Testimony and Statement for the Record

Bruce Schneier
Founder and Chief Technical Officer
Counterpane Internet Security, Inc.

Hearing on
"Overview of the Cyber Problem—A Nation Dependent and Dealing with Risk"

Before the
Subcommittee on Cybersecurity, Science, and Research and Development
Committee on Homeland Security
United States House of Representatives

June 25, 2003
2318 Rayburn House Office Building

Mr. Chairman, members of the Committee, thank you for the opportunity to testify today regarding cybersecurity, particularly in its relation to homeland defense and our nation's critical infrastructure. My name is Bruce Schneier, and I have worked in the field of computer security for my entire career. I am the author of seven books on the topic, including the best-selling *Secrets and Lies: Digital Security in a Networked World* [1]. My newest book is entitled *Beyond Fear: Thinking Sensibly About Security in an Uncertain World* [2], and will be published in September. In 1999, I founded Counterpane Internet Security, Inc., where I hold the position of Chief Technical Officer. Counterpane Internet Security provides real-time security monitoring for hundreds of organizations, including several offices of the federal government.

Cyber Risks

When I began my long career in computer security, it was a marginal discipline. The only interest was from the military and a few scattered privacy advocates. The Internet has changed all that. The promise of the Internet is to be a mirror of society. Everything we do in the real world—all of our social and business interactions and transactions—we want to do on the Internet: conduct private conversations, keep personal papers, sign letters and contracts, speak anonymously, rely on the integrity of information, gamble, vote, publish authenticated documents. All of these things require security. Computer security is a fundamental enabling technology of the Internet; it's what transforms the Internet from an academic curiosity into a serious business tool. The limits of security are the limits of the Internet. And no business or person is without these security needs.

The risks are real. Everyone talks about the direct risks: theft of trade secrets, customer information, money. People also talk about the productivity losses due to computer security problems. What's the loss to a company if its e-mail goes down for two days? Or if ten people have to scramble to clean up after a particularly nasty intrusion? I've seen figures in the billions quoted for total losses from Internet epidemics like Nimda and the SQL Slammer; most of that is due to these productivity losses.

More important are the indirect risks: loss of customers, damage to brand, loss of goodwill. When a successful attack against a corporation is made public, the victim may experience a drop in stock price. When CD Universe suffered a large (and public) theft of credit card numbers in early 2000, it cost them dearly in their war for market share against Amazon.com and CDNow. In the aftermath of public corporate attacks, companies often spent more money and effort containing the public relations problem than fixing the security problem. Financial institutions regularly keep successful attacks secret, so as not to worry their customer base.

And more indirect risks are coming as a result of litigation. European countries have strict privacy laws; companies can be held liable if they do not take steps to protect the privacy of their customers. The U.S. has similar laws in particular industries—banking and healthcare—and there are bills in Congress to protect privacy more generally. We have not yet seen shareholder lawsuits against companies that failed to adequately secure their networks and suffered the consequences, but they're coming. Can company officers be held personally liable if they fail to provide for network security? The courts will be deciding this question in the next few years.

This hearing was convened to address another type of risk: the risks of our nation's critical infrastructure that is largely in the hands of private companies. One of the great challenges of cybersecurity is the interdependencies between individual networks. The security decisions one company makes about their own network can have far-reaching effects across many networks, and this leads us to different sorts of risks. I call these ancillary risks because their effects are ancillary to the particular network in question. Ancillary risks abound in cyberspace. For example, home computer users are at risk of attack and of having their machines taken over by others, but an ancillary risk is created when their attacked and taken-over computers can be used for further attacks against other networks. Vulnerabilities in software create a risk for the corporation marketing that software, but they also creates an ancillary risk for those who use that software in their networks.

The cybersecurity risks to our nation are largely ancillary; because our critical infrastructure is largely in the hands of private companies, there are risks to our nation that go beyond what those private companies are worried about. The telephone network has value to the telephone companies because that's how they get revenue, and those companies will secure their networks to that value. But the network has value to the country as a nationwide communications structure in addition to that, and there are ancillary risks as a result of that. Companies put themselves at risk when they purchase and use insecure software, but they also cause ancillary risks to everyone else on the Internet because that software is on a common network. These ancillary risks turn out to be critical to the current insecurities of cyberspace, and addressing them will give us the only real way to improve the situation.

As risky as the Internet is, companies have no choice but to be there. The lures of new markets, new customers, new revenue sources, and new business models are just so great that companies have flocked to the Internet regardless of the risks. There is no alternative. Governments feel the same sorts of pressures: better ways of interacting with citizens, more efficient ways of disseminating information, greater involvement of citizens in government. The Internet is here to stay, and we're going to be using it for more and more things regardless of the risks. This, more than anything else, is why computer security is so important.

