The Economics of Cybersecurity: A Practical Framework for Cybersecurity Investment

AFCEA Cyber Committee[1]

The Problem

Conservatively, $15 billion is spent each year by organizations in the United States to provide security for communications and information systems.[2] [Market Research 2013; Gartner 2013] Despite this investment, the economic impact of cyber breaches continues to be very large. In 2009, President Obama estimated the economic impact of cyberattacks at over $1 trillion/year or about 6% of the Gross Domestic Product (GDP) of the United States. [Obama 2009] More conservative estimates of the economic impact from cyberattacks put the losses at a hundred billion dollars per year. [Reference: Friedman 2011; Gorman 2013] Regardless of the exact dollar impact, loss due to attacks against cyber systems is clearly a significant drain on the US economy and organization resources.

Most government and private sector organizations have been increasing their spending for protecting their cyber assets. These investments constitute about 5% of the information technology budget for most companies.[Gartner 2013] Despite the large and growing investment in cyber security (6% compound annual growth rate according to Market Research), many organizations struggle to determine how much to invest in cybersecurity as well as where these investments should be made. The security company Symantec observed that targeted cyberattacks increased 42% in 2012 over 2011. [Symantec 2012] Verizon's 2013 Data Beach Investigations Report (DBIR) report notes that three quarters of the 2012 attacks employed attack techniques that were low or very low difficulty to launch. [Verizon 2013]

> *The fact that the majority of reported damaging attacks have little sophistication leads to the logical question of why the $15 billion annual investment in cyber security protection by United States Organizations is not more effective in successfully blocking these unsophisticated attacks.*

[1] This paper is the result of collaboration among the members of the Economics of Cybersecurity Subcommittee of the AFCEA Cyber Committee and a set of outside advisors. The principal author is John Gilligan. Other contributors include: Robert Dix, Charles Palmer, Jeffrey Sorenson, Tom Conway, Wray Varley, Gary Gagnon, Robert Lentz, Kenneth Heitkamp, Phillip Venables, Alan Paller, Jane Holl Lute, and Franklin Reeder.

[2] In this paper, the terms 'communications and information systems', 'cyber systems' and 'information technology' are used to describe the wide range of general purpose computing and communications systems as well as embedded computing technology (e.g., computing capabilities embedded in machines, control systems, and mobile devices).

As an additional statistic, Verizon notes that 66% of successful breaches were not discovered until a month or more after the breach occurred and that most of the breaches (69%) were discovered by external sources. Based on these findings, reasonable questions for senior executives to ask include the following:

1. Is the investment by the United States in cybersecurity being appropriately applied?
2. Should organizations invest more or less in order to provide adequate security?
3. Where should organizations invest to gain the biggest economic return?

This paper attempts to help answer these questions in a manner that provides a clear focus and priorities for cyber security investments. In particular, the paper outlines an easily-understood framework for focusing investment in cybersecurity. In turn, examination of the framework is used to develop a set of straightforward investment principles that can guide an organization's cyber security investment strategy as well as help assess incremental investment in cyber protection capabilities.

Background

Over the years, a number of organizations have developed models to help guide investment by organizations in cyber security. Despite these efforts, there is no generally accepted model or set of investment principles that organizations can use to guide cybersecurity investments. A partial explanation for this state is recognition that quantitative economic analysis efforts in the field of cybersecurity have been hindered by the lack of solid data about the number and impact of cyberattacks. Even though organizations are increasingly willing to acknowledge attacks, many cyber-attacks go unreported. Moreover, as the Verizon's 2013 DBIR notes, most attacks are discovered weeks or months after the attack and, more importantly, they are most often discovered by external parties. This would seem to strongly suggest that data presented in cyberattack reports are just the tip of the iceberg. The appendix to this paper provides a brief summary of some economic models that have been developed to assess investments in cyber security.

