round Table is being audio cast to folks who are calling in and they can only hear you if the microphone is on. And if you forget to turn it on you may see me pointing at my ear to remind you. On the other hand, please turn your microphone off when you stop speaking, as we can only have a limited number of them on at a given time. Finally, if you use abbreviations or technical terms, please explain them the first time you use

them, as some of us are a bit less familiar.

I should note that while my colleagues and I will be asking questions and may express tentative views, anything any of us says represents at most only our personal views and does not represent the view of the staff as a whole or of the Commission. I should note as well that we'll be making a transcript of this Round Table which will be posted on the CFTC website. And, finally, we will also be making the video from this feed available eventually on YouTube. Previous videos have accumulated hundreds of views and I imagine this will.

Okay. I think I would like to wait maybe for five minutes because we have one or two panelists who we're still waiting for, so if you could give us five minutes and then we'll still begin a few minutes early if that's okay.

MR. WASSERMAN: Okay. Bill, if I could turn to you to tell us a little bit about the context in which we're operating here in terms of cyber threats.

MR. NELSON: Sure. A little bit about my organization, because as you said a lot of people don't know what FS-ISAC stands for, but we're the Financial Services Information Sharing and Analysis Center. Information sharing and analysis is our middle name. We've been around since 1999. I joined the organization in 2006. I have to tell you, in 2006, there wasn't a lot of information sharing going on at the time. In fact, if a member shared some threat information, we'd literally throw a party; it was such a rare event. That's changed, a lot of it has changed because the attacks have grown more frequent; we're seeing some of the same attacks, but the criminal attacks are still there, cyber criminals. Hacktivists were something new that really emerged I'd say in the probably 2009-2010 timeframe. And some nation state attacks have hit the financial services sector too and other sectors as we know. What we do is, a member that has an incident occurring, they typically will share that with the other members often on a distribution list, an

email distribution list, and then we do more research on that attack or incident, and we push that out as an alert to the rest of the members. We also work very closely with our government partners, including the FBI, Treasury, Department of Homeland Security, and others. We push out joint products from time to time. I think our contribution many times in those joint products are, what are the risk mitigation recommendations to address that particular risk out there. You know, the types of things that we share, or what we typically call threat indicators, are things like an attacking IP address, it would be a subject line in an email that's used for social engineering to trick you to click on a link to a malicious site. It might be the malware itself, the executable file, to look for that and delete it. We don't share personal identifiable information at all. We're really strictly about sharing attack data, threat indicators.

One of the challenges we have in our system, or really any sharing of information, is

the bad guys can get in your system within a matter of seconds or minutes, and there's a really time-scale challenge here. How do we get that information out and how do you get it into your system today to block it or to delete the malware. Getting in takes seconds, minutes. Discovering it, doing something about it, can take hours, days, weeks. And that's something that we've actually teamed up with the DTCC on an automation project to address that issue. Instead of taking that long to really try to do machine to machine sharing so it can go right into your security systems to block the attacker.

You know, looking ahead at the threats is something that I think we're doing a better job. I think a lot of times you -- maybe as a regulator you see this sometimes -- you're reacting to yesterday's threats. We really need to address the future. And we are very concerned about some of the things that the Commissioner mentioned. The Sony attack for instance was destructive malware. We've done a number of

exercises in the last couple of years looking at destructive malware and data integrity issues. We've done that with the sector, with the regulators, and will continue to do those. There are a number of exercises planned for this year including what's called Quantum Dawn 3, also Hamilton Vault, and a number of drills that we're doing all year long working with our government partners and industry.

I should mention that we do work with regulators sometimes on a membership basis too. The Federal Reserve, the FDIC, the OCC, are members of FS-ISAC. You may want to consider membership at CFTC. That's my plug for FS-ISAC. The only thing we ask or require is that if it's very sensitive information, and we have a way we call the traffic light protocol that we rank and classify all the information we share -- and that's become a standard I think within government too; FBI and Treasury use that -- that the information not be shared with examiners. We just have your critical infrastructure people look at

it. So that will be the only requirement.

That's really all I had. Just a kind of description of what we're doing. The membership also has grown. We've added 1,500 new members. These are organizations, not individuals; organizations in the last year. We affectionately call it the membership tsunami. It really started because of the FFIEC regulators including Federal Reserve, FDIC, OCC, really pushing membership in the FS-ISAC as part of your defense and depth of strategy you should have.

That's it.

MR. WASSERMAN: Thanks. And I'd like to turn at this point to Steve Chabinsky who is General Counsel and Chief Risk Officer for CrowdStrike to basically discuss the types of cyber threats that the financial industry, in particular financial infrastructures, are currently facing.

MR. CHABINSKY: Thank you very much. And first I'd like to thank the CFTC itself for its vision and for preparing this Roundtable today

and bringing these important issues to the table.

The threat landscape evolves constantly and we ended up I think over the last few years being a bit surprised at how it's evolved for the financial industry. We've, of course, always known that the financial industry is ripe for attack, or intrusion I should say, from criminal groups, always after the money. It makes all the sense in the world. You know, going back to the old Willie Sutton apocryphal statement -- I'm not sure if it's true or not, right -- why do you rob banks, because that's where the money is. And we saw really quite good resiliency from the banks, meaning that the financial crime that we tend to see focused on the user accounts, the weakest point in the chain, getting passports and the like, doing man in the browser attacks, where it's the end user whose computer ends up being infected so the passwords are taken. And then from the bank's perspective, the transaction looks normal, it's being accomplished through user credentials.

There has been a shift, however, as

we've seen both in terms of attempts at destructive attacks roundly attributed to Iran and DDoS against banks, potential motivations of course being the political landscape, reflecting then that there's another force that's going on here, meaning that political will ends up becoming a motivator for the attacks that could be against financial institutions. It's not just about the money any longer.

And then of course we've seen nation states that are quite interested in intellectual property, including trading algorithms, and stealing either by insider access or now attempting remotely. At CrowdStrike, we have seen interest in the financial industry both by the nation states of China and Russia, as nation states looking to penetrate in order to get intellectual property and an understanding of either how the markets are working, or how the systems are structured. And in the worst case actually creating a beach head in case there becomes more political division, which would be

the most touchy of cases because what we've seen more recently is there has also been an increase in destructive attacks against networks, which creates in the minds of many, whether it's criminals realizing that destructive attacks could be used for extortionist purposes, nation states recognizing that it could be used for political will, or in the worst instance, terrorists recognizing that they could do destructive attacks to accomplish their political goals.

So when we're looking at this threat, I think it's important to recognize that the old threats remain, meaning the use of computer intrusions to conduct fraud, but we are concerned with protecting our clients against more deliberate, more pervasive, more stealthy intrusions that are not meant to be noticed and that don't have the traditional indicia of an intrusion, meaning fraud that accompanies it that eventually -- there's only so much fraud that could occur before you start noticing there's a problem. Not the case with nation states that

might have access within your systems for quite some time and are looking to remain there without note.

MR. WASSERMAN: So, Michael Daniel, who is Special Assistant to the President, White House Security Coordinator, I was hoping I could ask you to give us the administration's view of the context for these cyber threats that we're facing.

MR. DANIEL: Sure. So I think there are a couple of ways that you can frame that question up, but I think there are two in particular, one of which Steve was actually just alluding to which is, there is sort of two trends that we're actually watching, one of which is the emergence of cyber and cyber capabilities as a key tool of state craft. It is becoming part of the arsenal, if you will, of pretty much all states, and the capabilities that used to be restricted to those with very high-end capabilities are now sort of proliferating out to more and more states. So on one axis you have sort of the expansion of this capability as a tool of state craft and obviously

many countries have discovered that it is apparently a very useful tool, and so are rapidly building up their capabilities. At the same time I would also say that the cyber threat is becoming broader, more sophisticated, and more dangerous, all at the same time. Broader because we keep hooking more and more stuff up to the internet. The internet -- one of the catch phrases in today's cyber world is "The Internet of Things", but pretty soon, you know, your coffee maker, your refrigerator, your car, they're all going to be threat vectors. So we thought doing -- Steve and I thought doing cybersecurity in a world of wired desktops was hard, now we're going to have to do it in the big data mobile cloud, where everything is sort of connected and interconnected. So that threat surface is now incredibly more diverse. Second, all of the actors in this space are more sophisticated, and I don't just mean on a technical basis, although that's very true. Certainly the days of the simple phishing expeditions with the Nigerian Prince who would

really like you to help him out -- I'm sure your spam folders are still full of those, but most people have moved beyond that. So certainly the technical capability of the adversary has evolved. But what is actually more important is their organizational capacity has evolved. Organized crime has moved into this space and is applying all of the principles that they have learned in many other venues. Nation states themselves are getting themselves more organized. So there's a level of organizational capacity. Somebody the other day actually used the term "the industrialization of hacking" which is actually probably a good term for it, the sort of applications of the principles of division of labor and other things to what hackers are doing. So while certainly the hacker in his pajamas living in his mother's basement is still a threat, that's not actually the primary one that we're concerned about.

And then lastly, it's also apparent that the actors in this space are willing to take

actions that they weren't previously willing to do. So the threat has become more dangerous. You know, five to ten years ago this conversation was largely about the digital equivalent of graffiti, the defacement of websites and other things like that, but now, clearly you have actors that are not only willing to steal PII and commit fraud, but actually willing to carry out destructive attacks like what we saw with the attack on Sony Pictures Entertainment. So certainly in that respect, the threat is actually more dangerous and has a greater potential for causing harm, not just to individuals, but to the country as a whole.

MR. WASSERMAN: So I'd like to turn to Leo Taddeo who is the Special Agent in Charge of the FBI's Cyber Division. And, Leo, if you could tell us a bit about how law enforcement and the intelligence community are working together to help the private sector meet these threats.

MR. TADDEO: Well, thank you, thank you. First of all it's a pleasure to be here. Thanks for the opportunity to address the audience on

this important topic. I want to make two points and then I'll talk about how we are collaborating with the private sector. The first is I couldn't agree more with Mr. Daniel that the threat is more dangerous, more sophisticated, and more capable. But I would say this, when you talk about it in the context of pen testing and other methods of hardening your system, they are rational in that they will avoid hardened targets. They will go to the weakest of the group. So as those responsible for protecting networks that think about this problem, it's very important for them to realize that pen testing as part of a larger framework is critical to hardening your system. Not because they will be perfectly protected, but because if they're protected enough, sophisticated adversaries will look elsewhere.

The second is to point out that you're not alone; you're interdependencies will also affect your overall performance and capability to run your businesses. I'm responsible for the area of New York City which houses a large part of the

financial infrastructure, but that sits on top of the very complex and fragile urban infrastructure, meaning subways, water, electricity, communications, all of the things that a business will need to respond to one of these attacks. So if you have a response plan that depends on people being at a particular location, if the adversary can shut down a subway, if the adversary can shut down a 911 system, you may not have the people you need to actually respond to these emergencies.

So with those two points I'll talk about the public-private coordination. We have learned, as well as the Secret Service, that it's imperative to listen to the network operators to find out what's important to them. So for the FBI and the Secret Service, the first priority is to not create more of a negative impact when we show up than the actual adversary is creating. So we have very carefully listened to network operators to determine what it is we can provide that is of use. And as Bill Nelson pointed out, there are indicators that we collect that we are sharing

with the private sector through the FS-ISAC and through other means. We are very careful to dispel myths about what it means to cooperate with law enforcement. One of them is that network operators lose control of the investigation when the FBI or the Secret Service shows up. That's not true. We work in close collaboration with the general counsel. We know that it's important to stay in business, continue operating, keep those systems up. We don't show up with raid jackets and evidence tape to shut down networks in order to conduct our investigations. So over the last few years, I think law enforcement has done a very good job at changing the way it interacts with the private sector in order to create a positive net effect when we show up. And the main reason we do that, of course, is to create a deterrent, to actually attribute these attacks to the adversary, but also because we want financial institutions and others to call us when they have a problem. There are some surveys out there that show that we're not getting called as often as we should be.

There are times, of course, where state law or federal law requires notification to the government, but in cases where notification is not required, we'd still like to be called. Not getting called means we're blind in certain areas.

So we have a number of reasons to interact more effectively with the private sector. First and foremost is to be more effective, but second is to increase the amount of information that the private sector is willing to provide to us. So I think we've gone through that evolution. We have a long way to go. There are a number of government avenues where you -- or government outlets for this information. I think we need to do a better job of bringing that all together. I think the Administration is doing very important work in bringing that together under a threat integration center that will make it -- give us a common operating picture of the threat. So I think while we've made a lot of progress, we still have some work to do.

MR. WASSERMAN: And Brian Peretti is the

Director of the Office of Critical Infrastructure Protection and Compliance Policy at Treasury, and I know him very well as the leader of the FBIIC. And, Brian, if you could tell us about FBIIC and specifically about how the financial sector regulatory agencies are working together and with the private sector to address some of these issues.

MR. PERETTI: Thank you. I really would like to thank CFTC for really getting this panel together and getting this whole day together.

This is something that is near to my heart to be able to continue to move in this area, to increase the cybersecurity of the sector as a whole, and especially the important role that the futures industry plays within that space. If futures doesn't work, many other things don't work, and it's a wholly interconnected system. And the more we can make all the parts more secure, the more resilient it's going to be overall.

After 9/11, the Treasury Department,

working with the other federal regulatory agencies, created an organization called the Financial and Banking Information Infrastructure, the FBIIC, and housed it within my office, the Office of Critical Infrastructure Protection and Compliance Policy. The purpose for that was to really be able to focus on operational risk issues between the different regulators so that we have a forum to discuss these issues. One of the problems we had after 9/11 itself was that the infrastructure was damaged in New York and we didn't have a natural forum to be able to get together and discuss these key issues. The FBIIC has been very helpful in many instances. Going forward from there, the northeast power blackout, hurricane Katrina, the pandemic flu issues, and then now the cybersecurity situation going on.

The role for the Treasury and the FBIIC is really to help coordinate, foster, and facilitate information sharing amongst the federal financial regulators and the state regulators. Our goal is not to be able to dictate to anybody

how to be able to actually do anything, but to be able to bring together the best minds and the best conversations to be able to help advance the industry to increase the resiliency of the sector as a whole. Our goal is to really try to figure out what are the gaps that exist within the private sector or the public sector, and then try to figure out how to fill them together. We hold monthly conference calls and we hold joint meetings with the FS-ISAC and the FSSCC, the Financial Services Sector Coordinating Council, to foster these discussions, to continue to identify what the issues are that are going on within the sector, and then how to work closer together.

One of the key issues that we've seen is that the state of information sharing is not where it should be between government to private sector, private sector to the government, and between private sector firms between each other. The challenge we see is that there is still the concern of folks sharing information, concerns about the information being shared, and how it's

being shared. We're working with a lot of the private sector firms to first really identify what these concerns are and then figure out how to be able to remove them or limit any problems that may arise from them.

One of the challenges we see is that a lot of private sector firms aren't participating in the information sharing dialogue. And that's a concern because, if you're a network defender, how are you getting the best information possible? The FS-ISAC and some other information sharing organizations are really the key to be able to bring the information to the network defenders in a way which is understandable to them and being able to help them in a way that's going to be very beneficial. The project that's being worked on with DHS and some private sector entities tied to the STIX and TAXII delivery mechanism of information, is machine readable that can go directly into your system, and is something that we've really been striving very hard for. And in fact, Treasury is now sharing their information

specifically in that format to government and private sector organizations like the FS-ISAC.

But the challenge we have still is entering into this conversation and figuring out what the gaps are. How do we perfect this and get this better? We're never going to get to, I think, 100 percent perfection in information sharing, but we still see that there are probably areas where we continue to make it be better. And CFTC has been very helpful in this in convening forums like this and having discussions with the financial industry, the futures industry directly, to be able to go forward with this dialogue to figure out where the issues are and how we plug those issues.

MR. WASSERMAN: So, Brian, I'm going to follow up just for a second because I mentioned I want to get all terms defined. You mentioned something about STIX and that seems like a fairly new term, maybe even from today. If you could tell us a little bit about that.

MR. PERETTI: Yes, yes. I don't know

what STIX actually stands for as an acronym.

MR. CLANCY: It stands for Structured Threat Intelligence Expression. And its companion standard is called TAXII, Trusted Automated Exchange of Indicator Information. Those are both developed out of research from DHS by MITRE Corp.

MR. PERETTI: And the key for that is to be able to push out information in ways that can then go directly into systems to be able to have them used by network defenders in a much quicker format. So in the past what would happen was that if there was information out there either from the private sector or government, it usually made it either into an email or a PDF which would then be sent out. Somebody would have to look at the document, type it all in or cut and paste it, and then run it against their system. Sometimes, of course, somebody would do a fat finger and put in some incorrect information and you would have a problem, which the time in which it was sent out, from the discovery of the information to the time it was deployed, could be a very long period of

time. Now the information is going to be shot out in a much quicker format and much quicker through some trusted systems to be able to go into the receiver's network defense and be able to help plug those gaps in a more real time thing.

The key with the STIX and TAXII was that it was developed, as Mark said, by DHS and MITRE with input from the private sector. So it wasn't a format that we created ourselves and said, here it is; it's something that public and private came together to agree to, to address a problem that was identified. So as I was mentioning about trying to fill the gaps, this is one of those areas in which we had a true public-private partnership to make the sector more resilient.

MR. ORTLIEB: Is it largely to just address the time scale problem or is it also to address other issues?

MR. PERETTI: I guess it's first the time scale problem, and second, the reliability of the information. If somebody has to translate it and retype it in, there is always going to be a

potential problem in that translation. Here, it's going to be coming from trusted sources and then moving through the system. Before deploying on the network defense side, you still may want to run it against other things, but a lot of that time is now collapsed from where it was before for a much longer period.

MR. DANIEL: Yeah, and one of the other advantages of it is that it's a common format that can be used not just between the government and the private sector, but across the private sector. For example, across industries because the fields are common to the structure. It also enables you to -- you know, previously, as Brian was saying, mostly what was being shared were what we called "flat files", meaning, since they were documents in excel spreadsheets, and the STIX format actually enables you to share that in a format that the machines can automatically ingest and populate and run statistics on and do other kinds of queries. So it both enables the sharing of information, but it also enables the archiving of

that information in a way that enables us to do trend data and other kinds of analysis much more effectively on it.

The other thing that it does too by starting to move in this direction, one of the exercises we have ongoing with the financial services industry is actually breaking those STIX fields down and identifying what if any are the privacy concerns with sharing that data. And so in that structured format that allows you to very easily see whether or not there would be PII that could even be in that field, if it's even allowed to be part of that field, which makes it much easier to set up decision rules about whether or not to share that. And so you can create automated rules for handling that and protecting PII more effectively using that format. And so it enables you to figure out what you don't have privacy concerns about because there is no PII in there, and so that makes that sharing much easier. And it allows us to identify the fields where if there are privacy issues we can try to work out

the policies and rules to enable that information to be protected or stripped out so it's not shared with the government, or protected once it arrives at the government, and those kinds of things.

MR. WASSERMAN: So by PII I think you mean personal --

MR. DANIEL: Personally identiable information.

MR. WASSERMAN: Thanks.

MR. PERETTI: And so, you know, this is something that, as we enter into this dialogue more with the private sector and hear more from the panelists going forward, is exactly what type of information do they want and what does a network defender need to make their system more secure. We've heard pretty resoundingly that personally identifiable information doesn't help a network defender protect their system. They want TTPs, tactics, techniques, procedures, that bad guys are using. They want malware hashes that they can run against their system to identify potential intrusions or other problems with their

networks. But the personally identifiable information is not helpful because it's not something they can run against their system and, you know, the less they get that, the better. And that goes into the procedures we're trying to create to make sure that information is scrubbed out way before it can even be potentially, even accidentally, disclosed.

MR. WASSERMAN: And so you referred to a malware hash. If you could tell us a little bit about that and how those get used.

MR. PERETTI: So a malware hash is best to be explained as a fingerprint that a certain code would look like, and you can use that fingerprint to run against your network to see if that fingerprint is somewhere in your system otherwise. And so being able to identify a specific malware is in your system, running a malware hash speeds up the process quite a bit. So those are things in which there are known intrusions. Malware hashes are created from that. You then share it with other firms, they get that

running against their system to be able to pull that information up quickly and see if there is a problem.

MR. WASSERMAN: Okay. Just one thing as you were talking about that, so it sounds like we're getting information out to the private sector in ways that they can use it. What do you see folks doing so far in terms of how they are using it? How fast are we moving in the direction of taking this useful information and moving to acting on it?

MR. PERETTI: So I will defer to some of the users of that information, DTCC or Morgan Stanley, about how effectively they're seeing that information. We're providing information out to the private sector, and we're getting feedback from them to modify our processes going forward, and using that as a virtuous feedback loop to be able to continue to get the information better. I don't think the information is perfect where we are now, but we continue to try to make the process better going forward, and not only from

our side, but organizations like the FS-ISAC provide additional analysis onto the roll of modified information that we push out to even make it more beneficial to their membership.

MR. WASSERMAN: And just to make sure we're clear, so tell us just very quickly about the difference between the FBIIC and the FS-ISAC.

MR. PERETTI: So the FBIIC is a government-only group consisting of 18 federal and state regulators who coordinate homeland security, cybersecurity, other issues that are going forward really from the operational risk perspective. So the area which disaster recovery, disaster prevention, really if you look at the whole in this framework, cybersecurity framework, addressing all five of those key categories.

On the private sector side, there is a Financial Services Sector Coordinating Council who works on issues in the same area, but kind of a forum for discussion tied to policy considerations and other issues regarding the same concepts. The FS-ISAC is described as the operational arm of the

FSSCC, to be able to really be able to push out information and be able to work with their membership to really try to increase the resilience of the sector as a whole.

MR. WASSERMAN: So I think I'm going to ask if Mark and Gerry can take up Brian's invitation to talk about how the information is getting used. And so Mark Clancy is Chief Executive Officer of Soltra, which is a joint venture between FS-ISAC and DTCC, and is the DTCC Managing Director for Technology Risk Management. If you could start us on that.

MR. CLANCY: Sure. And I think what I might do is create some bridges between the description of the threat, the information sharing, and then position us toward the testing topic of which this panel is about. And I'll start with a really bad analogy. So about five years ago, if I was sitting around the room talking to our colleagues, we'd say, you know, our job isn't to outrun the bear, our job is to outrun the other guy. And the assumption there is, if

our threat at the time is we had a single adversary, which is criminals, they were trying to steal things, which quite simply was not hugely impactful particularly in the futures space. The reality, however, is there's more than one bear. And so as we ran away from one bear, we ran into another. And so we had to understand the threats that we faced particular to the types of businesses that we are. So DTCC for example, we operate systemically important financial market utilities including a swap data repository. That faces a very different type of threat than a retail payment system, like the kinds of things that criminals are going after in the case of Home Depot. And much to what Steve mentioned, those attributes of nation state, either espionage activity or potentially destructive type activity. Those are sort of primary concerns for a market infrastructure utility. The reason that we need to know about the threats and have the technical information about what's happening is it gives us the context in what controls matter and how those

controls are working, or unfortunately sometimes not working. And so the way that I look at it, and very specifically, we have an operational need to consume the data and see if similar activities occur in our environment in specific fact, in general pattern, and then understand if our controls are effective against countering or minimizing impacts from those threats. And, you know, here in a panel at the end of the day, the reason that business continuity is on the agenda -- because one of the components is how to create resiliency so that if adversaries get into the environment, cause some harm, that we can continue to operate markets successfully.

So the linkage for me is the threat informs what we need to know, that dictates the information we need to have to respond to the threat, which then leads us to the controls. We need to have an environment to be nimble in responding to either recovering, preventing, detecting, or recovering from such an attack. And to tie it to testing specifically, what we have

done is we've looked at that information and so, for example, you know, several years ago in an infrastructure like DTCC, we saw the things that happened because we are attached to the internet. And so we had basic controls and good hygiene and became a relatively hard and frankly boring target to those adversaries because they couldn't turn it into cash in their wallet. That has evolved over the last five years. And so we see people knocking on the door with intentions other than stealing money. And what that's forced us to do is to proverbially, you know, knock on the door, try to push in the door, lift open the windows, to see what exposures we have in our environment before somebody does it for us, an aggressor.

And so the concept of testing that we look at is informed by the threats that we face, how do our countermeasures, our controls, our operation capabilities stack up against the techniques, tactics, procedures bad guys use, the specific malware of the month club that they subscribe to and those kind of things. And what

controls are the most important in our environment. There's some great research done with the Australian Signals Directorate and the NSA that looked at government intrusions and what controls, if they were in place, reduced the attacker's ability. They published a large report, but their top four controls said if you patch systems well, if you patch applications well, if you white list software, meaning only authorized programs are allowed to run, and you remove administrative access as much as possible, you can stop 85 percent of intrusions from succeeding. That's very easy for me to digest. I can do four things and make 85 percent of my problem much smaller. I'm going to make sure that I have testing assessment and measurement against those things which then also maps up against the threats we face and the threat data that we process.

So that's sort of a very long way to take us to the testing topic, but I think that puts some of the context. And the tests that we

perform, I think it's important to know there are really three types of things I roll into testing. There are assessments, which are periodic tests based on business condition changes, threat landscape, we're launching a new product, those kind of things. There is actually testing, which is episodic, so every quarter we do a test of X, Y, and Z to make sure it works. And then there's measurement, and that's really continuous. We measure our systems every day to understand, are they performing as we expected. We do this in the IT space; we also do this in the security space. And when we talk about testing broadly, it's actually important to recognize there are those three subcomponents. You have panels later talk about vulnerability and penetration; I would put those in the testing because they're episodic, you know, when we release a new application or every quarter, or whatever the frequency might be. But I think measurement and the assessments are equally important in that overall testing regime because they tell you where to focus and they give

you that sense of, are those top four controls working at the operating level we need to prevent that 85 percent of the intrusion problem.

MR. WASSERMAN: So we're going to turn back to some of the specific issues around testing in a few minutes. I do want to finish setting the table here though and I think I'm going to turn to Gerry Brady who is a Managing Director at Morgan Stanley and their Chief Information Security Officer. And, Gerry, if you could tell us a little bit about how a successful attack on critical financial market infrastructures could affect the U.S. financial system.

MR. BRADY: Sure. Thank you. And I think profoundly. I think the short answer is easy, profoundly impactful, but that's because of a couple of things here today. One, the nature of the threat actors, the diversity and danger of the threat actors, but probably even more so, the interconnected nature of financial services firms. That diagnostic of what exactly is going on and whether an incident is even occurring at the

moment, whether that is something that is a threat actor or naturally occurring. It could be very difficult to diagnose in the event that those attacks are destructive in nature or affect information in ways that may foul systems, but I think the broader difficulty here is that that interconnected nature not only makes diagnostics very difficult, but goes back to something Leo said before about the weakest link in the equation. Unfortunately the weakest link in the equation is always part of our ecosystem. That interconnected nature means we care a lot about our peer firms, we care about exchanges, we care about clearing houses, we care about technology providers and supply chain. That's a very difficult diagnostic to do in terms of test, but in terms of the information sharing, a lot works very well right now around the intelligence community to private sector, and private sector amongst itself in order to enrich that information and get accurate pictures of exactly what's going on and what threats we have to deal with. But

following onto that, the coordination of instant response is extremely difficult. If you imagine how difficult it is to deal with national disasters when it occurs across the street, that coordination is really difficult to do when it's uncertain what kind of instance it is or when the instance goes unknown for a period of time. That makes diagnostics difficult and recovery very difficult. This is where it gets to probably the most difficult part of the equation which is that it's likely that a number of these attacks will be successful, [and] they'll have profound impact on the financial services ecosystem. It's likely that bad guys will target the weakest link in the system which may be outside of our visibility, difficult to coordinate, maybe not on U.S. soil. Said coordination piece is probably what is most impactful. It means that, at times, despite our knowledge of intelligence or activity, somewhere in our ecosystem there may be some good actor who is not aware of that activity, and coordinating recovery is very difficult. It just yields a

very, very complicated situation of difficult discovery, difficult diagnostics, and difficult response.

That's where I think that the information flow today works fairly well as it gets better and better every day, but marshalling that to coordinate response, that's something that's in its early days now. It's probably the most impactful part of all of this. If you miss in terms of intelligence or discovery or knowledge of an incident, you still get a chance to make up for that on recovery, but right now I think recovery is probably -- managing incidents and recovery, that's probably the most difficult thing they got going and that's an awful lot of coordination, an awful lot of complexity.

MR. WASSERMAN: At this point I'd like to throw it open a bit because I'd like to spend just a few more minutes on -- I like the way you put that, what happens if the bear -- what an interesting avatar to be using these days -- what happens if the bear gets the financial sector and

in particular financial market infrastructure. If anyone else would like to jump in on that. Brian?

MR. PERETTI: So I think, you know, that the bear analogy is interesting also for another way. Of course the purpose of that joke was to outrun the other person. And in the financial sector you can't outrun the other financial firms out there. So if you become more secure and your counterparty is less secure, you haven't really reduced your risk because that risk is just going to be transferred through to the rest of the sector and cause additional problems. So this information you're sharing is really important and why we see a lot of firms being more interested in doing this because protecting the system as a whole is now much more important than just protecting my system by itself because the way in which risk can be transferred through. And of course this goes down through the supply chain. So if you're buying goods and services from somebody and then they are plugged into your network and they're not secure enough as they

should be and that risk is now inside your system. You may never have known that, assuming that the product that you are buying was secure enough against any kind of cyber issue. So we're hearing, you know, continually about the use of the cybersecurity framework to be able to not only judge your own firm, but to be able to talk to your supply chain and be able to question them in a way which hasn't been done before, to see how secure they actually are, and to see if they take cybersecurity to the same level as you do, and then using that in your buying decision, if you have potentially different parties to buy from, so that you could look at your whole risk profile and see, you know, is this an aspect where risk is going to be transferred to me because a vendor or somebody else didn't take the appropriate level of mitigation to that risk that was out there.

MR. WASSERMAN: So you're talking about this in terms of how folks might deal with their counterparties or their vendors, but as a regulator one of the -- in particular under our

statute -- one of the things we very much need to do is look at the costs and the benefits of the things that we're mandating. And we will eventually talk a little bit about costs, but in considering benefits, what I think I'm hearing you say is that there may be some issues that go beyond the specific folks who we're regulating in terms of what might happen if the bear gets them. And is that, sort of, correct?

MR. PERETTI: So, you know, the issue tied to the construction of any system is that I personally don't know of any financial institution in the country who builds the system all by themselves. You buy parts to put together, right. Your computers are made by whoever is making your computer, your softwares are being made by other companies, and you're putting this all together to make a system as a whole in which you -- what you call your company. And then all that has to work under your initial risk management program. So as you move forward on these issues, of course, you're looking at trying to figure out how to

minimize risk as much as possible. But that gets, I think, to your issue tied to penetration testing or other testing that Mark brought up. You can only mitigate what you know. And as we're looking at this more and more, the more information we have being shared, the more insight you have into what risk your firm is actually taking on. And once you understand what the real lay of the land inside of your system, you're then going to be able to better allocate your resources to mitigate that risk that's most important.

So as I mentioned within this framework, right, the first thing is to identify what's key out there, and to be able to identify what your key aspects are, is to look inside your system, figure out what is most important to you, and how you're protecting it. And that of course is looking at, as Mark was saying, the testing against your systems to see how secure they are, what connections are being plugged into them, and how the overall security of your firm is being graded.

