“Developing a Framework to Improve Critical Infrastructure Cybersecurity”

RedSeal Networks is the leading provider of network security risk analytics, providing an analytic platform for businesses and government agencies to visualize their security architecture, continuously audit and monitor IT compliance, and eliminate cyber threats.

In “**Developing a Framework To Improve Critical Infrastructure Cybersecurity**,” the National Institute of Standards and Technology (NIST) requested information to help identify, refine, and guide the many interrelated considerations, challenges, and efforts needed to develop a framework to reduce cyber risks to critical infrastructure (the ‘‘Cybersecurity Framework’’ or ‘‘Framework’’). In the Request for Information, NIST included a **Request for Comment** section that encompassed three areas of focus:

Current Risk Management Practices
Use of Frameworks, Standards, Guidelines, and Best Practices, and
Specific Industry Practices

In its role as a trusted advisor to customers focused on their infrastructure security, RedSeal has gained insight into their practices, needs, and challenges, and has developed technology to assist them in assessing and securing that infrastructure. This response is from this perspective as a trusted advisor.

Current Risk Management Practices

NIST solicits information about how organizations assess risk; how cybersecurity factors into that risk assessment; the current usage of existing cybersecurity frameworks, standards, and guidelines; and other management practices related to cybersecurity. In addition, NIST is interested in understanding whether particular frameworks, standards, guidelines, and/ or best practices are mandated by legal or regulatory requirements and the challenges organizations perceive in meeting such requirements. This will assist in NIST’s goal of developing a Framework that includes and identifies common practices across sectors.

1. *What do organizations see as the greatest challenges in improving cybersecurity practices across critical infrastructure?*

In a word, complexity.

Modern IT infrastructure has grown faster than our ability, as operators, to track all the complex interactions in these systems. The technologies developed outside critical infrastructure migrate inexorably across into regulated environments, eroding older “air gap” approaches to control and defense. Individual infrastructure components are increasingly complex, but the interconnected web of networked elements greatly magnifies this problem. For example, a typical enterprise network has millions of paths for traffic resulting from the thousands of lines of configuration on dozens to hundreds of

network devices. We cannot understand and control critical infrastructure through redoubled human effort, so we have to embrace automation to ensure correct implementation of known security controls.

Existing efforts to enable automated assessment of cyber-defenses (for example, the SCAP standards) have built a substantial platform, but still fall short when it comes to assessing systems and their complex interactions. Checklists of audit controls are necessary, but not sufficient. To truly understand the security posture of our critical infrastructure, we have to be able to automate the assessment of the systems as a whole, to uncover, understand, and prioritize defensive gaps and to ensure defensive readiness ahead of the next shift in cyberwarfare techniques.

2. *What do organizations see as the greatest challenges in developing a cross-sector standards-based Framework for critical infrastructure?*

Consistency. Risk measurement is still in its infancy for today's rapidly evolving IT infrastructure, and ensuring comparable measures across the organization proves to be very difficult. The field of risk is well established in other areas – insurance, portfolio management, etc. However, these tools are difficult to apply to cyber infrastructure, since there is not enough data on the effectiveness of defensive security measures. This is compounded by the fear of sharing created by the stigma of security incidents together with the difficulty of being sure to measure those metrics which actually impact security, rather than those that are simply easy to measure.

3. *Describe your organization's policies and procedures governing risk generally and cybersecurity risk specifically. How does senior management communicate and oversee these policies and procedures?*

RedSeal builds automation software to analyze and measure cybersecurity risk. Using an analysis of network connectivity and endpoint vulnerability, the method determines attack paths and the security implications of those paths based on defined security standards for device configurations, network security zones, and host vulnerability accessibility. Building on SCAP standards like CVE and CVSS, the method computes the ease of exploit of given assets exposed to direct or indirect, external or internal attack.

Senior management consumes the results of this attack simulation in the form of dashboards, which display both the current overall attack risk of the infrastructure and the recent trend. That is, the measurements presume a motivated attacker, and the metrics then measure how easily that attacker will be able to break in, and how far they will be able to reach.

