

ONE HUNDRED FOURTEENTH CONGRESS  
**Congress of the United States**  
**House of Representatives**  
COMMITTEE ON ENERGY AND COMMERCE  
2125 RAYBURN HOUSE OFFICE BUILDING  
WASHINGTON, DC 20515-6115

Majority (202) 225-2927  
Minority (202) 225-3641

March 27, 2015

**Dr. Herbert Lin**

Senior Research Scholar, Center for International Security and Cooperation  
Research Fellow, Hoover Institution  
Stanford University  
Encina Hall, C-236  
Stanford, CA 94305

Dear Dr. Lin:


Thank you for appearing before the Subcommittee on Oversight and Investigations on Tuesday, March 3, 2015, to testify at the hearing entitled "Understanding the Cyber Threat and Implications for the "21st Century Economy."

Pursuant to the Rules of the Committee on Energy and Commerce, the hearing record remains open for ten business days to permit Members to submit additional questions for the record, which are attached. The format of your responses to these questions should be as follows: (1) the name of the Member whose question you are addressing, (2) the complete text of the question you are addressing in **bold**, and (3) your answer to that question in plain text.

To facilitate the printing of the hearing record, please respond to these questions with a transmittal letter by the close of business on Friday, April 10, 2015. Your responses should be mailed to Brittany Havens, Legislative Clerk, Committee on Energy and Commerce, 2125 Rayburn House Office Building, Washington, D.C. 20515 and e-mailed in Word format to [brittany.havens@mail.house.gov](mailto:brittany.havens@mail.house.gov).

Thank you again for your time and effort preparing and delivering testimony before the Subcommittee.

Sincerely,

  
Tim Murphy  
Chairman  
Subcommittee on  
and  
Investigations

cc: The Honorable Diana DeGette, Ranking Member, Subcommittee on Oversight and Investigations  
Attachment

## Attachment 1—Additional Questions for the Record

### The Honorable Tim Murphy

1. Each witness provided a slightly different perspective on cyber threats and the challenge of cybersecurity, extending from the past, to the present and future.

a. Are there areas where you feel there is a common view or shared theme and what is it?

I'm quite confident that we all share the view that cybersecurity is an issue of true national importance, that we are not doing as well as we could against the cyber threats we face, and that the road ahead to improving the nation's cybersecurity posture significantly will be rocky and difficult.

b. If there was one fundamental message you want Congress and the public to understand about cybersecurity, what would it be?

There are tradeoffs to be made for better cybersecurity, and if you want better cybersecurity for the nation, you had better be willing to make tradeoffs against other things that you hold dear. As an example that I emphasized in my testimony, market forces virtually require that reducing time-to-market for innovative products and services takes precedence over building these products and services with good security from the start.

c. Are there specific issues or areas of this issue that do not receive an appropriate level of attention?

The fact that there are tradeoffs between cybersecurity and other desirable attributes or public policy outcomes is not discussed or appreciated adequately. For example, in recent months, the FBI has asserted the need to have access to encrypted information in pursuit of their mission, while much of the information technology industry has argued that building in the capabilities for providing such access to law enforcement authorities would reduce the security of their products and services in ways that are detrimental to their marketplace acceptance. Both side are being truthful, and there is no way to fully reconcile the two conflicting positions. The nation must make a tradeoff between two good things—and the need to make that tradeoff has been obscured by absolutist rhetoric on each side asserting that the other side has no facts or argument to support it.

2. As the promise of innovation connects more of our lives to cyberspace — from smart pacifiers to cars that communicate with each other — cyberspace becomes, in theory, a limitless attack surface.

a. How do we manage the risks presented by "smart devices" and the Internet of Things while also enjoying the benefits and convenience they offer to society?