A Cybersecurity Framework

The objectives in developing a framework for use in informing executives regarding the key economic decisions regarding cybersecurity were the following: 1) ease of comprehension, and 2) readily applicable by both small and large organizations and across a range of communications and information systems. The framework is the result of an analysis of the strengths and weaknesses of existing economic models as well as identification of the key priorities relevant to cybersecurity investment decisions. The first priority that is important to

investment decisions relates to the relative sophistication of a cyber threat.[3] As noted in the Verizon 2013 DBIR, the vast majority of reported successful cyber-attacks are relatively unsophisticated. Therefore, addressing these unsophisticated but damaging threats should be the first investment priority for organizations. The second key priority involves the relative criticality of the mission being supported and/or the data being stored or transported. Organizations should focus their cyber efforts to ensure that mission critical functions and data are protected. Therefore, their second priority should be to invest in cyber security measures that achieve this objective.

These two priorities can be displayed using the simple graph shown in Figure 1.

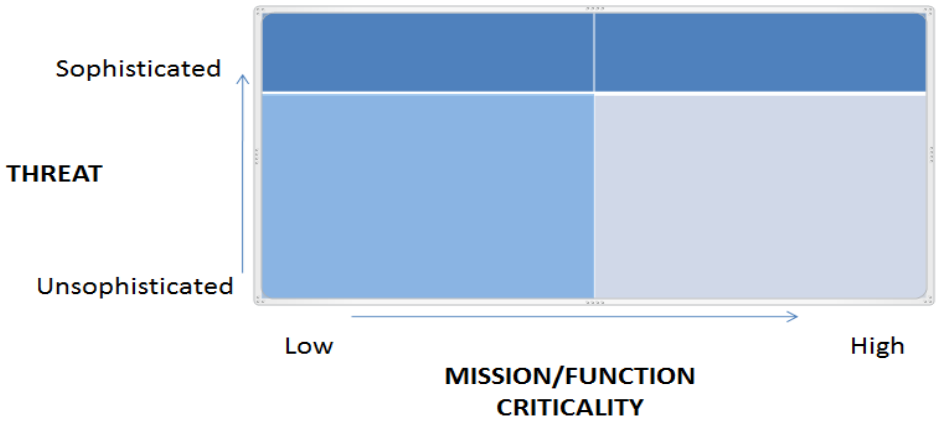

Figure 1: Foundation for Cybersecurity Framework

The horizontal dimension of the framework arrays mission and data criticality from 'low' (for example, an administrative function that would have low to no mission impact if unavailable or compromised) to 'high' (for example, core missions or critical data of the organization which if unavailable or compromised would effectively shut down the organization). The vertical dimension reflects the range of sophistication of cyber threats from 'unsophisticated' to 'sophisticated'. Examples of this range of threats might be, on the one end, exploitation of known but unpatched vulnerabilities perhaps using automated tools found on the web to, on the other end, exploitation of maliciously inserted or previously undiscovered vulnerabilities by an organized crime syndicate or a Nation State.

[3] A threat to cyber systems refers to persons who attempt unauthorized access to a cyber system device and/or network. This access can be directed from within an organization by trusted users or from remote locations by unknown persons using the Internet. Threats can come from numerous sources, including hostile governments, terrorist groups, disgruntled or poorly trained employees, and malicious intruders.

In analyzing Figure 1, a couple of important observations relating to investment decisions can be seen. Recalling that the 2013 Verizon DBIR found that 75% of the cyber-attacks are low or very low sophistication, we observe that investments that would reduce exposure to these threats (i.e., the lower portion of the graph in light grey) have the benefit of addressing the majority of attacks. It is acknowledged that attacks using less sophisticated techniques may or may not have the most significant economic impact. However, they do constitute the largest number of attacks, and experience has demonstrated that a potential attacker will usually take the path of least resistance. In addition, with regard to the other dimension of the figure, it is relatively easy to agree that organizations are interested in protecting their most critical data and functions as a first priority, and before making investments in protections for data and missions that have little or no mission impact. Therefore, investments that can protect higher impact data and functions shown on the right portion of the figure become an investment priority for an organization.

Based on the observations just noted a more robust view of the Cybersecurity Framework is shown in Figure 2 below.

