international risk governance council

WORKSHOP REPORT

CYBER SECURITY RISK GOVERNANCE

Swiss Re Centre for Global Dialogue, Zurich, Switzerland
29 – 30 October 2015

Sponsors

IRGC would like to express its gratitude for the support provided by its partners and sponsors. The principal partner was the Global Risk Center of Marsh & McLennan Companies, along with Oliver Wyman.
Support was also received from Swiss Re and AXA Technology Services.
The workshop was hosted by the Swiss Re Centre for Global Dialogue in Zurich, Switzerland.

MARSH & McLENNAN COMPANIES

OLIVER WYMAN

Swiss Re

AXA TECHNOLOGY SERVICES
redefining / service

“Organizations need to quantify cyber risk to enable better decisions about security investments and business management. Using the current best practices and new research initiatives, organizations can close the security gap by implementing the best controls and developing new tools to give defenders an edge in cyber security. "

Organizations are increasingly concerned about threats to data confidentiality, integrity and availability. When data are compromised and critical infrastructure and services are impacted, the cost to organizations and damage to trust and reputation can be huge. Most organizations use pragmatic solutions to address cyber-attacks, but much uncertainty remains about whether such solutions are able to address threats before they cause too much damage, whether the quantitative estimate of the potential impact (i.e. the risk) is accurate, whether investments for the protection of important assets are appropriate, and whether overall governance of the decision about cyber risk management is optimal.

On 29-30 October 2015, the International Risk Governance Council (IRGC) convened a workshop with Chief Risk Officers, Chief Information Security Officers, researchers in academia, government representatives, as well as experts from industry and insurance companies to discuss cyber security risk.

This report summarizes some of the discussions at the two-day workshop, together with other considerations in line with the theme. It was reviewed by some of the workshop participants, including Maya Bundt, Eric Durand, Marie-Valentine Florin, Claus Herbolzheimer, Marshall Kuypers, James Larus, Eireann Leverett and Richard Smith-Bingham. Neither participants nor reviewers were asked to endorse the content of this report. The responsibility for the final content of this report rests with the International Risk Governance Council.

Terms, organizations and cyber-attacks mentioned in this report:

Anthem	2015. Cyber-attackers executed a sophisticated attack to gain unauthorized access to one of Anthem’ parent company’s IT systems and have obtained personal information relating to consumers employees.
BSI	German Federal Office for Information Security
CIO	Chief Information Officer
CISO	Chief Security and Information Officer
CRO	Chief Risk Officer
ENISA	European Network and Information Security Agency
Experian	2015. Hackers stole the personal details of up to 15m T-Mobile US customers, because of a data breach at world’s largest credit checking company Experian.
German steel mill	2014. A blast furnace at a German steel mill suffered "massive damage" following a cyber-attack on the plant's network.
Home Depot	2015. Home Depot was the latest in the chain of large companies under a cyber-attack targeted at their payment terminals, where a security breach left approximately 56 million credit and debit card numbers exposed
IP	Intellectual Property
Michaels	2013-2014. Data on three million of the Michaels Stores customers’ payment cards were stolen in a breach over several months.
Netflix	Leading Internet television network
RSA	Provider of intelligence-driven security solutions
SANS Institute	Institute for information security training, certification and research
Snowden, Manning	Edward Snowden copied classified information for public disclosure in 2013 without prior authorization; Bradley Manning delivered U.S. government documents to individuals not entitled to receive them, in 2009-2010.
Sony “hack”	2014. Devastating cyber penetration that resulted in the release of confidential data.
Target	2013. Approximately 40 million credit and debit card accounts were impacted in a credit card breach.
Visa	Global payments technology company
Vodafone Greece	2004. Individuals either penetrated the network from outside or subverted it from within, and reprogrammed the software at the heart of the phone system, to monitor mobile phones.
Wyndham hotels	2008-2010. Recurring data breaches that provided a window into the regulatory risks and costs associated with a data breach.

Contents

Figures

Introduction: Cyber Security and Risk in Perspective

Cyber security was once managed with a prevention mentality, meaning that defense was primarily focused around denying attackers access to a system. As time progressed, security professionals learned that prevention was insufficient, given the large number of attackers, vulnerabilities, and the well-known rule that an attacker only needs to succeed once, while the defender has the impossible task of being perfect every time.

IT networks are no longer managed using the 'castle' model, but are recognized as similar to modern cities with roadways, seaways, and airways connecting individuals outside and inside the network. Securing a complex system involves using new strategies, such as incorporating detection and response along with prevention to manage networks. Overall, security should be approached with a cyber resilience mindset instead of just cyber defense.

Several other changes are occurring throughout the cyber risk environment. Executives have slowly moved from a mentality of 'securing' a system to a mentality where cyber risk is continuously managed. At the same time, cyber risk is no longer treated in isolation, and the connection of cyber security to core business operations is now appreciated by more executives. Cyber risk can no longer be compartmentalized, but must be incorporated into business processes and balanced with business opportunities in a holistic framework. Finally, risk assessment methods continue to improve. New research can quantify cyber impacts in dollar terms and can be used to communicate risk, prioritize security safeguards, and allocate resources. These changes are ushering in a new era of cyber security for organizations (i.e. companies, government bodies and other institutions), but challenges still remain.

Security and Risk Management

Industry has evolved to the current state of cyber security through observing the use and failure of different strategies over time. Security is not a binary outcome (a system is secure or not secure), but a continuous struggle of attackers and defenders. Security is also not something that is achieved, but instead something that is managed. Other domains take a similar perspective, such as the safe construction industry, which rates the effectiveness of jewelry safes by time (e.g. a 30-minute safe, or a 60-minute safe).[1] This directly addresses the issue with achieving a perfectly secure safe, which is impossible. However, a defender can increase the time that it takes for an attacker to gain access and if an alarm system enables a police response to occur quickly enough, the attackers may not be successful. It is useful to view cyber security in a similar way, where detection and response capabilities are both recognized as critical components to security.

A risk-driven strategy involves understanding and protecting assets while optimizing scarce resources. The risk-driven mentality also operates with the understanding that cyber risk cannot be eliminated, but only managed. There is a movement taking place away from cyber

[1] The time is an estimation of how long it would take an attacker of a certain level of sophistication to gain access to the safe.

security and towards cyber risk, which involves the management of a probabilistic process with uncertain prospects.

Risk and Benefit

Decision makers are also recognizing the need for balancing risk and opportunity. It is impossible to keep all attackers out while maintaining a usable IT system. The cost of cyber security is just one dimension of the cost of doing business and an overemphasis on security can hurt an organization's ability to be successful. More commonly, however, the benefits of IT are taken for granted and the risks involved with IT are not adequately addressed. Organizations may be unaware of what IT assets are valuable to attackers, or how one business unit's decision may introduce vulnerabilities that threaten other projects. Today, IT permeates an organization so completely that cyber risk cannot be separated from business risk.

Risk Assessment Methods

Most organizations have a very limited cyber risk management process, and those that do are very likely to use qualitative tools such as heat maps, which are not precise and can lead to poor decision making. The communication of risk is rarely in dollar terms, meaning that executives are unable to compare cyber risk to other types of risk across the business, and both awareness and accountability for cyber risk are often low.

Quantitative risk assessment methods that use dollar values are beginning to be developed and implemented by some researchers and organizations. These methods leverage historical cyber security incident data and other information sources to provide decision support for resource allocation, security investment prioritization, and to improve risk communication. Insurers are also developing mathematical models to price cyber risk, which could lead to a better understanding of effective best practices as well. More collaboration is needed between industry and academia, but significant progress is being made.

The Path Forward

Researchers in government, industry, and academia are all addressing different aspects of the cyber security conundrum. Still, cyber security is an extremely challenging field. The environment is rapidly changing, the wide range of attackers and targets makes modeling difficult, historical trends may sometimes be predictive (but not always), some risks are difficult to assess (e.g. reputational risk or IP theft), and aggregative risks are difficult to model (e.g. a widespread internet outage).

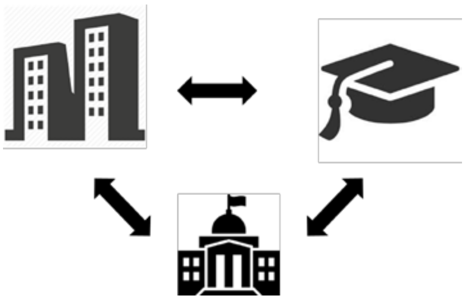

Figure 1: Collaboration between industry, government, and academia is needed.

