Gregory E. Shannon, PhD, Chief Scientist

The CERT Division, Software Engineering Institute, Carnegie Mellon University

Questions for the Record:

Hearing on “Cyber Threats and Implications for the 21st Century Economy,” before the Subcommittee on Oversight and Investigations, U.S. House of Representatives Committee on Energy and Commerce, held March 3, 2015

Attachment I-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?*

The continued success and growth in networked systems (aka, the Internet) drives the increases in cyber threats and risks. Design decisions made decades ago continue to reverberate in the cyber-security and privacy challenges we see today and will see as the Internet of Things unfolds in the next five to ten years.

Each entity (gov’t, corporate, individual) must mindfully manage their cyber risks with efficient mitigations. And, many/most entities consistently fail at such mindfulness.

Gov’t should consider how to best enable/encourage entities to be mindful about cyber risks and mitigations. Regulation often is a poor option since it usually encourages compliance rather than engaged consideration of risks.

b. *Ifthere was one fundamental message you want Congress and the public to under stand about cybersecurity, what would it be?*

Mindfully manage cyber risks; demand mitigations that are more efficient.

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

The need for research to be a part of the information sharing discussion. Not only do researchers need access to data to invent better tools, but we also need to include the sharing of vulnerability discoveries (before an attack) as well. Research is also needed to understand how to best digest shared information as well as what information is truly valuable/necessary to ingest.

Furthermore, research needs to be considered in major policy decisions – what is the evidence of likely efficacy for the policy? There are too many policies presented as "good ideas" or "the obvious/right thing to do" but are without supporting evidence that the policy would be followed or have the expected impact (often the desired impact isn't stated…).

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?*

Promote robust tool chains that automatically provide sound security and privacy without developer/programmer action. Encourage the development of such tool-chain ecosystems. The initial steps to improving this area would include defining and implementing a standard definition of "robust" across the relevant industries and standard ways to assess against this definition.

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?*

Companies will be overwhelmed (and possibly fail as businesses) if they don't mindfully manage the relevant cyber risks. Better understanding of which risks to manage and how to prioritize risk, along with budget constraints, will continue to be a huge challenge. We will see the problem of complexity continue to evolve and become more important. Companies will find it difficult to know which systems, vulnerabilities, and threats are important to maintaining the resilience and safety of the IoT. This is not only because of the complexity of the systems themselves, but also because of the rapid pace with which supply chain relationships can change.

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

Piloting new solutions in small scenarios and in situ, especially. By assessing the thoroughness and quality of the processes used to design, implement, and manage the risks of these products over time.

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?*

There are two aspects to assessing a vulnerability, both of which need to be periodically assessed since circumstances can change/evolve (sometimes quickly). (1) How are key assets affected by the vulnerability? Does this vulnerability enable a

path for an adversary (e.g., from HVAC controls to Point-of-sale terminals)?
(2) Are adversaries exploiting that vulnerability or have they shown interest in such the path enabled by the vulnerability?
Yes, the probability has to be weighted, though the more valuable the target the higher the probability. And, adversaries might know of or have imagined paths that the organization's security staff haven't considered.

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?*

It (the Internet) is everything in that everything is connected to it. So, just like electric motors used to be something you bought/used as a separate item (see 20th century Sears catalogues) they now are everywhere and you don't think about them. As the internet becomes more accessible to everyone and more of our lives/things are connected to the internet, it will become invisible only in that we simply won't notice it anymore, much like you barely think about the ability to make a phone call from wherever you are.

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

I agree that "seamless intelligence" is what we'll expect and experience. However, it'll take decades for the Internet to become "invisible," especially because we/society will struggle with new and old security and privacy challenges as the particulars of these technologies evolve. And, there will be failures, some painful, hopefully none catastrophic.

