Safety in Cyberspace

Vinton G. Cerf

Abstract: Safety in cyberspace continues to be an elusive objective. This essay explores various metaphors to aid thinking about the means by which safety might be increased. Notions such as cyber-fire-departments or cyber-police-departments as well as models drawn from public health scenarios are considered. The legal frameworks in which safety can be improved and international agreements adopted toward this end are briefly discussed. Users can also contribute to their own safety by adopting various practices that reduce vulnerability to cyber-infection and compromise.

The term *cyberspace*, first coined by novelist William Gibson[1] in 1984, has been absorbed into daily language and extended through the prefixing of *cyber-* to an extraordinary number of words. The effect is to imbue ordinary terms with the mystery and otherworldliness of the Internet, World Wide Web, and other artifacts formed through the use of computers and their interconnection. As the Internet and its World Wide Web application have expanded in geographic scope and use, especially by the general public, we have witnessed the emergence of a wide range of practices in the online environment that mirror their counterparts in the physical world. Of course, the Internet itself is realized in a physical way, but it also forms the basis for a virtual environment fashioned of bits rather than atoms. It is this artificial environment that Gibson labeled "cyberspace."

To distinguish various activities that entail the use of computers from those that involve the physical world, it has become common to label the computer-oriented activities as "cyber-activities" or, sometimes, "e-activities" or "e-objects." We speak of "email" or "e-books" and "e-commerce" as if they occupy a place in a virtual universe. We use our commonplace, real-world activities as models for their cyberspace counterparts. In many respects,

VINTON G. CERF, a Fellow of the American Academy since 1995, is Vice President and Chief Internet Evangelist for Google Inc. Previously, he served at MCI, the Corporation for National Research Initiatives, the Defense Advanced Research Projects Agency, and as a member of the Stanford University faculty. He co-invented the architecture and basic protocols of the Internet and is the co-recipient of the U.S. National Medal of Technology and the Japan Prize.

these models help us appreciate the significance and utility of the networked, online equivalents. Thus, "digital signature" is used to refer to a way of cryptographically "signing" a digital document, giving the sense that the digital signature is the equivalent of one rendered with pen and ink or other marking instrument.

While this terminology is convenient, it often has limited applicability and may even mislead our intuitions about the cyber-equivalents of real-world objects, practices, actions, and events. For example, the "owner" of a printed book is free to sell, loan, hypothecate, destroy, fold, spindle, or mutilate the work without intervention by the copyright holder. The owner is not generally free to reproduce and distribute copies unless such rights have been granted by the copyright holder. Now consider an electronic (digital) book that might be downloaded into an electronic book reader, laptop, "netpad," or other display device. The copy of the book held in the display device may not be transferable to another party, in part for technical reasons (the formats of various book reading devices may be incompatible, for example, or there may not be a mechanism for transfer between even compatible devices) or because such transfers are not authorized by the copyright holder.

While some electronic book distribution services permit users to hold multiple copies of a purchased book if the copies are placed solely on devices owned by the user, transfers among reading devices owned by distinct parties are atypical. This example alerts us to the potential hazards of uncritically adopting models to guide our thinking about policies applied to activities undertaken in cyberspace. Not every familiar action in the real world has an accurate analogue in cyberspace.

Although there is no single, agreed upon definition of *cybersecurity*, it is important to recognize that the term covers more than notions of security specific to the Internet. We have many devices and appliances that rely on computers and communication systems that are not necessarily part of the Internet. That these systems are vulnerable in varying degree seems unquestionable, even when they are not part of or even episodically connected to the Internet. Various kinds of digital media (DVDs, thumb drives, memory sticks) can be the bearers of bad software (malware) or potential infection resulting from a visit to a malware-bearing website.