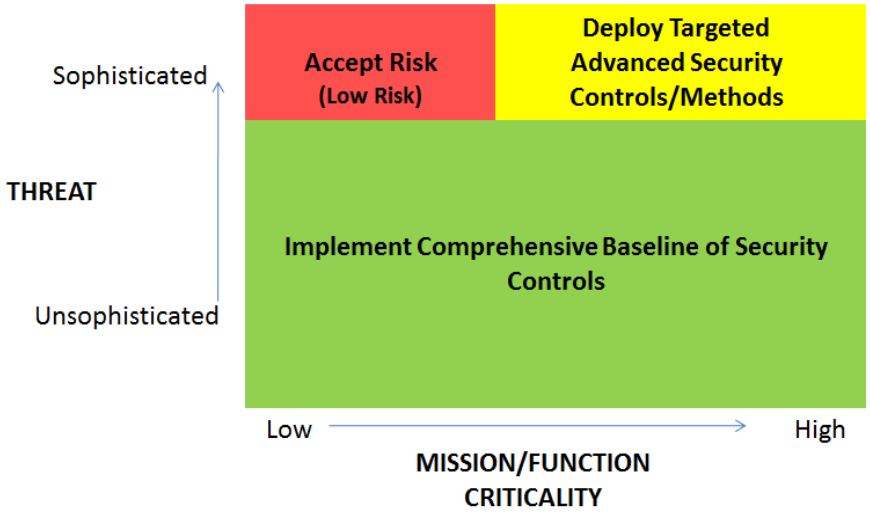

Figure 2: The Cybersecurity Framework

The Cybersecurity Framework shown in Figure 2 highlights several important principles that should guide cybersecurity investment decisions. The first investment principle is based on the lower portion of the figure and is as follows:

Investment Principle #1: Implementation of a comprehensive baseline of security controls that address threats that are of low to moderate sophistication is economically beneficial.

This investment principle is based on several factors. First, substantial research from the US National Security Agency (NSA) as well as other government and private sector companies [Lewis 2013] has shown that a relatively small number of security controls that avoid, counteract, or minimize security risks can be effective in protecting organizations against the overwhelming majority of cyber threats that are low to moderate sophistication. There are several examples of such a foundation or baseline of security controls including the Critical Security Controls (CSC) developed as a collaborative effort in the United States [CSIS 2013] and the Australian Department of National Defence's (DND) Top 35 controls. [DND 2012] Interestingly, recent analysis done by the Australian DND has shown that a very small subset of the controls, specifically the first four of the DND Top 35 controls listed below, are effective in preventing 85% of the targeted cyber intrusions addressed by the DND Defence Signals Directorate [DND 2013]. Verizon's 2013 DBIR also endorses organizations deploying these baselines of controls as appropriate against 75% of threats.

Investment Principle #1 also asserts that a baseline of security controls is economically beneficial. That is, a comprehensive baseline of controls has a sound economic basis. This economic benefit typically extends beyond security. Descriptions of baseline controls sometimes refer to them as "implementing sound management practices". Upon examination, the Critical Security Controls and the DND Top 35 controls are found to be focused on ensuring that sound information technology management disciplines are enforced. As an example, the Top 4 of the DND controls consists of the following disciplines:

1. Restricting user installation of applications (called "whitelisting")
2. Ensuring that the operating system is patched with current updates (especially security updates)
3. Ensuring that software applications have current updates
4. Restricting administrative privileges

Each of the DND Top 4 controls are fundamental configuration and system management principles needed for security, but they also have other benefits that help ensure high operational availability and ease of administration of networks and systems. These controls would not be implemented by an organization solely to improve security. These controls should be implemented in every communications and computing system environment to improve both operational effectiveness and cost efficiency, in addition to improving security. It has been well established in the information technology community that systems and networks that implement sound systems and network management principles (i.e., "good cyber hygiene") are less expensive to operate and have enhanced availability. The Cybersecurity

Framework highlights that these sound cyber management disciplines also significantly improve cybersecurity.