Overall, there is a consensus that if these issues are not addressed soon, organizations might scale back their use of IT resources. Historically, IT has been incredibly useful from a convenience perspective, but not very secure. There may be a reversal of priorities in organizations that recognize that convenience at the expense of security may not be a good trade-off. However, continued collaboration initiatives may lead to new technologies and better management strategies that enable organizations to continue to operate in a safe and secure environment.

1. The Evolving Threat Landscape

Cyber security pervades all business risks. News coverage is saturated with different types of attacks exploiting vulnerabilities in devices ranging from pacemakers, nuclear power plants, power grids, cars, hospitals, stock markets, and homes. The complexity of software is increasing dramatically as well, evidenced by the increase of lines of code used to send a man to the moon (~350,000 lines of code) compared code used to operate a vehicle (~100,000,000 lines of code). The threat types are also evolving rapidly. 2014 and 2015 saw a number of large data breaches reported in the media including Target, Home Depot, Anthem, and the Office of Personnel Management. The threat landscape is likely to evolve in the future, so that the theft of intellectual property (IP), corporate extortion, and attacks against control systems may become much more common. It is critical for organizations to know their threat landscape.

However, getting a sense of the threat landscape can be challenging. Many publically available reports are published every year, but many are either low quality or have an ulterior agenda aimed at marketing a vendor’s product. Neutral bodies have attempted to address this issue by publishing independent threat assessments. Two such reports are ENISA and the German BSI annual report of the "State of IT Security in Germany".

Top 5 threats listed in two 2014 reports		
Rank	ENISA[2]	BSI[3]
1	Malicious Code	Targeted attacks / APTs
2	Web-based attacks (drive-by exploits)	Cyber-attacks on websites and services
3	Web-application / injection attacks	Insufficient protection of networked industrial control systems
4	Botnets	Increasing complexity of IT systems through integration of mobile devices and external service providers
5	Denial of service attacks	Inadequate patch management and use of out-of-date

[2] ENISA Threat Landscape 2014 Report https://www.enisa.europa.eu/activities/risk-management/evolving-threat-environment/enisa-threat-landscape

[3] BSI “State of IT Security in Germany 2014” Report
https://www.bsi.bund.de/DE/Publikationen/Lageberichte/lageberichte_node.html

		software

The value of these reports should not be ignored, although the publication of the underlying data could contribute to the usefulness of these data. While additional challenges regarding anonymization would be needed to release more details about hacking incidents, the publication of more information would aid researchers who are currently limited by reporting biases, incomplete records, and differing definitions of hacking events.

An illustration of the challenges can be seen in the fact that the ENISA and BSI threat lists do not match.[4] Each report uses a different taxonomy and method for categorizing and assessing threats, since there is no well-accepted standard. Further, organizations need to tailor threat intelligence to their organization and cannot always use the threat rankings of other entities.[5] Organizations also need to consider other more mundane threats to security that are typically not reported, such as the value of information leaked in the form of job postings, press releases, and public forums. For example, an organization might have a job posting for a network administrator with experience managing Juniper firewalls, making it very easy for an attacker to infer aspects of the organization's internal security.

Public reports released by governments still have significant value, and demonstrate strong commitments to cyber security. For example, the publication of the German steel mill hack was generally regarded as a critical step forward in anonymized public disclosure of cyberattacks, and has significant benefits to the industry as a whole.[6] Governments can continue to be a valuable source of intelligence for organizations.

Thinking like a Hacker

Most aspects of cyber security result in an asymmetric advantage for the attacker, and defenders need all the help they can get. One underutilized resource is the information about the attackers that is available on the darknet, which is a collection of forums and markets where buyers can purchase guns, drugs, stolen information, or access to computers. Observing the underground economy for illicit IT services provides information about how and why attacks occur. This is an especially useful perspective for insurers, who can compare the market rates for types of stolen information and use this to estimate impacts to organizations that are hacked (for example, credit cards are typically sold for approximately $12, while stolen information including the name, address, social security number, and other personal information of an individual is typically approximately $30).

[4] The threats also overlap on each list. For the ENISA list, decision makers may have a difficult time prioritizing these threats since botnets are used for denial of service attacks.

[5] For example, consider an organization that has no industrial control systems. The 3rd threat from BSI is irrelevant to that organization.

[6] Kim Zetter, “A Cyberattack Has Caused Confirmed Physical Damage for the Second Time Ever”, Wired, http://www.wired.com/2015/01/german-steel-mill-hack-destruction/

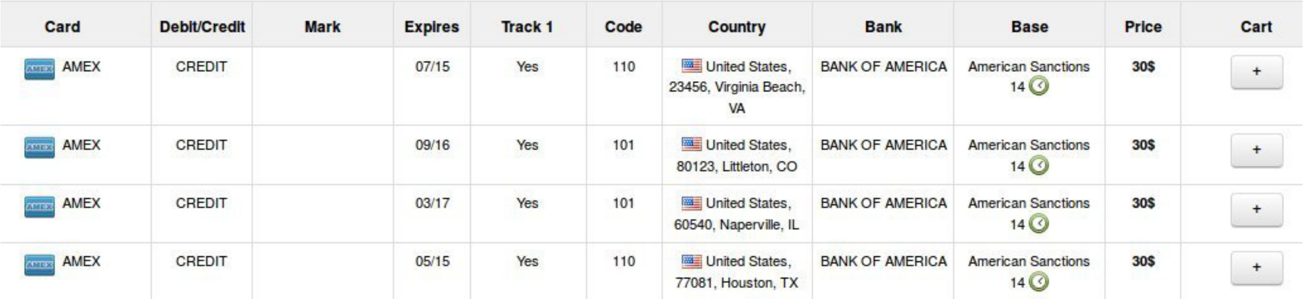

Card	Debit/Credit	Mark	Expires	Track 1	Code	Country	Bank	Base	Price	Cart
AMEX	CREDIT		07/15	Yes	110	United States, 23456, Virginia Beach, VA	BANK OF AMERICA	American Sanctions 14	**30$**	+
AMEX	CREDIT		09/16	Yes	101	United States, 80123, Littleton, CO	BANK OF AMERICA	American Sanctions 14	**30$**	+
AMEX	CREDIT		03/17	Yes	101	United States, 60540, Naperville, IL	BANK OF AMERICA	American Sanctions 14	**30$**	+
AMEX	CREDIT		05/15	Yes	110	United States, 77081, Houston, TX	BANK OF AMERICA	American Sanctions 14	**30$**	+

Figure 2: Screenshot of stolen credit cards for sale.

Viewing the underground economy can give rough estimates of the security ecosystem. The attackers are running a business, which happens to be the failure of the customers of insurance. Mining the underground market for information about cyber-attacks (for example, the price of a ransomware exploit) is useful for insurance companies.

Uncertainty Abounds

Assessing the threat landscape is challenging because of the uncertainty that permeates nearly all of the aspects of cyber security. For example, the impact of reputation damage resulting from a hack is still hotly debated. Discussions at the workshop highlighted that some experts believe that reputation damage is potentially the most costly impact of a cyber incident. For example, the hack against Vodafone in Greece in 2004 significantly damaged the organization's ability to do business. However, academic research thus far has shown a weak correlation between a publically announced data breach and stock market returns. Typically, a statistically significant decrease is found for roughly two days after the breach announcement, but not longer, suggesting that investors do not consider data breaches to be a threat to future profits.[7] Another frequently cited example is the poor performance of Target after hackers stole over 40 Million credit cards, although the drop in profits has also been attributed to Target's failed expansion into Canada. Other experts present at the workshop indicated that reputation damage depends on the sector. For example, failure of critical infrastructure to deliver vital services, such as energy or health, would probably negatively impact profit in the long term. The impact of reputation is thus likely to be highly dependent on the specifics of an organization, the market share, and the ease of consumers to take their business elsewhere.

Another difficulty is the connection between historical trends and the future. Some research shows that the rate of *certain* attack vectors is relatively constant or slowly changing over

[7] See:
- Campbell, K., Gordon, L., Loeb, M., and Zhou, L. "The economic cost of publicly announced information security breaches: empirical evidence from the stock market," Journal of Computer Security (11:3) 2003, p 431.
- Cavusoglu, H., Mishra, B., and Raghunathan, S. "The effect of Internet security breach announcements on market value: Capital market reactions for breached firms and Internet security developers," International Journal of Electronic Commerce (9:1), Fall 2004, pp 69-104.
- Kannan, K., Rees, J., and Sridhar, S. "Market reactions to information security breach announcements: An empirical analysis," International Journal of Electronic Commerce (12:1), Fall 2007, pp 69-91.

time, while others can change rapidly[8]; denial of service attacks tripled in bandwidth in a short period of time in 2013.[9]

The threat of supply chain attacks is also largely unknown. The complexity of global supply chains makes it difficult to verify components and supply chain security is now a recurring concern for all businesses. Software is written in a globally dispersed way, making it very difficult to verify code and guarantee that a few malicious lines were not inserted.

