From: Joe St Sauver, Ph.D., PO Box 3504, Eugene OR 97403 (jstsauver@gmail.com)
To: Diane Honeycutt, National Institute of Standards and Technology,
100 Bureau Drive, Stop 8930, Gaithersburg MD 20899 (cyberframework@nist.gov)
Date: April 8th, 2013
Re: "Developing a Framework to Improve Critical Infrastructure Cybersecurity"

Dear Ms. Honeycutt:

Thank you for the opportunity to comment on NIST's recent solicitation regarding "Developing a Framework to Improve Critical Infrastructure Cybersecurity," www.gpo.gov/fdsys/pkg/FR-2013-02-26/pdf/2013-04413.pdf

I'd like to begin by commending NIST and the Department of Commerce for undertaking this crucial work. Securing critical infrastructure in cyberspace needs to be a top priority for America, but unfortunately there's been all too-little progress to-date, either here at home, or abroad (e.g., in the countries of our trading partners & allies).

My interest in the cyber security of critical infrastructure dates back a decade, and can be seen in things such as my invited talk for FBI InfraGard on SCADA security and critical infrastructure.[1]

Although I am a cyber security practitioner and involved with many cyber security-related organizations and projects, including having served on a recent FCC CSRIC botnet working group, due to the compressed timeframe and the sometimes contentious nature of questions relating to this topic, and to keep this simple, I'll be offering my comments today in an individual capacity, and not on behalf of any organization or other entity.

I. DEFINITION OF CRITICAL INFRASTRUCTURE AND RELATED IMPLICATIONS

(1) The Generic Definition Of "Critical Infrastructure:" As noted on page 13024 of the Federal Register solicitation, your focus is on "reduc[ing] cyber risks to critical infrastructure," with "critical infrastructure" defined in footnote 1 as having the meaning given in 42 U.S.C. 5195c(e), to wit:

> "systems and assets, whether physical or virtual, so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, national economic security, national public health or safety, or any combination of those matters."

How that definition of "critical infrastructure" maps, *in practice*, to decisions about what is and isn't in scope is key to the remainder of this discussion, so let's begin there. Most would assume (probably correctly) that civilian "critical infrastructure" includes things like:

- Energy production and delivery infrastructure (nuclear power plants, hydroelectric dams, gas pipelines, coal-fired power plants, petroleum refineries, the electrical power grid, etc.)
- Food production infrastructure (including things like major farms and feedlots/stockyards, grain elevators, seed and agrichemical supply companies, food processors, etc.)
- Potable water collection and distribution (aquifers and reservoirs, water purification facilities, water distribution pipelines, etc.), as well as water used for industrial processes and irrigation
- Wastewater/sewage collection and treatment, as well as sanitation services

[1] "SCADA Security and Critical Infrastructure," http://pages.uoregon.edu/joe/scadaig/infraguard-scada.pdf

- Emergency services (police/sheriff, fire, ambulance service, and other first responders), as well as prisons and other correctional facilities, as well as custodial facilities for the criminally mentally ill

- Medical infrastructure (hospitals and medical centers, pharmaceutical production and distribution, blood banks, etc.)

- Transportation and logistics (airports, passenger and cargo airlines, railroads, the highway system and long haul trucking, strategically important bridges and tunnels, container ports and cargo ships, postal service and private package delivery companies, etc.)

- Manufacturing facilities (including both legacy traditional manufacturing facilities, such as chemical plants, as well as cutting edge electronic facilities and defense industrial plants, etc.)

- Commercial/retail distribution infrastructure (warehousing and distributors, POL tank farms, etc.)

- Financial infrastructure (commercial and investment banking, the stock exchanges, the payment card industry, clearing and settlement facilities, major retail banks, etc.)

- Communications infrastructure (physical network infrastructure, phone switches, cellular infrastructure, carrier hotels, major colocation facilities, transoceanic cable landing sites, etc.)

That said, I am also of the belief that there is at least one DHS-identified critical infrastructure sector that frankly does *not* actually warrant that designation, at least when it comes to the 42 U.S.C. 5195c(e) definition: "National Monuments and Icons," while historically/symbolically important, are not critical to our national survival.