An example of where implementing basic security controls has been shown to be beneficial economically as well as from a security perspective is an Air Force project implemented in 2005-2006. During the project, the United States Air Force deployed a standard configuration of Microsoft's operating system, browser, and Office Suite to over 500,000 Air Force desktops and laptops. Included with the standard configuration was the use of available Microsoft capabilities that supported the automated distribution and installation of system updates and patches. A parallel effort also addressed Microsoft Windows-based server configurations. The motivation for the project was that the NSA had identified that improperly patched software was the dominant cyber threat to the Air Force (i.e., greater than 70% of attacks). [Reference – NSA presentation to John Gilligan, then Air Force Chief Information Officer] The Air Force project included testing of all applications against the standard configuration and minor modifications to some software. At the completion of the project, the Air Force was able to demonstrate improved resilience and agility against cyber threats as a result of (1) the securely configured system, (2) restricting system administration privileges to a small number of authorized users, and (3) automation of updates and patches. In addition, however, the standardization of configurations resulted in the following additional benefits: improved overall system availability; far fewer help desk calls; and a significant reduction in system problems and resulting outages. Finally, and most importantly for this paper, the project's enforcement of standard, more secure operating system configurations as well as the automation of configuration patches and updates resulted in and lower operating costs achieved through a significant reduction in several thousand system administrators. [Heitkamp 2013] In this case, the operating cost reductions easily dominated the cost of the project implementation. Said another way, the disciplined redeployment of the standard Microsoft core software capability, limiting system administration privileges, and automated update capability resulted in an immediate positive return on investment (ROI) even before considering the positive economic impact of reducing cyber breaches.

An additional insight into economic costs of implementing baseline security controls was found when looking at organizations that had separate, parallel organizations and associated tools for ensuring security and (separately) for ensuring efficient operation of an organization's systems and networks. A superficial examination of this type of organization and investment structure might provide the impression of increased confidence that the investment in a separate staff and tools for security and for IT system operations was a sound one. However, upon examination, t was found that the security organization typically invests in baseline security controls and processes that are mostly overlapping and redundant to the baseline controls (e.g., for configuration management, patching, asset tracking, etc.) being implemented by and

the organization's system and network operations staff. Based on observations in the course of preparing this paper, this should not be a surprising finding. Of concern, however, is the observation that these overlapping controls can result in increased complexity and gaps in security effectiveness that often result in a weakening of the collective set of baseline controls and resulting in reduced effectiveness against even unsophisticated attackers. Moreover, the redundant efforts and investments in baseline controls are expensive and almost completely wasteful.

Examination of the top portion of the graph in Figure 2 addressing more sophisticated attacks leads to additional investment insights. Specifically, the observation suggests the need to focus investments when addressing sophisticated threats. Not all of an organization's capabilities and information are of equal value. And, not all security risk mitigations are of equal benefit. This observation leads to the following investment principle:

Investment Principle #2: Focus security investment beyond the baseline controls to counter more sophisticated attacks against the functions and data that are most critical to an organization.

This investment principle would appear to be self-evident and not merit additional discussion. However, interactions with various organizations in the course of developing this paper found that while most organizations would agree with this principle, they often did not, in fact, have a process for implementing it. In practice in these organizations, security controls were applied equally against all data and functions. While IT organization leaders acknowledged that they should put priority on protecting their more critical functions and data, it became quickly apparent in some cases that the architecture and governance of their organization's cyber systems had not been structured to permit cyber protection efforts to focus on the most mission critical elements. It was an embarrassing realization—a true "aha" moment for some of the organizations.

As a result of discussions with a variety of organizations, a corollary to Investment Principle #2 is the need to establish within an organization's IT architecture the ability to segregate or partition mission critical from less critical functions and data. This approach results in protection "enclaves"[4] where different security protection schemes can be implemented following sound economic principles for each enclave. While the concept of enclaves is simple

[4] Protection enclaves can be implemented through configuration of cyber systems and data to group those with similar protection requirements. It is noted, however, that few organizations have undertaken the prerequisite effort to analyze their systems and data to permit identification of protection requirements. The US Department of Defense has implemented a coarse-grained version of enclaves with the DoD classification and related clearance system.

intellectually, the determination of what systems and information should be assigned to the protection enclaves is not an easy task, nor is it one that is primarily a task for the IT team. An initial investment is required to determine mission criticality or financial value of systems and data. Once an organization is able to identify logical groupings of systems and data, a phased implementation of the enclave architecture is suggested focusing first on the most critical functions and data. Moreover, enclaves can be added, expanded, or reconfigured over time as an organization make decisions to extend more robust security controls and surveillance to a larger set of data and functions as well as to respond to ever changing threats.

An additional Investment Principle addresses the need to specifically accept risk:

Investment Principle #3: For sophisticated attacks, an organization should accept the security risk of not protecting functions and data that are of lowest impact to the organization's mission and where cost exceeds benefits.