The insider threat is another concern to many organizations and the true scope and magnitude of the problem is not well understood. Certain incidents create the perception that malicious insiders often occur (e.g. Snowden and Manning), but the large media attention may cause a reporting bias. However, it is clear that accidents (e.g. social security numbers are emailed to the wrong recipient) are much more common than malicious insider attacks.

Studying the Attacker

Cyber is unique from earthquake, weather, and many other insurable risks in the respect that cyber often involves an adaptive adversary. Therefore, relying on historical trends and blocking one attack vector is unlikely to be effective. However, insurance has experience dealing with other 'adaptive' adversaries, for example in offering insurance for typical crime (e.g. stolen artwork) or even in the case of insuring travelers against hostage taking and ransom. It is also important to note that some adversaries may adapt when a specific target is valued, but a significant proportion of cyber-attacks are against the weakest entities.[10] Therefore, for retailers with credit card information, hospitals with medical information, organizations with intellectual property that can be ransomed via crypto ransomware[11], or organizations with computing resources that can be used for a botnet, the attacker may not adaptively target their resources, since many victims exist. However, some organizations with specialized IP are more likely to face an adaptive adversary and will need to take specific measures to defend against their unique adversaries.

Studying the attacker is an important factor in being able to anticipate the threat environment. If statistics about the attackers’ behavior are available, they can be analyzed using a rich set of tools (similar to sports, where the statistical tendencies of each team are well-known). If statistics do not exist (i.e. the threat is new, or attacks are sparse), adversarial risk analysis can be used to study the attacker.

[8] See for example the two reports mentioned in note 2 and 3

[9] Specifically, the maximum rate of most DoS attacks was around a maximum of around 100 Gb per second, until 2013, when the maximum flowshot up to over 300 Gbps.

[10] Note that the bear anecdote is often used here, i.e. that you do not have to outrun the bear, you have to outrun the slower person next to you who the bear will eat. However, shouldn’t the group be defending against the bear collectively?

[11] Crypto ransomware involves malware that encrypts a user’s hard drive, and then demands a ransom be paid. If the ransom is not paid, the user’s files cannot be accessed because they are encrypted.

Adversarial Risk Analysis

Risk analysis has a long history of reducing the failure probability (or improving the security) of a wide range of technical and human systems. The basic framework involves modeling the system and its failures and then studying the effect of safeguards to minimize the probability of failure given limited resources. Risk analysis is rooted in systems analysis and probability theory and has been successfully applied to a number of areas, including offshore oil platforms, earthquakes, dams, medical devices, satellite architectures, spaceflight systems, and counter-terrorism.

Figure 3: Probabilistic risk analysis has addressed failure in a number of technical and social systems with applications in many fields, including oil platforms, earthquakes, dams, medical devices, satellites, and asteroids.

In cyber security, the presence of an adaptive adversary may require adversarial risk analysis. Here, game analysis is used to study what a defender knows about the adversary, including who they are, what they want, what they know, and what they have. The defender's decisions are dependent on the context, potential loss, the attack detection time, and what they are able and allowed to do.

Cyber security involves many of the same issues as counter-terrorism, including the involvement of a range of adversaries (script kiddies, criminals, and nation states) and how they operate. Key to the analysis are the details of the target, such as what they have, the vulnerabilities, and the countermeasures. One issue is the asymmetry of cyber-attacks, since an adversary can easily attack many victims simultaneously. Further, defense is often expensive while attacks are relatively cheap. Increasing the cost to the attacker is one strategy for reducing the rate of attacks, and might involve hacking back, stricter prosecution of criminals, or other deterrent mechanisms. But attribution and cross-border issues would complicate the task and more work is needed to determine which of these strategies are effective. Furthermore and for example, hacking back entails the known danger of impacting innocent individuals and its ethics and legality is politically loaded. Inflammatory rhetoric aside, hacking back does lead to important questions involving attribution confidence, cross-border legality, and many other issues, so other less dangerous strategies should be preferred, such as tricking the attackers into wasting their time and effort, or stealing useless goods. Stricter prosecution could be a more reasonable deterrent, but has its own challenges. Hacking laws have not caught up with technology in the sense that robbing a bank electronically may not carry the same sentence as robbing a bank physically.[12] Other thought leaders have called for a frame shift on IT responsibility, suggesting that leaving an unpatched server that criminals can use to attack others is akin to leaving loaded guns in your front yard, which society does not tolerate. Another issue is that many hackers are youth, and rehabilitation is likely a much better strategy than long jail terms. Finally, jurisdictional issues currently create massive barriers due to extradi-

[12] Note that while these two scenarios are different (physical theft typically involves a threat to cause physical harm with a weapon), sentences for computer crimes are generally short.

tion limitations. The lack of international guidelines and treaties on cyber crime allows criminals in certain countries to operate with impunity.

In practice, the implementation of an adversarial risk analysis involves identifying weaknesses, ranking threat scenarios, identifying and prioritizing countermeasures, and setting priorities in intelligence gathering and analysis. Defenders must look to the attackers to study their motivations, actions, and behavior. The relationship between attackers and defenders, like in terrorism risk, needs to be understood to better prevent and prepare for attacks. Useful analogies including a parasite-host relationship observed in biology could indicate new solutions to the cyber security problem. Co-evolution and predator prey-models all contain useful lessons in this aspect of the phenomenon.

2. Anticipating Threat Detection

From both a regulator and an industry perspective, early threat detection is needed to minimize damage and to prioritize security measures. To achieve this goal, organizations should work to generate incident data and to analyze incident data so that the nature of cyber-attacks can be better understood.

The availability of incident data is crucial. Security incident descriptions are often recorded in organizations for workflow or auditing purposes. These incident management systems are treasure troves of intelligence. Many organizations with mature cyber security programs are already collecting these data, and some governments require mandatory reporting of incidents in this style.[13] This has led to a slow accumulation of real world data on the types of attacks, along with their frequency and impact.

Data scarcity is a significant problem for many cyber researchers[14]. Insurance companies collect information on claims, but these data are seen as highly valued since they can give an insurer an edge in pricing policies. Therefore, data collection outside of insurance is important. Recognizing the barriers to data recording is important as well. User reporting of phishing emails is actually an important detection mechanism at most organizations, but the rate of user reporting is typically low.[15] One organization was able to trace this back to poor communication with its workforce; employees would report a suspicious email and never hear anything back, so reporting rates dropped. Communicating with employees rebuilt the collaborative relationship so that workers knew that forwarding suspicious emails was important.

[13] See US-CERT Federal Information System Report Guidelines. https://www.us-cert.gov/incident-notification-guidelines
[14] https://www.cambridgecybercrime.uk/ was formed to combat this problem
[15] For example, an employee should forward a phishing email to abuse@company.com

Cyber security incidents also need to be analyzed, so that the trends can be used to make better decisions. Hacking has changed over time. Email worms are becoming less common while targeted phishing attacks have risen over the past several years. Small businesses were historically safe from cyber-attacks, but criminal organizations began actively exploiting them around 2012 and attacks continue to rise.[16] However, evidence also shows that certain trends in cyber risk are remarkably stable. For example, the number of lost laptops appears to be constant over time. Similarly, academic research has shown that the rate and severity of publically disclosed data breaches involving personal information have not changed substantially over a 10-15 year period.[17] Overall, certain aspects of cyber risk are likely to be predictable, while others will change as attackers and defenders evolve. Insurers, organizations, and governments need to identify which is which.

Predictability

Cyber risk is inherently heavy-tailed and volatile, meaning that attacks occur irregularly and across a huge range of severities. There is no single measure of severity. As a result, resource planning becomes difficult because the resources needed to investigate cyber security incidents may change by huge amounts between months. Further, this creates issues because the lack of large cyber security incidents over long time periods creates the impression of effective security, although luck may be the cause.

Increasing Resilience

Many threats can be anticipated based on an organization’s characteristics. Once these threats are modeled, the security of an organization can be improved by implementing a variety of safeguards.[18] Broadly, security measures can be improved in three areas.