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?*

Companies can certainly mindfully mitigate insider threats. But no, they won't be able to prevent/eliminate them.
No matter how much money a company invests in security software, training, and other cybersecurity measures, they still remain vulnerable to the insider threat due to the fact that trusted insiders are needed for an organization to achieve its mission and to do so, need to be granted authorized access to critical assets. While this is true, it does not mean that an organization is unable to take steps to reduce the likelihood that an insider could cause harm. Many recent high-profile incidents were caused by an insider who intended to cause harm, whether that be an

individual who: stole information from an organization; stole money or defrauded an organization; sabotaged the organization; or disclosed classified information causing harm to the United States. Malicious insiders should be a threat recognized by an organization when building its protection strategies but also they must recognize the threat posed by non-malicious insiders who could cause harm, without intent.

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

Yes and yes. Insiders, including current employees, contractors, and other trusted business partners, to whom an organization grants authorized access to its critical assets, including its facilities, people, technology, and information, do have the ability to harm the organization, but the vast majority do not pose a significant threat because most lack the access and the motivation to cause harm. It is a widely accepted security best practice to limit authorized access to the minimum number of assets necessary for someone to do his job. By doing so, an individual does not pose a threat to everything in the organization. An organization should consider identifying and protecting its critical assets from all threats, both external and internal, with the internal threats being posed by only those with authorized access. By monitoring the asset, anomalies of access and modification can be alerted, triaged, and investigated. But not all insiders are a threat.

To protect against the unintentional insider harming the organization, a combination of technical and administrative controls should be implemented in addition to requiring regular security awareness training. Training should focus on making insiders aware of their responsibility in protecting the organization's assets, including specific techniques, tactics, and procedures used by adversaries to gain access into the organization, allowing potential compromise of its assets. Employees should be made aware of the fact that they can be a target and that targeted social engineering attempts may be made possible because of the information they make publically available.

Carnegie Mellon University's Software Engineering Institute offers multiple options for training and implementing Insider Threat programs in organizations. These programs prepare organizations to meet the intent of Executive Order 13587 -- Structural Reforms to Improve the Security of Classified Networks and the Responsible Sharing and Safeguarding of Classified Information. More information on the various programs can be found at https://www.cert.org/training/ and is available in multiple publications including the Addison-Wesley book, *The CERT Guide to Insider Threats*, and *The Common Sense Guide to Mitigating Insider Threats*.[1]

The lack of validated, efficient mitigations is a significant challenge. However, researchers are looking at mitigation strategies (organizational and technical) that have other benefits to organizations. For example, well-engaged staff very rarely are insider threats (even unintentional threats) due to their dedication to and mindfulness of their organization's mission. So, efforts to increase employee engagement with their work are expected to decrease insider threats from those employees.

[1] Common Sense Guide to Mitigating Insider Threats is found here: http://resources.sei.cmu.edu/library/asset-view.cfm?assetID=34017

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

Organizations must recognize the significant challenge when attempting to calculate and evaluate the cost benefit of security measures. Most organizations are limited by resource constraints, including limited time, people, and money to implement the optimal security solution. In addition, security implementations, if implemented at too high of a level, approaching a 100% solution, may prohibit an organization from achieving its mission. Organizations should strive to implement a solution that protects its assets to the greatest practical extent; identify and choose to accept the risk of implementing less than the optimal solution; and protect the organization from threats that originate from outside and inside the organization, including both those that are malicious and non-malicious.

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 Dr. Lin's testimony?*

Sharing information lets defenders know the current active threat vectors, especially if some threat is active “at scale”. However the gap is only reduced if those receiving the information first understand it and second, know how to effectively take action – which is often not the case. Companies that have high levels of cyber security tend to already be aware of the threat, those who are unaware are also the ones who lack the institutional knowledge (or budget) to do anything about it anyway. So a large and detrimental gap exists in terms of cyber skills and budget that information sharing does not fix.

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

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

Yes and no.

Yes, they’re easy; there are commercial products that support such approaches, and security staff have training on how to use them.

No, savvy adversaries know how to thwart signature-based defenses.