Broadly interpreted, cybersecurity refers to safety from deliberate or accidental harm in connection with the use of computer-based systems, whether networked or not. How many of us have experienced a hard-disk failure and suffered consequential damage, despite the likelihood that the failure was random, neither induced by any malicious attacker nor triggered by a network-based attack? Then there are the seemingly ever-present "bugs" in software that cause data loss or other detrimental failures. For the purpose of exploring some of the potential risks to consumer use of computing systems and responses to them, the term *cyber-safety* seems somewhat more apt. This term may help induce a broader mindset for thinking about the conditions needed to protect the users of software systems from malfunction and harm, even when the harm is not the consequence of intended malice.

In the twenty-five years since the Internet was operationally launched (January 1, 1983) and Gibson's book appeared (ironically, in 1984), a number of metaphors have been applied to frame thinking about software and the challenges it poses. Not long after personal computing became popular, hackers found ways to make computers ingest software that users did not intend to have loaded into their machines. The general term for this unintended ingestion, *infection*, is drawn from the obvious biological metaphor. The bad software was often called *malware*, shorthand for "malevolent software." In the period before the Internet became widely available, a common means of propagating malware was to place it on physical media such as floppy disks, CD-ROMS, and, later, memory sticks and thumb drives. When targets "read" these media, the malware would enter their computers, where it would carry out its function regardless of user awareness or intent.

The computing community adopted several terms for such malware, including *virus*, *worm*, and *Trojan horse*. Viruses are fragments of software that propagate through some vector (physical media, attachments to electronic mail, or downloads from websites). Worms are programs that propagate themselves, often through the Internet or by installing themselves on physical media. Trojan horse software may be drawn into a target computer by direct action of a hacker or through a virus or worm. The Trojan horse software typically stays resident and hidden until conditions cause it to activate (for example, by receipt of a message from some Internet source, or by some local condition in the computer in which it is lodged).

In recent years, among the most common ways for personal computers to become infected is through the use of browsers that help users "surf" the World Wide Web. In its earliest incarnation, the World Wide Web served up Web pages that were mostly text and imagery, but the last decade has seen the development of high-level programming languages such as Java[2] or JavaScript[3] that can be interpreted ("executed") by or through browsers to increase the level of interactivity and functionality of a Web page. When running on certain browsers and operating systems, these high-level language tools can be used to introduce viruses, worms, and Trojan horses into a target computer. Just as we are vulnerable to infection in the biological real world, our computers are vulnerable to becoming infected with malware in the cyber-world.

The objectives of infection may vary from entertainment without harm to considerable damage and data loss and everything in between. In some cases, the purpose is to obtain information possessed by the user of the computer that has become infected: passwords, usernames, financial account information, and corporate information are examples of the targets for this kind of attack. In other cases, the purpose may be to take control of the computer and use it to mount denial-of-service attacks against other computers that are reachable on the same network (not necessarily the Internet, per se) or to generate huge quantities of spam (unwelcomed commercial email). Sometimes, the target is extortion ("Pay me or I will cause all your information to be wiped out") or some other coercive purpose ("I have found incriminating email messages or images and will publish them unless you take the following actions"). Human ingenuity and malfeasance are alive and well in the cyber-world.

Many of the assets used by Internet attackers are harvested from the world's nearly two billion Internet users: namely, their laptops and desktops. In addition, Web servers have been compromised. The mechanisms of compromise vary, and this essay is not intended to catalog even a fraction of them. In principle, users of the

World Wide Web are frequently subject to compromise when they visit infected server sites that inject malware into their computers simply as a consequence of visiting a particular website.

The compromised machines, especially laptops and desktops, become "zombies" controllable by the attacker. Collections of zombie machines, called "botnets" (for "robot networks"), could include hundreds of thousands to millions of computers. The "botnet generals" control these assets remotely, using them to launch denial-of-service attacks, generate spam email, or carry out other nefarious actions. Ironically, many of these actors have no desire to disable the Internet. Rather, they are parasites seeking monetary gain from their control of these resources. Botnets are often rented out for use by third parties. More ironic, botnet generals jealously guard their ill-gotten resources. To defend against a takeover by a competing botnet, their resident malware sometimes tries to detect attempts by other infecting sources to capture the machine. One wonders whether it is possible to invent a "good botnet" we could all join that would protect us from the "bad guys" – rather like a vaccination! Employing such an idea would be more proactive than using current virus detection software, which generally tries to recognize malware before it is activated (executed) and usually relies on static signatures of previously recognized viruses, worms, and Trojan horses.