1) Defensive measures

Preventing an adversary from gaining access to a system limits the damage that an adversary can inflict. The strategies to improve defenses are well known and include designing more secure software or implementing an efficient patching program. Other strategies include whitelisting software to prevent unauthorized applications from communicating outside the network, access control and improved authentication to verify users, and compartmentalization to limit an attacker’s ability to move throughout a network.

2) Proactive measures

While the majority of security tends to be defensive, some proactive security measures can be implemented as well. For example, organizations might create false databases, or false network connections so that attackers spend time attacking systems that are not actually related to a business’s core network. The improved virtualization of machines has recently

[16] Brian Krebs. “Uptick in Cyber Attacks on Small Businesses“ http://krebsonsecurity.com/2012/08/uptick-in-cyber-attacks-on-small-businesses/

[17] See section 4 for citations.

[18] For an excellent resource see: Shostack, Adam. Threat modeling: Designing for security. John Wiley & Sons, 2014.

made this strategy feasible, so that an organization could run a virtual network on a laptop that distracts attackers from focusing on the organization's real network.

Another controversial strategy could involve preemptive action against cyber adversaries. However, this area is not well defined legally and includes a significant amount of risk that innocent agents will be impacted. The issue of attribution is also a limiting factor, since organizations would have to interpret an attack correctly. Many researchers scan the Internet on a regular basis, but these actions are sometimes interpreted as network reconnaissance. Another proactive approach would be to add an intelligence arm to security operations to learn about attackers' plans. Finally, having a credible deterrent in the form of threatened legal action could also reduce the frequency of certain types of attacks.

3) Retroactive security

Even once a security breach has occurred, organizations may have many options to deal retroactively with the fallout. An illustration of this involves recent discussions about preventing attacks not at the infiltration stage (when an adversary gains access to a network), but at the exfiltration stage (when an adversary attempts to take information from a network). Several security professionals have recently advocated that monitoring traffic leaving the network should be prioritized over monitoring incoming traffic. However, this strategy is highly dependent on the domain. The attackers who stole millions of credit cards from Target purposefully exfiltrated credit cards in batches that were transmitted during normal working hours so as not to trigger anomaly detection systems that might be triggered due to the large amounts of information leaving the network.

Even if information is successfully exfiltrated, organizations have many options for dealing with the response. For example, banks have assessed the cost of canceling credit cards after a data breach, and have found that certain cards (usually foreign-issued with no credit limit) are much more likely to be bought on the underground market than others. Certain banks have even chosen to simply buy back their customer's stolen information on the underground marketplace because the asking price was cheaper than the cost of reissuing cards. Giving security teams more creative license (within ethical and legal boundaries) may bring clever and cost-effective solutions.

3. Responding to Threats through Improved Security Measures

Organizations face many challenges in cyber systems, but academic research, industry initiatives, and government programs are actively developing new technical solutions that will improve the countermeasures and response to attacks. There are two types of gaps in cyber security, which can be addressed using short-term solutions or long-term planning.[19]

[19] Much of the following section was inspired by Clark, D., T. Berson, and H. S. Lin. "At the Nexus of Cybersecurity and Public Policy: Some Basic Concepts and Issues." (2014). This book was published by

Short-Term Initiatives

The first gap is the difference between the current cyber security posture of organizations and reality. Many best practices and technologies exist for cyber security that are not fully implemented or widely deployed and used. Closing this gap involves efforts to disseminate known best practices in cyber security and to use them both effectively and widely.

Long-Term Initiatives

The second gap is the difference between the strongest security posture that is possible today, with known practices and technologies, and the security posture needed to address the most advanced threat currently (and as they evolve). Closing this gap involves research efforts to develop new knowledge about cyber security.

3.1 Closing the Best Practice Gap

Many attackers currently have an advantage simply because organizations are not using best practices. Stronger efforts to disseminate information and expertise about best practices are needed so that existing methods and technology can be applied more efficiently and widely. Closing the best practice gap does not necessarily require new technical knowledge of cyber security, but rather the application of existing technical knowledge. Also, more research about the nontechnical nature of cyber security risk, including studying the culture, human factors, and management processes is needed to improve the understanding of how to better promote deployment and use of existing knowledge. This includes the economic or psychological factors affecting the implementation of known practices and techniques, enhanced educational efforts to promote security-responsible user behavior, incentives to build organizational cultures where higher degrees of security awareness are shared, and problem-specific research focused on developing solutions for urgent cyber security problems.

A solution to a cyber security problem can be implemented in industry provided it is effective, robust against a variety of attack types, inexpensive, easy to deploy, and easy to use. Furthermore, it must not significantly reduce other functionalities in the system into which it is integrated. Problem-specific research thus also includes developing better knowledge about how a technical solution can be deployed.

Implementing existing technology

Many organizations have still not yet implemented proven solutions to improve security. SANS publishes a list of 20 security controls that are a good start for many organizations.[20] Common safeguards include full disk encryption, network monitoring, strong authentication, compartmentalization, and the capability for logging and forensics.

the National Academies Press and is an excellent resource, and is also available for free download at National Academies Press website.

[20] See https://www.sans.org/critical-security-controls

Security Culture

Developing an organizational culture around security is critical. If employees view security as an impediment to work and do not recognize the importance and value of safeguards, shortcuts will be taken that introduce more vulnerabilities. For example, increasing the requirements for password complexity often leads to employees writing passwords on sticky notes. Organizations also need to build encourage accountability for cyber security (in a similar way that public safety campaigns focus on 'if you see something, then say something'). Employees that are empowered to report suspicious websites, phishing emails, or other indicators of compromise can enhance organizational security overall.

Metrics

Metrics and data are needed to monitor the progress of security. Currently, there are still rich debates about what security metrics are most useful to an organization, but the field is moving towards more repeatable, objective, and quantitative measures of security.

Situational Awareness

Organizations need good situational awareness not just of their threats, but also of their networks and assets. Many organizations do not have visibility into what devices are on their network at any given time, which significantly increases the difficulty of managing network security. Further, many organizations have not identified their core assets. For example, the databases, intellectual property, or computing resources that attackers are after may not be well documented. One organization failed to recognize that unsecured web servers were useful to adversaries as hosting sites for illicit activity, and found that an attacker had commandeered a neglected web server for over six months. Taking inventory of resources, mapping them to attacker's goals, and assessing the criticality of business operations can help ensure that attention is concentrated on monitoring the right assets.

Compartmentalization and Redundancy

Segmentation of a network and backups are an important tool for improving security. One specific area that is being addressed is the availability of data and services, which is a major concern in industry. For example, complete data centers might be replicated in different areas of a country to be redundant against natural disasters. However, while redundancy reduces the risk of unavailability, it may increase the risk overall by giving an attacker more opportunities to access the full system. Compartmentalization and building redundancies (independent backups) have traditionally been a strategy, but its effectiveness might be limited by the complexity of interdependent systems.[21] Depending on how it is used, the cloud can give either more or less exposure to cyber risk: security may be easier to implement in a cloud environment, but the standardization of security safeguards makes the system more vulnerable to dependent failures.

[21] For a discussion of one solution to the correlated failures, see Zhai, Ennan, et al. "Heading off correlated failures through independence-as-a-service." Proc. of OSDI. USENIX (2014).

3.2 Closing the Attacker-Defender Gap

We know that the ideal, strongest cyber security posture would still be inadequate against today's high-end threat, let alone tomorrow's. Even with the application of best practices, attackers currently still have an advantage. Closing this knowledge gap requires substantial research as well, which takes place in academia. New designs, technologies and approaches are needed. Fundamental applied and exploratory technical research in cyber security focuses on preparing future technologies as well as organizational and human arrangements to strengthen defenses against an evolving threat. Closing the attacker-defender gap calls for research that targets specific identifiable cyber security problems and also builds a base of technical expertise, to increase the ability to respond quickly in the future, when threats that are unknown today emerge.

Academia and industry partnerships have a long history of creating new technical solutions to security challenges. For example, privacy and security have long been thought of as a zero-sum game, but recent research has shown strong cryptographic tools that preserve both. Homomorphic encryption allows encrypted data to be mathematically manipulated without allowing the agent performing the calculations to learn anything about the underlying data.[22] Zero-knowledge proofs can be used to determine if a statement is true, without revealing any other information about the statement.[23] Verifiable computation and multiparty computation allow computation to take place without trusting that the agents doing the computing are honest. These new developments have far-reaching impacts and will fundamentally change certain aspects of cyber security.

