Organizational Security Concerns: Analysis and Recommendations

Andrew Bruce, CISSP, PMP, FITSP-D
CTO, RiVidium Corporation
andy.bruce@rividium.com
21 August 2010

Executive Summary

Computer security breakdowns in the news remind us how companies are vulnerable to many types of failures – logical, physical, and administrative. In this paper, we review a number of these news stories to see how they could affect our own organization's security posture. To some, security simply means guarding the computers. While the importance of physical computer security cannot be overstated, security far "transcends technology."[1] In this paper, we posit that specific security breaches are best stated as failures in the organization's high-level Security Policy (or lack thereof). As eminent security analyst Mich Kabay points out, the security policy "govern[s] how an institution's information is to be protected against breaches of security."[2] A properly implemented security policy provides both the formal effort to demonstrate due diligence to our customers (example: use of Bell-LaPadula[3] as a security model) as well as creating the security-aware employee mindset for preventing security problems in the first place. A security policy allows us to deliver on the basic security tenets of *confidentiality*, *integrity*, and *availability* (otherwise known as the CIA Triad[4]).

Audience and Topics

We address this paper to our fellow corporate officers, corporate information managers interested in reviewing and evaluating possible security solutions to common infrastructure issues, general security practitioners curious to see how we have addressed these common issues within a real-world environment, and students interested in learning about the fascinating and complex world of security management. To keep the paper well-focused, we concentrate on just the following high-level problems:

1. **Physical access** – Think of this as "Security begins at home." We discuss physical access to the facility and how this can be improved.

2. **Logical access** – How our corporate users and business partners access our systems, both to execute functions (such as time sheet submission) as well as to retrieve the raw data necessary for making informed decisions (such as Web service access).

3. **External attacks** – The ways that others may attack our computing resources, and the steps we have taken to mitigate these vulnerabilities.

4. **Internal protection failures** – The problems that can occur from the accidental or malicious disclosure of information by authorized users (the proverbial "inside job").

The above list is far from exhaustive — indeed, an exhaustive list would consume several thousand pages and still be incomplete. However, it serves to highlight a specific set of common scenarios that every organization must address, allowing us to use inductive reasoning[5] to see how specific security failures ultimately relate back to overall security policy problems we can address.

Table of Contents

Executive Summary 2

Audience and Topics 2

1.0 Physical Access 1

1.1 In the News 1

1.2 Problems and Mitigation 1

1.3 Recommendations 4

2.0 Logical Access 5

2.1 In the News 5

2.2 Problems and Mitigation 6

2.3 Recommendations 10

3.0 External Attacks 11

3.1 In the News 12

3.2 Problems and Mitigation 12

3.3 Recommendations 17

4.0 Internal Protection Failures 17

4.1 In the News 17

4.2 Problems and Mitigation 18

4.3 Recommendations 22

Summary 23

Appendix A: All Recommendations 24

Appendix B: Recommendations by Priority with Estimates 26

1.0 Physical Access

Protecting the physical computer provides the most basic and fundamental level of security and information assurance for an organization, and must be the primary concern of the overall security policy. The Parkerian Hexad[6] extends the classic CIA Triad and identifies *possession* as a critical underpinning for all other security elements; without possession, the organization loses both the *utility* and *availability* of the missing equipment. Additionally, the *confidentiality* and *integrity* of any data stored on the missing equipment are at great risk. Finally, consider what can happen if data from a stolen device (like a laptop or cell phone) contains readable information about our customers and partners. Aside from the regulatory non-compliance nightmares this would produce, the stolen data could be modified and published as authoritative from a rogue Web site. In that case, we lose even the final Parkerian principle of data *authenticity*.

1.1 In the News…

School's Out[7] – On June 15, 2010, a 15-year-old student was arrested for stealing $35,000 worth of computers from his local high school in Milville, NJ over a period of several months. The boy had repeatedly broken into the school; on the last occasion, he was caught simply because a police officer happened to be observing him as he was walking away from the school with a large cart filled with personal computers!

Hospital horror[8] – On May 28, 2010, Cincinnati Children's Hospital reported that a laptop containing 61,000 patient records was stolen. The records were password-protected, however, they *were not encrypted*. Data on the machine contained Personally Identifiable Information (PII), including medical records and patient histories. The administration rushed to assure the media that Social Security numbers, credit cards, and telephone numbers were not in the lost data (although with an address or simply a name, telephone numbers can easily be found as the spokeo.com[9] Web site makes disturbingly apparent).

Trouble in Tennessee[10] – On June 9, 2010, Tennessee officials reported that more than 10,000 names and Social Security numbers were on a stolen laptop. While the officials have opened a call center and have pledged to provide identity theft protection for affected individuals, there was no word on whether the data was encrypted or even password-protected — one can assume that neither was true.

1.2 Problems and Mitigation

1.2.1: Problems in the Physical Domain

The common thread to all of our news stories relates to the organization's failure to implement the basic security policy requirement of physical plant and human safety. In this age of small, portable, and at-risk equipment, physical loss becomes our smallest area of concern. We can easily replace a stolen PC or laptop.[11] The real cost is the loss of control over information; to use the Parkerian term[12] this equates to the loss of *control or possession* with the corresponding access to data contained on the stolen equipment.

Also, consider a troubling side note to the article on the Millville school theft: the crimes occurred over a multi-month period and the equipment was obviously stolen for resale — *what happened to the data stored on those machines?* Presumably, the criminals receiving the stolen goods could access this data, so we can add the possibility of privacy leakage or privileged information leakage to this event. However, the officials in charge of the investigation make no mention of this possibility. From a security policy point of view, we see that how physical theft easily leads to the related problem of managing the possible compromise of stored data – and security policies regarding data management (such as Personally Identifiable Information or PII) can be *required by law.* As many companies have discovered, once data has been compromised then the effects can be ongoing and pervasive; data loss truly qualifies as the "gift that keeps on giving."[13]

1.2.2 Standards to the Rescue

In 1988 Ronald Reagan made famous the phrase "the ten most dangerous words in the English language are, 'Hi, I'm from the Government and I'm here to Help'."[14] We are not here to argue the political implications of that statement, but we can say that, insofar as it pertains to physical security, that appears not to be the case. In fact, the federal government has invested significant time and money into developing a complex and exhaustive set of standards that can help to ensure the safety of computer systems.

1.2.2.1 FIPS 200[15]

The Federal Information Processing Standards (FIPS) publications serve as an excellent starting point for physical security, especially FIPS 200 (Minimum Security Requirements for Federal Information and Information Systems). Section 3 within this publication (Minimum Security Requirements) provides a set of seventeen high-level areas targeting physical, logical (technical), and administrative controls. The physical control areas identified by this standard include:

Media Protection (MP) – Organizations must protect information, control access to that information, and must ensure that the physical media used for information storage is properly sanitized prior to reuse. From a physical viewpoint, we are concerned primarily with the protection of the physical data storage media as well as the proper disposal of that media.

Physical and Environmental Protection (PE) – Organizations must control physical access to computer processing systems, ensure that the computer systems are protected from natural threats, and provide for a safe and effective physical environment (such as proper Heating, Ventilation, and Air Conditioning or HVAC).

System and Information Integrity (SI) – Organizations must protect their systems by applying software patches, must report system security breaches, and must monitor their systems. While these ostensibly sound like logical (technical) controls, that's only half of the story. For these logical controls to be implemented, the Operations group must define and implement necessary patching schedules as well as build and maintain a proper staging environment where patches and updates can be verified prior to production release. A better way to think of SI: providing the *physical infrastructure* necessary to allow the logical controls to be performed.

1.2.2.2 NIST SP 800-53[16]

The FIPS 200 standards guide provides a high-level description of the areas that must be protected as well as what should be protected within those areas, but it provides little detail to assist the implementer. However, the National Institute of Standards and Technology (NIST) created Special Publication (SP) 800-53: "Information Security," which provides more specific guidance not only on the standards to be met but also on the controls to be implemented. In fact, an organization could do much worse than simply to start at the beginning of Appendix F (Security Control Catalog) and analyze / implement each of the controls as necessary for their specific situation. The net result of this approach is a secure organization based on best practices and proven techniques.

NIST SP 800-53 contains two critical and targeted chapters on Fundamentals (security controls introduction, structure, and purpose) and Process (implementation guidance). The document covers numerous scenarios and helpful advice. For example Section 2.3 delineates three types of controls:

Common Controls – These are controls provided to a consumer (such as the Payroll department) by a third party (for example, the Information Technology Network Management group). These controls provide an *inheritable framework* for the entire organization, and allow a true corporate *security policy* to be provided. (A *security policy* includes the overall security posture for the organization, specifically the directives, regulations, rules, and practices guiding an organization's information management to protect against security breaches.)[17]

System-specific Controls – This set of controls includes those that enable individual systems to function effectively. For example, a common control for an organization may require physical access to all computer systems to use strong authentication such as smartcard logon. Individual systems may add to this requirement that of secure system access (such as is mandated for computer systems authorized for CLASSIFIED data storage within the U.S military).[18]

Hybrid Controls – Controls in this category include elements both from the common (organization wide) category as well as the system-specific category. A physical control in this category might be one pertaining to physical break-in response; the security policy (a common control) could provide the overall guidance in terms of management notification while specific procedures could be deemed system-specific. The standard notes that hybrid controls lend themselves well to function as templates an organization can use for further control customization.

