Title
Ten Years Past and Ten Years from Now

Author
Bishop, Matt

Publication Date
2010-06-01

Ten Years Past and Ten Years from Now

Matt Bishop
Department of Computer Science
University of California at Davis
Davis, CA 95616-8562 USA
bishop@cs.ucdavis.edu

***Abstract*—Ten years ago, computer security was an arcane discipline that many academics did not see as an interesting or deep research area. Today, that perception has changed. Information assurance and computer security touch every aspect of our lives, and the interconnections with more traditional academic disciplines such as analysis of algorithms, operating systems, and network protocols have convinced even the most hardened skeptic that the academy has a place for computer security. In this talk, we will examine the changes in computer security, how it has touched our lives and our society over the past ten years, and speculate what the next ten years will bring.**

Keywords-information assurance, computer security, past, future

I. INTRODUCTION

The past ten years have seen a revolution in the way people look at computers and the digital age. Before then, the average person saw computers and the Internet as an unmitigated boon, extending their world across borders and into areas that they could not previously venture into. When people want to buy a book, they can look on-line for one they would enjoy. The books and music we remember from childhood, which we thought would be just that—memories—now can be found with web searches. The glut of information available has helped us become more knowledgeable about our world, our environment, and ourselves.

Whether the knowledge we gain is good or bad, pleasant or unpleasant, is a different matter. Along with the ability to learn about our world comes the ability of others to learn about us. We may not like what they learn. What they learn may in fact be false. It may allow them to cause us damage, financially or otherwise. What makes these possibilities more threatening is the reach of the Internet over which this information travels. In the past, the people who learned about us were in our communities—physical, work-related, and other. We usually had relationships with people who knew about us, or whom strangers would get information from. But now, people can gather this information without us knowing who they are, or indeed that they are gathering the information.

Security and privacy concerns have grown with the growth of computing and networking. These concerns were not new. Indeed, the first report in the United States to identify computer security as a serious problem, the Ware Report, appeared in 1970 [8]; the Anderson Report [9] followed shortly. This led to the development of the Bell-LaPadula model [10], Saltzer's and Schroeder's design principles [11], and the RISOS and PA projects [12,13] that studied how to examine systems for vulnerabilities—and with that, the field of computer security began.

Over the past ten years, interest and work in the field of computer security has exploded. It changed from a relatively obscure academic discipline, with commercial work driven by government interests and requirements, to an everyday part of the fabric of our lives. When people purchase something over the Internet, they look at the lock on the web browser to be sure the connection over which they enter credit card information is secure. When citizens cast votes on DREs (which are computers), they ask whether those computers can be tampered with to change votes. The principles have not changed. The mechanisms have evolved, as have computers, our infrastructure, and other technology. And all this is far more visible than in the past.

Perhaps much of this visibility is due to the consequences of failures in computer security becoming widely known. For example, identity theft, the stealing of personally identifying information and using it to impersonate (or spoof) another, is widespread and considered a problem by non-technical people. A common piece of advice—given in several countries—is to check your bank and credit card statements and credit report to ensure that you have not been financially compromised in this way. As another example, in the United States, electronic voting systems are now widely distrusted, in large part due to a series of studies that found severe security and assurance problems in those systems [14-18].

This prominence is refreshing. It is alerting people to the fallibility of computers, and to the problems attendant to our society, and us as individuals, relying on them so much. It is also enriching the field of computer security by bringing in elements of other fields such as law, public policy, and other environmental areas; psychology, human factors, and other people-related subjects; and political science, sociology, and other social factors. Protecting computers and data, and assessing and demonstrating their trustworthiness, is not solely a technical discipline any more.

The goal of this paper is to examine computer security as a technical discipline used in a non-technical world. Over the past ten years, how have changes affect the way we use computers? How has the nature of the infrastructure supporting them, and other systems such as the power grid, changed—if at all? And throughout, how do these changes affect the people who use, and rely upon, computers and the infrastructures that support them?

II. How We Use Computers

Over the past ten years, the use of computers in everyday life has grown exponentially. So have the possibility of problems.

A. The Past Ten Years

Government had always both embraced advanced, sophisticated computer use and resisted it. The military, space exploration, and other branches that require massive amounts of computation or obtaining information from places that people cannot venture used computers for calculations, control, and communications. This pace increased as technology improved. Sensor networks gathered information; automated airplanes and drones came into widespread use. Space vehicles relied on information sent from their earth-bound control stations, and transmitted data to those stations.