MR. WASSERMAN: But you're talking in terms of looking essentially at my own firm and essentially assessing, you know, what I need to do. What I think I was hearing though before, in terms of this -- and now, in terms of the interconnectedness, is that if my firm is harmed, if my firm is affected, not only will there be impacts to me but also to my counterparts around. And if I'm an infrastructure, I'm thinking maybe that might be even more pronounced.

MR. PERETTI: Once again as I mentioned before, I don't know of any financial firm who's an island in and of itself and isn't connected into the rest of the sector as a whole. And the futures industry is an important part of the overall U.S. financial system. And so we at Treasury care about the entire financial system as a whole, and even all the individual parts that make up that whole system. And our goal is to be able to share as much information as possible to make the entire system resilient and try to figure out how to make sure problems out there do not

cause additional damage than what may happen at one firm. And, of course, the more we can increase resilience and reduce the overall risk is something that will be very beneficial for us.

MR. WASSERMAN: Please, Steven.

MR. CHABINSKY: Let me start with the proposition that we do when we go into testing, whether it's vulnerability testing, but more importantly, penetration testing: act under the presumption that the bear will get in, right. The first step is prevention. And I think, you know, Leo's point is well taken that for opportunistic crimes, the bad guys move on if you're secure. I mean, so for criminals, if they could just as easily commit a fraud with somebody else, they will. That's not the case with targeted attacks where a specific firm or an exchange or a company is absolutely being targeted. We see this routinely where the bad guy is there to stay. They will come back time and time again. If you notice that they are there, you know it's a long-term engagement. You will be in hand-to-

hand combat with them in perpetuity. And so those targeted attacks we have to view in that way.

Now one of the beauties about penetration testing is, when we conduct penetration testing, certainly we're trying, as Mark very eloquently described the scenarios, to make sure that as much that can be prevented is prevented and preventable. But then there's a second part.

MR. WASSERMAN: Let me interrupt just for a second because penetration testing, for the benefit of everyone here, if you could just give us a moment on what that is and then tell us about problems.

MR. CHABINSKY: It's a good point. Typically you look at your system in two different ways. One, what are the controls that you put into place, right. Have you patched your system, what processes do you have, what technology have you deployed, what physical restrictions have you placed with locks on doors. So there's a whole assessment of your capabilities, but penetration

testing is how would you then react if someone then tries to actually intrude into your system. So it's different from the setup, as someone had described. It's the difference between putting the alarm on the house and the video camera versus actually monitoring those and actually being able to detect when the bad guy comes in with those in place. So in penetration testing, the idea is really this notion of detection. And when you're first trying to prevent, of course, that's great, but then you have to move to what you're seeing, as the security industry has spent a lot of effort moving then toward rapid detection, containment, and mitigation. And that gets really to the answer of your question of what happens if the bear gets in. The hope is, you notice immediately, you contain it quickly, and mitigate it before too much harm. As an analog in the physical world, think about an air bag, right. I mean what happens if your car gets into a crash, right. You know you'll try to prevent that best you can, but if it does, you want something to go

off quickly and contain the damage so that the harm is reduced. And the same is true here, right. So we're trying to work with our customers including exchanges to make sure that that time to detection is not the industry standard of hundreds of days, but microseconds, and that then you could contain it so that any damage is limited, maybe in the best case, only to a reconnaissance phase where the bad guy was looking at your system, jiggling the handle, but didn't then get actually to do anything, whether to see anything, put anything on your system, and so that you're contained immediately. And what good penetration testing looks like someone in the private sector is going into a network, starting as stealthily as they can, mimicking the exact methodologies that known attackers use. And then if they could get through at that level of stealth, then they start becoming a little bit more noisy to see where in the chain your entity is able to pick up that detection, recognizing what's good, what's working, those best practices, and then seeing

where that gap was between the best hackers, your detection capabilities, and resolving them quickly, easily, inexpensively typically, and then moving on. So again the question of what happens if they get into the system, that's what it's all about and that's why we're doing the penetration testing.

MR. TAYLOR: Steve, let me ask you a quick follow up question to that. And I'd love to hear from other people on the panel who I think will have something to say about it. You mentioned that you have some exchanges as clients, and there is a sort of difference between the situation of exchanges and clearing houses and, you know, some other pieces of the world here. Trading systems, to a good extent, and clearing systems, even to a greater extent, aren't internet facing. And it's possible for some people to feel, well, with multiple fire walls and we don't face the internet, we're safer than the average bear, to test the analogy. Is that true? When you're mostly not internet facing are you

vulnerable despite that, and in light of that, what kind of testing do you need for the purposes you're talking about?

MR. CHABINSKY: Yeah, that's a great question. And the answer is you may be less vulnerable to the common criminal, but unfortunately you're quite vulnerable to targeted attacks. And I'll tell you why. There are two different reasons. And this is true across infrastructures. So not only in the financial industry, but if you look at other critical infrastructures. One thing is that there tends to be an interconnectivity now between what we would call the internet technologies, the IT world, the corporate enterprises, and the operations technologies, the OT world. And although there are firewalls and, you know, there are ways to isolate that, become a little bit more technical I think than we want to get into here, to control those different domains. What we are seeing in our experience as pen testers is the ability to get onto the enterprise system and then to move

within the system. And we could talk a little bit more, perhaps it's better for the next panel, of how we escalate our privileges. So starting out in a system that would not otherwise have access to -- so the normal user might not otherwise realize they could access other parts of your system, including platforms, to be able to then gain the passwords and credentials by being in that system and moving up your capabilities, before you know it, viewing trades and the like with the ability to view, alter, delete trading information. And so that's one way, the fact that there ultimately is interconnectivity.

The other, even in situations where there is not connectivity, we've seen in very well protected areas like the military, where the Department of Defense suffered malware infections on its SIPRNet, the secret internet protocol routing network, because of the use of thumb drives, where there was a thumb drive that was used -- not a thumb drive, this happened on multiple occasions -- on an internet facing

computer that had been infected, where the malware, which was installed to our knowledge by an intelligence service, was then actually programmed to look for removable media. And it would hop onto that removable media, and then when it was placed in another computer, it would hop off, kind of go around like a road trip, right. You know, your first stop, get out, take pictures, send them home. And then if they could not figure out how to get back out because that computer is not on the internet, go back onto removable media, recognizing that it would be able to get back out. Similar to being in prison and looking for that laundry truck, right. And so that's a great example of how the bad guys evolve to even recognize that there are these isolated systems, but there are still ways to get through that air gap. And that doesn't even begin to discuss the problems of supply chain where the hardware that's being used to create those isolated systems can already be infected. And being on the internet might suggest, at best, that confidential

information would not be able to get out readily. I mean there's this thumb drive issue we just discussed that would show it could be, but it completely discounts destructive attacks or integrity attacks which don't rely on further communication between the hackers and the victim. They can be preprogrammed with something that we in the industry would typically now, for well over a decade, call the logic bomb, meaning it's preprogrammed to do something, and that something might not be taking information, it might be leaving something behind, and in the worst case scenario something quite destructive of the nature that we've seen. We saw a couple of years ago a company in the energy field wake up to find 30,000 of their 40,000 computers had been wiped clean through one of these types of malware attacks. And, of course, when we work with our clients, that is to them every bit as important if not more important than the potential loss of information through confidentiality. It's the difference between a privacy, you know, a data privacy

problem and a data period problem.

MR. ORTLIEB: If I can just jump in here one second. So I think what you're talking about though is -- I'd like to jump back to something that Brian talked about earlier, and that is that we're still looking at this in a microcosm sense, right, of just within that one business. So even if let's say DTCC removed themselves from the internet, Morgan Stanley is still connected. And so that's kind of what we wanted to talk about there is, again that bubble in the rug like analogy where I think if you push it down in one place, and it pops up somewhere else. And even if DTCC solves their problem, it doesn't necessarily mean that Morgan Stanley can't necessarily cause one for them. And so that's the interconnectedness I think that we're looking at in asking about testing in that area.

I mean I take your point. I mean two things. One, regardless of who protects themselves there's also, you know, the viewing of confidence in the market, so that's one other

aspect that is not even based on the interconnectivity from a technology perspective, but if parts of the industry start suffering, it doesn't bode well for the markets in general. But with respect to interconnectivity, just because you're all on the internet together I think that, you know, there are vendor issues that we've seen in the past where certainly anyone in the financial industry does have constant connections and those are increasing with respect to the exchange of data, so those are entry points and egress points, you know, to the point you're making. From the perspective of penetration testing, obviously you're not reaching outside of your clients, you know, network to test outside, but what you are doing is looking at those connectivities and determining how you're monitoring what's coming in and out of your environment, regardless if it's with someone else within your ecosystem or any other website or customer facing site. So we've seen that customer facing websites can be an infection point that

would propagate through your networks. So it really in that way matters less about who you're connected to than the fact that you are connected and have to be monitoring those points.

MR. NELSON: Yeah, just to add to Steve's point, actually Gerry and I were talking about this earlier. I think when we get to the response phase, typically when we think of back up we think of like a 9/11 scenario and having a back up site, you know, 700 miles away and, you know, hot back up and all that, but what if the malware has been in there for a while. What if it's infecting lots of different systems including your back systems and you go to back up. Kind of like in the Sony situation where I don't know -- I still don't think they've produced financials for third quarter. You know, so those types of issues are of great concern. And I don't know if, Gerry, you want to continue to comment on that, or Mark.

MR. CLANCY: I was just going to add, so I think something Steve said earlier, before he painted the whole sky black, which I'm really good

at as well, there are a very, very, very, very -- add six more verys -- large number of attackers who can attack something that's directly attached to the internet. There are a much, much, much, much -- five muchs -- smaller number of people who can do what Steve mentioned. And so part of what I do as a market infrastructure operator and an operator of a private network is I look at what is my exposure to everyone on the planet who has an IP connected device, including a refrigerator. And then what I worry about is the people who are going to research, plan, plot, and come up with that level of sophistication because I become their most important target. And the premise that I have is that the bulk of my controls are to make it so I don't have to worry about the billion internet users, I have to worry about the several dozen groups that Steve and his team track closely as well as Leo and FBI, right. So part of it is the mitigation of the internet channels to reduce the amount of bad guys I have to worry about so I can focus on the ones who really are willing to

spend that time and energy and basically climb over the proverbial wall as opposed to those who are stopped by the wall and the controls being effective, right. And the distinction between those two is very important. I can do a pretty reasonable job. I won't claim perfection of stopping the billions of attacks to the internet. I am fairly certain that the well-funded, highly motivated people will unfortunately have some success. Which then takes us to the resiliency discussion that follows. So we just have to kind of split that out. It's a bit of a simplification, but I think it's a very important distinction to make. And so air gapping, as this is called, of having two separate networks that aren't connected to a very sophisticated attacker, is not a huge barrier, it just costs them more time and money. To the average adversary, if there is such a thing anymore, it's enough to stop them. And that distinction is very important.

MR. WASSERMAN: So I wanted to follow up on something else that Steve mentioned which is,

you know, you mentioned there's a couple of dozen bears let's call them --

SPEAKER: (off mic)

MR. WASSERMAN: There you go. Bears, pandas, whatever. And one of the -- Steve mentioned two words that really resonate with me which is market confidence. And what are the potential market confidence issues that we have to -- that you do worry about in the event that there is a successful intrusion?

MR. CLANCY: So in classic information security kind of terms, we use CIA: Confidentiality, integrity, availability, as the sort of moniker. And for market infrastructure, integrity is the most important thing. And slightly behind but close to it is availability. And unfortunately very down the ladder is confidentiality because if confidentiality is lost, markets can function. If integrity is not assured, if we don't know it, markets can't operate. And so from an operation of the markets, the integrity piece is the most important

objective. And when you look to tie it back to sort of where you're going in the rulemaking space, if you look at the body of best practice most of best practices are tailored for protection of confidentiality. And it's not to say that some of those controls don't help support integrity, but you actually need to look at different things and emphasize different things to ensure integrity. And the state of most of the intrusions that we talk about publicly are mostly those things where confidentiality has been lost historically, with the direction in these destructive attacks which are attempting either to take integrity or availability out of the scenario.

So I think that's the other piece that you need to focus on is, for market infrastructure, maintaining that integrity, that data is correct, we know who owns what, the prices are good, that's the most important thing. And if unfortunately, data is disclosed about activity in the markets, that is a survivable event from a

resiliency perspective, but if we don't know who owns what and what their positions are, then there are no markets.

MR. WASSERMAN: Gerry, I think you wanted to contribute on this.

MR. BRADY: Sure. Just to chime in on the interconnection, this notion. Most of this mythology is based on people's belief that either technology is acting exactly as they expect, or that people are acting exactly as they expect. When people say two networks are disconnected that's because they believe that some technology, whether it be a firewall, a switch, or some administrative technology actually works. When flaws happen, networks become interconnected. People rarely act -- even if they're honest and well-intentioned -- rarely act predictably and sometimes they'll join networks together out of convenience, to fix something, or in error. So this interconnectedness is difficult in that it's not just networks, it's people, sometimes it's information. Pricing feeds are a really good

example. But to Mark's last statements, some of that integrity issue comes down to confidence. Do you have confidence that I can figure out who I need to pay at the end of the day and that I get the right wire instructions to get that money there, is trading occurring with the right pricing, is the trade being attributed to the right individuals. All of that is confidence around whether or not you can manage integrity of systems, that referential data is correct, and the business really operates the way you expect it to be. If your counterparties don't believe that you're going to get payments at the end of the day for some object you bought, it's unlikely that transaction will occur. That turns into a crisis of confidence. I think that's how the overall ecosystem gets affected by integrity problems that lead to confidence problems.

But all of the conversations around this not being possible, or that things aren't connected, or that you're not sitting right next to that family member who is also running from the

bear, that stuff is all a very, very, very connected environment and you, your family, your lawyer, and your employer are all kind of running as a pack. So there's an awful lot of mythology around whether we all have the same threat, who's attacked, that it's more likely that your back door that's not well protected is someone else entirely. And that's where the interconnectedness, if you get past the mythology, that lack of connectedness is possible, that's when you get down to actually diagnosing the problem.

MR. DANIEL: So I would just echo that, you know, I have yet to find a situation where a network was truly actually disconnected. I can't tell you how many conversations I've seen where the head of some organization is saying yes, that network is completely disconnected, and then their CIO or their CISO whispers in their ear, oh, except for those two other lines we installed to do maintenance. You know, there are always the exceptions that get put in there.

When you talk about risk and systemic risk, there is the risk in the sector, but one of the things that we actually discovered, the President's Executive Order from a couple of years ago, 13636, required the Department of Homeland Security to do an analysis of the places in our ecosystem where we had a single point of failure that could cause catastrophic damage. And it turned out that that analysis was really hard to do. And it turns out that we have an estimate of some of that, but we don't actually really understand how all of our critical infrastructures actually interconnect. The system is both chaotic and complex in the mathematical senses of those terms. And so we really don't actually know what happens if you really had an impact on a significant part of the financial sector, and then how that might impact the power sector, which would then subsequently impact the transportation sector, and how this could ripple through. We really don't actually understand that right now.

And we're starting to see the emergence

of some other areas that we didn't even expect to see risk in, but are rapidly becoming areas of risk. And some of these are in the internet utilities. So the most noteworthy of these from this past spring was the Heartbleed vulnerability in the secure socket layer utility that is used by like everyone for everything in some form or another. And it turns out that this particular piece of software was essentially developed open source by like a fairly under-resourced organization, and it has this massive vulnerability that had been sitting there for an extremely long period of time that some researchers finally discovered. And I anticipate that we will see more and more of those kinds of vulnerabilities emerge in the utilities that underpin what's going on in the internet. And that is a risk that is very difficult to identify. You have to have the -- this again gets to the resilience question. You know, it's very hard to identify those ahead of time. You have to have the ability to rapidly respond when one of those

emerge and actually be able to address and patch your systems and quickly get ahead of it. In the case of Heartbleed, from the time the researchers identified that and published it to the time that it was actually incorporated into malware sets and malware tools that we were watching was approximately 18 hours.

MR. WASSERMAN: Leo.

MR. TADDEO: Part of my responsibility in the special operations division is to conduct offensive operations as well. So we look at some of the best in the world at protecting their networks. So Gerry mentioned the human factor. Now when you think about penetration testing, that's probably the one vulnerability that is overlooked when we are discussing this topic. Many of us talk about configurations of networks and how they're connected and whether or not they are hardened to a certain degree, but the real professionals, my folks who are trying to get into these systems, are looking at the people who operate them. So as you develop your penetration

testing protocols, don't leave that out. Don't leave out testing the people who have their hands on the systems and who have closest access to them.

MR. CLANCY: So just to pick up on what Leo said, so one of the tests that's common to do for financial institutions is test your employees to see if they click on phishing messages. It's a very common technique that attackers use. And one of the firms that does this produces benchmarks for their clients about it. In the average financial company, 40 percent of the staff will click on the malicious link. Companies who are good get down to 20 percent. And the best companies get to single digit percentages. But if you send enough of those messages, a single digit percentage is going to yield fruit. And so the challenge is that very important piece, both in terms of social behaviors, you know, recruiting, etcetera, but we also have to recognize there is a sort of asymptotic limit as to where you can get -- at the best performing organization, it's still

not going to be zero. And so that human element is always part of the design. So as you define your control infrastructure and your testing regimens, you need to test and probe the understanding and measurement of that along with the, okay, when it fails, what happens. Because it's going to fail. I mean people will eventually click on the link either because the attackers are so good at making it so compelling, or the person is just not paying attention and they just want to see what this new notice from my payroll company is all about.

MR. WASSERMAN: So as we're coming -- we've got about 10 minutes left. What I'd like to do -- this is the panel that's supposed to sort of set the table for the rest of the day. And we're going to be discussing in the following two panels some specific types of testing, penetration testing, which has been discussed so far, vulnerability testing, key controls testing. And actually I sort of want to raise a sort of more high level question which is: are these the right

things that we as regulators should be looking at, are there other things we should be looking at in terms of testing as we are looking at what kind of rules that we might, you know, we might be imposing? So if anyone could jump in on that.

MR. BRADY: You know, I think we touched a couple of times on, you know, will you actually win these battles, and sometimes recovery is really where winners emerge, or at least losers emerge. Recovery testing is very important and I think in today's world, the attacks we see that are the most frightening to us are certainly the ones that are either destructive in nature or information contaminating in nature. Very few people do disaster recovery in business continuity testing in good ways that really address an adversary causing that outage and recovery testing in that vein. That's very, very important. The subtle sort of issues of either contamination on data recovery, testing those practices are very, very important as well. Typically part of business continuity, but these days that crosses

over very nicely into cyber threats as well.

MR. WASSERMAN: So let me try and draw you out on that. So when you're talking about data recovery testing, are you talking about essentially okay, so what happens if there is a loss of integrity?

MR. BRADY: Knowing that there is a loss of integrity, knowing that there's lost integrity in feeds that you receive from outside parties. So being aware when there is a disturbance in the force, whether it be pricing or other information coming from the outside world, and knowing what that impacts in your shop and having recovery plans to recover from that. As an example, information feed that's no longer wholesome and trustworthy or knowing when low and slow attacks occur with data that may be very difficult to reconstruct.

MR. WASSERMAN: I'm sorry, low and slow attacks?

MR. BRADY: Meaning that data changes subtly over a long period of time, not as simple

as data being deleted, but historical data being tampered with in ways that are non-intuitive and maybe not things you directly test. Being able to recover the very short-term data, the last hour in getting to a safe state, or being able to look at long-term data that may fuel anything from a risk model to other operational practices, and knowing that you can get back to some known state. Detecting subtle changes in data is very, very difficult. And not only do people think about that in terms of external influence, you know, a trading model stops working or other sort of external events, it's very, very difficult to get down to a level where you can detect subtle changes in information and have something useful to do about it in terms of rolling back the clock and knowing what that affected.

MR. CHABINSKY: To pull on the thread that Gerry is discussing, the penetration testing and the vulnerability assessments you were talking about are one part, but then we like to talk about something else that's called tabletop exercises.

And I think that's really where Gerry is getting to, right. How do you actually put this knowledge into practice, what happens when you really do detect something, when do they contact the regulator, right? Who is involved? Do they know to contact their general counsel, you know, is crisis management at the table, what are the forensic firms, outside counsel? How do you actually play this out upon the detection? And there's a difference where you used to discuss -- unfortunately in the government we would put out a lot of documents all the time that people refer to as shelfware, meaning they never got used but you could check the box that you created it and it's there. I unfortunately realize that's not just a problem for the government, right. So a lot of people create shelfware where they might be able to check the box saying they have an incident response plan, but they really haven't tested it. And the idea of understanding what that looks like from things that are very subtle to, do you have phone numbers at home to call people as opposed to

having to log into a system that might not be operable. To actually use your computers to discuss with the FBI, I have an intrusion, when the bad guy might be on that computer looking at that. And so the activities that surround detection and containment and response are equally significant. And I would add that as another topic of discussion for you.

MR. WASSERMAN: So just -- I'm going to follow up because -- so the people you see in front of you are the folks whose responsibility is basically to draft and propose to the Commission rules that our infrastructures would have to follow, and as well folks who would be going out and looking and seeing -- essentially looking at the infrastructures to see whether they're following these rules.

So in order to avoid, you know, what the Chairman was talking about in terms of just basically getting employment for IT and to -- or the way you put it very nicely, you know, shelfware -- what are the things that we should be

doing? How should we be going about writing those rules? How should we be going about examining those infrastructures to best do our job?

MR. CHABINSKY: Well, I mean, you know, you start with a dialogue with industry and I'm not telling you anything you don't know, but the question is, how is that already occurring and, you know, to find the best practices that are out there. The regulators are all, you know, positioned already on site. I would go about that instead of first thinking about rule making, of understanding what's already working, and what that looks like. And I think that you'll find, as we have found out in the field, that a lot is really working very well. In fact what's not working is the exception, and that's why it's so good to have people come in and just narrow those small gaps. And I think you'll find that as well. And I think that there might be instruction and guidance that you can put out that would show what "good" looks like. And I think that that will come easier than you believe.

MR. CLANCY: Just one maybe different thought on that topic, right. So as a security person, what I want to know is: how does my production environment perform against stresses, things injected by attackers, etcetera. And as an operator or an infrastructure, the last thing I want to do is stress that infrastructure and cause it to break in an unexpected way. And so the real challenge in all of the testing discussions of any kind is, how do you balance that tension, because really the only thing that we care about is the production infrastructure that we operate facing the markets, and that's the first thing we want to test and the last thing we really want to test because we don't want to cause a failure through our testing, but at the same time we want to prevent a failure by someone else trying to induce it. And so the challenge has always been, how do you look at -- and this is really to abuse the analogy -- individual links of the chain without actually testing the whole chain at once. And that's I think what you'll hear for the rest of

the day, that's sort of the hard problem that no one has unlocked yet is, how do you test the whole chain so that if there is one link and you pull it apart you don't actually break the chain. And I think that's the piece as you go to the rule making side in doing these component pieces, that's the hard problem that quite frankly no one has come up with a good answer to. There have been discussions about, you know, market wide exercises and other kinds of things. That's an attempt to figure out, how do you test the whole chain, but logistically that is incredibly complicated. And quite frankly every non working hour is already consumed with testing all the individual links. There's not a lot of time on the clock to test the chain all together. And I think you'll see that as you go through the panel today. That's sort of the underlying theme of why this is so difficult and hard to come up with very precise rules around you should do X, Y, and Z.

MR. WASSERMAN: One question out of curiosity: when we're talking about testing, are

there advantages and disadvantages as to how much of this are things that can be done by the entities themselves, how much of this are things that basically need to be done by independent contractors?

MR. BRADY: There are advantages to both, but neither can stand alone. I think there's an awful lot of testing that needs to happen internally first hand. I think there's an awful lot of testing that needs to happen from the view of an outsider, and in particular maybe against particular threat models that you might find independent parties a little more adept at practicing. At times, you want to look at testing as something that comes from the eye of a particular threat actor to understand how you sort of size yourself against them. But to the relationship between this and rule making, the size of the institution matters a lot in terms of how impactful they are and what kind of testing makes sense. Scaling it down to smaller institutions for the utilities makes that

especially difficult. I wonder if rule making here as opposed to guidelines, sort of which is going to make more sense, but both kinds of testing are required and they get you to different places. One gets you something you can do more chronically to know the state of systems that you already know very well otherwise, and the other gets -- the independent testing gets you the ability to roll all of it up to, how does all of that perform against a particular threat actor. So you need both; it gets you to different places though.

MR. NELSON: Yeah. I would just add that I think you need a risk assessment first and really assess your risks and figure out what you want to test. And then I would -- I know in our case, FS-ISAC, we're going through that whole process or we're using an outside party to do it. We're doing our own risk assessment first. We'll be running it by our board next week, our risk assessment, and then go back to see what we need to test.

MR. BRADY: Both risk assessment and also scenario analysis. So what are those risks and what are those scenarios where they get exercised; both get better a lot if you're going to target the balance of the work you do.

MR. WASSERMAN: Leo?

MR. TADDEO: So for the Commission I'll try to give a perspective of a government official who has tried to operate in this area, with this industry, with this problem. The financial industry does not act like a typical victim in this particular threat area. They are very well financed and they don't often complain like a normal victim would. So if they had been the victim of a bank robber walking in with a gun, we would get the call right away. So you have to approach the problem, I think, a little differently. And we've somewhat learned the hard way that if you're not adding value you're going to be in a position of having to compel cooperation. And we of course have subpoena power and we have other ways to make it highly

uncomfortable for a financial institution to not cooperate, but I think the best results we've gotten is when we've gone in and helped solve a problem in a way that was a net benefit.

So that's only my two cents as someone trying to operate in this space with this level of sophistication, this level of financial capability, this level of legal capability, and going in and not being able to just say this is what we want and this is the date you're going to give it to us.

MR. WASSERMAN: I think we have run out of time. I would like to again extend my thanks to all of the panelists. I think this has been a very, very helpful conversation, certainly helpful to us, and hopefully helpful to everyone out there. So thank you again. We're going to break for 15 minutes and come back at 11 o'clock.

(Recess)

MR. TAYLOR: This is our second panel of the day addressing a topic that was very well set up by panel one. We're now going to turn to two

of the most important types of testing that critical infrastructures might do, namely vulnerability testing and penetration testing.

A couple of administrative things first. There is an opportunity for members of the audience to ask questions. There was a three by five card on your seat, and there are also more of them on the table over here to my left. If you would like to send in a question, write it down. There will be a gentleman passing through the aisle periodically who can bring them up here. We're happy to have you enter into the dialogue this way.

I would remind the panelists, if I may, when you are going to speak, please turn on your microphone by pressing the button in front of you, and if you would, when you are done speaking, turn it off. There is the possibility for feedback and so on, if too many of us have the mic on at the same time. And the goal here is a dialogue between the panel members. The more you respond to each other, the more useful this is going to be

for us.

Well, let me start by turning to Kevin Greenfield, who is director for bank information technology at the Office of the Comptroller of the Currency, and is here also because OCC is an important part of the FFIEC. To start with, a question that I'm going to pose to all of the panelists. Since we're focusing on vulnerability testing and penetration testing, what do those two types of testing mean to your organization, and what do you think the costs and benefits are associated with this?

And Kevin, if you could say a word about what FFIEC is and the role it plays in the best practices arena, that would be great.

MR. GREENFIELD: Sure. The FFIEC is the Federal Financial Institutions Examination Council, and it's composed of member agencies representing the Office of the Comptroller of the Currency, the Federal Reserve Board, the FDIC, the NCUA, as well as now the CFPB, and then, representatives from the liaison committee. So,

representation from all the banking regulatory agencies in the United States.

And essentially, we're charged with supervision of the financial institutions that we individually charter, whether those have a national charter, whether they are a state bank that are a member of the Federal Reserve, or a state chartered bank, non-member or a credit union. In addition, on an interagency basis, we do supervise the critical technology service providers to the financial industry.

MR. TAYLOR: So Kevin, what do vulnerability testing and penetration testing mean for FFIEC and OCC? And what do you think the costs and benefits are for financial sector infrastructures?

MR. GREENFIELD: Sure. And with this, I always say depending on if you ask two technology professionals to define these, you'll get three different definitions. So, I've always used a good analogy when speaking with bank management or some of the executives from the regulatory

agencies of describing vulnerability assessments as looking at the security of your home. And with a vulnerability assessment, you'll be scanning and reviewing to ensure software updates are in place, patches are in place on a timely basis; that network components are configured properly; there are no known vulnerabilities present in application software.

So often, I say it's -- using the home security analogy is you're checking to make sure all the doors are locked, the windows are locked; that the doors are thick enough; that the security system is on and the batteries are charged. Doing that scanning and assuring yourself that all of the controls you've set are in place and operating properly.

When looking at penetration testing, the scope of the penetration test is very different, and that's where, as opposed to looking to make sure all of the security components are in place, I'm paying someone to try to break into my house, to try to break through that security, so I can

test and get a level of assurance that the security parameters that I've determined and I've set in place are actually adequate.

From the FFIEC's point of view, for security of a banking network, the use of both -- the need for both vulnerability scanning and penetration testing is essential. Financial institutions need to constantly be scanning their environments for known vulnerabilities and correcting those, for ensuring that they know everything that's present in their network and it's configured up to the latest security standards set by the institution.

And as for penetration testing, you will never know how strong your security is until you try to break it yourself and try to bypass. And I've often used the phrase that if you're not testing to see how strong it is, I guarantee you, somebody else is.

MR. TAYLOR: Let me turn to Jerry Perullo, who is chief information security officer at ICE, and I'll ask the same question, Jerry, and

maybe take vulnerability testing first.

MR. PERULLO: Sure. So, I mean, I'll definitely echo what Kevin said as far as the definitions. They were bang on. More specifically, I'd say vulnerability testing can often be passive, while penetration testing is active. Vulnerability testing -- you know, he used the analogy of making sure that your windows were locked. You don't necessarily have to try to smash one to do that. So, the exploitation piece is a big differentiator, and that's where penetration testing comes in.

In scoping, there's a big difference, as well. With vulnerability assessment, you can certainly scope fairly effectively, so you could take -- for example, I'm representing several subsidiaries that are under regulation here today. I could easily scope a vulnerability assessment to one of those. I could scope it to one network or to one network or to one data center.

Penetration tests will really suffer if you try to limit it like that with a scope,

because the bad guys aren't worried about the scope. They'll get in any way that they can. So in penetration testing, not only is it not beneficial to try to limit the scope, but it's sometimes near impossible, because if you give a tester a general target, they're going to find any way they can to get in, and that's generally a positive thing.

To touch on the costing just a touch, because I know you asked about that, David, a few times, it's very hard to pull numbers out, and maybe even not a useful exercise in some cases. But I think one thing that's important to denote is that vulnerability scanning generally lends itself more to automation.

So, you can put some infrastructure in place and begin doing vulnerability scanning, and you can have a lot of automated systems that will learn about the latest configurations and see that they're in place. And penetration testing, on the other hand, is usually more manual. There's a human involved. They have to pretend to have

malice and to try to model what an adversary would do. So, those are usually more atomic engagements where someone will come in, do something at a point in time and wrap it up.

MR. TAYLOR: Let me turn to the other side of the table. Steve Chabinsky, you're general counsel and chief risk officer at Crowdstrike, and I know you do some of the testing for clients. What do these types of testing mean in your world?

MR. CHABINSKY: I actually think Kevin and Jerry did a great job of defining them. Right? And this notion of passive versus active is a nice way of looking at it, as well. Right? Making sure that your system -- really, these analogies, I think, to houses are right on target. Right? Did you close the windows, lock the windows, lock the doors? You know, what does your perimeter look like?

That's very different than saying, okay, now we've got everything in place. We're ready. Right? This is the best we think we could do.