Section 2.2.3: Mitigation

As stated above, NIST SP 800-53 provides Appendix F (Security Control Catalog) which identifies a number of guidelines for implementing all of the control areas. Some of the key points that specifically address physical controls include:

Category	*Code*	*Name*	*Comment*
Media Protection	MP-4	Media Storage	"The information system uses cryptographic mechanisms to protect and restrict access to information on portable digital media." Interpret this as support for *data-at-rest*

			(DAR), a critical element of guarding physical equipment in the event of loss.
	MP-5	Media Transport	"Protection and control of media during physical transport." This could have assisted the British Ministry of Defense as it "lost" USB sticks over a multi-year period from 2003 to 2007.[19]
Physical and Environmental Protection	PE-2	Physical Access Authorizations	"Use strong (two-factor) authentication." We cover this methodology in more detail further in this paper, for now the reader should keep in mind that requiring this type of control for after-hours school entry would probably have prevented the Millville high school computer theft.
	PE-3	Physical Access Control	"The organization provides physical boundary enforcement, such as manned ingress and egress points." Individuals are authoritatively identified prior to being allowed facility access, once again an effective deterrent.
	PE-6	Monitoring Physical Access	"The organization uses devices (such as video cameras or motion detectors) that can detect and record access." Monitoring is useful, but only if the organization has a defined structure in place to *react* to intrusions.
	PE-8	Access Records	"The organization maintains and reviews visitor access records persistently." Taking this approach can aid greatly in forensics investigations after an incident has occurred.
	PE-16	Delivery and Removal	"The organization controls, tracks, and verifies correct receipt of equipment from vendors and customers as well as equipment deliveries made to third parties." Implicit in this requirement is that of effective inventory management.

Table 1: Selected NIST Physical Security Controls

1.3 Recommendations

The following recommendations are specific to physical security for a small organization:

1. **Define a security policy** – In order to handle growth, an organization needs to create an overall security policy. This overall policy guides all types of controls: administrative (e.g. employee hiring and termination), technical (e.g. data management and privacy controls), and physical

(e.g. employee safety and facility planning). As the organization grows, the security policy grows with it.

2. **Know the inventory** – Setup a strong inventory management system. Label and record each piece of equipment along with the responsible party. Perform periodic reviews of this inventory to verify that all equipment can be identified and verified. Inventory management is top priority along with defining the organization-wide security policy.

3. **Protect data backups** – Create a secure location for backup data, and ensure that delivery to and from that secure location occurs through well-managed channels. This external backup can be done very inexpensively for small organizations and may be as simple as having the Chief Security Officer store backup media in a fire-resistant safe at home.

4. **Enable strong authentication** – Invest in a Public Key Infrastructure (PKI, discussed in more detail further in this paper) to allow the distribution and management of smart cards. Provision the smart cards to contain a biometrics reading (such as fingerprint or iris) and require the use of both (the card and the reading) to permit access to the facility. For system access (logical controls), this strong authentication is quite important and should be setup as soon as possible.

5. **Secure the facilities** – In addition to properly hardened doors and windows and a burglar alarm from a reputable monitoring company, investing in a closed-circuit TV at exit points (along with warning signage) can help to deter both external and internal unauthorized access.

2.0 Logical Access

Logical access relates to "soft" or "technical" controls[20] put in place to protect computer systems or facilities. Examples of these controls include passwords, thumbprints or other biometric data readers, and smart cards. The organization's security policy should define the high-level requirements for data classification as these requirements drive the type of controls necessary to protect specific systems. (For example, administrative logon access to a publicly-facing network server should be more difficult to obtain than for an internal workstation.)

2.1 In the News…

Hotmail in Hot Water[21] – On October 6, 2009, thousands of Hotmail.com mail users were told to change their passwords after their account details were posted online due to a phishing attack[22]. Each user had received one or more fraudulent emails that directed her to open a page resembling a real Web site, where the user was prompted to enter password and security information. Once entered, the fraudulent Web site then added the user to a list of compromised accounts; one such list (with user accounts from A to B) was publicly shared on the site pastebin.com. Microsoft acknowledged that numerous other compromised lists could exist.

40 pence for your thoughts[23] – On April 14, 2009, reporters shared the troubling finding that stolen online banking passwords can be had for as little as 40 pence — cheaper than a can of soda. These

passwords are routinely harvested from unsuspecting users by key loggers[24] installed as a result of unsafe Web surfing. These key loggers capture not only credit card numbers, but also user names and even the "safe" 3-digit security code from the back of the credit card. One especially troubling note from the article indicates that attackers are actively targeting Facebook and MySpace users who divulge significant personal information. The word *teenager* comes to mind...

Trite Twitter[25] – On July 15, 2009, Twitter was reported as leaving a "search" server with the default password of..."password." While the company did not report any personally identifiable information (PII) as being at risk, the perceptions around the exposure bring out the real fears that people have for storing data "in the Cloud;" namely, that data (and servers) are only as safe as the managing organization. Additionally, as Han points out in his case study on cloud computing, when problems such as Twitter's do occur there is a real question on jurisdiction and legal recourse for the injured parties.[26]

2.2 Problems and Mitigation

Logical access depends on how users access computer systems (basically, the logon required). In order to access a computer system, users must be identified (example: providing a user ID), verified (example: entering a password), and authorized (that is, what functions the logged-on user can perform). Of the different types of verification in common use, passwords are the most common,[27] and in this section we focus on the problems associated with passwords as a verification technique.

2.2.1 The Problem with Passwords

While passwords are the most common logon control, multiple passwords are the bane of every user's life. Passwords are truly the most problematic and least effective of user access controls: they can be shared, forgotten, guessed, or reused. System administrators add to this problem by emphasizing aggressive password expiration policies. Ironically, this emphasis itself leads to insecurity because passwords do not exist in a vacuum. A given user may have dozens of passwords for different purposes. Each password has its own requirements; for example, Microsoft Hotmail allows very strong passwords of long length and containing special characters while MyCheckFree.com allows only eight letters or numbers. (In the author's case, he has a special-purpose program[28] just to manage the *hundreds* of passwords he is required to use — the password program itself guarded by yet-another-password.) The result of these these multiple and conflicting password requirements? Users do not choose and use strong passwords, but rather passwords that can be managed and remembered. Some system administrators seek to address this problem by automatically generating strong passwords for users, but that has a trivially easy workaround as well: the user simply writes the password on a sticky note attached to the monitor or stuffed in a purse or wallet.

One effective mitigation for the password problem relates directly back to the organization's overall security policy, specifically in regard to *user awareness training*. Policies created without a defined training plan are doomed to irrelevance, so by instituting an appropriate Usage Policy[29] as part of the overall security policy we help to ensure that passwords are as strong as possible.

2.2.2 Some Common Password Alternatives

In our section on Physical Security, we touched upon two-factor authentication and indicated that this provided "strong" authentication by using any two of the following:

Something you **are**	A reliable long-term physical trait such as fingerprints or the shape of blood vessels at the rear of the eye. In a word, "static biometrics."[30]
Something you **do**	Handwriting analysis, voice recognition; basically, a biometric characteristic that can change ("dynamic biometrics"). Behavior-based biometrics offer less accuracy than static biometrics due to the variability inherent in the readings involved.
Something you **have**	An identity card, or perhaps a "synchronous token" that flashes a code synchronized to the remote system.
Something you **know**	Our old friend "passwords" (and its near relation "passphrases"[31]) fall into this category.

Table 2: Authentication Types

Of the three non-password-related elements, each has its relative strengths and weaknesses. One key element to a successful authentication scheme is that of user acceptance,[32] while another is how well the scheme handles various errors: *Type 1* (false acceptance rate — a bad guy got in!) vs. *Type 2* (false rejection rate — the user cannot logon).[33] We provide some of the more common comparisons below:

Authentication Type	**Pros**	**Cons**
Behavior-based (e.g. voice recognition, "something you do")	High user acceptance rate	Time-consuming (slow throughput rate) Higher Type 1 errors (less accurate)
Ownership-based (e.g. tokens, "something you have")	Very difficult to crack (the token is never used by itself) Easy to setup and integrate	Users can lose the token Users may need multiple tokens for access to different systems (especially when maintained by different organizations)
Characteristic-based (e.g. iris scan, "something you are")	Depending on the biometric chosen, can be extremely accurate	Privacy concerns Hygiene concerns (for example, hand-geometry scanners require users to put

		their hands onto a shared device) Safety concerns (retina scans are widely, albeit erroneously, believed to use a laser with subsequent danger to the user's eye)

Table 3: Authentication Types: Pros and Cons

2.2.3 Passfaces[34]

One interesting alternative to traditional passwords is based upon the well-known human ability to identify known faces very quickly and reliably. In this model, users are allowed to specify a number of pictures (faces) during logon setup. At logon time, the user is presented with a three-by-three square that includes eight incorrect faces (picked from random composite images) as well as one of the previously submitted faces. Studies have shown that a human being can select the well-known face extremely quickly. As an added security feature, the system can subtly alter the shape and coloring of the faces presented so that they are never quite the same for each logon; humans consistently demonstrate that they still have no trouble identifying the known face from strangers. Such an approach solves numerous problems associated with passwords. First, it removes the memorization (left-brain) aspect of logon and instead concentrates on the more hard-wired right-brain facial recognition capabilities, making it extremely difficult for malicious parties to "get the password" (keyboard logging becomes useless). Second, by making the logon screen very brief we lower the attack surface even if another user sees the screen; it can be very difficult to identify which face was chosen. Simple techniques (such as making the selection keyboard based rather than mouse based) makes the selection process virtually uncrackable.

2.2.4 Public Key Infrastructure (PKI)

In order for the biometric and other authentication mechanisms discussed above to be implemented most effectively, organizations typically create a *public key infrastructure* or PKI. This type of infrastructure ensures that every entity (user, program, or computer) has a unique "certificate" generated from a centrally-trusted "certificate authority." (X.509v3 is the latest version.)[35] These certificates have two aspects — the *public* aspect ("public key") which can be shared with everyone, and the *private* aspect ("private key") which must be guarded by the entity. Losing the private key effectively negates the value of the certificate assigned to that entity.