These applications raised many security issues, all well known. The most surprising occurred when video from a drone military aircraft was intercepted and broadcast [19]. Techniques for securing video transmissions are widely known and used; yet they were not used. Indeed, many computers and networks exist in hostile circumstances, and this example showed what could happen when basic security precautions were not applied.

Assurance issues also arose. One set of issues sprang from the data on which systems operated being untrustworthy. A NASA probe sent to Mars crashed because the probe expected units to be English units, and the control system transmitted metric units—so the probe's rate of descent was faster than expected [20]. A second issue was the lack of assurance of many systems. Government systems around the world have been successfully attacked, and have failed.

This raised the issue of what standards government systems should meet. Many countries use the Common Criteria [21], which provide a basis for evaluating functionality using Protection Profiles, and evaluating assurance, using the Evaluated Assurance Level (EAL) descriptions. During the past ten years, the Common Criteria began to mature as protection profiles for many types of security-related systems, such as firewalls, were introduced, and the meaning and requirements for the higher EAL levels were refined. Independent, accredited test laboratories determined conformance to the profiles and EALs, contrasting with practice in the past. And different countries signed agreements with one another to recognize each other's accreditation, making the Common Criteria a truly international effort.

Branches that deal with sensitive or critical information, or that are centered in traditional technology, tended to use computers sparingly. For example, the United States Social Security Administration uses computers for its large databases, and recently developed plans for providing a computerized interface for consumers to examine their accounts [22]. Many of these branches realized they needed to upgrade their technology in order to provide better service to government personnel (such as legislators and the executive) and to the public. A number of laws defined requirements that these agencies had to meet, and many governments set up new departments (or tasked existing departments) to refine the requirements as needed.

One specific goal of these laws was openness—to provide the public ways to access data about them held by the government, such as pension data or medical data. Such data needed to be protected, and the assurance issue again arose: how do you demonstrate that the data is "safe", for the appropriate definition of "safe"?

Bureaucracies that dealt with compliance did so in the usual manner of bureaucracies: write documents describing what has been done, and then use the information in the documents to evaluate the systems. The United States government's Federal Information Security Management Act (FISMA) mandated extensive reporting. The key question is whether the documentation matched the reality of the systems. Also, the time that security officers spent doing the paperwork limited the time they could spend actually securing systems. The problems posed by this situation became apparent near the end of the decade, and some agencies began to experiment with a more practical approach. They probed systems looking for security holes, and examined how systems were configured. Other agencies simply mandated a specific configuration and system, ensuring conformity to a particular set of requirements.

Connectivity increased drastically, not only within government but also (and more importantly) in the commercial world. Commercial institutions suffered from many of the same problems as government institutions: growth of automation and combining networks. The problems that industry faced, however, had two additional aspects that affected how they dealt with potential security issues.

The first was publicity. As identity theft increased, and people's personal information was exposed more often, governments passed laws requiring companies to notify people whose personally identifiable information was on a compromised computer. Several cases of this type of exposure quickly led to improvements, amid much embarrassment.

The second was financial liability. In many countries the government cannot be sued unless it consents to be sued. But in most countries, individuals *can* take private companies to court. Further, many companies sell stock on the open market, so the value of their stock rises or falls as the company's reputation changes. To such companies it is advantageous to protect data, because of the possible consequences to the shareholders—the company wants to maximize dividends, and typically this means avoiding bad publicity or costly lawsuits.

Another trend has security implications for commercial firms as well as society. We are seeing a *convergence* of communications media. As a simple example, the University of California at Davis has some areas with poor cellular phone coverage—but excellent wi-fi coverage. Many telephones can now detect the poor cell phone coverage, and switch to Voice over IP using the wi-fi coverage. Further, the voice telephone system records messages when someone does not answer the telephone, and those can be sent to a recipient over the Internet or using SMS. Thus, varied technologies can be combined to achieve a single purpose, and to increase the availability of resources.

Convergence raises interesting security questions. For example, an organization may have physical control of its voice telephone network. It does not have physical control of the SMS networks or the Internet. What implications does this have for sensitive calls? On a broader scale, convergence implies that information from a wide variety of sources and networks (of all kinds) is available. This raises a problem of aggregation—given two pieces of data, from which one can conclude little, what deductions will someone be able to make by combining that data? Considering our web browsing, our on-line shopping habits, and our interaction with customer service departments are available to data aggregators, privacy is at risk—as, in some cases, are more tangible things.