The security problems posed by smart devices and the Internet of Things are not different in principle than those posed by other computational devices. However, smart devices and the IOT pose a problem of scale—along with hundreds of millions or even billions of new Internet users

coming online in the next decade or so, smart devices and the IOT have the potential to make the security problem much, much worse. In addition, widespread adoption of IOT devices suggests that they will be relatively inexpensive, which gives vendors less ability to build in cybersecurity capabilities. And they may well be installed in place where they are not regularly updated, which means that security patches are less likely to be installed frequently. Finally, the public attention given to cybersecurity for smart devices and the IOT by vendors and advocates is certainly greater today than it was in the past for other computing devices, but whether security *practices* have actually changed significantly is an open question.

Managing cybersecurity risks is largely a matter of persuading or incentivizing vendors and users to pay more attention to cybersecurity issues, and I prefer the use of market forces to do that over direct regulation. Harnessing market forces to this end means that something must happen that adjusts the market forces in that direction; leaving it all to the vendors and users to decide on their own what they want to do regarding cybersecurity is a recipe for inaction on the security front. But there is no consensus on the steps that might be needed to adjust market forces. For example, some people believe that liability of cybersecurity defects would help to push vendors to pay more attention to cybersecurity; others believe that liability would dampen innovation significantly. For every measure proposed to harness market forces, good reasons can be assembled to oppose it.

- b. As more devices connect to cyberspace and interact with one another, what challenges does this present for how security professionals or companies anticipate potential vulnerabilities or risks?

More connected devices mean a more complex environment to analyze for security risks. Today's analytical tools are inadequate for this task. Security companies and professionals will have even greater difficulty in undertaking such analysis in the absence of better tools.

- c. How do we assess the security of individual products relative to the security of the system as a whole?

Today's systems are composed of a variety of components. Even when those components are individually known for sure to be trustworthy, there is no guarantee that the system as a whole is itself trustworthy. And the trustworthiness of individual components is difficult to assure as well.

- d. In such an interconnected world, how do you draw the line between a potential vulnerability and a realistic vulnerability? In other words, just because something is possible, how important is it to assess the probability that it will occur?

It is important to distinguish between two different types of threat. One type (call it Type A) is a threat that does not change in reaction to a change in the target's defenses. Typically, the perpetrator of a Type A threat does not care very much about success against a particular individual target, but rather relies on statistical likelihood to succeed. The canonical example is a criminal trying to steal credit card numbers—he does not particularly care whose credit card numbers he obtains, only that he obtains as many as he can get. A second type (call it Type B) is a threat that does change in reaction to a target's defenses. The perpetrator of a Type B threat cares very much about success against a particular target—for example, the CEO of a major

defense firm. When that CEO's security people deploy defenses to thwart a particular kind of attack, the Type B threat will morph into something else and the perpetrator will try again. And the perpetrator will try repeatedly until successful.

The difference between Type A and Type B threats is that against a Type B threat, one must address essentially all vulnerabilities, independent of likelihood of exploitation. Why? Because the Type B threat can take advantage of any vulnerability. By contrast, when facing a Type A threat, probability does matter because by definition, the Type A threat can only take advantage of a given set of vulnerabilities and will not change. Thus, addressing the vulnerabilities most likely to be exploited by a Type A threat has significant value in the sense that such action will extend the period of time in which a Type A threat will be unsuccessful.

3. Quite a few respected technologists — at Google, and also at the Institute of Electrical and Electronics Engineers Computing Society — have theorized that in the future, the Internet will be so integrated into our daily lives that it will become "invisible" and provide "seamless intelligence."
  - a. Can you expand a little more on how exactly a world with an "invisible" Internet would work?

Today's electrical power system provides an analogy. There are a huge number of devices in our homes and office and factories that use electricity. But for the most part, the electrical infrastructure is invisible to us—except when it fails. The vision offered by Google and the IEEE CS is an appealing one of information at the automatic beck and call of any device whose operation can be improved or is enabled through the use of the appropriate information. And the Internet is expected to be the platform through which such information is delivered at the appropriate times.