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

It’s a community response, which is important in order to make everyone feel like they own at least part of the problem and the solution.

Information sharing is one way defenders can accelerate dynamic response to threats – in minutes/hours vs. today’s days/months.

While important, it does not truly fix much, considering the vulnerability is the weakest link – a point of entry that would not know what to do with such information even if they had it.

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

Yes, if we can correlate reports with organizational behaviors. We're just starting to see some studies on such data, but that work is typically done by private organizations with special access to a small amount of data.

It can if we allow the research community access to information – to allow for better innovation and cyber security solutions. If we included vulnerability disclosure (the discovery of a vulnerability *before* it is exploited) into the discussion, then that would greatly improve the landscape.

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?*

Not easily.

At the micro level it's very difficult given the lack of data to measure the collectively experienced impact of security practices. However, it is possible for an organization to calculate the economic benefits of using technology versus the risks that it poses to the organization. We (CERT) are working on a model that takes into account the impact of a cybersecurity event on the outcomes of the organization's mission. Using this model results in a dollar amount, or cost, of the impact. This amount can then be compared to the cost of the technology that could detect, protect, respond to, and/or recover from the event.

At the macro level the State Department is looking at this (The Office of the Chief Economist), and I've seen an interesting presentation by Mellissa Hathaway on how cybersecurity investments appear to impact GDP (in an October 2014 talk for OAS titled Lessons Learned in the Design of National Cyber Security Strategies, http://www.iadb.org/en/news/news-releases/2014-10-22/cybersecurity-workshop-for-latin-america,10957.html).

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

Yes, it is possible to calculate the cost of less tangible benefits. These calculations may be organization- or sector-specific, but they can and should be included as part of the overall cost benefit model. As one example, an enterprise risk management program should include determinants of, quantification of, and ways to manage reputational risk when appropriate to an organization.

c. *Is there value in this?*

Absolutely. There is not only value in doing this, but it is critical for some organizations to include these as part of effective risk management programs. While the actual values and categories will differ between organizations, the process for building and implementing a risk management program should include weighted values to compare the benefits of cybersecurity capabilities against the cost.

Cybersecurity and privacy involve significant positive and negative externalities, which are part and parcel of policymaking. We continue to improve our understanding of these dynamics, especially as more and more critical infrastructure is (overly?) connected to the Internet.

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 won't get in, should the emphasis shift to a focus on resilience rather than prevention?*

The emphasis should shift to a focus on resilience – and prevention is but one piece of a resilience program. Given the current state of IT technologies, you have to assume adversaries can/will get in, or are in, your systems.

Resilience encompasses identifying the most critical assets to an organization's mission – these assets can be people, information, facilities, or technology. Once the critical assets are identified, there should be a balanced approach across protection, detection, response, and recovery so that an organization can continue to provide service or meet its mission DESPITE the disruption or cybersecurity event.

So, (1) we need to work to efficiently make "getting in" more difficult and (2) ensure that once "in" it is difficult to significantly disrupt operations. Unfortunately, we have limited efficient mechanisms for either – R&D is needed both by government and the private sector.

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

Resilience is important because we cannot control the ever-changing and evolving threat landscape, but we can control our actions to protect, detect, respond to, and recover from incidents. Resilience is going to be the basis of economic and social survival – we need to avoid fragility and brittleness under "failure" --we must be able to recover. Resilience takes into account events, incidents, and threats not intended to disrupt technology, but the important things connected to technology. This includes not only failures in technology itself, but also the actions of people, failures in process, and even natural disasters that can disrupt organizations. Approaching cybersecurity as another potential operational risk provides better potential to incorporate practices into the organization's risk management process as a means to resilient operations. In recognition of the need to shift to resilience, CERT developed the CERT-Resilience Management Model (CERT-RMM)[2] as a foundation for a process improvement approach to operational resilience management. CERT-RMM is a maturity model that defines the essential organizational practices that are necessary to manage operational resilience. An organization can use CERT-RMM to determine its capability to manage resilience, set goals and targets, and develop plans to close identified gaps. By using a process

[2] http://www.cert.org/resilience/products-services/cert-rmm/

view, CERT-RMM can help an organization respond to stress with mature and predictable performance. Actively used derivatives of CERT-RMM include the Department of Homeland Security's Critical Resilience Review (CRR), used for assessing an organizations cybersecurity practices in ten select domains of practice.

