4 One View of Protecting the National Information Infrastructure

Eugene H. Spafford

As we think about the threats to information technology, we may wonder what the real threat is to the public. After all, disruption of eBay, Amazon, Google or online chat groups does not seem like much of a menace. It is actually the amplification of other threats that is disturbing. Imagine if a few months from now we suddenly have an outbreak of a new variety of hemorrhagic fever. Cases begin to appear in emergency rooms at a dozen hospitals in a dozen cities around the country. Within two or three hours after the story hits the national news wires everyone becomes aware of the outbreak and we begin to marshal our resources. But the Internet goes down, along with 50 percent of the national phone network, and neither comes back up. Most people who boot up their computers get a message that says "Death to the great Satan," and then their disk is erased.

How many would be able to respond effectively without phones or network resources? What kind of panic would ensue? While thinking about this question, consider that our government has decided that cost is more important than quality. They use a monoculture computing system that has a compromised (and some would say,

---

Eugene H. Spafford is professor of computer science, professor of philosophy, and director of the Center for Education and Research in Information Assurance and Security at Purdue University. *This chapter is based on remarks delivered at the 27th Annual AAAS Colloquium on Science and Technology Policy held April 11-12, 2002, in Washington, DC.*

minimal) immune system. It is being used for weapons guidance, national defense, government, and communications. Most people use the same system on their personal and business computers. Currently we are seeing new computer viruses and worms, targeted at that platform, reported approximately once every 75-90 minutes, on average. Extrapolating from the last 12 years of data, we may be seeing a new virus appear for this platform once every 30 minutes by sometime in 2003.

This system is the one on which we base our defenses, our economy, and much of our scientific enterprise. It is increasingly under attack from malicious software, in addition to a continuing litany of crashes, bugs and associated patches. And I have not even touched on the problem of hackers and automated attack tools, which may actually pose a larger threat than viruses as time goes by. The infrastrucure is built on shaky ground.

One recent study conducted in cooperation with the Federal Bureau of Investigation revealed that companies lose, on average, over $1 million each per year from computer misuses and computer crime. Worldwide, as much as $1 *trillion* may be lost in downtime and damages each year. Not only is poor security costing us real money, it is also harming our national competitiveness.

Looking at the current state of the practice in research makes the picture even more dismal. Nationally, we are producing, on average, about seven new Ph.D.s a year with in-depth knowledge of information security obtained at the leading educational institutions in this field. Based on recent trends, we can expect that two to four of them, on average, will return to their home countries each year, and perhaps three will take positions in academia.

Nationwide, we have no more than 100 (and perhaps as few as 60) faculty who have real training and expertise in this area and who are teaching in higher education. If a particular group of about 10 of them decide to retire or leave, we might lose two-thirds of our research centers in this area. Currently, the Federal government is investing no money in these centers to keep them running, so a lack of other resources might well have the same end result.

As best as I can tell, the total amount of money available this most recent fiscal year for *basic* research in information security was about $2 million (through the National Science Foundation); a great deal of

money is being spent on acquisition and development of technology for security, but that is money spent on extensions of known methods rather than basic research.

That is, in brief, the current state of information security in the United States.

Background

A little background may help us understand where some of the problems are and potentially where some of the solutions lie. Part of our current situation has to do with the speed at which information technology has progressed.

Forty years ago, we saw the creation of the first academic program in computer science as a discipline. This happened at Purdue University, with a few others created shortly thereafter. Thirty years ago, there were no large-scale networks. Everything was mainframes. Twenty years ago, the ARPANET had 231 nodes on the network (which was considered huge); today, although we have no way of getting an exact count, it is estimated there are as many as 300 million hosts connected to the Internet. Ten years ago, the World Wide Web protocol was developed. Commercial use of this network started only about seven and a half years ago. The World Wide Web now supports tens of billions of dollars worth of electronic transactions each year. The rapidity of the development of all this technology has certainly played a role here, as has some of the changing nature of the goals of the systems.

As we add more systems and as the growth continues, we also see a number of interesting emergent effects. This is not surprising given the increase in the complexity and size of the network. The amount of traffic that we see on the backbones of the networks has been doubling approximately every 90 to 120 days. That is an incredible increase in the amount of traffic. And the number of people estimated to be online has been doubling about every eight to 10 months for the last decade. Try to imagine any other kind of environment where you can continue that population growth reliably. It has strained our ability to cope with what is going on.

The Problems We Face

As I mentioned above, we have a significant shortage of professionals in computing science in general. At the undergraduate level we are producing only a fraction of the needed personnel, with some estimates indicating several available positions for each new graduate. Trend figures from the Computing Research Association's Taulbee Survey indicate that we have had a slight decline in the number of graduate degrees awarded in this field. At the Ph.D. level, nationally we are only producing enough Ph.D.s to stay even with the current number of computer science faculty in major universities—and there are certainly not enough graduates to satisfy the need to grow departments. In addition, currently, 57 percent of the graduate students in computer science in this country are not U.S. nationals. Only about 12 percent of our students are female. Other underrepresented minorities are an even smaller percentage, unfortunately.