It is important to note that the creators and propagators of malware may wish to keep their targets unaware of infection – remaining invisible and unidentifiable as the source of control – in order to use the infected computers to generate income. Their purpose may not necessarily be to destroy or disrupt but to abuse access to computer resources they lack the legal or moral rights to use.

In keeping with the metaphor of infection, we might consider what steps consumers can reasonably take to minimize the risk that their computer-based devices become infected with malware. That is, how can they practice good *cyberhygiene*? Commercial software systems called "anti-virus packages" examine incoming files and messages arriving via email, on physical media, or through browsers to try to detect and filter out known or suspected malware. While not 100 percent capable of detecting every infection, such systems provide a partial defense against the ingestion of malware.

Anti-viral defenses cannot be static because the makers of malware evolve their software in response to these defenses. Users of anti-viral software must install regular updates to stay current with known viruses, worms, and Trojan horses. More recently, these commercial offerings have begun trying to detect malware never seen before (so-called zero day attacks) by detecting various forms of anomalous behavior.

Search companies such as Google also attempt to identify websites housing malware that might be employed to infect a user's computer. As the search engine forms the index of the World Wide Web by "crawling" through all the Web pages on the Internet, the indexing software also looks for possible malware on each site and page visited. If malware is suspected, the crawling software can make note of it. For example, if a user of the Google search service clicks on a hyperlink leading to a site that Google suspects is infected, a bright-red warning page shows up on the screen, notifying the user of the potential hazard of visiting that website.

If a website operator believes that the malware report is incorrect, he or she is invited to visit the StopBadWare site[4] for consultation and assistance to verify the presence and assist in the removal of any malware.

Users can also contribute to increased cyber-safety by observing good safety and security practices. Passwords should be long enough and sufficiently complex to eliminate guessing as a means of "cracking." Incredibly, some users choose the word *password* to "protect" their access to online services. Users should avoid the use of birth dates, addresses, common words, or even common word combinations as passwords. For example, random, pronounceable passwords that alternate consonants and vowels, mixed together with digits and special symbols (such as "horif@mi837" or "5atogesi#37"), can form useful bases for stronger protection.

In the past decade or so, new mechanisms for strong authentication have emerged. Devices with small liquid crystal displays are used to generate cryptographically random six- to ten-digit numbers serving as temporary passwords that expire within tens of seconds. Even if these passwords are detected by malware or seen by other users, they are not reuseable. The use of such techniques is sometimes referred to as "two-factor authentication" because users must offer not only a username and conventional password before access to a service or system can be authenticated, but also the cryptographically generated value from a physical device they have in their possession.

People do forget their passwords. One of the values of the two-factor scheme is that the random password generator does not have to remember anything; it continually generates new passwords (in synchrony with the cooperating server). Some online services prompt users to create "secret" questions to be answered with "secret" answers. Examples include "mother's maiden name" or "name of your pet." One of the problems with such methods is that the answers may be easily discovered through a World Wide Web search. Users need to be very creative about their choices of "secret" questions and answers. It might even be sensible to use incorrect information (for example, don't use your mother's real maiden name). Of course, that tactic requires users to remember the false information, which may be harder than remembering the password they forgot!

Some online services will send users a new password (or their old one) as an email message to an email address that users have configured into their account information for that service. That system is flawed by the fact that if the email account has been compromised, it can be used as the avenue through which many other accounts can be penetrated. An attacker can visit the target site and assert that "I have forgotten my password," prompting the service site to send password information by email, which can then be intercepted by the attacker who has, by some means, obtained the user's email-access password(s).

