1

INTRODUCTION

CHAPTER OUTLINE

National Cyber Threats, Vulnerabilities, and Attacks 4
Botnet Threat 7
National Cyber Security Methodology Components 9
Deception 11
Separation 13
Diversity 16
Consistency 17
Depth 19
Discretion 20
Collection 21
Correlation 23
Awareness 25
Response 26
Implementing the Principles Nationally 28
Protecting the Critical National Infrastructure Against Cyber Attacks 29
Summary 32
Chapter Review Questions/Exercises 33

Somewhere in his writings—and I regret having forgotten where—John Von Neumann draws attention to what seemed to him a contrast. He remarked that for simple mechanisms it is often easier to describe how they work than what they do, while for more complicated mechanisms it was usually the other way round.

Edsger W. Dijkstra[1]

National infrastructure refers to the complex, underlying delivery and support systems for all large-scale services considered absolutely essential to a nation. These services include emergency response, law enforcement databases, supervisory control and data acquisition (SCADA) systems, power control networks,

[1] E.W. Dijkstra, *Selected Writings on Computing: A Personal Perspective*, Springer-Verlag, New York, 1982, pp. 212–213.

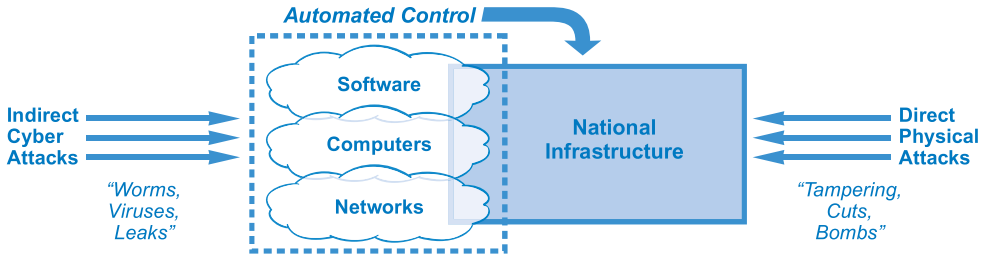

Figure 1.1 National infrastructure cyber and physical attacks.

military support services, consumer entertainment systems, financial applications, and mobile telecommunications. Some national services are provided directly by government, but most are provided by commercial groups such as Internet service providers, airlines, and banks. In addition, certain services considered essential to one nation might include infrastructure support that is controlled by organizations from another nation. This global interdependency is consistent with the trends referred to collectively by Thomas Friedman as a "flat world."[2]

National infrastructure, especially in the United States, has always been vulnerable to malicious physical attacks such as equipment tampering, cable cuts, facility bombing, and asset theft. The events of September 11, 2001, for example, are the most prominent and recent instance of a massive physical attack directed at national infrastructure. During the past couple of decades, however, vast portions of national infrastructure have become reliant on software, computers, and networks. This reliance typically includes remote access, often over the Internet, to the systems that control national services. Adversaries thus can initiate cyber attacks on infrastructure using worms, viruses, leaks, and the like. These attacks indirectly target national infrastructure through their associated automated controls systems (see Figure 1.1).

A seemingly obvious approach to dealing with this national cyber threat would involve the use of well-known computer security techniques. After all, computer security has matured substantially in the past couple of decades, and considerable expertise now exists on how to protect software, computers, and networks. In such a national scheme, safeguards such as firewalls, intrusion detection systems, antivirus software, passwords, scanners, audit trails, and encryption would be directly embedded into infrastructure, just as they are currently in small-scale environments. These national security systems would be

[2]T. Friedman, *The World Is Flat: A Brief History of the Twenty-First Century*, Farrar, Straus, and Giroux, New York, 2007. (Friedman provides a useful economic backdrop to the global aspect of the cyber attack trends suggested in this chapter.)

connected to a centralized threat management system, and incident response would follow a familiar sort of enterprise process. Furthermore, to ensure security policy compliance, one would expect the usual programs of end-user awareness, security training, and third-party audit to be directed toward the people building and operating national infrastructure. Virtually every national infrastructure protection initiative proposed to date has followed this seemingly straightforward path.[3]

While well-known computer security techniques will certainly be useful for national infrastructure, most practical experience to date suggests that this conventional approach will not be sufficient. A primary reason is the size, scale, and scope inherent in complex national infrastructure. For example, where an enterprise might involve manageably sized assets, national infrastructure will require unusually powerful computing support with the ability to handle enormous volumes of data. Such volumes will easily exceed the storage and processing capacity of typical enterprise security tools such as a commercial threat management system. Unfortunately, this incompatibility conflicts with current initiatives in government and industry to reduce costs through the use of common commercial off-the-shelf products.

National infrastructure databases far exceed the size of even the largest commercial databases.

In addition, whereas enterprise systems can rely on manual intervention by a local expert during a security disaster, large-scale national infrastructure generally requires a carefully orchestrated response by teams of security experts using predetermined processes. These teams of experts will often work in different groups, organizations, or even countries. In the worst cases, they will cooperate only if forced by government, often sharing just the minimum amount of information to avoid legal consequences. An additional problem is that the complexity associated with national infrastructure leads to the bizarre situation where response teams often have partial or incorrect understanding about how the underlying systems work. For these reasons, seemingly convenient attempts to apply existing small-scale security processes to large-scale infrastructure attacks will ultimately fail (see Figure 1.2).

As a result, a brand-new type of national infrastructure protection methodology is required—one that combines the best elements of existing computer and network security techniques with the unique and difficult challenges associated with complex, large-scale national services. This book offers just such a protection methodology for national infrastructure. It is based on a quarter

[3] Executive Office of the President, *Cyberspace Policy Review: Assuring a Trusted and Resilient Information and Communications Infrastructure*, U.S. White House, Washington, D.C., 2009 (http://handle.dtic.mil/100.2/ADA501541).

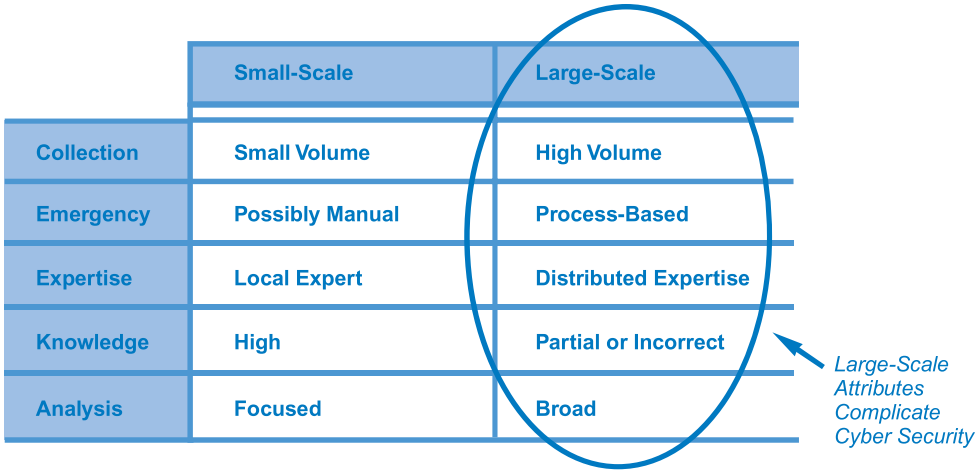

Figure 1.2 Differences between small- and large-scale cyber security.

century of practical experience designing, building, and operating cyber security systems for government, commercial, and consumer infrastructure. It is represented as a series of protection principles that can be applied to new or existing systems. Because of the unique needs of national infrastructure, especially its massive size, scale, and scope, some aspects of the methodology will be unfamiliar to the computer security community. In fact, certain elements of the approach, such as our favorable view of "security through obscurity," might appear in direct conflict with conventional views of how computers and networks should be protected.

National Cyber Threats, Vulnerabilities, and Attacks

Any of the most common security concern—confidentiality, integrity, availability, and theft—threaten our national infrastructure.

Conventional computer security is based on the oft-repeated taxonomy of security threats which includes confidentiality, integrity, availability, and theft. In the broadest sense, all four diverse threat types will have applicability in national infrastructure. For example, protections are required equally to deal with sensitive information leaks (confidentiality), worms affecting the operation of some critical application (integrity), botnets knocking out an important system (availability), or citizens having their identities compromised (theft). Certainly, the availability threat to national services must be viewed as particularly important, given the nature of the threat and its relation to national assets. One should thus expect particular attention to availability threats to national infrastructure. Nevertheless, it makes sense to acknowledge that all four types of security threats in the

conventional taxonomy of computer security must be addressed in any national infrastructure protection methodology.

Vulnerabilities are more difficult to associate with any taxonomy. Obviously, national infrastructure must address well-known problems such as improperly configured equipment, poorly designed local area networks, unpatched system software, exploitable bugs in application code, and locally disgruntled employees. The problem is that the most fundamental vulnerability in national infrastructure involves the staggering complexity inherent in the underlying systems. This complexity is so pervasive that many times security incidents uncover aspects of computing functionality that were previously unknown to anyone, including sometimes the system designers. Furthermore, in certain cases, the optimal security solution involves simplifying and cleaning up poorly conceived infrastructure. This is bad news, because most large organizations are inept at simplifying much of anything.

