Cyber-risk, Standards, and Best Practices

The electric power industry needs a transparent, funded, independent, dedicated, focused Best Practices effort. If we want to achieve appropriate mitigation levels to protect industry infrastructure against cyber attacks we should do no less.

By Paul Feldman and Dan Hill

The subject of cybersecurity is not only here to stay but will grow in importance over time. The literature is already filled with summaries of various attacks of all varieties—right up to nation-state mini-attacks[1] such as the North Korean 2014 attack on Sony. The literature is abundant with suggestions as to what to do to protect against cyber attacks—from the simple "don't click on unknown email links"—to the sophisticated response that requires a small army of experts to implement.

Above the clutter of attack and defense articles and reports, superstructures are emerging to put all this into some context and, as usual, we rely on tried-and-true mechanisms that arguably have been successful in other areas: standards and enforcement.

Cybersecurity Standards

Paul Feldman *is a director and past chairman of the Midcontinent ISO, where he chairs the Markets Committee and serves on the Information Technology and Governance Committees. Previously, Mr. Feldman served on the board of the Western Systems Coordination Council.* **Dan Hill** *is a board member of the New York ISO, where he chairs the Audit and Compliance Committee and serves on the Commerce and Compensation committee. Mr. Hill is retired from Exelon where he was Senior Vice President and Chief Information Officer. The views expressed here are Mr. Feldman's and Mr. Hill's and should not be attributed to any organization in which they serve.*

[1] The Sony attack was not particularly sophisticated, but was very effective from an attacker viewpoint, and devastating from the view of the attacked. Even so, it pales against the state of the possible in this area, and did not include simultaneous multi-site attacks that would likely form the core of an infrastructure attack.

he Executive Branch and Congress realize the importance of protecting the nation's critical infrastructure and have singled out energy as one of the critical industries that must be protected.[2] The authority and responsibility for making that happen, however, is not crystal clear, but involves an alphabet soup of agencies: DOE, NERC/FERC, DHS, likely NSA, and the infrastructure companies themselves.

US electric utilities had voluntary standards that Congress overrode in the Energy Policy Act of 2005. Congress made standards mandatory and ultimately empowered NERC to organize the energy companies to develop standards, and to audit and enforce them through NERC regional entities. In time, Critical Infrastructure Protection (CIP) standards emerged as a distinct set of NERC standards. Today they are the most often discussed, evolved, and more often violated than NERC's non-CIP standards[3].

Aside from the predominant process of NERC, utilizing voluntary utility representatives, recommending CIP standards to FERC, and having FERC approve or point out areas for reconsideration[4] by NERC, other players have also entered the standards business. The NRC[5] is responsible for standards, including cybersecurity, for the nuclear fleet. NIST[6] is involved in standards setting, but not auditing or enforcement. The NIST standards are mandatory for the federal government facilities, and the president has recommended that companies managing critical infrastructure employ them. And there are others.[7]

Do standards, backed by enforcement and penalties, work in the US power industry? We think there are possibly two answers, one for non-CIP and one for CIP. For non-CIP NERC standards, there is little debate in the utility industry that the Bulk Electric System (BES) (where almost all NERC standards apply) is more reliable than in the past. Changes to the BES components and technology happen at a glacial rate, so non-CIP standards can keep pace moving in tandem.[8] NERC standards, perhaps combined with organizational advances such as RTOs and technology advances such as PMUs,[9] seem to have made a positive

---

[2] In fact, the Presidential Policy Directive on Critical Infrastructure Security and Resilience calls out "energy and communications systems as uniquely critical"

[3] See http://www.nerc.com/gov/bot/BOTCC/Compliance%20Committee%202013/February_5_BOTCC_Open_Package.pdf, p. 16

[4] FERC does not actually have the authority to make a standard.

[5] See http://en.wikipedia.org/wiki/Nuclear_Regulatory_Commission

[6] See http://en.wikipedia.org/wiki/National_Institute_of_Standards_and_Technology

[7] See a collection of cyber references at http://www.EnergyCollection.us/457.pdf

[8] It can take NERC two plus years to proffer/change a standard.