A smarter practice may be to establish backup email accounts used rarely and perhaps predominantly for password recovery; but generally, allowing passwords to be sent by email can compromise security. Some systems rely on mobile phones, sending a text message with a new password. In other cases, though they are more expensive to implement, true out-of-band methods (postal mail or telephone calls) may be more effective; however, their disadvantages include delay for access or the problem of determining whether the caller in need of a new password is, in fact, whom he or she claims to be!

A commonly reported problem in cyberspace is identity theft. The thief manages to discover sufficient information through Web-surfing, possibly adding information gathered by other means (from *Who's Who* publications or alumni magazines, for example) to make a creditable (no pun intended) application for loans, credit cards, bank checking accounts, and other financial instruments. This problem suggests that consumers should be careful in choosing the information that they share, for instance, in social networking applications, blogs, personal Web pages, email messages, and other online means of communication. Financial institutions are hard-pressed to balance consumer safety with convenience and utility. Some brokerage houses, for example, will not accept email orders (out of concern for timeliness and reliability), and others will insist on voice or fax confirmations of orders.

Consumers must be thoughtful in disclosing personal information on the Internet. The organizations they belong to must be equally careful in deciding what to provide online (such as whether board or staff biographies should include email addresses, telephone numbers, or information about family members).

In their essay for this issue, Deirdre K. Mulligan and Fred B. Schneider explore in greater detail the public health analogy for cyber-safety. The notions of worms, viruses, infections, vaccinations, immune systems, and the like all derive from a biological metaphor for the relationship between software and the engines that interpret it. The idea of public health as a model for defense against computer malware has much to offer, at least as an organizing principle for considering responses to threats to cyber-safety and security.

The public health metaphor also suggests the many distinct but interacting "cyber-life-forms" in the software ecology. Interactions among these cyber-life-forms may occur through networks, including the Internet, or by physical transfers of data using memory sticks, thumb drives, CDs, DVDs, and other devices. Humans inject data into these systems through keyboards, cameras, tracking pads, microphones, and an increasing array of sensor systems. The idea that interacting software systems may produce something as complex as a biological ecology should give us pause as we think about protecting our society from deliberate or accidental malfunctions of the complex software systems that we depend on so heavily.

Protecting public health may involve quarantining, vaccination, and other preventative mechanisms; exhortations to manage diet and exercise regularly; or treatment of chronic illnesses through repetitive consumption of medication. When applied to software systems, however, many questions arise regarding proper analogies in the cyber-environment. What does it mean to quarantine a computer or computer-based system? Who can decide to quarantine and how is it enforced? What software vaccinations (that is, anti-virus software) are necessary? Which are effective? How are these validated? Is there an equivalent to the Food and Drug Administration and, if not, should there be? These and many other questions arise when bio-notions are applied to the cyberspace ecology. Establishing institutional practices and guidelines for safety in cyberspace will likely require thoughtful legislation and regulatory response if the metaphor is to prove practical and useful.

Two other metaphors have emerged in the past two decades: *cyber-police-* and *cyber-fire-departments*, which both fall under the rubric of *cyber-first-responders*.

Many kinds of attacks in cyberspace have analogies in the real world: stalking, fraud, denial-of-service, identity-theft, theft of goods or services, possession or sale of stolen goods, counterfeiting, drug trafficking, and so on. In addition, many conventional crimes use the Internet as an aid to perpetration. However, it is not always apparent that a cyber-attack is, in fact, deliberate or a violation of law. Moreover, laws are not uniform, varying among countries, within countries, and among intra-national jurisdictions. Indeed, in the case of the Internet or cyberspace more generally, jurisdiction may be a slippery notion, owing to the highly geographically distributed nature of the parties involved.