I was pleased to see that Section 9 of the recent Whitehouse EO on "Improving Critical Infrastructure Security,"[2] explicitly excludes "*commercial* information technology products or *consumer* information technology services" [emphasis added] from being considered to be "critical infrastructure," at least for the purpose of that EO. **For the avoidance of doubt, I'd like to see similar exclusions carried forward in any critical infrastructure cyber security framework defined by NIST.**

The actual language used should also be clarified/expanded. For example, "*commercial* information technology products" are excluded from the Executive Order's definition of "critical infrastructure," but does that also include major *free/open source* information technology products, too? If not, something like the freely downloadable and widely used open source name server package BIND[3] could potentially end up being considered part of what's "critical infrastructure." **Is free/open source software excluded from consideration as critical infrastructure, or not? I believe it *should* be excluded from discussions moving forward.**

Similarly, I'd note that the EO explicitly excludes *consumer* information technology services, but are *enterprise* information technology services still potentially in scope for consideration as "critical infrastructure"? If so, virtually every company in America might be faced with the daunting prospect of potentially operating "critical infrastructure" simply as a result of running a normal corporate email systems or other routine business systems. Or as another example, would *university* mail servers and networks also potentially remain in-scope as potential "critical infrastructure?" They're not "*consumer* information technology services" after all, are they? **Therefore, I'd also urge NIST to also consider clarifying whether "enterprise IT services" and "university IT services" remain potentially in scope as cyber critical infrastructure, or not. (I believe that they should be *out of scope*.)**

---

[2] http://www.whitehouse.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity

[3] https://www.isc.org/software/bind

Ambiguity when it comes to what is and isn't in scope is a source of great anxiety for those who may potentially end up being responsible for operating and protecting "critical infrastructure." Ultimately, if companies cannot clearly tell what is genuinely "critical infrastructure" and what is not, we run the risk of losing focus and ignoring that which is truly critical while we attempt to deal with what may frankly be inconsequential (or peripherally important) things. **What is and isn't critical infrastructure MUST be operationally better-defined.**

(2) Identify and Prioritize The Most Urgent Critical Infrastructure Areas: Noting the large number of areas that are likely genuinely critical infrastructure, as enumerated on pages one and two of this response, it is clear that critical infrastructure needs to be prioritized so that the most urgent critical infrastructure with the greatest vulnerabilities/the greatest need (or the greatest potential negative impact) can be addressed first. **Please be explicit in identifying and PRIORITIZING the most urgent critical infrastructure areas with cyber security vulnerabilities.** "Everything" can't be critical infrastructure and everything can't be a "top" priority. If virtually everything is listed as "special," in effect that "special" category becomes the new "average" or "normal" category, and a pointless designation.

"Over inclusivity" when it comes to "critical infrastructure"/potential terrorist targets is not a new problem. See for example the July 2006 article in the *New York Times*, "Come One, Come All, Join the Terror Target List,"[4] reporting on a DHS Inspector General's report that found that sites nominated for inclusion in a federal antiterrorism database reportedly included things like "Old MacDonald's Petting Zoo, the Amish Country Popcorn factory, the Mule Day Parade, the Sweetwater Flea Market and an unspecified 'Beach at End of a Street'," apparently in an effort to maximize local eligibility for future federal counter-terrorism funding.

A sharp eye needs to be kept on what gets designated as being "critical cyber infrastructure assets" -- and why. In particular, **I would urge you to adopt a target value for what fraction of all cyber infrastructure should be designated as "critical infrastructure," limiting that to no more than 5% of any class of cyber assets.**

(3) At Least In Telecommunications, Criticality Will Often Be A Function Of How Things Are Designed and Built, and How Customers Use Those Resources; Going for "Uber Availability" Will Come At A Cost: I am particularly interested in the telecom critical infrastructure sector, so in addition to my comments on critical infrastructure in general, in the preceding part, let me also specifically comment on telecom critical infrastructure cyber security.