While space considerations preclude a detailed discussion of PKI, we touch on the key elements below. PKI exists to solve a fundamental problem in public key key exchange: ensuring that the entities in the communication link are indeed who they say they are.[36] First, PKI depends on the interesting fact that certain mathematical operations can be performed in such a way that two keys are required for an encryption process: one key to encrypt, and a second key to decrypt. The key used to encrypt a value cannot be used to decrypt that same value. Several algorithms have been created since the 1970s to build upon this type of relationship.[37]

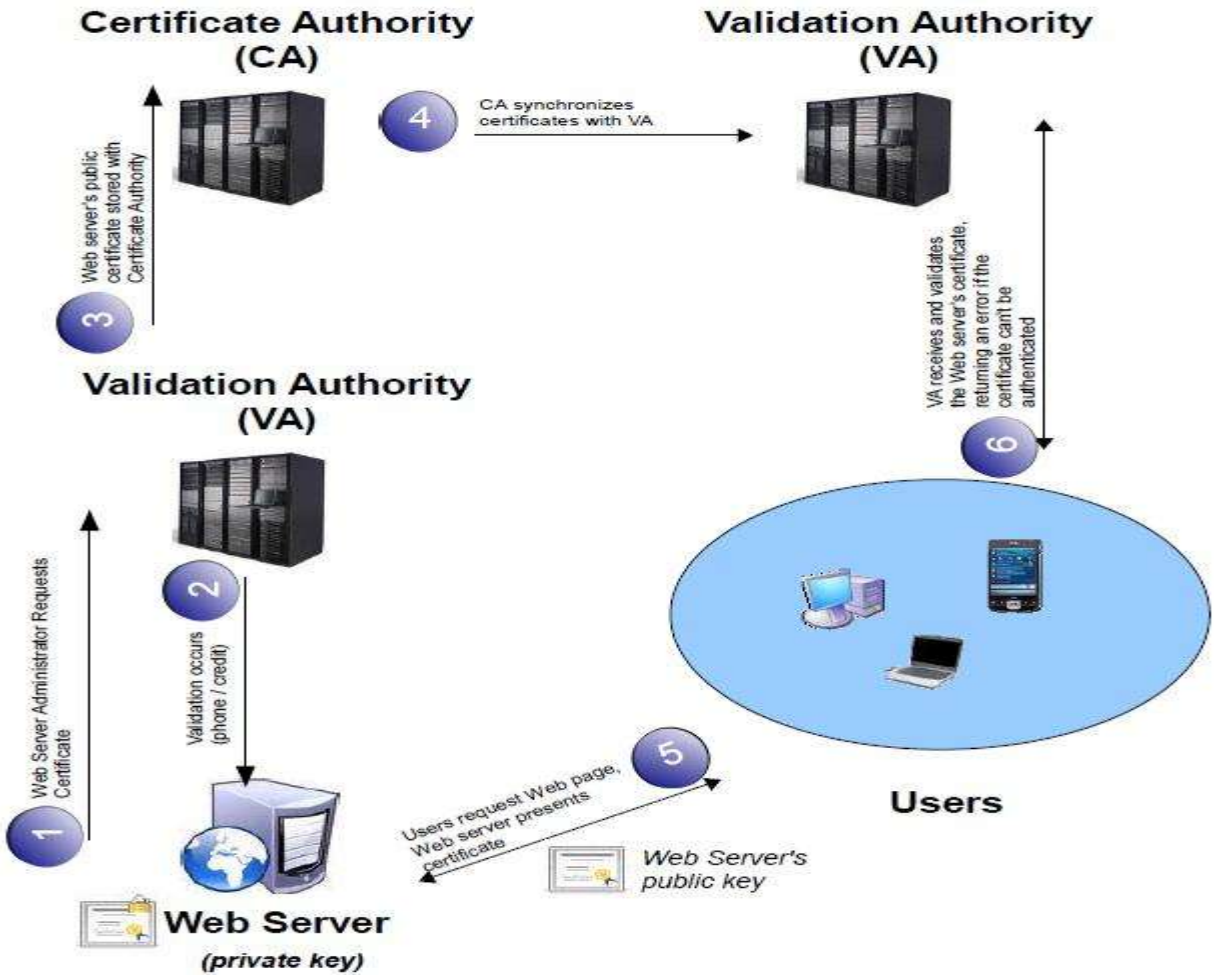

Drawing 1: Simple Public Key Infrastructure (PKI)

In a PKI framework, a central Certificate Authority (CA) serves as the issuer and holder of all certificates. When an entity (say, a new Web server) wants a new certificate, then that entity generates a *key pair* (public and private key) and submits the *public* key to a Registration Authority (RA). The RA verifies the authenticity of the request (as an example, GoDaddy.com uses a telephone call to verify the identity of an organization requesting one of their "Deluxe SSL" certificates)[38] and then submits the public key to the CA. The CA uses the submitted public key to generate a new certificate of the requested type (in our example case, a Web server certificate that allows Web browsers to connect) and returns that certificate to the requesting entity. This certificate (which does not contain the private key) can be provided to all clients who connect to our Web server; the clients issue a programmatic query to a Validation Authority (VA) to verify that the public key truly is from that Web server (and not an impostor).

The value of the PKI is shown during a typical Web site interaction:[39]

1. The user directs the Web browser (client) to navigate to a Web site (say, https://www.us.army.mil).

2. The Web browser first requests the certificate containing the Web server's public key from the desired Web site. The browser then uses a standards-based protocol to validate that the certificate is from a trusted source. As Millán et. al. point out, this protocol can be Online Certificate Service Protocol (OCSP) or Server–based Certificate Validation Protocol (SCVP), or any other protocol agreed-upon between the client and server. The only requirement is that the *Web browser client fully trusts the Certificate Authority* that issued the Web server's certificate.

3. As part of the certification verification, a number of checks occur.

    - First, the Web browser checks the *certificate revocation list* (CRL) to be sure that the certificate is still valid. If the Web server has been hacked (lost control over its private key) then the Web server administrator must notify the CA to *revoke* the Web server's related certificate; this adds the certificate to the CRL so that clients will no longer trust that certificate. Problems arise, of course, if the Web server administrator either does not know that the private key has been hacked or if the Web server administrator fails to notify the issuing CA (as still happens regularly today).[40]

    - Second, the Web browser engages in a handshaking protocol with the Web server that allows a *shared secret* to be exchanged. Basically, the Web browser encrypts a value with the Web server's provided public key and sends it to the Web server along with the Web browser's own public key. This encrypted value cannot be decrypted with the Web server's public key; it can be decrypted only with the Web server's private key. The Web server must decrypt the value, and then send it back encrypted with the Web browser's public key. The Web browser decrypts that returned value with its own private key and compares it to the original value sent — if they match, then the Web server is actually who it says it is.

This process guards against both *man-in-the-middle* and *spoofing* attacks. A man-in-the-middle attack is where a malicious party intercepts traffic between two entities; PKI guards against this because the man-in-the-middle may have the Web server's public key but will not have its private key. Therefore the shared secret exchange will fail. The same is true for spoofing (where a Web server uses an existing public key to pretend to be a different entity); once again, for this to work the spoofing party must have the private key. PKI provides a strong set of authentication functions allowing entities to exchange information safely, and is in fact the method used by all secure Web servers (those that require **https** to access them).

Further in this paper, we'll explore another useful approach for handling logical access, especially in the context of cross-network user access: *single-sign on* (SSO). SSO typically builds upon a PKI and leverages the strengths that PKI delivers.

2.3 Recommendations

For an organization to prove to outside auditors that it is serious about protecting logical access to systems and data, the organization must be able to provide a security policy that defines its overall security posture. This posture includes specific (“drill-down”) policies and procedures on how logical

security is applied at different levels within the organization, and these recommendations provide a starting point for the lower-level policies and procedures:

1. **Define a Security Domain** – Identify the organization's users, and ensure that these users are accounted for in an organizational hierarchy maintained as part of the network infrastructure. If using the Windows Server family of operating systems, Active Directory provides this capability along with the ability to organize user into logical units ("organizational units," or OUs) that can be used to map the organization in a number of ways (for example, by geographical location or by corporate structure).[41] The key part of a security domain is that it allows one to identify and define user permissions based on common characteristics (such as "Software Developers" vs. "Network Engineers"), and also provides the underlying foundation for a single-sign on (SSO) solution across the enterprise. While in our organization we do have an effective security domain defined, we have not considered how to define a common set of user attributes that will allow our security domain to exist within a larger identity federation (discussed in more detail further in this paper).

2. **Setup and Use a Certificate Authority** (CA) – As discussed above, certificates are an integral part of a public key infrastructure (PKI). By using a Certificate Authority, one provides the ability to associate multiple certificates with your end users. As an example, in our own organization we use our CA for all sorts of purposes, from identifying remote clients in our virtual private network (VPN) server to providing "code signing" certificates that vouch for the integrity and authenticity of software products we build internally. However, we can do more with our CA by integrating with other authorities of record (such as the Federal Bridge Certification Authority, discussed further below) to enable single-sign on across not only our organization but also with our partners and vendors.

3. **Implement strong authentication for access** – Using the CA mentioned above in conjunction with biometrics, configure all operating systems to require two-factor authentication that does not require passwords for entry. The best way to mitigate the plethora of problems associated with passwords is by eliminating the need for them at every logon. Enabling strong authentication should be considered high-priority.

3.0 External Attacks

A major focus of the organization's security policy is in preventing outside malicious agents from accessing our protected assets (computers and data) via external network attacks. These external attacks work in various ways; they can be made to access computer systems illicitly (such as a criminal trying to get unauthorized access to data on a system) or can be made to deny access to others (such as the common *denial-of-service* attacks that attempt to take existing computer networks such as Amazon.com offline from customers).

3.1 In the News…

Stock market Cyberwars?[42] – On June 13, 2010, investment legend Warren Buffet mentioned during an interview with CNBC the possibility that the so-called "flash crash" on May 6, 2010,[43] could have been caused by a computer attack on the major Wall Street houses. He followed up this terrifying statement with the reassuring news that he wasn't concerned in the slightest about the event. The prospect of a concerted attack upon our financial infrastructure certainly is not something of which to speak glibly.