The problem of aggregation enabling unwelcome deductions arises in social networks. Take Facebook as an example. It is quite common for prospective employers to look up applicants on Facebook to learn about them. As people tend to use social networks to talk to friends (and make new friends), much personal information becomes available to the viewer. If what the prospective employer sees is deemed unfavorable, the application is not considered further. The difference between doing a web search and a search on a social network is the personal nature of the information available—but that difference is also quickly eliding.

Finally, over the past decade, many people and small companies began using computers. In this case, the users are not system administrators and are not interested in learning the finer points of locking their system down. They simply want the system to work, be reliable, and be secure—although they may not be able to articulate a clear definition of "security." This complicates the problem of security immensely, because occasionally they must take action to protect themselves—and do not know what to do.

Two examples will make this problem clear. One of the goals of Microsoft Corporation's Windows XP Service Pack 2 was to "harden" Windows XP systems so they would be more difficult to compromise. But the effect of applying the patch was to block many ports that game programs, and other third party programs, relied upon. The average computer owner simply saw this as the Service Pack breaking the system. A second example comes from the web browsing interface on cellular phones. Because the screen is so small, it is usually impossible to put the full URL being visited on the screen, so the browsers elide the URL, usually by not displaying characters in the middle. Spoofing attacks may trick visitors into going to an unsafe page with a URL that looks the same as that of a safe page on the cell phone web browser [23]. How can people who are occasional users of computers and cell phones protect themselves?

To summarize:

- Connectivity, and the convergence of different systems, has greatly increased. The increased access often results in changes in paradigms that are inimical to security.
- Government and industry are moving towards a model of supplying service through the Internet. The intended clients are the public at large, whose computer skills and knowledge of security issues range from great to non-existent. As a result, there is an increased emphasis on understanding technology, even for those who do not *want* to understand it, or who simply cannot.
- Aggregation of publicly available data enables the drawing of inferences that either reveal information about people or provide incorrect, and possibly damaging information. The rise of social networking and the increase in data made available on the web have aggravated this potential loss of privacy.

B. The Next Ten Years

The next ten years are likely to bring dramatic changes to computing, and equally dramatic changes to computer security: how it is implemented and how it is perceived.

Convergence will accelerate. This will affect the practice of computer security in several ways. The primary way lies in the composition of domains with different security policies, and the need to resolve this conflict. Even if two domains are secure, their composition may not be secure. As the domains with the data that is converging may have wildly different security policies about how data may be shared and who may access the data, people may not understand what control they are surrendering by tying domains together.

Non-technical matters such as differing laws and legal jurisdictions complicate this. As an example, some countries require that any data crossing the border either not be encrypted, or the cryptographic key be registered with the police. Other countries simply forbid encrypted transmissions. Consider a multinational corporation with people who work all over the world. As the domains converge, how can the corporation protect its secrets from eavesdroppers (who may work for a government) yet obey the laws of all jurisdictions involved?

The problem of aggregation is a second issue. As data increases, and becomes available to more people, data aggregators will become better at building profiles of people. In some cases, this will be helpful and harmless. In other cases, the profile will reveal information that is correct but that the subject wishes to keep private. In perhaps the worst cases, the profile will be wildly inaccurate, but believed. Such erroneous descriptions could cause irreparable harm—for example, aggregating visits to web sites discussing illegal narcotics and inferring the person is looking for those drugs (but in reality, the user is writing a high school report about the effects of those drugs). Computer security and information assurance will examine new ways to deal with these problems.

An interesting question is how. Two paths seem possible, both drawn from handling covert channels. The first is to provide mechanisms to minimize exposure of information. But if the information has been exposed, experience shows that concealing it again is generally not possible. The second approach is to add noise, so that the aggregator cannot determine what information about the individual is true. This increases the difficulty of drawing accurate inferences. But it also increases the possibility of inaccurate inferences drawn.

Industry trade groups, or governments, will develop new standards governing the gathering of data—whether to protect the users or the aggregators is unclear at this point. An obvious question is ensuring compliance with these standards.

Compliance will be one of, if not the, most important topics in computer security. The days of checking compliance by reading documents will pass, and because the documents report what the writers want to report, or think is true, and often bear little to no relation to the reality of the systems and sites. Compliance validation will continue to shift to hands-on validation, including penetration testing and examination of configurations, software, and hardware. Standardized configurations will become common; this will enable rapid deployment of systems that will enforce standardized organization security policies. While checklists and documentation will still be needed, during the next ten years they will not be enough to demonstrate compliance with security standards.