The best one can do for a comprehensive view of the vulnerabilities associated with national infrastructure is to address their relative exploitation points. This can be done with an abstract national infrastructure cyber security model that includes three types of malicious adversaries: *external adversary* (hackers on the Internet), *internal adversary* (trusted insiders), and *supplier adversary* (vendors and partners). Using this model, three exploitation points emerge for national infrastructure: *remote access* (Internet and telework), *system administration and normal usage* (management and use of software, computers, and networks), and *supply chain* (procurement and outsourcing) (see Figure 1.3).

These three exploitation points and three types of adversaries can be associated with a variety of possible motivations for initiating either a full or test attack on national infrastructure.

Figure 1.3 Adversaries and exploitation points in national infrastructure.

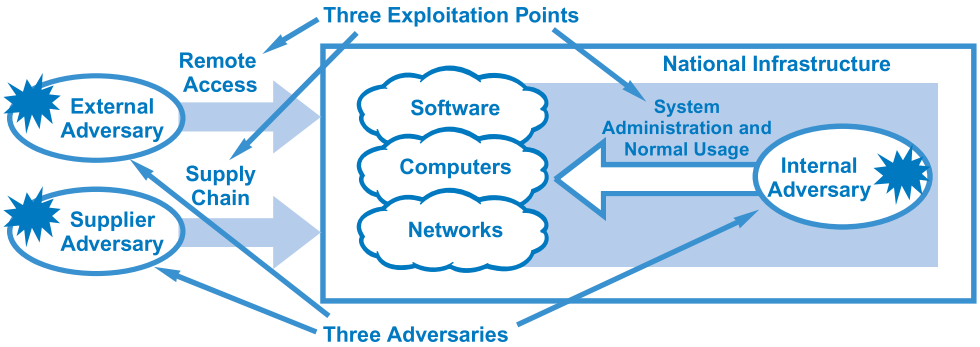

Five Possible Motivations for an Infrastructure Attack

- *Country-sponsored warfare*—National infrastructure attacks sponsored and funded by enemy countries must be considered the most significant potential motivation, because the intensity of adversary capability and willingness to attack is potentially unlimited.
- *Terrorist attack*—The terrorist motivation is also signifi cant, especially because groups driven by terror can easily obtain sufficient capability and funding to perform significant attacks on infrastructure.
- *Commercially motivated attack*—When one company chooses to utilize cyber attacks to gain a commercial advantage, it becomes a national infrastructure incident if the target company is a purveyor of some national asset.
- *Financially driven criminal attack*—Identify theft is the most common example of a fi nancially driven attack by criminal groups, but other cases exist, such as companies being extorted to avoid a cyber incident.
- *Hacking*—One must not forget that many types of attacks are still driven by the motivation of hackers, who are often just mischievous youths trying to learn or to build a reputation within the hacking community. This is much less a sinister motivation, and national leaders should try to identify better ways to tap this boundless capability and energy.

Each of the three exploitation points might be utilized in a cyber attack on national infrastructure. For example, a supplier might use a poorly designed supply chain to insert Trojan horse code into a software component that controls some national asset, or a hacker on the Internet might take advantage of some unprotected Internet access point to break into a vulnerable service. Similarly, an insider might use trusted access for either system administration or normal system usage to create an attack. The potential also exists for an external adversary to gain valuable insider access through patient, measured means, such as gaining employment in an infrastructure-supporting organization and then becoming trusted through a long process of work performance. In each case, the possibility exists that a limited type of engagement might be performed as part of a planned test or exercise. This seems especially likely if the attack is country or terrorist sponsored, because it is consistent with past practice.

When to issue a vulnerability risk advisory and when to keep the risk confidential must be determined on a case-by-case basis, depending on the threat.

At each exploitation point, the vulnerability being used might be a well-known problem previously reported in an authoritative public advisory, or it could be a proprietary issue kept hidden by a local organization. It is entirely appropriate for a recognized authority to make a detailed public vulnerability advisory if the benefits of notifying the good guys outweigh the risks of alerting the bad guys. This cost–benefit result usually occurs when many organizations can directly benefit from the information and can thus take immediate action. When the reported vulnerability is unique and isolated, however, then reporting the details might be

irresponsible, especially if the notification process does not enable a more timely fix. This is a key issue, because many government authorities continue to consider new rules for mandatory reporting. If the information being demanded is not properly protected, then the reporting process might result in more harm than good.

Botnet Threat

Perhaps the most insidious type of attack that exists today is the *botnet*.[4] In short, a botnet involves remote control of a collection of compromised end-user machines, usually broadband-connected PCs. The controlled end-user machines, which are referred to as *bots*, are programmed to attack some target that is designated by the botnet controller. The attack is tough to stop because end-user machines are typically administered in an ineffective manner. Furthermore, once the attack begins, it occurs from sources potentially scattered across geographic, political, and service provider boundaries. Perhaps worse, bots are programmed to take commands from multiple controller systems, so any attempts to destroy a given controller result in the bots simply homing to another one.

The Five Entities That Comprise a Botnet Attack

- *Botnet operator*—This is the individual, group, or country that creates the botnet, including its setup and operation.When the botnet is used for financial gain, it is the operator who will benefit. Law enforcement and cyber security initiatives have found it very difficult to identify the operators. The press, in particular, has done a poor job reporting on the presumed identity of botnet operators, often suggesting sponsorship by some country when little supporting evidence exists.
- *Botnet controller*—This is the set of servers that command and control the operation of a botnet. Usually these servers have been maliciously compromised for this purpose. Many times, the real owner of a server that has been compromised will not even realize what has occurred. The type of activity directed by a controller includes all recruitment, setup, communication, and attack activity. Typical botnets include a handful of controllers, usually distributed across the globe in a non-obvious manner.
- *Collection of bots*—These are the end-user, broadband-connected PCs infected with botnet malware. They are usually owned and operated by normal citizens, who become unwitting and unknowing dupes in a botnet attack. When a botnet includes a concentration of PCs in a given region, observers often incorrectly attribute the attack to that region. The use of smart mobile devices in a botnet will grow as upstream capacity and device processing power increase.

[4] Much of the material on botnets in this chapter is derived from work done by Brian Rexroad, David Gross, and several others from AT&T.

- *Botnet software drop*—Most botnets include servers designed to store software that might be useful for the botnets during their lifecycle. Military personnel might refer to this as an arsenal . Like controllers, botnet software drop points are usually servers compromised for this purpose, often unknown to the normal server operator.
- *Botnet target*—This is the location that is targeted in the attack. Usually, it is a website, but it can really be any device, system, or network that is visible to the bots. In most cases, botnets target prominent and often controversial websites, simply because they are visible via the Internet and generally have a great deal at stake in terms of their availability. This increases gain and leverage for the attacker. Logically, however, botnets can target anything visible.

Home PC users may never know they are being used for a botnet scheme.

The way a botnet works is that the controller is set up to communicate with the bots via some designated protocol, most often Internet Relay Chat (IRC). This is done via malware inserted into the end-user PCs that comprise the bots. A great challenge in this regard is that home PCs and laptops are so poorly administered. Amazingly, over time, the day-to-day system and security administration task for home computers has gravitated to the end user. This obligation results in both a poor user experience and general dissatisfaction with the security task. For example, when a typical computer buyer brings a new machine home, it has probably been preloaded with security software by the retailer. From this point onward, however, that home buyer is then tasked with all responsibility for protecting the machine. This includes keeping firewall, intrusion detection, antivirus, and antispam software up to date, as well as ensuring that all software patches are current. When these tasks are not well attended, the result is a more vulnerable machine that is easily turned into a bot. (Sadly, even if a machine is properly managed, expert bot software designers might find a way to install the malware anyway.)

A DDOS attack is like a cyber traffic jam.

Once a group of PCs has been compromised into bots, attacks can thus be launched by the controller via a command to the bots, which would then do as they are instructed. This might not occur instantaneously with the infection; in fact, experience suggests that many botnets lay dormant for a great deal of time. Nevertheless, all sorts of attacks are possible in a botnet arrangement, including the now-familiar *distributed denial of service attack* (DDOS). In such a case, the bots create more inbound traffic than the target gateway can handle. For example, if some theoretical gateway allows for 1 Gbps of inbound traffic, and the botnet creates an inbound stream larger than 1 Gbps, then a logjam results at the inbound gateway, and a denial of service condition occurs (see Figure 1.4).

Any serious present study of cyber security must acknowledge the unique threat posed by botnets. Virtually any Internet-connected system is vulnerable to major outages from a

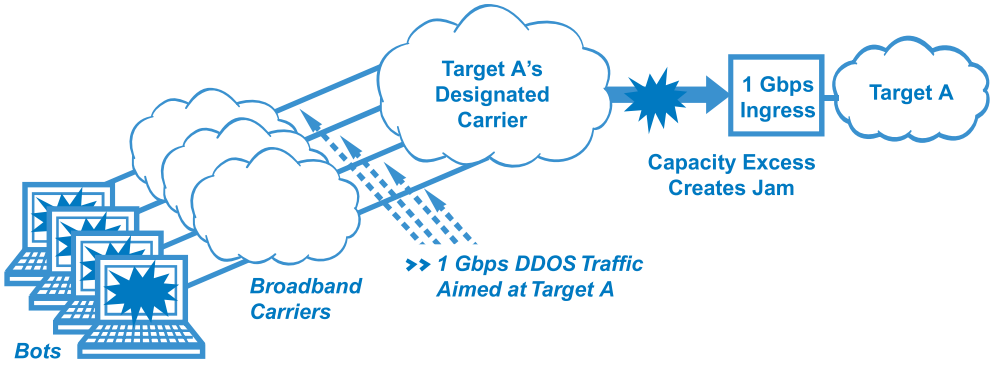

Figure 1.4 Sample DDOS attack from a botnet.

botnet-originated DDOS attack. The physics of the situation are especially depressing; that is, a botnet that might steal 500 Kbps of upstream capacity from each bot (which would generally allow for concurrent normal computing and networking) would only need three bots to collapse a target T1 connection. Following this logic, only 16,000 bots would be required theoretically to fill up a 10-Gbps connection. Because most of the thousands of botnets that have been observed on the Internet are at least this size, the threat is obvious; however, many recent and prominent botnets such as Storm and Conficker are much larger, comprising as many as several million bots, so the threat to national infrastructure is severe and immediate.