This outcome-oriented metric approach is distinct from many existing security metrics, which fall into a trap of measuring "busyness" and activity levels, or focus too much on compliance details, missing the broader context provided by attack simulation.

Trends in attack risk are easier to communicate, easier to understand by a non-specialist audience, and make it easier to correlate security investments with effectiveness.

4. Where do organizations locate their cybersecurity risk management program/office?

While the location of the cybersecurity risk management program/office varies across the spectrum of organizations RedSeal advises, it typically falls within the Information Technology or Finance functions of the organization. While there has begun to be high-level focus on Risk, raising it to the executive staff for some forward-looking organizations, most organizations have it reporting to either the CIO or CFO.

5. How do organizations define and assess risk generally and cybersecurity risk specifically?

The board looks at organizational risk overall, including financial assessments, workplace comp and liability risks, and so on. In this arena, cybersecurity risk has limited focus, without repeatable, quantifiable results to show. RedSeal is attacking this problem directly by providing a method that simplifies the complexity and offers a consistent measure to manage trends within the cybersecurity infrastructure.

6. To what extent is cybersecurity risk incorporated into organizations' overarching enterprise risk management?

It is acknowledged as an important issue and there is clear consensus that risk management needs to improve. Most organizations want realistic risk measurement of IT infrastructure, but have neither tools nor measures to do it effectively. In a competitive commercial environment, executive management fears losing their edge if they over-invest compared to their peers. This is a problem that standards and minimum requirements can address: ensuring that one company can afford to meet the minimum standards without fear of being undercut by a competitor willing to take excessive risks. Outside critical infrastructure, it can make sense to allow competitive downward pressure on expenses to find the right level of spending on attack prevention. However, once infrastructure becomes fundamental to broader national interests, the shareholders of each company alone cannot make optimal risk assessments in the national interest. Hence the role for standards: to allow critical infrastructure organizations to adopt the security automation controls they want to use to manage and reduce risk, but that they cannot adopt for fear of risk-taking competitors.

7. What standards, guidelines, best practices, and tools are organizations using to understand, measure, and manage risk at the management, operational, and technical levels?

Organizations are hungry for standards, guidelines, best practices, and tools to help address the challenge of cybersecurity. As a result, many available components are in active use, and organizations are working to find ways to integrate and correlate the

information already being gathered. For example, at the core of many vulnerability analyses is the National Vulnerability Database (NVD); automation tools already use the Common Vulnerability and Exposure (CVE) IDs from the NVD to determine severity of issues. Unfortunately the CVE and Common Vulnerability Scoring System (CVSS) is only one isolated standard and measurement that needs to be correlated with other system information (such as paths through the network) to determine mitigation priority and true risk. Tools that correlate this information are critical to the overall cybersecurity framework.

8. What are the current regulatory and regulatory reporting requirements in the United States (e.g. local, state, national, and other) for organizations relating to cybersecurity?

On the borderline between guidelines and regulation, the Data Security Standard promulgated by the Payment Card Industry deserves special attention. While not a legally mandated requirement, it serves as a strong example from the private sector of the measures that need to be taken. PCI DSS exists to protect one community (the card issuing banks) from the critical infrastructure they depend on (any merchants or card processors who hold on to cardholder data). In effect, the banks have said “if you’re going to have my money in your network, these are the rules you need to live up to”. The rules do not deliver perfect security, or freedom from risk. However, compared to other regulatory frameworks and guidelines, the rules have achieved a difficult balance. On the one hand, the rules in PCI DSS are quite technically specific, detailing what kinds of protections are mandated, and yet on the other hand, the rules contain some critical flexibility to allow different audited companies to innovate and improve. In that sense, PCI DSS is a powerful model – it is prescriptive, without being unnecessarily rigid. It is motivated by the (financial) interest of one group who depend on the basic security diligence of another.

Within formal, legal regulation sets, the only comparable, technically prescriptive example is the Critical Infrastructure Protection ruleset from the North American Electric Reliability Corporation (NERC CIP). Other legal standards come into play (HIPAA, SOX, GLBA), but generally these other regulation sets only demand that the end organization set their own technical standards for cybersecurity and IT infrastructure, and then prove they are following those standards. While well-intentioned, this gap is large – it means that different companies enforce wildly different levels of protection, while still meeting the letter of the law.