Quantifying the Risks

Quantifying the risks is difficult, because we simply don't have the data. Most of what we know is anecdotal, and what statistics we have are difficult to generalize. In summary, cyberattacks are very common on the Internet. Corporations are broken into regularly, usually by hackers who have no motivation other than simple bragging rights. There is considerable petty vandalism on the Internet, and sometimes that vandalism becomes large-scale and system-wide. Crime is rising on the Internet, both individual fraud and corporate crime. We know all this is happening, because all surveys, corporate studies, and anecdotal evidence agree. We just don't know exact numbers.

For the past eight years, the Computer Security Institute has conducted an annual computer crime survey of U.S. corporations, government agencies, and other organizations [3]. The details are a bit numbing, but the general trends are that most networks are repeatedly and successfully attacked in a variety of ways, the monetary losses are considerable, and there's not much that technology can do to prevent it. In particular, the 2003 survey found the following:

- 56% of respondents reported "unauthorized use of computer systems" in the last year. 29% said that they had no such unauthorized uses, and 15% said that they didn't know. The number of incidents was all over the map, and the number of insider versus outsider incidents was roughly equal. 78% of respondents reported their Internet connection as a frequent point of attack (this has been steadily rising over the six years), 18% reported remote dial-in as a frequent point of attack (this has been declining), and 30% reported internal systems as a frequent point of attack (also declining).

- The types of attack range from telecommunications fraud to laptop theft to sabotage. 36% experienced a system penetration, 42% a denial-of-service attack. 21% reported theft of proprietary information, and 15% financial fraud. 21% reported sabotage. 25% had their Web sites hacked (another 22% didn't know), and 23% had their Web sites hacked ten or more times (36% of the Web site hacks resulted in vandalism, 35% in denial of service, and 6% included theft of transaction information).

- One interesting thing highlighted by this survey is that all of these attacks occurred despite the widespread deployment of security technologies: 98% have firewalls, 73% an intrusion detection system, 92% access control of some sort, 49% digital IDs. It seems that these much-touted security products provide only partial security against attackers.

Unfortunately, the CSI data is based on voluntary responses to surveys. The data only includes information about attacks that the companies knew about, and only those attacks that they are willing to admit to in a

survey. Undoubtedly, the real numbers of attacks are much higher. And the people who complete the CSI survey are those experienced in security; companies who are much less security savvy are not included in this survey. These companies undoubtedly experience even more successful attacks and even higher losses.

The Honeynet Project is another source of data. This is an academic research project that measures actual computer attacks on the Internet. According to their most recent statistics [4], published in 2001, a random computer on the Internet is scanned dozens of times a day. The average life expectancy of a default installation of a Linux Red Hat 6.2 server—that is, the time before someone successfully hacks it—is less than 72 hours. A common home user setup, with Windows 98 and file sharing enabled, was successfully hacked five times in four days. Systems are subjected to hostile vulnerability scans dozens of times a day. And the fastest time for a server being hacked: 15 minutes after plugging it into the network. This data correlates with my own anecdotal experience of putting computers on an unsecured home broadband network.

At Counterpane Internet Security, we keep our own statistics. In 2002, we monitored several hundred computer networks in over thirty countries. We processed 160 billion network events, in which we uncovered 105 million security alerts. Further processing yielded 237,000 "tickets" which were investigated by our trained security analysts, resulting in 19,000 customer contacts from immediate security incidents. Assuming our data is representative, a typical company in the United States experiences 800 critical network security events—events requiring immediate attention—each year. At Counterpane we're smart and experienced enough to ensure that none of those events results in financial losses for the companies we protect, but most companies do not have such vigilant cyber guards.

Cybersecurity Trends

Several cybersecurity trends are worth highlighting. First, over the past few decades attacks on individual computers, early networks, and then the Internet have continually gotten more severe. Attack tools have gotten more potent, more damaging, more effective. Attacks that were once slow to implement are now automated. Attacks that used to be defeatable by a single mechanism are now adaptive. Viruses, worms, and Trojans are more elaborate and intelligent; malicious programs that years ago took weeks to spread across cyberspace, and last year took hours, today spread in minutes.

Second, over that same time period, the expertise required to launch those attacks has gone down. Many attack tools are easy to use. They have point-and-click interfaces. They are automated. They don't require any expertise to operate. "Root kits" are both easier to use and more effective.