The use of standardized systems, configurations, and software will force organizations towards a model of central administration. Vendors will follow suit; witness Microsoft's new operating systems validating that they are not pirated before being activated. The motivation behind this shift will be to reduce the number of vulnerable systems in use throughout the Internet. But the vendor shift will cause more problems than the organizational shift mentioned earlier.

The reason lies in the singular nature of an organization's security policy, and the multiple *de facto* security policies of ordinary users. An organization can decree a single policy, or a set of well-defined policies, and implement configurations to support them. Exceptions can be added as needed. But vendors must cope with systems configured as text processors, as gaming systems, as systems used in a variety of small businesses, and so forth, so vendors cannot simply push patches to systems without risking problems. So, different models for vendor administration of systems will develop. For reasons of liability and good public relations, these will evolve into opt-in systems, where the user must affirmatively join the administration program. This way, the vendor can ask questions to determine what software is safe to patch, and what software is problematic.

This leads to another area in which security will improve: communication with users. The "users" here are the average user, not those who are technologically savvy. Asking a literary agent to construct a security policy, for example, is absurd. The average literary agent is skilled with words, books, and negotiating contracts—and not with the attributes of security policies. Over the next ten years, people will study how to construct security policies, and from them configurations, that will protect ordinary users. In addition, notification of problems (such as the discovery of a computer virus in an attachment) will evolve from obscure messages like "Email-Worm:VBS/LoveLetter found in attachment; moving to quarantine" to "Bad data attached—deleting", with an option for the user to get details if desired. The paradigm will shift from a technologically oriented one to a human-oriented one.

This attention to communication will have a number of desirable effects. First, the number of non-secure systems will decrease. Second, as programs are adapted to take advantage of a deeper understanding of psychology and human factors, people can make more effective use of the software and of computers. Third, people will be less intimidated and confused by arcane messages, and so will be more confident that their systems are doing what they want, leading to a perception of improved security.

Finally, social networks will continue to grow in number, size, and complexity—not surprisingly, because humankind consists of social beings. Social networks will also expand from meeting places for friends to aids for therapy, medical treatment, and other functions that enhance the quality of life for groups of people. This information will be available to trusted people on the network—and now, a breach of any one of those accounts will lead the attacker to the sensitive information. Thus, this dissemination of personal information will cause personal information to spread. This means that either users of social networks will accept the loss of privacy, or the term "privacy" will be redefined to capture the needs that evolve in the face of ubiquitous social networking.

Underlying these uses of computers is a set of infrastructures, called the "computing infrastructure," that is evolving as the use of computers is evolving. Its security will evolve as well. We next turn to an examination of its evolution.

III. The Computing Infrastructure

"Infrastructure" is "a collective term for the subordinate parts of an undertaking; substructure, foundation" [1]. Generally, in computing, it refers to networks such as the Internet, and management and support software, hardware, and other entities that support computation. Note that "computing infrastructure" is actually imprecise, and its exact meaning depends upon the environment in which the relevant computing occurs. For example, computers that run life support equipment in a hospital (and which are not connected to any networks) have as infrastructure the people and procedures that maintain them, and the power grid that supplies the energy to keep the systems running. For home computing, the infrastructure includes the Internet and DNS servers, routers, and gateways that enable the World Wide Web to function.

A. The Past Ten Years

The growth of the use of the World Wide Web has made us realize our reliance on its supporting infrastructure and the problems in that infrastructure.

Issues of trust abound. The foundational IP, TCP, UDP, and other layer 3 and 4 protocols were developed to provide robust networking. At the time, security was focused on the end points; indeed, the first security-related RFC [2] dealt with host-level compromises. Thus, additional mechanisms must provide confidentiality and integrity in the face of determined attacks.

Many mechanisms have been tried; some worked, some did not. In the 1990s, Netscape developed the Secure Socket Layer (SSL) protocol to provide transport layer confidentiality and integrity. After some years of experience with this protocol, the IETF developed a successor, the Transport Layer Security

(TLS) protocol, and in the past decade the use of both SSL and TLS increased dramatically. Both use public key mechanisms to authenticate end points, so when one connects to a vendor, the vendor supplies a certificate, and the client's software can validate that the certificate is properly signed.

In computer security, one should always ask *exactly* what terms mean. Here, "properly signed" means that the certificate signature can be validated against a list of known signers—"issuers" or, more properly, "certification authorities" (CAs). These implement several different types of public key infrastructures (PKIs). At the beginning of the decade, may hoped that a single, cohesive PKI could be designed and implemented. Such a PKI system would simplify many aspects of certificate-based key management. Perhaps more importantly, such a structure would provide a framework for certificate recipients to determine the degree of trust they could place in the identity within the certificate—assuming the CAs followed the policies and procedures embedded in the framework. No such coherent framework emerged, and many "root CAs" currently exist, each with their own policies and procedures for validating the identity of those to whom they issue certificates.