9. What organizational critical assets are interdependent upon other critical physical and information infrastructures, including telecommunications, energy, financial services, water, and transportation sectors?

RedSeal works alongside many organizations in different sectors, including critical infrastructure providers. This gives us a view into the different levels of protection, as we assess defensive readiness of these networks.. We find there is always critical interdependency, caused by the rise in networking and its constantly increasing

complexity. Traditionally “air gapped” assets such as power control systems are no longer isolated. As the network spreads and connects more formerly isolated environments together, complexity increases, errors of configuration increase, defensive gaps are left open, and risk rises dramatically.

10. *What performance goals do organizations adopt to ensure their ability to provide essential services while managing cybersecurity risk?*

In short, organizations monitor and test their network defensive readiness. The best prepared organizations regularly test their own defenses, using either real penetration testing teams, or attack simulations, or both. Such continuous testing is the only way to understand the priority of defensive gaps, since all modern infrastructure is rife with weak points. The performance objective cannot be “zero attack surface” – this has been untenable for decades now. The objective has to shift to a risk-based, quantified assessment – how much attack surface reduction, in return for how much effort?

The good news is that cybersecurity risk can be actively measured and managed, using known techniques. In our experience, risk tradeoff calculations to closely assess the cost of a security measure against the risk reduction benefit are often trivially simple, because the mitigation steps are not hard. The problem, as noted earlier, is the complexity and scale of modern infrastructure – we simply make too many mistakes building out existing, proven controls, making us highly vulnerable to attack. The cost to repair the mistakes is generally very low – the only real challenge is to find and prioritize all these defensive gaps. Automation is a great help here, so the most mature organizations (generally those with the greatest criticality of their infrastructure) have deployed automated, continuous re-assessment of defensive posture, and they measure their performance directly as an outcome of this continuous testing. They can say that they are effective when they demonstrate, day after day, year after year, that it gets measurably harder to break into their infrastructure.

11. *If your organization is required to report to more than one regulatory body, what information does your organization report and what has been your organization’s reporting experience?*

As a provider of a network risk management system, RedSeal works with customers in regulated spaces to report on PCI, NERC CIP, internal audit controls, and other specialized cybersecurity risk mitigation regulations. One aspect of this work is providing customized reports targeted at addressing the needs of the regulations and the auditors reviewing the cybersecurity controls.

12. *What role(s) do or should national/international standards and organizations that develop national/ international standards play in critical infrastructure cybersecurity conformity assessment?*

PCI is an example of a powerful and technically specific standard that has moved many enterprises towards better network security architecture. While backed by significant

financial penalties for non-compliance and as a consequence of breaches, it provides an example of how a standard can help an entire industry move towards best practice in risk mitigation and avoidance.

As NIST works through the development of a Cybersecurity Framework, consideration of the value of best practices, the benefits of incident avoidance, and the importance of a focus on reducing and mitigating risk in communicating the framework to owners of critical infrastructure will be vital.

Use of Frameworks, Standards, Guidelines, and Best Practices

As set forth in the Executive Order, the Framework will consist of standards, guidelines, and/or best practices that promote the protection of information and information systems supporting organizational missions and business functions.

NIST seeks comments on the applicability of existing publications to address cybersecurity needs, including, but not limited to the documents developed by: international standards organizations; U.S. Government Agencies and organizations; State regulators or Public Utility Commissions; Industry and industry associations; other Governments, and non-profits and other non-government organizations.

NIST is seeking information on the current usage of these existing approaches throughout industry, the robustness and applicability of these frameworks and standards, and what would encourage their increased usage. Please provide information related to the following:

1. What additional approaches already exist?

There are many frameworks and standards for proper configuration of individual elements of IT systems, but few for the proper defense of whole systems composed of interacting parts.

Examples of element configuration guidelines include NIST publications on hardening of network equipment and endpoints, the CIS Benchmarks, and DISA STIGs, among others.