What would someone who actively is trying to bypass your, you know, security protocols be doing? And how could we do a better job of deflecting that?

So, I think you could have more analogies as well. Right? You know, is everyone properly positioned? You know, did you clean your weapons? Are they on the front? That's a lot different than saying okay, come at me. You know? And you could do it, you know, in any number of ways. But I think it really is this view of passive versus active that defines it.

MR. MCGONAGLE: David, could I just jump in? And sir, I have a question just about the type of penetration testing that you would be interested in doing. And you make a determination or you know, a client makes a determination as to whether the type of penetration should be external to the organization or specifically targeted.

I know you were saying the differential between you don't want to limit the scope. But aren't there areas where you would think about, is

there a risk protocol within the organization that I want to do penetration testing and not just test to see whether there are vulnerabilities, but to see how far into the system I can go.

So, can you just talk a little practically about how those decisions are made?

MR. CHABINSKY: Yeah, I think that's absolutely right. We're trying to make sure the client -- and the client is trying to make sure that they've looked at their risk, and they're trying to determine what the greatest harm is to their environment, and that that's what they're testing for. Right?

And in the cases that we've been discussing today, it's that production environment. Right? The operations piece. And it's either a look at it independently -- right? Just if you were already in that operations piece, whether it's from an insider perspective, or you know, just the ability to detect removable media in that environment, that would be one way of scoping it.

Or, you could expand the scope to say if you did not have access already to that environment, how is it connected to the enterprise environment? If you had a foothold in an enterprise computer at a normal user level, every regular employee, would you be able to escalate to get into our production environment? Those are typically the engagements that we deploy on.

MR. MCGONAGLE: And then, so how do you make the decision, when you're talking to a client, of what a recommendation is to the most effective types of penetration testing that they should consider? Is it, you have to coordinate first with a risk mitigation analysis that you know, qualifies or quantifies where you think the greatest degree of vulnerabilities are? Do you just let loose, and you know, go where it takes you?

MR. CHABINSKY: Yeah, (Laughs) there's definitely no letting loose in this environment. And it gets back to an earlier point.

We want to make sure that we're not

doing damage to the environment. The point is to make the environment more secure.

And we go about that in two ways. One is, dialogue with the client. In this case, the clients are very sophisticated and have done a lot of work. They have a sense of where they believe they have more weaknesses than other areas, and where their expertise is limited.

We were talking before on the earlier panel on what's the advantages of in-house teams versus outside vendors. Well, one of the advantages of the inside team is it's constant. It's perpetual. They know the systems. They know their users. They know their risks, and they get to study that continuously.

The advantage of the third party vendor is we're specialists. It's like the GP. Right? It's like having your doctor that you go to more routinely, but then you're going to want someone who is a specialist at understanding what not only the latest attack vectors are, but being able to compare it to all of their client base.

The in-house person understands their network environments. The outside vendor sees how this impacts a multitude of different clients, thousands across other industries that have similar types of architectures that could be similarly vulnerable. And we look at our intelligence database. What is hot at the moment? Who has an interest in the financial sector? What tools are they using against the financial sector? And we will actually replicate the activities that are taken that are not academic that we're actually seeing in other customer environments.

But to get to your point, it's definitely not a let loose. It's at that point, it's how are we going to make sure that we come into an environment where we understand where the operational components are. It's very much a scientific process. Everything we do is audited. It's logged. It's repeatable.

So at the end of the day, the way one would view a penetration testing report, if you've never had the opportunity to review one, is

probably reviewing what your worst nightmare would be if you read it in the paper, and someone had done it. It's kind of this eye opening moment, like oh my goodness. Yeah, you told us all the great things we were doing, but that at the end of the day, didn't prevent this. But then, it has a really happy ending, which is, this didn't really happen to you. You were smart enough to come in and look at it.

And here are the steps. We actually rank them in terms of low, medium and high risk and low, medium and high cost, so that the operators could then decide how they want to tackle you know, some of the environmental challenges that were noted. So, it's very much a coordinated activity.

And also, even though we talk about the idea of how do you break in, in the physical world, the analog doesn't really work, because to break through a glass window, you really break the glass window. We don't break glass windows in our environment. What we would do, as a way of an

example, is we plant flags. So, when we get somewhere, we don't take data, we actually create data. We'll put it in a file, and then, we'll alter the data we actually created and put it in another file, and then we'll retrieve the file we just created and altered to show that all of this could have occurred with something that was resident.

So, it has to be an environment that everybody would be comfortable with; really, is not going to break anything, but really is looking at what is the most risky environments -- what's the worst possibilities for your operating environment? How are you going to see if that exists? And how are you going to make recommendations so that reading this story turns into really, the best thing you ever did?

MR. TAYLOR: So, I take it, Steve, in a way, you're telling us it's important to have penetration testing both by the infrastructure itself and by independent outsiders?

MR. CHABINSKY: Well, you know, some of

this is a resource issue. Right? Again, you are -- the financial services industry has the good fortune of having an industry that has always been concerned about security. It's part and parcel of what the industry does. That's not the case with all of the sectors we operate for, many of which do not have budgets and have not traditionally had to focus on the security challenges that now are involved by being connected to an Internet that allows the world access.

So, I think you'll find that for the majority of the industry, they do have teams that are continuously monitoring the situation, whether they have an ability to do the penetration testing that we're talking about, as opposed to the vulnerability scanning, which is standard across the industry, differs between the clients.

MR. TAYLOR: That provides a very interesting segue, I think. I want to ask Dave Evans from the Bank of England, who's senior manager for sector and supervisory cyber support there, how the Bank of England approaches this

question of how do you set the scope for penetration testing? Do you break the windows, and so on? And could you explain a little bit about the CBEST program that the bank is doing?

MR. EVANS: Certainly. So, the Bank of England, a couple of years ago, started taking an active interest in the types of threats that we discussed in the first panel. So, moving away from cyber crime, e-crime, e-fraud, those sort of long established patterns and threat vectors that banks and financial institutions have to look at, they became concerned more about the destructive and disruptive types of attack.

And so, in the summer of 2013, our financial policy committee, which is similar to the U.S. FSOC over here, made a recommendation that we were to test and improve the resilience to those types of cyber attack to the core of the UK financial system.

But before we could sort of effectively test, or as we are sort of testing and improving, we also needed to be conscious of the fact that at

the time of the recommendation, the FPC were concerned, but they didn't know how concerned they should be. You know? There's a problem about how big was the problem.

So, we wanted to come up with a repeatable testing framework that incorporated all the sort of better practices that we've heard mentioned this morning, in terms of a penetration test, but we wanted to also include threat intelligence as a key component of that part. So, the actual driver behind the test is intelligence. So, that was from both a commercial and a UK government angle, as well.

So, we wanted to have those two components right at the heart of our testing. I mean, I should stress at this point that the testing framework, CBEST, that we've built, is not a panacea. It's not a fix-all. You can't expect to do one of these tests and you will suddenly become cyber secure or cyber resilient. It's a component.

The other thing that we've done with

CBEST is, we've built it, truly with openness and transparency at its heart, between the regulator and the regulated. There's a problem we have in the UK, in terms of supervisors do a fantastic job of regulating financial type issues. When you move into the operational space, it all becomes a different language, and it becomes a very different type of topic that needs to be supervised. So, we needed to educate our supervisors along the way.

What better way to educate a line supervisor for an individual firm than actually have them as part of the whole of the process? So, right from the outset, the regulator is there with the regulated entity, and they're both determining the scope of a penetration test. So, they're looking at, not the technology. That comes next.

But what is it that this organization does that if disruptive, would affect UK financial stability? And then, you start to understand the functions and processes that you want to focus on.

And then, you can ask the more meaningful questions. Well, what's the technology that allows those processes to happen? And where are the people and who are the people that are connected to all of those processes? So, you start building a scope.

But the regulator will have a view of what's critical that that organization does. The organization will have a view of what's critical. And perhaps, the Bank of England independently, is sort of looking at a financial stability angle, and the system as a whole might also have a slightly different perspective.

So, the scope of the penetration test in CBEST terms is that amalgamation of the three different viewpoints, so that we can have some confidence that as we start on the test, and it's doing exactly what Steven's said in terms of mimicking tactics, techniques and procedures of threat actors, is it's targeted on the right systems, and sort of for the threats we're talking about. It's talking about the right people, the

right processes. And it is informed through threat intelligence.

We take our regulators along every step of the way, not so that -- you know, with a view for any form of sanctions. It's so that when they get the report that Steve spoke about, and it talks about you know, low, medium, high, or maybe a signal in a red, amber, green, and you've got your costs associated, the regulator can take that information, and not take it in isolation and say, right, here's a bunch of red risks I need fixing.

They can put their other supervisory hat back on and go, right, where does this fit within the other risks I'm asking this organization to manage? And that allows them to provide some sort of proportional supervision of the firm. And we think that's really important, that we don't suddenly bolt on a whole new regime of cyber security supervision that ignores everything that's gone before it, because it is -- you know, it's a big worry. It's a big risk.

But can we sit here today and say, you

know, we'll tackle it to the detriment of other risks? I don't think we can. So, we need to try and put cyber security to the supervisors in a language and a format that they understand. And CBEST, we think, goes a long way to doing that.

MR. TAYLOR: Can you talk to us a little bit about why the financial policy committee thought all of this testing was so important, and what relation they thought it had to financial stability?

MR. EVANS: So, it really comes down the potential impact. You know? Again, back to the first panel, when they were talking about how the threat landscape has evolved. You know, destructive and disruptive types of attack on the UK critical national infrastructure, and because of the interconnectedness, and for all of the reasons that have been discussed in the previous session, you can start to see that there is now the potential for threat actors, for whatever reason, political, ideological, to just cause harm. You know?

And we do have critical national infrastructure in the UK financial system as, you know, every financial system has critical national infrastructure, to a greater or lesser degree. So, we need to protect those. You know, we need to understand the threats. We need to build adequate and appropriate protection, so that we can minimize that disruption.

MR. TAYLOR: Let me turn to Ann Barron DiCamillo, who is the director of US-CERT at the Department of Homeland Security.

MS. BARRON-DICAMILLO: Mm-hmm.

MR. TAYLOR: And Ann, if you can start by explaining those terms a little bit, that wouldn't be bad.

MS. BARRON-DICAMILLO: Okay.

MR. TAYLOR: But the question I wanted to pose for you is, in today's cyber security threat environment --

MS. BARRON-DICAMILLO: Mm-hmm.

MR. TAYLOR: -- what types of penetration and vulnerability testing are you

seeing that critical infrastructure should be doing?

MS. BARRON-DICAMILLO: Okay. So, there are a lot of acronyms at DHS, and I'll try to go through a couple of them, just so you'll understand what I'm talking about. So, US-CERT is part of the NCCIC, which is the National Cybersecurity Communications Integration Center. And that's why we use the term NCCIC. It's a little easier to roll off the tongue.

And our focus at US-CERT and NCCIC is on state, local, tribal, territorial, federal and in the 16 critical infrastructures. So, one of them happens to be financial services. We have a lot of interaction with the financial services through the FS-ISAC in working on event based activity, intrusions, other kinds of incidents, as well as providing (Inaudible) compromise from other activities that we're seeing across the critical infrastructure.

As you're all aware, there's a lot of activity currently around the healthcare industry,

as associated with the breaches that we're seeing. So, within the NCCIC, we're kind of that entity that kind of opens up the aperture and shares indicators that are happening within one sector across the others, so that they can ensure that they're protected, when and if that activity trickles to what they're dealing with.

So, from the incident response perspective, I think when it comes to vulnerability testing and penetration testing, you want to look to see what is actually hitting my sector. From activities that we've been engaged in, what aspects of my network are my vulnerability areas that are not currently being, I guess, robust enough to be able to thwart the kind of activities that we're seeing?

I think it was stated by Dave, that we're seeing more sophisticated actors going not so much for criminal activity, but focused on kind of more what we call nation state events, persistent threat types of activities, where they are not focused on stealing credit card

information from financial institutions. They're interested in disruptive or even destructive type activity.

So, you need to look and see, you know, when you look at the whole kill chain of an event, where am I most vulnerable within my own infrastructure. And then, focus both your vulnerability testing, even cyber hygiene kind of aspect. One of the things we do within NCCIC is we provide cyber hygiene evaluations for critical infrastructure partners, as well as federal entities.

And then from that, you can kind of get a picture and a landscape of the architecture, and better understand where do I then need to focus on vulnerability testing, and then, where do I actually want to focus on penetration testing to make sure that what I found in these other assessments is actually accurate; that it's not just a paper exercise; that you can actually evaluate that in real time and in a production environment, and making sure that you're not

disrupting those networks.

And so, from intrusions that we've been engaged in over the last 18 months, it seems to be, from our perspective, one of the highlighting areas that we always focus on when it comes back to the vulnerability testing and pen testing, network segmentation -- the lack of that.

A lot of the common controls that we see that are being exploited, patching of operating systems, patching of applications -- these are all things that you can evaluate in some of these assessments to see where your infrastructure and where your architecture is associated with that. You know, it's not a silver bullet. There's no silver bullet out there. And it's not the -- we're not making it hard enough on the adversary.

We're letting them get in with patches that have been available since 2012. You know, there's a paper that's about to come out, the top seven CVEs that we see being leveraged by adversaries. Some of them go back to 2009. And so, we want to make it harder.

And as you kind of, you know, look at cyber hygiene, best practices, and then, get into areas of your network that you can be most vulnerable in, as they get the foothold, as Steve said, and then escalating privileges. You want to be sure that they don't have the ability. You want to be able to contain it, and be sure that -- you know, you think you have these containers, but there's no way for them to leverage from one to another.

So, these are all kinds of best practices. But again, we see too many intrusions happening because they're not -- they're implemented, but then they're not monitored. They're not updated. And so, as vulnerability testing and as penetration testing can help you identify those gaps in your network based on what's actually happening, not only in this sector but other sectors, because things do start to trend, a lot of times we see the adversaries leveraging lower level targets as an entry to test out new techniques. And then, that translates to

the primary target.

I think another big trend we saw from this last year is third party partners, and the vulnerability that is exposure of the third party partner to the primary target, and ensuring that you treat your third party partners with the same types of security controls that you do your own employees. We saw a number of cases last year where that wasn't the case.

And so, the whole aspect of vulnerability testing, penetration testing in those environments with that kind of constraint associated with it, I think is what we're trying to help focus on. But you can't say enough about how important it is just to follow the best practices in cyber hygiene.

(Simultaneous discussion)

MR. WASSERMAN: So let me just -- quickly --

MS. BARRON-DICAMILLO: Sure.

MR. WASSERMAN: Two things.

MS. BARRON-DICAMILLO: Sure.

MR. WASSERMAN: You used the term cyber hygiene.

MS. BARRON-DICAMILLO: Mm-hmm.

MR. WASSERMAN: If you could tell us what that means.

MS. BARRON-DICAMILLO: Yeah.

MR. WASSERMAN: And also, you mentioned top seven CVEs.

MS. BARRON-DICAMILLO: Mm-hmm.

MR. WASSERMAN: And if you could --

MS. BARRON-DICAMILLO: We're about to put out a paper about that, and this is --

(Simultaneous discussion)

MR. WASSERMAN: Well, what does it mean?

MS. BARRON-DICAMILLO: I'm sorry? Oh, CVEs are Common Vulnerability Exploits. And so they're just -- CVEs, they're -- the Microsoft patch Tuesday -- they put out a patch, and it gets a CVE number. These are not -- these are commonly available vulnerabilities with the patch.

And so, what we've seen in the last 18 months is a trend associated with certain

intrusion activities. They all relate back to these top seven CVEs. I think 60 percent of them are -- these CVEs would have stopped that attack if they had been applied. And so, we're putting out a paper associated with our findings from that. And this is something that we're working with our partners internationally.

This is a UK-Canada-Australia and New Zealand, as well as the U.S., is all putting out a paper associated with this, because it's -- what's trending here is also trending in those markets, as well. We're all multi-national organizations, and so we have to share this information, shared responsibility.

And so, we put that out -- or we'll be putting that out at the end of this month so we can get back to the cyber hygiene. We don't want to let these adversaries get in because we didn't patch our system.

SPEAKER: And cyber hygiene is?

MS. BARRON-DICAMILLO: Oh, cyber hygiene is the common controls that we should all be

applying. So, think about patching of operating systems, patching of applications, reducing administrative privileges across your environment. We see way too many users that have God (sic) privileges. Why is that happening?

And then, network segmentation. Ensuring that you have segmentations between your networks, and that those enclaves are contained, so that if an adversary does get a hole in your DMZ (sic), they can't then use that to get into your secret sauce, or the keys to your kingdom.

And then, the other one we also preach as much as we can, because we see this being leveraged a lot, is white listing. So, application white listing. A lot of times, executables are running in an environment that should never be running as part of the malware drops.

So, if we're leveraging application white listing within those environments, and it's difficult to implement, it helps reduce the kinds of incidents that we respond to on a regular

basis.

MR. TAYLOR: You know, I think I heard from the last several speakers a need for threat intelligence that may be, at the highest levels, only available from you know, governmental sources; a need for penetration testing expertise, you know, the medical specialists that you call in that might be best found in a third party service provider; and a need for the kind of inside knowledge and expertise that really only the infrastructure itself will have. How do you put all three of those together?

MR. PERULLO: Do you mind if I comment on that a bit? And I also wanted to speak a bit about the CBEST program, because I think that there's -- a lot of answers to that question are in there.

So, we're very familiar with the CBEST program. I also represent some bank regulated subsidiaries in the UK. And we've been involved in the program since the very early days, and went over to London during the kick-off. And it's an

excellent methodology.

And the way that the bank structured it, there's really three parties involved. There's the regulated entity, of course. And there is your internal infrastructure subject matter expertise that you mentioned.

There is a third party. There's actually two. So there are two companies; private sector penetration testing outfit and an intel provider. And that's completely private sector. And then, there's the bank themselves as the regulators that are involved.

And the methodology is very intel heavy. It's very threat intel heavy. So, the idea is -- and I'll really dumb it down, and hopefully I won't speak out of turn on this, but it is what's been going on in your sector before? How have people broken into your peer institutions? And let's try the same thing against you.

And that's where the threat intel comes in. So, we went down that road. We looked at the methodology. It is different than what's been

done in the past. We've been doing pen testing for -- I can personally say for at least the last 14 years, so this isn't very new. But adding that threat intel component is, to a degree.

And we went down that road. We took -- the methodology is published. We retained two vendors that were on the approved list, because there's also an accreditation piece to the CBEST program. And we engaged them, and we said, okay, we want to conduct an exercise along the CBSET guidelines, and it was actually a six month long exercise. And I can tell you, we've never had a pen test that went over that long of a period, and part of that is that it's very opportunistic. We did not let anyone inside know.

They did an amazing job with social engineering. There were conference calls made with employees. It was an amazing level of detail it went through. So, it's a very good idea. It's very effective. It was well written, well designed. Some of the challenges, on the other hand with it, were with the fact that there's a

regulator involved. All right?

So in general, I think any private entity has a set of the examinations that you want to pass and a set of the examinations that you really want to fail. And when you do things -- you know, generally, when you have a regulatory examination or whether it's Sarbanes-Oxley or whether it's year end, the ultimate goal is to pass. You know? You've done all your homework. You've put your defenses in place. You've put your controls in place. And now, let's have them come in and let's talk about it. And the end goal is, you want to get a clean bill of health there.

When you do a pen test and when you bring in a company like Steven's, you want to fail. You know? Because you want to find out about the holes any way that you possibly can. And if you don't fail, you're going to lower your guard a little bit and see if you do.

So, I mentioned earlier that vulnerability assessments can be scoped pretty easily. Really, the scoping with vulnerability

assessments is the target. So, I could scope it down to a regulated subsidiary and say, okay, here's a vulnerability assessment for ICE Clear U.S. -- something like that. There's scoping possible on pen testing, too. But it's not on the target. It's on the threat actors.

So, it's let me scope just to what could somebody in Eastern Europe do. Go. All right. Now, let's step back. What could somebody at Morgan Stanley do to our company? Now, let's step it down. What could an internal employee do? And then finally, what could a privileged employee do? So, there's scoping involved, but it has nothing to do with the regulated entity.

So, there's an inherent conflict of interest potentially there if you bring a regulator in to the table, because if you're testing from an Eastern European adversary's perspective, does that mean that the CFTC doesn't have a remit there? If there are -- obviously, it means more of the target, but they're not going to be limited to ICE Clear U.S. They might poke

around through a UK subsidiary.

And if the CFTC was sitting at the table with them poking around, then how am I going to explain to the bank, in that case, why the U.S. CFTC was breaking into a London entity (Laughs)? So, there's a lot of benefit to the methodology, like CBEST, but it really behooves us as private sector to use these third parties that we contract directly. And I can say, Steven, you have my authority to break in through any way. Come in through Singapore subsidiaries if you need to. Whatever it may be.

And then when it's done, if you as regulators come in, and you do, and ask for results of these, we can look at it and say, ugh, well, this would be scoped out, because it's germane to a different subsidiary. But here's the things that are relevant to you, and of course we want to show them to you. And then, what you always ask for is, what are you doing about it. And we go through the remediation plans.

So, I just wanted to bring that in,

because the bank has definitely paved the way, and they're ahead of the curve, I must say, with this type of thing. But I think that there are lessons to be learned in these early days, and before you guys go running (Laughter) down the same path, I wanted to at least throw some experience out there.

MR. TAYLOR: Dave Evans, let me ask you, just following on from what Jerry was talking about, how does the bank address the role of the regulator, as Jerry is saying here? And what role does a remediation plan play there?

MR. EVANS: Yes. I'm quite happy to pick up those points.

So, the role of the regulator is -- to begin with, is very much an observer role. There's a number of people on the same team as myself that understand the CBEST process. They know how the phases should work; who needs to be involved, when and how, and everybody that undergoes a CBEST test will be assigned somebody from my team to monitor the whole of the process.

So, they're there to make sure that CBEST is sort of being adhered to as a process; that no steps are being missed; that the scope is still within the UK financial stability arena. Just to make sure that the test, whilst slightly different for each organization, are following the CBEST process.

But the regulator is there in the room to observe, to provide input to the scoping, provide input to what's critical for the organization. That regulator then may, if it's a multi-national organization, may elect to open up dialogue with overseas regulators to let them know it's happening, to perhaps ask if they want to be involved. And that has happened on a number of occasions.

And the regulator will typically go through the whole of the process very much in an observing capacity. They're there to understand threats to cyber security in a little bit more detail, and they're there to understand what it might mean for the organization they've been asked

to supervise.

But Jerry spoke about our remediation plans, and David, you just asked about remediation plans. So, at the end of the test, and we now know what you know, some of the issues are that have been identified, and perhaps, some weaknesses in the cyber security posture of the organization, well, then, we go back to good, old fashioned supervision. We have some issues. They need to be managed. How are you going to manage them? And let's agree what that remediation plan looks like.

Do we like the -- do we agree with the time scales that are put in place? Are there measurable milestones? And then, it does very much, go -- you know, it's handed off to supervision in large parts, because they will now have a program or remediation to monitor. And that's what supervisors do day in, day out. So, to begin with, very much you know, observing, learning. But at the end, hopefully, they've got something they're comfortable and familiar with,

and they take forward you know, with the regulated entity.

MR. TAYLOR: And does that -- and I want to ask this question of Jerry and Jerry, who get regulated. And I want to ask it of Kevin, who does a little bit of regulating. Does that help solve the problem that's inherently there for an infrastructure? Because as I think Steven was saying, or Jerry was saying earlier, you want to fail a penetration test.

But then, there's the issue of how does the regulator look at you when you fail the test? Does this remediation plan road -- is that the way to address this?

MR. EVANS: That's a -- likely, yes. I think the -- it's all in the matter of how this place -- how that interaction really happens. But I think that does really get you to the right road.

Remember that a lot of the other things we talk about around vulnerability testing and the broader topic of vulnerability management, those

are operational practices. Those are things you expect to hit a hundred percent every time. They're related to other key controls, like patch management and configuration management. And at some level, those are things that are occurring every day, every week, all year long.

And it's a practice where each of those tests that you have, like vulnerability testing, is expected to verify that the things that you are going to -- pushing patches and configuring systems are really happening; that those hygiene aspects are really well managed. Penetration testing just tells you -- the prospect of giving you the views that you may not have yourself in any other way, a view from an attacker's eyes, where you would like to torque that volume up to the place where you fail, and know exactly where that red line is where you need to have concerns.

It won't always result in deciding that you're going to remediate a particular vulnerability. It may be within your risk tolerances, and it may be acceptable, or it may be

an unattainable goal to avoid that. It also tests other things, like people, process above and beyond just technology. But it takes you down the road to a conversation where you can have a very, very practical discussion around whether or not that is a reasonable outcome; or that unexpected outcome in a pen test is still reasonable, whether technology controls the right response or whether that is an acceptable risk or some other counter measure makes more sense; financial insurance, something with a counterparty. But it leads you to a very useful conversation, that if it's managed well, can be very, very productive.

MR. PERULLO: Yeah, I'll add to that. So, I've seen firms before when -- that will go through draft iterations with the pen testing firm and try to edit the results. So, to say, no, that's not a high. I think it's a medium. Here's why it is, and go back and forth.

And the fear there is that if they are ultimately impregnated with a report that says there's a high vulnerability that they'll be

responsible for that, and it may be totally out of context. The way that a firm really, ideally would operate is to let Crowdstrike, or whoever it may be, have an external perspective -- go nuts. Prepare a report. And if you think it's high, let it say it's high and it's done, and we get the report.

When we internalize it, on the other hand, we may say, oh, that was a red herring. That was a honey pot system. That was one that we -- and a honey pot system is one that you allow to be exposed and compromised, so you can find intruders. Or, it may have been miscategorized. It may be something that the tester thinks is confidential data, but really, it's completely public data; that sort of thing.

So ideally, we internalize these reports. We look through it. Of course, we look at what was considered high first. But ultimately, we put our own categorization on it. And you know, we document why that is.

Likewise, during the actual pen testing,

there's an actual pen testing, there's a back and forth between the private entity and the pen tester in near real time. Hey, I'm about to try this. Oh, that's not really us. Sorry, you made a mistake. It's a totally different company. That does happen a lot in pen testing, by the way, where you'll hopefully - not get to the results phase, but they'll actually pick the wrong company name or something like that.

So, there's a lot of that real time back and forth. If I'm getting that back to -- and sorry, I keep using you, Steven (Laughter), but you're a perfect example with your firm.

MR. CHABINSKY: As long as you're not using me for those examples where they're getting it wrong, I'm fine.

(Laughter)

MR. PERULLO: All right, well, I'll dance around a little bit.

But if you say, hey, we found this. Is it a big deal? We'll immediately take a look, and so, no, no it's not. And that's the end of it.

No, no, it's not, it's appropriate there. If you're paired with a regulator, that's not going to cut it. You know? I can't be on the record with a regulator saying, no, that's not important. And that's the end of it.

I owe you a formal response, and it has to be on the record. It changes everything. So, it is very challenging. So, I don't think it fully addresses -- and I also don't think that it's possible to fully -- for a regulatory entity to really completely take off a hat. Right? I mean, at the end of the day, you can't say, well, I'm just in an advisory capacity here, and then later on, put on a supervisory hat and completely wash away everything. You know?

So, I think it's very challenging, and I don't think that it fully addresses it. And I think it's great and really, the right way to be completely engaged with private sector testers, internalized results, and then, engaged with the regulators directly later, when we're dealing with not just pen test results, but our entire spectrum

of assessments and controls during an examination.

MR. WASSERMAN: So Jerry, one note I would make -- I think maybe part of the way to harmonize that is, you're right. We can't just say, no, no, that's not important, because the question I would be asking as your regulator is, well, why do you say it's not important.

MR. PERULLO: Mm-hmm.

MR. WASSERMAN: And if you have a good answer, then, yeah, then we can move on. But I think you know, you would be put to the burden of explaining why it's not important. I don't think you would be put to the burden of fixing 100 percent of everything regardless of how important it is.

MR. TAYLOR: Kevin, just as a follow up on the remediation plan issue, does that feel like a solution to -- there are tests you ought to fail? And then, how does the regulator deal with that?

MR. GREENFIELD: Sure. And I'll tell you, as a supervisor, I look at the penetration

testing and vulnerability assessments very similar to business continuity tests, in that the most successful tests are the tests that do identify issues.

And very much, when we come in and we will do a thorough review of the penetration testing, the results and issues identified, and what are the remediation plans, we're not focused on a -- well, there was a vulnerability or a gap. That's an issue. That's a regulatory issue. We're looking at the risk management process in place for were the mechanisms in there in place to identify, which if you're finding ensuring your penetration -- your regularly scheduled penetration testing and vulnerability assessments, and assuming they're not things that should have been identified long ago, that is a process. That is an effective process.

And then, looking for that follow up as to how do you remediate and how do you prioritize? Because with these tests, there often are a number of issues, and they all can't be critical. We

look to the institution. How do you identify what are the highest risk issues? How do you remediate those for some of the medium and lower risks? How do you eventually address those, or make the conclusion that it's something that does not need to be addressed, and demonstrate that there is not a risk to the organization?

And that's what we're very much focused on during our supervision, is that there is an effective process in place, because at the end of the day, this is all about making sure the institution is secure against threats and vulnerabilities, and not a compliance checklist of did you do A, B and C.

Because what may be adequate scope of testing today will be completely inadequate a year from now, even six months from now, depending on the threat and vulnerability landscape.

MR. WASSERMAN: Just for one second, I want to turn back to Dave, though, because we're talking a lot about you know, third parties doing this. And I know part of the CBEST program is a

sort of accreditation process for third party vendors.

And I was wondering if you could tell us just a little bit about that and how you think that might be applied elsewhere, outside of the scope of the UK?

MR. EVANS: Yeah, sure. So, you know, you're exactly right, Robert. So, accreditation of third party providers was absolutely an essential process within the CBEST framework. You know, there's been some media reports where it's the Bank of England that have got a team that are doing it.

I can you know, put on the record today that I do not have the skills to do a penetration test, and the Bank of England hasn't got the technology to conduct them, either. So, it's definitely not us. We rely on you know, third party penetration testers, and we rely on third party providers of commercial intelligence, as well.

In the UK, we do have an industry body

within the penetration testing arena that's been in place for a number of years, and they're called CREST, the Council of Registered Ethical Security Testers. Now, they're closely aligned with GCHG. They deliver a number of penetration testing services with the GCHA seal of approval.

MR. TAYLOR: Dave, sorry. Can you explain what GCHQ is?

MR. EVANS: Okay, sorry. Yeah, so GCHQ -- that's the UK's NSA. So, it's the national authority for signals intelligence authority (Laughter) for the UK. And they're charged with looking for threats to national security.

So, this industry body, CREST, offers a number of existing penetration testing schemes, and they've all received GCHQ approval. So, they have history in auditing these companies. It's an industry group that you have to become a member of. You have to provide references. CREST has audit rights against the company. So, if any third party is a member of CREST, they've already reached a certain level. They've got security

clearances. There's certain criteria that have already been checked.

So, if we were to put something in place ourselves, that would take time, resource, and effort, and we'd just duplicate what's already been done by CREST, and we probably wouldn't do it to as high a standard as CREST. So, why invent something new if it's already there?

So, we leveraged CREST's experience. We raised the bar of what already existed in terms of penetration testing. So, the penetration testing companies can apply to be a member of the CBEST scheme, but their requirements will be higher than what is currently required for any other CREST scheme.