These two trends combine to exacerbate another trend: the rise of crime in cyberspace. The vast majority of cyberspace attacks are nothing more than petty vandalism: the Internet equivalent of spray painting. The attackers aren't after anything except a cheap thrill and bragging rights. Sometimes they're bored teenagers. Sometimes they're smart kids with no other outlet. But we're starting to see significant increases in real crime on the Internet. Criminals, who often don't have the computer expertise to break into networks, can employ these easy-to-use tools to commit crimes. Credit card thefts and other forms of fraud are on the rise. Identity theft is on the rise. Extortion is on the rise. At Counterpane, often the hardest job we have is detecting these criminal attacks among the hundreds of petty vandalism attacks. I expect this trend to continue as more criminals discover the value of committing their frauds in cyberspace.

On the defensive side of things, cyberspace is becoming less secure even as security technologies improve. There are many reasons for this seemingly paradoxical phenomenon, but they can all be traced back to the problem of complexity. As I have said elsewhere [5], complexity is the worst enemy of security. The reasons are complex and can get very technical, but I can give you a flavor of the rationale: Complex systems have more lines of code and therefore more security bugs. Complex systems have more interactions and therefore more potential for insecurities. Complex systems are harder to test and therefore are more likely to have untested portions. Complex systems are harder to design securely, implement securely, configure securely, and use securely. Complex systems are harder for users to understand.

Everything about complexity leads towards lower security. As our computers and networks become more complex, they inherently become less secure.

Another trend is the ineffectiveness of security products. This is not due to failures in technology, but more to failures of configuration and use. As amazing as it seems, the vast majority of security products are simply not implemented in ways that are effective. The blame could be laid on the products themselves, which are too hard to use. The blame could be laid on the system administrators, who often install security products without thinking too much about them. But the real blame is in the culture: security simply isn't a priority in most organizations. Security is routinely ignored, bypassed, or paid lip service to. Products are purchased because an organization wants to pass an audit or avoid litigation, but much less attention is paid to how they are used. It's as if a homeowner bought an expensive door lock and installed it in a way that didn't provide any security.

Along similar lines, the quality of software security is abysmal. Products routinely ship with hundreds or thousands of security vulnerabilities. Again, there are technical reasons for this. As a science, computer security is still in its infancy. We don't know, for example, how to write secure software. We have some tricks, and we know how to avoid some obvious problems, but we have no scientific theory of security. It's still a black art and, although we're learning all the time, we have a long way to go. But again, the real reason is that security isn't a priority for software vendors. It's far better for a company if they ship an insecure product a year earlier than a more secure product a year later.

The result of these trends is that security technologies are improving slowly, not nearly fast enough to keep up with the new insecurities brought about by the increasing complexity of systems. Every year brings more new attacks, faster-spreading worms, and more damaging malicious code. Software products—operating systems as well as applications software—continue to have more and more vulnerabilities. As long as the trends of increasing complexity and security's low priority continue, cyberspace will continue to become less secure.

Complexity is something we can't change. The only thing we can change is to make security a higher priority.

Cyberterrorism or "Digital Pearl Harbor"

There is one often-discussed trend that I do not see: the rise of cyberterrorism [6]. An essay I wrote on this issue is included as Attachment #1. I believe that fears about cyberterrorism, or the likelihood of a "Digital Pearl Harbor," are largely the result of companies and organizations wanting to stoke the fears of people and of the news media looking for sensationalist stories. Real terrorism—attacking the physical world via the Internet—is much harder than most people think, and the effects of cyber attacks are far less terrorizing than might seem at first. Cyberterrorism is simply not a problem that we have to worry about.

This does not mean that large-scale cyberspace threats are not a problem. A single vulnerability in a widely used software product can affect millions, and an attack that exploits that vulnerability can do millions of dollars of damage overnight. Attacks against popular Internet services, or critical information services that use the Internet to move data around, can affect millions.

While people overplay the risks of cyberterrorism, they underplay the risks of cyber-crime. Today credit card numbers are no longer being stolen one at a time out of purses and wallets; they're being stolen by the millions out of databases. Internet fraud is big business, and it's getting bigger.

And someday, cyberterrorism will become a real threat. Technology, especially technology related to cyberspace, is fast-moving and its effects are far-reaching. Just as some unknown attacker used the physical mail system to spread the anthrax virus, it is certainly possible that, someday, a terrorist may figure out how to kill large numbers of people via the Internet. But that day is not coming soon, and even then the same terrorist would probably have a much easier time killing the same number of people in a physical attack.

The Resilience of the Internet

Despite all of these risks, the Internet is reasonably safe from a catastrophic collapse. As insecure as each individual component or network that makes up the Internet is, as a whole it is surprisingly resilient. Often I have joked that the Internet "just barely works," that it is constantly being revised and upgraded, and that it's a minor miracle that it functions at all.