The reason behind this investment principle is that protection of low impact systems and data from sophisticated threats will require greater investment than the benefits that would be achieved. Unlike the baseline of controls addressed in Investment Principle #1, investment to address more sophisticated threats can be quite costly and not economically justifiable for low impact information and systems. However, an organization might expand their protection capabilities over time using a recurring return on investment (ROI) analysis. It is important to reassess the ROI analysis on a recurring basis for two reasons: (1) cyber control investment costs will normally decline over time as solutions become more widely available or integrated into standard commercial systems; and (2) the level of mission criticality for specific functions or data may change over time. The desired result of the recurring analysis would be to include in an organization's cyber protection realm additional mission functions and information as shown by the direction of the arrow in Figure 3 below.

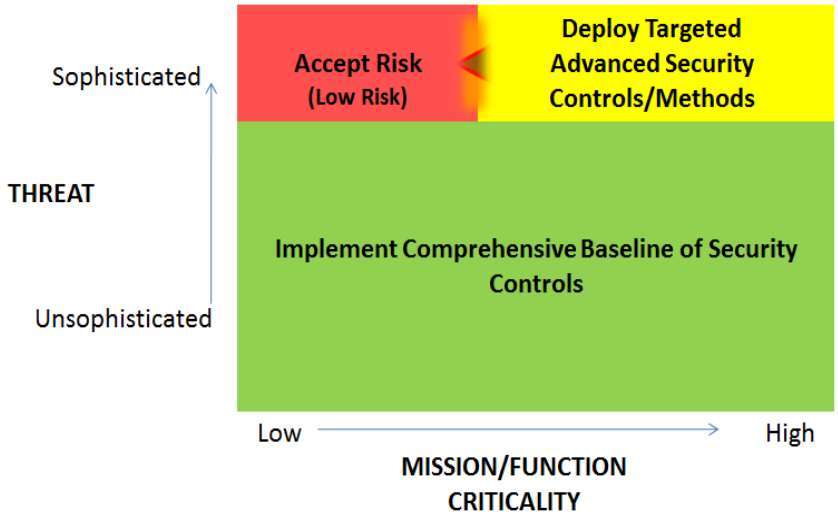

Figure 3: Reducing Risk Acceptance

Summary

The objective of the paper was to answer the following three questions:

1. Is the investment by the United States in cybersecurity being appropriately applied?
2. Should organizations invest more or less in order to provide adequate security?
3. Where should organizations invest to gain the biggest economic return?

With regard to the question of whether cybersecurity investment is being appropriately applied, based on the analysis conducted for this paper, one can conclude that much of the approximately $15 billion spent by US based organizations per year could be better focused. This is supported by the observation that the vast majority of attacks are originating from unsophisticated threats and that a straightforward baseline of security controls has been shown to be effective against these attacks. As shown in this paper, these baseline controls are not enormously costly, especially when considering that virtually all of the baseline controls are required to ensure efficient and effective operation of an organization's networks and systems. In some instances, baseline controls can actually result in significant savings as shown in the Air Force example. Implementing a baseline of controls reflects sound management of an organization's information technology. Similarly, the absence of an effective baseline of controls would appear to be an indication of inadequate management of an organization's cyber resources. Moreover, even for controls that help address sophisticated attacks, the research found that over time the cost of many of these more sophisticated controls have shown a consistent pattern of significantly reduced cost, thus permitting them to be economically implemented by a larger set of organizations.

With regard to the question of whether an organization should invest more to provide adequate security, the Cybersecurity Framework provides a "roadmap" for incremental investments. An assertion that can be made based on the research conducted is that most organizations do not need to invest significant additional resources to implement a comprehensive baseline of security controls (i.e., to implement the lower portion of the framework). Most organizations have already purchased commercial tools that can effectively implement the baseline controls with minimal additional investment. The experience of John Streufert as CISO at the State Department clearly demonstrated this for a large enterprise. [Streufert 2011] In short, most organizations do not need more tools but better discipline in the effective implementation and use of the already-purchased tools that implement baseline security controls. Most of the investment to implement the baseline security controls and associated governance is a one-time cost. Thereafter, recurring sustainment of the baseline security controls actually costs less due to the inherent higher levels of standardization and use of automation. This is admittedly not an easy task for larger organizations. However, it is a task that puts a premium on management and leadership abilities rather than on economic investment.