The cyber security and privacy trade-off is being analyzed in other research as well. Schemes to cooperatively model the security of a network without revealing sensitive details are being developed.[24] Other work analyzes the best way to catch criminals without compromising other users' privacy, with an emphasis on mass surveillance conducted by law enforcement.[25]

Researchers are also continuing to push the boundaries of what is known to develop new tools and techniques. Some research revisits the fundamental architecture and protocols that support the Internet. For example, today's Internet is full of hijacking that redirects packets, leading to an inefficient and fragile system. Packets traveling between two cities less than 100 miles apart may actually traverse the ocean three times before arriving at their destination. Further, small errors can have huge consequences, as illustrated by a small mistake in routing tables in Malaysia that resulted in a large regional Internet outage. New solutions include initiatives like SCION (Scalability, Control, and Isolation on Next-Generation Networks), which

[22] Applications include users who want to learn if their DNA contains a certain gene, but do not want the organization performing the tests to learn their DNA sequence.
[23] For example, someone might want to verify that I have enough money in my bank account to make a purchase, but not reveal how much money I have in my account.
[24] See Zhai, Ennan, et al. "Heading off correlated failures through independence-as-a-service." Proc. of OSDI. USENIX (2014).
[25] Segal, Aaron, Bryan Ford, and Joan Feigenbaum. "Catching bandits and only bandits: Privacy-preserving intersection warrants for lawful surveillance." 4th USENIX Workshop on Free and Open Communications on the Internet (FOCI'14). 2014.

enable high-availability point-to-point Internet communication.[26] Many of these technologies emphasize availability, security, flexibility, transparency, scalability, cheapness, and ease of deployment.

New research methods are moving away from a weakest link security model, where an attacker simply has to compromise the weakest entry point, to a strongest link security model. These new methods leverage multiparty computation so that users do not have to rely on a fragile central authority that could be vulnerable to attack. For example, in some proposed schemes, critical signatures could be signed by a diverse, heterogeneous, and scalable system.[27]

3.3 New Areas for Research

More research is needed to close the gap between attackers and defenders. The list below details several recommendations that could guide future research directions, and is motivated by section 4.4 of the report "At the Nexus of Cybersecurity and Public Policy".[28]

Include uncertainty

Since the nature and severity of the future cyber security threat are largely unknown, research on new designs must be modest about what is known about future challenges and be able to adapt to change. Attacks will create new exploits, techniques, or fundamentally different attacks that current defenses cannot address. The legal landscape could radically change in the next ten years, and major shifts including the internet-of-things and cloud services will occur. Research portfolios should balance work on known threats, and new threats that will likely emerge in the future.

Programmatic continuity

Stakeholders should support a substantial effort in research areas with a payoff over a long time horizon. Long-term research can have intermediate milestones, to allow for midcourse corrections. It must engage both academic and industry actors, and it can involve early collaboration with technology-transition stakeholders, even in the fundamental research.

Broad research agenda.

Cyber security risks are anticipated to stay. Given that few specifics about future risks can be known with high confidence, it is not realistic to imagine that cyber security risks in the future can be prevented or substantially mitigated with one or a few "silver bullets". Cyber security research must be conducted across a broad front. There is an iterative feedback between attackers and defenders. Attackers respond and change their tactic in response to action

[26] Zhang, Xin, et al. "SCION: Scalability, control, and isolation on next-generation networks." Security and Privacy (SP), 2011 IEEE Symposium on. IEEE, 2011.

[27] Syta, E., Tamas, I., Visher, D., Wolinsky, D. I., & Ford, B. (2015). Decentralizing Authorities into Scalable Strongest-Link Cothorities. arXiv preprint arXiv:1503.08768.

[28] See Clark, D., T. Berson, and H. S. Lin. "At the Nexus of Cybersecurity and Public Policy: Some Basic Concepts and Issues." (2014).

from defenders. With a broad research agenda, chances are increased that the time needed to develop countermeasures against these new attacks when they appear will be shortened.

Multi-disciplinary research

Cyber security encompasses computer science, economics, sociology, law, risk, and many other domains. Research agendas should include multi-disciplinary work about organizational, sociological, economic, legal and psychological factors as well as technological ones.

Academic and Industry collaboration

Academia and industry need to share knowledge, prototypes, and data. While the two groups approach problems differently, the perspectives are complementary (immediate short-term goals in business vs. long-term view –beyond 2-5 years in academia). Business looks for solutions, while academics work on newly imagined problems. It is thus important that the right problem is picked so that the solution has impact and influence. Academics should maintain awareness of problems relevant to industry so that new PhD theses are relevant to the real world.

Critical questions

Many areas of cyber security still need to be researched. For example, organizations are currently assessing the risks and benefits of moving to cloud infrastructure. In this case, the details of how to classify data, train users on security concerns, select a trusted cloud provider, and how to integrate cloud capabilities with in-house infrastructure and processes are not well understood.

Detection and analytic capabilities also have important unanswered questions. Should all organizations staff a security operations center, or is outsourcing for large security incidents more effective? How useful are data analytics capabilities such as Splunk? Organizations know that incident detection and response is important, but strategies to reduce the response time are still not fully developed. Privileged accounts are another segment of challenges. Security professionals would prefer to eliminate root accounts, but current technologies do not allow this. The understanding of the risk of centralization for different resources and processes is critical, but still being developed.

The economics of security best practices and technology adoption are not well understood, along with many other important questions. However, progress is occurring by implementing known state-of-the-art solutions in organizations, and developing new solutions.

4. Dealing with the Remaining Risk

Many risk assessment frameworks exist, generally at a high level. However, the real value of risk assessments is in the implementation and the specific details of how a local operator uses the assessment for decision support. This requires a high degree of specialization and few examples of how to do this in practice are publically available. Without these tools, organizations are left making multimillion-dollar investment decisions at the CIO and CRO level without any support.

A comparison of how business risks are managed to how cyber risk is currently managed by most organizations illustrates the limitations with the current approach. Many types of business risk are typically modeled from the 'top-down'. There are clear assets that are at risk (a new product, or money that might be taken away due to regulatory fines). On the other hand, many organizations are unaware of the IT assets that are most valuable to their organization, and this type of risk modeling often occurs from the bottom-up. For example, many organizations are still uncertain how impactful a business interruption would be to their profits, and uncertain about the valuation of IP or the reputation damage that could result from a data breach. Many organizations do not even know if website attacks or stolen devices represent a greater risk to their organization.

The management of cyber risk requires solutions that involve addressing cyber risk and business strategy simultaneously. Stakeholders need to connect IT resources with the success of business development, and model the processes that could be impacted by cyber security incidents. Once this is done, the importance of security measures is understood much more easily, because stakeholders can immediately see how and why a security issue impacts business objectives.

Probabilistic Approaches

The objective of cyber risk management is to minimize the risk with constrained resources. For organizations to make actionable decisions, the identified risks need to be translated into business terms (aka dollar values).

Given the rebalancing between prevention & security versus risk management, new tools are needed to assess the risk of cyber security incidents. Ultimately, managers want to know how much they can lose and the probability of occurrence. Without this quantification, executives are often left with the impression that more money is being invested in cyber security, but the organization is not being made more secure. While significant uncertainty surrounds cyber risk (for example, reputation damage, the cost of IP theft, and the effectiveness of security controls), considerable progress has been made and organizations often have the data and capabilities to implement a quantitative risk assessment process.

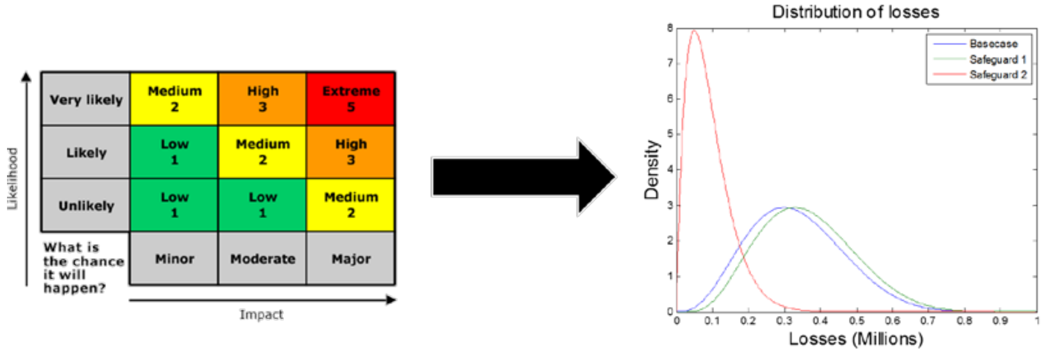

Very likely	Medium 2	High 3	Extreme 5
Likely	Low 1	Medium 2	High 3
Unlikely	Low 1	Low 1	Medium 2
What is the chance it will happen?	Minor	Moderate	Major

Figure 4: Risk management is moving away from qualitative risk matrices (heat maps) to probabilistic methods.