When it comes to providers of commercial threat intelligence, we were quite surprised when we looked, that before we started, there was no accreditation for commercial threat intelligence. You could have a shared, with a laptop, access to Google and build a web site and sell commercial intelligence. And you know, that is what some

people were doing.

There were then, some people that are very, very good at providing commercial threat intelligence. But how are we going to put the Bank of England's name to this process? How are we going to differentiate between them? Well, let's do exactly what's happened in the penetration testing world over the last sort of, 10 or more years, and let's build in some accreditation for the provision of threat intelligence.

So, there's now examinations. There's the whole CREST membership which needs to be reached by the firms. And whether you're a penetration tester or a threat intelligence provider, the people that have ultimate sign-off for the accreditation is our team. You know, we need to go and do a site visit. We will check references. We will ensure that in terms of the threat intelligence, it's being done ethically and it's being done professionally.

In terms of penetration testing, you

know, we'll insist that examinations are met; the right number of people have got the right number of security clearances; that your data holding and data destruction techniques are all in line with GCHQ approved standards, and ultimately, it's our call as to whether you know, CBEST accreditation is approved.

MR. TAYLOR: That's the perfect segue. Let me turn to Murray Kenyon from our GCHQ (Laughter) NSA, who leads the stakeholder engagement efforts for the Information Assurance Director of the NSA. Murray, can you talk to us about what lessons NSA has learned that are most relevant to our efforts to protect critical financial infrastructure?

MR. KENYON: Yeah, certainly. We'll offer some comments there. I might just leverage off that. You know, we are the United States' GCHG. But it probably would be worth just a few words about where we fit into the constellation of security experts and security service providers.

As part of Department of Defense, NSA

has clear authority to operate right in Department of Defense networks. And that's really where we cut our teeth. In addition, however, Executive Order 12333 and National Security Directive 42 give the director of NSA authority to provide assistance, technology assistance to civil authority.

We do not have, as in direct response to your question, we do not have the authority to work directly with critical infrastructure. However, when one of our government partners, and the Big Three either are here or have been here today, and from DHS, FBI and the Treasury have all been here earlier this morning.

When, for whatever reason, they determine that they could use our technical assistance, we then exchange some paperwork and attorneys nod in the right direction, and then we can go into partnership to provide a variety of technical services; design guidance, operations advice, in some cases, mitigation tools that we may have developed, and certainly, kind of in the

broader scale, incident response.

What we find in doing that, and I will tell you that again, most of our work has been done in U.S. Government networks. It's only been in the last, maybe five to six years, that under those requests for technical assistance from our government partners, we have started to work more and more in supporting their authorities to work with critical infrastructure.

And what we have found, really, has already been said in a number of ways today. We have found that repeatedly, it's poor basic network management, poor security practices that provide or allow the majority of intrusions to happen, and often, with some of the greatest consequences.

We believe that job number one has to be standardization and automation of patch management. That, far and away, is the one thing we believe that could make the most difference. Following close behind that, though, is the notion of administrative accesses. Ann mentioned this

already.

Preventing those pathways to escalating privileges by segmenting accounts containing losses, minimizing privileges consistent with work role are absolutely critical. And we find again and again that that is not implemented in many, if not most of the networks that we examine.

We can, through a variety of practices, contain an adversary's ability to maneuver by minimizing work station to work station communication. That's another thing that we often find simply left wide open, whether it's you know, one of our government partners or one of the industry affiliates that we work with. Ann mentioned, as well, ensuring that you can't have unexpected execution of applications on your network. Hardening those applications and then limiting their ability to execute is critical.

And finally, certainly, with our Department of Defense clients, we recommend again, that as much of this as possible be automated in such a way that a host mitigation package of some

sort is implemented that would include things like application white listing, anti-exploitation features, anti-virus cloud look up, a variety of other features.

And many of those features are, in fact, provided by the technology providers, but it's bringing them together in such a way that they can be automated and managed in a way that, I believe it was perhaps Steve said this morning, you know, we need to manage these things in micro seconds. Automation is the only way to do that.

I might just mention, as well, that taking the lead from GCHQ, and I would note that NSA's authorities and GCHQ's authorities are significantly different in some ways. But the CBEST program has given us some guidance to create what we're now calling the National Security Cyber Assistance Program. We are, in fact, accrediting U.S. Companies to do the kinds of network vulnerability assessments that in the past, we would have done.

But as good as we are, if I do say so

myself, we simply can't scale to the need. And so, about a year ago, we launched down the path of working with some industrial partners. And today, there are 10 private companies, including Steve's. I make no implicit or explicit recommendation there. But Steve's company is on the list, along with nine others, that have met the standards.

And I would note that the standards that they have achieved deal with U.S. national security systems. And by and large, national security systems are defined as those that handle classified information, or those that are used for military or intelligence purposes.

And I sat across the table, I believe it was from the Secretary of Commerce a number of years ago, and he kind of gave me the finger, and he said, you can't say that my networks aren't important to national security. I quickly said, sir, that's not what I'm saying at all. Despite the fact that you don't meet the strict definition of a definition derived in DoD, clearly our financial sectors --

So, while I have no authority to do so, I have expanded in my own mind, that definition to working in national security systems and other systems of national interest. Clearly, the financial system is one of those.

The National Security Cyber Assistance Program that we're working with seeks to accredit companies in four key areas: Intrusion detection, incident response, vulnerability assessment and penetration testing. So, it is right up the alley of what we're discovering or what we're discussing today.

I would also mention that the fact that NSA is a large agency, and much of the information that we have is not shared broadly, I would call out two exceptions to that. One, that we share a tremendous amount of information with DHS for their mission. And that, in various formats, is then shared with other government agencies, as well as with industry.

The other is, in the information assurance directorate -- because much of what we

do is working on unclassified networks, we seek to produce our knowledge and share our knowledge in unclassified format as often as we possibly can, while protecting proprietary information and PII and those kinds of things.

But I would draw your attention to the NSA.gov web site, the information assurance button. Much of what I have already talked about today is published on that web site, and it is available not only to CFTC, but to industry partners, as well. We have such things as our top ten mitigations; our top technology challenges or things that we're on there.

We have architectural guidance. We publish white papers. One that I picked up this morning; defensive best practices for destructive malware, published right there. And I don't know that everyone knows that. We've noticed that some of the adversaries know that, and we've -- But in any case --

And perhaps, finally, I would say that again, in the interest of sharing information

broadly and enabling others to do the missions that we often can't scale to, we have developed a program that we call Commercial Solutions for Classified, which is using entirely commercial technology to provide -- to design and build networks and then operate networks that are wholly and entirely composed of commercial technologies; no secret sauce from the government in them.

But if implemented correctly, we have approved those systems, those layered systems or composed solutions, as we call them, for classified U.S. government information. That involved is available on that same web site, so it would certainly be available for at least consideration by members of critical infrastructure.

MR. WASSERMAN: So, let me press just for a few moments, because I must say, this National Security Cyber Assurance Program, if I got it correct --

MR. KENYON: Cyber Assistance Program.

MR. WASSERMAN: Cyber Assistance. I'm

sorry. Assistance Program -- is fascinating.

Is that among the things that is on that public facing web site?

MR. KENYON: Yes, it is. Absolutely. And one of the newest developments there is that as of the 23rd of this month, we're going to open up a new round of applicant -- a new round of applications for additional companies to join that. It will have a portal online that initial application can be submitted. But we feel like we had such good success with that first round of companies, that it's time to expand the program.

MR. WASSERMAN: And I realize, of course, that this is probably -- that is not any endorsement of anyone who's on there, but -- and you sort of touched on this, but I'd like to press just a little bit harder.

How applicable would this be to someone who is looking at, you know, critical infrastructure from a regulatory perspective?

MR. KENYON: So, I'm not sure I can answer from the regulatory perspective. But in

terms of the service provided by these companies, it is essentially the same service that NSA would provide you, where you -- the operator, owner operator of a national security system, and you asked me to come in and help you ensure that you didn't have unpatched vulnerabilities, and in fact, to do penetration testing of a classified system, as an example.

MR. WASSERMAN: So, if we were hoping to have critical infrastructures protected and resilient at the highest achievable level, this would be one place to go?

MR. KENYON: I think it would certainly be a resource. Steve, you might be able to expand on that, as well.

MR. CHABINSKY: Well, I think certainly, from a vendor perspective, you want to ensure that any company you are working with has the proper credentials and is following processes that are recognized in the industry. This is one way of doing that.

Crowdstrike had to be positively

assessed in, I think it was 21 critical focus areas in order to achieve that type of accreditation. And that is not otherwise a standard that is -- there is no private sector standard, I should say. There is no accrediting body in the private sector that otherwise exists. So, it certainly is one place to look for a view of whether or not your vendor possesses qualifications that are consistent, not only with best practices, but rise to the level that would be necessary for national security systems.

(Simultaneous discussion)

MR. KENYON: Can I add something? And to be very clear, we are directing some of our federal government customers to those same companies.

MR. PERULLO: Yeah, I just wanted to -- I know this isn't about information sharing today, which is amazing, because almost all of these are about information sharing.

But we got drug into it a little bit there. So, just to quickly outline the

information sharing flow, and anyone can correct me if I have this wrong, in the States, and in particular, in financial services, we have a really good system in place, thanks to the FS-ISAC.

And so with intel, I have to assume but never will know that it ultimately came from, let's say, the NSA. It will flow through, let's say, ultimately, the NCCIC, and it will get to FS-ISAC and members such as us. And it's very effective and it's working really well. And we have some smaller groups within FS-ISAC where they can deliver targeted intel, as well. And it's working really well.

But the reason I wanted to bring it up is that when we get to accreditation, for example, what I think we need to steer clear of is the idea that you can only get the threat intel if you decide to participate in a certain program or something like that. And Dave, if I can ask you, in the UK in particular, there's been confusion about that; that whether or not you can only have

access to GCHQ intel if you sign up for CBEST, for example.

How do you avoid that tension to where you were holding back threat intel?

MR. EVANS: Sure. So, in part, it goes to GCHQ's roles and responsibilities. So, it's very much focused on UK national security. And what they're looking for is probably no different to any national signaling authority.

They're looking for threats to critical national infrastructure. So, there are far more organizations that operate in the financial services space that are not critical national infrastructure in the UK, than are.

So, as soon as we start drawing that distinction, then there's already a connection between the critical national infrastructure and what GCHQ has or may have access to. And those relationships already exist. What we've done through CBEST is improve the mechanisms in the relationships for those to work.

It does mean that -- that means there's

a large population of your financial services that may not directly benefit from anything that GCHQ has, but through the likes of FS-ISAC, which now has a European arm to it, and of course, lots of the companies that operate in the UK are multi national anyway, so are probably part of the U.S. Branch of FS-ISAC anyway, but we have some information sharing platforms that are led by the UK government, and they will also be taking feeds from you know, GCHQ and other government sources.

Although back to your point, Jerry, you need to participate in those, and you may receive information and just have to assume that that's where it's come from. You may never truly know.

So, CBEST is not there as a, hey, this is the only way you're going to get it, but there will be improvements made in your relationship with GCHQ by participating in the CBEST program, just because of the processes that are there that don't exist in any other mechanism.

MR. TAYLOR: Let me turn in a slightly different direction. And this is actually a

question that came in from the audience. Anybody feels so moved, you're still welcome to be doing this. But it's a topic we wanted to discuss, at any rate.

And the questioner directed to this to Jerry, Jerry and Steve, but anybody can join in. And it spun off of the fact that as Dave was relaying -- or no, actually, Jerry said CBEST took six months to do the whole penetration testing cycle for ICE --

So the question is, how long should a good pen test cycle take? And what's the optimal frequency without you know, breaking operations?

(Simultaneous discussion)

MR. TAYLOR: And this is critical for us, as well.

MR. EVANS: Dave, before I answer the question directly, I should be very clear. So, we did not have a CBEST engagement, per se, because we did not have the bank party to it. We followed along the methodology and we used the accredited testers and we went through the exact same steps

of it.

But I want to be clear on that, because I know the bank likes to be clear on that, as well. So in our case, it was really as long as it takes to break in. That's the most effective pen test. But as far as frequency goes, we actually have a huge number of things that qualify as a pen test. We have a whole -- and a lot of them fall into our application development software life cycle.

So, we have hundreds of software applications. We identify them. We tier them by exposure, so if they're external facing, they're higher priority to us. And we walk them through a life cycle that includes pen testing. So, some of those tests may be very micro engagements, if you will. Some of them are third party. Some of them are in-house. And there will be a lot of those going on.

But for any given cycle, usually, the bar is annual. That's what I hear a lot. And I think it goes back, again, to the whole thrust of

this engagement today. You know, from an examination standpoint, the questions are usually do you have a pen testing program and does it meet the standard. And I think it's an at least thing.

So, annual seems to be what's thrown around out there. We certainly strive to do it much more frequently than that. But if you came in, and the question was, can you demonstrate that you're operating a pen testing regime, we would want to make sure we could always show at least one, during that minimum cycle, which currently is annual.

MR. TAYLOR: So, you might say at least --

MR. PERULLO: At least.

MR. TAYLOR: -- as a minimum --

MR. PERULLO: Exactly.

MR. TAYLOR: -- annual is there. The question was also for Jerry and Steve. So, jump in.

MR. BRADY: Yeah. At some level, it's useful to drive these programs to multiple tests

throughout a year and multiple frequencies, depending on the things that are being tested.

Some of the things you test, call centers and people, in particular, benefit from long lived pen tests that are sort of low and slow like a bad guy might do, as well. But oftentimes, you're trying to mimic the behavior of bad guys, so you want to sort of use periodics that are useful.

Oftentimes, you're testing against some new technique or emerging technique that is useful to do on an off cycle. But you want a program, at least, in a calendar year that shows the amount of coverage across your infrastructure and people processes, and harp on the ones that change often or are frequent targets of activity. So, key control infrastructure. Things that are authentication systems, Internet facing, client facing and so on, yearly, makes a lot of sense.

For things that are more long lived and less significant, frequencies that are more sparse, things like two years and three years make

sense, but all part of a program where you can look at the end of the year and say, this is what I've got coverage of. This is the level of confidence I have, and it fits the bill of that assurance.

The things that make more sense today than they did maybe a few years ago are driving us around the intelligence theme of changes in activity, changes in themes that break those calendars. So, I think the traditional, I'm going to pen test every year doesn't make a good amount of sense now. Having a yearly program that shows a lot of coverage across the shop makes a lot of sense, and then using intelligence to prompt when those tactics need to change or the boundaries change or so on, makes sense.

But it's a broader program. It's hard to say every year, every two years. It's all of the above, in a program that makes sense and gets you coverage across the year, so you can speak to clients, regulators and your own desire to know that you're operating within your risk tolerances.

MS. STEWART: Can you just clarify --

(Simultaneous discussion)

MR. TAYLOR: Oh, go ahead.

MS. STEWART: Sorry. Can you clarify in that annual program, how much of that testing would be performed by a third party?

MR. EVANS: Sure.

MS. STEWART: -- and how much of it would be internal.

MR. EVANS: And lots of people have different preferences in this space.

I generally would prefer to see all penetration testing occurring by a third party, and lots of other kinds of control testing happening more frequently. Oftentimes, automation by first party kinds of testing. But we generally do penetration testing on a third party basis because it's very useful to get that external perspective; not tainted as an owner, and to know that it's independent and you can use it for more purposes; to demonstrate to our regulator, to demonstrate to our client that things are

operating the way they should.

Insiders oftentimes have different skills that are very useful for the recurring testing, the control testing, maybe not so much getting out and knowing more about attackers themselves today. So, I generally see splitting those two -- control testing, then automation when possible, internal staff, often, penetration tests majority or exclusively by third parties to the independent aspect.

MR. TAYLOR: Let me go to Steve, because the question did --

MR. CHABINSKY: Yeah.

MR. TAYLOR: -- but I'd also like to get Ann and Kevin to chime in on this frequency question.

MR. CHABINSKY: I think one thing that you're hearing is that there is not something called a penetration test. All right? There are different tests, depending on what's being tested.

You have web applications, external network scans, internal testing. And so, right

off the bat, there's this recognition that there's not just one -- you know, did you get your scan. Right? It's a question of what is being reviewed.

The other thing that you're hearing is that there is nothing really static in this space. The systems being tested are dynamic. They're changing constantly, based on software or network architecture. And the bad guys aren't static. They're dynamic.

I'll never forget a conversation between a CFO and a CISO where the CFO said, I just gave you all this money last year. How come you're asking for more? And the answer was, had the bad guys stuck to what they were doing, right, I actually wouldn't be asking for more. Or it just as easily could have said, had our architecture remained the same, I wouldn't have been. Right? So, you have two dynamic things.

And then, the length of time of a penetration test, of course, is going to vary based on what you're testing. But it also is based on what information is provided to the

tester. And that doesn't sound intuitive, but we engage with our clients, and the first thing we ask them is, how much of the work do you want us to do to get to this level.

You know? We heard discussion earlier about the rate of opening up a phishing email. Right? It might start out at 60 percent, then 40 percent, then 20 percent. You might get it down to single digits, but that single digit is not the digit zero.

And so, our first question is, do you want us to actually try to send the spear-phish and get someone to open it, or should we just save that time and you'll give us just a computer, and we'll open it from an external source, just to see if you can detect the malware coming into your environment and opening it.

And the more information we get from the client, the shorter the engagement and the less expensive the engagement. But it also isn't testing certain processes. Right? So, the notion I'm getting across is, when I'm thinking about

penetration testing, I typically think of something that's occurring in weeks, definitely not in months. But there's a time and place for different types of testing your systems.

MS. BARRON-DICAMILLO: I definitely concur with both Jerry and Steve. I think the frequency of pen testing is exactly what Jerry said. It's how long does it take for them to get into the network. So, it's not a hard and fast number that we see. Sometimes engagements are days, and sometimes they're, you know, a couple of weeks. But I also don't think they're months. I think if you have a pen testing engagement going on for months, then you probably need to get back to the criteria associated with what you're going after.

From a vulnerability testing, I definitely concur with Jerry. This should be automated as much as you can make it. Getting back to instant response engagements, and then the mitigation plans, I also think you know, the testing done post mitigation is so critically

important to almost kind of the auditing aspect.

One of the things, we're not the auditing function within DHS. We don't do that role. But many times, in the engagements that we have even been part of, we give them the mitigation plan based on the assessment from the intrusion, and they want us to come back and validate that. That has to be done by a third party. It can't be done by the team that provided -- that did the assessment and then, provided you the mitigation. So, that's another aspect of vulnerability assessment, you know, kind of the auditing function that needs to be captured and done by that third party. It can't be the people that were engaged in the assessment or internal. It needs to be done externally.

MR. BRADY: Just one thing to mention on that front. There is a separation of duties here that makes good sense. You don't really want to be testing your things that you designed or operate. So, your question on why some things are internally done and some are independently done,

separation of some duties is very important, not just because you may not be ethical in the way you execute, but you may not look for things that you didn't contemplate when you designed or operated when you're testing. And that's a very important and different perspective that an external tester brings to the table.

MR. GREENFIELD: Sure. And a lot of good concepts have been brought up in this discussion. I think the key thing is to make sure that when looking at penetration testing, you're taking a risk-based approach, looking at your environment. How many different applications are you running? Are those applications constantly being updated and changed? Are you on a leading edge operating system that is constantly getting new patches, new updates? Or, are you operating in a static environment with two or three basic products?

That's what's going to drive the frequency, scope and depth of a lot of your penetration work. If your environment doesn't

change, you have the same set of three products or four products that you update once a year, annual penetration may be sufficient.

But if you're working in an environment where your network is constantly changing, your products and services are constantly changing, that penetration testing, that vulnerability scanning needs to keep up with those changes. I know annual penetration testing is a guideline that many people follow, but if I conduct penetration testing, then I change my network environment or a new patch or a new vulnerability comes out three weeks later, am I going to wait another 300 plus days before I do that testing again?

So, it's very important when setting your standards that, what are you testing, why, and what is the risk to the organization? Because some penetration tests are two or three weeks focused on a specific application or focused on a specific segment of the network.

What hasn't come up is, there are

penetration tests where I'm not looking at your network. I'm calling in to your senior executives posing as a network administrator, trying to phish for passwords, for credentials. We see a lot of that being done now. That may be a longer term.

It's really how do you scope in the concept, being how do I structure my penetration tests to try to break in to the environment, try to identify gaps and controls from every aspect that a malicious actor would be taking and thinking of.

MR. TAYLOR: We have about 10 minutes left. This is from the point of view of, as Bob said, the people up here who, in the end, have to write something for the Commission. This has been tremendously valuable. And I want to be sure we come to the topic of setting an adequate scope for pen and vulnerability testing. We've touched on it in a bunch of ways, but I want to get the panel to draw it together.

But before we do that, there was one very important thing, I thought, in what Steve

said and a couple of people echoed. I think you said there's not just one thing that's a penetration test. It's a penetration testing program.

How do we describe the adequate penetration testing program that a critical infrastructure ought to have?

Jerry, you want to write your own rule (Laughter)?

MR. PERULLO: Yeah, I should have brought a copy of our policy on penetration testing (Laughter) and just fed it right to you.

I think that --

MR. ORTLIEB: I have that all right, so --

MR. PERULLO: Yeah, that's right, Jim. Jim has definitely seen it before.

And I'll talk about vulnerability assessment, as well. So, the scope is important. On the pen testing, at a minimum, it has to be looked at from a threat adversary standpoint. So, looking at outside the company, definitely, all

critical infrastructure should be doing pen testing from an external viewpoint entirely.

Where it gets more gray, I guess, is as you step inside the walls and you do internal penetration testing, or where you know, Steve mentioned allowing somebody to just assume you were phished, and then go from there. But -- or even take it down to, what if you had a rogue employee? What if you had an insider?

So you know, in general, we will exceed the bar that's set by regulation. We have our own motivators, too. It's aligned with, you know, our shareholders' concerns, certainly. So, wherever we set the bar for regulation is not where we're going to end. You know, that's not the end of it. It's just the beginning. So, I think that that external threat and making sure that we're looking at things from the complete outside is at least one place where that bar could be.

On the vulnerability scanning, there are similar analogs. So, we have a lot of automated vulnerability scanning going on, for example. And

from the outside, just looking for holes, it's as frequently as daily or weekly. If you step up that, then you start looking at internal vulnerability scans. So, these are viewpoints that the outside world wouldn't even have. So, you're looking for vulnerabilities if somebody were to break into the network.

And then, the final leg is what we call authenticated scans. So, now take a user ID on a critical infrastructure system that's already on there, and those are a lot of assumptions. So, assume a bad guy got all the way there, and then run a vulnerability assessments on these servers, and generate a long report of things that should be fixed.

So, there's a lot of context you apply to there when you internalize that and decide whether or not you're going to fix these things. So, I think from a regulation standpoint, looking at what's practical to -- what has happened in the past to actually compromise infrastructure is the most important point.

And going in and asking for a hundred percent of things that a system administrator with root privilege would be able to see, I think that's pretty far fetched and over-reaching. So, I think there's definitely a common ground, but it's not -- you know, it's not everything in one swoop.

MR. BRADY: I don't disagree. I think the problem is that these testing programs, like a lot of other security programs, don't stand alone very well. Trying to define what is critical infrastructure to you -- what do you depend on, things like authentication services, online platforms, authorization and so on, is one take at defining what is the extent of a set of tests that comprise your penetration test program.

And I'm looking at it from it an eye on attacker and describing what are the outcomes that you're most concerned about -- but it all comes down to coming up with some model for defining what you're protecting and what you're protecting against, and defining critical infrastructure that

represent key controls, like authentication authorization, the leakage controls, et cetera -- access controls and so on. Then, those outcomes that are unacceptable.

And putting that together into a program that both tests on an appropriate frequency with things that matter the most to ensure the controls are operating, and the things that you're trying to protect the most against new and emerging tactics or the threat actors that might go after those, that's what gets you to the right program.

I don't think you can call out the five things that you should pen test or the two threat actors you pen test against without starting off with that view of, what am I protecting, what am I protected against, and what is that key infrastructure that supports that whole security operation. And that's a little bit of that risk assessment and threat assessment that leads you down the path of putting together a real pen test program or a real vulnerability management program.

MR. TAYLOR: Would anybody else like to chime in on setting the scope?

MR. PERULLO: One thing we didn't talk about too much was remediation. So, I think from a regulatory perspective, it's very fair to ask about the work flow for findings. So, even if you say you have to have this type of test and it has to include these things, I think it's fair to ask us to talk a little bit about what we do with those findings to make sure they're all run to ground -- so the things that we do come up with.

And you know, that definitely happens already in inquiries. But, I think demonstrating that we have a program that gets eyes on things and gets them closed out timely is fair.

(Simultaneous discussion)

MR. EVANS: David. I mean, I just want to sort of clarify that. You know, over in the UK, we don't have a sense of how often a CBEST should run. My gut feeling is it will vary depending on the organization that's being tested.

It will vary on how that organization

changes over time. It will vary on how the threats to that organization changes over time. And it will certainly vary depending on how robust we think their approach in the individual sets of tests that Jerry and Steven have already outlined.

I mean, we've had analogies about bears and pandas and dragons (Laughter) and who knows what else. But there's another one we like to use, which is an airplane. And if you consider all of your security controls are components of an aircraft, you will have some wings, and you will have somebody that signs off to say these are definitely wings. They produce lift and they run on aircraft fuel. They're definitely wings. I can sign off to that.

And on a periodic basis, you're going to need your wing designer to say, yep, these are definitely wings. Somebody has got to build a fuselage. These have got to be made out to a light material. It's got to have a cockpit for the pilot, et cetera, et cetera. And periodically, they'll say, yep, it's definitely a

fuselage.

You might want some landing gear, I suppose. So you'll go through exactly the same process. But you're not going to get a passenger on there until you've bolted it all together and actually proven that the thing can fly. There might be wings, but they might be the wrong size. There might be a landing gear, but you might not have pumped up the tires.

The cockpit might not be big enough for the pilot, and you might have forgotten the tail plane completely. But you have received periodic updates from your designers and your -- you know, your controllers that everything is all right. It's not until you bolt it together periodically that you actually see that the whole thing works. We just don't know how periodic that needs to be.

MR. GREENFIELD: Okay, and just one other --

MR. TAYLOR: How --

MR. GREENFIELD: I'm sorry. Just one other aspect that hasn't come up; that it's one of

the bottom baseline fundamentals, but definitely needs to be something that's incorporated as a fundamental part of your penetration, your vulnerability assessment program. And that is, what is your asset management program?

And that's what are the components in your network? Do you know everything that's present in your network that you can ensure it's tested; that it's scanned and updated? One of the fundamental principles of security is, you can't secure aspects of the network you don't know exist. And in large organizations, that can be very difficult.

MR. MCGONAGLE: And just as a practical question in talking about the external testing. The NSA is thinking about opening up for additional entities to come in to get certain accreditation. You talked about accreditation with respect to the Bank of England.

So, my question is, just sort of the availability of third party vendors to do the work that they're being tasked with, and maybe, on an

increasing basis. I mean, do you see -- let's say particularly at the Bank of England, when you're doing these reviews, that there's a sufficient number of vendors that are available for selection to exercise these tests within a -- you know, a lot of time frames.

MR. EVANS: Yes. So, as it stands today, we have -- there's enough providers to meet the demand. We were quite concerned when we launched CBEST last summer that we might not have enough providers. But you know, that didn't come to the fore, so that was quite good.

But we currently have -- but of course, if this takes off and more sectors follow our lead, then there might be a contention between supply and demand. But not as we sit here today.

MR. CHABINSKY: We haven't seen that as a problem yet. And one of the reasons is because penetration testing can be scheduled, as opposed to incident response, which is very urgent and immediate, and you don't know what team you might have available in a location.

The really good part from a supply and demand perspective for penetration testing is, if it's done correctly, it's not we need you in this afternoon (Laughter) because we think we have a problem. That's not penetration testing. It's quite scheduled. It's far easier, therefore, for the vendor to make sure they have the right resources available at the time that's consistent with the client's demand, and we haven't seen a problem in that regard. If anyone sees a problem getting a vendor, they should contact me (Laughter).

MR. TAYLOR: All right. I hate to stop, because this has been enormously valuable from our perspective. But it is time for lunch. I guess people who are coming back in the afternoon will thank me for leaving the lunch hour to be an hour.

As we break, we will resume again at 1:30 with the next panel on key controls testing. Bob Wasserman has some tips, without endorsements, on where there's food close to here.

MR. WASSERMAN: Yes, yes. I'm not

accrediting anyone. (Laughter)

(Recess)

MR. WASSERMAN: I'd like to thank everyone for coming back so promptly from lunch, and thank as well our panelists. I'm going to mention a couple of administrative details which some folks may have already heard. Panelists need to press the button to activate the microphone when they start speaking because, both to make sure the folks in the room hear, but as well, we've got some folks connecting, dialed-in through audio and this is the only way they could hear. And if you forget to do that, you may see me pointing towards my ear. When you are done speaking though, if you could then turn the microphone off, because we can only have a limited number on at the same time. Members of the audience, there may yet be some of those question cards left on your seats and so if you do have questions, you can write them down, legibly please, and we will be picking them up periodically and taking them down and we will try

to seed them in, probably towards the end of the panel. There should be more question cards as well on that table over there. And what did I forget? That was it. Okay, in which event, our third panel here is on key controls testing, and I'm going to read a possible definition of key controls testing, and folks on the panel may well have something to say about that. And we're looking at it as assessment, in our case, of the registered infrastructure's operational and automated system controls to determine whether such controls are implemented correctly, are operating as intended, are sufficient to address all material identified vulnerabilities, and are enabling the registered entity to meet the regulatory requirements. And so I put that down again, just sort of as a marker, but it is certainly, well, open to question. And so, I think I would like to start with a question to Tom Millar. And Tom is the Chief of Communications for US-CERT, and you can tell us about US-CERT very briefly. But the question is about the types

of key controls that are most effective in protecting our focus, which would be towards financial market infrastructures.

MR. MILLAR: Well, US-CERT's the United States Computer Emergency Readiness Team; it serves as the National CERT for the U.S. and is part of the Department of Homeland Security. As the Chief of Communications, I support US-CERT in terms of sort of outreach and awareness activities, also customer engagement or constituent engagement with our information sharing and analysis partners, our international counterparts and so on. And the key controls we've seen, and this is from our incident response perspective, what we've seen lacking in various types of enterprises over the last two years, where we've been involved in quite a lot of on-site engagements, first of all -- network segmentation, for example, and sort of the rule of least privilege, are two of the things that people generally are not sustaining. I think a lot of people when they first initially design their

networks or stand up a new system, or endeavor to protect their data and their customers' data, are very disciplined about setting up limited accounts, segmenting their network appropriately, firewalling off their DMZ from production and such as that, but over time, as new systems are deployed or personnel turnover, these things get soft. And a great deal of our incident response engagements, especially where we've seen these massive PII breaches and other sensitive customer data breaches, we've discovered what we call super flat networks, which is to say that segmentation is not there. We've also seen that the rule of least privilege is generally not followed and that people are using workarounds, so that they can, from their perspective, get their job done easier.

MR. WASSERMAN: Okay, and one thing I should mention, folks, panelists, as you -- to the extent you use acronyms or technical terms, I'm going to press you if you would to give us some definitions. And so, we've heard about PII, but two of the things you mentioned are the DMZ and

rule of least privilege.