System-wide approaches are much rarer. Section 1 of the PCI Data Security Standard (DSS) is one such example, where an organization is required to show correct configuration of access between whole network zones, not just individual elements. Meeting this requirement through manual effort alone is especially challenging, since the complexity of a single network counter-measure can easily outstrip the ability of a human assessor to read and digest in full. This problem is then compounded by the complex interactions of the overall system. The only way to answer essential questions for cyberdefense, such as “how big is my attack surface?”, is to apply automated analysiss.

2. Which of these approaches apply across sectors?

Single-element hardening rules are generally applicable – for any IT asset, there are ways to set it up in an insecure fashion, so at least some rules are truly universal. Of course, the degree of security required varies based on context – networks used during, say, war-fighting or air traffic control have more stringent requirements than typical web sites. However, this specialization is a relatively minor aspect of the problem – there are far more examples of gross errors that no organization wants to have than there are specialized rules for extra security in specialized environments.

System-wide rules can readily be established for any sector, while the precise details vary. So, for example, a commercial network with credit card data in it needs to define a zone for storage of cardholder data and this zone needs to be protected with appropriate controls. A power network has different specifics, but following the same framework: they need to define zones of control for SCADA (supervisory control and data acquisition) equipment or similar critical assets, and then ensure the zone is appropriately built, maintained, and defended. The concept of system-wide zone-based analysis applies to any sector, while the specific zones and controls will vary by specific need.

3. Which organizations use these approaches?

PCI is used by any organization that accepts credit card payments, so it crosses many industrial sectors. NERC CIP is used by electrical power producers. In addition, most organizations with mature security and risk analysis functions are creating infrastructure controls that are similar to those of PCI and NERC CIP. DISA STIGs represent another group of (primarily) per-device hardening checklists, specialized to the military sector, but grounded in common best practices and well publicized configuration errors.

Attack simulation and risk measurement is used less often (since it is generally not mandated). It is pursued most often by highly sensitive infrastructure – civilian and military agencies, intelligence networks, power generation infrastructure, etc. That is, there are organizations who aim higher than the bare bones mandated protections – generally when they are less beholden to shareholders, or when lives are explicitly on the line if the infrastructure is breached. This is the kind of advanced preparedness that needs to extend to all critical infrastructure, including those parts owned and operated by the private sector.

4. What, if any, are the limitations of using such approaches?

Primarily that neither automation nor human effort alone is sufficient to deal with the complexity of the challenges. It’s imperative that the two work together to accomplish that which neither can do alone, with automation analyzing the mountain of information required to accurately reflect reality, and with humans to digest the resulting knowledge to determine the subtle implications of the output. That is, no “pure” automation solution will be sufficient – it’s essential to provide visualizations and reports to track defensive

posture, so that management and operations can do what they do best – steer by setting policy.

5. What, if any, modifications could make these approaches more useful?

The Federal effort sometimes known as “FISMA reform” is a move in the right direction– stronger requirements for Continuous Monitoring of the state of defensive readiness is a critical foundation. The core breakthrough in this area is to push for a speed of reporting that makes it impossible to use human effort. Human effort has been the cheap way out for too long – people are not good at this work, since we do not understand complex interactions of layered infrastructure. More emphasis on the automation of the assessment of controls is the way forward – humans to set policy, machines to evaluate compliance.

Also, a move up from “checklist” thinking to “systems” thinking is necessary. Checklists are adequate when the number of items is modest and comprehensible, but modern cybersecurity isn’t one of those areas. In fact, most of the components of cybersecurity are not comprehensible in isolation, much less when integrated with the other components. Today, too much compliance testing focuses on those components in isolation because the overall systemic view is so complex. Yet, the truth is that assessing them in isolation can create a false sense of security, since the impact of integrating the components into a system can effectively change the overall outcome.

6. How do these approaches take into account sector-specific needs?

In general, the concept of zone-based network security architecture is universally recognized as a best practice, while the specifics of the zone memberships and inter-zone rules are sector-specific.

7. When using an existing framework, should there be a related sector-specific standards development process or voluntary program?