In an ideal telecommunications network, nothing in that hypothetically perfect telecommunications network, *considered in isolation*, should actually qualify as "critical infrastructure." When properly designed and carefully built, you should be able to lose any individual server or any arbitrary network link or node, and have that outage be virtually unnoticeable due to engineered-in redundancy and resilience. Thus, at least for random (rather than intentional) outages, traffic should automatically and seamlessly route around down or damaged facility. That ability to route around damage, so that a network can continue virtually seamless operation, relies on redundancy and sufficient spare capacity to temporarily accommodate rerouted loads.

In the real world, however, if networks or systems are allowed to be built with single points of failure, or operators run with too little spare capacity, systems and networks may become brittle, and practical resilience may be lacking. Thus, in the real world, or in the face of intentional attacks, failures may still arise. The frequency with which incidents occur, and the time it takes to recover from them, is ultimately a statistical risk management question: just how rare an incident do you want to be able to ride out?[5]

If you go far enough "out in the tail" of the distribution of incidents, there is no question that you can always find some extreme scenario that will cause even the most robustly architected system to fail for some period of time. For example, it is not unheard of for both primary connectivity, and diverse redundant fallback paths, to be

---

[4] http://www.nytimes.com/2006/07/12/washington/12assets.html?_r=0

[5] This is much like doing flood planning. For example, are you willing to build in the 10 year flood plain? Or would you rather be more conservative (and potentially pay more) to build in a more desirable area that is less at risk of flooding, perhaps in the 50 year or 100 year flood plain, instead?

intentionally sabotaged by knowledgeable insiders during an acrimonious labor dispute, thereby causing an outage that would normally not occur.[6]

In the telecommunications industry, great attention is paid to weighing the likelihood and consequences associated with various potential failures so that commercially prudent measures can be employed to mitigate those risks when it is financially rational to do so,[7] and SLAs (Service Level Agreements) make sure that there are financial consequences if contractual obligations aren't fully met.

That said, we recognize that in some cases, while telecommunications infrastructure may be "critical" for some users, it may be distinctly "non-critical" for others. Telecom customers normally deal with this through what they purchase. For example, if Internet access is genuinely critical to a customer's business, a company may multihome (connect to the Internet through multiple upstream ISPs), and purchase multiple physically separate local loops (physical circuits from their location to their ISPs' points of presence), rather than relying on a single ISP and a single local loop for all their connectivity.

Getting that **improved reliability always comes at a cost.** As the saying goes, "You can get as many 9's [e.g., as much reliability] as you're willing to buy." In the example we mentioned in the preceding paragraph, imposing the complexity and cost associated with multihoming (and the use redundant local loops) on ALL customers wouldn't make sense from a telecommunication customer's point of view, at least if you're the sort of customer who can leverage the Internet when it's available, but can live without it (at least temporarily) when it isn't. In that case, accepting an outage will virtually certainly be less expensive than paying to avoid it.

This can be explicitly seen in things like GSA Networx pricing for Internet service for federal agencies.[8] For the purpose of this discussion, let's consider federal pricing for "regular" and "critical service level" 1 Gbps connections to the Internet, via Ethernet, here in the continental US:

-- CLIN 744084 (Internet Protocol Service Ethernet, 1 Gb/s, continental US, monthly recurring cost), available for as little as $8,424.32/month from one GSA provider (Level3)

-- CLIN 744253 (Internet Protocol Service Ethernet, 1 Gb/s, CRITICAL SERVICE LEVEL, continental US, monthly recurring cost), where the cheapest GSA provider option is $25,032.80/month (from Qwest)

Dividing $25,032.80/$8,424.32 we can see that getting "critical service level" support for the same level of connectivity costs almost exactly three times as much as just basic connectivity. That's a substantial premium.

I **cannot stress enough that all efforts aimed at improving telecom resilience/survivability come at a cost,** sometimes, as in this example, at quite a substantial cost. Imposing costs of this sort of magnitude on genuine critical infrastructure that relies on telecom may be appropriate and unavoidable, but the magnitude of these costs nicely underscores why overbroad designation of "critical infrastructure" may be economically toxic to the American business community.