Information security is an even smaller component of this population. Why? In large part, because there is huge market demand. It draws away people well before they go on to advanced study or teaching. It is a very small community to begin with and the market demand shrinks it even more. The demand is so great that many companies are hiring former criminals and confessed vandals as security experts. This is incredibly poor business sense (would you want a "reformed" arsonist to install your fire alarms?), but poor technical sense as well: most systems are so fragile that kids can become experienced system hackers without any real technical depth.

This incredible demand not only pulls people out of academia before they complete their training in information assurance, but it also has led to companies hiring individuals without sufficient training in basic computing. Many of the major software firms have been so desperate to get anyone who knew how to write a program that they have hired people who have only a high school diploma. In fact, they have hired some people before they finished high school whose entire education in software engineering may have been picked up in an introductory book. The code they write, their engineering and their designs form the foundation for our national (and global) information infrastructure. And there is a woefully small cadre of information assurance specialists to help shore it up. It is a grim picture.

High market demand is only one of the issues we must deal with. The speed of the market is another. In producing software, time to market is a critical business decision and simply one factor in the speed of the market. If you take six extra months to design and test your software carefully, you may be preceded to market by another firm in the same area and you would lose your opportunity to dominate. Getting there first is now more important than getting there correctly.

One result is that companies no longer do so much as beta test their software. Instead, they market it with disclaimers, anticipating that they can patch it in the next release. This is complicated by the fact that the Internet is a marvelous distribution channel, which can be used to disseminate software with no shipping cost and very little advertising. You simply put something up on a Web site; then people find and download it. That is also a very convenient mechanism for patch distribution—one can ship the next emergency patch over the Internet with almost no additional cost. It also means that because patches and configuration options are so available, there are no standard configurations to test to anymore.

There is no standard version of software anymore, partly because of the speed of the market. Demand trumps issues of quality and safety. We have all kinds of eager adopters who are enamored of the technology and who want the latest, greatest, and newest features. They are willing to download marginal software, install it, and run it in a risky manner simply to have it.

Technology trumps management in this regard. Trying to set policy to prevent people from downloading early-release versions and forcing them to run a standard configuration runs the risk of rebellion. Preventing employees from browsing Web sites during office hours is enough to cause them to leave companies and go to work elsewhere. If the employees are computer savvy, they can easily find new positions, so management is reluctant to enforce basic controls. This means we have real problems in security management.

One approach is to impose technology that works as a network nanny, as it were, to guard what people are browsing at work, or to set up firewalls to prevent users from downloading risky software. The problem with this technology is that you are trying to prevent technology-literate individuals from doing something that they want

to do. They can—and will—find ways around it.

Another problem involves the vendors. They are annoyed that management is trying to keep out their advertisements. As a result, they create protocols and methods to circumvent security.

Another issue is that of liability. Vendors produce goods they know are bad. Not only is there no feedback, there is no liability. The vendors are spending a considerable amount of money trying to get legislation passed to shield them from lawsuits. The Uniform Computer Information Transactions Act, a modification of the Uniform Commercial Code, is being lobbied very heavily by a number of manufacturers for passage at the state level. Virginia and Maryland have already passed it. This law allows the vendors to disclaim all liability and to actually prohibit individuals from writing anything critical about their software for publication. It is rather frightening how that is being pushed.

Cost is pushing organizations, including government, to adopt unsafe technology because it is inexpensive. This is akin to the U.S. Air Force buying cropdusters because they are cheaper than F-16s or the U.S. Navy buying bass boats with outboard motors because they are inexpensive. But that is effectively what we are doing with software. The next generation of Navy aircraft carriers is going to have all weapons systems, propulsion, and command and control run by the very same system that you use at home to browse the Internet and play computer games. This is the same one that keeps coming up with "blue screens of death," which take on new, grim meaning in a military environment. This is a problem, in part, because those systems are made to sell to everybody. They have the lowest common denominator features and the simplest policy. Demand, again, is part of this. Vendors have to sell to people who have no computer training, many of whom are actually functionally illiterate and could not understand the manual if they wanted to.

If companies put in security controls and turn them on, the volume of calls for help increases many fold. Right now the margin of profit is so slim on much of the software that doubling the number of calls for support would make the product lose money. To remain profitable, they turn off all the confusing features—including all of the security features—that might limit the interoperability of the systems.

We also have a preference for fads, which is part of the issue of

demand over quality. The excitement over wireless computing is an example. Little or no thought is given to the dangers behind that—how easy it is to disrupt, how easy it is to eavesdrop, and all of the various implications behind the lack of security in wireless networking. Another technology in vogue is browsing the Web on cell phones. The convergence that is going on here is dangerous. Although not always reliable in an emergency, cell phones could at least be depended on for stable software. Now companies are marketing, because of user demand, cell phones that allow you to download new features and actually run programs on them. Thus, we are now seeing viruses for cell phones. So you cannot depend on that platform anymore. We are introducing new means of instability because of user demand.