[9] See http://en.wikipedia.org/wiki/Phasor_measurement_unit

difference although it will take long-term analysis for the data to prove that supposition with great certainty.

Changes in the CIP physical security arena have been the target of recent attention, but the jury is still out on whether the standards process moves fast enough to adjust to physical security risks. The mismatch between standard-setting timeframes and technology evolution timeframes may present an ongoing exposure if standards alone are to guide our actions. A prime example is the recent rise of inexpensive, albeit high-functioning, drones and the risk they pose to facilities. However, the timeframe evolution of physical threats pale in comparison to cyber threats.

On the CIP cyber side, the picture is less clear. There are dynamics at work in the cyber area that do not apply to non-cyber BES operations. First, the cyber world is exactly the opposite of the non-cyber electricity world when it comes to speed. Cyber-attacks increase in sophistication on an almost daily basis. The attack surface is also constantly growing as the "Internet of Things" takes hold in the utility world on both sides of the meter.[10] Finally, the availability of attack weapons is growing and the underlying technology is simply code rather than advanced equipment. Sophisticated nation states can design their own attack weapons, while unsophisticated nation states can simply buy what they want from willing sellers; the culmination of the deal is a simple transference of code and digital directions rather than a physical shipment. What is clear however, is that a standards process that can only move at a two-year rate of change cannot compete with would-be cyber attackers whose tools can evolve daily.

While very few people have security clearance to actually know, it appears that the electric utility industry in the US has thus far avoided an outage due to a cyber-attack. No utility has admitted to such an outage. NERC senior management regularly claims no such outage has occurred each time a major vulnerability is announced or a successful attack has occurred elsewhere. Yet, every utility is attacked in some way every day, and there is little debate that there may well be a successful attack someday, we just don't know when or how bad.

Despite the record of success thus far, it would be unwise to declare success in an ongoing war where our defenses (if only Standards) evolve slowly and the enemies evolve quickly. If we rely on Standards we will eventually fail because of this simple mismatch.

Also, it is highly unlikely that the utility industry has been exposed to the full cyber-attack armory that exists. What would a real attack aimed at taking down the grid look like?

[10] Most electric utilities have segregated their OT (Operational Technology) and IT (Internet Technology) systems to provide a layer of separation that helps protect the OT systems. The authors are concerned with some of the architectural approaches being used – but save that subject for a future paper.

The main protection built into the management system for maintaining grid reliability is a combination of sufficient resources, sophisticated dispatch and component protection mechanisms. Sophisticated dispatch algorithms monitor and schedule operation of the grid on a five-minute basis, always protecting the grid from any single outage event: the so-called N-1 calculation. The N-1 criterion requires that the grid be able to withstand the loss of grid components resulting from any single event (typically a large generator or transmission line loss) and that component failures beyond N-1 are unlikely because such events are assumed to be uncorrelated.

We cannot protect the grid from every danger, but more resiliency, microgrids, stand-by generation, and distributed resources are complementary forms of protection.

Component protection mechanisms are built into the most expensive and hard-to-replace components so that they shut down (protect themselves) in the face of unexpected conditions that might lead to their destruction (e.g. misuse). These protection mechanisms are installed with the realization that it is better to have a temporary outage and restart the system than a long-term outage in which physical replacement of damaged grid components may have very long lead times.

These conditions are not the same in the cyber world, and it would be wrong to think about the supporting processes in the same way. A serious attack on US infrastructure would not target a single company or grid component, and such an attack would be well aware of built-in protection mechanisms. A serious attack would simultaneously attack multiple components on both cyber and physical dimensions. Even a small success in terms of the number of components taken off-line could at least temporarily take sections of the grid off-line, because the N-1 criterion does not protect against simultaneous multiple facility losses, only independent single facility losses.