PLEASE try to be very restrained in what you designate as "critical infrastructure", particularly in telecommunications, because whatever you designate as "critical" will immediately become substantially more expensive than if it wasn't so-designated.

---

[6] See for example, "California Telecom Knocked-Out By Low-Tech Saboteur," http://www.redorbit.com/news/technology/1669300/california_telecom_knockedout_by_lowtech_saboteur/

[7] In saying this, yes, that implicitly means that sometimes the right decision is simply to accept the risk, simply because the cost of mitigating a particular risk may far outweigh the expected loss associated with accepting that vulnerability.

[8] http://www.gsa.gov/portal/category/25318

II. WHAT IS THE "CYBERSECURITY FRAMEWORK FOR CRITICAL INFRASTRUCTURE?" IS IT PRACTICALLY DO-ABLE?

On page 13025 of the Federal Register, the solicitation notes that,

> Given the diversity of sectors in critical infrastructure, the Framework development process is designed to initially identify *cross-sector* security standards and guidelines that are immediately applicable or likely to be applicable to critical infrastructure, to increase *visibility and adoption* of those standards and guidelines, and to find potential *gaps* (i.e., where standards/guidelines are nonexistent or where existing standards/guidelines are inadequate) that need to be addressed through collaboration with industry and industry-led standards bodies.

The solicitation then elaborates that the Framework will:

> [...] incorporate voluntary consensus standards and industry best practices to the fullest extent possible and will be consistent with voluntary international consensus-based standards when such international standards will advance the objectives of the Executive Order. The Framework would be designed to be compatible with existing regulatory authorities and regulations. The Cybersecurity Framework will provide a prioritized, flexible, repeatable, performance-based, and cost-effective approach, including information security measures and controls to help owners and operators of critical infrastructure and other interested entities to identify, assess, and manage cybersecurity-related risk while protecting business confidentiality, individual privacy and civil liberties. To enable technical innovation and account for organizational differences, the Cybersecurity Framework will not prescribe particular technological solutions or specifications. It will include guidance for measuring the performance of an entity in implementing the Cybersecurity Framework and will include methodologies to identify and mitigate impacts of the Framework and associated information security measures and controls on business confidentiality and to protect individual privacy and civil liberties.

While that's an admirable vision, it likely isn't one that's likely to be able to be accomplished.

(1) Cross Sector Security Standards/Guidelines

In thinking about a potential "Grand Unified Theory" for critical infrastructure cyber security, which is basically what it sounds like you're seeking, I'd note that other disciplines have historically tried to find similar underlying universally applicable frameworks.

For example, there is a management theory that asserts that a good manager can manage any type of business, whether that business is bottling beer, selling commercial paper in the financial markets, running an opera company or overseeing the construction of government facilities overseas -- the principles of good business are (theoretically) universal and domain-specific expertise really isn't needed, or so the "universal process approach" theory goes.[9] We detect a similar perspective in this call to create a universal, or overarching, cross sector framework for critical infrastructure cybersecurity.

We believe that approach is likely a recipe for failure.

When it comes to cybersecurity for critical infrastructure, there are important differences between the various critical infrastructure sectors, differences driven by the unique engineering constraints associated with each subject matter domain. **It is utterly essential that domain-specific unique engineering constraints be gathered, understood, and respected, since in many cases they were discovered and have evolved as a matter of life safety, or in order to avoid serious and expensive real world financial impacts.**

---

[9] college.cengage.com/business/kreitner/foundations/1e/students/appendices/appendix_a.html#The%20Univer

Does this mean that there are absolutely now universally applicable cyber security principles for critical infrastructure? No. There are a few general principles that can probably be applied cross sector, but those things are so basic and so simple as to be rather underwhelming as a potential core catechism for this work:

i) *Uptime is critical.* If we consider the three objectives of the classic cyber security C-I-A trinity, *availability* is the dominant consideration in critical infrastructure cyber security, far more so than *confidentiality* or *integrity*.

ii) If you want availability, *build in redundancy and extra capacity* throughout to avoid any potential single point of failure and to ensure that you can safely accommodate loads shifting from part of your infrastructure that may be down to backup capacity. Test these mechanisms to ensure that they actually work as designed. Recognize that this enhanced availability service may be costly to deploy.