If the metaphor of a cyber-police-department is to be of value in organizing a means for protecting citizens from assault in cyberspace, many questions will have to be answered. What actions are considered violations of law? In which jurisdictions? Are there applicable extradition treaties? Are there agreements for international or interjurisdictional cooperation among law enforcement agencies? Applying existing law to cyber-crimes across many jurisdictions would require a great deal of work. It is also important to recognize that not all incidents that appear to be attacks are the result of deliberate intent. Cyber-law-enforcement will need tools and nuanced facilities to distinguish crimes from accidents or harmful but unintended mistakes.

Now imagine that your home or office building is ablaze. Your first instinct is to call not the police department but the fire department. The objective is to put out the blaze as quickly as possible. The private sector's potential inadequacy to respond to a serious fire is captured by the image of a homeowner standing before his burning residence holding a garden hose. To fight the fire, he needs someone with a big hose, a pump, and a lot of water. A cyber-fire-department might be called to help respond to a cyber-attack that the victim is not able to cope with using locally available tools.

This metaphor has some useful features. The fire department's primary job is to put out the fire. After the fact, it may try to determine the cause of the blaze. If it appears to be deliberate, that is, arson, law enforcement agencies may then be called on for further action. Moreover, the fire department often will offer to inspect buildings for fire code violations or hazards and may also participate in the development of fire codes to help protect the community from poorly designed buildings.

This metaphor, too, raises many questions. Who is permitted to call the cyber-fire-department to fight a cyber-attack? Is it purely voluntary? Can a company call the cyber-fire-department as an anti-competitive tool to disrupt a competitor's business? What is the cyber-fire-department allowed to do to systems that are under attack? The conventional fire department is allowed to break doors, windows, roofs, and walls in its effort to quell a blaze and/or rescue endangered parties. What about the cyber-fire-department? What authorities would need to be granted, and by whom, to allow for effective operation? These and many other questions remain unanswered and are made more complex by the possibility that a cyber-fire is burning across international or other jurisdictional boundaries.

In the commercial sector, perhaps the closest relative to cyber-emergency-response is the customer service department of consumer equipment outlets. The Apple Store comes to mind as an example. At Google, there are Tech Stops populated by engineering experts who analyze malfunctions, debug configurations, and help consumers return to productive use of their computer-based equipment.

Many of the metaphors for dealing with cyber-emergencies draw on the idea of first responders. It may be that the notion of first response is applicable to various situations that will later evolve into much more complex, potentially internationally coordinated responses. For the most part, very few institutions have been crafted on the basis of the models suggested above; thus, their utility as guides for increasing safety in cyberspace remains to be explored.

Responses to risks in cyberspace fall into three broad categories:

1) technical responses to prevent harm and preserve safety;

2) post-hoc detection and punishment regimes; and

3) moral suasion.

We know that none of these approaches, nor even any combination, can absolutely guarantee the preservation of safety and protection from harm. All are capable of mitigation in some degree. For this reason, most responses to potential hazards, threats, and vulnerabilities are treated as risk-management problems.

It is not within the scope of this essay to identify every technical effort bent on preventing harm in cyberspace, but a few illustrative examples may be helpful. In the realm of standards creation, the Internet Engineering Task Force (IETF)[5] and the World Wide Web Consortium (W3C)[6] have begun to introduce mechanisms for improving the security of the Internet's infrastructure. For the most part, these mechanisms operate without the need for user intervention. For example, many Web-based services operate by encrypting the data flowing between the user's browser and the serving website. These protected "tunnels" are set up automatically, generally without user action. Employees using organizational resources remotely (whether from home or from Starbucks) may be required to use what are called Virtual Private Networks, created by end-to-end encryption of the data exchanged. The W3C has developed encrypted versions of the HyperText Transport Protocol (HTTP and HTTPS) that use underlying Internet security standards to provide confidentiality.

Recently introduced mechanisms for securing the Domain Name System – called Domain Name System Security Extensions (DNSSEC)[7] – have been implemented by the Internet Corporation for Assigned Names and Numbers (ICANN) and other agencies that manage Internet domain names. (An example of a domain name is www.icann.org.) Again, operating invisibly to users, the system allows the user's laptop or desktop software to verify that it has the correct Internet address for the destination computer that the user is seeking to reach. These methods mitigate sophisticated attacks that alter the integrity of the directory of locations found in cyberspace.