Yes, definitely. For one example, PCI, while a solid approach for cardholder data, does not apply directly to power generation, which is where NERC CIP contributes. The same is true across other critical infrastructure providers. The bulk of what is needed is well understood, and common. Studies such as the Verizon Data Breach Investigation Report (released annually) clearly show that the majority of breaches use well-understood techniques, for which the mitigation strategies are well known. The problem is that we fail to implement these controls consistently – the chances of overlooking one exception are too high in a vast infrastructure, especially when relying on human diligence in audits and assessments. The attackers use automation, “twisting doorknobs” on a grand scale, locating any missed weak spots. This is the core idea that is not sector specific – that we need automation of the assessment of our defensive readiness. The sector-specific adjustments are details, such as the kinds of assets to protect (databases are not the same as trading floors, and neither resembles a nuclear

power plant), and the appropriate controls. What all sectors share is complexity, and a need for automated evaluation of attack readiness.

8. What can the role of sector-specific agencies and related sector coordinating councils be in developing and promoting the use of these approaches?

Mapping concepts like “zone definition”, “zone defense”, and “compensating controls” from a general framework down to a specific sector-required set.

9. What other outreach efforts would be helpful?

This section intentionally left blank

Specific Industry Practices

In addition to the approaches above, NIST is interested in identifying core practices that are broadly applicable across sectors and throughout industry.

NIST is interested in information on the adoption of the following practices as they pertain to critical infrastructure components:

• Separation of business from operational systems;

• Use of encryption and key management;

• Identification and authorization of users accessing systems;

• Asset identification and management;

• Monitoring and incident detection tools and capabilities;

• Incident handling policies and procedures;

• Mission/system resiliency practices;
• Security engineering practices;
• Privacy and civil liberties protection.

1. Are these practices widely used throughout critical infrastructure and industry?

Restrictions on access are essential to the protection of critical infrastructure. However, typical industrial practice without regulation tends to be weak. Organizations may enforce some basic authentication, VPN access, application login, and some logging. However, adoption is expensive and complex – profit-making companies can often see such efforts as money sinks. Even if technology is acquired, the personnel to run it and keep protections in place are expensive and increasingly hard to come by. (It’s a common remark that Information Security has “negative unemployment” as an industry

– there are not enough trained experts to go around.) As a result, technologies are often only partially deployed, ineffectively configured, and weakly monitored. Regulation has been one of the few ways to “level the playing field” to ensure that profit-making companies can meet standard levels of protection without worrying that their less secure competitors, while taking a risk, are also beating them. After all, our capital-centric system actively encourages risk-taking behavior!

2. How do these practices relate to existing international standards and practices?

No comment

3. Which of these practices do commenters see as being the most critical for the secure operation of critical infrastructure?

The most direct way to limit risk to critical infrastructure is to simulate attacks and identify weak points. None of the specific controls listed are a “silver bullet”–a system with excellent user identification, but poor or non-existent log monitoring, is not going to survive. Balance is required–all controls cost money during design/build, but even more during operation, and resources are finite. Therefore it’s imperative to balance spending–to invest resources where they are most needed, to shore up weak defensive areas that make it too easy for cyber attackers to cause damage. This is why integrated testing of defenses is essential–without it, organizations are too easily distracted by one or a few favorite security offerings.

Hence the ideal framework does not establish proscriptive lists of “technology X must be used”–such lists are inflexible and are rapidly outdated. (Even PCI DSS–one of the most technically sophisticated regulatory specifications–has to be updated frequently to keep pace with shifting attackers and shifting defensive methods.) Rather, the framework should encourage automation and outcome-centric assessment of defensive readiness. Organizations using the framework should be able to justify their defensive activities by measuring the ease of compromise of their infrastructure and driving this measure down.

4. Are some of these practices not applicable for business or mission needs within particular sectors?

The framework should ideally focus on organizational self-assessment of defensive readiness and gap analysis. If the gap analysis never shows a need for a given technology, it need not be adopted. The standard should explain clearly how to test, not which answers give the best answer to the test. The best answer is very likely to vary by sector, for example, banking tends to use extremely up-to-date technology (with equally modern vulnerabilities), while typical power distribution utilities rely on technology several generations older. As a result, a measurement of defensive readiness for banks would indicate greater need for awareness of zero-day or recent vulnerability discoveries, while the same framework and attack simulation approach for utilities may

indicate that network-layer segmentation may have far higher criticality, to eliminate all access to older–and possibly unpatchable–equipment.