The Internet has seen examples of what many people have in mind when they think about large-scale attacks or terrorism, only they've been the result of accidents rather than maliciousness. Telephone switching stations shut down as the result of a software bug, leaving millions without telephone service. Communications satellites temporarily malfunctioned, disabling a nationwide pager network. On 9/11, the World Trade Center fell on much of lower Manhattan's communications network. What we've learned from these episodes is that the effects are not devastating and they're only temporary; communications can be quickly restored, and people adapt until they are restored.

Additionally, random events are still much more damaging than malicious actions. In the closest example of a cyberterrorist attack we've experienced, Vitek Boden hacked into a computer network and released a million liters of pollution into an Australian estuary. His damage was cleaned up in a week. A couple of months later, a bird landed on a transformer in the Ohio River valley, causing it to blow up; this set off a chain reaction that released about ten times as much sewage into the river. The cleanup was much more expensive and took significantly longer. Even today, random birds can do significantly more damage than the concerted effort of someone intent on damage.

Security and Risk Management

Companies manage risks. They manage all sorts of risks; cyber risks are just one more. And there are many different ways to manage risks. A company might choose to mitigate the risk with technology or with procedures. A company might choose to insure itself against the risk, or to accept the risk itself. The methods a company chooses in a particular situation depend on the details of that situation. And failures happen regularly; many companies manage their risks improperly, pay for their mistakes, and then soldier on. Companies, too, are remarkably resilient.

To take a concrete example, consider a physical store and the risk of shoplifting. Most grocery stores accept the risk as a cost of doing business. Clothing stores might put tags on their garments and sensors at the doorways; they mitigate the risk with technology. A jewelry store might mitigate the risk through procedures: all merchandise stays locked up, customers are not allowed to handle anything unattended, etc. And that same jewelry store will carry theft insurance, another risk management tool.

An appreciation of risk management is fundamental to understanding how businesses approach computer security. Ask any network administrator what he needs cybersecurity for, and he can describe the threats: Web site defacements, corruption and loss of data due to network penetrations, denial-of-service attacks, viruses, and Trojans. The list of threats seems endless, and they're all real. Ask senior management about cybersecurity, and you'll get a very different answer. He'll talk about return on investment. He'll talk about risks. And while the cyber threats are great, the risks are much less so. What businesses need is adequate security at a reasonable cost.

Given the current state of affairs, businesses probably spend about the right amount on security. The threats are real and the attacks are frequent, but most of the time they're minor annoyances. Serious attacks are rare. Internet epidemics are rare. And on the other side of the coin, computer security products are often far less effective than advertised. Technology changes quickly, and it's hard to mitigate risks in such a rapidly changing environment. It is often more cost effective to weather the ill effects of bad security than to spend significant money trying to improve the level of security.

Externalities and Our Critical Infrastructure

If companies are so good at risk management, why not just let them manage their own risks? Companies can decide whether or not to have a guard in their corporate offices, install an alarm system in their warehouses, or buy kidnapping insurance for their key executives. Shouldn't we simply let companies make their own security decisions based on their own security risks? If they don't care whether they buy and use insecure software, if they don't bother installing security products correctly, if they don't implement good cybersecurity policies, why is that anyone else's problem? If they decide that it's cheaper to weather all the Internet attacks than it is to improve their own security, isn't it their own business?

The flaw in that argument is the reason this hearing was convened: the ancillary threats facing our nation's critical infrastructure. The risks to that infrastructure are greater than the sum of the risks to the individual companies. We need to protect ourselves against attack from an enemy military. We need to protect ourselves against a future where cyberterrorists may target our electronic infrastructure. We need to protect the underlying economic confidence in the Internet as a mechanism for commerce. We need to protect the Internet above the risks to individual pieces of it. Companies are good at risk management, but they're only going to consider their own risks; the ancillary risks to our critical infrastructure will not be taken into account.

One easy example is credit card numbers. Company databases are regularly broken into and credit card numbers are stolen, sometimes hundreds of thousands at a time. Companies work to secure those databases, but not very hard, because most of the risk isn't shouldered by those companies. When an individual finds that his credit card number has been stolen and used fraudulently or, even worse, that his entire identity has been stolen and used fraudulently, cleaning up the mess can take considerable time and money. The company secures the database based on its own internal risk; it does not secure the database based on the aggregate risk of all the individuals whose information it stores.

Software security is another example. Software vendors do some security testing on their products, but it's minimal because most of the risk isn't their problem. When a vulnerability is discovered in a software product, the vendor fixes the problem and issues a patch. This costs some money, and there's some bad publicity. The real risk is shouldered by the companies and individuals who purchased and used the product, and that risk doesn't affect the vendor nearly as much. When the SQL Slammer worm spread across the Internet in January 2003, worldwide losses were calculated in the tens of billions of dollars. But the losses to Microsoft, whose software contained the vulnerability that the Slammer used in the first place, were much, much less. Because most of the risks to Microsoft are ancillary, security isn't nearly as high a priority for them as it should be.

