

CYBER SECURITY SERIES
FEBRUARY 2013

Foreign Policy
at BROOKINGS

Bound to Fail: Why Cyber Security Risk Cannot Simply Be “Managed” Away

Ralph Langner and Perry Pederson

Bound to Fail: Why Cyber Security Risk Cannot Simply Be “Managed” Away

Ralph Langner and Perry Pederson

Executive Summary

Rather than a much-needed initiative to break the legislative deadlock on the subject in Congress, President Obama’s new executive order for improving critical infrastructure cyber security is a recipe for continued failure. In essence, the executive order puts the emphasis on establishing a framework for risk management and relies on voluntary participation of the private sector that owns and operates the majority of U.S. critical infrastructure. Both approaches have been attempted for more than a decade without measurable success. A fundamental reason for this failure is the reliance on the concept of risk management, which frames the whole problem in business logic. Business logic ultimately gives the private sector every reason to argue the always hypothetical risk away, rather than solving the factual problem of insanely vulnerable cyber systems that control the nation’s most critical installations.

The authors suggest a policy-based approach that instead sets clear guidelines for asset owners, starting with regulations for new critical infrastructure facilities, and thereby avoids perpetuating the problem in systems and architectures that will be around for decades to come. In contrast to the IT sector, the industrial control systems (ICS) that keep the nation’s most critical systems running are much simpler and much less dynamic than contemporary IT systems, which makes eliminating cyber vulnerabilities, most of which are designed into products and system architectures, actually possible. Finally, they argue that a distinction between critical and non-critical systems is a bad idea that contradicts pervasiveness and sustainability of any effort to arrive at robust and well-protected systems.

Obama and Cyber: From Offense to Defense

There is little disagreement that a major characteristic of the 44th presidency is the growing role cyber played in the context of national security.¹ President Obama's first term was marked by the incredibly quick—yet mostly silent—buildup of the world's largest cyber firepower, including an actual “*bits on the ground*” operation in a hostile country (Stuxnet's cyber sabotage of Iranian nuclear research). But even while much of the United States' offensive cyber activities have been justified by the government as (active) defense, the case can be made that several significant threat agents, especially non-industrialized adversaries and non-state actors, will hardly be impressed by such deterrent cyber force.

Thus, President Obama has put special emphasis on cyber protection of national critical infrastructure at the beginning of his second term. The new presidential executive order titled “Improving Critical Infrastructure Cybersecurity” is a timely acknowledgement that the Pentagon's capabilities are not sufficient to protect the nation's most critical systems.² The order says:

“The cyber threat to critical infrastructure continues to grow and represents one of the most serious national security challenges we must confront. (...) It is the policy of the United States to enhance the protection and resilience of the Nation's critical infrastructure and to maintain a cyber environment that encourages efficiency, innovation, and economic prosperity while promoting safety, security, business confidentiality, privacy, and

civil liberties. We can achieve these goals through a partnership with the owners and operators of critical infrastructure to improve cybersecurity information sharing and collaboratively develop and implement risk-based standards.”

Unfortunately, this new order is set up to fail. By promoting voluntary action by the private sector supported by information sharing on cyber threats and risk-based standards, the executive order doesn't deliver on a fresh approach. Efforts to address the very same problem by similar means go back to the Clinton administration and have not resulted in any measurable improvements.

As it stands, critical infrastructure protection is an area where private companies are expected to assume much more responsibility—and even pay the cost—for national security. While it is comfortable to think that the private sector would be willing and able to solve the problem, either on their own or with the help of the government, so-called public-private partnerships, experience has shown otherwise.³ For example, well-intended government efforts such as the National SCADA Testbed (NSTB) elicited a bulk of design vulnerabilities in the control system products that are used to control the nation's most critical installations.⁴ While these vulnerabilities were passed to the vendors in question, they mostly went unaddressed. Control system vendors don't see a business opportunity in improving the cyber security of their products if it's not a selling proposition. Corresponding programs by the U.S. Department of Homeland Security that identified critical design vulnerabilities such as the Aurora vulnerability (physical destruction of electrical generators by cyber) or the Boreas vulnerability (permanent

¹ David E. Sanger, *Confront and Conceal. Obama's Secret Wars and Surprising Use of American Power* (New York: Crown, 2012).

² Barack Obama, “Improving Critical Infrastructure Cybersecurity,” Executive Order, February 12, 2013, <http://www.whitehouse.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity>.

³ James P. Farwell, “Industry's Vital Role in National Cyber Security,” *Strategic Studies Quarterly* (Winter 2012): pp. 10–41. <http://www.au.af.mil/au/ssq/2012/winter/farwell.pdf>.

⁴ “Common Cyber Security Vulnerabilities Observed in Control System Assessments by the INL NSTB Program,” U.S. Department of Energy, Idaho National Laboratory, November 2008, http://www.inl.gov/scada/publications/d/inl_nstb_common_vulnerabilities.pdf.

disabling of controllers by loading manipulated firmware) have not been similarly addressed by the majority of owners/operators or vendors even five years after they have been documented.