- b. Do you agree with these predictions? Why or why not?

I agree with them in the sense that I don't believe that the vision is fundamentally impossible to achieve, and that the world depicted could be a desirable one—provided that other concerns are adequately addressed, such as security, resilience, and privacy. But whether these other concerns will in fact be adequately addressed is anyone's guess. Standing in the way of achieving this vision, entirely apart from the very formidable technical challenges, is the need for a societal consensus about the right balance between many desirable qualities of this world. And at this stage, I don't see what that consensus might be or how it might emerge.

4. No matter how much money a company invests in security software, training and other cybersecurity measures, they still remain vulnerable to the insider threat. This can range from the intentional actor — such as a disgruntled employee stealing information or letting the bad guys in — to inadvertent actors — such as an employee clicking an infected link in a targeted phishing email.
  - a. Will companies ever be able to prevent internal threats — employees lowering the proverbial draw bridge — regardless of whether their actions are intentional or unintentional?

No. They may be able to detect such individuals after they have done their dirty work, but if the individual(s) in question is or are willing to bear the consequences of being caught, there is no way to thwart entirely the insider threat. Note also that a goal of many outsider threats is to achieve insider status, a goal often reached through social engineering against an insider.

- b. if it can never be eliminated, does it come down to managing risk? Are there proven strategies to minimize this risk?

Yes, it is a matter of risk management. And some strategies for minimizing risk are better than others. But the optimal strategy for any given organization is highly dependent on the nature of that organization; one size does not fit all in the world of risk management. In other industries, risk management techniques have evolved to address the insider threat. For example, in the financial industry, techniques such as double-entry accounting, separation of duties, and periodic audits have emerged as useful risk management techniques. All of these techniques involve some degree of business process redesign, which is inevitably one of the

- c. How significant of a challenge is this to those evaluating the cost benefit of security measures?

It's a huge challenge. For example, those trying to do cost-benefit analysis of security measures need a lot of data to make their assessments. That means they need to collect it, and identify specific ways in which security measures can fail. Sometimes such data is available; more often, it is not. Even when it is, analysts need to know how to use it. As a rule, it is easier to quantify what a security measure costs, although one must be careful to account for non-obvious costs such as loss of convenience, availability, and so on. But it's much harder to quantify how many of those bad things were prevented from happening because of the security measures or for some other reason.

5. Information sharing, though it has its benefits, is still a reactionary solution. Someone has to first suffer an attack before that threat information can be shared, and oftentimes the attackers change their signatures from target to target.

- a. How does information sharing help reduce the gap between cybersecurity capabilities and threats to cybersecurity described in your written testimony?

Information sharing can help in a number of ways. For example, knowledge that you have been attacked may alert me to the possibility that I have been attacked, or am about to be attacked. The former case may prompt me to do an investigation to determine if in fact I have been attacked, an investigation that I would not otherwise do. The latter case may prompt me to take additional defensive measures, above and beyond those that I would have taken in the absence of such information. At a technical level, information sharing may help me to pinpoint the threat that is attacking me.

There is a lot of focus on signatures when it comes to information sharing.

i. Are signature-based defenses effective? Why or why not?

Signature-based defenses are effective against a certain kinds of threat but not against other kinds of threat. An example is that not all signatures of possible threats are known in advance of that threat's strike. A zero-day vulnerability is one whose existence was not known prior to an attacker's use, and thus against which no specific defense could be mounted. Note that signatures are only a subset of information that might be shared.

b. How does information sharing fit into the broad picture of the cybersecurity challenge?

i. Does it offer opportunity beyond improving our defensive capabilities?

If one is willing to include under the rubric of "defensive capabilities" issues such as attack detection and remediation (as would be proper), sharing has value far beyond attack prevention; these include detection; remediation. Moreover, information sharing of all kinds is at the heart of successful collaborations. Sharing threat information may not directly contribute to business collaboration, but it may be the first step towards establishing a trust between two organizations that would enable them to share other kinds of information.