This brings us to the fundamental problem of cybersecurity: It needs to be improved, but those who can improve it—the companies that build computer hardware and write computer software, and the people and companies that own and administer the small networks that make up the Internet—are not motivated to do so.

More specifically: Our computers and networks are insecure, and there every reason to believe that they will become less secure in the future. The threats and risks are significant, and there is every reason to believe that they will become more significant in the future. But at the same time, because much of the risks are ancillary, software and hardware manufacturers don't spend a lot of money improving the security of their products and private network owners don't spend a lot of money buying and installing security products on their networks.

In economics, an externality is an effect of a decision that is not part of the decision process. Most pollution, for example, is an externality. A factory makes an economic decision about the amount of pollution it dumps into a river based on its own economic motivations; the health of the people living downstream is an externality. A welfare mother makes a decision whether to marry someone or live with him without marriage partly based on the economics of the welfare system; the societal degradation of the institution of marriage is an externality. Ancillary cyberrisks are an example of an externality.

There are several ways to deal with externalities. They can be regulated through a legal system: Laws and regulations which prohibit certain actions and mandate others are a way to manage externalities. They can

be internalized through taxation or liabilities, both of which provide economic incentives to take externalities into account. Sometimes societal norms modify externalities. And so on. The particular mechanism chosen will depend on politics, but the overall goal is to bring the various externalities into the decision process.

I believe that externalities are the fundamental problem of cybersecurity. The security of a particular piece of the Internet may be good enough for the organization controlling that piece, but the external effects of that "good enough" security may not be good enough for the nation as a whole. Our nation's critical infrastructure is becoming more and more dependent on a secure and functioning Internet, but there's no one organization in charge of keeping the Internet secure and functioning. Our software has very poor security, and there is no real incentive to make it better. We are increasingly vulnerable to attacks that affect everyone a little bit, but that no one has enough incentive to fix.

Recommendations

This fundamental problem of cybersecurity is much more an economic one than a technical one. Our nation's computer infrastructure could be much more secure if the business incentives were there to make it so—if the externalities were internalized, so to speak. Asking companies to improve their own security won't work. (We've tried this repeatedly; it's doomed to failure.) Trying to build a separate government network won't work. (The whole point of cyberspace is that it is one large interconnected network.) Hoping technology will improve won't work. (It doesn't matter how good the technology is if people don't want to use it.)

The basic capitalist and democratic business process is capable of improving cybersecurity, but only if the proper incentives are in place. My general recommendation is that you pass laws and implement regulations designed to deal with the externalities in cybersecurity decisions so that organizations are motivated to provide a higher level of security—one that is commensurate with the threat against our nation's critical infrastructure—and then step back and let the mechanisms of commercial innovation work to solve the problems and improve security. Specifically:

1. Stop trying to find consensus. Over the years, we have seen several government cyberspace security plans and strategies come out of the White House, the most recent one this year [7]. These documents all suffer from an inability to risk offending any industry. In the most recent strategy, for example, preliminary drafts included strong words about wireless insecurity that were removed at the request of the wireless industry, which didn't want to look bad for not doing anything about it. A recommendation that ISPs provide personal firewalls to all of their users was likewise removed, because the large ISPs didn't want to look bad for not already providing such a security feature. Unlike many other governmental processes, security is harmed by consensus. Cybersecurity requires hard choices. These choices will necessarily come at the expense of some industries and some special interests. As long as the government is unwilling to move counter to the interests of some of its corporate constituents, huge insecurities will remain.

2. Expose computer hardware, software, and networks to liabilities. I have written extensively about the effect of liabilities on the computer industry [8]; one of my essays is included as Attachment #2. The major reason companies don't worry about the externalities of their security decisions—the effects of their insecure products and networks on others—is that there is no real liability for their actions. Liability will immediately change the cost/benefit equation for companies, because they will have to bear financial responsibility for ancillary risks borne by others as a result of their actions. With liabilities firmly in place, the best interests of software vendors, and the best interests of their shareholders, will be served by them spending the time and money necessary to make their products secure before release. The best interests of corporations, and the best interests of their shareholders, will be served by them spending the time and money necessary to secure their own networks. The insurance industry will step in and force companies to improve their own security if they want liability coverage at a reasonable price. Liability is a common capitalistic mechanism to deal with externalities, and it will do more to secure our nation's critical infrastructure than any other action.