At the core of such continued failure is a methodological flaw. It is rooted in overstressing the concept of “risk,” which also acts as the backbone of President Obama’s suggested strategy. We have talked about risk and risk management so often that we have become convinced of the utility of the concept rather than giving it a thorough examination and reality check. For critical infrastructure protection, risk management is a recipe for failure.

How Risk Management Separated from Security

Can risk be effectively managed? The sober reality is that in respect to the cyber security of critical infrastructure, there is no empirical evidence that a risk-based approach, despite its near decade of practice, has had any success. In fact, all of the data suggests that we’re losing the cyber security battle in the IT space as well as the ICS space.⁵ The best way to understand why the concept of risk continues to be used anyway is to trace its origins.

Several decades ago, IT security experts realized that it had become practically impossible to fully secure their systems—largely due to growing system complexity.⁶ The logical strategy was to prioritize systems with respect to their importance (or

value) and to factor in the cost of security countermeasures. The cost/benefit approach would then result in situations where some risk would be “accepted” because mitigation appeared to cost more than the perceived impact of a cyber attack. The idea of risk management in IT caught on quickly, partly because it looked like something that would fit well into an environment where complex algorithmic solutions were everyday business. Over time, a useful heuristic turned into complex math. It was believed possible to predict accurately the future, and to allow for a calculation of mitigation cost versus cost of consequence, in which decision makers would ultimately be able to derive whether specific risks should be mitigated or simply “accepted” in a spreadsheet exercise.

The implied prediction is that negative cyber consequences will only materialize in the identified areas, and because the appropriate mitigation has been applied, they will not happen at all. All this comes at a cost, which is lower than the hypothesized cost of a materialized incident, thereby making the business case. In other words, risk management is not a technical approach but a business approach. It teaches to identify and apply risk mitigation strategies that ultimately result in a lower cost than the cost of a potential incident. Business decision makers quickly realized a remarkable difference between

The sober reality is that in respect to the cyber security of critical infrastructure, there is no empirical evidence that a risk-based approach...has had any success.

⁵ Mark Fabro and Zach Tudor, “What Went Wrong? A Study of Actual Industrial Cyber Security Incidents,” 2012 ICSJWG Spring Conference Presentation, http://www.us-cert.gov/control_systems/icsjwg/presentations/spring2010/02-ZachTudor.pdf.

Leon Panetta, “Remarks by Secretary Panetta on Cybersecurity to the Business Executives for National Security, New York, October 11, 2012, <http://www.defense.gov/transcripts/transcript.aspx?transcriptid=5136>.

“Cyber Intelligence...setting the landscape for an emerging discipline...,” Intelligence and National Security Alliance, September 2011, http://www.insonline.org/i/d/a/Resources/Cyber_Intelligence.aspx.

“2012 Data Breach Investigations Report,” Verizon, 2012, www.verizon.com/enterprise/databreach.

⁶ Chris Hall et al., “Inter-X: Resilience of the Internet Interconnection Ecosystem,” European Network and Information Security Agency, April 11, 2011, <http://www.enisa.europa.eu/activities/Resilience-and-CIIP/critical-infrastructure-and-services/inter-x/interx/report>.

cost of consequence and cost of mitigation: The latter displayed in red numbers on the balance sheet, with the former, some hypothetical number in a projected worst-case future, never shown at all. Business logic encourages risk-taking over the spending of resources to avoid one of many potential futures. And when it comes to cyber security, many private asset owners and operators of critical infrastructure are quite happy with risk management since it provides a rationale for doing nothing.

What's Wrong with Cyber Risk?

How could a seemingly useful concept become so blatantly abused? A closer look reveals several conceptual flaws that explain the failure.

Exactly how far are we looking into the future?

Several important assumptions underlie the risk-based approach: First, that it is possible to correctly identify all exploitable vulnerabilities for any given target of evaluation; second, that it is possible to measure the threat level, which implies correct identification of all potential attackers, and correct assessment of their motivation and capabilities; and third, that it is possible to correctly quantify all cost consequences resulting from a successful exploitation, and all of the cost of mitigation. If for any given target of evaluation any of these parameters doesn't represent the full story (for example, because vulnerabilities have been overlooked or cost of mitigation failed to include substantial cost items), the risk assessment for such target of evaluation might turn out to be grossly incorrect. On a larger scale, where risk assessments of multiple targets of evaluation are used to prioritize mitigation efforts, such prioritization can again be grossly misleading.

When it comes to incomplete data entered into the risk formula, that's not even the biggest problem. The concept of risk is predictive, since it arrives at assumptions about future events and their cost. Risk management is an attempt to control the future based on predictions of what the future will look like. Therefore, any determination of risk implicitly assumes a predictive timeframe. Unfortunately, cyber security experts rarely specify if their timeframe under consideration is a month, a year, a decade, or the lifetime of the target of evaluation. Failing to be specific on timeframe alone makes any risk assessment non-verifiable—usually without explicitly saying so. Even the National Institute of Standards and Technology (NIST) cautions that the validity of risk assessments are bounded in time, but does not offer a solution to this intractable problem.⁷ For example, when calculating cost of consequence it makes a difference if the projected negative outcome manifests once every ten years or once per week. Not only is the timeframe a complicating factor, so is the scope. Case in point, the 2003 blackout demonstrated how consequences can cascade well beyond the confines of a single organization.⁸ In a similar way, the cost of mitigation is largely influenced by timeframe, as some security controls, most notably the testing and installation of security patches and anti-virus updates, for example, must be performed periodically, in which case mitigation costs multiply over the lifecycle.