Quantifying cyber risk is well recognized as critically important. In boardrooms, organizations have quantified other risks in dollar terms for decades, although cyber risk is still most often evaluated on qualitative scales. New techniques in industry and academia, paired with new data about cyber security incidents are now resulting in quantitative risk management frameworks.

First, organizations often have useful data that have historically not been utilized to inform cyber risk assessments. Industry reports and academic literature have significantly contributed to basic facts regarding cyber security incidents as well. Currently, significant literature exists that suggests that the rate of cyber-attacks is more constant than previously thought, and the impact distribution is relatively stable as well.[29] Other literature has studied reputation damage by observing the change in stock prices after a data breach is disclosed and found that the effect of data breaches lasts only about two days (as discussed earlier). Individual organizations often have good visibility into cyber-attacks as well, and many organizations record data in the form of incident management systems. These databases can be leveraged to quantify cyber-attack trends over time, which inform probabilistic models of cyber security.

The incorporation of probabilistic methods is very important. Academic research has shown that cyber impacts are heavy-tailed, meaning that most cyber incidents are small, and that large incidents occur rarely but can be orders of magnitude more impactful than the average incident. In these types of systems, the expected value can be extremely misleading, since it combines frequent small losses with rare large losses.

A basic probabilistic framework involves modeling an organization, their assets, attackers, and vulnerabilities. Next, the rate and severity of cyber security incidents are assessed probabilistically. Here, organizations can use internal data, industry reports, or expert opinions as inputs.[30] Once the cyber security incidents are modeled, organizations need to determine the impact to the organization, taking care to model the monetary effects of different cyber events. Here again, probability theory is critical given the uncertainty involved in the outcomes. For example, some cyber breaches were widely reported (e.g. Target), while others are not widely known because of the public's interest relative to the news cycle (e.g. Michaels). Significant uncertainty is involved with the cost estimates, but organizations al-

[29] Several papers have found that data breaches are not rapidly increasing in frequency or severity. See:
- Maillart, T., and D. Sornette. "Heavy-tailed distribution of cyber-risks." The European Physical Journal B 75.3 (2010): 357-364.
- Edwards, Benjamin, Steven Hofmeyr, and Stephanie Forrest. "Hype and Heavy Tails: A Closer Look at Data Breaches. "Workshop on Economics of Information Security". (2015)
- Wheatley, Spencer, Thomas Maillart, and Didier Sornette. "The Extreme Risk of Personal Data Breaches & The Erosion of Privacy." arXiv preprint arXiv:1505.07684 (2015).

[30] Again, a major emphasis is that high quality statistics on the rate and severity of cyber security incidents are increasingly available, and suggest that cyber security incidents are much more patterned than previously believed. In fact, standard risk assessment methods are sufficient for broad classes of attacks, although certain scenarios (adaptive adversaries) may still require other methods.

most certainly have enough information to give rough estimates, which is sometimes all that is needed.

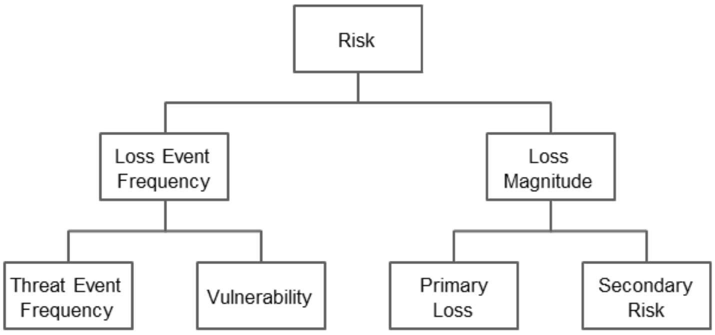

Figure 5: A method for assessing cyber risk, taken from the FAIR framework.

Once the cyber incidents and the costs are modeled, organizations can obtain risk curves to compare different types of attacks such as malicious emails, website attacks, or lost devices, and compare the value of different security safeguards. This is useful at two levels. First, a decision maker can quickly observe which class of attacks contributes the most risk. This prevents biases that might lead a decision maker to believe that certain attacks are more significant than they actually are. For example, many executives are concerned about malicious insiders, although evidence suggests that while mistakes may cause losses, employees that intend to cause harm are very rare. Website attacks are significantly less compelling than malicious insiders and get less attention, but are actually much riskier at many organizations. Second, a decision maker can assess the value of different security safeguards. For example, organizations might consider implementing a segmented network architecture to limit an adversary's ability for lateral movement and privilege escalation in a network. They can use a probabilistic model to observe how much risk is reduced, and decide if that investment is worth.

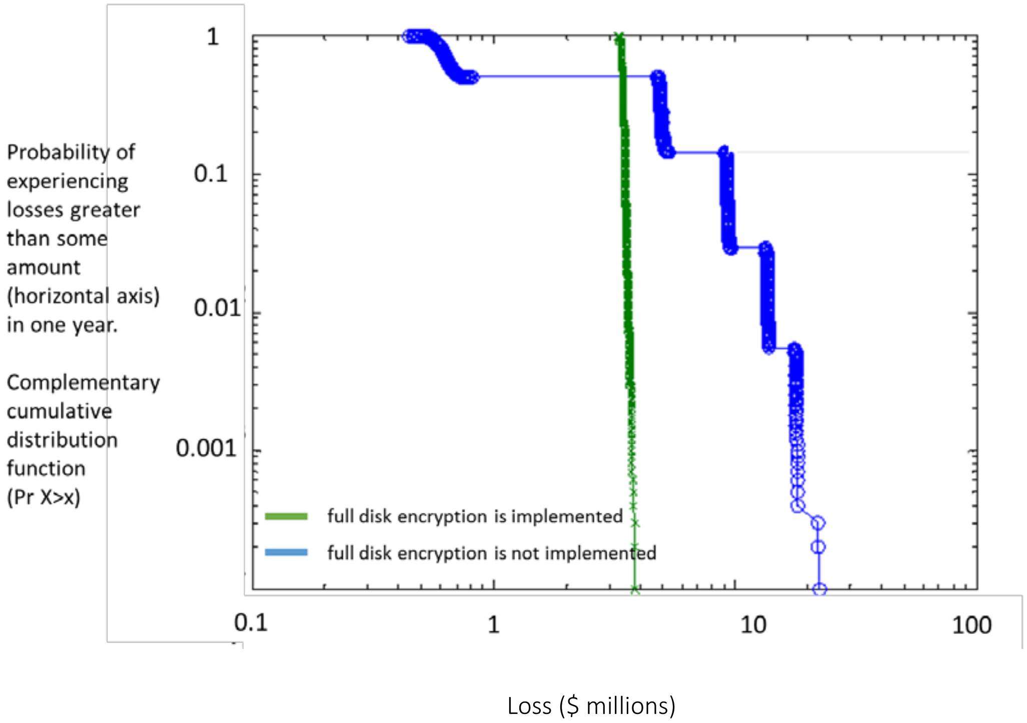

Figure 6: An example of a probabilistic risk analysis of yearly losses due to lost devices (laptops and phones) at an organization. The vertical axis shows the complementary cumulative distribution function, or the probability of experiencing losses greater than some amount (horizontal axis) in one year. The blue line shows the risk when full disk encryption is not implemented; a large proportion of the time losses are small, but large losses rarely occur if a laptop with lots of sensitive information is stolen. The green line shows that when full disk encryption is implemented, losses are consistently at $3-$4M, since the cost of the encryption software increases the loss, but large incidents are eliminated. Using this chart, a decision maker can choose if full disk encryption is a good investment. Taken from Kuypers and Pate-Cornell.

Organizations have already begun to recognize the benefits of probabilistic risk management and to apply these methods.[31] The application of risk management to cyber has involved developing a quantitative understanding of cyber risk, driving the risk management process with analytics, using optimization tools paired with a risk management mentality, and improving the simple communication of risk.

In the end, a risk assessment process should be highly tailored to the organization, and the risk assessment should not be confused with the decision. Each decision maker's risk attitude may result in different management decisions about the cyber risk. The quantification process improves transparency and communication by rigorously defining the attacks and their business impacts. All of this drives a culture and communication change.