3. Secure your own networks. Fund programs to secure government networks, both internal networks and publicly accessible networks. Only buy secure hardware and software products. Before worrying about the security of everyone else, get your own house in order. This does not mean that it's necessary to redo what is already being done in industry. The government is a consumer of computer products, like any large corporation. The government does not need to develop its own security products; everyone's security is better served if the government buys commercial products. The government does not need to create its own organization to identify and analyze cyber threats; it is better off using the same commercial organizations that corporations use. The threats against government are the same as the threats against everyone else, and the solutions are the same. The U.S. government, specifically the Department of Homeland Security, should use and improve the resources that are available to everyone, since everyone needs those same resources.

4. Use your buying power to drive an increase in security. U.S. government procurement can be a potent tool to drive research and development. If you demand more secure products, companies will deliver. Standardize on a few good security products, and continually force them to improve. There's a "rising tide" effect that will happen; once companies deliver products to the increasingly demanding specifications of the government, the same products will be made available to private organizations as well. The U.S. government is an enormous consumer of computer hardware, software, systems, and services. And because you're using the same commercial products that everyone else uses, those products will improve to the benefit of everyone. The money you spend on your own security will benefit everyone's security.

5. Invest in security research; invest in security education. As the market starts demanding real security, companies will need to figure out how to supply it. Research and education are critical to improving the security of computers and networks. Here again, use your financial muscle to improve security for everyone. Research and education in this important field need to be increased. The benefits will be beyond anything we can imagine today.

6. Rationally prosecute cybercriminals. In our society, we rarely solve security problems by technical means alone. We don't wear body armor or live in fortresses. Instead, we rely on the legal system to rationally prosecute criminals and act as a deterrent to future crimes. We need to beef up law enforcement to deal with real computer crimes. This does not mean charging sixteen-year-old kids as adults for what are basically 21st century pranks; this means going after those who commit real crimes on the Internet.

Conclusion

None of this is easy. Every computer company you bring into this room will tell you that liabilities will be bad for their industry. Of course they're going to tell you that; it's in their best interests not to be responsible for their own actions. The Department of Homeland Security will tell you that they need money for this and that massive government security program. Of course they're going to tell you that; it's in their best interests to get as large a budget as they can. The FBI is going to tell you that extreme penalties are necessary for the current crop of teenage cyberterrorists; they're trying to make the problem seem more dire than it really is to improve their own image. If you're going to help improve the security of our nation, you're going to have to look past everyone's individual self-interests toward the best interests of everyone.

Our nation's cybersecurity risks are greater than those of any individual corporation or government organization, and the only way to manage those risks is to address them directly. I strongly recommend that you put the interests of our nation's cybersecurity above the interests of individual corporations or government organizations. The externalities of rational corporate cybersecurity decisions are hurting us all. It's the job of government to look at the big picture and the needs of society as a whole, and then to properly motivate individuals to satisfy those needs.

Thank you for the opportunity to appear before your committee today. I would be pleased to answer any questions.

References

[1] Bruce Schneier, *Secrets and Lies: Digital Security in a Networked World,* John Wiley & Sons, 2000.

[2] Bruce Schneier, *Beyond Fear: Thinking Sensibly About Security in an Uncertain World,* Copernicus Books, 2003.

[3] Computer Security Institute, "2003 CSI/FBI Computer Crime and Security Survey," 2003. http://www.gocsi.com/press/20030528.html

[4] Honeynet Project, :Know Your Enemy: Statistics," 22 July, 2001. http://www.honeynet.org/papers/stats/

[5] Bruce Schneier, "Software Complexity and Security," *Crypto-Gram,* March 15, 2000. http://www.counterpane.com./crypto-gram-0003.html

[6] Bruce Schneier, "The Risks of Cyberterrorism," *Crypto-Gram,* June 15, 2003. http://www.counterpane.com./crypto-gram-0306.html

[7] White House, *National Strategy to Secure Cyberspace,* Feb 2003. http://www.whitehouse.gov/pcipb/cyberspace_strategy.pdf

[8] Bruce Schneier, "Liability and Security," *Crypto-Gram,* April 15, 2002. http://www.counterpane.com./crypto-gram-0204.html

ATTACHMENT #1

The Risks of Cyberterrorism
Bruce Schneier

Reprinted from: *Crypto-Gram,* June 15, 2003.
http://www.counterpane.com./crypto-gram-0306.html

The threat of cyberterrorism is causing much alarm these days. We have been told to expect attacks since 9/11; that cyberterrorists would try to cripple our power system, disable air traffic control and emergency services, open dams, or disrupt banking and communications. But so far, nothing's happened. Even during the war in Iraq, which was supposed to increase the risk dramatically, nothing happened. The impending cyberwar was a big dud. Don't congratulate our vigilant security, though; the alarm was caused by a misunderstanding of both the attackers and the attacks.