National Cyber Security Methodology Components

Our proposed methodology for protecting national infrastructure is presented as a series of ten basic design and operation principles. The implication is that, by using these principles as a guide for either improving existing infrastructure components or building new ones, the security result will be desirable, including a reduced risk from botnets. The methodology addresses all four types of security threats to national infrastructure; it also deals with all three types of adversaries to national infrastructure, as well as the three exploitation points detailed in the infrastructure model. The list of principles in the methodology serves as a guide to the remainder of this chapter, as well as an outline for the remaining chapters of the book:

- *Chapter 2: Deception*—The openly advertised use of deception creates uncertainty for adversaries because they will not know if a discovered problem is real or a trap. The more common

hidden use of deception allows for real-time behavioral analysis if an intruder is caught in a trap. Programs of national infrastructure protection must include the appropriate use of deception, especially to reduce the malicious partner and supplier risk.

- *Chapter 3: Separation*—Network separation is currently accomplished using firewalls, but programs of national infrastructure protection will require three specific changes. Specifically, national infrastructure must include network-based firewalls on high-capacity backbones to throttle DDOS attacks, internal firewalls to segregate infrastructure and reduce the risk of sabotage, and better tailoring of firewall features for specific applications such as SCADA protocols.[5]
- *Chapter 4: Diversity*—Maintaining diversity in the products, services, and technologies supporting national infrastructure reduces the chances that one common weakness can be exploited to produce a cascading attack. A massive program of coordinated procurement and supplier management is required to achieve a desired level of national diversity across all assets. This will be tough, because it conflicts with most cost-motivated information technology procurement initiatives designed to minimize diversity in infrastructure.
- *Chapter 5: Commonality*—The consistent use of security best practices in the administration of national infrastructure ensures that no infrastructure component is either poorly managed or left completely unguarded. National programs of standards selection and audit validation, especially with an emphasis on uniform programs of simplification, are thus required. This can certainly include citizen end users, but one should never rely on high levels of security compliance in the broad population.
- *Chapter 6: Depth*—The use of defense in depth in national infrastructure ensures that no critical asset is reliant on a single security layer; thus, if any layer should fail, an additional layer is always present to mitigate an attack. Analysis is required at the national level to ensure that all critical assets are protected by at least two layers, preferably more.
- *Chapter 7: Discretion*—The use of personal discretion in the sharing of information about national assets is a practical technique that many computer security experts find difficult to accept because it conflicts with popular views on "security through obscurity." Nevertheless, large-scale infrastructure protection cannot be done properly unless a national culture

[5] R. Kurtz, *Securing SCADA Systems*, Wiley, New York, 2006. (Kurtz provides an excellent overview of SCADA systems and the current state of the practice in securing them.)

of discretion and secrecy is nurtured. It goes without saying that such discretion should never be put in place to obscure illegal or unethical practices.

- *Chapter 8: Collection*—The collection of audit log information is a necessary component of an infrastructure security scheme, but it introduces privacy, size, and scale issues not seen in smaller computer and network settings. National infrastructure protection will require a data collection approach that is acceptable to the citizenry and provides the requisite level of detail for security analysis.
- *Chapter 9: Correlation*—Correlation is the most fundamental of all analysis techniques for cyber security, but modern attack methods such as botnets greatly complicate its use for attack-related indicators. National-level correlation must be performed using all available sources and the best available technology and algorithms. Correlating information around a botnet attack is one of the more challenging present tasks in cyber security.
- *Chapter 10: Awareness*—Maintaining situational awareness is more important in large-scale infrastructure protection than in traditional computer and network security because it helps to coordinate the real-time aspect of multiple infrastructure components. A program of national situational awareness must be in place to ensure proper management decision-making for national assets.
- *Chapter 11: Response*—Incident response for national infrastructure protection is especially difficult because it generally involves complex dependencies and interactions between disparate government and commercial groups. It is best accomplished at the national level when it focuses on early indications, rather than on incidents that have already begun to damage national assets.

The balance of this chapter will introduce each principle, with discussion on its current use in computer and network security, as well as its expected benefits for national infrastructure protection.

Deception

The principle of *deception* involves the deliberate introduction of misleading functionality or misinformation into national infrastructure for the purpose of tricking an adversary. The idea is that an adversary would be presented with a view of national infrastructure functionality that might include services or interface components that are present for the sole purpose of fakery. Computer scientists refer to this functionality as a *honey pot*,

Deception is an oft-used tool by law enforcement agencies to catch cyber stalkers and predators.

but the use of deception for national infrastructure could go far beyond this conventional view. Specifically, deception can be used to protect against certain types of cyber attacks that no other security method will handle. Law enforcement agencies have been using deception effectively for many years, often catching cyber stalkers and criminals by spoofing the reported identity of an end point. Even in the presence of such obvious success, however, the cyber security community has yet to embrace deception as a mainstream protection measure.

Deception in computing typically involves a layer of cleverly designed trap functionality strategically embedded into the internal and external interfaces for services. Stated more simply, deception involves fake functionality embedded into real interfaces. An example might be a deliberately planted trap link on a website that would lead potential intruders into an environment designed to highlight adversary behavior. When the deception is open and not secret, it might introduce uncertainty for adversaries in the exploitation of real vulnerabilities, because the adversary might suspect that the discovered entry point is a trap. When it is hidden and stealth, which is the more common situation, it serves as the basis for real-time forensic analysis of adversary behavior. In either case, the result is a public interface that includes real services, deliberate honey pot traps, and the inevitable exploitable vulnerabilities that unfortunately will be present in all nontrivial interfaces (see Figure 1.5).

Only relatively minor tests of honey pot technology have been reported to date, usually in the context of a research effort. Almost no reports are available on the day-to-day use of deception as a structural component of a real enterprise security program. In fact, the vast majority of security programs for companies, government agencies, and national infrastructure would include no such functionality. Academic computer scientists have shown little interest in this type of security, as evidenced by the relatively thin body of literature on the subject. This lack of interest might stem from the discomfort associated with using

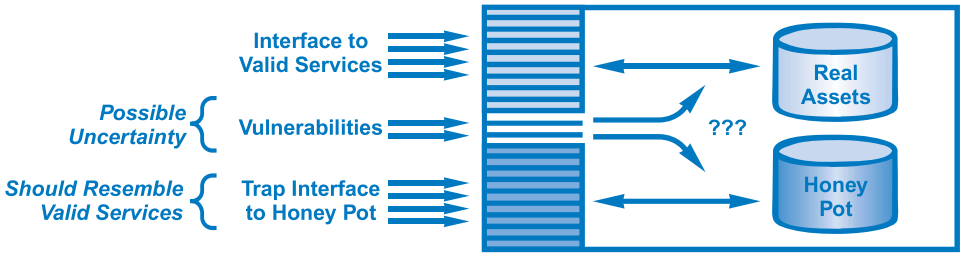

Figure 1.5 Components of an interface with deception.

computing to mislead. Another explanation might be the relative ineffectiveness of deception against the botnet threat, which is clearly the most important security issue on the Internet today. Regardless of the cause, this tendency to avoid the use of deception is unfortunate, because many cyber attacks, such as subtle break-ins by trusted insiders and Trojan horses being maliciously inserted by suppliers into delivered software, cannot be easily remedied by any other means.

Deception is less effective against botnets than other types of attack methods.

The most direct benefit of deception is that it enables forensic analysis of intruder activity. By using a honey pot, unique insights into attack methods can be gained by watching what is occurring in real time. Such deception obviously works best in a hidden, stealth mode, unknown to the intruder, because if the intruder realizes that some vulnerable exploitation point is a fake, then no exploitation will occur. Honey pot pioneers Cliff Stoll, Bill Cheswick, and Lance Spitzner have provided a majority of the reported experience in real-time forensics using honey pots. They have all suggested that the most difficult task involves creating believability in the trap. It is worth noting that connecting a honey pot to real assets is a terrible idea.

Do not connect honey pots to real assets!

An additional potential benefit of deception is that it can introduce the clever idea that some discovered vulnerability might instead be a deliberately placed trap. Obviously, such an approach is only effective if the use of deception is not hidden; that is, the adversary must know that deception is an approved and accepted technique used for protection. It should therefore be obvious that the major advantage here is that an accidental vulnerability, one that might previously have been an open door for an intruder, will suddenly look like a possible trap. A further profound notion, perhaps for open discussion, is whether just the *implied statement* that deception might be present (perhaps without real justification) would actually reduce risk. Suppliers, for example, might be less willing to take the risk of Trojan horse insertion if the procuring organization advertises an open research and development program of detailed software test and inspection against this type of attack.