There are ongoing efforts to create additional standards and practices to protect the routing system in the Internet from malfunction or deliberate misrouting. As in most of the previously mentioned tactics, users are largely oblivious to and do not need to take action to benefit from these methods.

In the Internet environment, one of the earliest institutional responses to the problem of cyberspace hazards was the formation of the Computer Emergency Response Team (CERT) at the Software Engineering Institute of Carnegie Mellon University.[8] The Information Processing Techniques Office of the U.S. Defense Advanced Research Projects Agency funded the creation of CERT in response to the so-called Morris worm.[9] CERT initially focused its efforts on identifying vulnerabilities in the UNIX operating system and its many derivatives. It has since broadened its mandate to include facilitating the formation and coordination of national, local, or private Computer Security Incident Response Teams (CSIRTs), including the formation of the US-CERT.[10]

A Forum of Incident Response and Security Teams (FIRST)[11] emerged from this early initiative, linking many of the incident-response team systems together in an information sharing, and sometimes incident-response coordinating, community.

National law enforcement agencies around the world have expanded their scope of operation and attention to include cyber-crime and forensic assessment of security incidents that may have inimical origin. The American FBI has set up an extensive cyber-crime response system,[12] as have many other federal government agencies. In addition to the intelligence agencies, the U.S. military has formed a new Cyber Command (USCYBERCOM)[13] as part of the U.S. Strategic Command. Its mission statement reads:

> USCYBERCOM plans, coordinates, integrates, synchronizes, and conducts activities to: direct the operations and defense of specified Department of Defense information networks and; prepare to, and when directed, conduct full-spectrum military cyberspace operations in order to enable actions in all domains, ensure US/Allied freedom of action in cyberspace and deny the same to our adversaries.

In carrying out this mission, the Cyber Command coordinates its efforts across all military departments and other security agencies, including the U.S. Department of Homeland Security, among many others.

A few private firms have attempted to offer *cyber-insurance*: that is, policies that pay off when some terrible cyber-event occurs. Such offerings could potentially play a role in managing risk more effectively. In a mature industry, the insurance company would have a wealth of cases of incursions and other bad events and have both the means and the incentive to identify weaknesses in existing practices. Just as a company might offer fire insurance only if a home security system or smoke alarms are installed, a company might offer cyber-insurance only if a firewall and an intrusion detection system are in operation.

In general, insurance companies would find it financially attractive to discover and encourage various methods of risk management for computer installations and insist that appropriate measures be used either as a requirement for receiving insurance at all, or as a requirement for receiving lower rates.

Note that insurance companies are well placed to determine whether proper procedures have been followed after an event has occurred; thus, costly monitoring of compliance might not be necessary on an ongoing basis. Audit trails of good practice would likely be a natural consequence of thoughtful implementation of good security in business settings.

In this model, insurance companies take on a role as knowledge repositories for risk management techniques, along with an incentive system for ensuring accuracy of their knowledge: they have to pay damages only when they are wrong. Historically, insurance companies operated in this manner in the development of building codes; there is no reason why they could not play a similar role in the future.

At this point, the constraining factor seems to be expertise. Insurance companies, and the actuaries who work there, know very little about the relevant economic risks. There have been noteworthy attempts by economists and computer engineers to share expertise about appropriate forms of risk management, but this field is in its infancy.

Even a cursory search of the World Wide Web turns up many examples of national and international efforts to define cybercrime, create response regimes, and develop tools to defend against attacks.