An unhealthy Apple[44] – On June 9, 2010, Apple confirmed that over 114,000 iPad owners suffered exposure of both their email addresses and their unique ID for authenticating on the AT&T network. The list of names exposed in this way included such well-known figures as Les Hinton (CEO of Dow Jones) and Rahm Emanuel (White House Chief of Staff to President Barack Obama). The attack was made by a group calling itself Goatse Security[45] by making unauthenticated sequential requests to a script running on AT&T's Web site.

Breaches are Looking Up[46] – The 2010 infoSecurity Report by PriceWaterhouseCoopers (PWC) provides voluntary responses from industry leaders on the overall state of security and their organization. The report indicates that 92% of large respondents had a security incident over the last 12 months as compared with 72% from 2008. Within that sobering statistic, we see more bad news: 61% of these same large respondents identified significant network break-in attempts compared to just 31% from 2008.

3.2 Problems and Mitigation

External attacks occur constantly against computer networks, and as we've seen above these attacks are only growing. Attacks, however, do not occur randomly against all computers within an organization; instead, specific attack vectors exist. Primarily, those computers and devices physically connecting one network to another are at risk. Let's look at the types of physical networking equipment that are at risk, why they are at risk, and discuss how to mitigate these risks.

3.2.1 The Seven Layers of Network Connectivity

In order to understand the types of external attacks that can occur against a given network, one must first consider the logical way that information is shared between two systems. Over the course of many years, the Open System Interconnection (OSI) initiative started in 1978 by the International Organization for Standardization (ISO) defined what is called the *OSI Model* defining seven "layers" for network communications.[47] All network equipment works at one (or more) of these layers, and all external attacks target characteristics unique to individual layers. We'll very briefly discuss each layer here as well as their vulnerabilities.

3.2.1.1 Layer 1 – Physical

This layer defines the physical electronic signals sent across a cable (or through the air) between two connected devices. Attacks at this layer generally include tying into the physical connection (for example, splicing into a network cable or eavesdropping on wireless network traffic). To mitigate this attack surface, physical cabling can be shielded and routed through secure conduits; this both protects

physical access as well as preventing data emanations from being read by eavesdroppers. Wireless communications can be encrypted and wireless access points can be required to support strong authentication (to prevent clients from connecting to rogue wireless access points masquerading as a legitimate access point).

3.2.1.2 Layer 2 – Data Link

This layer wraps the "frames" (individual data packets) that are sent over the physical connection. Errors in physical transmission are detected at this level and automatically corrected. Associated with this layer is the Media Access Control (MAC) code, a 48-bit code provided by the device manufacturer that uniquely identifies this specific device from all others. As Marro notes, attacks at this layer generally result from "insider problems"[48] and lead to traffic overload (and thus denial-of-service). Additionally, impersonation via "ARP poisoning" (also referred to as MAC poisoning) can occur. *Address Resolution Protocol* (ARP) poisoning occurs because Layer 2 requires a physical ("MAC") address in order to send a network message between two machines. If Machine "A" simply has an Internet (IP) address, then it must learn the physical address. A special message is sent out using ARP; an attacker can listen for these messages and simply try to be the first machine to respond to them. From that point forward on the local network, messages will be sent to the attacker.

To mitigate attacks at Layer 2, one can limit the number of MAC addresses that can be "learned" (basically imposing an upper bound on the number of stations that can send data on this network segment). Additionally, disabling unused connection ports can significantly reduce the attack surface area. On small networks, the ARP impersonation attack mentioned above can be mitigated by using static IP to MAC translation tables (thus avoiding the use of Address Resolution Protocol at all).

3.2.1.3 Layer 3 – Network Layer

This layer provides the two required elements of the modern Internet: an *Internet Protocol* (IP) address (either 4 octets in IPv4 or 8 octets in IPv6) that identifies the destination machine, and a *port number* such as port 80 for standard Web page requests via the Hypertext Transfer Protocol (HTTP). By definition this layer involves the capability for messages to traverse networks (to be *routed*, as we will discuss later). This routing occurs by having individual frames make "hops" from computer to computer until the frame arrives at its ultimate destination (further below we discuss timeouts). This layer has specific vulnerabilities to three different types of attack:[49]

> *sniffing* – As data messages are sent across multiple networks, all machines involved in the forwarding process have the opportunity to intercept and attempt to read the message data.
>
> *spoofing* - Each frame has a unique IP address that identifies the sender; malicious individuals can simply modify this "source IP" address to make it seem as though a message is arriving from a trusted network.
>
> *modification* – As with sniffing, malicious software can be used to attempt to modify the frames as they are "hopped" across the Internet.

At this layer, mitigation consists in requiring security and authentication between two connections. Security encrypts messages between two computers so that sniffing and modification attacks fail, while

authentication ensures that each computer is actually who they say they are. Public Key Infrastructure (PKI) provides this type of authentication and is widely used on the Internet as we saw above.

Another effective mitigation technique for both the sending and receiving sides is the use of *packet filtering firewalls*. These firewalls look at the source and destination addresses in each transmitted message as well as the communication port, and can apply filtering rules setup by the system administrator. For example, to block default "telnet" access to a network,[50] a packet filtering firewall can simply watch for and deny any packet with a port of "23" in it.

3.2.1.4 Layer 4 – Transport Layer

This layer defines end-to-end connectivity either by *connection-oriented* or *connectionless* protocols. In a connection-oriented protocol (such as Transmission Control Protocol, or TCP), requests from one device are guaranteed to arrive at the other device and to have a response returned. Think of it like a package mailed with delivery confirmation; the Post Office guarantees that the package will be delivered and that a confirmation is returned to the sender (or the original package if it cannot be delivered). Compare that to connectionless protocols (such as User Datagram Protocol, or UDP); these protocols do not guarantee delivery. They are most like a standard first-class letter where a best-effort is made to deliver the mail.

Mitigation at the transport layer includes *circuit-level firewalls*, which look at connections made using TCP or UDP. These connections can be analyzed with administrator-configured rules to allow or disallow the connection; once a connection has been allowed then no further analysis occurs.

3.2.1.5 Layer 5 – Session Layer

This layer goes one step beyond guaranteed delivery of individual messages as defined in the transport layer: instead, an entire *session* is created for sending and receiving multiple messages. Virtual Private Networks (VPNs) typically use this layer to pass data.

Attack opportunities at this layer include forged certificates (discussed in the Public Key Infrastructure topic above), whereby a malicious server attempts to fool the sender into thinking that it has reached a different (and desired) destination. For example, in 2001 the commercial Certificate Authority VeriSign incorrectly issued twp digital certificates to fraudulent entities purporting to be Microsoft Corporation.[51] A "secure" network connection between a Web browser and a Web server using one of those fraudulent certificates would have led the user to believe that she was actually connected to Microsoft when such was not the case!

3.2.1.6 Layer 6 – Presentation Layer

This layer exists to translate specific types of messages from the running application (such as Web requests from the browser that uses the HTTP protocol) to and from the lower-level data encoding schemes. The presentation layer is susceptible to indirect attacks; a network message purporting to be in one format (such as HTTP) could be deliberately malformed so that attempts by the presentation layer to interpret the message could cause a fault. Thus, attacks at this layer fall into the denial-of-service category. Mitigation includes keeping operating system and third-party application software patches up-to-date.

3.2.1.7 Layer 7 – Application Layer

This is the layer that most end-users are familiar with; it includes Web browsers and email clients. Data messages sent between applications (such as an email client to a mail server) are encoded using an "application-level" protocol (in our email example, most often this would be Simple Mail Transfer Protocol or SMTP) and then submitted to the lower layers for physical transmission to the ultimate destination. The receiver (our mail server) then decodes the data contained in the message and performs the requested action. Using our email example, this action might include forwarding the mail message to its ultimate destination.

The application layer is a prime target (perhaps *the* prime target) for external attacks. Attacks on most of the other network layers tend to be geared toward denial-of-service (DOS); consider an attack at Layer 3 (Internet layer) by generating massive amounts of inbound requests to overwhelm a Web server. At the application level, however, it becomes possible for the attacker to take advantage of specific and sometimes well-known flaws in the receiving application itself. These flaws can occasionally be extremely destructive and can lead to allowing the attacker to take control of the machine itself (such as occurred against the popular free remote-management application "VNC Server" in 2002).[52]

Attacks against the application layer typically fall into the *malformed message* category, where an attacker especially crafts a network message to force a fault in the receiving program. Keeping in mind that many application programs have source code readily available on the Internet, criminals can study this source code to discover weaknesses. One common scenario resulting from this study is the *buffer overflow* attack, where a malformed network message containing more information than the application expects to receive. This can cause the application to fail in such a way that the attacker can either force protected data to be revealed or the attacker herself can achieve remote control.

Effective mitigation at the application layer includes the use of *application-level gateways*. These are special firewalls that analyze not just the raw network traffic but also the application-specific contents of each message entering or exiting the network. For example, an application-level gateway can analyze email traffic to look for and eliminate spam before it enters the network. Additionally, application-level gateways can provide *proxy services* whereby the actual network addresses of internal machines are hidden. This effectively eliminates an attacker's chance to attack internal machines without first compromising the proxy server itself (generally a much more difficult feat to accomplish).

3.2.1 Network Connection Equipment Types

The seven OSI layers listed above all run on physical devices. To understand the attack vectors better and to apply mitigation more effectively, we must understand the relationship of the OSI layers to basic network equipment types.[53]

Hubs – These devices connect a group of nodes (network-capable computing devices) together at Layer 1 (physical layer); each node connects to the hub via an RJ-45 connection[54] called a *port*. Communication packets called *frames* are used to send data across the network; when a hub receives a frame from one node it automatically sends that frame to all other connected nodes. Hubs provide no data security and generally result in excessive network traffic. They should not be used in a production environment to connect networks together.