Separation

The principle of *separation* involves enforcement of access policy restrictions on the users and resources in a computing environment. Access policy restrictions result in separation domains, which are arguably the most common security architectural concept in use today. This is good news, because the creation of access-policy-based separation domains will be essential in the

protection of national infrastructure. Most companies today will typically use firewalls to create perimeters around their presumed enterprise, and access decisions are embedded in the associated rules sets. This use of enterprise firewalls for separation is complemented by several other common access techniques:

- *Authentication and identity management*—These methods are used to validate and manage the identities on which separation decisions are made. They are essential in every enterprise but cannot be relied upon solely for infrastructure security. Malicious insiders, for example, will be authorized under such systems. In addition, external attacks such as DDOS are unaffected by authentication and identity management.
- *Logical access controls*—The access controls inherent in operating systems and applications provide some degree of separation, but they are also weak in the presence of compromised insiders. Furthermore, underlying vulnerabilities in applications and operating systems can often be used to subvert these methods.
- *LAN controls*—Access control lists on local area network (LAN) components can provide separation based on information such as Internet Protocol (IP) or media access control (MAC) address. In this regard, they are very much like firewalls but typically do not extend their scope beyond an isolated segment.
- *Firewalls*—For large-scale infrastructure, firewalls are particularly useful, because they separate one network from another. Today, every Internet-based connection is almost certainly protected by some sort of firewall functionality. This approach worked especially well in the early years of the Internet, when the number of Internet connections to the enterprise was small. Firewalls do remain useful, however, even with the massive connectivity of most groups to the Internet. As a result, national infrastructure should continue to include the use of firewalls to protect known perimeter gateways to the Internet.

Given the massive scale and complexity associated with national infrastructure, three specific separation enhancements are required, and all are extensions of the firewall concept.

Required Separation Enhancements for National Infrastructure Protection

1. The use of network-based firewalls is absolutely required for many national infrastructure applications, especially ones vulnerable to DDOS attacks from the Internet. This use of network-based mediation can take advantage of high-capacity network backbones if the service provider is involved in running the firewalls.

2. The use of firewalls to segregate and isolate internal infrastructure components from one another is a mandatory technique for simplifying the implementation of access control policies in an organization. When insiders have malicious intent, any exploit they might attempt should be explicitly contained by internal firewalls.
3. The use of commercial off-the-shelf firewalls, especially for SCADA usage, will require tailoring of the firewall to the unique protocol needs of the application. It is not acceptable for national infrastructure protection to retrofi t the use of a generic, commercial, off-the-shelf tool that is not optimized for its specific use (see Figure 1.6)

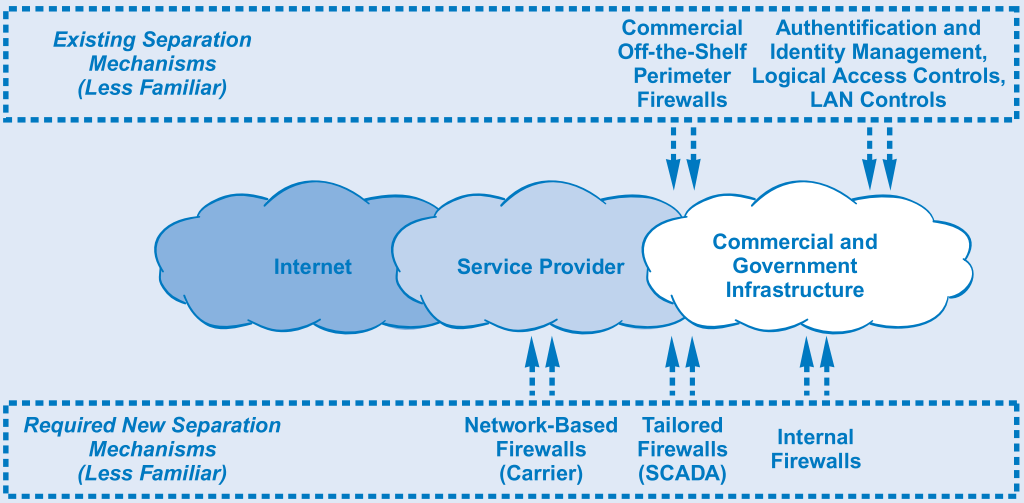

Figure 1.6 Firewall enhancements for national infrastructure.

With the advent of cloud computing, many enterprise and government agency security managers have come to acknowledge the benefits of network-based firewall processing. The approach scales well and helps to deal with the uncontrolled complexity one typically finds in national infrastructure. That said, the reality is that most national assets are still secured by placing a firewall at each of the hundreds or thousands of presumed choke points. This approach does not scale and leads to a false sense of security. It should also be recognized that the firewall is not the only device subjected to such scale problems. Intrusion detection systems, antivirus filtering, threat management, and denial of service filtering also require a network-based approach to function properly in national infrastructure.

An additional problem that exists in current national infrastructure is the relative lack of architectural separation used in an internal, trusted network. Most security engineers know that large systems are best protected by dividing them into smaller

systems. Firewalls or packet filtering routers can be used to segregate an enterprise network into manageable domains. Unfortunately, the current state of the practice in infrastructure protection rarely includes a disciplined approach to separating internal assets. This is unfortunate, because it allows an intruder in one domain to have access to a more expansive view of the organizational infrastructure. The threat increases when the firewall has not been optimized for applications such as SCADA that require specialized protocol support.

Parceling a network into manageable smaller domains creates an environment that is easier to protect.

Diversity

The principle of *diversity* involves the selection and use of technology and systems that are intentionally different in substantive ways. These differences can include technology source, programming language, computing platform, physical location, and product vendor. For national infrastructure, realizing such diversity requires a coordinated program of procurement to ensure a proper mix of technologies and vendors. The purpose of introducing these differences is to deliberately create a measure of non-interoperability so that an attack cannot easily cascade from one component to another through exploitation of some common vulnerability. Certainly, it would be possible, even in a diverse environment, for an exploit to cascade, but the likelihood is reduced as the diversity profile increases.

This concept is somewhat controversial, because so much of computer science theory and information technology practice in the past couple of decades has been focused on maximizing interoperability of technologies. This might help explain the relative lack of attentiveness that diversity considerations receive in these fields. By way of analogy, however, cyber attacks on national infrastructure are mitigated by diversity technology just as disease propagation is reduced by a diverse biological ecosystem. That is, a problem that originates in one area of infrastructure with the intention of automatic propagation will only succeed in the presence of some degree of interoperability. If the technologies are sufficiently diverse, then the attack propagation will be reduced or even stopped. As such, national asset managers are obliged to consider means for introducing diversity in a cost-effective manner to realize its security benefits (see Figure 1.7).

Diversity is especially tough to implement in national infrastructure for several reasons. First, it must be acknowledged that a single, major software vendor tends to currently dominate the personal computer (PC) operating system business landscape in most

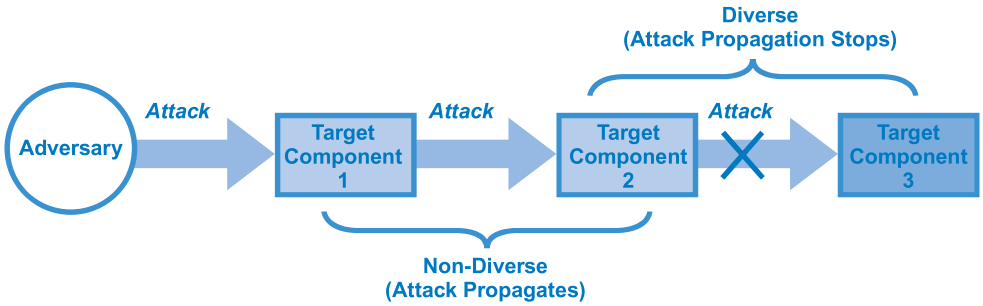

Figure 1.7 Introducing diversity to national infrastructure.

government and enterprise settings. This is not likely to change, so national infrastructure security initiatives must simply accept an ecosystem lacking in diversity in the PC landscape. The profile for operating system software on computer servers is slightly better from a diversity perspective, but the choices remain limited to a very small number of available sources. Mobile operating systems currently offer considerable diversity, but one cannot help but expect to see a trend toward greater consolidation.

Second, diversity conflicts with the often-found organizational goal of simplifying supplier and vendor relationships; that is, when a common technology is used throughout an organization, day-to-day maintenance, administration, and training costs are minimized. Furthermore, by purchasing in bulk, better terms are often available from a vendor. In contrast, the use of diversity could result in a reduction in the level of service provided in an organization. For example, suppose that an Internet service provider offers particularly secure and reliable network services to an organization. Perhaps the reliability is even measured to some impressive quantitative availability metric. If the organization is committed to diversity, then one might be forced to actually introduce a second provider with lower levels of reliability.

Enforcing diversity of products and services might seem counterintuitive if you have a reliable provider.

In spite of these drawbacks, diversity carries benefits that are indisputable for large-scale infrastructure. One of the great challenges in national infrastructure protection will thus involve finding ways to diversify technology products and services without increasing costs and losing business leverage with vendors.