The Cybersecurity Framework and the Investment Principles are also useful in helping organizations focus additional investments that address more sophisticated threats and produce sound return on investment. The Framework reflects the practices of leading organizations to focus investments against sophisticated threats to ensure cyber protection for critical missions and data. Moreover, the encouragement to specifically accept risk in low impact areas and to implement protection enclaves reflects both sound economic as well as effective cyberdefense strategy. Finally, the Framework and Investment Principles are easy to understand and can be used by organizations of any size.

A subsequent paper, "Extending the Cybersecurity Framework", addresses extensions to the Cybersecurity Framework in order to increase focus on more sophisticated threats. The extended framework relies on the necessary foundation described in this paper.

As a closing note, it is recognized that his paper focused primarily on technical measures for cyber security. These technical controls must be a part of a comprehensive cybersecurity program that also addresses trained people, adequate policies, and appropriate processes.

Appendix A: Models to Assess Investment in Cyber Security

This Appendix summarizes models and strategies examined during the research that have been developed to guide organizations in investing in cyber security.

At the University of Maryland, a cross disciplinary team of faculty members has had a decade-long focus on understanding the economics of cyber security investments. Professors Lawrence Gordon and Vernon Loeb have published several papers outlining the results of their research in this area. [Gordon 2002] Their analysis has looked at areas such as the impact of investments in cyber security measures, the cost of responding to security breaches, as well as the impact of a publically-acknowledged security breach on stock valuation. They developed an economic model for cyber security based on an analysis of organization spending as well as marginal effectiveness and return on investment of cyber investments. In the absence of other comparable economic analysis, the Gordon-Loeb model has become the "gold standard" in the area of cyber economic models. Among the many findings of their research, the Gordon-Loeb model makes two important conclusions:

1. Incremental additional investment in security provides additional benefit by reducing the potential of successful attack up to a point. Beyond this point, there is diminishing (or no) additional benefit for additional investment.
2. An organization should not invest more in cybersecurity protection measures than 37% of the expected potential loss due to successful cyberattack.

While the Gordon-Loeb investment model has merit, it also has some inherent limitations. The model does an excellent job of describing general principles for cyber investments. However, the model is very complex for practical application by most organizations. Moreover, it has some assumptions that, while essential for the model, were found by this review to limit its broad applicability. The conclusion reached is that while the Gordon-Loeb model reflects solid academic research, it may not be well suited for direct implementation as an investment model by an organization.

The National Institute for Standards and Technology (NIST), the United States government organization responsible for providing standards and guidance in the area of cyber security (and many other areas), has also developed a model for guiding investment in cybersecurity countermeasures. This model is elaborated in a series of guidebooks. Most fall under the "800 Series" publications. Specifically, NIST Special Publication 800-37 [NIST 2010] and 800-53 [NIST 2009; NIST 2010(2)] define the cybersecurity Risk Management Framework (RMF) including a method for assessing the implementation of controls to mitigate risk. The underlying approach for determining the proper investment according to the guidebooks is to identify the appropriate security controls that mitigate an organization's cybersecurity risks as determined

by a thorough risk assessment. The NIST guidebooks provide instruction on this process. Federal agencies are required to implement the NIST guidance; many non-government organizations have also found benefit from implementing NIST guidance. The NIST Risk Management Framework does not attempt to characterize risks or benefits in economic terms.

The NIST Risk Management Framework (RMF) provides extensive guidance (well over 2,000 pages). However, the foundation step of the NIST approach is for an organization to perform a thorough cyber risk assessment. That is, the NIST Framework is based on the assumption that organizations are able to adequately perform an accurate assessment of their cyber risk posture. In the RMF, residual risk (R) is calculated as a function of threats (T) exploiting cyber vulnerabilities (V) offset by countermeasures (C) or $R = (T \times V) - C$. Unfortunately, many (perhaps most) organizations have difficulty in performing an adequate risk assessment. The reasons for this are primarily twofold. First, most organizations have limited insight into the detailed nature of the threats (T) against their cyber environment. In fact, many threats are classified and not broadly shared due to potential impact on areas such as stock valuation and brand name. In addition, the extraordinarily complex nature of modern cyber systems makes determining vulnerabilities (V) a very complex effort requiring extensive knowledge of system of systems architectures and sophisticated systems engineering expertise. As a result, many organizations that employ NIST's RMF struggle with implementation. The conclusion reached by this analysis is that, despite the comprehensive nature of the NIST Risk Management Framework and its unassailable underlying logic, it has not proved practical for organizations who are struggling to determine where to invest in cyber security and, in particular, how much investment in cyber security is warranted.