MR. MILLAR: Well, right, thank you for asking me to clarify. When we say rule of least privilege, this is for example, if you work on a corporate network where your computer is issued to you by your employer, you may have noticed from time to time, you are not allowed to install that software that you need. That is part of the rule of least privilege. And sometimes it can be inconvenient, which is why people usually work around it. The idea is that you do not have any more privileges to do things to your data or your system, than is absolutely necessary for you to do your job. And this is always a -- well it's frequently contentious between work force or systems administrators in a technical environment and the security personnel. It's very important to adhere to this because what we see is that, as soon as you give somebody the ability to install whatever software they want, sometimes people will just click on that email and install whatever software the bad guy wants. And this has been a

-- this has sort of been the soft underbelly for a lot of institutions that we've had to work with over the years.

The other aspect of the DMZ comes from the term for demilitarized zone, which is, anyway, the point being that basically where your web server is, that anybody in the public can access from their phone, from any random device, should not be the same place that your financial management system is. Those things should be in different parts of your network and protected with different controls. And that's also extremely important, because again, what we've seen is, many institutions and organizations that will allow that line to become blurred, and all of a sudden the place where their public web server, which is available to anybody with an internet connection, and the place where parts of their, perhaps their financial management systems and the back end, or their systems containing privacy data, are actually, network-wise, in the same territory, which is very dangerous, which you probably

understand.

MR. WASSERMAN: Would anyone else like to jump in on this one?

MR. GREENFIELD: Yes, when we talk to controls and network controls, Tom brought up a very good point about things change over time, and one of the key controls that we focus on is change management controls, is making sure that over time, a network environment will evolve and change, software operating systems are updated, how are you ensuring that those changes are understood, documented, approved, tested, before they go into production? People's roles change and as that occurs, making sure that access is changed as those responsibilities are changing for their functions, and to the point of rule of least privilege, if I move on to a new job responsibility, all my capabilities on the system for the prior job should be removed, and then only what's needed for the new job responsibility added on. Often we'll see, you'll just continue to collect additional authority to do your job, but

your previous responsibilities or capabilities haven't been removed.

And then the other key aspect under that change management control concept is not just systems, but operational processes. There are a lot of controls that are not necessarily technology, but operational in nature. As those processes change, is someone making sure that those control structures don't degrade or disappear altogether over time?

MR. WASSERMAN: Ron?

MR. ROSS: I would agree. Least privilege is certainly one of the most important of the key controls that we need to be concerned about for the reasons that Tom talked about, and change management also. I think that one of the other big ones that is responsible for a lot of our discomfort today is least functionality. That's the other major area that we really don't do a very good job at. It has to do with complexity. And when you talk about all the talk about testing, whether its vulnerability testing

or whatever kind of testing you're doing, the sheer complexity of the networks and the systems we're building is almost unmanageable today. And it's largely because the very basic principle of security, least functionality, we violate every day. And it has a lot to do with the technology and how we're driven toward all of the great new technology. I use the analogy, if I was at a movie theater right around Christmas time, and on the screen the guy says, "you can download an app that will tell you the optimal time to go to the restroom during this movie," and that's a metaphor for where we are today. We are consumed by the technology to the point where we cannot buy enough of it, and that complexity is building from the hardware to the operating system, to the middleware, to the applications. And the result of that is that we end up having networks that are largely indefensible. And so going back to those fundamentals, like in football, no matter how fancy your playbook is, blocking and tackling always come first. And so those fundamentals,

least privilege, least functionality, change management, and all of those things, those are going to be discussions for the leadership and the culture of organizations that are going to be responding to the things that you're going to be working on, and that's going to be a big issue out there. Because it's hard to change the culture, as Tom was just talking about.

MR. WASSERMAN: So when we're talking about controls, is there a way -- are we looking at automated controls, manual controls, all?

MR. ROSS: They're actually, in the NIST Special Publication 800-53, we used to have, what we talked about, three categories -- management, operational, and technical. Many of the technical controls that you would deal with, access control mechanisms, identification, authentication, two-factor encryption -- all those things, the firewalls, those are largely buried in the commercial products that you buy, the operating systems, the databases, the network devices. There's another class -- two classes of controls

called management controls and operational controls. Management control is doing a good risk assessment. It's a management level activity. Operational controls might be something like developing a contingency plan. What happens when the malware brings down your system? What do you do? What's the plan B? Most organizations today unfortunately are not getting that plan B up front, tested and evaluated so they can understand that they can go to a backup and have that resiliency of the critical mission. That's really what we're talking about, a resiliency.

MR. WASSERMAN: And so I noted Ron, you are of course a Fellow at the NIST, National Institute of Standards and Technology, and you mentioned a publication, 800-53, which I've come to learn is very important in this area. So if you could perhaps tell us a bit about 800-53 and how we can apply it in the area of financial market infrastructures.

MR. ROSS: Well 800-53 is one of our foundational security guidelines that we produced

under our responsibilities under the Federal Information Security Management Act of 2002, recently updated in 2014. And in that, it's a catalog. I call it the great parts bin of security controls. It ranges across 18 different families, everything from access control, identification, authentication, incident response, education, training. It is a full spectrum of controls. There's about 860 in the catalog and it's part of a risk management framework that we publish that really guides our customers on how to select the right controls for the mission that they're conducting, in this case, financial operations, the environment in which they operate, and the technologies that they're using. And so the risk management framework is a flexible framework. It's not every organization, every company, every agency, even within our own federal government; they don't end up with the same sets of controls, because their missions are very different. But the framework allows you to customize and tailor, and that's what would be

advisable for every sector to figure out what is essential for them, and use the framework accordingly.

MR. WASSERMAN: So actually we've been speaking a bit about key controls, we've talked earlier of course, about the threats that folks are facing, particularly in the financial sector, and I was wondering if any of the representatives of the infrastructures or members of infrastructures might comment on essentially how they see key controls frameworks.

MR. CLANCY: So this is Mark. I'll start. So from my perspective, I think the fundamental challenge in the cyber risk domain is, everything works at the aggregated level and which you communicate around a PowerPoint, and it all goes wrong in the detail and in the environment. And the real issue is the difference between those two points, right? And so in our infrastructure, we have thousands of systems with thousands of pieces of software, with lots of functionality and lots of privileged people. And that minimization

theme is definitely one that we subscribe to, but admittedly, we completely struggle with, because all the things we buy and consume aren't built that way. And so, for example, in the minimization of privilege, we focused on those people in our environment who are the most privileged, whose access rights could cause us significant harm, when we started a project we called the IT-300 after the movie the 300 Spartans. We didn't actually know how many we had when we started the project several years ago. The number turned out to be smaller than 300, but we didn't know and so the first thing was like, what are those things that could really hurt us if they're abused? What are they, where are they, who has them, and why? And then we've been successfully narrowing down the people who have them, the circumstances which they have those rights, and the mechanisms they can use to get to unlock those rights and pull it out of the vault. We have a process we call break glass, which is named after the metaphor of the fire extinguisher

in the hallway with the glass. You break glass, pull it out, and use the fire extinguisher. And the reason you do that is you want to know if somebody used it, so you make sure you can recharge it, because they might need it again, right? We're doing the same thing with administrative access for our most important access rights. To go to Tom's point, we started looking at the way that we access our systems, and the average size of the operating system was measured, usually 30 to 50 megabytes in size. The tool that we use to get to our privileged credentials is one and a half. Now it's just a lot smaller so there are less things you can do to make it go wrong. All right, so that minimization piece is very important. So when we look at our key controls, we look at those things that keep our system integrity high, so I mentioned earlier in the first panel, patching of application vulnerabilities, the white listing of software -- you know we're a big proponent of using virtualized desktops where all of the software is

described by the system, not by the user, which gives us a whole lot of advantages in terms of repair and remediation, and that removal of access rights. And in traditional thick desktop environments, removal of access rights is very, very hard. And that's why these virtualized environments and these separate administrative environments are so important, because that breaks the chain of the feature creep. You know, do I really need to be able to open a Word document and browse the internet when I'm at the command line updating a system? And the answer is no, but what the attackers have found is because I've connected those things historically, they have an attack channel they can exploit. And that's really what's been happening, and so some of the earlier panels talked about that in the threat side and the vulnerability (inaudible). So from our perspective we look at those kind of things.

Then the second order of key controls are those things that test the effectiveness of whether our processes work. I will say one

failing of the cybersecurity experts has been, we're very good at adding capabilities. All of the security tools we buy have focused on this anomaly detection, meaning, if something weird happens we tell you, and zero of them let you know if you're collecting all the data that tells you whether or not you have the anomaly. And so there's a structural problem in the tools that we procure. And so me, as an end customer, now has to build apparatus to inject signal into all of those tools to see if they're actually functioning normally. I have to do the same thing with my operational processes, and I have to do the same thing with my management processes. And that's, from our perspective, not where we have been, but where we need to go with key controls testing, is to inject that signal and that noise into the environment and make sure our stimulus response to it is appropriate based on what those things are.

MR. WASSERMAN: So one thing, a number of folks have, through the course of the past two panels, mentioned the term white listing.

MR. CLANCY: Yes.

MR. WASSERMAN: If you could tell us about that.

MR. CLANCY: Sure. So white listing is simple in concept and hard in execution. It's basically saying, here are the 27 software programs that should be on our workstation for people to do their job, and the 28th can't run because it's not on the list. All right, and there's technical enforcement mechanisms, but it's the intersection of sort of policy and technical implementation procedure to enforce it. And the idea is that, instead of trying to stop everything that's bad, only let the things that are known to be good, run. And that's a very powerful concept and quite frankly, the reverse of where the security industry came from. And we've always been about enumerating and stopping badness, not about defining goodness. And I think, what was it, two years ago? Symantec published a report that indicated there is more malicious software than there is good software, and has never come

back. There is not more good software by count than there is malicious software, and I don't think that will ever change.

MR. WASSERMAN: So Kevin Greenfield is Director for Bank Information Technology at the Office of the Comptroller of the Currency, which is a very important part of FFIEC, the Federal Financial Institutions Examinations Council. I got it right, okay. This begins to sound like it ties into something you mentioned on the last panel, which is, you need to have an inventory. So is the white listing approach the solution there?

MR. GREENFIELD: White listing is a common control that we do see, but what the inventory is, again going to the simple concept of knowing what you have so that you can secure it. White listing takes it a step further as to say, and then, this is what you're allowed to operate on it, and if, to your point of, if it doesn't fit within that list, there are technical controls that prevent it. A key point I've heard that I

also, when we look at examining large complex institutions, a key control and, I thought Ron did an excellent job with identifying the management, the operational, and technical controls, is, with the management controls, we've been talking a lot about minimization and least privilege and one of the key controls that we focus on, especially in larger complex organization, is the idea of having an architecture and architecture strategy. And the reason why that's important is, some of the vulnerabilities that we've seen is, as technology evolves, many technology environments in large organizations will just be built on top of the existing structure, and to have a defined architecture strategy and program where older software, older, that's not supported, older network components are retired as new are being added on, and that you stick to an environment that you can secure, is very important. And I'll bring back in an earlier panel, the concept, that of an airplane, bringing all the parts together and pulling it all together. Well, you've got

your airplane, but you wouldn't take a DC-9, a propeller driven airplane, and over time, well let's add some jet engines to it. Let's clamp a radar on it. We don't have a first class section. Let's expand the fuselage. Over time, that airplane's not going to fly very well. The same with network environments and securing them, is, if you're continuing to build on older software, older network components that are no longer supported, you open up the organization to vulnerabilities.

MR. ROSS: Bob, can I -- there's one other one that I think, we missed it, is strong identity management and authentication. That's something, like two-factor authentication again, these are technologies that are proven to stop significant attack vectors and again, the passwords and all of the nightmare of managing all those. The two factor is a clear solution that really, really helps reduce lots and lots of vulnerabilities that end up in these successful cyber-attacks.

MR. GREENFIELD: I'll weigh in a bit. So when I think about this discussion, I think a lot about what I know at least the Commission already asks about and so I think that's pretty settled. But more -- so I think more about well, what are you not asking about, from going through evaluations. So if I omit anything, it's not because it's not important. It's because you're already covering it. So what I don't see a lot of are really controls that are key against advanced threats, and you know, when I try to think about well, what are the best controls against that, whatever I say will certainly damn me by omission, but security awareness is absolutely huge, and social engineering training and social engineering testing. I'm not asking for more examination in those areas, but those are key controls, because at the end of the day, the human is always going to be the weakest link. Along the same lines, there's a lot of, you know, Mark mentioned, there's more malicious software than benign out there, so the anti-malware controls -- you know,

we don't get a lot of questions about how those actually operate. The more kind of generic approach is, principle of least privilege is certainly key. When we pressure test, we find a lot of the malware has privilege escalation routines built into it. So in other words, we're moving local administrator privileges from every machine, really only defending us against the software that counted on having local administrative privileges, but apparently that wasn't all the software out there -- all the malware more importantly that was out there. So a lot of that one-size-fits-all evaluation and examination can be very taxing, laborious in trying to be exhaustive, but totally missed the mark when it comes to, you know, what an adversary would actually come through on.

On the asset management side of it, you know I also think, I'm sure that you would really like to know, on the inside, during an examination, what do we leave thinking, that was a waste of time, that doesn't really keep me up at

night, and really focusing on that. And then the other side which I just kind of went through is, what did they not ask about that I really wish they would for my personal interest and critical infrastructure. Asset management is a noble goal and it's part of every program and not just information security, but operationally generally. But it is a very challenging bar, to know about every single asset. In reality, the way that infrastructure defends itself, is to carve out entire segments and say, well, this whole segment isn't even going to have, you know, be able to knock on the door. This isn't going to have any access, and so prioritizing asset inventory in that segment is going to go way down the list, all right. So when we get questions about well, what's on the wireless network, and our answer is, we assume it's bad, so then wireless network can't touch to anything in production, we're ready to move on. And that's why that's not a focus area.

So I guess the analog would be, you would keep the DC-9 around for the parachuters,

you know, for the guy who's running a weekend shop out of there. There's a time and a place for things. You're right. You wouldn't bolt it onto the commercial flights and throw an MD-88 engine on it. But you know you got to be really careful, and you got to look at the actual environment before you bring any of these controls through, and a lot of them that try to be exhaustive, end up being a disproportionate usage of examination time.

MR. WASSERMAN: How do you -- I'm trying to tie together that, because you're right. You want to, obviously -- doing things on a risk basis has some very real and important advantages. On the other hand, we learned back in panel one that air gap is a myth, and so when you say, oh, this can't touch anything, are you really sure?

MR. GREENFIELD: Yeah, so two other pieces on that. One thing that we found is very valuable internally is starting with a threat objective assessment, so, you just answer the questions, well what are the bad guys really

after, or what could they be after? And starting there, and I think that examinations should follow that same path. So like identity management is something that we hear about a lot in the space, especially from vendors, and there are a lot of companies that are represented in this room that have north of 100,000 employees and that is a huge challenge. But if you look at the environment that you're in and you find that this company has 3,000 employees, it's probably not near as high of a challenge, and therefore you shouldn't be looking for the same controls and for the same solutions to that problem. So I think just like we internally can start with the threat objectives and then work backwards to what controls are important, that examiner should do the same thing. And even if you're working in the same industry, different institutions are going to have different scale, size and business models. And if you start with those threat objectives, then you'll get down to what are the controls that you really should be asking about there. And it's going to be

different from the last assessment.

MR. CLANCY: I would just add and this goes back to some of the earlier comments. I think you also have to understand the two things that lead to the conditions where these exposures exist. So as it relates to Tom's comment about flat networks, there is a financial gravitational pull to a flat network because they are much less expensive to operate on a day to day basis. You don't have to do as many changes; you don't have to do as much testing, all those kinds of things. So there's huge advantages to having a flat network, which comes with the risks. And so the challenge is how do you sort of optimize the risk management side with the partitioning and the segmentation with the cost efficiency because we all have to operate the stuff and still figure out how to pay the bills, right? So that's sort of one tug.

The second, to go to Kevin's comment, is there's a human incentive structure built into these processes that we're trying to counter.

People will do anything they need to do to get access to the thing they need to go use. And if they don't need it anymore, they have no incentive to do anything, right? And so you have to intersect those two sort of just facts of life in terms of how companies operate. And so part of the control regimen and why controls become as important, is they're trying to address the gravitational pull of those two realities, right? That you want to keep things that are as cost efficient as possible and that are simple because they are easier to manage, you screw it up less on those kind of things. And then human nature, you know -- I definitely figure out how I get access for this person to go do this thing I need him to do, but if they only need to do it once, what incentive do I have to say, oh get rid of it? And the answer is usually none, and why the backup checks of reviewing and reasserting and you know, when they change jobs, removing access and those kinds of things are so important, is because they counter that very human nature of, if I have to do

it, I'll figure it out. And then if I don't have to do something, I won't do it.

MR. WASSERMAN: So let me -- let us move a bit from a discussion of the key controls themselves, to the issue that we, I think, were looking at from a regulatory perspective, which is the testing. And Jerry, I think I was going to perhaps start with you. What does key controls testing accomplish? How does that mitigate risks?

MR. PERULLO: Yeah, I think Mark really spoke to it pretty well earlier, that it is a level of maturity. So first you get the controls in and then you start testing them. And I have been getting questions about that already. So okay, your intrusion detection systems are in, that's great. How can you show that they're actually operating? And our first response to that, and I'm speaking more generically, is to use existing testing we already have, and then make sure that our controls pick it up. So in other words, not to avert testing just for the sake of testing these six controls, but rather say, we

know we already have this testing going on, let's see if it was reflected in the controls.

The next step would be actually having periodic ticketed and documented tests of specific controls. And you know, it really is -- it is important because a lot of times, you don't know that something is still operating. You know it's very easy for things to more or less just get turned off for different reasons. In general, a lot of security controls are what we would call passive, meaning if they go down, business doesn't stop for very good reason. But they also don't get the attention. Back to aligning motivations as Mark was just mentioning a minute ago, if the core system goes down, people are going to raise their hands right away. If the intrusion detection system goes down, they probably aren't. So I think it is very important because a lot of times when you have an incident, then the ultimate answer is, oh, well, that thing stopped working six months ago and nobody knew. So some analogs in the physical space, I mean you really have to

do that to have any control be effective. I can use an anecdote in house. One of the ones that we have is, for conference room phones to not be able to allow internal calling from the outside, because when you have that, people inevitably put auto-answer on it and then the world can listen in on your conference calls. So you can turn that all off, but you actually have to have somebody walk around every conference room once a quarter and try it out, in order for it to know that it's actually real. So yeah, you do need that type of thing with any control. You need some kind of periodic testing. But it's just a matter of maturity and it's definitely far beyond actually getting the control in.

MR. WASSERMAN: So when you're saying it's a level of maturity thing, I mean, when we're dealing with the sort of infrastructures that we're regulating, is that level of maturity reasonable to expect?

MR. PERULLO: Not comprehensively. I think we all need to get there. So I don't think

it's totally off the plate. But I think that you know, just knowing the industry, that first the control comes. There's a lot of technology, like behavioral insider threat detection, things like that that are so new and the reason I'm stressing the maturity is that first it has to be pressure tested and vetted. Then it has to be adopted and installed. And it's not until after that the controls testing gets put into place. So for things that have been around awhile, absolutely. But it needs to come after the expectation of the control to begin with.

MR. TAYLOR: Let me follow up on that for a second. Ron mentioned in 800-53, there are 800 and some controls listed and a critical infrastructure is going to have large numbers of key controls both for automated processes and for manual processes. I think Jerry, you're suggesting, you need a way to separate out, let's call them the significant key controls, or the most important key controls. How do you do that and what tells you what they are? How does that

relate to risk analysis or intelligence for instance?

MR. PERULLO: Yeah, that's absolutely right and it goes back to what I was calling a threat objective assessment. And I think that by doing that with a broader audience and saying, so for example, pretty much data theft has been dominating the headlines. We're probably here today because of credit card theft, even though it has nothing to do with anybody in the room. But as a result, in certainly the vendor space and really anybody involved in cybersecurity, has a bent towards that, towards data exfiltration. And as a result, controls that are stressed are often about data leakage prevention, or encryption at rest and that sort of thing. And I think practically speaking, you need to walk in and do that threat objective and say, is data exfiltration for this entity a top concern? It's always somewhat of a concern, but is it a top concern? And if not, then maybe those aren't the controls that are key. Is availability a concern?

Probably more often in this room. And therefore, denial of service, which has not been a big part of examinations at all I can say, should definitely be a lot higher up there, right? And so that would be a key control. So start with that assessment of the threat objectives for the entity under review.

MR. WASSERMAN: And I think we were talking earlier, someone mentioned integrity as even the highest of the goals. I mean, I guess you had mentioned some of these things, that there is a maturity issue and I guess what I think I hear you saying is, to expect you to have certain key controls tested in the very near term may be difficult because essentially the science needs some time to develop. I mean is this the sort of thing that we would address through some sort of implementation timeline? How do we, in other words, looking at it as regulators, trying to basically have rules for how critical infrastructures need to be protected, balance -- you know, giving you goals that are achievable but

making sure that they are sufficiently rigorous, that we're addressing the risks, and that you are?

MR. CLANCY: So maybe picking up that theme and tying a tiny bit back to the NIST 800-53 framework, you know the way we look at it is you have a maturity of control. So at some point you start as new and you have nothing. Then you start building things, and over time, those things mature, and to go back to sort of the access review side, in an immature organization, the control finding is, you don't review access when people change jobs. In a somewhat mature organization, you don't review changes timely, or you don't get to all the systems or whatever. And in the most mature organizations, what you're discussing is, Fred changed jobs and he still has his access, but he doesn't need it in his new job. What's going on, right? And the level and depth of that conversation evolves as your maturity of technical and operational capability evolve, and your understanding of the risk becomes better, right? And so if you take a control framework

that has 860 controls or another model that we use, I think it's 400 something -- each control has a current maturity state and a target maturity state. And every organization is trying to mature the controls that are most important, but that mix of which 860 apply to a DTCC versus ICE versus a Morgan Stanley, are going to be different, because what we do is different. So as a swap data repository, all right, we have a different impact if the swap data repository is unavailable than if the trading system's unavailable and it's the only market venue where that trading can occur, and so availability may be different. In our case, we're custodian of records of what happened and so the integrity of that data is very important. So we would try for our control footprint and where we have our most high maturity to those things that are more direct to the business we're in and it's not to say we wouldn't do anything on the other 860 controls, but they may not be targeted for peak maturity. And what we were trying to do, and this is using a different framework than this, but

just trying to get the aggregate level of our controls to a certain maturity objective, which means some are five out of five and some are two out of five, but in aggregate, you sum them all up. Our overall picture was what our target was of being four out of five, on this particular scale and I'm using a generic model. But that type of target, and so the dialog that needs to occur is based on our assessment of risk of functions that we have, which of these subsets of controls, either by category or by specific ones that we prioritize for maturity, and are we either there or progressing to our target state? That's the kind of thing we look at. And so when we do our testing, we're trying to figure out two things -- one, are we there yet? And more importantly, did we regress back to some lower state of maturity because the control broke down because it decayed over time or whatever happened. So we're trying to assess those things -- the where should we be, the where are we, and how do we get to where we want to go.

MR. WASSERMAN: Ryan, can I turn to you?

MR. LIBEL: Absolutely. I just want to first, I think, echo what we've heard a lot of here. I think some of the key concepts that I'm hoping that people are hearing are things about depths of controls. So we talk a lot about what are the key controls. I think that's influenced a lot by, what are the risks the organization faces. I think, what are the key things that they are trying to prevent from happening or limit by going to the risk dialog of, what are they trying to reduce the impact, if it is likely it's going to happen, that's going to involve a depth of controls. I think key controls we've historically, I think if you went back to things like SOX-IT, you would talk a lot about change and config, which would be a lot of the most fundamental blocking and tackling. I think that's gotten onto some of the other things now when you come more purely into what the information security world has become more worried about, rightly so, on the things we're talking about all

morning -- your vulnerability and patch, your pen testing. I think we didn't talk much here about some other fundamental things and that things will happen, so your key controls around how do you detect, how do you limit, how do you respond? I think all those are in the mix of what we would have to come to, I think, a good joint understanding of what do we mean by key controls, to be helpful to each other, but I think that for each company, what those key controls are, will come back to again, what are the risks that you feel that you are facing and which ones are the most key for you to operate most effectively? It's a complicated space.

MR. WASSERMAN: Kevin, from a regulatory perspective, how do you see this in terms of, specifically, when we're talking about testing? What do you see as the appropriate scope?

MR. GREENFIELD: We focus on, how does any institution map out, again, for any critical operational process, what are the key control points? What are those actual controls? And how

is the testing sufficient to gain a level of assurance that those controls are operating effectively? So one of the things we try to emphasize is the maturity of that control and the testing process is very important. Maturity of testing can range from, take a sample of five, let's use user access levels, test them to make sure the users have that level of access, but it really doesn't give you a whole level of assurance. We look to financial institutions to better identify in their testing, how to be more intelligent about their testing, using automated tools as well as focusing on what are the areas of most likely control gaps, or to highlight control gaps. So for example, in the user access example, we would say, don't test 525,100 users. Identify the users who have changed their jobs over the last six months and go and test those. Those are more likely to be the ones where it will surface if that review process is working or not. So we look for the -- absolutely expect there to be testing in place and expect that testing to be of

a sufficient level to gain that level of assurance.

MR. WASSERMAN: So one of the things you were saying is testing sufficient to gain a level of assurance that the controls are operating properly. Would you also be looking at, that the controls are sufficiently comprehensive?

MR. GREENFIELD: Absolutely. And again, something that we look for a lot, what we commonly see a lot of financial institutions doing is mapping operational processes and identifying those key control points, single points of failure in the process and highlighting those as critical controls that need to one, be included in the scope, and then tested on a regular basis.

MR. WASSERMAN: And then when you're doing some of these other types of testing, like we were discussing before in terms of penetration and vulnerability, would the results of those tests feed into essentially the key controls you are looking for?

MR. GREENFIELD: Essentially, and this

is where focusing on that remediation plan that we had discussed in an earlier panel but as part of remediation, identifying what was the cause of that vulnerability, or that gap that was identified in the penetration testing and looking towards, it was introduced by new software and you couldn't have identified it before implementation, that's one issue. But if it's a failure of proper training of employees, staff being able to circumvent standard change control processes, identifying that root cause and then mapping back to what was the control that was in place that should have prevented that gap being in place, with the understanding that there will be times when no, that could not have been reasonably anticipated.

MR. WASSERMAN: Ron?

MR. ROSS: When you talk about testing of key controls in the scope or the value, really you're making an assumption that there's been a set of key controls to find. I mean that's -- if you go out to the sectors and you're going to be

asking them to do a set of tests, that's going to kind of work in the back door, saying, this is what we're going to look at, so the implication is that they have applied the control that you're going to be testing, just by making that statement. The value of testing in general, is really tied to the controls per se, because if you pick the wrong set of controls, all the testing in the world is going to be throwing money down a black hole. And so it really is important to start -- the risk management framework starts out with an assumption of, what is the mission and the business that we're trying to achieve. And the controls that are selected are based upon that mission protection. So we select whatever controls, management, operational, technical, are necessary to protect the mission and the business. Those controls are then implemented and then, after that process is complete, we go to the assessment. We call it assessment. There are lots of different things in the assessment process you can test. You can evaluate, you can audit,

you can inspect different names. But the basic idea, the words that you said earlier, to see if the controls are implemented correctly, operating as intended, in producing the desired effect, to make sure that your security policy is enforced, and that the mission has a high degree of probability of success. So focusing on the threat space, in some sense is like chasing your tail, because the threat space is out there. We know what the capabilities are -- the adversaries. Anybody with a laptop computer, and a couple hundred thousand dollars or maybe a million if they can rustle up the money -- they can go out and buy these very sophisticated attack tools today. So how we build our infrastructure, and that gets back to Tom's original point about network segmentation, that assumes that we look at all of our assets, and we can figure out, hey, what stuff goes into my safe deposit box and what stuff do I leave in my house? And there is an air gap. The air gap is not dead. If the air gap were dead than network segmentation would be

meaningless. We segment because we want to prevent that escalation when the adversary comes in the front door, getting through the whole house. And that can only happen through good architecture and engineering. Again, those are part of the controls. So the point I think I'd like to make is that whatever you call them, whether they're key controls or whatever the name is, they have to be comprehensive. Because if you spend all of your time on access controls, and by the way, we worry about confidentiality, integrity and availability -- if I compromise my passwords or my credentials, that's a non-disclosure issue but the adversary then gets in, changes something in the system which causes it to crash and you lose the availability. So these are all interrelated objectives. And the controls are built to support all of those. And so you know, some people say, hey, access controls or encryption are the most important things. Well what happens when that 10 percent of the cyber-attacks that we know are going to get

through, get through, and bring down your system? Where is your -- is your contingency plan a critical or a key control? I think it is, because your system is going to be breached at some point. Everybody's is going to be breached. And that's an important part, of what do you do to maintain resilience in a world where you depend upon the technology, but yet, it's very vulnerable because we're susceptible to cyber-attacks because of how we built this whole infrastructure to begin with. So those are just some considerations that I think are important.

MR. WASSERMAN: So let me grab on to something that you said early on, which is that you first have to identify the mission of the institution. And I'm wondering if one of the, at least from where I'm sitting, one of the key controls, or whether it's a set of key controls or a type of key control, is looking at whether the set of key controls is sufficient to essentially protect against threats to that mission.

MR. ROSS: It's even higher level than

that. I liked Mark's example because he articulated exactly how the risk management framework was built. He talks about what their mission is and then he says they go through and they pick, they select a set of controls which are targeted to their mission. And he assumes that every organization is going to have different levels of maturity. So the way he described it is that you know, some things you do in a very mature organization. Other ones you don't do because the organization is just starting out. They don't have that level of institutional security that's built in through all the organizational processes. So I think if you're looking at an organization, do they have the maturity to start with the mission, and are they going through a thoughtful process to select their controls, and see what they end up with, or are they just throwing stuff out there and seeing what happens? That's a very different way of looking at it.

MR. WASSERMAN: So let me move that on to Mark, because you had mentioned how you select

a set of controls that's relevant to the mission. Is one additional control on that, sort of looking back at that selection, to determine whether it is appropriate and sufficiently comprehensive?