Consistency

The principle of *consistency* involves uniform attention to security best practices across national infrastructure components. Determining which best practices are relevant for which national asset requires a combination of local knowledge about the asset, as well as broader knowledge of security vulnerabilities in generic

infrastructure protection. Thus, the most mature approach to consistency will combine compliance with relevant standards such as the Sarbanes–Oxley controls in the United States, with locally derived security policies that are tailored to the organizational mission. This implies that every organization charged with the design or operation of national infrastructure must have a local security policy. Amazingly, some large groups do not have such a policy today.

The types of best practices that are likely to be relevant for national infrastructure include well-defined software lifecycle methodologies, timely processes for patching software and systems, segregation of duty controls in system administration, threat management of all collected security information, security awareness training for all system administrators, operational configurations for infrastructure management, and use of software security tools to ensure proper integrity management. Most security experts agree on which best practices to include in a generic set of security requirements, as evidenced by the inclusion of a common core set of practices in every security standard. Attentiveness to consistency is thus one of the less controversial of our recommended principles.

The greatest challenge in implementing best practice consistency across infrastructure involves auditing. The typical audit process is performed by an independent third-party entity doing an analysis of target infrastructure to determine consistency with a desired standard. The result of the audit is usually a numeric score, which is then reported widely and used for management decisions. In the United States, agencies of the federal government are audited against a cyber security standard known as FISMA (Federal Information Security Management Act). While auditing does lead to improved best practice coverage, there are often problems. For example, many audits are done poorly, which results in confusion and improper management decisions. In addition, with all the emphasis on numeric ratings, many agencies focus more on their score than on good security practice.

A good audit score is important but should not replace good security practices.

Today, organizations charged with protecting national infrastructure are subjected to several types of security audits. Streamlining these standards would certainly be a good idea, but some additional items for consideration include improving the types of common training provided to security administrators, as well as including past practice in infrastructure protection in common audit standards. The most obvious practical consideration for national infrastructure, however, would be national-level agreement on which standard or standards would be used to determine competence to protect national assets. While this is a straightforward concept, it could be tough to obtain wide concurrence among

all national participants. A related issue involves commonality in national infrastructure operational configurations; this reduces the chances that a rogue configuration installed for malicious purposes, perhaps by compromised insiders.

A national standard of competence for protecting our assets is needed.

Depth

The principle of *depth* involves the use of multiple security layers of protection for national infrastructure assets. These layers protect assets from both internal and external attacks via the familiar "defense in depth" approach; that is, multiple layers reduce the risk of attack by increasing the chances that at least one layer will be effective. This should appear to be a somewhat sketchy situation, however, from the perspective of traditional engineering. Civil engineers, for example, would never be comfortable designing a structure with multiple flawed supports in the hopes that one of them will hold the load. Unfortunately, cyber security experts have no choice but to rely on this flawed notion, perhaps highlighting the relative immaturity of security as an engineering discipline.

One hint as to why depth is such an important requirement is that national infrastructure components are currently controlled by software, and everyone knows that the current state of software engineering is abysmal. Compared to other types of engineering, software stands out as the only one that accepts the creation of knowingly flawed products as acceptable. The result is that all nontrivial software has exploitable vulnerabilities, so the idea that one should create multiple layers of security defense is unavoidable. It is worth mentioning that the degree of diversity in these layers will also have a direct impact on their effectiveness (see Figure 1.8).

Software engineering standards do not contain the same level of quality as civil and other engineering standards.

To maximize the usefulness of defense layers in national infrastructure, it is recommended that a combination of functional and procedural controls be included. For example, a common first layer of defense is to install an access control mechanism for the

Figure 1.8 National infrastructure security through defense in depth.

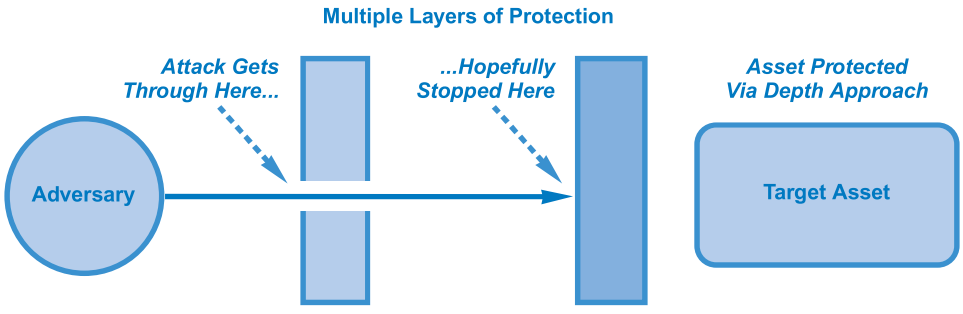

admission of devices to the local area network. This could involve router controls in a small network or firewall access rules in an enterprise. In either case, this first line of defense is clearly functional. As such, a good choice for a second layer of defense might involve something procedural, such as the deployment of scanning to determine if inappropriate devices have gotten through the first layer. Such diversity will increase the chances that the cause of failure in one layer is unlikely to cause a similar failure in another layer.

A great complication in national infrastructure protection is that many layers of defense assume the existence of a defined network perimeter. For example, the presence of many flaws in enterprise security found by auditors is mitigated by the recognition that intruders would have to penetrate the enterprise perimeter to exploit these weaknesses. Unfortunately, for most national assets, finding a perimeter is no longer possible. The assets of a country, for example, are almost impossible to define within some geographic or political boundary, much less a network one. Security managers must therefore be creative in identifying controls that will be meaningful for complex assets whose properties are not always evident. The risk of getting this wrong is that in providing multiple layers of defense, one might misapply the protections and leave some portion of the asset base with no layers in place.

Discretion

The principle of *discretion* involves individuals and groups making good decisions to obscure sensitive information about national infrastructure. This is done by combining formal mandatory information protection programs with informal discretionary behavior. Formal mandatory programs have been in place for many years in the U.S. federal government, where documents are associated with classifications, and policy enforcement is based on clearances granted to individuals. In the most intense environments, such as top-secret compartments in the intelligence community, violations of access policies could be interpreted as espionage, with all of the associated criminal implications. For this reason, prominent breaches of highly classified government information are not common.

Naturally, top-secret information within the intelligence community is at great risk for attack or infiltration.

In commercial settings, formal information protection programs are gaining wider acceptance because of the increased need to protect personally identifiable information (PII) such as credit card numbers. Employees of companies around the world are starting to understand the importance of obscuring certain aspects of corporate activity, and this is healthy for national

infrastructure protection. In fact, programs of discretion for national infrastructure protection will require a combination of corporate and government security policy enforcement, perhaps with custom-designed information markings for national assets. The resultant discretionary policy serves as a layer of protection to prevent national infrastructure-related information from reaching individuals who have no need to know such information.

A barrier in our recommended application of discretion is the maligned notion of "security through obscurity." Security experts, especially cryptographers, have long complained that obscurity is an unacceptable protection approach. They correctly reference the problems of trying to secure a system by hiding its underlying detail. Inevitably, an adversary discovers the hidden design secrets and the security protection is lost. For this reason, conventional computer security correctly dictates an open approach to software, design, and algorithms. An advantage of this open approach is the social review that comes with widespread advertisement; for example, the likelihood is low of software ever being correct without a significant amount of intense review by experts. So, the general computer security argument against "security through obscurity" is largely valid in most cases.

"Security through obscurity" may actually leave assets more vulnerable to attack than an open approach would.

Nevertheless, any manager charged with the protection of nontrivial, large-scale infrastructure will tell you that discretion and, yes, obscurity are indispensable components in a protection program. Obscuring details around technology used, software deployed, systems purchased, and configurations managed will help to avoid or at least slow down certain types of attacks. Hackers often claim that by discovering this type of information about a company and then advertising the weaknesses they are actually doing the local security team a favor. They suggest that such advertisement is required to motivate a security team toward a solution, but this is actually nonsense. Programs around proper discretion and obscurity for infrastructure information are indispensable and must be coordinated at the national level.

Collection

The principle of *collection* involves automated gathering of system-related information about national infrastructure to enable security analysis. Such collection is usually done in real time and involves probes or hooks in applications, system software, network elements, or hardware devices that gather information of interest. The use of audit trails in small-scale computer security is an example of a long-standing collection practice that

introduces very little controversy among experts as to its utility. Security devices such as firewalls produce log files, and systems purported to have some degree of security usefulness will also generate an audit trail output. The practice is so common that a new type of product, called a *security information management system* (SIMS), has been developed to process all this data.

The primary operational challenge in setting up the right type of collection process for computers and networks has been twofold: First, decisions must be made about what types of information are to be collected. If this decision is made correctly, then the information collected should correspond to exactly the type of data required for security analysis, and nothing else. Second, decisions must be made about how much information is actually collected. This might involve the use of existing system functions, such as enabling the automatic generation of statistics on a router; or it could involve the introduction of some new type of function that deliberately gathers the desired information. Once these considerations are handled, appropriate mechanisms for collecting data from national infrastructure can be embedded into the security architecture (see Figure 1.9).

What and how much data to collect is an operational challenge.

The technical and operational challenges associated with the collection of logs and audit trails are heightened in the protection of national assets. Because national infrastructure is so complex, determining what information should be collected turns out to be a difficult exercise. In particular, the potential arises with large-scale collection to intrude on the privacy of individuals and groups within a nation. As such, any initiative to protect infrastructure through the collection of data must include at least some measure of privacy policy determination. Similarly, the volumes of data collected from large infrastructure can exceed practical limits. Telecommunications collection systems designed to protect the integrity of a service provider backbone, for example, can easily generate many terabytes of data in hours of processing.