Advances in cryptography changed many of the parameters of the existing public key systems and introduced new algorithms for protecting data. The number of bits in the modulus for RSA doubled; flaws were found in several cryptographic hash functions, and the Advanced Encryption Standard (AES) replaced most uses of the Data Encryption Standard (DES). The first practical identity-based encryption scheme [3,4] was developed.

These mechanisms are tools for improving the assurance of the infrastructure. A good example of their use lies in the attempts to harden the Domain Name System (DNS) server infrastructure. *DNS cache poisoning* is an attack where a bogus DNS record is appended to a legitimate one being sent in response to a query. Both the legitimate response and the bogus response are cached. When the victim tries to resolve a name served by the bogus record, the data in the bogus record is used. A protocol called DNSSEC supports digitally signed DNS records. As the bogus record is illegitimate, it will either be unsigned or signed with the wrong key, and hence not validate properly. The resolver could then reject the cached record as untrustworthy. Unfortunately, DNSSEC has not been widely adopted for a variety of reasons, including complexity and the overhead induced by early versions of the protocol.

Two versions of the IP protocol are in use. By far, the most common is the IPv4 protocol, released in 1981. A successor version, IPv6, was released in 1998 and provided many security enhancements. During the last decade, IPv6 began to spread—slowly. Systems offered network stacks to handle both IPv4 and IPv6, and some of the security mechanisms in IPv6 (known collectively as "IPSec") were made available for IPv4.

Because of the increase in sophistication and scale of network-oriented attacks (such as 'botnets and other large-scale distributed attacks), forensics and the ability to trace packets to understand how attacks work became more important. The research area of packet traceback increased in importance. This led to concern with accountability and attribution, and in particular how to determine the *ultimate* source of attacks. Various forensic and analytic techniques have been developed, largely on an *ad hoc* basic, to overcome many of the problems that lack of accountability and attribution mechanisms cause. Those areas also are increasing in importance.

When important legal processes rely upon the Internet, the need for these attributes (accountability, attribution, integrity, and so forth) becomes clear. As an example, in California an effort began to enable on-line filing of documents associated with the purchase of a house. These documents, once filed, establish ownership; they cannot be removed. (So, for example, if a court rules the sale invalid, a copy of the court order is also filed—but the now-invalid sale documents remain in the file.) Thus, if these documents change after signature, for example in transmission, the results would be at best a legal nightmare. Clearly, this requires a hardened infrastructure that enables the *naming* and authentication of the individuals and organizations as well as the integrity (immutability) of the signed documents.

Interestingly, the focus on securing the transmission of those documents at first overlooked the need to secure the endpoints—the hosts—on which these documents were generated and signed. The infrastructure aspect of this issue is that many of the endpoints were to be general-purpose systems that themselves relied on the infrastructure to be configured and to download the programs that would be used in the generation of the documents. Thus, a Trojan horse or other malware could result in bogus documents being recorded even though the document was not altered in transmission. It is simply altered after signing but before transmission.

The concern about the security of the Internet reflected concerns about other large-scale networks, most notably those using Supervisory Control And Data Acquisition (SCADA) systems to manage power, water and natural gas distribution, and other management functions. SCADA systems are typically very simple—they were designed decades ago—and both the systems and the protocols they use are extremely vulnerable to attack [24]. The past decade has seen a dramatic increase in the interconnection of these networks to general-purpose networks (such as the Internet) to ease management and data gathering and processing. The security problems such availability could pose has led to work in hardening SCADA systems and protocols, and examining how connecting highly sophisticated data collection and regulation devices such as smart meters affects the security of those SCADA networks. In particular, the integration required to have the old technology work with the new technology, and *vice versa*, must take security into considerations—and often, the precise policies involved are not well defined (or even defined). This is the composition of different domains mentioned in the discussion of convergence, above.

A key problem in all this work has been testing. Protocols and work that on paper appear to scale often fail miserably when deployed, so researchers need to test their work on testbeds with complexity and size comparable to the networks for which those solutions are designed. This has led to the development of some testbeds that can simulate large-scale networks [26,27], and near the end of the decade an international testbed called "Global Environment for Network

Innovations" (GENI) began to grow [28]. In security, such testbeds are particularly important because testing defenses requires that attacks be simulated or actually launched (to see if the defense actually works)-and doing so on a production network such as the Internet could have catastrophic consequences on people who were not part of the experiment.