IT specialists do not think like business people and vice versa. Dialogue between these two groups is essential, to understand the trade-offs between cost and risk, benefits now vs. risk later. The explicit risk attitude of an organization can be modeled with a probabilistic framework, and its strategic business decisions can be assessed. For example, certain companies might opt to release a customer application that will experience fraud, but the losses due to fraud might be outweighed by benefits of growing a user base.

5. Insurance

Cyber insurance has the potential to improve overall cyber resilience by offering support to businesses. In summary, cyber insurance serves to transfer part of the residual risk, after a thorough risk management process has been performed by business, to drive best standards through risk-based pricing for policies, and to act as a consolidated entity that can aid in information sharing. However, many challenges currently exist. Underwriters may not have enough technical data or expertise to assess cyber risk, organizations may not understand the coverage in a policy (for example, whether other non-cyber policies may offer overlapping coverage, or if important gaps exist), and all stakeholders seem unsatisfied with policy pricing (businesses think it is too high, while insurers think it is too low). Further, coverage is always capped, meaning that businesses might be limited in the amount of risk they can transfer.

[31] For example, Risk Lens (www.risklens.com) offers probabilistic risk assessments based on the FAIR framework.

Cyber Insurance

Cyber insurance is an emerging market and sufficient data on claims do not yet exist. This is in contrast to car insurance for example, where indicators of risk are well understood, and a young man will be expected to pay more for car insurance than a middle-aged woman. In cyber insurance, the risk indicators are still in development. The insurance applications that organizations fill out are not consistent across different insurance companies and brokers, a sign that it is unclear what are the most accurate "predictors" of cyber risk among, for example, the size of the organization, the sector in which it operate, or its security budget. Insurers are slowly gathering more data to improve risk measures, and they recognize the value of claims data as a competitive advantage, hence they are unlikely to make it publically available at this time.

Cyber risk insurance can be roughly divided into three categories, for which the aggregation of the risk needs to be analyzed:

1. Business interruption (Denial of Service, large scale internet outage, deletion of large databases, etc.)
2. Breach of privacy information or financial data (credit cards or medical records stolen)
3. Physical damage to infrastructure (cyber-attacks leading to damage of cooling system at a manufacturing plant)

Many factors contribute to the risk involved within each of the categories above. For example, the line of business, the industry, the region impacted and the type of coverage all influence the risk to the insurer and reinsurer. To better model the impact of possible catastrophic events, a scenario approach is useful. Several scenarios could lead to widespread business interruption. For example, a general malware attack could impact a large segment of industry, an attack on a major cloud provider could affect many customers, or a long scale outage of the Internet could make many services unavailable. Further, the propagation or downstream effects of a cyber-attack are important to analyze, especially given the question of the role of geographic distribution of the covered insureds.

Reinsurance and Aggregation Risk

Reinsurers are primarily concerned with the aggregation of cyber risk, which can be caused by an event impacting:

- either a large number of organizations
- or a large number of different insurance covers for the same insured.

In the first case, the effects of a cyber incident may not be severe at the individual level to result in a high aggregation risk. If an incident is widespread, the damage may still accumulate to a level where reinsurance is impacted. For example, a coordinated cyber-attack against many companies in one sector, or a wide-scale vulnerability that impacts a large proportion of devices could result in a huge number of impacted customers.

In the second case, the effects of a cyber incident could trigger reinsurance because they can be far reaching by disrupting ordinary business operations and are not limited to pure cyber impacts. For example, a dedicated attack to a large financial institution could create an accumulation of data privacy breaches and business interruption, possibly coupled with a D&O claim. Similarly, a cyber-attack that interferes with the power grid would cause widespread disruption because of the power failure, not just the cyber-attack itself.

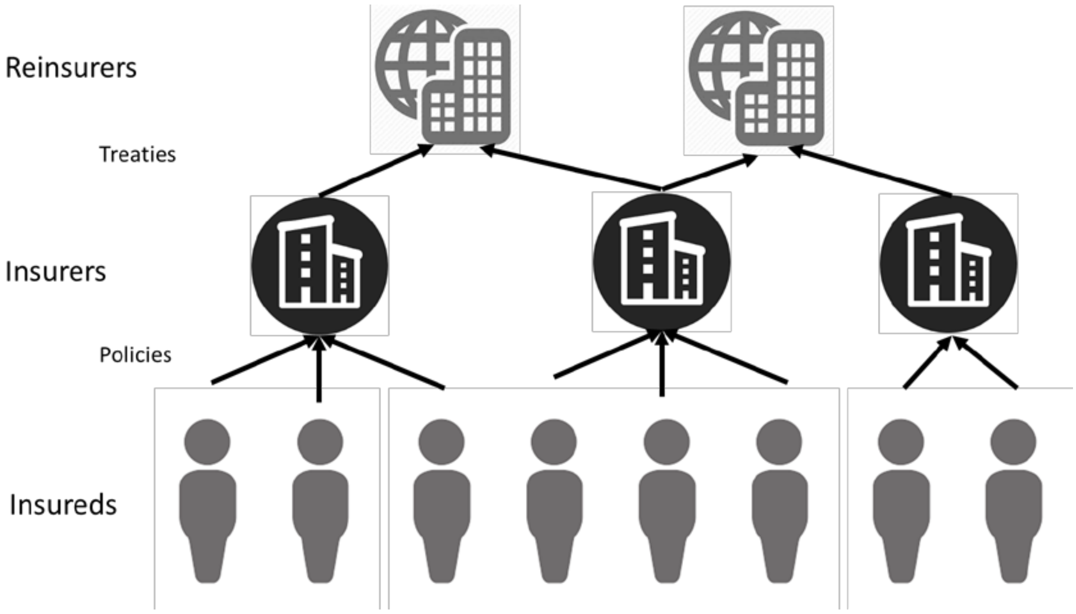

Figure 7: Simplified structure of insurance markets. Companies obtain insurance policies with an insurer. The insurers then write treaties with a reinsurer to protect against events that involve a large number of insured organizations, which could lead to insolvency.

The general mechanism for reinsurance is that multiple policies from an insurer get bundled into a treaty, which multiple reinsurers may cover. Critical pieces of information need to be known in order to quantify the accumulation risk to reinsurers. For example, reinsurers would like to know the type and characteristics of the software that an industry segment uses. Potential dependencies to the same originating cause include operating systems, firewalls, cloud providers, database products, and many others. If an industry is too homogeneous, then a single failure can result in widespread damage. Reinsurers have a unique view of this problem, and may be able to encourage industry diversity. In the retail sector for example, it is better to rely on several different payment card processors instead of only one.

Driving cyber security improvements

Beyond providing loss coverage, the insurance sector sees its role in driving improvements in risk management and cyber security. Once the key features of cyber security are determined, these can be used to set policy prices and to drive changes in organizations. An analogy with the field of fire insurance reminds us that insurance companies have incentivized organizations to install smoke alarms and sprinkler systems as effective devices to reduce fire risk, thus obtaining lower rates or higher covers, and avoiding the need for government regulation. In a similar way, flight recorders ("black boxes") have helped to understand the cause of airplane accidents, thus reducing the number and severity of accidents. These examples suggest that insurance could drive improvements in cyber systems as well, once it has identified which are the key security measures and through the incentive of adequate pricing.

6. Governing Cyber Security Risk – Important Trade-offs

Cyber security encompasses many stakeholders across business units, disciplines, industries, and national boundaries. Finding the correct governance level involves balancing many trade-offs.

The importance of cyber security has slowly gained traction over the past five years to a point where a cyber security briefing is now part of quarterly reports in many boardrooms. However, since cyber risk is often assessed in qualitative terms (low, medium, high) instead of dollar terms, cyber risk is often not compared to other business risks. This difficulty creates limitations for oversight and governance.

Governance of cyber security is itself an evolving domain. Executives are rarely fired if a hurricane causes significant losses, but they can still be held responsible for massive cyber-attacks, as detailed by the Target hack and an attack against Wyndham hotels. However, the issue of culpability is important; while executives have no control over an attacker, the board is responsible for overseeing IT security. For example, after the Sony hack and the Target hack, information leaked that both organizations had a culture that persistently ignored security concerns. Sony had actually been hacked several times before, but security was not prioritized in the organization. Therefore, while some organizations are simply unlucky, others are more vulnerable if cyber security is not taken seriously.