It is worth considering that in other domains of discourse, efforts have been made to minimize or even inhibit the "militarization" of commonly shared resources. For example, Article IV of the 1967 Outer Space Treaty[14] specifically rules out the placement of nuclear or other weapons of mass destruction in space. In his lengthy treatment of this subject,[15] Detlev Wolter of the UN Institute for Disarmament Research sets the stage for further elaboration of international protections against the harmful use of outer space. In a similar vein, the international UN Convention of the Law of the Sea[16] facilitates common agreements on how the world's oceans, and events on and under them, are to be treated by signatories to the convention.

One can readily imagine the potential for a similar convention regarding uses and practices in cyberspace. One might find it necessary to focus at first on the common, publicly accessible Internet because the term *cyberspace* covers much more virtual ground than the Internet alone. Reaching common agreements about unacceptable behaviors or practices on the Internet could form a basis for reciprocal cooperation in the enforcement of national laws or the protection of the world's networked citizens from abuse and harm. This is not to overlook the many concerns about definitions, permitted actions, coordinating mechanisms, and the like that such an endeavor would create.

The ability to detect abuse and make use of forensic tools to identify perpetrators is the second of the three legs of cyber-safety named above. International cooperation seems appropriate and even necessary if we are to achieve any measure of security and safety in our use of computer-based systems.

Apart from defense against abusive practices, multilateral treaties can create a basis for improved electronic commerce and an increasing sense of safety, or at least protection, in the online environment. We might try to agree on the means by which digital signatures can be treated as the legal equivalent of ink signatures on a paper contract. The technology of and the rules by which cyber-certificates are issued to validate identity in cyberspace could enhance consumer, business, and government confidence in cyberspace. Cooperation among law enforcement agencies, financial institutions, and other commerce-enabling entities has the potential to significantly improve cyber-safety and the growth of cyber-transactions.

Finally, one can imagine that an additional element of such multilateral treaties might include commitments for educating the general public, the private sector, and governments at all levels about best practices and behaviors promoting safety.

There is much to be gained through voluntary practices that improve safety in cyberspace, including the use of strong authentication mechanisms, anti-virus practices, good cyber-hygiene, and international cooperation on improving safety and security in cyberspace. Absent from this essay is a discussion of the general problem of software "bugs" and the ways in which they can create unsafe conditions and even have fatal consequences. Much more research is warranted to make software more reliable.

In addition to seeking formal mechanisms and agreements that promote the general welfare of citizens dependent on the Internet and computer systems, informal cooperation mechanisms among service and software providers also can provide powerful means of response. It is in the realm of law enforcement and diplomacy, however, where formality can enable the protection of cyber-safety. All these tools will be needed if we are to realize the benefits and minimize the hazards of the increasingly complex, powerful, but potentially brittle virtual worlds we have created in cyberspace.[17]

ENDNOTES

[1] William Gibson, *Neuromancer* (New York: Ace Books, 1984).

[2] See http://www.java.com/en.

[3] See http://www.java.com/en/download/faq/java_javascript.xml.

[4] See http://www.stopbadware.org.

[5] See http://www.ietf.org.

[6] See http://www.w3.org.

[7] See http://en.wikipedia.org/wiki/Domain_Name_System_Security_Extensions.

[8] See http://www.cert.org/meet_cert.

[9] See http://www.symantec.com/connect/articles/brief-history-worm.

[10] See http://www.us-cert.gov/index.html.

[11] See http://www.first.org.

[12] See http://www.fbi.gov/about-us/investigate/cyber/cyber.

[13] United States Strategic Command, U.S. Cyber Command fact sheet, October 2010, http://www.stratcom.mil/factsheets/Cyber_Command.

[14] See http://www.state.gov/www/global/arms/treaties/space1.html.

[15] Detlev Wolter, Common Security in Outer Space and International Law (Geneva, Switzerland: United Nations Institute for Disarmament Research, 2006), http://www.unidir.org/pdf/ouvrages/pdf-1-92-9045-177-7-en.pdf.

[16] See http://www.un.org/depts/los/index.htm.

[17] David Clark, Hal Varian, and Harry Wingo were most helpful in reviewing and recommending changes and additions to earlier drafts of this essay.