iii) *Recognize that system lifetimes are unusually long.* If you're building a plant today with a thirty year lifespan, the control systems you install during construction may be the same control systems you're using in 2043. That is dramatically different than the enterprise or consumer world, where you might cycle through ten different laptops or fifteen different smart phones, or more, during that same thirty year period.

iv) *Never connect your critical infrastructure network to the public Internet, or to your enterprise corporate LAN.* That said, always assume that your "non-interconnected network" will somehow become accessible to unauthorized parties, perhaps as a result of an unknown and insecure wireless access point that someone informally added to the critical infrastructure network without authorization. Remember that hackers can and will also try simple but effective expedient tricks such as dropping malware-infected USB thumb drives in a facility's parking lot. If remote access must be offered, require use of a VPN and multifactor authentication, strictly limit the IP address ranges that can successfully access that channel, and log all access closely.

v) Don't discount the importance of *physical security controls*. We *are* talking about *infrastructure*, after all! If an attacker can obtain physical access to your systems or network, that's often all they'll need to be able to successfully attack that infrastructure.

vi) *Humans play critical roles in critical infrastructure systems*, both as operators and as potential insider threat actors. Train personnel carefully, and ensure you have an effective personnel reliability program in place.

vii) Design so that if/when you fail, you *fail safe* (e.g., you fail into a default safe/secure condition). For example, in a nuclear plant, the default response to many potentially serious incidents is to do an emergency shut down of the reactor (insert the control rods to make the reactor sub-critical, or "SCRAM" the reactor).

viii) *Logs* can provide vital details about incidents -- if you collect them and someone actually looks at them.

ix) There are some *low probability/high impact potential threats* that may involve cyber critical infrastructure, such as electromagnetic pulse (EMP) attacks.[10] Mitigating these threats can be done via relatively modest incremental investments, if management is aware of the vulnerability and makes mitigating these issues a priority. In all too many cases, though, unfortunately "it hasn't happened yet" gets wrongly generalized to be "it *can't* happen, ever." We cannot let a failure of imagination leave us unprepared for high consequence cyber security events.

x) Recognize that control system security may be subject to *unique hard constraints* that typically will not be present in the enterprise or consumer environment. For example, a real time operating system, as used by infrastructure, may have strictly scheduling limits and latency bounds that a consumer or enterprise PC doesn't. At the same time, *some control system environments may also be effectively immune to some threats* that continually threaten consumer or enterprise PCs -- for example, when's the last time you heard of malware targeting QNX, a popular control system operating system, eh? It just doesn't happen.

---

[10] "Electromagnetic Pulse," http://pages.uoregon.edu/joe/infragard-2009/infragard-eugene-2009.pdf

(2) The Business Case Issue: As you formulate your critical infrastructure cyber security framework, please remember that attending to cyber security isn't just a matter of technology, it's also a matter of economics.

Remediating vulnerabilities isn't costless for critical infrastructure operators. It costs money to hire cyber security specialists, and to invest in hardware and software tools. In a competitive environment where management is rewarded based on its ability to maximize shareholder value, investing in discretionary security measures may not make much sense if the rest of the industry isn't doing likewise.

Whatever you propose either must make objective financial sense, or be subsidized at least to the point of being cost neutral, or be a uniformly imposed requirement. Voluntary measures that don't make economic sense will simply be ignored by industry.

(3) Practical Operational Security vs. Compliance Activities: There is an inescapable tension between practical operational security and compliance activities.