5. Which of these practices pose the most significant implementation challenge?

In general, reactive technologies to sense and react after a breach are necessary, but are the hardest to deploy and most expensive to maintain, due to the high level of analyst skill and attention required. Even the best signal-reduction technologies still flag a vast number of incidents to investigate for possible response. This area continues to evolve, but no ultimate answer is yet in sight.

At the risk of sounding simplistic, it's true in this arena that an ounce of prevention is worth a pound of cure. Defensive readiness pays off many times over, most directly in incident reduction and avoidance of compromise. Secondary benefits also accrue–a high level of defensive readiness requires well-mapped infrastructure with detailed knowledge of inventory of assets and roles. This information, prepared ahead of an attack as part of defensive risk assessment, is highly valuable during an attack, speeding up both response and recovery.

6. How are standards or guidelines utilized by organizations in the implementation of these practices?

For most organizations, the security professionals observe far more problems that need to be fixed than can be fixed with the available finite resources. However, their funding is limited, and so prioritization is necessary. Standards and guidelines are an immense assist in the internal budgeting process, helping make clear what constitutes "due diligence" or "sufficient security" or "industrial sector norms". Routinely, technologies wanted by security teams for the inherent security benefits are purchased using justification documents that focus entirely on the compliance or standards impact of the technology. This is why guidance from standards bodies can have such huge impact on which defensive approaches are used.

7. Do organizations have a methodology in place for the proper allocation of business resources to invest in, create, and maintain IT standards?

Maintenance of IT standards is by far the hardest part of this area. Most organizations have copious amounts of policy, documented in binders on shelves and in document control systems. Standards are often a mix of internal guidelines and external frameworks. However, they are generally not effective in production unless there is clear payoff in doing so. As a result, many internal standards achieve little traction. External standards, especially ones known to be adopted by sector competitors, are the most fruitful, because they can escape the concern of over-spending compared to peers.

8. Do organizations have a formal escalation process to address cybersecurity risks that suddenly increase in severity?

Most organizations cannot detect that cybersecurity risk levels increase, because they cannot measure them in the first place. This is another key reason why ideal frameworks focus on predictive and proactive measurement of defensive posture. Without a quantified, repeatable methodology for assessment, organizations are at the mercy of hearsay and FUD (fear, uncertainty, and doubt) when it comes to thinking that risk has increased.

Note also that much of the change in cybersecurity risk is self-induced, that is, errors in operations and configuration have repeatedly been shown to be involved in breaches. (For one source, see the annual Verizon Data Breach Investigation Report, where it's clear that the overriding majority of breaches could readily have been prevented through the consistent application of already known controls.)

9. What risks to privacy and civil liberties do commenters perceive in the application of these practices?

Defensive posture assessment requires understanding of three things: the as-built infrastructure, its changing weaknesses, and the varying activities of bad actors. The former two–how the organization is set up and what kinds of vulnerabilities are exposed–can be assessed without any risk to privacy or civil liberties. Fortunately, these two factors are the higher priority in today's critical infrastructure. Our defenses are generally weak and growing in complexity. Automated assessment of this posture, before assessing current activity levels of miscreants, is the most urgent national defensive priority.

10. What are the international implications of this Framework on your global business or in policymaking in other countries?

No comment

11. How should any risks to privacy and civil liberties be managed?

No comment

12. In addition to the practices noted above, are there other core practices that should be considered for inclusion in the Framework?

As emphasized earlier, the most important aspect of the Framework should be its focus on the automation of assessment of defensive posture, instead of whatever array of controls and technologies the organization has found necessary. Automated assessment of likely attack paths is by now a well established practice, with clear benefits in both incident reduction and rapid, accurate response when incidents do occur.

Dr. Mike Lloyd, CTO
Stephen S. Hultquist, CIO
RedSeal Networks, Inc.
2540 Mission College Blvd.
Santa Clara, CA 95054, USA

408-641-2200