Additionally, computer security as a discipline has not focused much on scientifically rigorous experiments. Typically, data is gathered for a particular domain, or it is simulated, and used as the basis for deriving results. The problem is that the data is often unavailable to others, inhibiting verification of results, or the data is substantially different than that others would encounter. McHugh [5] ably discussed the nature of the problem, using testing of intrusion detection systems as his example. By the end of the decade, experiments in computer security became a topic of great interest.

Finally, during the decade, the public became aware of the fragility of many components of the infrastructure. Identity theft, long considered something that "rarely happened," became common enough that financial institutions such as banks implemented programs to protect their clients. The use of "smart meters" in San Francisco sparked resistance [6], and the rise of "click fraud" and spam brought home to average users the risks of trusting the Internet infrastructure completely.

To summarize:

- The assurance of the infrastructure currently is not suitable for highly secure computations and applications. Part of the problem is that components have varying levels of trust. Indeed one entity may trust part of the infrastructure completely, and another entity not at all. This causes problems in composition.
- The injection of sophisticated technology into an infrastructure that was not designed to support it causes problems of security because the features and power of the new technology cannot be controlled by the old technology. This problem effectively layers security mechanisms onto systems (here, infrastructure) not designed with security as a primary focus, or designed with limited threat models that do not take into account new threats.
- The focus on protecting the infrastructure often overlooked that the end points communicating over the infrastructure needed protecting—and that protection relied on the infrastructure.
- Ordinary users learned that the Internet infrastructure was sometimes untrustworthy. The basis for this was twofold: first, occasional perceived failures in network security (for example, due to invalid certificates); second, because components of the infrastructure gather information that the user did not expect or want gathered.

B. The Next Ten Years

As with the use of computers, the next ten years may see dramatic changes in the infrastructure.

The first change will be the realization that many of our dreams for the infrastructure are simply infeasible. In particular, the societal complexity of implementing a standard, universally accepted PKI will prevent that from ever happening. As a simple example, under what conditions will the People's Republic of China accept a PKI system that legitimized a CA run by the government of the Republic of China, and *vice versa*? Instead, the current system of multiple PKIs, each with its own root-level CA and (possibly) subordinate CAs will continue. This allows different PKI models to coexist—and this preserves one of the Internet's longest-enduring traditions.

Anonymity may seem incompatible with certificate-based PKIs—after all, a certificate binds an identity to a public key—but in fact it is not. People often wish to remain anonymous. One reason is to protect themselves from retaliation, for example by an employer or a repressive government. Another reason is to spare themselves embarrassment in the community. Yet they will want to ensure their messages are not altered. Thus, they need a PKI with varying degrees of subject identity validation. Such a model exists in PGP, called the "web of trust." Essentially, anyone can sign anyone's certificate, and indicate the level of trust in the validation of identity (ranging from "untrusted" to "ultimate trust"). One interesting feature of the lack of a central CA for the web of trust is that what is "ultimate trust" for one signer may correspond to "medium trust" for another. Under such a model, anonymity is easy to achieve. Under a hierarchical model, one must use a CA that issues certificates without vouching for the correctness of the identity.

Anonymity protects people, but it also helps people hide from the consequences of their actions. Attackers use this feature to keep from being caught. Thus, there is great interest in attribution for a variety of reasons. Attribution of packets to origins can help law enforcement trace attacks to their origins, and apprehend the malefactors—but it can also be used to trace those who speak unwelcome truths or unpopular opinions. Perhaps most interesting is that complete attribution is probably equally undesirable for law enforcement. For with complete attribution, malefactors could detect that law enforcement is looking at their network connections and systems! As it is probably impossible to determine who is a law enforcement officer entitled to non-attribution (indeed, the likelihood of North Korea recognizing a law enforcement officer from the government of South Korea is low), simply requiring attribution solves little [25]. Policy work on attribution, and the consequences of attribution, will grow.

This is an exemplar of a much larger trend, that of social impacts of security policies and mechanisms. As another example, forensic logging and auditing within the infrastructure, which requires attribution of various sorts, is likely to improve drastically as advances in technologies enable the recording and storage of more data as well as faster real-time analysis of data. This will help administrators detect and counter attacks with increasing effectiveness as the next decade passes. It is axiomatic that the law changes slowly, though, because society needs constancy. Thus the legal standards for extracting, presenting, *and interpreting* the data in court will begin to change, but not change much by the end of the decade.