MR. CLANCY: So the short answer is yes, and there's a much longer answer that goes with it, but yes, we are informed as we make our decision about what controls we think are important based on our past experience and our understanding of the threat environment that exists, right, which I describe as the projection of the future experience we may have. And the challenge in designing this is controls; they eventually run out of room, right? So we have a control to do access reviews and we eliminate unintended access but it doesn't help us if somebody abuses the access they're supposed to have. So we also have to recognize for every control, there's sort of a maximum amount of effectiveness that individually can do and whether it's through accident or malicious act, there will be things that over top the capability of that

control. And that's where some of the adjacent controls then can help you. So if I have a highly segmented network, somebody's authorized to do something and they want to pull data out but they can't yank it out, well then it's not going to happen, right, so it's that -- and we call it defense in depth, although I have a lot of challenges with that wording but that concept of it's no single control, there's no silver bullet. There's no single thing that makes everything stop. The trick in this, and the hard part, and one comment that I've made in other forums is, we don't really have a good sharing mechanism to receive information back about what happened and when controls failed at other institutions. And to beat on the airplane analogy, which seems to be in addition to bears, the theme of the day, is when there is a crash, we don't get the lessons learned back from the crash unless it was our plane. And the only reason we get the lessons learned from our plane is because we did the investigation and we figured out that the wings

were frozen and that's why it crashed. And so I think one thing that we need to talk about in this policy discussion and as a regulatory discussion is, how do we get those aggregated anonymized lessons learned and evidence that says when the access review control failed, this led to this type of event. And this happened twice this year or 17 times this year, with this set of consequences. Because that can better inform other institutions as to, these are controls you should go look at because there have been incidents in your neighborhood and the analogy I use is, when all your neighbors decide to get alarm systems because somebody's house has been robbed and then they buy a new TV and the house gets robbed again, their response to that stimulus changes their behavior. And then they put in a lighting system and they do other things. They don't rely just on the door lock anymore. And it's that type of piece, and so we can talk about frameworks and those pieces and they're very helpful, but I think you have to also inject the

real world data of what's happening and how that changes over time. And I think for me, what we're trying to do is, we have a current set of controls and we set a target of where we want to go, and yes, that target has things maturing from our current state, but the reason for that is, we project forward what we assume is going to happen to us, and how this threat landscape is escalating, I guess is the best way to put it. And now we have to expect new things showing up at our doorstep. And what we're going to do to position ourselves and admittedly, a few things we're catching up on too, because we lost focus, lost attention, or didn't prioritize something, it turned out to be important because maybe Jerry or Jerry's firm had a problem with this and we say oh wow, we better go jump on that. And for us for example, denial service capabilities for one, because we made an assumption that since we have a private network, everything important goes over a private network, except that wasn't actually the behavior of our customers. And so when you had a

threat of those types of attacks intersected with that decay over time, not because we didn't have a private network, but because usage migrated to the public network and we didn't really notice it. We had to re-pivot what our capabilities had to be to protect that public network, which we actually thought was less important. It turned out to be more important than we originally assessed and so we had to pivot. And so it's sort of a constant tuning mechanism and you take the experience of what happens to yourself or hopefully to others, and you learn from it and you adjust the maturity targets and then the capabilities of what you put your resources towards.

MR. WASSERMAN: So it sounds like some of what you're saying ties back to what Bill Nelson was talking about earlier for FS-ISAC, that essentially we want to promote that kind of anonymized sharing of results.

MR. CLANCY: Yeah, the FS-ISAC is very much about sharing the technical bits of what attackers use to do their attack. This is a

little bit different, as what are the circumstances that led to the attack being successful. There's a slightly different pivot than where we emphasize that sweet spot is today.

MR. BRADY: That's root cause analysis right?

MR. CLANCY: Yes.

MR. BRADY: And what control failure caused the incident to occur, but I think you wanted to --

MR. MILLAR: Well I wanted to jump in and say, if anybody remembered, that's exactly how I opened up, trying to explain US-CERT's contribution to this. And where we're trying to head with the FS-ISACs and Information Sharing Analysis organizations, hopefully is encouraging more of that type of sharing, because over the past, say three years I think, we've gotten actually through our Cyber Security Information Sharing and Collaboration Program -- did not use the acronym -- that we've gotten much better at sharing those, as Mark put it, the technical bits

and or bytes and or kilos thereof, describing how you can detect a threat that may have hit one of your partners' competitors, somebody in the vertical or somebody in a completely different industry but who shares a risk factor with you, or is a similarly, perhaps, appealing target to certain adversaries. We've gotten much, much better at that, but the next challenge is to try and figure out how can we best anonymize what has happened to certain institutions and organizations that we've worked with and bring that lesson back in a way that's actually digestible or as we say, achievable. Maybe practical is actually the plainest word, for other institutions of varying size, because what is a great control for a Fortune 50, is not going to be the same for small to medium businesses. And we see that from our government perspective when we look at commissions versus cabinet-level departments, right? So we have our quintiles, as we put them, and we have to line those up as well. You have some with like DHS, 280,000 employees and then you have some we

call small and micro-agencies that we also have to defend. And they have about 50 in some cases. So those are all challenges that we're familiar with, already in the public sector and now trying to figure out how do we apply, because the same things apply. We don't want to spill any -- we don't want to share anybody's dirty laundry after we've done an on-site incident response engagement. That doesn't do anybody any good if we're just calling people out for failing, which is what it sounds like. What we want to do is say, if this is what broke down and this is what we've recommended then that actually appears to have fixed it or minimize that risk going forward, this is what we recommend everybody else do. And that's tying actual incident response, especially applied to sort of the severity of impact that happened to the organization, applying what we learned from incident response and putting that forward towards recommendations of which controls should be focused on. And that's very much an evolving process and I think we're starting to

build out a lot of the trust infrastructure between public, private, and all the parties therein. Also with privacy and civil liberties organizations that there's not going to be something resembling collusion going on while we do this, that it will be above board, and that we're going to do this in a responsible fashion but that's also scientific and rigorous, that it's not just anecdotes, which is kind of where we're at today. We did, you know two dozen, somewhere between two dozen and 50 incident response engagements. Most of them looked kind of like this, and then we move that forward and push it out there. But is that really scientific? Does it help Ron write a better 800-53?

MR. WASSERMAN: So we do, of course, already require risk analysis. We do, on the other hand, sometimes see either controls that are in place but haven't been tested, or maybe that aren't doing the things that they're supposed to. And I guess the question is, is that simply a matter of maturity? Are there specific things

that can be put in place to more reliably address these issues?

MR. GREENFIELD: Now I think with that important concept that's applied in banking organizations, when looking at the adequacy of the control environment, is the three levels of defense model and looking at the business line itself. Myself as a business line owner, I own the function. It's incumbent upon me to make sure the controls are comprehensive and effective for the process I manage. But to an earlier comment, if the controls I'm focused on are not the right controls, it does not matter how effective they are if they're not addressing the correct risks, having that next level of defense being an independent risk management function, which is very familiar with the process but independent of my reporting line, that is looking and providing a credible challenge of, do I have the right controls in place? Am I managing the correct risks? And providing that level of challenge, that's something that, through some of the market

disruptions, was identified as a common theme of not having that credible challenge outside of the business line to some of the processes, practices, and controls. And then the third level of defense being the independent audit function that is completely independent and reports to a board of directors and tests and confirms that the controls are adequate, but in that manner, having those three levels allows that check and balance to ensure not only are the controls in place, but are they the right controls? And are they being tested on a sufficient frequency and sufficient depth?

MR. WASSERMAN: So I'm going to follow up on a couple of the things you've said, but one of them you mentioned is terms of independence of the testing. And Brian, I was going to ask you, looking at it from the perspective of an infrastructure, are there some types of controls, key controls testing that are best performed in-house? Are there others that may be best performed by third parties?

MR. LIBEL: Yeah.

MR. WASSERMAN: How do you guys look at that?

MR. LIBEL: Well I think Kevin can see my notes here, because I was going to say the same thing. I think one of the very key concepts to think about, there is those lines of defense when it comes to the controls. And I think just to echo what Kevin said and to play it back to get into your question. That first line of the infrastructure, so the business line, are probably technology operations and development to (inaudible), responsible for operating and adhering to those controls, week in and week out. Some independent but knowledgeable set of expertise that's in there doing that risk management and maybe compliance-type function of, are these things really being followed through on? Do they appear to be effective? Probably focusing that testing then on where knowledgeable areas of risk would be from being on the inside, but again, reporting somewhere different in that

organization, so for example, not having, not checking his SDLC -- I'm sorry -- Software Development Life Cycle being followed, solely by someone that's reporting to a development manager. Is there some independent function inside of your technology group that's overseeing that?

And then that final layer of defense is that pure independence, probably standards based more likely, etcetera, maybe internal audit and likely some external expertise that is doing that pure assurance independently from the outside. To the kinds of testing that I think lend itself to each side, and in a general way, I think it's a hybrid model, in all honesty, in that, things we've spoken about already here today, lend themselves to an external party. Penetration testing clearly takes a great deal of advantage from subject matter expertise, skill sets, and also independence of having an external party do that, and in all honesty, see the world perhaps from a viewpoint that you do not, as I'm coming in from the outside. Similarly, when perhaps testing

very standardized things, that you want an independent assessment and some kind of a certification on a routine basis again, either perhaps your internal audit group, or an external party. Some things that definitely lend themselves to internal testing or using existing staff, would be things where you do need a great deal or expertise perhaps in the systems themselves. They're deep. They're complex. There's a lot to understand. You're probably going to need to have some staff that understands that involved in some cases just to really know what they're looking at. Or in other cases where the level of access that you would need in order to really see and understand things you wouldn't want someone from the outside to have.

One thing I would offer that's kind of again, the hybrid approach between some of these things that we've said is, we've talked a bit about penetration testing and we focused entirely for the most part on external parties doing that. There are very good practices that are about also

complementing that by some form of an internal red team or someone -- someone who knows, who has those skills, or is a collective of people who has those skills, but know your network enough to press in the right place. These are some of the things that would be thought of.

MR. WASSERMAN: So let's talk just for a second on that, about key controls testing in particular. Is there perhaps a certain level of periodic key controls testing which might benefit from having an external viewpoint?

MR. LIBEL: So we kind of do all three, and I think I mentioned that a little earlier. We'll do it again. So things that involve measurement, like every day we check who has access to what or what vulnerabilities are in the system -- those are always done internally. Those things that we do episodically that are truly testing, we do a mix. Sometimes we do them; sometimes we hire somebody to do them. A lot of the times we do them ourselves. But those independent assessments of what's really going on,

those work with outside parties. And the reason for that is, you want to do two things. You want to make sure you have coverage, so you identify blind spots. When I found the internal processes and where internal control testing tends to break down, is you get myopic and like, well, this is what we always looked at. And you narrow your scope intentionally or accidentally, and the outside party comes and looks at it from a different frame of reference and says, well what about all this stuff over here? And for whatever reason, your process evolves to the point where he missed it. So I am a big advocate of what I call hybrid, where you do both. You do some things internally and you do some things externally, and the intersection of the two get you better coverage than either one of them would do on their own. So I'd caution against saying it should always be external testing of this kind. I actually think the hybrid is the best piece, because no one knows your environment better than you, and nobody knows the -- doesn't know your

environment better than you, so they ask questions you forgot to ask because you included your thinking because this is how you always did it.

MR. WASSERMAN: Kevin and Tom can I get your perspective?

MR. GREENFIELD: I completely agree with that perspective of you need to have the expertise because every institution's unique. The institutions we're talking about are very large and complex. You have to understand how the operations work and it's going to be someone internal to your organization, but to that point, they're only going to look at that through the view of what they're familiar with. They're going to test what they know, where that external third party will come in with a completely different view, and more often than not, views of how other similar institutions have their control structures, and will ask things from a different perspective. And then back to, regardless of internal or external, that independence, making sure that the persons conducting the testing are

independent of the function. The first line of defense should still be doing its own testing, but when you're going to get that level of assurance, you need that level of independence, which could be someone who's part of the organization or a third party.

MR. WASSERMAN: Tom?

MR. MILLAR: It may be a little bit of a darker spin on all of this but everybody's organization is being penetration tested right now by independent external evaluators. And we live with it.

MR. BRADY: For free.

MR. MILLAR: Right, for free, which means they don't work for you.

MR. WASSERMAN: The formula's maybe not as good.

MR. MILLAR: Right. Their independence is perhaps a little extreme.

MR. ORTLIEB: Their information sharing is zero.

MR. MILLAR: Depends. That's not always

the case. There are a couple of the gray hats or white hats out there who are automatically blowing up iPhone apps to make sure that they do correct certificate validation, and they will tweet to your bank about the problem they found. And we talked to that guy and told him not to do that again, because we were sponsoring him at the time. (laughter) But there is good work being done that is independent and for free. But yeah, they have to have a way sort of like to communicate and coordinate in a trusted fashion, with the people that they found a problem in your system with, right? Sometimes we actually get to broker those communications and it gets pretty interesting. But overall, the point I was trying to make was, especially if you're evaluating, you should have a hybrid approach. Obviously we completely agree with

that and try to encourage that everywhere we go. But when evaluating how much to spend perhaps, or how much to invest in periodic independent external penetration testing, you should always consider or remind your decision makers to consider the fact that it's already being done by people who do not have your interests at heart, and you'd rather find out from people under contract than from people under contract to somebody else, perhaps.

MR. ROSS: There's a common theme I think that you might be sharing and it goes back to the notion of complexity, these complex systems. When you talk about, how do you do a test, that talks to whether the governance level of an organization, the senior leadership, is enforcing what Kevin talked about -- a good enterprise architecture where the architecture

itself, the basic constructs, drive you to consolidate, optimize, and standardize the infrastructure that you're building, because if you look at -- there was a defense science report about two years ago, and they asked the question, could the United States military survive a massive cyber-attack and still defend the country? That's a pretty important question. And in that study they described three classes of vulnerabilities. The first ones we all deal with all the time, the known vulnerabilities. The big companies, every Tuesday, we call it patch Tuesday, because they announce the latest patches. Those are known vulnerabilities that are patched. The second level were the unknown vulnerabilities that we all have. Those are the source of zero day exploits. And the reason why zero days are exploiting now is because --

MR. WASSERMAN: And a zero day exploit is?

MR. ROSS: A zero day exploit is when somebody, a threat, exploits a vulnerability that

they know you have, but you don't know you have. And once they exploit it, now you know you have it. And that goes to the known vulnerabilities stack. The third class --

MR. ORTLIEB: And that's if their exploitation occurs in such a way that it's made aware to you.

MR. ROSS: Yes.

MR. ORTLIEB: You're made aware of it.

MR. ROSS: Yes. And when it's detected and all of that, of course. And then the third level is the vulnerabilities that are actually created within your infrastructure, your organization, by the advanced persistent threat. They penetrate. They establish a presence. Now if you look at it, the two-thirds of those vulnerabilities are totally off our radar. That's why all the talk about chasing vulnerabilities and doing all the vulnerability scanning and testing, and every time you think you've closed down the last vulnerability, I'll find ten more. Why? Because the complexity of the systems we're

building. And that is a cultural issue. That's an institutional issue that we're going to have to get our arms around, and all the testing in the world is not going to solve that problem. There's a glass ceiling on testing. It doesn't really fix the basic architecture. It doesn't really change the complexity level. And therefore, when you look at an operating system of 50 million lines of code, and there are a certain percentage of weaknesses and deficiencies in that code, this is why we have literally thousands of security vulnerabilities in the software and the things that we're deploying. And nobody, even the best among us, can deal with that complexity and chase those things down one by one. The only way you solve it is to go back to the things that Kevin talked about -- good architecture, good engineering, and mandate that from the top. And what kind of a test can you do to make sure that the organization is enforcing that? That's an important question I think.

MR. CLANCY: So just maybe to expand

this zero day definition a little bit; where it came from in time is, you had vendors announcing here's a patch. And there were a number of days from when they announced the patch to when bad guys were exploiting it. And I think we heard earlier, that's not days anymore, it's hours -- you know, 10, 12 something hours, from when a patch is released, these people are reverse engineering and figure out how to attack it. The zero day was when somebody disclosed the presence of vulnerability publicly, but there was no fix for it. It's been sort of morphed to also include those vulnerabilities that an attacker never disclosed and exploited and then, because they're exploited, now you tell people, hey, this thing's broken, and that causes disclosure. So there are a few other pieces in there.

The other thing on the glass ceiling on testing -- yes, to a point I would agree. And the point where I disagree is that there are different categories of adversaries, and the most sophisticated adversaries, if you test and remove

100 percent of your known vulnerabilities, they're still going to come over that wall. They're going to go above the glass ceiling or whatever you call it, but there are a large number of adversaries that if you close all the known holes, they have to move on. And so again, I sort of mentioned this in the first panel, it's sort of that difference of who are you worried about? If you're only worried about the most advanced attackers, then testing only gets you so far, and what you're really trying to do is increase their work cycle and their energy and expense to attack you, but you're not going to necessarily stop them, and so resiliency response becomes extra important there, as opposed to sort of the commodity threats as we started to call them, where if you get that high level of hygiene -- somebody mentioned that in an earlier panel, where there are very few of the known holes, either configuration platform or architectural vulnerabilities. Those attackers are not going to be very productive, and the work effort required

for them to breach your environment exceeds their capacity to supply resources. And so they go away. And so you've got to actually do both. The question is, how do you tell when you've tested and you got it to enough and you've taken those people out of play and now you just need to focus on the detection and response and resiliency for those more advanced attackers. And that's not an easy thing to determine.

MR. WASSERMAN: So I'm going to spend a few minutes now on something that is of very big concern to us. As I mentioned earlier, one of our responsibilities as a regulator in terms of promulgating regulations, is to consider issues of costs and benefits. And I think we've talked a lot about the benefits of key controls testing and the importance. But one of the things that we're supposed to do is, to the extent practicable, and the practicability may be very relevant here, we need to estimate costs, and so I'm hoping, and I think I may start with you Mark, having some experience in this area, how could we go about

estimating the types of the costs that would be involved in a properly scoped program of key controls testing?

MR. CLANCY: Yeah, and so there's no answer to this question, but I'll give you the parameters of how you drive to an answer. If you look at a single component like application vulnerability testing -- it's a function of how many applications do you have, and I price it -- it's like buying a car. And unfortunately most of the time in financial infrastructures, we're in the luxury car market for costs. So these are our expensive automobiles. We're buying, in a company like ours, we have several hundred applications. And so the frequency of testing -- so if we're testing every app twice a year, which we do for a subset of our apps, it's like buying 200 cars a year.

MR. WASSERMAN: Right.

MR. CLANCY: Some are Chevy's and some are Ferrari's but you got that kind of range. And so the gist of that -- one testing regimen can add

up quickly, and that's why also, the hybrid piece matters because, and this is where the analogy doesn't work -- it's cheaper for me to use my own resources than always go outside, but I want to have the mix of that expertise. So that's one piece. The way I look at it is, if I look at my team, roughly a third of my resources spend their time doing control testing. And so whatever my budget is, 33 percent of that, that's about what we spend on control testing. If Jerry and I and Jerry, we've been surveying other financial firms and the amount of money spent on this topic varies greatly. We haven't found the perfect measure of what is a reasonable amount to spend and what are the leading companies doing versus the trailing companies, but spending in this range is roughly between one and five percent of IT spending. And it's hard to translate. People who spend five may spend less on IT, so it's kind of hard to get a comparable metric and maybe Jerry, you want to talk about some of the work we're doing to get those benchmarks there. But this is a significant

part of the op ex of a security function, is control testing, be it pen testing, vulnerability testing, control testing, whatever it might be. It's a huge part of the run rate of a security org.

MR. MCGONAGLE: Just before you go on, to Mark, can you go back to when you talked about the one to five percent? What costs associated, or is it just for testing or is costs all that --

MR. CLANCY: No that's for INFO SEC broadly. As a percentage of IT, it seems to be between one and five. There are lots of factors, and this is for financial market infrastructures and not so much retail institutions, mainly because those organizations don't have a lot -- we tend to be smaller human scale than say a large retail bank. But there were some people reported as high as 20 and some people said a half a percent. The main issue is there's no standard of accounting of, well these are the nine things that I include and here's the 22 things that I include into that spend. So for example, when I do the

math, we exclude patching of systems. That's done as an IT function. We exclude pushing out of fire wall rules. That's an IT function, right? Whereas somebody else may measure their environment and say well this is part of the security function, and so it's very hard to come up with good spending guides. I know Jerry, I know you've been doing a lot of work on this.

MR. PERULLO: Yeah, I have. We've -- the challenges that Mark mentioned are very real. So not only do we have different definitions of what information security spending is but we all have hugely different definitions of what IT is, no less IT spending. So we -- is software development in IT or not? I mean that's a huge chunk of a lot of companies and it's not always in IT. There's nothing in GAAP that says IT. There's nothing at all in financial statements. So one of the things that, and this is, you know, we're still testing this out to try to get more meaningful metrics, is to go against the entire organization's operational expenditures, because

that is something that's published and standardized, at least for public companies it's published, but it's standardized everywhere, and when we've looked at that, and we've kind of beta tested this within the CHEFS groups, at the Clearing House and Exchange Framework for our forum, and it's generally within the one to three percent of an entire company op ex is spent on what we'll call information security op ex. And it's tough to -- and then we have to have a very strict definition of what information security is, as Mark mentioned. So in our organization, we do run fire walls in the group and that's a huge piece of it. When you go to depository institutions and you have fraud, is that included or not, and that's generally a very big spend. Is identity management in or out? So I wouldn't put a lot of stock in any of the metrics unless you know exactly how the numerator and the denominator are both defined, and that everybody agrees on it.

MR. WASSERMAN: And just to be clear, when I hear op ex, I assume that means operational

expense?

MR. PERULLO: Yes. Versus capital expenditures. Well, and then that's why, since development is often a capital expenditure, that's why Mark was alluding to, or you were at least alluding to the fact that sometimes dev is inside or outside of IT, and there is a lot of operational expenditure associated with software development as well. But is it even in IT, no less, and then IT has op ex or cap ex and everything else. What is IT? It's not a standard thing.

MR. WASSERMAN: What I'm hearing is, different institutions are going to measure these metrics very differently, and therefore, it's going to be very difficult to get some sort of standardized estimates.

MR. PERULLO: So we'll try but, so to get back to your general question about how expensive this is, my quick answer is that it's very expensive. So controls testing is expensive. Mark pointed towards application security which is

a big area for all of us. It's a relatively new area I'd say. I got to give Mark a lot of credit. I know DTCC was doing a lot more in that space than most others for many years, but in any event, application security, the lifecycle of it involves, at least in our institution, five different phases of testing for any given one of the hundreds of apps that have been mentioned, so there is static code analysis, dynamic code analysis, vulnerability assessment, penetration testing, just on the app level, and then design reviews which are very iterative and a human going through the architecture of something. That's a lot of work. Not much of that can be outsourced. The pen testing can. And that's a lot of hours and that's a lot of time. And that's just within application development. So, and if I think about other controls testing that we do, the one that comes to mind for me is account recertifications. That takes a huge amount of time, and going through any sensitive access and gaining a list of the people who are authorized, and is that still

accurate as it was a quarter ago? That does take a disproportionate amount of operational labor.

MR. WASSERMAN: So what I'm hearing is --

MR. CLANCY: And just on that, and that expense is mostly borne outside of the INFO SEC org, because we have every manager in the company review the access to their staff twice a year, four times a year, whatever it is, because of the risk.

MR. PERULLO: That's true.

MR. CLANCY: And so a lot of those costs are not captured in the operating expense line in the INFO SEC org, even though the organization is bearing those costs.

MR. WASSERMAN: So what I'm hearing is, at bottom, there are a whole lot of costs being basically incurred right now under the current rule set.

MR. LIBEL: Yes.

MR. PERULLO: Yeah, and I can tell you, just throwing out a little tidbit of info that

might be interesting to think about, if you go back to what I was talking about as a unique key control that's not as widely tested these days, social engineering, so it's a phishing testing of employees. It's a lot more economically feasible to do phishing testing than account recertifications. Which one of those is more important to defending against real threats today? I think it's overwhelmingly weighted in one direction versus the other, towards the phishing testing. And just to qualify that, the reason why I'm downplaying recertifications in this case, is because when an organization recertifies access, the overwhelming majority, say 90 percent of the applications, aren't accessible from the outside anyway. So if there's an old account on there, there's 10 other controls piled on that would have blocked it from being useful anyway. So I think that while it is very expensive, if we can whittle down what those key controls really are and just emphasize on those, it might be a lot more reasonable.

MR. WASSERMAN: One question, we got a really excellent question from the audience, and I want to raise that to the panel as a whole. And they say, there are a lot of different tests going on for different purposes, SOX, financial statement audit, external parties asking for assurance, internal audit -- how can you leverage or I would say harmonize, combine, synthesize -- how can you do that for all of these assessments that are going on, to make sure you've got good coverage of all key controls?

MR. CLANCY: So I'd add one more into that. Ever increasingly our clients are asking these very same questions of us and so in the case of DTCC, we actually created an entire team that deals with all those pieces. So we've a combination of adding resources and reorganizing resources to deal with that. Because we literally have a regulatory exam going on every week. We get about, I think about ten customer inquiries a week, about various controls, and then we have all the external audits and all those other things

happening, in addition to the testing and regimens that we have. So we actually have had to create organizational capacity to just deal with the volume of all these inquiries globally. And the good news about that is, now that we've consolidated, we can get some re-use, where before they were all fresh like they never happened before. I will admit to anybody in my company who's listening, we still got a long way to go there, but that, from a model perspective, sort of building that knowledge base of what's happening, what tests, what things have been asked for, those pieces have been helpful. But it is quite difficult today, because there is, and I don't expect there to be -- there's no real harmonization of what people ask for across those different groups of inquiry, be they regulatory exams or audits or compliance inquiries or our customer inquiries. They all send to the a la carte right now. So we've been trying to figure out how do you create standard frame so you can answer these questions, once consistently, as

opposed to 300 times with little nuances and twists to them.

MR. PERULLO: And if I can add to that, one of the problems is that we're in the same boat. We have a dedicated team and the minute they're done with quarterly recertifications at the end of this month, they can get back on the customer inquiries, so Jerry, you'll have to wait on this. But in any event, there is a huge volume of them, customer inquiries, regulator inquiries, and industry group inquiries and everything in between. One of the problems isn't in the lack of consistency in the questions, but rather in that everyone wants them responded in their bespoke format. So you and I shouldn't point to the CFTC because there's a much more finite universe of regulators so that's more manageable, believe it or not. But one customer may ask, well you know, what is your penetration testing strategy and another one may say, what is the frequency of penetration testing for you? And both of them will have a spreadsheet for it and they'll expect

your team to fill out the spreadsheet. And if we all -- we have got to almost a treaty, if you will, of saying, let us assert our controls in our language and let us start any inquiry with that. Well here's a description of our environment. Before you even give me your questionnaire, look through this, put some time into it and map it, and then if there are any holes, let's talk, certainly. But what we're getting instead is that you see there's a lot of, depends on who has the bigger lever, so in any relationship, any customer vendor, someone has more leverage and their questionnaire always stands. Right now, we just want you to fill out the spreadsheet, that's it. And part of that is because they outsource it three levels deep and the person that's actually asking you has no idea what you even do anyway. But if we could just get used to that idea of well, let me get back a generic response that the customer keeps and reuses, map it, and then just fill in the gaps, maybe there'd be some hope.

MR. WASSERMAN: Ryan?

MR. LIBEL: If I could just say ditto. (laughter) I think that would be the shortest answer. Yeah, we're facing all of the exact same things and I think when it comes to one of the challenges I think woven within that, is the different frameworks that everyone is looking to use, so here on the panel, we have a fellow from NIST, we are dealing with our international regulators in another world, internal audit will see it under another framework, so a lot of the work that we've been trying to do, is to weave it into an overall control framework that we in technology use to mesh that together, to decide which controls we feel are most effective for us that also then boil those down to the common denominators and allow us to essentially risk rank those

and decide where are we spending our time. To Mark and Jerry's comments on the, whether it be regulators or customers, etcetera, and the dynamics of how that goes, in trying to come up with, I think, putting dedicated teams around it, having homogenized responses only to need to fill out the spreadsheet anyways, yes, it's a common challenge, and I think something that if we're able to find a way past and some more common language, etcetera, would probably help overall.

MR. CLANCY: And just to add, to the extent we don't do that with efficiency and it's beyond what we need to test the effectiveness of our controls, that's taking away from resources that defend our networks against attack. And so there's this big tradeoff problem that we have to make, is, we clearly have to provide transparency

to market regulators, sort of clients to auditors, etcetera, but there's a price that we're paying for that, and the ability of us to then marshal resources to defend our network. Because it's not a cost-free transaction.

MR. MCGONAGLE: And I know we were bumped up against time on this panel but --

MR. WASSERMAN: Four minutes.

MR. MCGONAGLE: Okay, good. So just on the question of the testing that the agency does, of the interaction that the agency does with our market participants is confidential, non-public discussions, right? And the sensitivity around the testing that's being done can't be underscored. But I wonder then about, is there some stamp or certification or some imprimatur about the testing that you're able to leverage in some way? You know this morning we had the Bank of England talking about their testing. Is there any utility in having a -- well, we've been, you know, subject to testing requirements by X and therefore that uniform standard gets you out of

having to respond from multiple inquiries of the same ilk?

MR. CLANCY: Noble goal -- it hasn't happened. I mean, we would love that. It's the proverbial holy grail of security assessment, trying to vet artifacts. But I've not seen one. We've tried as industry to come together and do some of these things. They work for a little while and then they sort of fragment and decay on their own. Everyone's like, I need one of these things. I have this additional question. I do know that some of our industry groups are trying to pull that up again and use some of the auditing standards, and what would be a common agreed reference artifact. I'm optimistic that we're looking at it. I'm also cynical that it's going to produce the outcome, because I've seen this happen a few times, but we have to keep trying, because that is where we need to get to, is that standard measurement and assertion that people can get confidence when they read the artifact that actually is the ground truth and they get an

understanding. That's what we all need. We just haven't figured out how to do it yet.

MR. PERULLO: And it may provide some assistance if the Commission or other commissions put a little bit of weight behind one of them. One of them that's out now that's kicking around is probably what Mark is alluding to, is SOC 2 plus NIST standard to taking the cyber-security framework.

MR. WASSERMAN: Okay, you need to define terms, right?

MR. PERULLO: Oh, God knows what they stand for.

(laughter) So the SOC 2 is an AICPA, is an accounting, at the end of the day, a CPA standard. Yeah, you know what that one is. An audit -- Standard --

MR. GREENFIELD: Service Organization Control.

MR. PERULLO: Control, so it's a control auditing standard. Long before it was cyber-

specific. It hasn't been cyber-specific very much. So that was already a standard. It was the old SAS 70. Somebody talked about what that stands for. And NIST has a cyber-security framework that they've released fairly recently, and so there's a group going on within SIFMA. I'll leave somebody else to fill that one in -- a work product there to try to come up with a way to enhance this SOC 2 auditing standard to actually map to those NIST controls. So hopefully that will be valuable for customers but if that would -- and you know, it would be one thing if the CFTC for example were to say yeah, that's great, but it would be great if that actually bought institutions something by complying. So if our lives were easier in some capacity, again, under a regulator, because we comply with that, then we would certainly drive towards it. And then once everybody at this table was on it, maybe the customers start to gravitate towards it as well.

MR. CLANCY: And SIFMA is the Securities Industry and Financial Markets Association.

MR. WASSERMAN: Good. Well I think we have run out of time, so I thank the panel once more again, really meaty, really really helpful. We are going to reconvene at twenty minutes after three.

(Recess)

MR. TAYLOR: All right and welcome to the last panel of the day on a very important topic, business continuity and disaster recovery testing, although as you'll hear some of our panelists may have another term or two to suggest in this space. I don't know if this topic is quite one topic to rule them all, but in a way it can embrace all the types of things we've been talking about all day.

A couple of administrative things at the very beginning: Panelists, if you will, when you want to talk, press the button on your mic. When you're done talking, please press it again to turn it off because the system will make funny noises if too many of us have our mics on at once. There are question cards, little 3x5 cards, over on the

table here and if people in the audience have any questions they'd like to send up to us, you are welcome to do that.

We are not going to have any extended set of closing comments at the very end of this, so our goal is going to be to actually get you out of here at 4:50, which I know some people catching planes and trains and so on will probably appreciate.

And I do want to say that we have one panelist, Randy Sabbagh, who's Senior Recovery Engineer for Schwab Technology, who's with us on the phone. Randy, can you say hi so I know it's working?

MR. SABBAGH: Yeah, this is Randy. Hello, everyone. How's it going?

MR. TAYLOR: Wonderful, thank you. Well, let me start this panel with the general question of what -- I'm going to ask it in a way that might be a little surprising, but we had a prep call with the panelists and they thought this term could be useful -- what does enterprise

resilience testing, which is sometimes called business continuity disaster recovery testing, mean to your organization and how has that changed in response to recent changes in the threat environment?