Switches – Similar to hubs, switches go one step further to identify each attached node uniquely a manufacturer-supplied code (the Media Access Control code, or "MAC Address"). Frames arriving at the bridge are analyzed at Layer 2 (data link layer). This layer uses the MAC address to determine the frame's destination. Because the majority of network frames are point-to-point (computer to computer) a switch generally results in lower network traffic than a hub. However, switches provide very little security because every frame is simply a "blob" of data. For this reason switches are not used to connect two networks together except where both networks are highly trusted.

Bridges – Bridges work at the same Layer 2 as switches (the data link layer where the MAC address is located). However, bridges are designed to connect networks efficiently (at least up to a point). Bridges work by looking at both the source and destination MAC addresses associated with each frame and "learning" where different nodes are located. Over time, this can result in a highly efficient network traffic pattern. As networks grow, however, the number of different communication paths (source to destination) grows exponentially. Also, bridges suffer the same lack of security as switches; bridges should never be used to connect an untrusted network (such as the Internet) directly to another network (such as the corporate backbone).

Router – Routers provide the same functionality as hubs and switches, but with the added capability of *routing* frames between networks and by generally operating at Layer 3 (internet layer). For example, if a Web browser makes a request to www.google.com then that means the router on the local network must automatically forward that request out to the Internet. Routers accomplish this task by looking not only at the raw MAC address, but at the destination IP address. The request to www.google.com actually gets translated to an Internet Protocol (IP) address via the Domain Name System protocol, or DNS. (DNS is a critical component of modern computing, but space precludes a discussion of it here.) Thus, the router does not see the destination as www.google.com but instead as the resolved IP address (such as 74.125.157.147).[55] Routers work by using rules (*routing tables*) that allow the router to decide whether a particular frame should be addressed to the local network or should be forwarded ("routed") to another network. Routers provide the ability to perform security, and are frequently used in production networks. We discuss routing in more detail further in this paper as it relates to network segregation.

Section 4.2.2: Modems

Modems (modulator-demodulators) were invented in the 1950s and by the 1960s allowed telephone lines to be used to transfer data.[56] Modems work by modulating digital data (electrical data) into analog data (sound waves) and sending that sound across copper telephone wires to be received by a modem on the other end (which demodulates the analog data back into digital data). While modern modems are no longer analog (for example, a cable modem connected via a fiber-optic line to the cable provider), they still serve the same purpose of receiving data from an external source and transmitting that data to an internal network node (generally a router, but a modem can also be connected directly to a PC).

Modems themselves have been a direct attack vector in the past, especially with regard to *war dialing* (dialing blocks of telephone numbers to discover modems that might be listening). However, today modems are generally connected to a hardened network server with a firewall that prevents unauthorized connections to be made to an internal computer system. The best way to keep a modem

secure is to ensure usage of a late model with software patches applied. Most Internet service providers include modems into the standard subscription plan, and automatically manage these modems on behalf of the consumer.

3.3 Recommendations

1. **Invest in application-level gateways.** This provides by far the biggest bang for the buck, because a modern gateway can analyze numerous types of network traffic. Some solutions (such as Microsoft's Threat Management Gateway 2010 product) offer subscription-based services where sophisticated rule sets are constantly updated to reflect the latest known network attacks and suspected Web addresses. These gateways are very inexpensive for their functionality; the TMG 2010 product can be deployed for a reasonably large organization (up to 750 persons) for less than $6,000 for the enterprise-level license.[57]

2. **Segment your network.** Rather than having all server and client machines on a single network, separate machines based on their logical function. For example, put the Software Development group on a separate network from the Marketing group. Also, create special "demilitarized zones" (DMZs) that contain public-facing servers (such as a corporate Web server). We discuss this in more detail in the next section. The zones themselves can be setup inexpensively, although the hardware to support full network segmentation can be expensive.

3. **Disable unnecessary network traffic.** By default, all types of traffic can occur over a network. Most of the time, this is not what is desired; instead, only explicit types of network traffic should be allowed. While firewalls can detect and ignore unwanted network traffic, it is possible and in many cases easy to program lower-level network switches and bridges simply to avoid passing on unnecessary network traffic in the first place. Because of the low cost factor and ease of this work, it should be considered high-priority ("low-hanging fruit").

4.0 Internal Protection Failures

In this last analysis of organizational security policy failures, we look at how a failure of *network segregation* can allow unauthorized or malicious users to access sensitive data and to share this data with third parties. A formal security policy can help to prevent this type of failure by making the safe storage of protected data a part of the corporate mindset from the beginning of network and infrastructure design.

4.1 In the News…

No Hooah[58] Here[59] – On June 12, 2010, an Army Intelligence analyst was jailed for leaking classified information regarding an ongoing investigation into Chinese cyberattacks against Google to the Wikileaks whistle blower Web site.[60] The analyst had actually shared this information with a well-known hacker via email, and had apparently revealed classified information in several other incidents. As of this writing, this investigation has not been completed.

National Insecurity?[61] - On April 15, 2010, a former top official at the National Security Agency (NSA) was indicted for providing classified information to a journalist. The official had, over the course of 2006 and 2007, provided source material to the reporter (known only as "Reporter A") for numerous articles on the NSA. Not content with that level of involvement, the official was even helpful enough to offer editing and reviewing services to the reporter. Not mentioned in the article were any reasons for the official taking these steps (we see the Opportunity and the Means, but what was the Motive?).

4.2 Problems and Mitigation

4.2.1 Problem Space

The news articles above highlight problems stemming from unauthorized and uncontrolled access to network data. Our thesis is that such problems ultimately derive from a faulty organizational security policy, and that a properly specified security policy would address internal data protection via a verifiable strategy. Of the many ways to protect data internally that could be specified in the organization's security policy, in this paper we focus on *network segregation*. Network segregation allows an organization to treat its various departments much like a ship at sea treats compartmentalization: the idea is that a breach in one area stays contained within that area. Additionally, connection points between network "segments" (groups of computers connected to one physical station) can be hardened and centrally-managed, allowing full disconnection support in the event that a serious breach is discovered. One practical application of this approach is to segregate a production network environment (such as a corporate Web server) from a beta testing network. Both networks may be available from the Internet, but you definitely want to ensure that problems on the beta testing servers have no impact on your production network.

One key problem to solve with network segregation is the issue of *cross-network access* for authorized individuals. (A 2010 report from Arbitron[62] shows that we hold the Internet as the "most essential" medium in our lives, so any network segregation must ensure that required access can still occur.) For example, consider two groups: Marketing and Product Development. Normally, Marketing has severely limited (or zero) access to Product Development computers (and vice versa). However, for a new e-commerce site under development a Marketing individual may have legitimate need to access a protected resource in the Product Development group. We'll discuss how this can be accomplished using a single-sign on approach; the beauty of single-sign on is that as the problem becomes larger (such as allowing a partner vendor to access an internal Web service), well-defined and implementable exist.

4.2.2 Routing

Above, we discussed the *why* of network segregation, and now we address the *how.* With a segregated network, the first problem to solve is the physical ability to *route messages* between the networks. Routing itself is a function of the Internet layer (Layer 3) within the OSI model we presented above; this layer defines the actions necessary to send a message from one station (and IP address) to a another station (also identified by its IP address). Where both stations exist on the same network segment, both stations identify themselves automatically using address resolution protocol (ARP). However, consider the case where a station wants to communicate with an external host (say, www.google.com). After the external name (www.google.com) has been translated to an IP address like 74.125.227.20 via the

Domain Naming System protocol (DNS, not covered here) the problem still remains on how the network message can actually be sent between the two stations. Somehow, we must establish a *virtual circuit* to the station represented by that other IP address.

This same problem (internal IP address connecting to IP address on a different network) affects not just an internal station trying to connect to Google, but trying to connect to *any other segregated network* in the organization. So by solving the problem for one case we effectively provide the strategy for solving for all cases.

The key to the routing problem is the *network gateway*. A gateway is really nothing more than a default route used to transfer unhandled network messages. For completely self-contained networks using only static IP addresses there are no unidentified routes to handle; thus, no need for a gateway (router). Normally this is not the case, and we do need to handle traffic to other segregated networks as well as the Internet. Depending upon the operating system, one uses **routing tables** to control how network messages to external IP addresses should be handled. On Windows, we can use the **route** command as shown below:

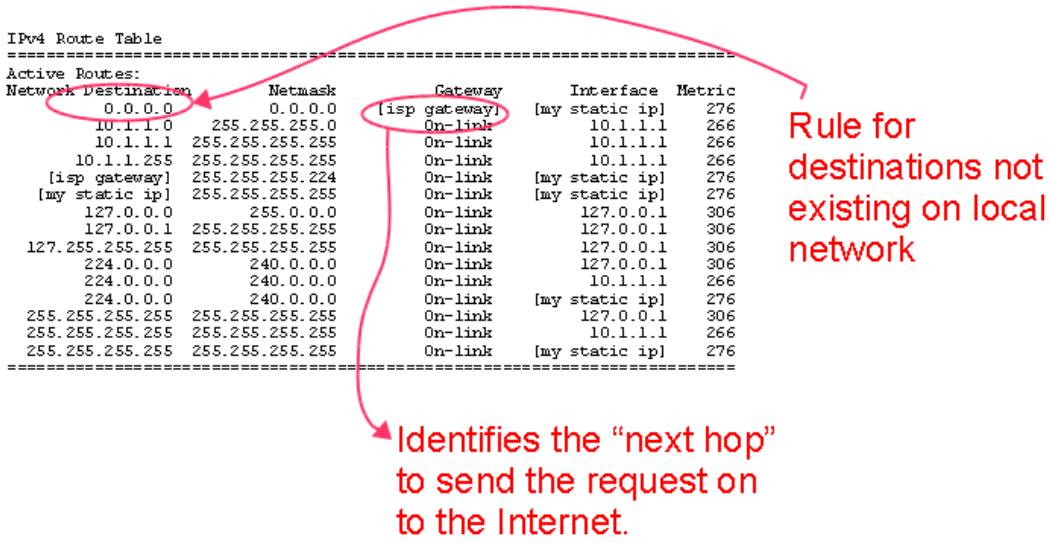

Drawing 2: Network routing

The key point from above is the line with a "Network Destination" of "0.0.0.0" that provides the route to use for IP addresses not physically associated with stations on the current network. In this case, we see that network traffic will be routed to the Internet Service Provider (ISP) gateway. That process of transferring a network message to that next gateway is called a "hop." The receiving gateway in turn checks to see if the IP address exists on its local network and delivers it if so; otherwise the process repeats to the next gateway. To prevent endless cycling of undeliverable messages (for example, to non-existent IP addresses) each network message has a Time-To-Live data field associated with it (TTL, defaults to a value of 255). This TTL is decremented by one upon each hop and when the TTL expires (becomes zero) the network message is discarded.