Three experts, four opinions

The reliability of a method is the degree to which independent people arrive at identical (or at least very similar) results when applying the method to the same subject. It appears that the risk-based approach to cyber security has very low reliability. There is usually some dispute as to the "real"

⁷ Ronald S. Ross, "Guide for Conducting Risk Assessments," National Institute of Standards and Technology, September 17, 2012.

⁸ "Final report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations," US-Canada Power System Outage Task Force, April 2004, <https://reports.energy.gov/BlackoutFinal-Web.pdf>.

risk level for any given target of evaluation—that is, for the few risk assessments that are actually debated by independent parties. For example, vendors of control system products usually arrive at very low risk levels for their products, compared to consultants who may work for vendors of cyber security products. Besides such obvious conflict of interest, there is an inherent problem because consultants have to use worst-case assumptions, leading to conclusions that are easily dismissed as unrealistic or pessimistic.

In most cases, the gambling coin in the discussion of risk assessment is threat. Since there is no scientifically sound way to measure threat (which would involve quantifying the motivation and capability of potential adversaries and predicting their actions), it is always possible to dismiss the threat part of the risk equation as minimal. Every seasoned ICS security consultant has experienced owners/operators asking the naïve question, “Why would anybody cyber-attack us?” They are expressing the all too common assumption that if a plausible motive cannot be established, it will not happen. The argument is sometimes backed up by the fact that a rather small number of cyber attacks against control and safety systems have been reported to date. However, even noting that the number of cyber attacks against ICS is increasing, there is a natural bias to under-report such events given the potential to affect a company’s stock price.⁹ Risk parameters are far from hard data that could be objectively measured (which is the simple reason why risk is assessed rather than measured). This can produce significantly different results that

different assessors arrive at for the very same target of evaluation. Such differences usually become a prominent subject of debate when vendors of ICS products are confronted with unfavorable cyber security features of their products.

What are we really referring to when talking about risk?

Abstract arithmetic for calculating risk may appeal to some, suggesting that risk is an exact science, yet most if not all the existing models have not been validated.¹⁰ It is also reasonable to suspect that the more complex risk calculations get, the less they can actually be linked to empirical reality. The stock trader Nassim Taleb, author of the popular book *The Black Swan*, recently commented that in the last century, financial decision makers favored the philosophy: “If there is a formula in it, don’t use it” when thinking about inherent risk in specific portfolio strategies and finance products.¹¹ This trend of the last century underwent a profound reversal when the banking industry and the financiers of Wall Street embraced an extremely complex risk formula. There is little argument that blind acceptance of the output of a risk calculation ultimately contributed to the global financial crisis of 2008.¹²

The validity of a measurement method is determined by how accurately it measures what it claims to measure. Applied to cyber risk, the question is, does the risk-based approach really measure the likelihood of experiencing a cyber attack? This is certainly an important question for any government regulator and even more so

⁹ Eric Byres, David Leversage, and Nate Kube, “Security incidents and trends in SCADA and process industries,” *The Industrial Ethernet Book* vol. 39, issue 2 (May 2007): pp. 12-20. http://www.mtl-inst.com/images/uploads/datasheets/IEBook_May_07_SCADA_Security_Trends.pdf.

¹⁰ Vilhelm Verendel, “Quantified security is a weak hypothesis: a critical survey of results and assumptions,” in *Proceedings of the 2009 workshop on New security paradigms workshop*, Oxford, UK, September 8-11, 2009, <http://portal.acm.org/citation.cfm?id=1719030.1719036>, pp. 37-50.

¹¹ Nassim Nicholas Taleb, *The Black Swan: The impact of the highly improbable* (New York: Random House, 2007).

¹² Felix Salmon, “Recipe for Disaster: The Formula That Killed Wall Street,” *Wired* vol 17, issue 3, (February 23, 2012), http://www.dpwireless.net/BackPage/TheSecretFormulaThatKilledWallStreet_200903.pdf.

for any executive who bases investment decisions on a risk assessment. For example, what exactly does a “security control” control? The general problem is that there is no empirical co-variation between security controls and reported cyber security incidents in critical infrastructure; doing nothing (in terms of security controls) does not necessarily result in security incidents. In other words, quantifying security has not been proven as valid, despite concerted efforts in many fields such as computer science and economics.¹³

Any forward-looking assumptions that are based on past behavior are prone to a well-known fallacy; because nobody has as yet done it is no proof that it cannot be done. As a case in point, the vulnerabilities exploited by Stuxnet had been dormant for over 10 years.¹⁴ While it may be tolerable to simply “wait” for any incident to materialize as proof of risk, there are facilities, for example in the nuclear industry, where such a posture is not viable. Vulnerability analysis in lab environments can be used to reliably identify real vulnerabilities that simply “wait” to be exploited. Sitting it out, while remaining silent about the vulnerabilities rather than fixing them, usually referred to as “security by obscurity,” has been considered a “best practice” for decades. Such practice urgently needs to be abandoned.