6. Is it possible to quantify the benefits of the Internet and information technology relative to the cost of security?

a. In other words, is it possible to calculate the economic benefits of these technologies relative to the economic costs of cybersecurity, including prevention and response in the event of a breach?

There is a substantial cottage industry devoted to making such calculations, but I'm sorry to say that their estimates differ from each other significantly. Moreover, their methodologies are highly questionable. So from my perspective, I can say that I've never seen a comparison of these costs that I believe or have any faith in.

b. How about the social, cultural or other less tangible benefits?

These benefits are even harder to quantify. As an example – what dollar value would one put on the ability of Americans to communicate freely with each other through electronic means? One could estimate the revenues of the telecommunication industry and use it as a proxy for the monetary value that we as a nation assign to communications. But as a U.S. citizen, I value my First amendment rights even when I am NOT communicating with others—and there's essentially no amount of money that anyone could give me that would compensate for the loss of such rights.

c. Is there value in this?

I interpret the question to mean – is it valuable to quantify the value of the Internet and information technology and the costs of keeping them secure? Yes, I think there is value in the

exercise because the exercise forces the analyst to consider various factors systematically, but it will be important to refrain from treating the numbers that emerge from such an analysis as anything more than suggestive. Any particular analysis is likely to omit some very important factors and to have very large uncertainties about the estimates it does make.

7. Discussions about cybersecurity often focus on prevention or keeping actors out of system

- Is this the right way to approach this issue?

a. If there is no guarantee the bad guys will not get in, should the emphasis shift to a focus on resilience rather than prevention?

In the long run, both prevention and resilience have meaningful roles to play, but I agree with the thrust of the question that the value of resilience is underappreciated. Today's world is largely one of perimeter defense, a paradigm in which you can cleanly separate "inside" from "outside". The boundary between inside and outside is the perimeter, which is where most defensive efforts are concentrated. But with this model, an attacker that is successful in penetrating the perimeter then has free rein inside the system, with little to impede his efforts. Consider, for example, that in any organization, the information technology on which it relies is the end point in a long supply chain, and compromises to supply chain security enable the attacked to be present "inside" even before the perimeter of the system is established. Savvy organizations are learning to operate in a compromised information technology environment, with all that such operation implies, though perimeter defenses are still valuable in reducing the scale of the problems they face inside the perimeter.

I also note that even defining a perimeter is often problematic in a world in which the components that make up the system originate in myriad places over which the system owner or operator has no control. Even an individual hardware chip may have circuitry inside that comes from many different and possibly untrustworthy sources.

b. Why is the concept of resilience important to effective cybersecurity?

Consider the value of file backups. You back up your files so that if a file is accidentally deleted, you can retrieve a recent copy and not lose most or all of the work that went into creating it. It's not a big step to imagine a bad guy deleting your precious file deliberately, but even in this case, the backup has significant value. In this context, backup is a part of effective resilience-based cybersecurity.

Backups are not free, however. You need to expend some time and effort to perform a backup. So you work somewhat less efficiently because you don't entirely trust your environment—that is, you operate under the assumption that environment may be (probably is) compromised.

The same general lessons apply to any other aspect of resilience. You pay something in time and effort to preserve some essential functionality—you hedge against disaster.

c. How does resilience support a risk-based approach to cybersecurity?

You select specific features of a resilience architecture depending on what you care about most. For example, a bank might care much more about preserving the integrity of its data than its

confidentiality—that is, it would be much worse for a bank to have its records scrambled so it did not know what was in the accounts of every depositor than for those records to be revealed to the outside world. Neither is good for the bank, of course, but under these circumstances, the bank might well provide extra support for resilience efforts to enhance data integrity than data confidentiality.

8. Dr. Lin, in your testimony you said that "complexity is the enemy of cybersecurity."

a. Is it possible to reduce this complexity?

i. If yes, what are the consequences?

ii. If no, why not?