As I mentioned in my February 2nd, 2013, comments on the Kingdom of Saudi Arabia's seventh draft report on "Developing a National Information Security Strategy For the Kingdom of Saudi Arabia,"[11]

> In every country there is a tension between operational security (as practiced by technical cyber security staff, including network engineers, system administrators, and other people that are sometimes summarily referred to in the West as the "geeks" or "techies"), and regulatory compliance (an activity largely the province of attorneys and compliance specialists/report writers, sometimes dismissively referred to as "suits" by Western geeks/techies).
>
> Obviously, resources (executive sponsorship and oversight, budget, personnel positions, etc.) dedicated to regulatory compliance won't be available for operational security, and vice versa. Thus, it is critical to get the balance between these two competing priorities right. In at least some cases, I've seen evidence that some governments, including the government of the United States, have put too much emphasis on regulatory compliance and too little on actual operational security. Quoting the Congressional testimony of Alan Paller of the SANS Institute,[12]
>
> > *Four terribly damaging processes were institutionalized in the aftermath of FISMA and GISRA (the Government Information Security Reform Act that predated and is essentially the same as FISMA). These wasteful processes slowed down our defenses and threw away billions of dollars that were acutely needed to protect systems. They forced federal chief information officers to defer investments in enterprise security because their security budgets were being consumed buying 3‑ring binders full of reports that were out of date when delivered and had no discernible impact on security. To implement GISRA and FISMA, the government created an audit process that regularly results in misleading reports to agency heads and Congress. That flawed process was adopted by the Inspectors General, as well, who also are producing reports that answer the wrong questions. GISRA and FISMA rewarded ineffective behaviors and created a cadre of people who call themselves security professionals but who proudly admit they cannot implement security settings on systems and network devices or find a programming flaw. Most of these paper-warriors have no depth of understanding of current threats, cannot do an effective risk assessment, nor select the right controls to protect systems against the increasingly sophisticated attacks. If the federal government were the only organization being impacted by the FISMA‑flaws, that would be bad enough. But*

[11] http://www.mcit.gov.sa/NR/rdonlyres/514E7B51-5710-46D9-9EC5-2D78BC2E1219/0/NISS_Draft_7_EN.pdf

[12] "Testimony of Alan Paller, Director of Research, The SANS Institute, Before the Subcommittee on Government Management, Organization, And Procurement of the Committee on Oversight and Government Reform, Hearing on "Federal Information Security: Current Challenges and Future Policy Considerations" March 24, 2010, PDF page 2. http://oversight.house.gov/wp-content/uploads/2012/01/20100324Paller.pdf

> *increasingly state governments, radically short of money, are being forced to spend scarce funds on reporting rather than security. Even worse, the electric power industry has been caught up in the culture of compliance created by FISMA. The head of security at a major southern power company told me last Friday, "I had to hire a writer rather than a security person because writing compliance reports is seen by management as more important than actually securing the systems." FISMA has perturbed the entire security job market. In the federal contractor community, writers who know a few words about security and federal regulations now make 50-80% more money than the people who actually secure systems and networks and applications. It is as if we paid the compliance staff at a hospital more than we pay surgeons. The best and brightest technical people are being forced into compliance roles because they want to keep their jobs and earn more money.*
>
> *[* * *]*
>
> *Both the guidance for implementing FISMA and the guidance for auditing compliance are focusing on out of date, ineffective defenses. What we need instead is a process that directs agencies to focus their cyber security resources on monitoring their information systems and networks in real time so that they can prevent, detect and/or mitigate damage from attacks as they occur. And oversight must be focused on the effectiveness of the agencies' real-time defenses.*

My concern is that the new NIST critical infrastructure cyber security framework (like the Saudi Arabian draft NISS), has the potential to degenerate into yet another regulatory compliance regime, diverting resources and attention from actual operational cyber security issues. **Please, I urge you, devote the majority of your efforts toward actually improving operational cyber security for critical infrastructure, relegating compliance activities to a secondary/supporting role to the extent that resources are available for that supporting priority.**

To understand what I mean by "actual operational cyber security activities," consider "Project SHINE," a project undertaken by Bob Radvanovsky and Jake Brodsky of InfraCritical.[13] Their effort used the Shodan search engine to identify about 7,200 U.S. control systems directly connected to the Internet with weak, default or non-existent login credential requirements, systems that any malicious actor could also potentially have identified and abused. Those systems, and similar systems abroad, are now the focus of notification and remediation efforts by the DHS Industrial Control Systems Cyber Emergency Response Team. That project is a perfect example of the sort of critical infrastructure cyber security work that will actually result in a tangible improvement in critical infrastructure operational systems. We need more of this sort of work, and less in the way of compliance report writing.