These attacks are very difficult to execute. The software systems controlling our nation's infrastructure are filled with vulnerabilities, but they're generally not the kinds of vulnerabilities that cause catastrophic disruptions. The systems are designed to limit the damage that occurs from errors and accidents. They have manual overrides. These systems have been proven to work; they've experienced disruptions caused by accident and natural disaster. We've been through blackouts, telephone switch failures, and disruptions of air traffic control computers. In 1999, a software bug knocked out a nationwide paging system for a day. The results might be annoying, and engineers might spend days or weeks scrambling, but the effect on the general population has been minimal.

The worry is that a terrorist would cause a problem more serious than a natural disaster, but this kind of thing is surprisingly hard to do. Worms and viruses have caused all sorts of network disruptions, but it happened by accident. In January 2003, the SQL Slammer worm disrupted 13,000 ATMs on the Bank of America's network. But before it happened, you couldn't have found a security expert who understood that those systems were dependent on that vulnerability. We simply don't understand the interactions well enough to predict which kinds of attacks could cause catastrophic results, and terrorist organizations don't have that sort of knowledge either—even if they tried to hire experts.

The closest example we have of this kind of thing comes from Australia in 2000. Vitek Boden broke into the computer network of a sewage treatment plant along Australia's Sunshine Coast. Over the course of two months, he leaked hundreds of thousands of gallons of putrid sludge into nearby rivers and parks. Among the results were black creek water, dead marine life, and a stench so unbearable that residents complained. This is the only known case of someone hacking a digital control system with the intent of causing environmental harm.

Despite our predilection for calling anything "terrorism," these attacks are not. We know what terrorism is. It's someone blowing himself up in a crowded restaurant, or flying an airplane into a skyscraper. It's not infecting computers with viruses, forcing air traffic controllers to route planes manually, or shutting down a pager network for a day. That causes annoyance and irritation, not terror.

This is a difficult message for some, because these days anyone who causes widespread damage is being given the label "terrorist." But imagine for a minute the leadership of al Qaeda sitting in a cave somewhere, plotting the next move in their jihad against the United States. One of the leaders jumps up and exclaims: "I have an idea! We'll disable their e-mail...." Conventional terrorism—driving a truckful of explosives into a nuclear power plant, for example—is still easier and much more effective.

There are lots of hackers in the world—kids, mostly—who like to play at politics and dress their own antics in the trappings of terrorism. They hack computers belonging to some other country (generally not government computers) and display a political message. We've often seen this kind of thing when two

countries squabble: China vs. Taiwan, India vs. Pakistan, England vs. Ireland, U.S. vs. China (during the 2001 crisis over the U.S. spy plane that crashed in Chinese territory), the U.S. and Israel vs. various Arab countries. It's the equivalent of soccer hooligans taking out national frustrations on another country's fans at a game. It's base and despicable, and it causes real damage, but it's cyberhooliganism, not cyberterrorism.

There are several organizations that track attacks over the Internet. Over the last six months, less than 1% of all attacks originated from countries on the U.S. government's Cyber Terrorist Watch List, while 35% originated from inside the United States. Computer security is still important. People overplay the risks of cyberterrorism, but they underplay the risks of cybercrime. Fraud and espionage are serious problems. Luckily, the same countermeasures aimed at cyberterrorists will also prevent hackers and criminals. If organizations secure their computer networks for the wrong reasons, it will still be the right thing to do.

ATTACHMENT #2

Liability and Security
Bruce Schneier

Reprinted from: *Crypto-Gram,* April 15, 2002.
http://www.counterpane.com./crypto-gram-0204.html

Today, computer security is at a crossroads. It's failing, regularly, and with increasingly serious results. I believe it will improve eventually. In the near term, the consequences of insecurity will get worse before they get better. And when they get better, the improvement will be slow and will be met with considerable resistance. The engine of this improvement will be liability—holding software manufacturers accountable for the security and, more generally, the quality of their products—and the timetable for improvement depends wholly on how quickly security liability permeates cyberspace.

Network security is not a problem that technology can solve. Security has a technological component, but businesses approach security as they do any other business risk: in terms of risk management. Organizations optimize their activities to minimize their cost * risk product, and understanding those motivations is key to understanding computer security today.

For example, most organizations don't spend a lot of money on network security. Why? Because the costs are significant: time, expense, reduced functionality, frustrated end users. On the other hand, the costs of ignoring security and getting hacked are small: the possibility of bad press and angry customers, maybe some network downtime, none of which is permanent. And there's some regulatory pressure, from audits or lawsuits, that add additional costs. The result: a smart organization does what everyone else does, and no more.