An alternative approach to statistical probability that relies on historical data is the logical cause-and-consequence model. This model assumes a logical relation of cause and consequence, moderated by uncertainty. Risk is an intermediary concept between fate and uncertainty. Where cause and consequence are deterministic, the whole notion of risk becomes inappropriate.

The logical link between root cause, moderating factors, and effect can be established either by experiment (demonstrated causality) or by statistics (inferred causality). If neither is possible, it is inappropriate to talk in terms of risk because no empirical prediction can be made. In other words, predictions without the possibility for empirical verification are as useless as astrology. While it may be difficult for computer scientists to admit, cyber security is not yet a science.¹⁵

For cyber security, it turns out that what is often referred to as risk is just “uncertainty,” as any moderating factors appear to be unknown or immeasurable.

“Maneuver speed” of risk mitigation in critical infrastructure environments

Threats and vulnerabilities are moving targets. New threats and vulnerabilities can pop up overnight.

A basic assumption of risk mitigation is that after having performed a risk assessment, security controls can be implemented in ample time before an incident materializes, putting defenders and attackers in some kind of race where the defenders try to outsmart attackers by anticipating their moves and implement working countermeasures in time.

Even if that would work to some extent, industrial environments are not designed for rapid reconfiguration, making it practically impossible to implement mitigation for new threats in short order. Formerly unsuspecting systems on the plant floor can become “critical” all of a sudden due to the discovery of new threats or new vulnerabilities, but it may take years to

¹³ Verendel, “Quantified security is a weak hypothesis.”

¹⁴ Ralph Langner, *Robust Control System Networks: How to Achieve Reliable Control After Stuxnet* (New York: Momentum Press, 2012).

¹⁵ D. McMorrow, “Science of Cyber-Security,” The MITRE Corporation, November 2010, <http://www.fas.org/irp/agency/dod/jason/cyber.pdf>.

develop and/or implement proper protection. The latter is an important aspect of ICS security; it could be viewed as the security “maneuver speed” of control system installations.¹⁶ In the control systems world, it is not possible to mitigate vulnerabilities within days or even weeks.¹⁷

The Stuxnet computer virus gives an example: An asset owner intending to mitigate the risk of a copycat attack would need at least months best-case and years worst-case.

Where warning time will predictably always be far short of adequate, preparedness must become a strategic priority.

For instance, industrial control systems in power plants can be reconfigured only once per year in a maintenance window where the

plant is shut down, which is usually referred to as annual outage. After plant restart, it will take another year until further configuration changes can be made, quite a huge—and little known—difference to the IT world where security patches can be rolled out within hours in a virtualized environment.

The maneuver speed needed to respond to new vulnerabilities and threats in a control system installation should reflect an organization’s willingness to accept the impact of compromise. In other words, if you can accept the impact of newly discovered vulnerabilities, then an ICS environment with slow maneuver speed causes little consternation. However, because this timeframe is extraordinarily large for typical environments in critical infrastructure, extending well beyond a decade in the nuclear setting because of the required testing and certification process, a reactive approach to cyber security has little chance for success. It

would make sense only for threats that can be identified more than a decade in advance.

It is worthwhile to point out the relationship to President Obama’s executive order, which relies on threat information sharing, because threat intelligence can be useless if there is no practical way to act on it. The uncomfortable reality is that the majority of asset owners in critical infrastructure, and maybe even those within the U.S. Department of Homeland Security who are responsible for assisting them, would have no idea what to do when learning that a significant cyber attack was imminent. Until this changes, the authors suggest to put less emphasis on information sharing. Where warning time will predictably always be far short of adequate, preparedness must become a strategic priority.

Where risk went wrong

In light of such conceptual problems, let us see how well intentioned organizations have missed the mark by using or promoting risk-based methodologies.

In a tacit admission that risk-based cyber security presented more problems than it solved, the North American Electric Reliability Corporation (NERC), early in the development of the Critical Infrastructure Protection (CIP) standards, opted to establish a set of criteria to identify critical assets.¹⁸ The NERC-CIP criteria makes sense from an electrical engineering perspective, for example, in identifying generating plants that equal or exceed 1,500 MW in a single interconnection as “Critical Assets.” Although one can view the NERC-CIP standards as having integrated risk into their

¹⁶ Industrial control systems are the digital devices that are used to control and monitor the actions of physical equipment such as pumps, valves, motors, or burners in realtime. Those systems are not computers running the Windows operating system and don’t have keyboards, monitors or hard drives attached. For an introduction to industrial control systems and their particular security challenges see Joseph Weiss, *Protecting Industrial Control Systems from Electronic Threats* (New York: Momentum Press, 2010).

¹⁷ “NSTB Assessments Summary Report: Common Industrial Control System Cyber Security Weaknesses,” Idaho National Laboratory, May 2010, <http://www.fas.org/sgp/eprint/nstb.pdf>.