In such a case, the protection mechanisms built into individual pieces of equipment may allow a black-restart,[11] but an attacker would likely know that and presumably have a plan to address that dimension as well. We cannot protect the grid from every danger, but movements toward more resiliency, microgrids, stand-by generation, and distributed energy resources are complementary forms of protection.

The reason for pointing out these issues is that standards are unlikely to save us from a concerted sophisticated cyber-attack. And yet, we cannot tolerate such an attack, as it would change the very fabric of our society, and in ways too horrible to describe.

[11] See http://en.wikipedia.org/wiki/Black_start

However, in addition to standards and frameworks as guidance, we have the familiar well-honed concepts of Risk Management.

Risk Management

Risk Management is an old idea that maintains its profound relevance in decision-making today. Much has been written on the subject and there are many well defined frameworks and materials to assist in applying this important principle.

The unofficial Internet Security Glossary[12] defines Risk Management as "The process of identifying, measuring, and controlling (i.e., mitigating) risks in information systems so as to reduce the risks to a level commensurate with the value of the assets protected."

Many are familiar with some form of the Risk Equation. Most board members have been on the receiving end of an explanation of how a nice graph of a company's Risk Analysis chart was developed. Most have also seen the famous Risk Management Equation:

Risk = [Threat] x [Probability of Threat] x [Impact if Successful]

Some will claim they actually use the equation. Some will claim that is notional and meant to guide further work. The goal of the equation, of course, is to understand risk exposure, and to mitigate the risk. Expenditures on mitigation that is less than the risk exposure seem prudent.

In the cybersecurity world however, while not deterring its widespread use, the equation is of little value beyond thinking about its ramifications. The "Threat" is an occurrence of an attack that we have not experienced, and can come in many forms that would fit under the heading of persistent, pervasive, military grade attacks. The "Probability of (such a) Threat" is most likely very low, but unknowable. The "Impact of a Successful Threat" is so large as to be unthinkable.[13] What remains is a very, very low probability times a very, very high negative impact; the multiplication of one by the other, coupled with confidence limits, would cover a very large dollar span indeed—so large as to be quite useless when comparing this to various cost levels of mitigation.

Traditional approaches to Risk Management fall short in the cybersecurity example but nonetheless expose a real issue: How do we know when we have done all that we should when there is not a quantitative approach that can adequately inform our decisions?

Best Practice Technology

Standards contribute to being cybersecure but cannot make us cybersecure. (Actually, it is impossible to be absolutely cybersecure.) Standards take too long to develop and implement, and the rate of technology change evolves the threat vectors in ways standards cannot anticipate. We should continue with

[12] See http://tools.ietf.org/html/rfc4949

[13] Setting aside the impact to human live, it is instructive to understand that the normal Distribution outages we experience absent any cyber attacks is in excess of $100B today.

Standards but not place our faith in them as the only approach or answer.

Likewise, traditional Risk Management is not a great help in determining how much protection is enough – i.e. we cannot realistically weigh the consequences against the expenditures and put forth a properly-sized level of protection.

Still, we must address the very real issue and to do that, the industry should pursue Best Practices. Wikipedia includes a good discussion of the term, but the core definition is helpful to frame the issue: "A best practice is a method or technique that has consistently shown results superior to those achieved with other means, and that is used as a benchmark."[14] Best Practices can evolve quickly, both to address the quickly evolving threats and to address the quickly evolving technologies available to address those threats. Where we thought ABC was not a threat yesterday, but is today we can assemble experts to put forth a Best Practice. Where XYZ technology was not available yesterday to protect against a threat, but a better technology exists today so experts should weigh in on the appropriateness of the new technology.

Admittedly, a Best Practice approach is not easy. Who is responsible? What if they are not compliant with Standards? What if it is deemed too expensive to provide the highest level of protection?

In the US utility industry, while there is a consistent view at the top corporate levels that pursuit of Best Practices is necessary, a complete infrastructure to deliver cybersecurity Best Practices in real-time is less developed.

NERC will not claim that turf. While ES-ISAC[15] is an industry-centric information clearinghouse it is not really involved in determining or defining Best Practices.