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

Resilience is an advanced form of risk management. Resilience not only takes into account the risks, but it focuses on the impact to the critical few assets so that limited resources can be applied to ensure that an organization can still meet its mission even through disruption. Resilience management enables organizations to transform uncertainties into measurable operational risks and then to efficiently manage those risks while maintaining operations.

d. *In your written testimony, you mentioned the Cybersecurity Capability Maturity Model (C2M2) in your submitted testimony. Can you expand upon this program and how if differs from other options, such as the NIST Framework?*

The Electricity Subsector Cybersecurity Capability Maturity Model (ES-C2M2)[3] is a joint Department of Energy (DOE) and Department of Homeland Security (DHS) effort to help energy sector organizations determine their cybersecurity posture. The ES-C2M2 comprises a maturity model, an evaluation tool, and voluntary DOE facilitated self-evaluations. The model is a collection of cybersecurity practices grouped in ten domains and arranged according to maturity levels. Measures of performance are applied to each domain. The evaluation tool allows an organization to compare its security practices against the criteria of the ES-C2M model. These scores can be compared to target levels of performance to determine gaps in cybersecurity capabilities.

The ES-C2M2 was preceded by the DHS Cyber Resilience Review (CRR).[4] The CRR also measures cybersecurity posture by means of a capability maturity model (CERT-Resilience Management Model[5]), an automated self-assessment evaluation tool, and voluntary DHS facilitated evaluations. The CRR was designed to be applicable to all critical infrastructure sectors, and does not contain the sector-specific tailoring of the ES-C2M2. The CRR also examines the maturity of cybersecurity practices organized into ten distinct domains. In both the ES-C2M2 and the CRR maturity is defined as the institutionalization of cybersecurity practices and processes. Institutionalized practices and processes are more likely to continue to operate effectively during a time of organizational stress (e.g. cyberattack). This examination of maturity differentiates the ES-C2M2 and CRR from more traditional assessments of cybersecurity in which conformance to a standard practice is the only element being measured (e.g. NIST 800-53, ISO/IEC 27001:2013, etc.).

The NIST Framework for Improving Critical Infrastructure Cybersecurity, more

[3] http://energy.gov/oe/cybersecurity-capability-maturity-model-c2m2-program/electricity-subsector-cybersecurity

[4] https://www.us-cert.gov/ccubedvp/self-service-crr

[5] http://www.cert.org/resilience/products-services/cert-rmm/

commonly known as the NIST Cybersecurity Framework (NIST CSF),[6] arranges cybersecurity practices into a of hierarchy of Functions, Categories, and Subcategories. Categories are analogous to the domains of the ES-C2M2 and the CRR. Subcategories are roughly equivalent to the specific practices contained within the domains of the ES-C2M2 and CRR. The NIST CSF is not a maturity model and does not evaluate the institutionalization of practices and processes. The completeness of prescribed practices is the exclusive focus of the NIST CSF. The framework does apply a progression of Implementation Tiers to measure the integration of cybersecurity risk management activities. These should not be confused for the measures of process maturity found in the ES-C2M2 and CRR. Using the NIST CSF does not preclude an organization from also applying the ES-C2M2 and CRR. The evaluation methods can be used in combination. Both the ES-C2M2 and CRR assessment packages include detailed correlation of results to the criteria of the NIST CSF, so an organization can use those assessments to determine if it has met the criteria of the NIST CSF.