Secure technologies will become more numerous and pervasive throughout the infrastructure. SSL and TLS will continue to coexist, and the use of DNSSEC, IPSec, and IPv6 will increase. The increase will be slow at first, and then more rapid as the decade ends. My reasoning is that societal and non-technical considerations, rather than technical ones, will drive their adoption. Were security considerations the primary driver, all three protocols would used be far more than they are now. But as government and corporate (and, perhaps, individual) needs for secure and trustworthy communication increases, mandates from organizations will require a change to those technologies that can provide better security and trust. How well the change will work, and whether the provided security services will meet the needs, no-one can predict.

The basis for this security (and other new) technology will lie in its management. Current management practices in most institutions view security as a necessary evil, and something that is to be overridden when it is inconvenient or obstructs work. These will change, because security will be seen as symbiotic with the functions of the organization. This is a result of the institutional imperative, which states that "every action or decision of an institution must be intended to keep the institutional machinery working" ([7], p. 49).

As the complexity of managing the different aspects of an infrastructure increases, so will the need for management tools. In the next decade these tools will become unified, so that one interface enables many different types of controls, including both security configuration and testing. The security of these tools themselves will be critical, as will the assurance that these tools will perform properly. This echoes another trend, namely teaching secure design and implementation of software. This material will expand greatly in scope and, by the end of the decade, be considered as important a supportive discipline as proper language is in essay and literature courses.

Finally, the concept of virtualization, and in particular virtual networks, will provide a mechanism for enhancing the security of the infrastructure. Consider the infrastructure supporting a set of virtual networks, with virtual routers, gateways, and other support. Then, if one virtual network is compromised, other virtual networks are not affected. Hence, a flooding attack, or a compromise of a router, will become more difficult because the virtual networks will be limited in bandwidth. This "virtual infrastructure" is currently infeasible over a wide area because of the limits that the technology imposes. As the technology improves, this barrier will recede.

IV. Conclusion

Any prognostication of the future is dangerous. One can extrapolate existing trends, but an unexpected event, discovery or development can make such predictions wrong. Thus, the above speculations should be treated as just that: informed guesses that may, or may not, be accurate.

What is certain is that security will continue to be a people-oriented subject. The theme of technology, society, and individuals interacting, often in undesired ways, has permeated the problems and solutions discussed here. In a sense, computer security seeks to solve a "moving target" because our notions of security and privacy evolve along with society. Indeed, co-existing societies may define them differently; contrast the Soviet Union's notion of security being primarily economic with the United States' notion of security being primarily about personal rights. Which view is "right" or "wrong" is not the point. That the mechanisms that support the different notions of "security" differ is the point.

Rather than fear these conflicts and view them as threats, societies that have experienced them have grown as ideas that did not work were discarded, ideas that did work were adopted, and from the failures and successes new ideas were synthesized. Societies that rigidly cling to one idea, and suppress conflicting ones, tend to collapse; those that adapt tend to survive and, indeed, prosper.

Perhaps this is the key lesson of the past ten years, that people have different ideas of security that must co-exist. No single notion of "privacy" or "security" will dominate. Like people's opinions and beliefs, "privacy" and "security" will have different, conflicting definitions, and so differing and at times conflicting mechanisms must provide privacy and security. How they will interact, and the results of those interactions, is a question the answer to which no-one can predict.

Acknowledgment

The author gratefully acknowledges the support of the National Science Foundation under grants CNS-0716827, CNS-0831002, and CNS-0905503. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

References

[1] J. A. Simpson and E. S. C. Weiner. The Oxford English Dictionary. Second Edition. Oxford: Clarendon Press, 1989.

[2] B. Metcalfe. The Stockings Were Hung by the Chimney with Care. RFC 602, December 1973. Available at http://www.rfc-editor.org/rfc/rfc602.txt.

[3] D. Boneh and M. Franklin, "Identity-Based Encryption from the Weil Pairing," Advances in Cryptology—Proceedings of CRYPTO 01, pp. 213–229, 2001.

[4] X. Boyer and L. Martin. Identity-Based Cryptography Standard (IBCS) #1: Supersingular Curve Implementations of the BF and BB1 Cryptofystems. RFC 5091, December 2007. Available at http://www.rfc-editor.org/rfc/rfc5091.txt.

[5] J. McHugh. Testing Intrusion Detection Systems: A Critique of the 1998 and 1999 DARPA Intrusion Detection System Evaluations as Performed by Lincoln Laboratory. ACM Transactions on Information System Security 3(4) pp. 262–294 (2000).