¹⁸ “Reliability Standards for the Bulk Electric Systems of North America,” North American Electric Reliability Corporation, January 13, 2013, http://www.nerc.com/docs/standards/rs/Reliability_Standards_Complete_Set.pdf.

criteria thinking, an artificially constructed perimeter that exists only in the mind of the defender does not seem like an effective defense.

As another example, the following is a definition of risk-based decision making from Appendix C of the Department of Homeland Security's Risk Lexicon: "Risk-based decision making is defined as the determination of a course of action predicated primarily

Agencies are expected to make sound decisions on unsound (non-scientifically validated) methodologies.

on the assessment of risk and the expected impact of that course of action on that risk."¹⁹ The basic assumption embedded in this and all risk formulae is that

unknown future events of an unknown frequency, unknown duration, unknown intensity, from an unknown assailant, with unknown motivations, and unknown consequences are quantifiable. Consequently, if one thinks s/he can measure the risk, the mistaken conclusion is that one can manage the risk.

Using the NIST as exemplar, the recently published Special Publication 800-30 Rev 1 "Guide for Conducting Risk Assessments"²⁰ provides a strong endorsement of risk assessment methodology. Although the NIST guide provides some caveats to consider, these do not seem to impede its adoption by many agencies as the cautions are lost in the cadence of the march toward adoption of risk-based approaches. There are many ways to manage risk, but if an organization envisions

the first step as making a big investment in calculating risk (and its constituent threats, vulnerabilities, and consequences), it may not be making the best investment of its cyber security resources. Nonetheless, agencies are expected to make sound decisions based on unsound (non-scientifically validated) methodologies.

The best example of risk quantification gone wrong is in the financial market and, what's worse, it became a shared global experience. David X. Li, a brilliant Ph.D. with multiple advanced degrees, created a formula known as a "Gaussian copula function." Banks, rating agencies, and regulators alike adopted this formula to measure risk until everything went wrong in the so-called financial crisis of 2007-2008 (dubbed "The Great Recession" more recently). Some could see the problem coming and were less enamored with this new risk formula as noted by Taleb, "Anything that relies on correlation is charlatanism."²¹ This sentiment is echoed by Salmon, "It was a brilliant simplification of an intractable problem," and further emphasized by Li himself, "The most dangerous part is when people believe everything coming out of it."²²

If the Concept of Risk is the Problem, What's the Solution?

There are many approaches to addressing the issue of cyber security for critical infrastructure, however, addressing the policy first provides the greatest advantage.²³ While there are many lim-

¹⁹ "DHS Risk Lexicon: 2010 Edition," Department of Homeland Security Risk Steering Committee, September 2010, <http://www.dhs.gov/xlibrary/assets/dhs-risk-lexicon-2010.pdf>.

²⁰ Ross, "Guide for Conducting Risk Assessments."

²¹ Taleb, *The Black Swan*.

²² Salmon, "Recipe for Disaster."

²³ Steven R. Chabinsky, "Cybersecurity Strategy: A Primer for Policy Makers and Those on the Front Line," *Journal of National Security Law & Policy* vol 4, issue 27 (August 13, 2010): pp. 27-39.

Cyberspace Policy Review: Assuring a Trusted and Reliant Information and Communications Infrastructure, The White House, May 29, 2009, http://www.whitehouse.gov/assets/documents/Cyberspace_Policy_Review_final.pdf.

Barack Obama, "Remarks by the President on Security Our Nation's Cyber Infrastructure," Washington, DC, May 29, 2009, <http://www.whitehouse.gov/video/President-Obama-on-Cybersecurity#transcript>.

"Cybersecurity Two Years Later," Center for Strategic and International Studies, January 2011, http://csis.org/files/publication/110128_Lewis_CybersecurityTwoYearsLater_Web.pdf.

itations as to what policy can achieve, in terms of preventing all manner of bad things from happening policy isn't helpless. It has every right, or even duty, to employ a pragmatic best-effort approach for staying out of trouble.

In the following sections, we outline three basic principles that might guide policy-based approaches to address the cyber threat to national critical infrastructure without referring to the concept of risk. The main emphasis of our approach is to provide a framework for real-world protective action that may help to break the habit of discussing or attempting to measure cyber risk as a handy excuse for doing nothing. The three principles, which together form the critical infrastructure cyber protection triad, are: one, the primacy of politics over economics; two, a focus on practical efforts to fix design vulnerabilities; and three, pervasiveness rather than restricting cyber security efforts to "critical" systems. In short, it's *politics, practicality, pervasiveness*, or PPP.

Principle 1: Primacy of Politics

Critical infrastructure protection is a political issue, it doesn't necessarily generate profit.