8. In Dr. Lin's written testimony he stated that "complexity is the enemy of cybersecurity."

a. *Do you agree with this assessment?*

Yes, given the way cybersecurity is practiced today, complexity does make cybersecurity harder, more expensive, and less effective.
However, I believe that complexity here is more about a lack of understanding about what systems can/should/could do and how adversaries might breach them.

b. *Is it possible to reduce this complexity?*

Yes, by creating technical ecosystems that have security and privacy properties built in (e.g. to tool chains) so that only a few specialists need to fully appreciate the security challenges and have the tools and knowledge to ensure that whole ecosystem is protected. If we can do this, then our adversaries will have to work much harder to disrupt cyberspace.

i. *If yes, what are the consequences?*
ii. *If no, why not?*

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 Lenovo/Superfish revelations. This raises questions as to whether the digital certificate model is providing an adequate level of security for users of the Internet.**

Note: CERT is organizing a workshop for this summer on operational security challenges and opportunities for the global PKI, especially certificate authorities.

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

There are no proofs of correctness for the software that implements these protocols and there’s been limited formal analysis of the protocols as used in practice. Hence,

[6] http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214.pdf

adversaries can readily discover and exploit gaps.

When it comes to authentication and encryption, there are weaknesses with the current PKI model used by SSL. When validating SSL, the trust anchor lies in each certificate authority (CA). There are a few things to keep in mind:

- Your browser or OS chooses the "trusted" CAs, not you.
- Any CA may issue a certificate for any domain.
- The weakest CA determines the strength of the whole PKI.

http://www.net.in.tum.de/fileadmin/bibtex/publications/papers/ssl-landscape-trento.pdf

There are currently more than 100 trusted CAs across modern platforms. For the PKI architecture to work, each one of these CAs must provide due diligence to:

1. Not get hacked (DigiNotar, Comodo)
2. Not get tricked
3. Follow the Certification Practice Statement (CPS) policy that they have published
4. The CPS (and any other certificate issuance and verification processes) must be sound

http://googleonlinesecurity.blogspot.com/2015/03/maintaining-digital-certificate-security.html

PKI centers around the use of trusted third parties. As it is currently implemented, Internet-scale PKI requires trust in all 100+ Certificate Authorities, with little defense should one or more be untrustworthy whether due to error, sloppy business practices, or malice.

When a user visits a site over HTTPS and their browser does not indicate a certificate problem (i.e., when it all works), at best that means that the certificate received was issued by one of the root CAs that is trusted by the browser. Due to the point-to-point nature of SSL and its associated PKI architecture:

- End-to-end encryption is **not** guaranteed.
- End-to-end authentication is **not** guaranteed.

i. *How significant are these weaknesses?*

Significant. Certifications are the foundation of trust transmission on the Internet.

Each of the requirements of CAs outlined above have been violated at some point. That is, CAs have been hacked, tricked, and been found to violate their own CPS policies. The result of these incidents is that users' expectations of encryption and authentication are violated. Traffic that should be protected by SSL could be spoofed, monitored, or altered by an attacker.

ii. *Can these weaknesses be eliminated or adequately mitigated?*

Yes, but it will take a sustained technical R&D investment as the weakness are numerous.

In its current form, the SSL PKI has an architectural design that prevents the weaknesses from being eliminated. There are some things that can help, though (see below).

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

For the most part, no. There is some market pressure from web browser and operating system vendors who require CAs to meet certain standards in order to be included in browsers and operating systems.

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

Software vendors that include the Certificate Authorities in the trusted root CA stores of their respective software are currently the primary oversight. For example, if a CA violates the policy that Mozilla holds them to, then Mozilla can choose to remove the trust in that CA (https://groups.google.com/forum/m/#!msg/mozilla.dev.security.policy/czwlDNbwHXM/amxjB32uY8AJ). Other software vendors such as Apple, Microsoft, and Google have the same control over the certificates that are included in their own trusted CA list. The CA/Browser Forum (https://cabforum.org/) is an organization that provides guidelines that CAs may choose to follow. However, participation in this consortium is strictly voluntary.