The same economic reasoning explains why software vendors don't spend a lot of effort securing their products. The costs of adding good security are significant—large expenses, reduced functionality, delayed product releases, annoyed users—while the costs of ignoring security are minor: occasional bad press, and maybe some users switching to competitors' products. Any smart software vendor will talk big about security, but do as little as possible.

Think about why firewalls succeeded in the marketplace. It's not because they're effective; most firewalls are installed so poorly as not to be effective, and there are many more effective security products that have never seen widespread deployment. Firewalls are ubiquitous because auditors started demanding firewalls. This changed the cost equation for businesses. The cost of adding a firewall was expense and user annoyance, but the cost of not having a firewall was failing an audit. And even worse, a company without a firewall could be accused of not following industry best practices in a lawsuit. The result: everyone has a firewall, whether it does any good or not.

Network security is a business problem, and the only way to fix it is to concentrate on the business motivations. We need to change the costs; security needs to affect an organization's bottom line in an obvious way. In order to improve computer security, the CEO must care. In order for the CEO to care, it must affect the stock price and the shareholders.

I have a three-step program towards improving computer and network security. None of the steps have anything to do with the technology; they all have to do with businesses, economics, and people.

Step one: enforce liabilities. This is essential. Today there are no real consequences for having bad security, or having low-quality software of any kind. In fact, the marketplace rewards low quality. More precisely, it rewards early releases at the expense of almost all quality. If we expect CEOs to spend significant resources on security—especially the security of their customers—they must be liable for

mishandling their customers' data. If we expect software vendors to reduce features, lengthen development cycles, and invest in secure software development processes, they must be liable for security vulnerabilities in their products.

Legislatures could impose liability on the computer industry, by forcing software manufacturers to live with the same product liability laws that affect other industries. If software manufacturers produced a defective product, they would be liable for damages. Even without this, courts could start imposing liability-like penalties on software manufacturers and users. This is starting to happen. A U.S. judge forced the Department of Interior to take its network offline, because it couldn't guarantee the safety of American Indian data it was entrusted with. Several cases have resulted in penalties against companies who used customer data in violation of their privacy promises, or who collected that data using misrepresentation or fraud. And judges have issued restraining orders against companies with insecure networks that are used as conduits for attacks against others.

However it happens, liability changes everything. Currently, there is no reason for a software company not to offer more features, more complexity. Liability forces software companies to think twice before changing something. Liability forces companies to protect the data they're entrusted with.

Step two: allow parties to transfer liabilities. This will happen automatically, because this is what insurance companies do. The insurance industry turns variable-cost risks into fixed expenses. They're going to move into cyber-insurance in a big way. And when they do, they're going to drive the computer security industry...just like they drive the security industry in the brick-and-mortar world.

A company doesn't buy security for its warehouse—strong locks, window bars, or an alarm system—because it makes it feel safe. It buys that security because its insurance rates go down. The same thing will hold true for computer security. Once enough policies are being written, insurance companies will start charging different premiums for different levels of security. Even without legislated liability, the CEO will start noticing how his insurance rates change. And once the CEO starts buying security products based on his insurance premiums, the insurance industry will wield enormous power in the marketplace. They will determine which security products are ubiquitous, and which are ignored. And since the insurance companies pay for the actual liability, they have a great incentive to be rational about risk analysis and the effectiveness of security products.

And software companies will take notice, and will increase security in order to make the insurance for their products affordable.

Step three: provide mechanisms to reduce risk. This will happen automatically, and be entirely market driven, because it's what the insurance industry wants. Moreover, they want it done in standard models that they can build policies around. They're going to look to security processes: processes of secure software development before systems are released, and processes of protection, detection, and response for corporate networks and systems. And more and more, they're going to look towards outsourced services.

The insurance industry prefers security outsourcing, because they can write policies around those services. It's much easier to design insurance around a standard set of security services delivered by an outside vendor than it is to customize a policy for each individual network.

Actually, this isn't a three-step program. It's a one-step program with two inevitable consequences. Enforce liability, and everything else will flow from it. It has to.

Much of Internet security is a common: an area used by a community as a whole. Like all commons, keeping it working benefits everyone, but any individual can benefit from exploiting it. (Think of the criminal justice system in the real world.) In our society we protect our commons—our environment, healthy working conditions, safe food and drug practices, lawful streets, sound accounting practices—by legislating those goods and by making companies liable for taking undue advantage of those commons. This kind of thinking is what gives us bridges that don't collapse, clean air and water, and sanitary

restaurants. We don't live in a "buyer beware" society; we hold companies liable for taking advantage of buyers.

There's no reason to treat software any differently from other products. Today Firestone can produce a tire with a single systemic flaw and they're liable, but Microsoft can produce an operating system with multiple systemic flaws discovered per week and not be liable. This makes no sense, and it's the primary reason security is so bad today.