John Rappa, who's President and CEO of Tellefsen & Company, I'll turn to you first. And would you explain a little bit what is meant by shifting to the term enterprise resilience testing?

MR. RAPA: Sure, David, thank you. Taking more of a holistic approach --

MR. TAYLOR: By the way -- sorry. You all on this side might want to lean into your mics a little. I don't know why, but it's harder to hear that side. It's not you.

MR. RAPA: Okay, thank you. I think taking a more holistic approach of what's been your traditional business continuity management program that covers both the technology side and the people side and extending it and encompassing under it information security and cybersecurity in

the context of what we've been talking about today. So we're talking about the resiliency of your people and your processes should you have a cyberincident.

One of the things I don't think I heard earlier today, which I think is important not necessarily in the context of testing, is what is the awareness at the C-suite level of infosec and cybersecurity strategy and tactics? Imagine your CEO, COO, your CTO, even your Chief Compliance Officer, as this becomes and has become more in the media and in everybody's face every day, do they really understand what is going on? When the CTO says well, don't worry. We've got good firewalls and content filters and stuff like that. We're okay. Do they really understand what that means, what's behind that, and what types of questioning and interrogation is done at the C-level in the organization?

We've been talking about testing -- penetration testing, vulnerability testing -- but certainly a war room exercise, what's been called

a table top. Now, I find with my clients when I try to sell them a table top exercise, it doesn't fly. When you sex it up and you say a war room planning exercise, well, you get the testosterone going. But the ability to come in and put together a scenario that the following just happened: We've got a theft of data. We've got a corruption of data. What do you do? What's the thought process? You've got incident management teams in place. What's the involvement? What is the group dynamic between them when something like this occurs?

These things are quite valuable because you can do them without breaking things necessarily and is one additional type of test that you can do, but you need to mix it up. You can't keep doing the same thing over and over again. Whether it's the same penetration test or the same table top or whatever, you've got to mix it up. And when you start to plan these things, you've got to think deviously. We're at war here. People are coming at us and many people have said

today what they can do and you can let your mind just trek through this stuff. But if you're going to plan some of this stuff and you're going to look at your environment, you've got to think deviously.

MR. TAYLOR: Let me open this question to anyone on the rest of the panel who'd like to chime in. What's your concept of enterprise resilience testing? And you might touch on just a bit what does that sort of testing, what should that sort of testing, accomplish and maybe even touch what's going to be the next question -- if you're going to do it, how do you determine the scope that's needed?

MR. GIST: I would like to agree with everything John just said with one other important component and that's your supply chain. Your resilience is completely dependent on your suppliers and who you supply information to as well in order to maintain your service agreements. And if you don't have a good notification or incident management process not just internally,

but getting a phone call from those critical suppliers, that can put your resilience and your customer obligations at risk as well. So I wanted to put that out there.

And to help I guess advance the conversation on your second point, there is no test in a box. You need a series of tests, whether it be on the industry level, a group of companies coming together, a group of market utilities coming together, table top exercises, you need a testing program that is relative to the points spoken to before need to be risk based. You can or certain components could be done on an annual basis, but once again, if that is not where the risk is, some evaluation of that should be stated and some rationale should be documented as to why you have shifted your perception or devoted your resources to a specific area. And all of that knowledge on how to do that comes in my opinion down to one word and that's intelligence.

On the threat environments, what threat actors are doing, all the things we've heard

today. On the industry level for testing, it's been pretty much centered around 9/11-type events. The threat landscape has changed. We did not have the same type of threat activity from cyber and other nation-state threat actors and other highly sophisticated organizations that we do today. So most testing has evolved or needs to have these other additional components; not to say that physical testing because of 9/11-type events aren't important, of course, they are. We still have fire, flood, earthquake, and we change our technology components or our processes around all the time. You need to make sure that when you plug something into the wall, the light bulb is going to go off. So that will never go away. But you need to be able to say what the holistic picture is of what your risk landscape is based on intelligence and defining a series of threat scenarios that you can define those exercises against.

MR. TAYLOR: Randy Sabbagh on the phone. Would you like to weigh in on this? What should

enterprise resilience testing accomplish? How do you scope it? And since Greg brought it up, let me throw in how do you get the right intelligence component into it?

MR. SABBAGH: Actually, one of the most key components of this whole thing is making sure that the people who are going to be making these decisions have been trained and have an easy-to-use process to be able to manage these types of things. One of the things that -- they're affectionately known as the three P's -- you plan, you practice, and you prevail. And firms that take the time to do planning, but also practice and train their folks to be able to quickly make these decisions based on sometimes conflicting information or minimal information are the ones that are going to succeed. If you get into a situation where it's analysis paralysis, you may not be able to make a decision. You may be severely compromised.

But I think the key to it is making sure that you have a framework in place where people

know what numbers to call, where to go, what is expected of them, and also a framework for being able to make a decision quickly. If you know you can just basically say here are these potential scenarios. If this happens, this happens, this happens, this is what we need to do. But it also needs to make sure that whatever you come up with, it's not so full of technical jargon that you're actually completely excluding the people from the business side who are probably the more important part of the equation because they're the ones that are dealing with keeping our business up and running. Technology is an enabler, but to people that are actually running the business are the ones that are really making the money and are the ones who really have to wind up making the decision.

But I think the key to it is practicing and also making sure that you have your underlying framework for being able to do incident management and incident response. That's what is going to be key to having a successful testing program. And

as John and Greg basically said, your scenarios can be just about anything. But it's one of these things where it should be something that they can relate to from the business. I've been in some exercises where the scenarios they came up with made absolutely no sense and people just stopped listening.

So, again, it's look at your business. Where are your weak spots? Identify them and then say all right, I've got to train say 25 people. Let's bring them into a virtual EOC and then really throw a monkey wrench into this thing based on this one scenario. But the planning and the practice are going to be the absolute key things that are going to show the firms that are going to be able to respond quickly and effectively. That's it for me.

MR. TAYLOR: Let me press just a little bit, Randy, and then I'll do the same with the rest of the panel. You said some very interesting things in there about scope, and Greg was saying a minute ago you don't need just a single test. You

need a testing program. How do you determine the scope that's needed for this sort of testing for a critical infrastructure today?

MR. SABBAGH: Again, it's knowing your business. For some people their critical infrastructure is actually externally hosted. So for somebody it's okay, we've lost Rackspace, we've lost Equinix, or it isn't running. What are we going to do? Another scope is -- again, because we're seeing more and more regulations around vendor resilience supply chain, it's okay -- we've lost our market data provider. What are we going to do? Everybody else is able to trade except for us because we lost our circuit to X.

So again, it's looking at your business, knowing your business, and then giving us something that could potentially happen as opposed to something that's just so off the wall that people just won't be able to relate to it. Hopefully, I answered that question.

MR. TAYLOR: Let me turn this to the rest of the panel, the scope question. How do you

determine the requisite scope for critical infrastructure for a testing program?

MR. RAPA: So if you look -- and we've got exchanges and clearinghouses here -- you look at your traditional production systems that run the exchange, the clearinghouse, et cetera. You need to look at those as key, but also think about the fact that you've got an active directory. You've got a shared drive. You've got your Internet backbones, your phone system. You lose any of those, your shared drive gets hacked. Look what happen to Sony. That stuff is as valuable as what's in your clearinghouse systems. So you need to look at holistically the entire enterprise and do testing either on component or business unit levels and then across the enterprise. And then we've done between FIA and SIFMA, we've done industry tests the last 12 years that touch on this relative to the fact that I think Greg or someone said no one's infrastructure is static. You're adding new products, new features and functions. You're upgrading technology. No one's

environment is static. So you're testing every year to make sure that they work as specified and as expected.

MR. TAYLOR: David LaFalce from DTCC. David is Global Head of Business Continuity and Crisis Management there. You had a comment I believe.

MR. LaFALCE: I'm going to add a couple of things. I agree with everything Greg said. I agree with everything everybody said. I think we are at a juncture and a turning point. Largely over the last decade we've been very concerned about what's called kinetic events on the business continuity end.

MR. WASSERMAN: By which you mean?

MR. LaFALCE: Physical events, so storms, transportation outages, things like that, 9/11 events. So we're at a point now where -- this is kind of a perfect forum and a perfect time for this -- cyber and business continuity are kind of intersecting right now and we've got to determine what's next.

So the next is for me when I think of resilience, I think of it a bit differently. How can you flatten that curve of impact? So if you can go ahead and by rote, meaning by normal practice, go ahead and instead of using vendor #1, this month we're going to use vendor #2. Instead of using data center #1, this month we're going to use data center #2, thereby so you brought up the idea of active directory that may not be a thing that's tested. But sure as if you're going to be operating out of that other data center, it's going to be tested over a prolonged period of time.

Other aspects barring that, the idea of having integration between the event -- so let's say it's an inject of evil into your systems. That's difficult to go ahead and test via table top. So you almost have to go ahead and say hey, we're going to preface this by injecting evil into a lab and then we'll see what is necessitated from either recovery or a recovery and resumption point of view then after.

So I mean I think the key pieces to add are by rote, how much can you normalize on a regular operational basis, and then the idea of integration.

MR. WASSERMAN: And just to be clear, when you say "injecting evil," was it evil?

MR. LaFALCE: So I'm not speaking in terms of specters or anything like that. But my friend, Kevin Mandia who uses this term often, says malware, viruses, worms, things like that.

MR. TAYLOR: So, Chris Kinnahan, who's Associate Chief Information Security Officer for security operations at the Treasury Department, has a comment I believe.

MR. KINNAHAN: Yes, so I was going to say I think John said something really key earlier, which was how devious can you make your scenarios because that's really what we're coming down to. And what David had said earlier about we'd spend a decade going over what happens if a hurricane hits, what happens if an earthquake hits. Well, cyber events are very, very different

in the sense that it's a planned, thought-out, methodical thing. We never actually practice an earthquake and a hurricane and a whatever else all at the same time because the likelihood of that naturally happening is very slim.

But with cyberattacks, a lot of scenarios focus around okay, we found someone. They came in on this one particular vector, whatever else. The scenario needs to be they've been in my network for four years. So what can they do for four years? What happens when you can't trust anything that's online? So we've built a lot of systems that are very redundant, that synchronize very quickly, that are always available, but that can also be a hindrance in a cyber exercise. So what happens when they flip the bit? They corrupt some data that quickly synchronizes and all of a sudden you have four corrupted copies instead of one.

And I think that's really what it comes down to when you talk about scope. We need to really start thinking about how bad can it be?

And I know that's not a popular thing necessarily. I know it kind of goes a little bit against the okay, well, maybe the businesses won't see that as a likely scenario, but what we're seeing is that it is actually becoming a likely scenario.

MR. TAYLOR: John Rapa?

MR. RAPA: I think to Chris's point, yes, you have a scenario where data's corrupted and you've got three or four grandfathered copies there. You've got your business unit thinking about the fact that you can't open the doors this morning or this afternoon. What are we going to do tomorrow? Well, it may take us a lot longer because we've got to make sure all four copies are clean. So suddenly I'm not going to be able to open tomorrow. Who do I have to call first?

So some of these things are important to get your wheels spinning with your management team and your business heads.

MR. WASSERMAN: So at an earlier panel people had raised the loss of data integrity as perhaps the most serious thing, and here we are at

enterprise resilience. So I'm going to ask the sixty-four-whatever question, which is how do you plan for addressing a circumstance where you've lost data integrity?

MR. LaFALCE: So you get -- there's a cost benefit, right? So at some point in time for an enterprise like ours, you're likely failing forward instead of backward. So you're likely saying hey, everything's that cleared and settled prior to now may be no good. And so the idea is that that becomes the new benchmark and you have to actually fail the markets forward and reconcile forward, which is an interesting concept as you can see by your face.

MR. WASSERMAN: Did I mention I work in clearing?

MR. LaFALCE: But think about it, so if, in fact, the evil's been in there for longer than a period of time where you've cleared and settled a bunch of stuff, that now becomes your new baseline. Unless you have a DeLorean and a flux capacitor -- did everybody get that reference --

you can't go backwards anymore. So you have to fail forward.

So the concept of what we've been toying with is what -- so the cheapest thing in technology now I'm going to ask is probably memory, right, is storage. If we go ahead and ask our participants to store things, their native data, longer than the clearing and settlement period, then we have these native copies of data that we could possibly run through that become the golden copy again. It's a huge rule change. It'll be pushed back. But we're now thinking of the extended enterprise and maybe that's something that is the logical path forward.

MR. WASSERMAN: So what I'm hearing you say is that part of the solution there is through the rules of the infrastructure; you can basically look to your counterparties, your members --

MR. LaFALCE: The rules as an SRO.

MR. WASSERMAN: -- yes, as a self-regulatory organization and, therefore, you can pass rules that your members have to follow

and essentially so that they're maintaining information, which would be distinct from yours hopefully. That might be the solution there.

MR. LaFALCE: Correct. So that goes back to -- let's say we're taking data from Citi. It might be highly unlikely that data from Citi and data from JPMC and data from Morgan Stanley are all corrupt. Maybe we just can narrow it down to a singular institution if the corruption is coming on the submission side. So now we're into isolating where the evil may be coming from.

MR. TAYLOR: So implicit I think, David, in what you were saying is that business continuity and disaster recovery testing needs to have some focus on how to recover sort of when the inevitable happens. Would other people like to weigh in on how do you deal with that aspect of this?

MR. GARLAND: Thanks, David. I think the broader question is --

MR. TAYLOR: By the way, this is David Garland from CME Group.

MR. GARLAND: Thank you. I think the broader question is how do you -- we've talked about a lot of specific events. We just talked about data integrity. How do you plan for any eventuality? You can exercise. You can table top through any number of worst-case scenarios as John said, you know, the end of the world is coming. But not to beat a dead bear analogy again, but for one more time today, you can't tell which bear is coming to attack you. How do you plan for all of them? And I think a helpful way to do that is to plan for unavailability of people, systems, and facilities. And if you do that and you align the -- and this goes back to your scope question -- the scope of your testing with what the company thinks is its current risk environment and what it's most fearful of at the time, you can align those things correctly and then plan for them regardless of what comes to attack you.

MR. TAYLOR: There was a piece that relates to that in what some people were saying earlier. I don't want to go too far beyond before

teasing out a bit more, and that was that there needs to be an intelligence component here in terms of current threat in setting the scope for what's adequate testing for critical infrastructures. How do you all think that can be accomplished? How can the critical infrastructures get the intelligence component that's needed here?

MR. GIST: I think that happens on multiple levels. The FS-ISAC is a fantastic resource. Some people have private clearance authorizations to attend classified briefings sponsored by Homeland Security or Treasury. I personally don't think there are enough people with those classifications given the number of people in our industry that are involved in trying to design and think about threat scenarios that need to be tested.

I think that just the analysis of current media, the use of industry groups that bring people together to talk about what other companies or firms are facing in a very informal

environment, off-the-record conversations to talk about this happened to me last week. Have you seen something like this? It's building your trust network within industry as well to say I see something or I remembered this or reading about this on an FS-ISAC bulletin or I heard about this through Treasury. Maybe I need to pick up the phone and call somebody. That's how the first step in remediation would start taking place; just tell somebody that you think something's going on.

MR. LaFALCE: I think that Greg touches at least on the vectors for getting that intelligence correctly. We're in an interesting -- the DTCC as well as probably the rest of the clearinghouses -- are in an interesting kind of predicament. There's nothing anybody individually can probably gain from what we have in our stores. So probably somebody who's looking to attack us is looking for the secondary effect of taking down the economy. I mean it would probably be -- and I know this is ill defined and I'll use arrow quotes around this -- almost an "act of war" for somebody

to come after DTCC. And so to prepare for something like that is probably difficult because some of the strategies that may be utilized are not terribly public yet. And even talking about those strategies internally based on what we may find out during briefings may in and of themselves land us in a heap of trouble.

So we do go to that eventuality, as David Garland was saying, we do go to that ultimate eventuality from an impact point of view and then work backwards from there as far as the scenarios go.

MR. TAYLOR: John, then Chris.

MR. RAPA: Have you read the Tom Clancy novel, Debt of Honor, about 15 or 18 years ago? To Greg and David's point, a lot of the success we've had with the FIA and the SIFMA testing relies on whatever the secret sauce is, what I call the hub-and-spoke effect, the exchanges and the good relationships they have in the clearinghouses with their members.

And so if bad activity is determined,

things are percolating, information is percolating around, there are ways that the exchanges and the clearinghouses communicate with their members today already, those pipes, those relationships are there. I don't see that changing. I see that as part of the critical success factors of our resiliency also.

MR. TAYLOR: Chris?

MR. KINNAHAN: So going back to something David and Greg said about access to classified information and threat intelligence and all that. What I would say to that is there's a lot of open source information that is enough for the purposes of what we're talking about to come up with creative scenarios. And I think one of the things that we need to do is engage our technical staff at the lowest levels to say if you going to try to bring us down, what would you do? Because there's a lot of different ways that we would never think of at the higher levels that they'd be like I wouldn't even bother doing that. It would be really simple. I'll just knock out

our DNS servers or I'll just do this or that, which is an underlying technology that we would maybe not think about.

And so I think engaging at all those levels and running through those kinds of just thought exercises of okay, how many of us actually spend half a day in a room thinking about how to take down our companies without going to jail. But it's like we don't really do those types of exercise, but we should be.

MR. WASSERMAN: I would just observe -- I mean I think you're right that one of the possibilities you need to look at from the perspective of a critical infrastructure is what we were discussing earlier about nation-state actors. And to a certain extent you can say well, look, the resources of a nation-state actor are such that they can ultimately get through. I don't think, though, you can go too far down the council of despair. Ultimately, it is your responsibility as critical infrastructures to do what can be done, realizing that certain things

cannot be prevented. But then I guess part of it is going back to the old concept of business continuity and disaster recovery, okay. If the penetration testing that we've done is insufficient to protect us, okay, we've been penetrated, our data integrity is lost, now what do we do to recover from that?

MR. LaFALCE: I don't disagree with that and I don't think that I had implied that I disagreed with that before. I think that ultimately goes to what Greg was stating before, which is -- or sorry, David was stating before -- which is you lost this capability. Independent of how you lost it, what are you going to do? I completely agree.

I think what we've got to, though, couple with this now is in all honesty that's a very 2004 way of thinking I think because that hinges largely on again the kinetic and physical events.

What we're talking about now is -- so if you're going to kind of hold the firm to the 2

hour requirement, that 2 hour requirement was for full recovery to the end to maximum allowable downtime. Now you've got to add the component of the unknown, which is I've got to go find out what happened -- again, we're talking about a cyber event -- I've got to find out what happened. I've got to remediate what happened, and then I've got to recover. And that's a very different rubricking calculus than existed on 2003/2004's white paper.

MR. WASSERMAN: And while I will remind everyone of what I said at the very beginning of this roundtable, which is that anything anyone up here says is not necessarily the views of the staff of the --

MR. LaFALCE: I should echo that on behalf of DTCC also.

MR. WASSERMAN: Yes, I think you're right that if you've lost data integrity, recovering within 2 hours may be impracticable. But, nonetheless, you've got to say well, what can you do?

MR. LaFALCE: Agreed. I don't disagree with you at all.

MR. TAYLOR: Bob, by the way, we're chuckling up here at the Tom Clancy comment because we've been saying in FBIIC meetings and elsewhere for some years that you've had everything Mr. Clancy foresaw in Debt of Honor, including planes flying into buildings, with the exception of the destruction of the data integrity of the whole financial sector. Everything else he predicted has come true.

MR. ROST: I just want to add one other dimension to this. We've been talking about business continuity and disaster recovery as if your business is attacked, it goes down, or it is a hurricane and you lose capability. The bigger problem today with cyberattacks is the exfiltration.

We're losing literally hundreds of millions of dollars from intellectual property just going out of these systems. Every day you read about another cyberattack that's either going

after one of the health care things and compromising Social Security numbers or birthdates. There are unintended consequences downstream. You're looking just fine from this point of view. You're up and running. You're more valuable to the adversary in your upstate than your downstate because that's all about the resources, the value of the information that they're stealing.

So I think we have to be a little bit more nuanced on how we look at business continuity and disaster recovery. What are we recovering from? How much reputation can you withstand? Every day on the front page of the Washington Post is a different cyberattack. That's a form of resiliency, too, because the company has to be able to withstand itself in today's modern world with those kinds of things going on and have a good story to tell. What is due diligence? What did you do to prevent that? I think everybody understands that there's no perfection today, but there will be serious questions asked about what

did you do to prevent the exfiltration? And it's very different than losing capability. Both are bad, but sometimes we lose sight of the other dimension to this problem.

MR. TAYLOR: In light of all that and to pull us back to focus on what sorts of testing should the critical infrastructures be doing, let me ask whether you all think comprehensive end-to-end enterprise resilience testing is needed. And David LaFalce, I'll turn to you first, but I'd like others to jump in.

MR. LaFALCE: So I wrote down notes just so I wouldn't use acronyms. So I want to throw out two definitions first, and we've talked about both of them. The first is recovery. And so to me recovery is a purely technology term. It's the taking of what I'm going to call the compute environment and bringing it someplace safe to operate. Resumption is then operating that for business purposes. So when I think of end-to-end, I think of resumption. That's what it means to me and that's what we've adopted at DTCC also.

So the answer's simply yes. We do need to do this type of testing. I do think that a weekend exercise is highly synthetic in my mind. I mean what I would rather see is a move towards resilience, which again as I stated before is operating out of particular environments for extended periods of time. Because let's face it, we all have bits and pieces of our operations that happen once a month. And to go ahead and bring something up in another environment for a weekend may not test that once-a-month activity.

So I'm a big proponent of the concept of -- I guess it's somewhat of an active-active model, the idea of we're going to operate out of data center #2 for a period of time and operations center or people center #1 for a period of time. I think that's where we should be moving towards.

MR. TAYLOR: Let me follow up with that and then I do want to have others chime in. If critical infrastructures did that or if the Commission in some way said critical infrastructures, you need to do that, how much of

that would be different from the BC/DR testing that goes on today? And maybe we come back to this, but I do want to know what would that do to costs?

MR. LaFALCE: I think that you'd be -- I'm going to get to that answer. So I think that you'd be in a much more resilient environment. Don't forget, we're all bound by physics. So in 2001 through September of 2003 when the white paper came out and the 2 hour timeframe was bestowed upon us all, you're still bound by physics. So we all have multiple data centers in the same geographic region with something offsite, so something far away in an asynchronous mode.

MR. TAYLOR: Almost everyone.

MR. LaFALCE: I mean I think if you get to this resilience model and you do something like back off of the 2 hours and again look at the extended enterprise. I have an asynchronous environment. But maybe I couple that with the extended enterprise that we talked about before, so data exists for an extended period of time at

certain other places. You get into an asynchronous mode where maybe things are not replicated as quickly, so an inverse relationship right now between physical resilience and cyber resilience because of the replication problem.

You may extend the maximum allowable downtime let's say to 3 hours. The recovery may be just as long. Now you have a data reconciliation issue -- not issue, but you've got to go ahead. It's longer data reconciliation and then you get to resumption maybe within the 3 hour timeframe. But I think overall you're looking at a much more resilient sector.

MR. WASSERMAN: Let me press on that just for a second. If you're talking about -- I mean data reconciliation, you're talking about a question of integrity, yes?

MR. LaFALCE: No, no. With data reconciliation I'm purely talking about data loss at that point in time.

MR. WASSERMAN: Oh, I see. So you're talking about transactions in flight?

MR. LaFALCE: Yes, transactions in flight are lost or theoretically lost because the replication is now asynchronous.

MR. WASSERMAN: And is what you're saying that -- I mean is there a material difference in the ability to recover and resume in 3 hours versus 2?

MR. LaFALCE: Are you talking about material difference on whom? On the firm or on the sector?

MR. WASSERMAN: Each.

MR. LaFALCE: I personally think no. I think that in the greater scheme of things, if you go ahead and do a cost benefit analysis and say, guess what, I don't have two centers within 45 miles of each other anymore and I have one center here and another center here and maybe I put a data bunker somewhere else just in case there's a targeted attack, I think you're looking at a much more resilient play overall for the sector. Then I think you're also looking at -- yes, I don't think in the greater scheme of things people are

going to worry between 2 hours and 3 hours.

MR. WASSERMAN: Just trying to make sure I'm understanding. What I think I'm hearing you say is if you have greater distance than speed of light, it's the law, no regulator can change it. So then greater distance increases resilience. What I think I'm hearing you say is if we were to increase the mandated recovery time objective that would make it easier to have these greater distances.

MR. LaFALCE: Yes.

MR. WASSERMAN: Then I'm not sure whether you were saying yes or no. Is the necessary increase in recovery time objective something on the order from 2 hours to 3, or from 2 hours to 4? What is it that you're asking for in terms of the change in recovery time objective?

MR. LaFALCE: 3 hours was semi-arbitrary for a company like ours. For somebody like Greg who's got huge amounts of data, 3 hours may not do anything for him. But 3 hours allows you to be asynchronous. So if you have 120 minutes

currently and you go ahead and say all of that is taken up by recovery, and I don't have to worry about data reconciliation because I have a synchronous environment, but that means I still have to be within the same geography because I need a synchronous link. If I go ahead and extend it out and say now I need 3 hours because I have more data reconciliation required, I'm able to move away from the anchor that is that in-region physics.

MR. TAYLOR: And I think implicit -- Bob, it's not only 3 hours might help with geography, but it helps because you need to add this data reconciliation piece because you're asynchronous.

MR. LaFALCE: Right.

MR. TAYLOR: Let me turn this to the rest of the panel, and I do want to get a response from as many of you as can. And I'll say this; several of the gentlemen who are on this panel are very deeply involved in the annual FIA Business Continuity Disaster Recovery testing that goes on

now. So they know of what they speak. If we were to move towards the model that we've been discussing, how much of it is different than what goes on now? To do that piece as well as the other stuff?

MR. GIST: To one of the points Dave was raising, in terms of 2 to 3 hours, I think it's important to try to define what the entire recovery time objective and what the entire recovery point objective process is. Some companies don't start the clock on recovery until they've thought about what the incident is that just happened. So they take that first half hour to say, OMG, something's happened. Let me think about it before I call my technology people and say let's failover. You've just taken 30 minutes away from your recovery time objective. Somebody may say they've seen something on CNN and immediately pull the plug or pull the trigger on something and they have 30 minutes of additional recovery. There isn't a single industry standard on how that occurs yet.

So I think some base lining in that area needs to happen, and once again it depends on firm and capability. I don't think there's any way about it. It's just the way so many of our firms have grown organically or through acquisition.

For some of the larger firms to the data recovery and reconciliation point, it may take hours to try to figure out and reconcile systems with the supply chain just as well as your own internal systems. I don't think there is a process in industry that takes a health check to say that this component is here, this component is here, okay everybody can flip the switch on and everything will start again synchronously at the same point in time. That process doesn't exist yet.

So I would be cautious as to what you are defining that objective to be within that time window because back to the testing point, we haven't tested that yet to make sure that it's operational and we are capable of doing it.

MR. TAYLOR: John?

MR. RAPA: I would add to this and argue that depending on when an incident occurs during the day and you've got to synchronize more data, I would argue that an enterprise like ICE and CME and DTCC, you're pumping 10 to 20 terabytes of data a day between trading and clearing and the rest of your pipes. So depending on what time of the day something occurs, it could take longer. And then the ancillary effect on, okay, your members and key service providers and everybody else, it just ripples out from there.

MR. TAYLOR: Let me raise the question of how best testing needs and operational impact can be balanced in this area. Assume for a minute that the critical infrastructures do or are required to do what we've come to agreement on is the optimal, adequate, enterprise resilience testing that ought to be going on. How do you balance need for testing versus operational impact?

MR. GIST: I would say that there are so many different levels of requirements for testing.

There's the threat environment. There's what you have committed with your internal auditors to do. There is what testing with your partners. There is testing with your third-party suppliers. There are not enough days on the calendar to get all of this testing done with the threat environment constantly evolving, using the same people all the time.

So you have to figure out these green zones, if you will, as to when you can do this type of testing. And the more complex the threat environment is, and it's getting more complex. I'm not saying that we shouldn't do it, it's just that the boundaries of green zones that we have are a very scarce resource and that's where the operational impact is. But to help free up or create some of that green zone, perhaps one suggestion Dave made in terms of resilience in making sure you can connect the pipe that operates for an extended period of time is a possible solution. But that may not solve all of the scenarios that you need to plan and test against.

MR. GARLAND: I would agree with everything that Greg said. And in addition to talking about timing and how many tests one would be expected to do throughout the year, I think it's also important to look at the types of tests. So I think it's a really simplified answer, but at the very highest level balancing operational needs and testing should just look very carefully at not introducing any additional unnecessary risk. I think it was said earlier in the day, production systems are of the upmost importance. We cannot be introducing risk by doing testing for testing's sake. As long as the testing is responsive of a threat environment or a risk environment, which we feel we need to deal with, that's a good balance. But introducing risk just for testing is something we need to be very careful about, especially with the number of tests that Greg mentioned coming from all different ends.

MR. TAYLOR: In light of something that Greg alluded to, which is the wide variety of types of tests you need to do and the broad

variety of other parties you might need to test with, your vendors and so on, do you all feel that there is a need for a coordinated multiple entity or even sector-wide type of disaster recovery testing? Randy Sabbagh, let me turn to you on the phone first for that, but then we'll open it up to the panel.

MR. SABBAGH: I've actually been -- Greg and I have been co-leading the industry testing program for SIFMA for what is it, 7 years, 8 years now, Greg?

MR. GIST: 8 years.

MR. SABBAGH: 8 years. The question is whether you can actually perform end-to-end testing. It all boils down to what the clearing cycles for a system are. A lot of us have mixed technologies. Some of the stuff obviously -- I mean some of us are running ancient mainframe-based systems that are very clunky, but work, and then we have also other options as well.

The issue that you get into is if you get into something that's got like a T+3, how on

earth are you going to be able to take a base system out for 3 days in order to actually run your testing, especially if you have downstream processing that impacts a lot of systems internally as well as externally? The challenge is going to be just trying to figure out a way to really simulate the full process from end to end in a very limited time period and in a way that is not going to expose the firms to risk. We have had situations in the past where a number of firms who were using one clearing side firm, by accident, they opened up a trading queue and actually processed I think a large number of real trades that were in the queue for Monday execution and that caused some problems.

But, again, in order to do the testing you have to really understand the entire system from beginning to end. This was something like for us when we were working a commercial on commercial paper and Greg can chime in on this one. It took us a year to try to understand how that thing worked with [garbled] and then figure

out how on earth we were going to be able to test with it, and we discovered there really was no way. So there's a significant amount of effort that's got to be done to understand how these systems work and talk to one another and then how you can make them do what they need to do.

So, again, there is a need for it, but currently with a lot of these off cycles, like some are T+1, some are T+3, some are T+2. It does make end-to-end testing very difficult if you're trying to do holistic-type tests for the business line. Hopefully, I answered that question.

MR. TAYLOR: John?

MR. RAPA: And to add to Randy's point, we've tested -- the industry tests we've done are on a Saturday. Why? Because no markets are open on Saturday. But firms, exchanges, the infrastructure providers, they've got to get ready for Sunday night trading. So you've got a window where you've got to safe-store everything, get everything ready on Friday, open up on Saturday morning for testing, do order entry, get fills

back on that, get to the clearinghouse, pull stuff off the clearinghouse, exercise other systems, and then roll everything back with in the case of the FIA test, 62 firms, 24 exchanges and clearinghouses.

So you can't process the trades that were done by the order entry part of the test all the way through the entire plumbing, including the back-office systems like GMI and Rolfe & Nolan. You can't do that because there's not enough time. That's the challenge.