(4) Dead End Trails That Should Not Be Pursued: At the same time I want to emphasize the importance of empirical work that results in tangible improvements, like the Radvanovsky and Brodsky study mentioned in the previous paragraph, I want to also call out some areas that are, I believe, dead end trails. Please try to avoid wasting your time on these areas:

-- *Attempting to Import A Traditional Risk Management Approach For Use in Critical Infrastructure Cyber Security:* Traditional risk management, which often serves as the foundation for all subsequent cyber security activity in the enterprise security space, will fall apart in the critical infrastructure space. Why? Risk management relies on a historical database of incidents and experience to empirically estimate the likelihood, and the likely impact, associated with any given security exposure. For example, why do banks not issue two factor authentication (2FA) tokens to U.S. bank customers, customers who are likely to be phished? Answer: banks don't issue 2FA devices to customers because banks know that it is cheaper to just "eat" occasional phishing loses, rather than paying for 2FA tokens and their support, or losing irritated customers frustrated at having to jump

[13] "Thousands of Industrial Control Systems at Risk," *Information Week*, January 11th, 2013, http://www.informationweek.com/government/security/thousands-of-industrial-control-systems/240146091

through additional security hoops. Banks have the empirical data, collected over time, that allows them to make those judgments in a rational way. Critical infrastructure cyber security does not. The number of critical infrastructure cyber security incidents is still miniscule in comparison to consumer or enterprise cyber security incidents. Even if you were determined to adopt a risk management framework, you simply don't have the data to do so in a uniform way.

Critical infrastructure cyber security also has 2nd order effects that would be difficult for a critical infrastructure provider to properly assess in attempting to do risk management. For example, consider the 9.0 Tohoku Japan earthquake of April 11th, 2011. As I noted in my April 19th, 2011 talk,[14] early estimates were that it might cost up to 25 trillion yen (roughly USD 309 billion) -- but those costs ***do not*** include lost economic productivity due to power outages, or the broader impact of the nuclear crisis. Given the magnitude of those direct and second order costs, clearly the "risk management" paradigm badly failed the citizens and corporations of Japan, and of the world -- had the Japanese known just how bad Tohoku was going to be, surely they would have done more to try to prevent it.

-- *Attempting to Push Traditional Enterprise "Solutions" To The Control System/Critical Infrastructure Space:* Another example of a rather pointless direction that some might consider would be a potential push for "cyber hygiene" as part of critical infrastructure cyber security.

Remember that in many cases, even in the enterprise space, antivirus isn't working well, since the bad guys/gals can create new malware variants, or repack existing malware in new hard-to-detect ways, faster than the good guys/gals can create and test new antivirus "signatures." With antivirus struggling (if not in its death throes) in the enterprise space, why would it make sense to try moving this failing paradigm to critical infrastructure? Answer: it doesn't. Antivirus is increasingly like a young child's favorite (and just about worn out) blanket: it may provide psychological solace, but it has so many holes it doesn't really provide much protection against the cold hard world. (On the other hand, a program centered around whitelisting only approved applications might make a lot of sense in this sort of space)

Similarly, some might suggest a program of aggressive patch management, but remember that many devices in the control system space are inaccessible or nonupgradeable, and any patches that are potential candidates for application may need to be thoroughly vetted and approved by vendors and/or regulators before they can be applied -- assuming the vendor is even still in business twenty or thirty years after a plant was built, and assuming that a twenty or thirty year old device can even run code that may be orders of magnitude more complex than what was common when it was produced.

-- *Tools and Metrics:* Lastly, I suspect that there may be a push to collect tons of statistics, if only to validate the presumably dire condition of the industry, and of course, once we know things are bad, we will need a new "magic bullet" or "miracle tool" (rather than diligent hard work) to magically make everything right.