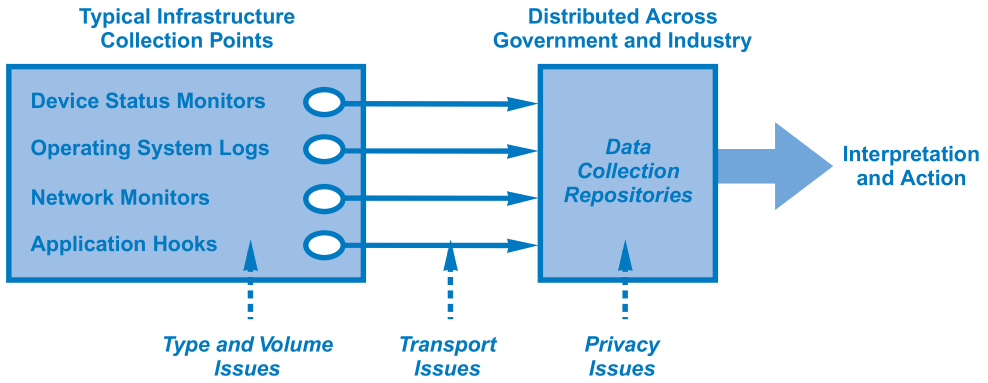

Figure 1.9 Collecting national infrastructure-related security information.

In both cases, technical and operational expertise must be applied to ensure that the appropriate data is collected in the proper amounts. The good news is that virtually all security protection algorithms require no deep, probing information of the type that might generate privacy or volumetric issues. The challenge arises instead when collection is done without proper advance analysis which often results in the collection of more data than is needed. This can easily lead to privacy problems in some national collection repositories, so planning is particularly necessary. In any event, a national strategy of data collection is required, with the usual sorts of legal and policy guidance on who collects what and under which circumstances. As we suggested above, this exercise must be guided by the requirements for security analysis—and nothing else.

Only collect as much data as is necessary for security purposes.

Correlation

The principle of *correlation* involves a specific type of analysis that can be performed on factors related to national infrastructure protection. The goal of correlation is to identify whether security-related indicators might emerge from the analysis. For example, if some national computing asset begins operating in a sluggish manner, then other factors would be examined for a possible correlative relationship. One could imagine the local and wide area networks being analyzed for traffic that might be of an attack nature. In addition, similar computing assets might be examined to determine if they are experiencing a similar functional problem. Also, all software and services embedded in the national asset might be analyzed for known vulnerabilities. In each case, the purpose of the correlation is to combine and compare factors to help explain a given security issue. This type of comparison-oriented analysis is indispensable for national infrastructure because of its complexity.

Monitoring and analyzing networks and data collection may reveal a hidden or emerging security threat.

Interestingly, almost every major national infrastructure protection initiative attempted to date has included a fusion center for real-time correlation of data. A fusion center is a physical security operations center with means for collecting and analyzing multiple sources of ingress data. It is not uncommon for such a center to include massive display screens with colorful, visualized representations, nor is it uncommon to find such centers in the military with teams of enlisted people performing the manual chores. This is an important point, because, while such automated fusion is certainly promising, best practice in correlation for national infrastructure protection must include the requirement that human judgment be included in the analysis. Thus, regardless of whether resources are centralized into one physical

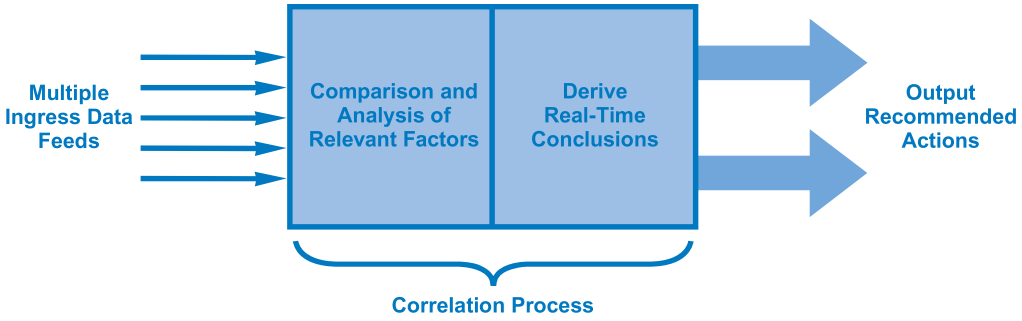

Figure 1.10 National infrastructure high-level correlation approach.

location, the reality is that human beings will need to be included in the processing (see Figure 1.10).

In practice, fusion centers and the associated processes and correlation algorithms have been tough to implement, even in small-scale environments. Botnets, for example, involve the use of source systems that are selected almost arbitrarily. As such, the use of correlation to determine where and why the attack is occurring has been useless. In fact, correlating geographic information with the sources of botnet activity has even led to many false conclusions about who is attacking whom. Countless hours have been spent by security teams poring through botnet information trying to determine the source, and the best one can hope for might be information about controllers or software drops. In the end, current correlation approaches fall short.

What is needed to improve present correlation capabilities for national infrastructure protection involves multiple steps.

Three Steps to Improve Current Correlation Capabilities

1. The actual computer science around correlation algorithms needs to be better investigated. Little attention has been placed in academic computer science and applied mathematics departments to multifactor correlation of real-time security data. This could be changed with appropriate funding and grant emphasis from the government.
2. The ability to identify reliable data feeds needs to be greatly improved. Too much attention has been placed on *ad hoc* collection of volunteered feeds, and this complicates the ability for analysis to perform meaningful correlation.
3. The design and operation of a national-level fusion center must be given serious consideration. Some means must be identified for putting aside political and funding problems in order to accomplish this important objective.

Awareness

The principle of *awareness* involves an organization understanding the differences, in real time and at all times, between observed and normal status in national infrastructure. This status can include risks, vulnerabilities, and behavior in the target infrastructure. *Behavior* refers here to the mix of user activity, system processing, network traffic, and computing volumes in the software, computers, and systems that comprise infrastructure. The implication is that the organization can somehow characterize a given situation as being either normal or abnormal. Furthermore, the organization must have the ability to detect and measure differences between these two behavioral states. Correlation analysis is usually inherent in such determinations, but the real challenge is less the algorithms and more the processes that must be in place to ensure situational awareness every hour of every day. For example, if a new vulnerability arises that has impact on the local infrastructure, then this knowledge must be obtained and factored into management decisions immediately.

Awareness builds on collection and correlation, but is not limited to those areas alone.

Managers of national infrastructure generally do not have to be convinced that situational awareness is important. The big issue instead is how to achieve this goal. In practice, real-time awareness requires attentiveness and vigilance rarely found in normal computer security. Data must first be collected and enabled to flow into a fusion center at all times so correlation can take place. The results of the correlation must be used to establish a profiled baseline of behavior so differences can be measured. This sounds easier than it is, because so many odd situations have the ability to mimic normal behavior (when it is really a problem) or a problem (when it really is nothing). Nevertheless, national infrastructure protection demands that managers of assets create a locally relevant means for being able to comment accurately on the state of security at all times. This allows for proper management decisions about security (see Figure 1.11).

Interestingly, situational awareness has not been considered a major component of the computer security equation to date. The

Figure 1.11 Real-time situation awareness process flow.

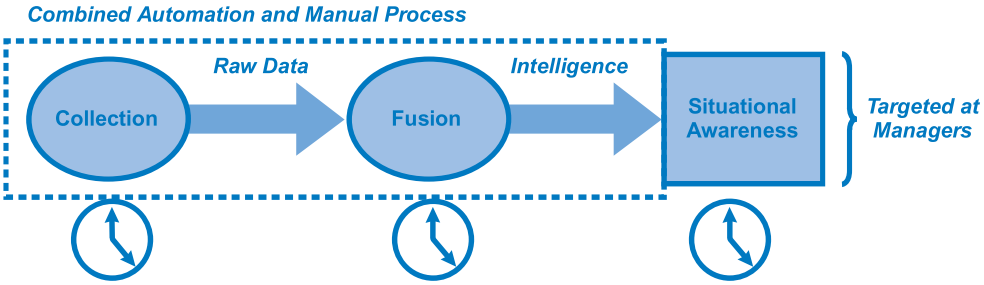

concept plays no substantive role in small-scale security, such as in a home network, because when the computing base to be protected is simple enough, characterizing real-time situational status is just not necessary. Similarly, when a security manager puts in place security controls for a small enterprise, situational awareness is not the highest priority. Generally, the closest one might expect to some degree of real-time awareness for a small system might be an occasional review of system log files. So, the transition from small-scale to large-scale infrastructure protection does require a new attentiveness to situational awareness that is not well developed. It is also worth noting that the general notion of "user awareness" of security is also not the principle specified here. While it is helpful for end users to have knowledge of security, any professionally designed program of national infrastructure security must presume that a high percentage of end users will *always* make the wrong sorts of security decisions if allowed. The implication is that national infrastructure protection must never rely on the decision-making of end users through programs of awareness.

Large-scale infrastructure protection requires a higher level of awareness than most groups currently employ.