As described in Section 3 above, network message routing can be considered a prime attack vector – after all, this routing allows *all* software applications to establish remote communications (not just authorized applications). To combat this threat, we typically establish port-level shutdowns at Layer 2 (Data Link Layer) so that only well-defined ports like HTTP (80) or HTTPS (443) can be used. Additionally, we use application-level gateways as mentioned above to scan network message contents to detect malformed messages. Finally, we can use Layer 3 (Internet Layer) technology to look at the IP addresses (and external names) so that we block connection attempts to known bad addresses (such as known lists of adult Web sites, which is the approach used by tools such as NetNanny).

4.2.3 Single-sign On

As we mentioned in our discussion on PKI above, *single-sign on* (SSO) exists to enable access for authorized users across network boundaries. Basically, the problems we've discussed in this section highlight an underlying issue with network infrastructure as it currently exists; namely, that of *identifying the user*. In a perfect world, authentication security would be delegated to a strong, well-known, and completely trusted service. All Web sites, as an example, would use this centralized service to perform logons and would leverage the common *identity* provided by the logon service. (In fact, within the U.S. military, the Army Knowledge Online or AKO service performs precisely this function.)[63] However, that makes the incorrect assumption that multiple Web sites have an understanding of who a user is based on an external organization.

Let's consider some practical examples. Microsoft Corporation offers their “Passport” service whereby individual Web sites can allow a user to logon based on the Microsoft Passport credentials. Of course, Microsoft isn't the only organization to offer this service — Amazon.com also offers a shared logon service via their Amazon Web Services (AWS) platform. Additionally, Google offers their own identity service – and many others do so as well. Added to this complexity is the simple fact that the vast majority of commercial Web sites do not avail themselves of any type of shared logon service, while individual business organizations (such as logging into a work computer) invariably have separate logons created for each user.

The simple problem with multiple identities is that for each identity a user has (such as one for Google Apps, another for Amazon.com shopping, another one for online banking, yet another one for Facebook, and so on) the user *must be authenticated separately*. So if the user has a total of twenty online logons (not at all uncommon today) then the user has the equivalent of twenty identities. Each of those twenty identities requires a separate logon, thus a separate logon verification process. This leads to a tremendous duplication of effort along with significant management overhead required by each and every user on the Internet.

To combat this ever-widening set of online identities, many organizations and governments are looking hard at the concept of single-sign on. In this model, each time a user authenticates to an application, the application uses a shared service to perform the actual authentication (normally, businesses set this up by creating an organizational-wide user directory that all applications agree to use). However, the problem remains that this solves the logon issue just for that one company. One fascinating new entry in single-sign on is that of federation (and confederation) of user directories. Basically, this means that

cooperating organizations agree to honor user logons; as a practical example, the Euro6IX IST European research effort has linked numerous academic institutions together so that a student logon to one effectively provides a logon to any of the cooperating colleges. In the U.S., online efforts are well underway as shown by the Federal Bridge Certification Authority (FBCA), which ties together more than twenty organizations.[64]

One common problem brought out with single-sign on is that of *authorization*, or what an authenticated user is allowed to do. The fear is that by allowing authentication to occur through a central source (such as the FBCA effort mentioned above), computer administrators may have to contend with unintended privilege escalation (users being granted more permissions than they need to do their jobs). Nothing could be farther from the truth; in every single-sign on scheme there is a definite disconnect between the act of authentication (identifying and verifying a user's identity) and authorization. By effectively outsourcing the authentication aspect of system logon, all system administrators are really doing is simplifying the computing experience for their end users.

4.2.4 Single-sign On Options

Numerous application exist in the commercial and business worlds for single-sign on. For example, the Army Knowledge Online (AKO) portal (https://www.us.army.mil) offers a single-sign on service that is used by many Army systems (example: after logging into AKO, one can logon to an Army site like https://ea2f.army.mil without being challenged again). Let's take a quick overview of popular single-sign on solutions in use today:

Kerberos – Developed by MIT starting in 1983, this solution uses secret-key technology (all client computers must have a unique stored password with the Kerberos Authentication Server or AS). Numerous operating systems support Kerberos natively; the Windows operating system has supported Kerberos since the release of Windows 2000. (As with all secret-key technologies, secure *key exchange* becomes difficult as the network grows.) As a critical element of Microsoft's Active Directory, Kerberos can provide single-sign on support for the enterprise.

PKI – Kerberos simply does not work where a remote computer is not part of an organization's security domain and thus does not have a trusted password stored with the Kerberos Authentication Server. Enter the concept of *public key infrastructure* discussed in Section 3 above. As described there, a PKI allows two arbitrary computers to exchange keys securely by using the services of the PKI infrastructure (thus, the PKI infrastructure itself serves the Kerberos function of the Authentication Server). PKI can be extended through *federation* where different Certificate Authorities establish trusts between each other. While this type of single-sign on is powerful and quite secure, it can be difficult to manage and is still subject to attack (especially social engineering attacks on the Certificate Authorities themselves).

Federated Identity Management (FIM) – FIM exists to allow managing the user life-cycle when an SSO solution allows logon between organizations. A number of standards bodies are working on this problem (for example, the WS-Trust working group provides standards on how *tokens* representing *users* can be passed between organizations).[65] FIM works by providing the standards-based services necessary for a service provider to identify users authenticated by an identity provider. The biggest problem in this scenario, of course, is *trust* (as Smith puts it: “difficult to gain and easy to lose.”).[66] The two current

approaches for FIM include Microsoft Cardspace (users store identity information locally in "InfoCards") and OpenID (users store identity information in a trusted store on the Internet), time will show which approach becomes the industry standard.

Pervasive Trust Management (PTM) – Designed for mobile users, this model builds upon PKI when the devices are connected in conjunction with an "evidence-based" trust model that continuously scales the trust factor between participants.[67] New participants in the community are initially trusted minimally (neutral rating) and earn or lose trust depending upon their actions. By explicitly designing this system for unconnected support, this type of SSO allows for a *friend-of-a-friend* trust model (I trust John, John trusts Frank whom I've never met, therefore I trust Frank enough to grant him minimal access to my system). While targeted toward MANET (Mobile Ad-hoc Networks), this trust model shows promise for enabling "instant membership" into online Web sites.

4.3 Recommendations

1. **Establish single-sign on internally.** As segregated networks separate employee applications (such as HR and Payroll), trying to maintain multiple internal accounts for all workers and automated processes will prove unworkable. Kerberos can be quite effective for single-sign on and requires little effort to setup (although it can require much effort to implement fully). In our case, we need to integrate our user directory with our payroll provider (Paychex) to eliminate the need for extraneous user logons simply to perform payroll function. We should be able to enable this integration without extra cost beyond opening up our directory to Paychex using the secure Lightweight Directory Access Protocol (LDAP) network port number 686.

2. **Federate with partners.** Once established, determine who your partners are (Vendors? Customers? Others?) and determine how they will communicate with your network. In our case, we need to federate the Army Knowledge Online (AKO) user directory with our internal directory to enable our partners to logon to our networks using their Army ID. This requires us to develop an Active Directory Federation Services (ADFS) integration module; additionally, we will need a Memorandum of Agreement (MoA) with the AKO group.

Summary

In this paper, we have looked at current events and addressed four main topics of organizational security: *Physical Access* (protecting the physical plant and storage media), *Logical Access* (we focused on the problem of passwords), *External Attacks* (we focused on the ISO "protocol stack" to identify these attacks and how to stop them), and *Internal Attacks* (we focused on how we segregate our organizational computer networks while still allowing access via single-sign on). We addressed a number of of problems (along with mitigation strategies) and provided specific recommendations for each area. In the Appendices for this paper, we provide a a list of our recommendations along with estimates for implementation cost.

The overarching message we discovered is the importance of defining a corporate vision and ensuring that organizational efforts map back to that vision. In short, defining and delivering upon an overall *security policy* that defines what we as a company need to do to protect our investment both in hardware and software. All other initiatives follow directly from (and in support of!) this top-level effort. Finally and most importantly, by implementing a formal security policy we help to ensure that we protect our customer's data in the best possible way.

Appendix A: All Recommendations

In this appendix we highlight all of our recommendations in a numbered list.

1. **Define a security policy** – In order to handle growth, an organization needs to create an overall security policy. This overall policy guides all types of controls: administrative (e.g. employee hiring and termination), technical (e.g. data management and privacy controls), and physical (e.g. employee safety and facility planning). As the organization grows, the security policy grows with it.

2. **Know the inventory** – Setup a strong inventory management system. Label and record each piece of equipment along with the responsible party. Perform periodic reviews of this inventory to ensure that all equipment can be identified and verified. Inventory management is top priority along with defining the organization-wide security policy.