MR. LaFALCE: And to Randy's point, it's a completely synthetic test then because you're taking a T+3 cycle, compressing it into T+8 hours and then saying, yes, this is real.

MR. WASSERMAN: Although how much in -- and I'm going to show my ignorance about operational issues. In the futures industry, how much is T+3?

MR. RAPA: Three business days.

MR. WASSERMAN: Oh, what I'm saying is in the -- I realize in securities, but in the

futures industry, how much is T+3?

MR. GARLAND: I mean one is settlement date. It's the same day.

MR. WASSERMAN: It becomes a little bit easier I guess from that perspective.

MR. ORTLIEB: There are some commodity swaps that are 2, but yes, they are zero 3 days.

MR. TAYLOR: Let me turn the focus a little bit and ask whether there are best practices and standards out there? Which ones would be the most relevant here in the context of futures industry infrastructures and BC/DR testing? I want to be clear. I'm not asking about federal agencies or private sector market participants in general. Our focus for thinking about a rule that we might write is really on the critical infrastructures that we regulate.

Let me turn to Ron Ross from NIST first and then I'll come over to Chris Kinnahan for input.

MR. ROST: As far as disaster recovery contingency planning, we have two different

sources of guidance I guess you would call it. We have a special publication in our 800 series that deals with contingency planning for information systems. The special pub number is 800-34, good general guidance focused on IT primarily.

We also have one of our 17 families of security controls that deals exclusively with contingency planning, everything from developing the initial contingency plan to alternative communications capability, alternative storage sites, alternative processing sites. And, again, all of that is focused on the information system as being the core of the capability that we want to try to sustain during this disaster, whatever has happened.

So those are the two sources that we provide as far as continuity of operations contingency planning.

MR. KINNAHAN: What Ron said. So, yes, there's obviously a lot of published best practices out there. I think this goes back to -- I don't know if it was Greg or someone else that

said this much earlier. There's a point where there has to be industry collaboration and just information sharing amongst the different parties to sit there and say what has been working for you, what has been working for me. Best practices that are published are great, real world examples as to what actually works, which hopefully eventually works back into best practices, is even better.

MR. TAYLOR: John?

MR. RAPA: So you look at, especially since 9/11, 2001, 2, 3, the SEC, FINRA, the CFTC have come out with best practices in this area. Core principles under the CEA in Dodd-Frank touch on this. But if you look at even ISO-27002 as a standard, no one size fits all. So what works for CME doesn't necessarily work for OneChicago or the TowerXchange by scale and by size. And I think if you're looking at that, you need to look at more of a principles approach based on best practices. And, again, one size doesn't always fit all.

MR. WASSERMAN: So granting that, what

I'd like to try and do is -- and some of you folks are, of course, from the futures industry and financial sector. How can we take these standards, many of which are at a very high level for IT in general, and as David was saying narrow down and try and find what are the most relevant points to not fitting all, but fitting the market infrastructures in the financial industry more generally and in the futures industry more specifically?

MR. RAPA: I think Ron touched on a couple of these things, but if you look at the identification and mission critical systems information, backup and recovery of electronic and hardcopy data, alternate communication with clients and vendors, communication with regulators, there's about nine or 10 areas in the best practices that kind of span both the SEC and the CFTC equivalents. But I think if you look at those as a starting point, then you can go from there.

MR. ROST: I wanted to pick up on

something Greg said earlier. I think the supply chain and the -- we're doing a lot of outsourcing now. So if you've got an alternate telecommunications provider that you're depending on for your backup and they are vulnerable to some of the similar things that you're vulnerable to, that's a supply chain issue that we found during 9/11 when a lot of the cell service went down. There was a common core of that communications facility into the Trade Center's that impacted lots of people that weren't expecting it.

So we actually advise or have guidance that says you have to kind of run this a couple of layers into your supply chain to make sure that you're not bringing their vulnerabilities into your disaster recovery plan that could be impactful.

MR. WASSERMAN: So what I'm hearing you say, Ron, is it sounds like we have to have a balance. On the one hand we don't want to do the one-size-fits-all approach. On the other hand maybe in certain ways we're not as unique as we

think and there are, in fact, a number of commonalities that do tend to fit even our special part of the world.

MR. ROST: I think the reason that is the case is because all of us use pretty much the same information technologies. So we're kind of all working from the same threat space. We're all looking at the same basic architectures and the way we deploy our systems. And so there is a lot of commonality between the financial sector, the energy sector, because we're all kind of using these same little computers with hardware, software, firmware, and applications.

MR. SABBAGH: If I can add -- this is Randy. One of the things also to keep in mind -- and this panel that Greg and I were on with John Eckert, the lead auditor for the Office of the Comptroller of the Currency, the one thing that I caution people on is a lot of people say oh, it's not my problem anymore. It's in the cloud. I have outsourced. It's their problem. People need to understand, especially on the business side,

that just because you are no longer running it, you have actually magnified the risk because you have picked up the risk of the vendor along with yours. People need to understand that and factor that in in their planning.

One of the things we look at as potential we call domino effects. If we have a vendor, what we call a medium- risk vendor, go down in a function, what is it going to do to the rest of the firms that ripple out? In many cases I know lots of firms that have seen where one vendor will go down and take down the entire operation and very quickly without anybody actually realizing it.

So that's one of the things I think you have to factor in when you're looking at not only the technology aspect of the thing, it's also your feeds, services that are being provided, people that are actually using ACMD to host their frontends. If you're using AWS, Salesforce, a lot of things could really hit you.

Another thing also to keep in mind,

especially from the regulatory standpoint, is a lot of the firms that are on this call are actually very large. One of the things that we've had challenges on is when stuff that we've been developing as guidelines for the -- SIFMA put together a really nice workgroup on third-party risk management -- you have to factor in that a lot of the firms are small- to mid- size and do not necessarily have the resources to be able to do this stuff. So anything that the industry itself can do as far as frameworks, things to look at, things to worry about that don't terrify these firms and just make them feel more comfortable with the type of planning they're doing, it's going to add to the overall resilience of the industry because in the end we all really rely on each other. And if somebody goes down, they could take down a whole bunch of other firms as we've noticed in like flash crashes and things like that.

So, again, it's just not the ecosphere of your technology. It's also everything that

wraps around it and all the people that have to connect to it. Some people say oh, it's only mine, how do I manage it internally. There are people outside who could really do a lot of damage to you without you even knowing it because of all the supply chains and everything else and the interconnections that we have.

MR. WASSERMAN: That's really a good point. I should note, our rules currently state that while you may outsource functions, that does not relieve you of responsibility. I don't see that one changing anytime soon.

MR. KINNAHAN: Going back to Ron's point about supply chain, it kind of reminds me of this story. In a previous life I went out to a backup data center we had. It was a DR site. And I said oh, that's really interesting and who's that over there? And they said, oh, that's so-in-so's DR site. And I said, who's that? That's so-in-so else's DR site. And I said, what's that? And it's the only hotel within 20 miles. And all of a sudden they're like so we all have to timeshare

when we have DRs, got it.

And so I think when you make the point about outsourcing to cloud providers and things like that, yeah, they can maybe absorb my failure. But this goes back to the natural disaster scenario, right? Usually you're worried about okay, I have a hardware failure. It does not necessarily impact other vendors, or I have an earthquake, or I have a storm. With the current cyber landscape, they could take out whole sectors. They could take out whole areas, in which case then you're sitting there saying okay, well can the cloud provider absorb all of our traffic?

MR. LaFALCE: One of the -- a great parallel was remember the RSA breach from years ago. That was obviously a supply chain breach, but then they have a supply chain issue because then they have to reissue all of those tokens. So what's their throughput for something like that? That's just something really simple, right? And Randy touched on it and I think everybody touched

on it.

The interesting aspect is that -- and we found this out in 2010 I'm going to say when we did the last supply chain working group. Everybody's running the same stuff, so there's a hundred -- and I just wrote a note to Randy the other day to discuss this tomorrow -- there's a hundred of the same boxes and a hundred of the same applications that everybody has in their shop. And so maybe we just as a sector concentrate on them and trickle down that type of knowledge to the smaller institutions because they're part of our supply chain and ecosystem. They can't necessarily do it themselves.

MR. TAYLOR: Let me turn to a question that we talked about in the prep we would ask. And having listened to this discussion, I think I have to explain a little because I think the answer's going to be more complicated than we thought.

The question was we were going to ask about was the optimum frequency for BC/DR testing.

But the thing is I think I have heard you all saying BC/DR testing is not just a simple thing in a box. It's not just how often should we do the current FIA test. I've heard that some components of the testing that ought to be going on for critical infrastructures ought to be ongoing perhaps. They ought to be this month we're in data center #1 and next month we're in data center #2. And it's not when do you do it, but in a sense it's always going on. But I don't think anybody's advocating giving up the connectivity test that FIA leads now. There's another piece and you can go on from there.

So if you can, take a shot at first of all what are the major pieces, the higher level pieces, of the BC/DR testing that critical infrastructures in our world ought to be doing. And then in light of that, what's an optimum frequency? And if you would, remember we're thinking about doing a rule that at this table -- sorry, is going to be aimed at David -- is just going to aimed at the markets and the clearing

organizations, not at firms despite the fact that CME can't test without firms. It's a complicated answer, but with that in mind, would everybody take a shot at this.

MR. GARLAND: Sure, so I think your introductory remarks to that question, David, are spot on and it's a slightly more complicated answer. I think the short answer is that it depends, the frequency of testing. And the longer answer is what kind of testing do you want to do? What are your desired outcomes? Who's involved? What are the risks associated with each test? And then --

MR. TAYLOR: Let me throw one thing into that because we've done some preliminary thinking about this. And what I'm going to say here is not going to surprise anybody at the table I don't think. A way of saying what kind of testing should go on is testing that's sufficient to allow the critical infrastructures to fulfill their regulatory responsibilities; that is, to recover and resume and operate in spite of what might

happen.

MR. GARLAND: Sure, so we can start with the biggest test, the FIA test, which I think the cadence of annual testing has worked very well for the industry thus far. If you look at the percentage of participating volume, it's I believe north of 90 percent. John can correct me on that. But the other tests that you're speaking about -- and this is not necessarily the end-to-end testing with every firm and every piece of the futures industry's part of our critical infrastructure, but there are alternate worksite exercises.

There are smaller DRU unit testing you can do on small -- I think it was said earlier in the day the individual links of the chain rather than the whole chain, which can be more ongoing and reduce the people-spend on doing such a large industry-wide test. The table tops that we talked about earlier that John mentioned are key in addition to the actual failover tests. We've spoken several different ways about how it's important that the people who are making decisions

around why you failed over are doing that and they have got their muscle memory working in thinking about what needs to be done in the event that something happened that resulted in this failover.

And then there are exercises with partners and with external agencies, with various government agencies that can go on. So, again, it depends, but it really depends on what you're looking to accomplish with what the current risk environment looks like and also who your partners are in testing and how you can organize all that together.

MR. TAYLOR: John?

MR. RAPA: So, to David's point and again when we started the industry testing back in 2004, the idea was that firms like Citi that belonged to 10 or 20 marketplaces potentially would have to test 10 or 20 times over the course of a year. We put one common date together to get an economy to scale and again, people's infrastructures whether it's exchanges, the clearinghouses, the firms, the key vendors,

they're constantly changing over the course of the year. So recovery testing of your systems, your infrastructure, combined with business continuance of taking your key staff or selected staff to alternative worksites and having them manage the test, do the order entry, do the operations side of the clearing, from an alternate site. Conditioning them to do that, a byproduct of which is a need for cross-training and augmenting the big industry test with things that like Greg and David are talking about during the course of the year do two or three other key exercises. A war room scenario drill, individual tests run by IT on parts of the infrastructure. Everyone does a combination of these things and they change it up and you've got to constantly evolve over the course of time because the markets are evolving, products are evolving, the technology is evolving.

So it's not just one thing, but clearly the amount of planning to do an industry test is not trivial.

MR. SABBAGH: With industry testing you

have to do something with it to make it worth the while of firms to take part. I mean it's more of a -- somebody could say to me -- because one of the things we were looking at is for the SIFMA industry test. If somebody said to me you've been doing the same thing for 10 years. It's down to the point where we need to do it. I think you have to be able to make sure that when somebody's taking a look at this thing that it is worth their time and their effort to take part in it because they see benefit out of it as opposed to just finger painting-type stuff that you've done 10 years in a row. People just stop paying attention to things like that.

MR. GIST: To further Randy's point on that, one of the number one complaints we've gotten in recent years about the "SIFMA connectivity test" is that it's not reflective of the real world environment anymore. In 2001 technology was more tightly coupled geographically along with people, so a single incident could do serious operational damage to your firm. With the

diversity of people, geography, and technology these days to have firms operate from backup to backup is not reflective necessarily of a real world scenario. So that's one of the things that has helped or is one of the drivers I should say to try to help industry evolve testing in order to make it more real world oriented.

MR. TAYLOR: As a quick follow up to that, David LaFalce and some of the rest of you were talking earlier in the session about the need to shift focus beyond just kinetic events to cyber events. Would doing that help address the staleness, assuming there is as you were referring to?

MR. GIST: There are different issues with that. The primary one in my mind is the recovery time objective. What are you going to do if somebody corrupts your system or you have a corrupt piece of data? I've been in table top exercises where the participants have said that we need to stop operating because we don't know the extent.

And I'll bring another analogy into this. I'm tired of hearing about bears and birds and airplanes. It's the patient. Information security needs time to diagnose the patient, and business continuity needs to figure out the right type of life support to put the patient on while information security is trying to cure the disease. Information security needs time in many instances to cure the disease, so business continuity is not going to say let's activate our life support or our backup system until information security has adequately defined the disease to make sure it hasn't been spread into other organs of the body. So that's one of the things driving why BC-DR testing is coming together, just for these very purposes.

MR. LaFALCE: To Randy's point and to Greg's point, there's only so much reality you can -- the good thing about kinetic events is that they're easy to go ahead and conduct in real time where life is imitating art and not in the inverse. The problem with cyber so far is that

they're very much -- if you're going to frontend -- so what we do now is we do integrated exercises and they're largely kinetic-based, but a couple of years ago or last year what we did was -- so before a loss of region exercise we went ahead and said there's an EMP --

MR. WASSERMAN: Electromagnetic pulse?

MR. LaFALCE: Yes, thank you. And it must have been a big one because it knocked out most of Brooklyn and it came across the river and knocked down Manhattan also. But it also knocked out everybody's phone systems that would need to be part of the support. So we did the usual A, C, E alphabet and said you technology folks, you can't participate. It's a cyber, but it still has a kinetic element to it to make it as real as possible.

It's very difficult, at least in my mind and maybe Tom Clancy's got a better idea, but it's very difficult to go ahead and make this cyber exercise as it transitions into disaster recovery real. You can do it on paper. We can table top

it. But the only way I see making it real is to have a lab right next to it because you can't do these things on production systems. Have a lab right next to it and say, okay, based on what happened in the lab, based on the evil we injected into the lab, when would be the time we failover the production systems?

MR. WASSERMAN: Let me press you on that just for a second, David, because earlier on we were talking about well, the way you would recover from a loss of integrity is you go to your participants, your members, and get the information from them.

MR. LaFALCE: That's a future state. That's not a current state.

MR. WASSERMAN: Ah, because it strikes me that wouldn't that be the test that you presume for some reason or other your data has been corrupted and you can't fix it very quickly and so you need to go to --

MR. LaFALCE: So we've gone through this before. So there's the timing element. What data

do you want to get from the participants? Right now don't forget most rules say once you receive acknowledgment of settled or acknowledgment that we've gone ahead -- at least in our world, sorry -- acknowledgment that we've gone ahead and acted as the counterparty, you can delete your trades. You don't have to store that information anymore. So there's a rule change --

MR. WASSERMAN: Let me just press on that because our rules, and I can't believe the SEC's rules are that different, are you need to keep information related to your business for 5 years.

MR. LaFALCE: But there's a difference between information and playable data. Those are very different things. By the way, we may request that they keep it until settlement. I frankly don't remember our rules, but still there's that time component to it. I will tell you, and I would guess -- and I'm not trying to cause problems for CME or Greg -- but there's a very big difference between the data they have that's

replayable in the immediate sense and the data that they've got archived after a period of time.

MR. TAYLOR: We may need to look at that a bit. Let me turn -- we've got roughly 15 minutes left and I want to raise what for us as Bob has been saying in some earlier panels is an important point. When we're thinking about rules, it's incumbent on us to think about not only benefits of something that might be required, but about costs. And I have to preface this with the same thing I did the frequency question because we've sort of teased out here a picture of what BC/DR testing as it ought to be done might be and it's not one simple thing. So it makes this more difficult.

So let me ask it this way. Can you estimate the cost of the BC/DR testing that critical infrastructures ought to be doing? If not, why not? And if you can, can somebody take a shot at what are we talking about here?

MR. RAPA: I'll give you some feedback on some numbers we got a couple of years ago that

we commented to the SEC about Reg SCI that's on industry testing.

MR. WASSERMAN: You mean Reg S-C-I?

MR. RAPA: Yes, thank you. Sorry. The estimated number of man-days involved in planning and executing industry tests. They involve various types of skills -- operations managers, operations specialists, application engineers, network managers, network engineers, IT managers, information security engineers, business continuity managers, and key service providers. For exchanges and clearinghouses, between 175 and 200 man-days; for FCMs and key service providers, 80 to 85 man-days; and for the equivalent of SEPs or SDRs, 20 to 25 man-days. Planning, executing, postmortem. And as someone said earlier on the second or third panel, these resources aren't cheap.

MR. TAYLOR: Is that a picture of the cost of testing that's already going on today?

MR. RAPA: That's happened in the past, the past few years, yes. These estimates we put

together about 2 years ago.

MR. TAYLOR: Can any of you take a shot at if, for instance, we were to write a rule that established some minimums for a modernized-type of BC/DR testing for critical infrastructures that was different than what already is in place, what kind of cost would be involved there and how might it differ from the cost that we already know?

MR. LaFALCE: It's the additional costs. That's the key here, the delta between. I would think that if you're looking at having to involve maybe participants more than just a connectivity point of view, so it's an operational test. I'd say it's double. I mean if it's $250,000 per test for us, my guess is it's half a million dollars per test.

MR. WASSERMAN: I'm sorry, but do you mean that the total cost, including both you and your members?

MR. LaFALCE: No, that's just the hosting firm.

MR. WASSERMAN: Okay, so the hosting

infrastructure.

MR. LaFALCE: Yes.

MR. WASSERMAN: So you would double yours. And then I guess let me turn to John. Those estimates that you had, I take it those were on a per-firm basis?

MR. RAPA: Bob, again, exchanges and clearinghouses, FCMs, there were three different layers I gave you there. So exchanges and clearinghouses, between 175 and 200 man-days. And for FCMs and key service providers, 80 to 85 man-days.

MR. WASSERMAN: Each? Each or --

MR. RAPA: Each, each, yes.

MR. WASSERMAN: And are those numbers -- I mean is that sort of something that we might use as a basis looking at our world?

MR. RAPA: Yes, I would think it's certainly a data point.

MR. TAYLOR: And David Garland, I assume you could tell us if current testing is costing CME approximately 175 person-days -- we'd probably

have to modernize the term here -- and that doubled how much is -- what's the dollar figure for a person-day so we could do the math?

MR. GARLAND: I mean I think a similar question was asked in an earlier panel and the answer was there's no good answer. I would say that it would be an extremely substantial commitment is the best I can give you based on what we know today. And this is just the industry-wide testing you're talking about. This isn't all the other testing we talked about before -- alternate worksite, telecommuting, emergency communications testing, table tops -- the man-hours that are involved in those as well.

MR. TAYLOR: So is it essentially not really possible to quantify all of those additional components?

MR. GARLAND: I think it's an extremely substantial commitment is the best I can give you. It would be difficult to quantify.

MR. TAYLOR: No, I take it you say it would be extremely substantial, but is it

difficult or impossible to put any kind of dollar figure on the word substantial?

MR. GARLAND: I'm not entirely sure how we would go about doing that. I think, again, it would depend on the table top. Are we talking about 30 people? Are we talking about 60 people? How many agencies are involved? Are their partners involved? When we talk about DR testing, is it an internal test or are we just using for argument's sake a 100 IT resources out of region to do this? Or are we testing with partner exchanges or our customers in which case these numbers can grow exponentially.

MR. ROST: I think it's an impossible question to answer because it depends on the scope of the test that you're defining. How many people are involved, the skill levels, the extent of the -- how much you're exercising that contingency plan? How many different pieces? Unless there's a standardized scenario that you're going to come up with, even then you're going to have different entities providing different levels of effort

because there's no standardized amount they pay people for these different jobs that they're hiring.

So it's incredibly difficult, just like when you ask how much does it cost to do a FISMA set of tests on systems. It depends on what security controls you're using and how often you're testing and the level of effort you're going into each of those tests. So I think it's impossible to put a number on that.

MR. LaFALCE: So let me tell you how I put a number on it. We just went through this exercise for another acronym agency and we had kind of a clean slate because at DTCC we use one methodology for exercising and then at Omgeo, which we just absorbed, we had another methodology and the delta was what we looked at.

MR. TAYLOR: As a final question -- and I'm smiling to myself because in light of all the discussion we've had, I don't know if any question is impossible, but this might be a difficult one let's say. But in light of the discussion we've

had, all the different types of testing that might go into the sort of BC/DR testing program that would be adequate for resilience for critical infrastructures, what we've said about frequency for different pieces of that program, and what we've said about costs for different pieces of that program or not said, how should regulators address the resiliency testing that would be sufficient to protect critical infrastructures in today's cybersecurity threat environment? And that I would think -- and I'm speaking just for myself now, the same disclaimer as Bob gave a little earlier -- it might involve more setting of high-level principles and some minimums than diving at all too far into the weeds for granular particulars. But even with that in mind, how can we best address this to ensure that the critical infrastructures are, in fact, resilient enough today?

MR. LaFALCE: I keep beating this bear or horse or whatever the metaphor we want to use is, this dead horse. But if the ultimate goal is

resilience, again, I'm still a proponent that a resilient operating model is the best one, an active-active situation. If the ultimate goal is resilience, maybe testing's not necessarily the path to it. Maybe rethinking about how a company operates their production environments and things like that on a regular basis and looking at those controls or edicts that have been issued around that, maybe that's the best path forward. Again, I think these tests are good, but I think ultimately they're pretty synthetic.

MR. ORTLIEB: How's it auditable then? How can it be auditable at the end of the day? So if you do have a resilience goal, how can I --

MR. LaFALCE: If you have a resilience goal, you look at probably -- your key metric is your reporting mechanism. So the idea of what events have you seen? What are the root causes of those events, things like that? Beyond that up time, I don't know. I haven't thought down that far yet.

MR. ORTLIEB: You see what I'm getting

at, though, right?

MR. LaFALCE: No, no. I get it.

MR. ORTLIEB: You have to have a measurable goal that not only you are implementing for yourself, but that we then would say, okay, we're holding you to that standard. So without that yardstick, we're stuck on a straw man that we can't --

MR. LaFALCE: I get it, but then I would urge you to rethink is testing really the measurable goal of resilience?

MR. ORTLIEB: That's what I'm saying. So if you want to replace it with X, what's X, and then is it auditable and measurable?

MR. LaFALCE: Yes, I agree. I think that working backwards from that may be a logical pursuit.

MR. TAYLOR: I was going to say, I see some heads nodding and Ron, yours was one.

MR. ROST: I'm agreeing with David a lot. I think we put too much stock in testing, especially when you're looking at when we test our

systems in the federal government, we do these security control testing exercises and we get point responses back. You test this control, you get a response back. You test this one, you get a response back. It's like the -- I hate to use the airplane again, but we've got different pieces of the aircraft being developed and nobody's put them altogether yet. So the fact that I've tested all my controls individually and they're all doing just fine, that system still could be very vulnerable for the collective action together. They're not --

MR. TAYLOR: Excuse me, but I think I've heard you and the rest of the panel say in the real world, it's impracticable to put the whole airplane together and test it because people want to trade.

MR. ROST: Well, you want to do that at least one time.

MR. ORTLIEB: In real life, though, everything is testable. Remember that.

MR. ROST: At the end of the day,

though, the aircraft analogy does work because it is testable. We can just put one test pilot in there and say lift off and --

MR. ORTLIEB: But that's an operational exercise.

MR. ROST: But before all that final operational testing occurred, there was a lot of thought into the design, the development of that aircraft, best practices, the materials that were used to develop the aircraft. So by the time they get to that last phase, there's a high level of confidence that it's going to be resilient.

I'm not sure by doing these individual tests we're going to get that same type of payback, if you will. That's why I was thinking about what David was saying. It's worth exploring because if I can express the type of properties that exist within one of these critical infrastructures, having a good enterprise architecture, as one of the people said earlier in the last panel, making sure I have a good contingency plan, looking at that plan, taking it

down and doing the different scenarios, that gives you greater confidence that the organization has done the most important things in a cyber world to reduce their susceptibility to the cyberattack, which could either result in exfiltration or a loss of capability. And that may be much more valuable than these individual tests that really you can never run this thing full out from what everybody's saying. Now, you guys are the experts on that.

MR. WASSERMAN: Let me press on that just for a second here because I think what we're talking about here is not specific individual tests that we would require. As Jim Ortlieb was saying, ultimately as regulators we have to be able to verify what folks are doing because I can guarantee you one thing, if we go to the registrant and we say are you doing enough, I know what the answer is. Yes. Great, but how do we define enough and how can we on a principle basis because honestly I don't see how we can get to the level of deep detail that say you folks at NIST

can do because as I understand it, that's a constant effort on your part and that's not practicable. How can we establish principles and similarly what we can do from an audit or review perspective is not testing down to the level ourselves because honestly, there's a resource constraint. We then get into arguments. And so how can we set up principles that would promote the resilience and that we can then review in some kind of reliable way?

MR. TAYLOR: David?

MR. GARLAND: I can say that when setting up those principles, there's a couple of things that I think would be helpful to consider. The first of which is that -- and it was said earlier. We all use the same IT infrastructure, but every little piece that every firm has within even just the futures part of the financial services critical infrastructure does a different thing. So it's important to look at that and understand that one size doesn't fit all.

Additionally, when we talk about

resilience -- and this goes back to the very beginning of the panel when you talk about enterprise resilience -- what are you trying to prevent? Again, we're not focused on testing. We should be focused on preventing disasters. So when you look at a principle that firms should aspire to, I think it's important to look at that risk and say how are you addressing it either through testing or some of the other ways we discussed?

MR. TAYLOR: I think with that -- no, if it's another comment, go ahead, John.

MR. RAPA: I'm just going to add one more thing to David's. We talk about supply chain, supply chain disruption. What's the most valuable part of your supply chain? People. So clearly you want to understand the people preparedness if you have a disruption. And based on the nature of the disruption, how do you respond to the incident? How well prepared are your people to continue the business and from where and how? We talked about a number of

different scenarios throughout the course of the day here. The people side is important as well as the technology and you need to take that into consideration with whatever you do.

MR. WASSERMAN: Granting that, again though, we're coming back to the problem we have, which is how can we set principles that you can then assess and we can examine your assessment of to verify what's going on and verify that you're meeting the goals that honestly I think everyone here acknowledges. It's in your interest, right? These are your businesses. And so I think it's a concern of ours that we establish the right principles and that they're really the right principles. But then there needs to be some ability then to have the private firms measuring whether they're meeting them and us to be able to examine that. And I guess my question is is there some way that we can do that that, again, gets to the right results?

MR. LaFALCE: Now, this is -- I'm not trying to take work away from NIST, obviously. If

you focused on things like design principles and then the order of the validation or the metrics or the litmus test was how the firm operates under those principles. So in our world as part of our settlement operations they rotate their schedule between New York and Tampa. That's one of the most resilient things we have. One of the metrics could be the amount of days, successful days of settlement or something like that, out of each site. If it's 20 percent and 80 percent, then that's not the balance we're looking for, successful settlement days out of data center A or data center B. Maybe those are the litmus tests or those are the KRIs or KPIs -- key performance indicators -- that would necessary to measure adherence almost to the design standards.

MR. TAYLOR: Seeing no further flags, I think we've reached --

MR. GIST: I have one more.

MR. TAYLOR: Greg?

MR. GIST: I agree with everything that everybody has said, but there's still a part of my

gut that says how do you know based on -- and my magic word is intelligence -- that you're hitting the right things?

I think to one of the points Randy Sabbagh made, there are so many firms spanning financial services, not just futures, that don't have the resources to do those things. Treasury is sponsoring a two-year exercise through the FSSCC on a series of various cyber exercises, Quantum Dawn 3 being one of them, that are meant to target different size firms in different scenarios in different capacities with different incidents in each one. Each one is its own unique incident and I think one is international. As a matter of fact, I think they're doing one with U.S., Canada, and the Bank of England.

The problem with that structure is that it's only available to the FSSCC membership. If you could figure out a model that the government or -- I don't know how to translate this into something that's operational -- but to look at that model and be able to lift it and create that

recipe for firms to say I can do this, but I can't do that. If you could figure out how that recipe plays into the futures industry, I think that would be very beneficial to everybody.

MR. TAYLOR: Well, I think we've reached the end of the roundtable with the last panel, and I would invite Chairman Massad to say a few words to conclude.

MR. MASSAD: Well, David and Bob, I really should let you conclude. But let me just say I've been able to be here for quite a bit of this. I had to be in and out on this panel, but the day was really incredible. I mean the amount of expertise we had gathered at this table over the course of the day was really, really impressive.

So I just mostly want to thank all of you for being here, for contributing your time and your knowledge. It seemed to me that each panel we probably could have spent the whole day with each panel if not more and benefitted a lot, but it gives us a lot to think about. And I just want

to underscore in terms of at least how I think about this and I think the staff, we're not trying to write rules or set requirements just to show that we've written rules or set requirements. We're trying to figure out how we can really add value here. I think the discussion was very helpful in that regard in terms of thinking about how do we build on best practices? How is it collaborative? How does it help facilitate information sharing? So you've given us a lot to think about and, again, just thank you.

MR. TAYLOR: And thanks to everyone for coming.

MR. WASSERMAN: So the good news is that we've had some really incredibly good panels, and I'd like to second my appreciation to everyone who's participated. The good news is we've accomplished a lot. The other news is that we at this table have a very complex task ahead of us. On the other hand, you folks out there and the panelists and the broader industry have as well a very important responsibility both in terms of

helping to solve internally to the industry these very complex problems and as well assuming we do go forward and propose a rule to participate in the common process to help make sure we're getting it right. And so I think there's a lot of very challenging, but I think ultimately incredibly worthwhile work ahead of us. I mean I recall from the first panel just what's at stake here. We have just an increasingly complex environment where we're getting threats from incredibly able actors, including state actors. It is really is our duty to get this right. So thank you very much and I look forward to working with all of you.

(Whereupon, at 4:56 p.m., the PROCEEDINGS were adjourned.)

* * * * *

CERTIFICATE OF NOTARY PUBLIC

DISTRICT OF COLUMBIA

I, Stephen K. Garland, notary public in and for the District of Columbia, do hereby certify that the forgoing PROCEEDING was duly recorded and thereafter reduced to print under my direction; that the witnesses were sworn to tell the truth under penalty of perjury; that said transcript is a true record of the testimony given by witnesses; that I am neither counsel for, related to, nor employed by any of the parties to the action in which this proceeding was called; and, furthermore, that I am not a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

(Signature and Seal on File)

Notary Public, in and for the District of Columbia

My Commission Expires: May 31, 2018