A further advance that is necessary for situational awareness involves enhancements in approaches to security metrics reporting. Where the non-cyber national intelligence community has done a great job developing means for delivering daily intelligence briefs to senior government officials, the cyber security community has rarely considered this approach. The reality is that, for situation awareness to become a structural component of national infrastructure protection, valid metrics must be developed to accurately portray status, and these must be codified into a suitable type of regular intelligence report that senior officials can use to determine security status. It would not be unreasonable to expect this cyber security intelligence to flow from a central point such as a fusion center, but in general this is not a requirement.

Response

The principle of *response* involves assurance that processes are in place to react to any security-related indicator that becomes available. These indicators should flow into the response process primarily from the situational awareness layer. National infrastructure response should emphasize indicators rather than incidents. In most current computer security applications, the response team waits for serious problems to occur, usually including complaints from users, applications running poorly, and networks operating in a sluggish manner. Once this occurs, the response team springs into action, even though by

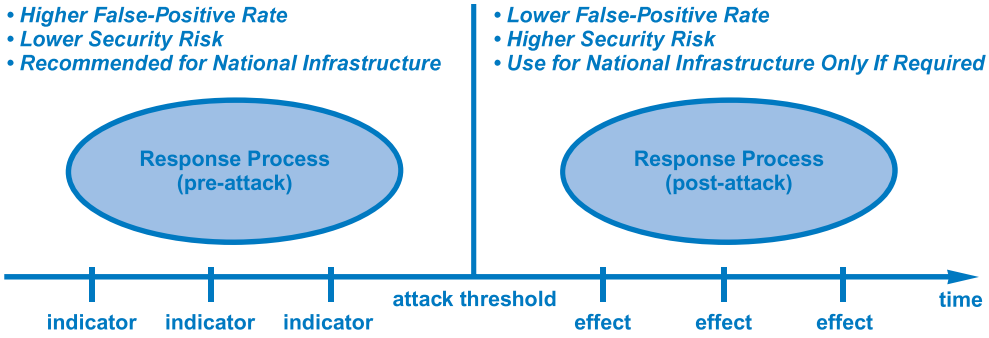

Figure 1.12 National infrastructure security response approach.

this time the security game has already been lost. For essential national infrastructure services, the idea of waiting for the service to degrade before responding does not make logical sense.

An additional response-related change for national infrastructure protection is that the maligned concept of "false positive" must be reconsidered. In current small-scale environments, a major goal of the computer security team is to minimize the number of response cases that are initiated only to find that nothing was wrong after all. This is an easy goal to reach by simply waiting for disasters to be confirmed beyond a shadow of a doubt before response is initiated. For national infrastructure, however, this is obviously unacceptable. Instead, response must follow indicators, and the concept of minimizing false positives must not be part of the approach. The only quantitative metric that must be minimized in national-level response is risk (see Figure 1.12).

A challenge that must be considered in establishing response functions for national asset protection is that relevant indicators often arise long before any harmful effects are seen. This suggests that infrastructure protecting must have accurate situational awareness that considers much more than just visible impacts such as users having trouble, networks being down, or services being unavailable. Instead, often subtle indicators must be analyzed carefully, which is where the challenges arise with false positives. When response teams agree to consider such indicators, it becomes more likely that such indicators are benign. A great secret to proper incident response for national infrastructure is that higher false positive rates might actually be a good sign.

A higher rate of false positives must be tolerated for national infrastructure protection.

It is worth noting that the principles of collection, correlation, awareness, and response are all consistent with the implementation of a national fusion center. Clearly, response activities are often dependent on a real-time, ubiquitous operations center to coordinate activities, contact key individuals, collect data as it

becomes available, and document progress in the response activities. As such, it should not be unexpected that national-level response for cyber security should include some sort of centralized national center. The creation of such a facility should be the centerpiece of any national infrastructure protection program and should involve the active participation of all organizations with responsibility for national services.

Implementing the Principles Nationally

To effectively apply this full set of security principles in practice for national infrastructure protection, several practical implementation considerations emerge:

- *Commissions and groups*—Numerous commissions and groups have been created over the years with the purpose of national infrastructure protection. Most have had some minor positive impact on infrastructure security, but none has had sufficient impact to reduce present national risk to acceptable levels. An observation here is that many of these commissions and groups have become the *end* rather than the *means* toward a cyber security solution. When this occurs, their likelihood of success diminishes considerably. Future commissions and groups should take this into consideration.
- *Information sharing*—Too much attention is placed on information sharing between government and industry, perhaps because information sharing would seem on the surface to carry much benefit to both parties. The advice here is that a comprehensive information sharing program is not easy to implement simply because organizations prefer to maintain a low profile when fighting a vulnerability or attack. In addition, the presumption that some organization—government or commercial—might have some nugget of information that could solve a cyber attack or reduce risk is not generally consistent with practice. Thus, the motivation for a commercial entity to share vulnerability or incident-related information with the government is low; very little value generally comes from such sharing.
- *International cooperation*—National initiatives focused on creating government cyber security legislation must acknowledge that the Internet is global, as are the shared services such as the domain name system (DNS) that all national and global assets are so dependent upon. Thus, any program of national infrastructure protection must include provisions for international cooperation, and such cooperation implies agreements between participants that will be followed as long as everyone perceives benefit.

- *Technical and operational costs*—To implement the principles described above, considerable technical and operational costs will need to be covered across government and commercial environments. While it is tempting to presume that the purveyors of national infrastructure can simply absorb these costs into normal business budgets, this has not been the case in the past. Instead, the emphasis should be on rewards and incentives for organizations that make the decision to implement these principles. This point is critical because it suggests that the best possible use of government funds might be as straightforward as helping to directly fund initiatives that will help to secure national assets.

The bulk of our discussion in the ensuing chapters is technical in nature; that is, programmatic and political issues are conveniently ignored. This does not diminish their importance, but rather is driven by our decision to separate our concerns and focus in this book on the details of "what" must be done, rather than "how."

Finally, let's look at how the ever-changing policy of the United States helps prevent or minimize disruptions to the critical national infrastructure. The implementation of the policy is crucial in order to protect the public, the economy, government services, and the national security of the United States.

Protecting the Critical National Infrastructure Against Cyber Attacks

Information technology has grown to provide both government and the private sector with an efficient and timely means of delivering essential services around the world. As a result, these critical systems remain at risk from potential attacks via the Internet. It is the policy of the United States to prevent or minimize disruptions to the critical information infrastructure in order to protect the public, the economy, government services, and the national security of the United States.

The federal government is continually increasing capabilities to address cyber risk associated with critical networks and information systems. On January 8, 2008, President Bush approved the National Security Presidential Directive 54/Homeland Security Presidential Directive 23, which formalized a series of continuous efforts designed to further safeguard federal government systems and reduce potential vulnerabilities, protect against intrusion attempts, and better anticipate future threats.

While efforts to protect the federal network systems from cyber attacks remain a collaborative, government-wide effort, the

Department of Homeland Security (DHS) has the lead responsibility for ensuring the security, resiliency, and reliability of the nation's information technology (IT) and communications infrastructure (see "An Agenda for Action in Preventing Cyber Attacks Methods" below).

An Agenda for Action in Preventing Cyber Attacks Methods

When completing the Preventing Cyber Attacks Methods checklist, the DHS specialist should adhere to the provisional list of actions for some of the principal cyber attack prevention methods. The order is not significant; however, these are the activities for which the research would want to provide a detailed description of procedures, review, and assessment for ease of use and admissibility. Current measures that must be adhered to in order to prevent future attacks and intrusion attempts, include (check all tasks completed):

1. Hiring additional personnel to support the U.S. Computer Emergency Readiness Team (US-CERT), DHS' 24 × 7 watch and warning center for the federal government's Internet infrastructure. US-CERT, a public–private partnership, operates round the clock to help government and industry analyze and respond to cyber threats and vulnerabilities.
2. Expanding the Einstein Cyber Shield to all federal departments and agencies. This will provide government officials with an early warning system to gain better situational awareness, earlier identification of malicious activity, and a more comprehensive network defense. The current version of the program is widely seen as providing meager protection against attack, but a new version being built will be more robust—largely because it is rooted in NSA technology. The program is designed to look for indicators of cyber attacks by digging into all Internet communications, including the contents of e-mails, according to a declassified summary.
3. Consolidating the number of external connections including Internet points of presence for the federal government Internet infrastructure (FGII), as part of the Office of Management and Budget's (OMB's) Trusted Internet Connections Initiative (TICI). TICI will more efficiently manage and implement security measures to help bring more comprehensive protection across the federal .gov domains.
4. Creating a National Cyber Security Center (NCSC) to further progress in addressing cyber threats and increasing cyber security efforts. The NCSC will bring together federal cyber security organizations by virtually connecting and, in some cases, physically collocating personnel and resources to gain a clearer understanding of the overall cyber security picture of federal networks.
5. Expanding the National Cyber Investigative Joint Task Force (NCIJTF) to include representation from the U.S. Secret Service

and several other federal agencies. This existing cyber investigation coordination organization overseen by the Federal Bureau of Investigation (FBI) will serve as a multiagency national focal point for coordinating, integrating, and sharing pertinent information related to cyber threat investigations.