When it comes to *statistics*, I'm generally supportive of collecting metrics about a problem, but please, *collect metrics passively*. Don't send out surveys or demand new reports asking people to guess about conditions they may not be adequately monitoring in the first place. Use existing monitoring capabilities of other U.S. government agencies to directly provide the administration with detailed information about cyber security threats to critical infrastructure. (E.G., you likely *already have* the data you might be asking for from other sources, you just don't know it, or perhaps you don't have access to it due to its classification.) Be sure to also put any statistical values in comparative perspective -- one of the things we saw in the botnet metrics report for CSRIC III WG7[15] was that often times people have a hard time interpreting statistics unless they can see how the US is doing to other countries, whether the comparative cohort is the Five Eye countries, the G-8, or all the countries of the UN.

---

[14] "The Tragedy In Japan: A 9.0 Earthquake, Massive Tsunami, Large Scale Radiological Contamination, …," http://pages.uoregon.edu/joe/japan-tragedy/japan-tragedy.pdf

[15] http://www.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG7_Report_March_%202013.pdf at pp. 62 and following.

When it comes to *tools*, *before endorsing the creation of a "new" tool, please do your best to ensure that there isn't an existing tool that can already do the job*. A classic example of this may be the federal Einstein network monitoring project -- I've yet to read a compelling justification in the open literature for developing and deploying Einstein monitoring boxes rather than using (or extending, if necessary) existing open source monitoring tools (such as Snort[16] or Bro[17]), or using commercially available network monitoring tools.

I know there's a strong allure to fielding your own proprietary tool, but you lose a lot by not taking advantage of the work that's already been done, and which continues to be done, on existing products.

As a second example in the tool space, relating to sharing information about cyber security incidents, is SES/CIF. I hope that every effort will be made to exploit (or extend) existing tools, such as SES/CIF,[18] which is already in wide use in trusted information sharing circles such as the Research and Education Network Information Sharing and Analysis Center (REN-ISAC),[19] before any new information sharing tool gets undertaken.

Encouraging use and/or refinement of an existing and developed tool of that sort makes far more sense, and will likely be far more productive sooner, rather than encouraging the creation of Yet Another Tool. (Let's all beware of the perils of the "Not Invented Here" condition)

(5) A Multinational Approach Must Be Pursued: John Donne wrote that "no man is an island." Well, so, too, no country's critical infrastructure is an island. We must work together with our friends and allies to ensure that our trading partners also have secure cyber critical infrastructure. The days when America is, and could be, self-sufficient without our international business partners are long gone. If we just gaze introspectively at our navel and ignore our neighbors, we may find that our critical infrastructure has failed not because of a direct attack on our own networks or systems or plants, but because of attacks on third party international networks or systems or plants upon which we nonetheless depend. We can't just worry about American critical infrastructure, we must take a multinational approach appropriate to the world's high degree of interdependency. An attack on Canada's cyber critical infrastructure or Germany's cyber critical infrastructure or Japan's cyber critical infrastructure could be just as devastating as an attack on our own. We need to work together to be successful in this critical effort.[20]

III. CONCLUSION

I believe NIST faces some formidable challenges in trying to create a framework for critical infrastructure cyber security, starting with some key issues relating to what is and isn't designated as "critical infrastructure." The goal of creating a single overarching grand unified theory of critical infrastructure cyber security is also fundamentally at odds with the reality of many unique critical infrastructure sectors, each with its own unique engineering constraints, and each of which likely needs to be treated as a "ship on its own bottom," rather than hoping that a rising cyber security tide will somehow magically raise all critical infrastructure "boats" equally.

I'd be happy to work with NIST in any way that might be helpful when it comes to this area. Please don't hesitate to get in touch if I can be of assistance in any way.

Sincerely,

/s/
Joe St Sauver, Ph.D.

---

[16] http://www.snort.org/
[17] http://www.bro.org/
[18] http://code.google.com/p/collective-intelligence-framework/ (For some early information about SES/CIF, see Wes Young's SES presentation from the first Data Driven Collaborative Security Workshop for High Performance Networks, May 2009; the agenda page is linked from http://security.internet2.edu/ddcsw/)
[19] http://www.ren-isac.net/
[20] A nice example of an international perspective on cyber security issues can be seen in things like MAAWG's recent outreach to India, see http://www.maawg.org/system/files/M3AAWG_India-Brochure-2012-05.pdf