Both the North American Transmission Forum and the North American Generation Forum encourage members to share information on what works and what does not, but they are closed to all but members and there is no publication of related information.

DHS dips into Best Practices when it delivers site audit advice mainly because DHS only cares about mitigation and not Standards.

EEI has a CIO Executive Advisory Committee that includes a significant effort in sharing best practices and real time experiences across the group. EEI and AGA jointly host a Technical Advisory Committee (TAC) in which best practices and real time experiences are shared.

The IRC[16] similarly collaborates on cyber matters for the ISO/RTO community.

[14] See http://en.wikipedia.org/wiki/Best_practice

[15] The Energy Sector Information Sharing and Analysis Center (ES-ISCA) is a part of NERC but has organizational and rules separation from the standards development and importantly the NERC enforcement process.

[16] See http://www.isorto.org/Pages/Home

EnergySec[17] can be tasked to put forth a Best Practice in a certain area. (Many consulting and security firms also have best practice offerings.) Companies can form Best Practices and sharing groups amongst themselves. There are many other companies, both in semi-organized and ad hoc groups, pursuing the same notion.

CIOs that have told their Boards that the company is safe need to find other words to express its Risk condition.

Sharing of Best Practices and the resultant enhanced security profile is most effective when trusted entities are similarly able to share real time information about threats or specific compromises – this is an area where significant opportunities remain. Arcane rules over confidential information exchange, the risk of FOIA requests, and concerns about not-required disclosures becoming public causing reputational or financial harm all tend to prevent companies from sharing specific threat or compromise information which would otherwise help protect others from similar threats. Various governmental bodies are evaluating potential solutions to close these gaps and enable more sharing; identifying and enacting such solutions are of critical importance.

While an unlikely circumstance, the idea that a Best Practice may not be compliant with a standard needs to be dealt with in a way that removes compliance risk from the company deploying the Best Practice. Just as Congress is looking into cyber-insurance and protection of companies willing to share cyber-information with other companies, the issue of Best Practice non-compliance needs an answer. In many cases, a new Best Practice may well be less expensive or cost-comparable to today's Standard but not implemented because of the regulatory (enforcement body) overhang. The idea is not to get a Standard for every possibility; that is not possible. Rather, the idea is to deploy the best technologies (considering deployment cost juxtaposed against cost of failure) to be cybersecure.

The issue related to cost is perhaps the most difficult one. It would literally cost a fortune to try to protect a company from all cyber threats. CIOs that have told their Boards that the company is safe need to find other words to express its Risk condition, because there is no such thing as being absolutely safe from a cyber-attack. A CIO may have done everything that's known to be done and everything they have been told by experts that they should do, but vulnerabilities remain.

CIOs must continuously pursue the notion of deploying Best Practices. Many CIOs are leveraging Capability Maturity Model (CMM)

[17] EnergySec is the National Energy Sector Cyber Security Organization. *See* http://www.energysec.org/.

type assessments to measure their maturity in various dimensions of cybersecurity. This is a particularly powerful tool with which to communicate with Boards of Directors. DOE has developed its Cybersecurity Capability Maturity Model (C2M2), and a sector-specific model, ES-C2M2, for use by electric utilities. Similarly, ESCC (Electricity Subsector Coordinating Council) has also created a cybersecurity CCM-type tool. There will be instances in which the cost to implement is beyond a company's financial capability, or greater than the benefit obtained, but those should at least get a good airing.

There is very little "secret sauce" an energy company can deploy that the bad-guys are not familiar with.

The authors suggest that more be done in this area. A Best Practices approach ought to be more organized, funded and transparent. There is very little "secret sauce" that an energy company can deploy that the bad-guys are not already familiar with. Transparency will aid in quick development and sharing of information amongst those that need to take action.

In place, or in support of today's ad hoc Best Practice approaches, the electric power industry needs a transparent, funded, independent, dedicated, focused Best Practices effort. If we want to achieve appropriate mitigation levels to protect industry infrastructure against cyber attacks we should do no less. ■