It is unquestionably true that some reductions of complexity are possible through more careful design given a set of performance requirements for a system. But in my judgment, by far the biggest driver of complexity is that we want our systems to do more and more. That is, our appetite for greater functionality in our systems is essentially unlimited. Any serious attempt to reduce the complexity of systems has to start out by someone being willing to say "no" to demands for more functionality.

Based on what you said about the complexity of a system increasing when additional components are connected to it, the "Internet of Things" is going to exponentially increase the complexity of the Internet.

b. What does this mean for the governments, businesses, and individuals that are going to use these connected devices?

I don't disagree with the implication that the IOT will make the internet much more complex, and thus more insecure. We can mitigate it to some extent, but for many people, the benefits of the IOT are not so compelling that it is worth the added insecurity that will result. I personally expect to be a late adopter of these technologies for exactly this reason.

c. How will this influence or reshape current cybersecurity practices?

I suspect the most important influences will be that it will increase the demand for people knowledgeable about cyber security

9. In the last few years, there have been several significant compromises and vulnerabilities discovered in regards to digital certificates and Certificate Authorities, two of the best well-known being the compromise of DigiNotar and the recent LenovoiSupertish revelations. This raises questions as to whether the digital certificate model is providing an adequate level of security for users of the Internet.

a. What are the weaknesses in the digital certificate model?



i. How significant are these weaknesses?

ii. Can these weaknesses be eliminated or adequately mitigated?

Certificate authorities exist to “certify” that two parties actually in remote electronic contact with one another in fact correspond to the parties’ expectations. John thinks he is talking electronically to George; a certificate authority certifies that John is really talking to George and not to Sam. In the absence of a reliable certificate authority, John might in fact be talking to anyone. A digital certificate is an electronic credential issued (in this case) to George that certifies that anyone relying on the certificate and communicating with George is indeed talking to George.

Therefore, a certificate authority must be trustworthy. Its technical protections must be robust enough to withstand attacks intended to compromise its certifications. Its management protections must be robust enough that it does not issue certificates improperly (e.g., does not issue a certificate saying “George” to Sam.) But over the years, the number of certificate of authorities has grown. With such growth the variance in the trustworthiness of the best and the worst has grown, and it is of course the least trustworthy certificate authorities that are the most vulnerable targets for compromise. For these untrustworthy certificate authorities, sloppy implementation of technologies and processes for issuing and managing certificates is common.

b. Are Certificate Authorities subject to any form of oversight?

i. If so, by whom and how does this function?

ii. If not, would enhanced oversight help address the weaknesses examined in Question I? Why or why not?

To the best of my knowledge, certificate authorities are not subject to government oversight as a general rule. Sometimes a government agency itself will serve as a certificate authority, and in such cases, it is by definition subject to government control, but that is not necessarily the same as being obliged to conform to a variety of standards.

I have not thought through the pros and cons of government oversight of certificate authorities. One immediate problem is that certificate authorities are international in reach (that is, internet users all over the world may come rely on a particular certificate authority) and yet the CA itself is subject to the jurisdiction of only one country. It would be easy, in principle, to set up CAs in places with weak or no oversight—an analogy is the maritime flag of convenience that allows ship owners to evade regulation associated with states more concerned about maritime safety. Users could choose to not use CAs that are not subject to adequate oversight, but as a rule that would require users to take specific action to do so—and many users would fail to take such action, and wind up trusting untrustworthy CAs.

c. Are there alternatives to the digital certificate model?

1. If so, what are they?
  - ii. If not, how can the current digital certificate ecosystem be improved?

There are a variety of mechanisms that could replace certificate authorities, but all of them have drawbacks as well as advantages. The fundamental problem is that the need for a trust mechanism cannot be avoided, and where trust is involved, trust can be betrayed. One can increase the difficulty of betrayal, but only at the cost of less convenience.