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

It is likely that enhanced oversight would improve both the operation of the certificate authorities, as well as the public trust in them. The nature of the oversight matters though. Possibly variants of oversight include:

- standards for CAs

(The CA/Browser forum already has Baseline Requirements: https://cabforum.org/baseline-requirements-documents/)

- penalties for CAs that act negligently

(The FTC fined Kredit Karma and Fandago for claiming to use SSL to secure customer data but not verifying server certificates: https://www.ftc.gov/news-events/press-releases/2014/03/fandango-credit-karma-settle-ftc-charges-they-deceived-consumers)

- audit requirements to demonstrate compliance

Note that many nation-states either directly operate or exert significant influence over trusted Certificate Authorities, which broadens the range of threats that must be considered in determining appropriate oversight.

c. *Are there alternatives to the digital certificate model?*

Yes.

i. *If so, what are they?*

A "web of trust" model like OpenPGP uses is an alternative to the PKI infrastructure used by SSL and TLS. The Monkeysphere project is an example implementation of such a model. It is important to note that a critical mass of adoption must be achieved for a web of trust model to be viable, and this does not appear to have happened with the MonkeySphere project.

Another model is Trust On First Use (TOFU).

Other models integrate somewhat with existing PKI see the following question.

ii. *If not, how can the current digital certificate ecosystem be improved?*

From a technical perspective, there have been some attempts to layer additional checks on top of the underlying SSL PKI, such as the Convergence and Perspectives projects. However, neither of these projects appear to have attained the proper support for widespread adoption and success.

Certificate Transparency (http://www.certificate-transparency.org/) provides an open, public framework that can detect mistakenly issued or maliciously acquired certificates issued by a certificate authority. It can also help discovery of certificate authorities behaving badly (i.e., maliciously issuing certificates). Certificate Transparency is backed by Google and is being developed further in the IETF Public Notary Transparency (trans) working group (http://datatracker.ietf.org/wg/trans/charter/) At present, Certificate Transparency appears to hold the most promise for improvement in the near term.

From an operational perspective, increased transparency with respect to what happens within certificate authorities can help improve the current ecosystem. Given the number of CAs that various software platforms trust, it is currently difficult, if not impossible, for an end user to assign a level of trust to the SSL PKI in general. While some organizations that provide CA capabilities may be trusted by the user, there are a lot of "unknown" CAs that users have likely never heard of and have no ability to judge whether they are performing their due diligence to protect the user's security.

Using DNS (may require DNSSec, which is not widely available)
- http://tools.ietf.org/html/draft-hallambaker-donotissue-02
- https://datatracker.ietf.org/wg/dane/charter/

Certificate Pinning, Google
- https://www.imperialviolet.org/2011/05/04/pinning.html
- https://www.chromium.org/hsts/
- https://tools.ietf.org/html/draft-ietf-websec-key-pinning-21

Google Certificate Catalog (–> Certificate Transparency?)
- http://googleonlinesecurity.blogspot.com/2011/04/improving-ssl-certificate-security.html
- http://www.certificate-transparency.org/

TACK
- http://arstechnica.com/security/2012/05/ssl-fix-flags-forged-certificates-before-theyre-accepted-by-browsers/
- http://tack.io/draft.html

Mutually Endorsing CA Infrastructure (MECAI)
- http://arstechnica.com/business/2012/02/ssl-fix-aims-to-mend-huge-cracks-in-nets-foundation-of-trust/
- https://kuix.de/mecai/mecai-proposal-v2.pdf

convergence/perspectives/observatory
- http://convergence.io/
- http://perspectives-project.org/
- https://www.eff.org/observatory

CERT is planning a workshop to discuss PKI issues, with potential DHS and IEEE support

10. In your written testimony you stated that we currently "'do not know" how to stop all serious cyberattacks while at the same time allowing for the efficient function of electronic commerce.

a. *Why don't we know? Is it that we currently do not have the technological expertise, or that technology has yet to evolve to a mature enough state?*

Both. Efficiency is the key. We know in theory how to make things secure, but the cover/overhead/usability is high, in some cases by many orders of magnitude.

b. *If the cause is that we do not currently have the technological expertise, how do we develop such expertise?*

More research, more training. Expansion of programs like scholarship for services.

c. *If the cause is that the technology is hot yet mature enough, what are the steps we need to take to accelerate that maturity?*

Access to operational data for researchers and measured pilot projects for developers/vendors.