3. **Protect data backups** – Create a secure location for backup data, and ensure that delivery to and from that secure location occurs through well-managed channels. This external backup can be done very inexpensively for small organizations and may be as simple as having the Chief Security Officer store backup media in a fire-resistant safe at home.

4. **Enable strong authentication** – Invest in a Public Key Infrastructure (PKI) to allow the distribution and management of smart cards. Provision the smart cards to contain a biometrics reading (such as fingerprint or iris) and require the use of both (the card and the reading) to permit access to the facility. For system access (logical controls) this strong authentication is quite important and should be setup as soon as possible.

5. **Secure the facilities** – In addition to properly hardened doors and windows and a burglar alarm from a reputable monitoring company, investing in a closed-circuit TV at exit points (along with warning signage) can help to deter both external and internal unauthorized access.

6. **Invest in application-level gateways.** A modern application-level gateway can analyze numerous types of network traffic. Some solutions (such as Microsoft's Threat Management Gateway 2010 product) offer subscription-based services where sophisticated rule sets are constantly updated to reflect the latest known network attacks and suspected Web addresses.

7. **Segment your network.** Rather than having all server and client machines on a single network, separate machines based on their logical function. For example, put the Software Development group on a separate network from the Marketing group. Also, create special “demilitarized zones” (DMZs) that contain public-facing servers (such as a corporate Web server).

8. **Disable unnecessary network traffic.** By default, all types of traffic can occur over a network. Most of the time, this is not what is desired; instead, only explicit types of network traffic should be allowed. While firewalls can detect and ignore unwanted network traffic, it is possible and in

many cases easy to program lower-level network switches and bridges to simply avoid passing on unnecessary network traffic in the first place.

9. **Establish single-sign on internally.** As segregated networks separate employee applications (such as HR and Payroll), trying to maintain multiple internal accounts for all workers and automated processes will prove unworkable. We need to integrate our user directory with our payroll provider (Paychex) to eliminate the need for extraneous user logons simply to perform payroll functions. We should be able to enable this integration without extra cost beyond opening up our directory to Paychex using the secure Lightweight Directory Access Protocol (LDAP) network port number 686.

10. **Federate with partners.** Once established, determine who your partners are (Vendors? Customers? Others?) and determine how they will communicate with your network. In our case, we need to federate the Army Knowledge Online (AKO) user directory with our internal directory to enable our partners to logon to our networks using their Army ID. This requires us to develop an Active Directory Federation Services (ADFS) integration module; additionally, we will need a Memorandum of Agreement (MoA) with the AKO group.

Appendix B: Recommendations by Priority with Estimates

In this section, we list our recommendations ordered by priority, along with an estimated time-line for implementation.

Priority	Recommendation	Time Estimate	Dollar Estimate (August, 2010)
1	Define a security policy	For our small organization, approximately 10 days for gathering information and writing the policies. Add another week for updating employee manuals with training.	In many cases no direct cost, but definitely lost work hours.
2	Know the inventory	Approximately 5 days to setup the simplest management system (spreadsheet based).	Minimal; in the simplest case the cost of a label gun while a single personnel resource builds the spreadsheet.
3	Secure the facilities	Due diligence in selecting an alarm takes up the great majority of time (plan at least 30 days with a few in-person quotes before deciding). A proper system should be from a reputable firm and include door / window monitoring as well as smoke and heat detectors. If considering biometrics / smartcard integration, be prepared to pay more and to purchase supporting software systems.	Once selected, for small businesses the alarm installation installs very quickly. The average cost[68] of a bare-bones system (doors, windows, smoke / heat) is US$1,000 and around US$50 / month for monitoring.. Securing with biometrics adds cost to the basic wiring and monthly monitoring. As an example, the ACTA-1K-FP[69] fingerprint reader with network capabilities retails at US$1,190 (one per entryway).
4	Protect data backups	30 days for installation and setup (not full time)	Barracuda Backup Server Model 690[70] (cloud-based backup storage with on-site appliance): US$7857. Includes 4TB storage.

Priority	Recommendation	Time Estimate	Dollar Estimate (August, 2010)
5	Segment your network	For a small network with separate lines for voice and data, approximately 10 days for network setup. Includes researching your network requirements and setting up routing for each virtual LAN (VLAN).	Dell 6224 24-port switch[71] is available from approximately US$1,300 new. For each VLAN, a separate router must be purchased and thus the cost extends by one server and operating system license.
6	Invest in application-level gateways	Initially, 1 day to setup "basic" (no implementation of corporate security policy beyond simple Web filtering). With a full implementation that requires client systems to meet security requirements, this can easily be a multi-month effort.	Microsoft Threat Management Gateway 2010[72] retails for US$5,999 for an enterprise license (per processor) and US$1,029 for the operating system. This is sufficient for organizations with up to 750 users depending on the power of the server used.
7	Disable unnecessary network traffic	Initially, 2 days for the network administrator to research and apply rules to the physical network switches (small network). This is an ongoing effort as network traffic is a function of installed system requirements.	No direct cost, although configuration can be time-consuming. Cisco Corporation publishes an online guide[73] for disabling / enabling port traffic; each network switch's documentation will be different.
8	Enable strong authentication	Initially, 2 days for the network administrator to setup network policy objects to require either smartcard or biometrics authorization. For biometrics, another week to ensure that all readings are taken.	Most modern operating systems come with the ability to use standard smartcard / biometrics readers automatically. The cost comes with the readers for each machine. For example, the Precise Biometrics 250 MC[74] integrated smartcard / fingerprint reader retails at

Priority	Recommendation	Time Estimate	Dollar Estimate (August, 2010)
			US$362.
9	Establish single-sign on internally	Unable to calculate. For each application (such as corporate Web portal), the appropriate solution must be found to integrate with the smartcard / biometric reader chosen for strong authentication.	Depending on the application, no direct additional cost (although the configuration may be time-consuming). As an example, Oracle Corporation offers full documentation on integrating their Oracle Access Manager[75] (which would be used to enable SSO to Oracle Financials internally).
10	Federate with partners	An ongoing, multi-month effort based on integrating the access federation with business drivers. In our case, integration with Army Knowledge Online requires at least a 60-day period to prepare the request, submit to AKO, and wait for access.	No direct cost (federation capabilities are provided with most server operating systems); approximately one local server must be dedicated to running federation software.

Table 4: List of recommendations with time and cost estimates

Reference List

Arbitron Incorporated and Edison Research. “The Infinite Dial 2010: Digital Platforms and the Future of Radio.” 2010. Available via free registration from the Arbitron Web site: http://www.arbitron.com/study/digital_radio_study.asp (accessed: July 27, 2010).

Bell, D. Elliott, and Leonard J. LaPadula. “Secure Computer Systems: Mathematical Foundations” in *MITRE Technical Report 2547, Volume I*, March 1, 1973.

Bosworth, Seymour, M.E. Kabay, Eric Whyne, eds., *Computer Security Handbook: Volume 1, 4th ed.* Hoboken, NJ: John Wiley & Sons, Inc., 2009.

Cooper, Mex. “Hot mail: stolen email passwords appear online.” *The Sydney Morning Herald*, October 6, 2009. Published online at http://www.smh.com.au/technology/security/hot-mail-stolen-email-passwords-appear-online-20091006-gjq9.html (accessed: July 25, 2010).

Department of Defense. “SSO Login Help.” *Data Services Environment*, https://metadata.dod.mil/mdr/help.htm?page=sso (accessed: August 12, 2010)

Freehling, Bill. “'Flash crash' result of a cyber attack?” *Fredericksburg.com*, June 13, 2010. Published online at http://fredericksburg.com/News/FLS/2010/062010/06132010/553954#axzz0uGYjGVF1 (accessed: August 19, 2010).

Forné, Jordi, Francisca Hinarejos, Andre´s Marín, Florina Almenárez, Javier Lopez, Jose A. Montenegro, Marc Lacoste, and Daniel Díaz. "Pervasive authentication and authorization infrastructures for mobile users.” *Computers and Security* 29 (2010), pp. 501 – 514.

Gill, Charlotte and Jonathon Weinberg. “Online black market where stolen bank passwords sell for as little as 40.,” *The London Daily Mail*, April 14, 2009. Published online at http://www.dailymail.co.uk/news/article-1169777/Online-black-market-stolen-bank-passwords-sell-little-40p.html (accessed: August 7, 2010).

Han, Yan. “On the Clouds: A New Way of Computing.”*Information Technology and Libraries* (June, 2010), pp. 87 – 92.

Harris, Shon. *CISSP All-in-One (AIO), 4th ed*. New York: McGraw-Hill, 2007.

Howe, Denis. *Dictionary.com*. http://dictionary.reference.com/browse/information island (accessed: August 8, 2010).

Headquarters, Department of the Army. “Information Assurance.” *Information Management (Army Regulation 25-2).* Washington, DC: March 23, 2009. Retrieved August 7, 2010, from the US Army Web site: http://www.army.mil/usapa/epubs/pdf/r25_2.pdf.

Ivory, Sean C. “Whisper32 Password Manager.” Published online at http://www.ivory.org/oldwebsite/whisper.html (accessed: August 16, 2010).

McCullen, Sean C. “Millville student charged with stealing $35,000 in computer equipment from school.” *New Jersey On-Line LLC*. Published online at http://www.nj.com/cumberland/index.ssf/2010/06/millville_student_charged_with_1.html (accessed: August 7, 2010).

Nadalin, Anthony Marc Goodner, Martin Gudgin, Abbie Barbir, and Hans Granqvist. “WS-Trust 1.3.” *OASIS*, http://docs.oasis-open.org/ws-sx/ws-trust/200512 (accessed: August 17, 2010).

O’Farrel, Peggy. “Missing records on stolen laptop from Cincinnati Children's Hospital.” *Cincinnati.com*. Published online at http://news.cincinnati.com/article/20100528/NEWS010701/5290339/Missing-records-on-stolen-laptop-from-Cincinnati-Children-s-Hospital (accessed: August 3, 2010).