Information Sharing and Data Collection

One important step that could be taken to improve cyber security risk assessments is the collection and distribution of anonymous, high-value data about security incidents. Germany has approached this problem with an IT security act aimed at aggregating cyber security incidents, with an emphasis on critical infrastructure. The goal of BSI is to act as a central data repository that can both collect data and distribute threat intelligence to other operators to improve cyber defenses.

BSI’s lead in creating a threat intelligence sharing center is key to the debate about who should lead data-sharing initiatives. Experts note that there are already large informal networks of information sharing, including in the financial services, universities, and technology companies. The government’s role in supporting or controlling these groups is uncertain, and different stakeholders have different views on the appropriate level of government involvement. The government may have access to additional data sources (classified intelligence) that can be very useful to organizations, but there is also a distrust of government involvement, especially after recent developments detailing government surveillance.

In debating the role of government and data sharing initiatives, the intended user is very important to consider. For example, it is uncertain how much aid governments can provide to sufficiently cyber-mature organizations. Instead, many information sharing initiatives are aimed at less cyber-mature organizations, who could benefit from government guidance in securing their systems.

Getting organizations to share incident data is difficult by itself. In many industries, small groups have informal information sharing networks, most notably in the financial services sector. Here, government regulation for incident reporting would likely meet pushback. It is still unclear if public data sharing initiatives will be effective[32], or if governments will need to fill a gap by supporting certain sectors or types of organizations (for example, small businesses that may not have the resources or knowledge to invest in good cyber security programs).

Regulation

Governments are the only entities that have the authority to quickly change the cyber landscape through regulation. However, governments typically move slowly, and the cyber field is evolving very rapidly. Care must be taken to enable cyber security, and not slow it down.

Risk frameworks and compliance can lead to organizations checking boxes instead of investing in cyber risk management programs. Compliance is difficult as well because of the heterogeneity of threats and victims, meaning a solution for one organization will not always work for the next. For example, Netflix might appropriately prioritize availability of its services, while Visa may focus on fraud detection. Regulation can lead to misallocation of resources. Also, compliance with standards and norms, especially if made compulsory, can increase homogeneity and make penetration easier, thus increasing susceptibility.

Governance in Organizations

Executives are becoming more aware of the cyber threat. Security is a challenging domain because if it is done well, it is invisible. Executives expect that hackers should not be able to get in, and the Chief Information Security Officer often only gets attention when something goes wrong. Organizations should seek to educate executives on the legal, technological, and policy implications of cyber security.

Summary of trade-offs relevant to future policy research			
Trade-off	Comment 1	Comment 2	Recommendation
Prevention/ containment versus recovery and resilience	Historically, prevention has been emphasized	Attackers can always get in, so defenders need detection and recovery capabilities	Cyber risk management should include both prevention and detection/ recovery mechanisms
Usability versus security	Historically, usability has driven technology adoption, at the expense of security	Over-emphasizing security can slow down business operations and adversely impact profits	Usability and security trade-offs should carefully be considered.
Compulsory versus voluntary reporting	More data would significantly help the community, but privacy concerns may be present	Organizations may not want to put effort into reporting unless it is compulsory	Certain systems should require reporting (critical infrastructure), others could be optional

[32] See VERIS, an open source incident reporting tool at www.veriscommunity.nets

Trade-off	Comment 1	Comment 2	Recommendation
Security versus privacy	Historically, enhanced security means compromising privacy (e.g. observing network traffic)	Privacy and security may not always be at odds, and new technologies balance these preferences	All stakeholders should be engaged to debate the security vs. privacy debate, with special attention paid to new solutions that have both
Government-led versus industry-led	Governments have more intelligence capabilities and compulsory powers, but may not be trusted	Many larger industries have information sharing that is more efficient than government, but smaller businesses may not have enough resources	A mixed approach is needed; industry initiatives can be implemented and government initiatives can support small businesses
Short-term vs. long-term	Investment in security measures (e.g. authentication, encryption, compartmentalization) typically concerns the short-term	Long-term security investment measures (e.g. systems redesign) are needed but not well known and can be expensive	Developing approaches that reconcile short- and long-term needs

Organizations should work on balancing different aspects of the cyber security problem. The table above lists several key trade-offs and a small sampling of comments, along with a recommendation for policy makers and academics to encourage more research. Cyber issues are complicated and involve a wide range of domains, technological issues, legal issues, and risks. Careful analysis of these trade-offs can lead to better decisions dealing with a diverse threat landscape.

Conclusion: Key Takeaways

Cyber security is a rapidly evolving field that impacts governments, organizations, and individuals. More work is needed in this area, but collaboration between government, academia, and industry (notably insurers) stands to deliver large benefits to the cyber security community.

Uncertainty permeates cyber systems, and scientific evidence is hard to obtain. Many hacks that attract high media attention are rare, and their impact may be exaggerated (hacks are often described with inflammatory rhetoric with warnings like 'malware is poised to cause catastrophic economic damage'). In some cases, this may be for good reasons such as drawing public attention to severe cyber risks to encourage collaboration to address it.

The reality of cyber risk is hard to assess, both over- and under- estimation co-exist. More collaborative research is needed to determine basic realities about cyber-attacks so that organizations can adequately address the threat. So far, cyber has not been an existential risk to large organizations, except in rare cases where businesses sustain reputation damage. For example, RSA, Target, Home Depot, Anthem, Sony, and Experian are all still successful businesses, despite experiencing severe hacking incidents. Small businesses, on the other hand, are much more vulnerable. IP theft and hacks that involve transferring large cash reserves outside the country have been especially dangerous to small business owners.
However, if attacks in the future target intellectual property, industrial controls systems, personal data, or some other asset and the provision of critical services such as energy, water, health or finance, losses may be both severe and lasting. This justifies devoting efforts and investments to improving both security measures and risk-based approaches.

Collecting data and sharing information about breaches and incidents is a critical step towards progress. However, how it can be done is subject to some controversy. For example: is it more efficient if it is under compulsory or voluntary schemes? Information sharing is already taking place in informal industry groups, but mostly involves large organizations (Fortune 1,000 organizations). The groups are typically invitation only and are likely more effective than government initiatives, although those also play an important role, in calling for transparency and acting as a neutral source of information. In addition, governments have a large role to play in helping small and medium size businesses, which may not have the resources to invest in an advanced cyber security program.

Quantitative risk assessments are still rare in industry, at least according to publicly available information. The vast majority of organizations use qualitative approaches, and huge potential exists to adopt powerful new probabilistic approaches that are being developed by industry and academia. These approaches emphasize risk management over a binary security mentality and support decision makers in their task to make well-informed trade-offs involving different business priorities. Probabilistic approaches also allow organizations to rigorously determine their level of risk acceptance and risk transfer, through insurance products.

Insurance companies can improve cyber security risk management by collecting data and driving incentives for effective security controls through policy pricing. Insurers and academics have some overlapping incentives, and data sharing should occur whenever possible.

Research into new security controls should broadly cover both short-term and long-term initiatives that either disperse best practices or work on novel technologies to shrink the attacker-defender gap. Research should be interdisciplinary and adaptive to address new issues that are not yet evident, given the rapidly evolving nature of cyber security.

In summary, **organizations need to quantify cyber risk to enable better decisions about security investments and business management**. The risk management process begins with knowing your assets, the threat landscape, your exposure, and your vulnerabilities. Improvements in organizational culture can also encourage the recognition that cyber risk management is integral to a successful business.

Using the current best practices and on-going research initiatives, organizations can address the security gap by (a) implementing the best controls, (b) systematically using the most advanced security tools and implementing new systems designs, which exist today and are being developed by academia and security firms, and give defenders an edge in cyber security, and (c) implementing risk-based approaches to minimize cyber risk impact on business.

Overall, tremendous potential exists. While cyber security presents many challenges, many smart, well-resourced, and motivated stakeholders are working quickly to defend organizations from cyber-attackers. By collaborating and continuing to advance the field of cyber security, organizations will continue to enjoy the benefits of technology while operating in a safer and more secure environment.

International Risk Governance Council

The International Risk Governance Council (IRGC) is an independent foundation based in Switzerland whose purpose is to help improve the understanding and governance of systemic risks that have impacts on human health and safety, on the environment, on the economy and on society at large.

International Risk Governance Council
c/o École Polytechnique Fédérale de Lausanne (EPFL)
CM 1 517
Case Postale 99
CH-1015 Lausanne
Switzerland

Tel +41 21 693 82 90
info@irgc.org

www.irgc.org