[6] A. Werner. PG&E Smart Meter 'Rebellion' Growing. CBS Broadcasting Inc. March 11, 2010. Available at http://cbs5.com/local/pge.smart.meters.2.1555294.html.

[7] R. Kharasch. The Institutional Imperative: How to Understand the United States Government and Other Bulky Objects. New York, NY: Charterhouse Books (1973).

[8] W. Ware. Security Controls for Computer Systems: Report of Defense Science Board Task Force on Computer Security. RAND Report R609-1. Santa Monica, CA: Rand Corp, 1970.

[9] J. Anderson. Computer Security Technology Planning Study. Technical Report ESD-TR-73-51. Bedford, MA: ESD/AFSC, Hanscom Air Force Base, 1972.

[10] D. Bell and L. LaPadula. Secure Computer System: Unified Exposition and Multics Interpretation. Technical Report MTR-2997 Rev. 1. Bedford, MA: MITRE Corporation, 1975.

[11] J. Saltzer and M. Schroeder. The Protection of Information in Computer Systems. Proceedings of the IEEE 63(9) pp. 1278–1308 (1975).

[12] R. Abbott, J. Chin, J. Donnelley, W. Konigsford, S. Tokubo, and D. Webb, Security Analysis and Enhancements of Computer Operating Systems. Technical Report NBSIR 76-1041. Washington DC: ICET, National Bureau of Standards (1976).

[13] R. Bisbey II and D. Hollingsworth. Protection Analysis: Final Report. Technical Report ISI/SR-78-13. Marina Del Rey, CA: Universioty of Southern California Information Sciences Institute (1978).

[14] RABA Innovative Solution Cell. Trusted Agent Report Diebold AccuVote-TS Voting System. Columbia, MD: RABA Technologies LLC. (2004).

[15] A. Yasinsac, D. Wagner, M. Bishop, T. Baker, B. de Medeiros, G. Tyson, M. Shamos, and M. Burmester. Software Review and Security Analysis of the ES&S iVotronic 8.0.1.2 Voting Machine Firmware. Tallahassee, FL: Security nd Assurance in Information Technology Laboratory, Florida State University (2007).

[16] A. Kiayias, L. Michel, A. Russell. And A Schvartsman. Integrity Vulnerabilities in the Diebold TSX Voting Terminal. Storrs, CT: VoTeR Center, University of Connecticut (2007).

[17] M. Bishop. Overview of Red Team Reports. Sacramento, CA: Office of the Secretary of State of California (2007)

[18] Project EVEREST (Evaluation and Validation of Election-Related Equipment, Standards, and Testing) Risk Assessment Study of Ohio Voting Systems: Executive Report. Columbus, OH: Office of the Secretary of State of Ohio (2007).

[19] S. Gorman, Y. Dreazen, and A. Cole. Insurgents Hack U. S. Drones. Wall Street Journal p. A1 (Dec. 17, 2009).

[20] J. Oberg. Why the Mars Probe Went Off Course. IEEE Spectrum 36(12) pp. 34-39 (Dec. 1999).

[21] Common Criteria for Information Technology Security Evaluation, Version 3.1 Revision 3. Available at http://www.commoncriteriaportal.org/

[22] L. Osterweil, L. Millett, and J. Winston. Social Security Administration Electronic Service Provision: A Strategic Assessment. Washington DC: National Academies Press (2007).

[23] Y. Niu, F. Hsu, and H. Chen. iPhish: Phishing Vulnerabilities on Consumer Electronics. Proceedings of the 1st Conference on Usability, Psychology, and Security, article 10. Berkeley, CA: USENIX Association (2008).

[24] V. Igure, S. Laughter, and R. Williams. Security Issues in SCADA Networks. Computers & Security 25 pp. 498–506 (2006).

[25] M. Bishop, C. Gates, J. Hunker. The Sisterhood of the Traveling Packets. Proceedings of the 2009 Workshop on New Security Paradigms pp. 59–70 (2009).

[26] R. Bajcsy *et al*. Cyber Defense Technology Networking and Evaluation. Communications of the ACM 47(3) pp. 58–61 (2004).

[27] B. Chun *et al*. PlanetLab: An Overlay Testbed for Broad-Coverage Services. ACM SIGCOMM Computer Communications Review 33(3) pp. 3–12 (2003).

[28] Global Environment for Network Innovations (2006). Available at http://www.geni.net.