11. In your written testimony you describe the need for "meaningful feedback" for what works if we are to encourage adoption of more effective safeguards or systems of security.

a. *Can you elaborate on what you mean by meaningful feedback?*

Does a practice or technology actually produce or correlate with operational security outcomes (e.g., fewer incidents). Individual organizations usually aren't large enough to measure such effects, but the gov't is if they can collect the data from sectors or nation wide.

b. *Can you give an example of how effective feedback could work in the real world for a small or medium business, given all the complexities and offerings in the marketplace?*

Yes. How much cyber-risk-management training is needed for a key staff member to efficiently mitigate common cyber threats at a Small and medium size businesses (SMB)? 1 Hour? 1 day? 1 week? 1 month? 1 year? Who should it be? If we could correlate the training practices of SMBs with their incident rates, that would then inform owners about what's an efficient investment – instead of just ignoring the problem wholesale. And, maybe the data shows that's there no efficient strategy other than ignore the problem – thought I doubt that.

c. *What progress has been made for developi ng a source of reliable information and measures of effectiveness?*

Limited. Though the NIST cyber risk management framework does offer hope.

d. *Is developing measures of effectiveness even possible in certain areas of cyberspace, given the size and complexities the networks?*

Absolutely yes. The focus has to be on outcomes, not behaviors or compliance. The question isn't "do we patch?", rather the question is "does patching reduce intrusions?" Also, since the adversary is adaptive, what works today might not work next year, so the outcome-based measurement of efficacy has to be continuous – this is one of the real efficiency opportunities for information/incident sharing.

12. In your written testimony you describe the need for medium-term solutions involving "richer data" to improve "situational awareness."

a. *What are the challenges to developing a better sense of situational awareness?*

Currently, access to data – and the privacy and liability concerns that accompany that.

Also, the tools and models to digest the data and present at a cognitively comprehensible view of the situation. A key situational awareness question is, who is trying to make what decisions? Research is continuing in this area. Too often "pretty visualizations" are seen as the answer without considering the decisions to be made.

b. How does improved situational awareness affect the cost of cyberspace safeguards and security practices?

With better situational awareness, less data would not only need to be shared – improving cost, but in theory solutions and tools could arise that allow for quick pinpointing of vulnerabilities and attacks, which would save time and money just in network flow analysis, forensics and response. Subsequently, if response is sooner damage is mitigated.

c. As improvements in awareness occur in some portions of cyberspace, how do we translate that to prevent development of new vulnerabilities as cyberspace technologies expand?

By seeing where adversaries continue to find success, vendors and customers can mindfully and more efficiently respond to systemic issues such as tools chains that produce vulnerable systems.

13. In your discussion of long term needs included in your written testimony, you note that there is no "silver bullet" but there are opportunities to increase the amount of energy requiredred for successful attacks.

This topic is the focus of my current technical research, so the answers are evolving, especially for non-technologists. I'll keep the committee staff informed of my progress.

a. *Can you please expand on what you mean by "energy based" barriers to cyber attacks?*

The goal is to create cyber infrastructure which requires an adversary to use a great deal of computing power (primarily electricity for the computers) to thwart/break a defense and cause a problem at scale (e.g., to the economy).

b. *What would be an example that a layman could understand?*

Encryption is already one example of a technology that we all use every day. Direct attacks on encrypted data are very expensive. This is known as "breaking the key." Only with lots of mega-watt years can you break most keys. Typically, the more long-lasting and central a key is, the key is designed to be hard to break.

c. *From your perspective, what progress is being made on this front and where is it most likely to develop -defense programs, private innovation?*

The ESCAPE workshop in June at CMU is meant to consider our technical progress (http://dimacs.rutgers.edu/Workshops/ESCAPE/announcement.html).