Existing approaches predicated on risk are bound to fail because they come with the promise of helping private corporations improve their bottom line by preventing cyber attacks. Framing critical infrastructure cyber security within the concept of risk automatically puts it into a business context and will inevitably result in business decision makers determining that they are better off to simply wait and do nothing (i.e., "accept the risk"). The authors of this article by no means assert that investments in the cyber security protection of cyber-physical systems that control and protect critical infrastructure could not in practice save money in any arithmetic sense. Rather, our intent is to make the case that investments in cyber security may or may not pay off; we just do not know. Moreover, the notion of

saving private corporations money has rarely been a factor in matters of true national security, and critical infrastructure protection certainly is a national security issue. For example: Do we build military capability based on how much money we can save? The question does not arise because it does not make sense. There was no debate whether the U.S. Department of Defense's activities in cyberspace would be beneficial for the economy—and rightly so.

For any individual business owner in critical infrastructure, protecting against cyber attacks can usually not be justified by the prospect of improving the bottom line in quarterly results. It would simply be justified by improving America's security. However, the objective of any commercial corporation is to maximize financial success. Improving national security doesn't necessarily pay off in hard currency. If cyber security would show in black numbers on the balance sheet, companies would have been doing it for many years, without the President having to touch the subject. Members of the U.S. Congress didn't realize this in 2012 when they rejected the idea of subsidies or tax benefits for asset owners implementing cyber security because they didn't want private sector companies to enjoy monetary benefits for simply "doing the right thing."

Like air traffic security, workplace safety (e.g., OSHA), environmental pollution, or the reduction of toxic substances like lead in commercial products (e.g., RoHS), the cyber security of critical infrastructure is not a mere technical or business issue but rather a political one. If such issues are left to the discretion of decision makers in private corporations, little progress will be made. On the other hand, new political priorities and regulatory rulemaking do create new market opportunities for the private sector.

Appropriate design changes in organizational security posture using reference architectures can be made without significant cost increases if the emphasis is on securely equipping new installa-

tions, which should be the highest priority.²⁴ While replacing or retrofitting insecure legacy technology is a huge task, there is no reason to continue procuring, installing and commissioning that same insecure product base every month in countless installations, which would be with us for another decade or more. This predilection toward status quo can also be viewed as a litmus test for how seriously the cyber security issue in critical infrastructure is being taken. Once the political will has reached a consensus to implement a change, its implementation is usually not approached radically but rather by attending it in new installations. For example, when the European Union decided that traditional light bulbs are bad because of their poor energy efficiency, the verdict was not to replace all the billions of light bulbs in use but rather to place a ban on selling legacy light bulbs. A similar approach with respect to existing insecure-by-design industrial control and safety systems would be a positive first step toward rectifying the problem.

Since there are more benefits from critical infrastructure protection for society than for individual business owners, we recommend federal subsidies or tax reductions for the implementation of appropriate programs.

Principle 2: Practicality

Fix the design vulnerabilities rather than hypothesize about threats

It doesn't require a risk assessment to come to the realization that flat networks without

defense-in-depth, non-hardened computer operating systems, poor programming technique, unauthenticated ladder logic and firmware loads, hard-coded access credentials, systems that don't survive a network scan, or incomplete and inaccurate documentation provide little protection against conventional malware, and zero protection against custom-built malware. Without having to look into a crystal ball (i.e., perform a risk assessment), such practice is wrong and dangerous and has little right to exist in critical infrastructure installations. It has been demonstrated many times in lab environments that popular control system products are highly cyber-fragile and insecure,²⁵ suggesting that they reliably exhibit the intended deterministic behavior only under a specific set of environmental conditions of which the owner/operator may not even be aware. Experience has shown that such unawareness also extends to vendors who, again, tend to put quality management only into product features that are considered selling propositions.

One problem that comes with the risk-based approach is the purported direct logical link between cause (threat), contributing factors (vulnerabilities), and consequence. Such linkage is often difficult to determine, urging security experts to provide exploitation scenarios that ultimately require psychic abilities. While the precise details of how, where, and when such vulnerabilities could be exploited, and what the resulting damage might be largely uncertain, prediction is not required for starting to reduce vulnerabilities.

²⁴ Fred Cohen, "Using architectural analysis tools for better protection decisions," Fred Cohen & Associates, October 26, 2011, https://ics-cert.us-cert.gov/icsjwg/presentations/fall2011/D2-20-1030am_Track3_Cohen_r_Title-ArchAnalTools.pdf.

Annie McIntyre, Blair Becker, and Ron Halbgewachs, "Security Metrics for Process Control Systems," Sandia National Laboratories, September 2007, http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/18-Security_Metrics_for_CS.pdf.

"NSTB Assessments Summary Report."

Keith Stouffer, Joe Falco, and Karen Scarfone, "Guide to Industrial Control Systems (ICS) Security," National Institute of Standards and Technology Special Publication, September 2007, http://industryconsulting.org/pdfFiles/NIST_Draft-SP800-82.pdf.

²⁵ Ludovic Pietre-Cambacedes, Marc Tritschler, and Goran N. Ericsson, "Cybersecurity Myths on Power Control Systems: 21 Misconceptions and False Beliefs." *IEEE Transactions on Power Delivery* vol 26, issue 2 (January 2011), http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5673737, pp. 161-172.

McIntyre, Becker, and Halbgewachs, "Security Metrics for Process Control Systems."