6. Reducing the potential for adversaries to manipulate IT and communications products before they are imported into the United States. In other words, the DHS specialist must work toward a stronger supply chain defense. To address this challenge, the federal government is exploring protections into the federal acquisition process and developing a multifaceted strategy to reduce risk at the most appropriate stage of the IT and communications product lifecycle.
7. Facilitating coordination and information sharing between the federal government and private sector to reduce cyber risk, disseminate threat information, share best practices, and apply appropriate protective actions as outlined within the National Infrastructure Protection Plan (NIPP) framework. For example, DHS created the Control Systems Vulnerability Assessment Tool (CSVAT) to help all critical infrastructure sectors assess certain policies, plans, and procedures currently in place to reduce cyber vulnerabilities and leverage recognized standards.
8. Leading the nation's largest cyber security exercise, known as Cyber Storm III, in the fall of 2010, bringing together participants from federal, state, and local governments; the private sector; and the international community in order to examine and strengthen the nation's cyber security preparedness and response capabilities in response to a simulated cyber attack across several critical sectors of this nation's economy. Cyber Storm III was built upon the success of previous exercises; however, enhancements in the nation's cyber security capabilities, an ever-evolving cyber threat landscape and the increased emphasis and extent of public–private collaboration and cooperation made Cyber Storm III unique. Cyber Storm III was the primary vehicle to exercise the newly developed National Cyber Incident Response Plan (NCIRP)—a blueprint for cyber security incident response—to examine the roles, responsibilities, authorities, and other key elements of the nation's cyber incident response and management capabilities and use those findings to refine the plan. Cyber Storm III (and the upcoming Cyber Storm IV in 2012) and other exercises help ensure that public and private sectors are prepared for an effective response to attacks against this nation's critical systems and networks.
9. Partnering with academia and industry to expand cyber education for all U.S. government employees, particularly those who specialize in IT, and enhance worksite development and recruitment strategies to ensure a knowledgeable workforce capable of dealing with the evolving nature of cyber threats.
10. Increasing funding for IT security through the president's FY 2012 budget for protection efforts against cyber attacks efforts across the federal government and the private sector.

Summary

This chapter discussed how pervasive and sustained cyber attacks continue to pose a potentially devastating threat to the systems and operations of the critical national infrastructure of the United States. According to recent testimony by the Director of National Intelligence, "there has been a dramatic increase in malicious cyber activity targeting U.S. computers and networks." In addition, recent reports of cyber attacks and incidents affecting critical infrastructures illustrate the potential impact of such events on national and economic security. The nation's ever-increasing dependence on information systems to carry out essential day-to-day operations makes it vulnerable to an array of cyber-based risks. Thus, it is increasingly important that federal and nonfederal entities carry out concerted efforts to safeguard their systems and the information they contain by looking at:

- Cyber threats to cyber-reliant critical national infrastructures.
- The continuing challenges facing federal agencies in protecting the nation's cyber-reliant critical national infrastructure.

Cyber-based threats to the critical national infrastructure are evolving and growing. These threats can come from a variety of sources, including criminals and foreign nations, as well as hackers and disgruntled employees. These potential cyber attackers have a variety of techniques at their disposal that can vastly expand the reach and impact of their actions. In addition, the interconnectivity between information systems, the Internet, and other infrastructure presents increasing opportunities for such cyber attacks. Consistent with this, reports of security incidents from federal agencies are on the rise according to the Government Accounting Office (GAO), increasing over 760% over the past 6 years. In addition, reports of cyber attacks and information security incidents, affecting federal systems and systems supporting the critical national infrastructure, illustrate the serious impact such incidents can have on national and economic security, including the loss of classified information and intellectual property worth billions of dollars. The Obama administration and executive branch agencies continue to act to better protect the cyber-reliant critical national infrastructures, improve the security of federal systems, and strengthen the nation's cyber security posture, but they are still falling short of their goals. In other words, they have not yet fully implemented key actions that are intended to address threats and improve the current U.S. approach to cyber security, such as:

- Implementing near- and midterm actions recommended by the cyber security policy review directed by the president.

- Updating the national strategy for securing the information and communications infrastructure.
- Developing a comprehensive national strategy for addressing global cyber security and governance.
- Creating a prioritized national and federal research and development agenda for improving cyber security.

Federal systems continue to be afflicted by persistent information security control weaknesses. For example, as part of its audit of the fiscal year 2010 financial statements for the U.S. government, the GAO determined that serious and widespread information security control deficiencies were a government-wide material weakness. Over the past several years, GAO and agency inspectors general have made thousands of recommendations to agencies for actions necessary to resolve prior significant control deficiencies and information security program shortfalls. The White House, the Office of Management and Budget, and selected federal agencies have undertaken additional government-wide initiatives intended to enhance information security at federal agencies. However, these initiatives face challenges, such as better defining agency roles and responsibilities, establishing measures of effectiveness, and the requirement of sustained attention, which government agencies have begun to provide. As such, the GAO continues to identify the federal government's information systems and the nation's cyber critical national infrastructure as a government-wide high-risk area.

Finally, let's move on to the real interactive part of this chapter: review questions/exercises, hands-on projects, case projects, and optional team case project. The answers and/or solutions by chapter can be found online at http://www.elsevierdirect.com/companion.jsp?ISBN=9780123918550.

CHAPTER REVIEW QUESTIONS/EXERCISES

True/False

1. True or False? National infrastructure refers to the complex, underlying delivery and support systems for all large-scale services considered absolutely essential to a nation.
2. True or False? Vulnerabilities are more difficult to associate with any taxonomy.
3. True or False? Perhaps the most insidious type of attack that exists today is the botnet.
4. True or False? The principle of deception involves the deliberate introduction of misleading functionality or misinformation into national infrastructure for the purpose of tricking an adversary.

5. True or False? The principle of separation involves enforcement of access policy restrictions on the users and resources in a computing environment.

Multiple Choice

1. The best one can do for a comprehensive view of the vulnerabilities associated with national infrastructure is to address their relative exploitation points. This can be done with an abstract national infrastructure cyber security model that includes three types of malicious adversaries, except which two:
   A. External adversary
   B. Remote adversary
   C. Internal adversary
   D. System adversary
   E. Supplier adversary
2. By using the abstract national infrastructure cyber security model, three exploitation points emerge for national infrastructure, except which two:
   A. Defined methodology
   B. Remote access
   C. Breach of contract
   D. System administration and normal usage
   E. Supply chain
3. The selection and use of technology and systems that are intentionally different in substantive ways is called the principle of:
   A. Consistency
   B. Depth
   C. Discretion
   D. Collection
   E. Diversity
4. The automated gathering of system-related information about national infrastructure to enable security analysis is called the principle of:
   A. Correlation
   B. Awareness
   C. Response
   D. Collection
   E. Recovery
5. To effectively apply the full set of security principles in practice for national infrastructure protection, several practical implementation considerations emerge, except which one:
   A. Commissions and groups
   B. Information sharing

C. International cooperation
D. Technical and operational costs
E. Current correlation capabilities

Exercise

Problem

A disgruntled former hospital employee with exceptional computer skills hacks into the hospital network from their home computer and plants a very aggressive computer virus into a Computer-Aided Facility Management (CAFM) system. The computer virus activates at 1:00 a.m., shutting down the Hospital Ventilation Air Conditioning (HVAC) system, security system, building automation, and patient medical monitoring system. Please explain how the hospital's cyber security team (CST) went about resolving the problem.

Hands-On Projects

Project

Trojan Horse e-mails sent from an intruder were targeted at specific organizations and people. The Trojan Horse e-mails, when opened, compromised the system and enabled the cyber attackers to infiltrate the internal networked systems. The cyber attackers then searched the systems and network for data files and exfiltrated information through the encrypted channels. On opening the document, a real document would display, while hidden activities are executed in the background. The possibility of applications crashing is extremely high. The following is an example:

- A reverse shell leveraging port 443 (secure sockets layer [SSL]) downloaded a command and control tools from a dynamic domain. Traffic was not SSL encrypted, but was obfuscated. Obfuscated code is source or machine code that has been made difficult to understand. Programmers may deliberately obfuscate code to conceal its purpose (security through obscurity) or its logic to prevent tampering or deter reverse engineering, or as a puzzle or recreational challenge for someone reading the source code.
- The intruder then gained access and conducted network scanning, data collection, and data exfiltration (military jargon for the removal of personnel or units from areas under enemy control by stealth, deception, surprise, or clandestine means, the opposite of infiltration).

So, how would your cyber security team go about identifying the intruder, the collection of tools used by the intruder, and recovering from the attack?

Case Projects

Problem

Let's look at a real-world scenario and how the Department of Homeland Security (DHS) plays into it. In the scenario, the United States will be hit by a large-scale, coordinated cyber attack organized by China. These attacks debilitate the functioning of government agencies, parts of the critical infrastructure, and commercial ventures. The IT infrastructure of several agencies are paralyzed, the electric grid in most of the country is shut down, telephone traffic is seriously limited and satellite communications are down (limiting the Department of Defense's [DOD's] ability to communicate with commands overseas). International commerce and financial institutions are also severely hit. Please explain how DHS should handle this situation.

Optional Team Case Project

Problem

A cadre of intruders leveraged their collective capabilities to mount a simulated coordinated cyber attack on a global scale. Although primary motives differed among the entities, a sophisticated network of relationships enabled the intruder to degrade Internet connectivity, disrupt industrial functions, and ultimately erode confidence in everyday communications. The intruder cultivated relationships with unaffiliated opportunistic intruders. Due to their critical nature and perceived vulnerabilities, the intruders specifically targeted several critical infrastructure sectors, along with state and federal agencies, the media, and foreign nations. Please identify the findings that were observed by the participants and observer/controllers through the implementation of this project.