To the best of my knowledge, an authoritative study on the strengths and weaknesses of the CA model and alternatives to it has not been performed; such a study could be performed well by the National Academies. (Full disclosure – I worked for the National Academies for many years, am in a state of semi-retirement from the Academies, and still consult for them from time to time.)

10. In your written testimony you describe how tradeoffs between security, innovation, and convenience are unavoidable.

- a. What is required to achieve consensus on tradeoffs? Is such a consensus possible?

Tradeoffs are hard for people to make. By definition, a tradeoff involves having more of X and having less of Y, when both X and Y are good things to have. The problem arises when you value X more and I value Y more. Making tradeoffs thus involves compromise, in which neither you nor I get as much of X and Y as we could have, and unfortunately we see today that compromise in the policy arena is often regarded as a problem rather than as an approach to a solution.

- b. Is there a way to narrow these tradeoffs, such as by developing a technology that is at once secure as well as convenient? How much more difficult is this kind of development?

The question above embeds an important insight—it is indeed often possible to do better on both security and convenience. But it is hard to do—harder to focus on two attributes

simultaneously than just focusing on one. This will increase the time needed for development—and will potentially delay the arrival of the new technology that is both convenient and secure. If a competitor puts out a product that is convenient and less secure first, it is likely that the company paying extra attention to security will lose in the marketplace. This does not mean it should not be done – only that it will happen less often than would be desirable.

11, In your testimony, you described a two-part goal for reducing threats in cybersecurity. The first is reducing the gap between average cybersecurity posture and the best possible cybersecurity posture. The second is research and development of the best possible cybersecurity posture,

a. Between these two goals, which is more attainable? Why?

The Part 1 gap requires the application of existing technical knowledge for better security. We may lack the nontechnical knowledge that would drive the further adoption of known security technologies and policies, but at least we know what some of these better technologies and policies are. The Part 2 gap is one where we don't even have the technical knowledge, let alone knowledge about how to drive adoption and use. So I think the Part 1 gap is easier to close.

b. Which is more critical to our long term economic success?

I think they are both critical, but I can't make the relative judgment you are asking me to make.

c. Are we making progress on either goal? If so, how and what is driving this change?

We are making progress in the sense that we are better at cybersecurity than we were a decade ago. For the Part 1 gap, for example, the last 10-15 years have seen the study of the economics of cybersecurity become a respectable field of research. Economics is a key driver of where and how cybersecurity technologies are adopted, and we are gaining some insights into the incentives and disincentives for cybersecurity. But taking action on these economic insights remains a problem, for reasons based on the lack of consensus regarding tradeoffs I mentioned

in my testimony. For the Part 2 gap, a variety of technically focused research has resulted in new technologies and approaches to cybersecurity that have some promise. But as I discussed in my testimony, what we ask of our information technology grows at a more rapid rate than our knowledge about how to remediate the accompanying security problems, and so despite these efforts, the gap continues to grow—though not as fast as it would in the absence of these research efforts.

Lastly, keep in mind that the skill and sophistication of the bad guys continues to grow. As I noted in my testimony, they do not simply wait around for the gap to be closed and then go home after it has been closed. They adopt new techniques, find new targets, employ new tactics and technologies for their dirty work—which means that improving cybersecurity is a long-term process rather than a one-time event.

### **The Honorable Markwayne Mullin**

1. It seems like whenever we start talking about the challenges that come with responding to any emerging industry or emerging threat, the issue of workforce development is front and center. With something like the engineering industry, we know we need to engage more students in STEM education, should we be treating the IT industry in the same way?

There isn't an IT corporate executive around who believes that the talent pool for IT workers is sufficiently deep and broad. The basic problem is that the skills and added value that different IT workers bring to the table differ enormously—and innovation in IT is driven by the best of the best, rather than by many workers of average talent working together. If this is true, there is a high premium on creating environments in which the best of the best can be identified and nurtured and persuaded to work in the IT industry.