This is especially true where those vulnerabilities are design flaws (or in some cases actual design characteristics with unintended consequences), as are now found in most control system architectures and products.²⁶ Using control system products with backdoors built-in by design and similar security flaws may not result in being hit by a cyber attack, and using a more secure product is no guarantee of not being successfully attacked. Nonetheless, using insecure products to control a nation's most critical systems is at the least intolerably negligent.

While readers with a background in IT may think that a vulnerability-centric focus on cyber security has failed, this argument can be countered in the industrial domain. The big difference is system complexity. Vulnerabilities multiply with complexity. In IT, progress is accompanied with growth in complexity of applications, networks, and connectivity. This has largely been facilitated by taking advantage of Moore's Law and the resulting decline of the cost of physical memory. The need to write clean and efficient code has been superseded by the ability to throw ever more memory and compute cycles at the problem. As an example, the Windows 7 operating system spans across roughly 100,000 files, and requires a minimum of 16 gigabytes hard disk space and one gigabyte RAM. That's just for the operating system, not including any real application processing work. Based on the market share, customers seem to see some real benefit in all

that complexity—if they are aware of the trade-off at all—which inevitably results in a piece of software that can only be secured (even then only temporarily) with the help of periodic security “patches.”

In the industrial domain, a similar paradigm didn't take hold. Average controllers that are installed today are not much more complex than controllers that were installed 30 years ago, and even top-of-the-line modern controllers look antiquated and simple compared to any smart phone. In terms of capacity, an average controller is comparable to an 1980s-style personal computer with little processing power and memory, the latter often being measured in kilobytes. Wide area networking for controllers is usually done for convenience rather than for operational reasons, resulting in bizarre effects such as thousands of control systems being exposed to the Internet and easily locatable via a specialized search engine.²⁷ An average controller has less than ten dedicated communication counterparts (operator panel, SCADA or DCS server, engineering station, peer controllers, etc.) that stay static over the lifecycle. Millions of controllers in use today, and brand new products that will be installed tomorrow, are no more complex than hardwired systems that served the very same purpose in decades past. As a matter of fact, many control systems in contemporary installations could still be implemented using analog electrical relays, wire, and welding. The more modern programmable version is usually

²⁶ “Hard Problem List,” INFOSEC Research Council, November, 2005, http://www.cyber.st.dhs.gov/docs/IRC_Hard_Problem_List.pdf.

“The Future of the Electric Grid: An Interdisciplinary MIT Study,” Massachusetts Institute of Technology, 2011, <http://web.mit.edu/mitei/research/studies/the-electric-grid-2011.shtml>.

Irshad Ahmad Mir, Mehraj-U-Din Dar, and S.M.K Quadri, “Towards the Application of Security Metrics at Different Stages of Information Systems,” *Journal of Global Research in Computer Science* vol 2, issue 2 (February 2011), <http://jgrcs.info/index.php/jgrcs/article/view/139>.

Pietre-Cambacedes, Tritschler, and Ericsson, “Cybersecurity Myths on Power Control Systems.”

SHODAN, <http://www.shodanhq.com/>, accessed February 2013.

²⁷ SHODAN, <http://www.shodanhq.com/>, accessed February 2013.

S4 2012, “Digital Bond's SCADA Security Scientific Symposium held in January 2012 in Miami Beach, FL,” <http://vimeopro.com/s42012/s4-2012/video/36494103>, accessed February 2013.

chosen for convenience. The number of individual controllers in an average power plant is several hundred—a size that is quite easy to manage and secure. In other words, these installations are not insecure because of system complexity, as most IT systems are. They are insecure because, until now, cyber security was never a concern in their design. It is not an engineering challenge to design cyber-secure products of such limited complexity, a fact that suggests market failure.

Principle 3: Pervasiveness

Don't restrict cyber security efforts to "critical" systems

There is no reliable way to predict which specific vulnerabilities will be exploited in a presumed attack, or what the specific consequences will be. In practice, determining the "critical" systems of a given installation has become a difficult to impossible task because a compromise of system components that appear non-critical in isolation may become critical when attacked simultaneously or in sequence. The reality is that hidden dependencies—both in cyber and in physical function—are the norm, and require an in-depth analysis by experts to discover. Not only does this apply on a component level within a specific installation but also on a macro scale. For example, Federal Energy Regulatory Commission (FERC) experts have observed that while an individual power plant might look non-critical when viewed in isolation, it may still play a crucial role for the reliability of the power grid in given situations.

Another drawback of the criticality concept is the suggestion that it is acceptable to use insecure products, architectures, and procedures in non-critical systems. Such a notion, which describes the reality on many plant floors, makes it very difficult to pursue a reasonable level of cyber security. It assumes that operators, maintenance staff, and contractors, who are not typically experts in cyber security, are capable of identifying critical systems accurately and understand that such systems must be treated differently from

non-critical systems where insecure means and methods would be acceptable. From an operational point of view, it is much easier to apply certified-secure designs and operations across the board. Contrary to common belief, such secure operations do not reduce operational reliability or safety.

The reality is that hidden dependencies...are the norm, and require an in-depth analysis by experts to discover.