From a policy standpoint, we are seeing, both in industry and in government, incredible investment in short-term solutions but almost none in long-term ones. We need to do basic research in issues of security, not only to develop better approaches, but also to build the next generation of researchers. Instead, most of the research investment is being spent on developing new methods of downloading patches for the same old buggy software, new methods of putting up firewalls to protect the same old buggy software, and new methods of virus protection for the same old buggy software. That is not going to advance us very far toward the next generation of technology.

We have a number of policy decisions that are being made by low-level technical people. They are designing protocols and systems to do what they think is interesting and to push the edge of possibility. Their work is finding its way into the marketplace and it is being adopted. We end up with protocols that are built simply to work, without any concern for issues of resiliency and accountability. For instance, right now, being able to determine where anything comes from on the network is next to impossible. As a result, we are seeing all kinds of emergent problems that were not anticipated and that we have very little ability to deal with. Spamming is one example. A recent study done in Great Britain revealed that 40 percent of the traffic that goes through their commercial Internet service providers is now spam. The Gartner Group has said that spam is basically doubling every four to six months. Spam alone may take down large portions of our network!

We do not have a lot of training going on in the area of regulation and law enforcement. In the last 20 years, almost everybody I have met from law enforcement who has learned enough about computing to do forensic analysis has been able to go into industry, working in software engineering or production, at two or three times the salary (and they do not get shot at as often). Few stay in law enforcement very long. Thus, we do not have the technology there, and the law is certainly not helping us. The criminal law lags behind, as it probably should, but it makes it difficult to do some investigations and enforcement.

We also have the major technology firms pushing for special interest legislation that actually hinders research. Many of the large intellectual content providers, such as Sony and Disney, have supported laws such as the Digital Millennium Copyright Act. This law makes it actionable for a researcher to perform many kinds of investigation into the weaknesses in copy protection protocols. If I were to do work in forensic technologies, I technically could be sued by any of these companies or arrested simply because I am investigating ways to break through security on malicious software—but which could coincidentally be used to circumvent their copy protection methods. Thus, if a company needed to reverse engineer a computer virus to derive a countermeasure, they would be violating Federal law—as would the vendors of any tools to support this effort. If the DMCA had been passed before 1999, most of the technology and efforts used to remediate Y2K problems would have been illegal!

Another bill has been introduced in Congress that would require hardware and software technology to prevent copying without approval from one of these companies (the Consumer Broadband and Digital Television Promotion Act: CBDTPA). It is being discussed as something to save the entertainment industry and promote the use of broadband, but, actually, it will severely weaken information security and reduce our ability to do research and communication.

So these are some of the factors–the high market demand, the push from industry, the interest in short-term results, and the primacy of cost over quality—that make this such a difficult domain. Few individuals work in the area of information assurance and there is little support for what we are doing. Our work is viewed as damaging to

those commercial interests or as increasing the cost of technology. So we are not supported so as to make a lot of progress or respond to some of the problems that are out there.

Steps to Solving the Problems

How can we make a difference? We can be better consumers. We can buy tools that have better quality and are better suited for what we need to do. For example, earlier I mentioned how the Microsoft family of software has tens of thousands of known viruses, and new ones are being reported at a rate of dozens per week. Macintosh OS 9 has fewer than 60 viruses in total, almost none of which run under native OS X (notwithstanding the viruses for Microsoft Word). Unix and Linux have about three. I leave it up to you to decide if this is the sort of factor that should make a difference in what someone should deploy in a security-critical environment.

Second, government and industry need to invest more resources in information assurance research and education. Otherwise, we are not going to see much progress towards the long-term solutions: we are going to be continually in the cycle of patching old problems.

Third, we need to see a significant, prolonged investment by government and industry in building up our research infrastructure, educational resources, and personnel base in information assurance. This requires funding both individual researchers and centers, and including sufficient resources to enable significant basic research.

Fourth, we have to start doing something to hold the vendors and the perpetrators of this mess responsible. Vendors know how to put in better quality. I could write another chapter on how 30 or 40 years worth of research has generated basic principles that are being ignored because they add to the time and costs involved with production. The same kinds of arguments that are being used now by the software firms are the ones that were used by some of the tobacco companies. We should not tolerate that.

And lastly, we need to understand that security is not an add-on. It has to be designed in, and pursued as an on-going goal of operation. There is very little we can do with existing systems now to add something on that will address the majority of risks. All we can do is increase the difficulty of someone exploiting obvious problems. But

security has to be built in as a fundamental, and that simply is not the culture. That is not the tradition and we do not have the infrastructure in place to fix that.

In conclusion, it is hard not to be somewhat pessimistic when being realistic, after looking at what has been happening over the last decade or so. If the consumers, vendors and government would all make quality and security a priority, we might begin to see a change; I'm afraid it may require a major disaster before that happens.

Acknowledgments

Thanks to Mike Atallah, Becky Bace, Matt Bishop, Steve Chapin, Simson Garfinkel, Steve Hare, Jim Hendler, Pascal Meunier, Ken Olthoff, John Richardson, Marv Schaefer, and Gene Schultz for comments on an earlier draft of this document. The opinions and interpretations expressed in this document are mine alone, and do not necessarily reflect the opinions of any of these individuals.