A number of private sector organizations have independently developed cyber investment strategies. For the most part, these are proprietary and not available outside the company. Interviews with corporate leaders, for example CEOs and members of several boards of directors, has indicated that even companies that have developed formal cyber investment models continue to struggle with assessing whether the investments are properly focused and are sound economically. Nevertheless, the appropriate investment in cybersecurity is increasingly becoming an important topic for the corporate boardrooms. [Ref PWC and NACD]

Recently, the Chief Information Security Officer (CISO) of one very large multinational company indicated to a gathering of corporate directors at a US DoD Defense Security Service (DSS) conference that they had invested more than $100M in cybersecurity measures over the past three years. The catalyst for the investment was a significant security breach resulting in loss of intellectual property and significant damage to corporate reputation as well as stakeholder confidence. The cybersecurity capabilities implemented by this organization seemed to be appropriate for the organization. The CISO demonstrated the effectiveness of the capabilities

through an analysis of several recent cyberattacks. The organization had also developed a set of metrics to show their board of directors that the $100M investment had resulted in significantly fewer cyber security breaches and a commensurate large reduction of post breach recovery costs. The metrics were used to support the conclusion that the investment had been sound economically when compared to their prior experience of cleanup costs following cyber breaches.

Research conducted in the course of developing this paper found few other models that could assist in answering the four economic questions raised in this paper. And, as noted above, even the best current models have some significant limitations. The remainder of the paper will outline a recommended cybersecurity framework for determining what investments are economically reasonable and where those investments should be made.

References

Market Research 2013: United States Information Technology Report Q2 2012, April 24, 2013

Gartner 2013: "Gartner reveals Top 10 Security Myths", by Ellen Messmer, NetworkWorld, June 11, 2013.

Obama 2009: Remarks by the President on Securing our Nation's Cyber Infrastructure, The White House East Room, May 30, 2009.

Friedman 2011: Economic and Policy Frameworks for Cybersecurity Risks, Allan Friedman, Center for Technology Innovation at Brookings, July 21, 2011.

Gorman 2013: "Cybercrime Costs Put at $100 Billion", Siobhan Gorman, The Wall Street Journal, July 23, 2013, p. A4.

Symantec 2012: Symantec Intelligence Report, 2012 Updated June 25, 2013.

Verizon 2013: 2013 Data Breach Investigations Report, Verizon.

Lewis 2013: Raising the Bar for Cybersecurity, James Lewis, Center for Strategic and International Studies Technology and Public Policy Division, February 12, 2013.

CSIS 2013: CSIS: 20 Critical Controls: Critical Controls for Effective Cyber Defense – Version 4.1. March, 2013.

DND 2012: Strategies to Mitigate Targeted Cyber Intrusions, Australian Government, Department of Defence Intelligence and Security, October 2012.

DND 2013: Top 4 Strategies to Mitigate Targeted Cyber Intrusions, Australian Government, Department of Defence Intelligence and Security, April 2013

Heitkamp 2013: Notes developed by Kenneth Heitkamp USAF (Retired) to summarize benefits achieved from Air Force Standard Desktop Project.

Streufert 2011: The Government Model: The State Department's Approach to Cybersecurity is so Innovative and Effective that Companies are Clamoring to Copy it. By Siobhan Gorman, the Wall Street Journal, September 26, 2011.

Gordon 2002: The Economics of Information Security Investment, Lawrence Gordon and Martin Loeb, University of Maryland, ACM Transactions on Information and Systems Security, November 2002.

NIST 2010: SP 800-37 Rev. 1. Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach, February 2010.

NIST 2009: SP 800-53, Rev 3. Recommended Security Controls for Federal Information Systems and Organizations. August 2009.

NIST 2010(2): SP 800-53 A, Rev 1. Guide for Assessing the Security Controls of Federal Information Systems and Organizations, Building Effective Security Assessment Plans.