Conclusion

Efforts to address the cyber threat to national critical infrastructure span several U.S. presidential administrations. The U.S. Congress has looked at the problem since 2002, before the establishment of the Department of Homeland Security, without arriving at substantial legislation. Lack of progress is not due to inactivity or unwillingness. For example, the last Congress had 61 hearings on cyber security before the issue was, again, buried in frustration and partisan quarrels. Such record of failure can hardly be explained by having attempted too little; it suggests that the metric of success was based on output (standards, policies, partnerships) rather than outcome (actually increased security or robustness). This paper posits that methodology played an important role in this outcome. A continued fixation on the concept of risk that is predictive, framing the problem and its potential solution in terms of business economics that cannot be backed by empirical evidence and producing improper pseudo-solutions for industrial control system installations, will almost certainly result in many more years of unproductive action. As it stands, in the fast-evolving world of cyber threats we do not have the luxury of waiting; we have already seen some very sophisticated cyber attacks.

An old adage says that if the only tool one has is a hammer then every problem tends to look like a nail. For more than a decade, the hammer of risk has been applied to the problem of cyber

security for critical infrastructure. The recent presidential executive order has attempted to refine the size, shape, or construction of the hammer in hopes of a better outcome.

Information sharing is not new, public-private partnerships are not new, and risk-based standards are not new. All of these measures have been applied without significantly altering the

It is better to be a well-armored and well-armed adversary than just a well-armed adversary.

balance that currently seems to favor those with nefarious intent. We can no longer simply do more or do better than what we did previously. We have to

accept the fact that what we have been doing is not working for critical infrastructure. We must let go of those actions that might make us feel better (i.e., security theater) and, rather, start to focus on protective action. Calculating the risk or probability of an attack accomplishes neither of these objectives.

Based on the President's remarkable use of cyber in his first term, it is justified to assume that the Obama administration is in a better position to finally arrive at measurable and substantial success in protecting the nation's most critical assets than anybody else. But that requires rethinking the problem rather than sticking with alleged "best practices" that don't cost money, just as we have seen it on the offensive side. While the debate on critical infrastructure protection has been stuck in the realm of risk management

for more than a decade, the impressive array of both defensive and offensive military cyber capabilities that came to the fore in just the first term of the Obama administration didn't suffer a similar unproductive fate. Within just four years, the United States became the first cyber superpower in history. A similarly scaled effort to protect critical infrastructure control systems is likewise needed. Even from a military perspective, reliance on deterrence (or active defense) while more or less ignoring protection (or passive defense) is equally questionable. It is better to be a well-armored and well-armed adversary than just a well-armed adversary.

More so than Congress, the President cannot shy away from the fact that this matter of national security is a political issue that ultimately costs money rather than makes money, a simple fact that the private sector understands and doesn't embrace for obvious reasons. Corporations certainly decry the burden of more regulation, yet somehow nuclear power plants, for which cyber security is regulated already, continue to operate and generate a profit. That regulation did cost industry additional money, and those costs are ultimately passed to the consumer. However, it was political reason that resulted in placing the priority on enhanced cyber security rather than on producing electricity at minimum cost. From a technical perspective, solid protection of cyber systems in critical infrastructure is indeed possible. It just needs us to reframe our understanding of the problem.

About the Authors

RALPH LANGNER is a nonresident fellow with the Center for 21st Century Security and Intelligence at Brookings. He is also director of Langner Communications, an independent cyber security consulting firm that he founded in 1988. In his consulting business, he accumulated 25 years of hands-on field experience in the cyber security issues of power plants, water treatment facilities, food and beverage plants, nuclear facilities, automotive factories, steel mills, and many more. Langner received worldwide recognition for his quick and comprehensive analysis of the Stuxnet malware. In addition to his consulting business, Langner is a frequent keynote speaker at international conferences on cyber security, national security, and critical infrastructure protection. He has been invited to share his insight by NATO, the U.S. Senate Homeland Security and Governmental Affairs Committee (HSGAC), the U.S. Nuclear Regulatory Commission, and the International Atomic Energy Agency. Langner also talked at TED and appeared on *60 Minutes* (CBS). He was featured in the movie documentary *Weapons of Mass Disruption* for the International Spy Museum and in *Wired* magazine.

PERRY PEDERSON began protecting critical infrastructure with the Department of Defense and continued that effort as the Director of the Control Systems Security Program at the Department of Homeland Security where he managed projects such as AGA-12 and the Aurora project. Currently, with the Nuclear Regulatory Commission, he is helping build the regulatory framework for cybersecurity at U.S. nuclear power reactors and has consulted with the International Atomic Energy Agency on applying security controls to digital instrumentation and control systems globally. He received the 2006 SANS Process Control / SCADA Security leadership Award and served as an inaugural member of the Governing Board for the Smart Grid Interoperability Panel for two years. Mr. Pederson is also a candidate for a doctorate in Information Assurance from the University of Fairfax. The views expressed in this article do not reflect the official policy or position of the U.S. Nuclear Regulatory Commission or the U.S. Department of Homeland Security.

BROOKINGS

1775 Massachusetts Ave., NW
Washington, D.C. 20036
brookings.edu