We're seeing progress both privately and with government research investments. There are multiple DARPA and IARPA programs addressing this challenge. And industry is starting to incorporate some of these technologies into their products, services, and business models.

d. *What in your view is the potential for the United States to achieve breakthroughs on this front, versus other nations?*

Very high. (1) we're willing to acknowledge the challenge and make investments.
(2) we have the R&D base to address the challenge.
(3) we can operationalize results through the highly-innovative parts of the software industry that are still predominately located in the US.

This can/should be an allied effort, though it's important the U.S. lead the way.

14. In your testimony, you talked about repositories of "pre-hardened" components such as programming libraries. Specifically, you said that such repositories would allow developers to access componen ts that have been tested and approved, therefore increasing the security and quality of the technologies they design.

Pre-hardening components is an idea that some of my fellow technologists have recommend. However, in my testimony, I was suggesting something (1) more general, (2) more easily adopted, and (3) that could facilitate such hardened repositories. In particular, I'm suggesting tool-chains (software frameworks, application programing interfaces (APIs), compilers, debuggers, editors, verification tools, etc.) that software engineers would use to create the artifacts (binaries) that one actually uses (executes) on a computer. We're already seeing this approach in frameworks, where security experts ensure that non-security engineers will create "safe(r)" artifacts. Those artifacts might be part of a library or it might be the actual application. These tool chains need not be regulated. An alternative approach is for the security evidence to be transparent in that how a tool chain assures a property is well documented, maybe even provable (in a mathematical sense).

One of the challenges with code repositories specifically is that they seem very similar to "code re-use" ideas promoted in decades past. Those "re-use" efforts often failed because of poor governance or business models. Hopefully, we have learned from "re-use" successes and failures.

a. *How would such a repository be created?*

These tool chains can be privately held by individual vendors, consortia, the government, as open source, etc. The governance of the tool chain and any repository is orthogonal to the technical security capabilities that it provides. A useful first public experiment might be SSL and PKI libraries and development tools since public key certificates are the foundation of security on the commercial Internet.

b. *What organizations, both public and private, would need to be involved?*

Public: NSF, NIST, DHS S&T, NSA-R, ASD(R&E)
Private: IEEE, ACM, ANSI, probably ISOC/IAB/IETF
Commercial: Key infrastructure providers willing to directly facilitate the approach

i. *Who would be responsible for such a repository?*

An existing or new non-profit, at least for the first steps.
Ideally open-source. A governance model (new or existing) would have to be worked out.

c. *Would a repository such as this be expensive to create and manage?*

i. *If so, how could those costs be managed?*

Initially as a community experiment; government could provide some funds/incentives, but involvement is voluntary.

d. *Who would be responsible for testing and approving the components that are made available through the repository?*

i. *Does such a repository raise liability concerns, if a "tested and approved" component is later found to be deficient?*

The model should be “tested with verifiable evidence reported”, that way no new liabilities are created for participants. The liability onus could remain on the vendors who incorporate the software into their products.

ii. *If so, how could the repositories address and account for that liability?*

Depends on the governance, and actions gov’t(s) could take.

e. *How exactly would developers use this repository?*

Download the artifacts and the accompanying tool chains and proceed in the context of their organization’s development practices.

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?

Yes, thought I thought we were (STEM and IT both are about technology, yes?).

IT is clearly a core enabling capability for operationalizing "STEM" innovations.
One challenge is that IT workers often have low status in their organization.
By raising the recognition of and respect for such workers, more students might see it as a rewarding occupation. Organizations need to highlight when IT staff have facilitated gains in productivity (and profits).