Page, Lewis. “MoD: We lost 87 classified USB sticks since 2003.” *TheRegister.com.* July 18, 2008. Published online at http://www.theregister.co.uk/2008/07/18/mod_secret_usb_sticks/.

Pfleeger, Charles P., Shari Lawrence Pfleeger. *Security in Computing, 3rd ed.* (Upper Saddle River, NJ: Prentice Hall, 2003).

Marro, Guillermo Mario. 2003, “Attacks at the Data Link Layer.” *Master's thesis*, University of California (Davis).

Microsoft Corporation. “Response to Inaccurate Crypto-Gram Article on VeriSign Certificates.” *Microsoft Technet*, http://technet.microsoft.com/en-us/library/cc751324.aspx (accessed: August 13, 2010).

Millán, Gabriel López, Manuel Gil Pérez, Gregorio Martínez Pérez, and Antonio F. Gómez Skarmeta. “PKI-based trust management in inter-domain scenarios.” *Computers & Security* 29 (2010), pp. 278-290.

Mills, Elinor. “Soldier leaked Google attack investigation details, hacker says.” *cnet news*, June 12, 2010. Published online at http://news.cnet.com/8301-27080_3-20007549-245.html (accessed: August 18, 2010).

Muskal, Michael. “Former NSA official indicted for allegedly leaking classified information to reporter.” *The Los Angeles Times (US Edition)*, April 15, 2010. Published online at http://latimesblogs.latimes.com/dcnow/2010/04/exnsa-official-indicted-for-allegedly-leaking-classified-data-to-reporter.html (accessed: August 17, 2010).

National Institute of Standards and Technology. “Minimum Security Requirements for Federal Information and Information Systems.” *Federal Information Processing Standards (Publication 200)*. Retrieved August 9, 2010, from NIST Web site: http://csrc.nist.gov/publications/fips/fips200/FIPS-200-final-march.pdf.

National Institute of Standards and Technology. “Recommended Security Controls for Federal Information Systems and Organizations.” *Information Security (Special Publication 800-53, Revision 3)*. Retrieved August 10, 2010, from NIST Web site: http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final_updated-errata_05-01-2010.pdf.

Potter, Chris and Andrew Beard. “Information Security Breaches Survey 2010 (Technical Report).” *infosecurity (Europe)*. PricewaterhouseCoopers: 2010. Available online at http://www.ukmediacentre.pwc.com/Media-Library/PwC-ISBS-report-2010-6bb.aspx (accessed: August 15, 2010).

Potter, Daniel. “Ten Thousand Tennesseans’ Names, Social Security Numbers on Stolen Laptop.” *Nashville Public Radio*, June 9, 2010. Published online at http://wpln.org/?p=18453 (accessed: July 29, 2010).

Reagan, Ronald. “Remarks to Representatives of the Future Farmers of America.” *Speech.* Room 450 of the Old Executive Office Building, Washington, DC, July 28, 1988.

rsmc. “VNC Man in the Middle Exploit Code.” *SecuriTeam*, http://www.securiteam.com/exploits/6S0040A6AW.html (accessed: August 17, 2010).

Skone. “Server Potentially Compromised – c99madshell.” *serverfault*, January 20, 2010, http://serverfault.com/questions/104721/server-potentially-compromised-c99madshell (accessed: August 20, 2010).

Smith, Don. “The challenge of federated identity management.” *Network Security* 9 (2009), pp. 7 - 9.

Stephenson, Peter R., ed. Information Security Essentials: Section 1. Auerbach Publishing, ISBN 978-1-4398-0030-0, 2009.

Tapper, Jake, Huma Khan, Alice Gomstyn, and Daniel Arnall. “Wall Street Nervous After 'Flash Crash' Thursday.” *Good Morning America*, May 7, 2010.

Tate, Ryan. “Apple's Worst Security Breach: 114,000 iPad Owners Exposed.” *ValleyWag.com*, June 9, 2010, http://gawker.com/5559346/apples-worst-security-breach-114000-ipad-owners-exposed (accessed: August 15, 2010).

Trustees of the Leland Stanford Junior University. “OU Design Suggestions.” *Stanford Windows Infrastructure*, http://windows.stanford.edu/docs/OUDesign.htm (accessed: August 3, 2010).

Wauters, Robin. “Another Security Tip For Twitter: Don't Use "Password" As Your Server Password.” *TechCrunch.com*, http://techcrunch.com/2009/07/15/another-security-tip-for-twitter-dont-use-password-as-your-password/ (accessed: August 8, 2010).

Wiederhold, Gio. "Acoustic Modem." *Computer History Exhibits Photo*. Stanford Computer Science History Display, http://infolab.stanford.edu/pub/voy/museum/pictures/display/index.htm (accessed: August 9, 2010).

---

[1] "Security Transcends Technology" is a registered trademark of the International Security Certification Consortium (ISC^2), https://www.isc2.org/. This organization provides the highly desirable Certified Information System Security Professional (CISSP) certification, the gold standard for certification in the field of Information Assurance.

[2] Seymour Bosworth, M.E. Kabay, Eric Whyne, eds., "Chapter 44.2.1: Security Policy Guidelines," *Computer Security Handbook: Volume 1*, 4th ed. (Hoboken, NJ: John Wiley & Sons, Inc., 2009), pg. 1148. Dr. Kabay's definition points out that without a foundational security policy, it is impossible for an organization to show a meaningful due diligence effort.

[3] D. Elliott Bell and Leonard J. LaPadula, "Secure Computer Systems: Mathematical Foundations," *MITRE Technical Report 2547, Volume I* (March 1, 1973). Available online at http://www.albany.edu/acc/courses/ia/classics/belllapadula1.pdf (accessed: July 31, 2010). Bell – LaPadula (BLP) defines a mathematical data security model that guarantees data confidentiality in all system states (when properly implemented). The BLP is the most widely recognized model in existence.

[4] Charles P. Pfleeger and Shari Lawrence Pfleeger, *Security in Computing*, 3rd ed. (Upper Saddle River, NJ: Prentice Hall, 2003), pg. 10. Dr. Pfleeger is widely credited with the first mention of the term CIA Triad in his first edition of this book (same publisher, dated 1989).

[5] Inductive reasoning, "Dictionary.com," *The Free On-line Dictionary of Computing*, Denis Howe, http://dictionary.reference.com/browse/information island (accessed: August 8, 2010). Definition: "reasoning from detailed facts to general principles."

[6] Bosworth et. al., *CSH,* "Chapter 3.1: Proposal for a new Information Security Framework." The chapter cited here is a reprint from Mr. Donn Parker's chapter "A New Framework for Information Security" in his book *Fighting Computer Crime* (New York: John Wiley & Sons, 1998).

[7] Sean C. McCullen, "Millville student charged with stealing $35,000 in computer equipment from school," *New Jersey On-Line LLC*, June 15, 2010. Published online at http://www.nj.com/cumberland/index.ssf/2010/06/millville_student_charged_with_1.html (accessed: August 7, 2010).

[8] Peggy O'Farrel, "Missing records on stolen laptop from Cincinnati Children's Hospital," *Cincinnati.com*, May 28, 2010. Published online at http://news.cincinnati.com/article/20100528/NEWS010701/5290339/Missing-records-on-stolen-laptop-from-Cincinnati-Children-s-Hospital (accessed: August 3, 2010).

[9] spokeo.com, "Not your grandma's phonebook," http://www.spokeo.com/. This Web site performs a number of Web searches ("crawls") to extract personal data of all sorts from various online "public" databases (specific

databases are not mentioned, no doubt from the corporate concern of outrage and backlash from the online community). Although spokeo.com offers "opt-out" capabilities for individual users, the danger is that information which people believe to be private is (technically, at least) actually in the public domain. The author can attest to the fact that this search engine identified his race, neighborhood makeup, number and names and *pictures* of children, estimated income and net worth, interests, and current educational endeavors. This is a frightening invasion of personal privacy and we highly support FTC complaints to force spokeo.com to cease and desist from their unethical business practices.

[10] Daniel Potter, "Ten Thousand Tennesseans' Names, Social Security Numbers on Stolen Laptop," Nashville Public Radio, June 9, 2010. Published online at http://wpln.org/?p=18453 (accessed: July 29, 2010).

[11] On the Web site *newegg.com* (http://www.newegg.com), one can perform a search for laptops based on a price range. As of August 12, 2010, a brand-new Lenovo laptop with 4GB of memory, a 250GB hard disk, and a dual-core processor was listed at $399.00.

[12] Bosworth et. al., *CSH*, page 105. The Parkerian Hexad quoted in our text aligns the classic CIA Triad with extensions: Availability with Utility (usability and usefulness), Integrity with Authenticity (completeness and validness), and Confidentiality with Possession (secrecy and control). The physical security concerns raised in this section primarily addresses the Confidentiality and Possession elements of the Hexad.

[13] In the interest of full disclosure, the phrase "the gift that keeps on giving" was registered on February 15, 1927 for "G & S: TALKING MACHINES AND PARTS THEREOF." The reader can visit the US Patent and Trademark Office at http://tess2.uspto.gov/ to search for this term as well.

[14] Ronald Reagan, "Remarks to Representatives of the Future Farmers of America" (speech, Room 450 of the Old Executive Office Building, Washington, DC, July 28, 1988).

[15] National Institute of Standards and Technology, "Minimum Security Requirements for Federal Information and Information Systems," *Federal Information Processing Standards (Publication 200)*. Retrieved August 9, 2010, from NIST Web site: http://csrc.nist.gov/publications/fips/fips200/FIPS-200-final-march.pdf. We reference the specific sections quoted from this source in our text.

[16] National Institute of Standards and Technology, "Recommended Security Controls for Federal Information Systems and Organizations," *Information Security (Special Publication 800-53, Revision 3)*. Retrieved August 10, 2010, from NIST Web site: http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final_updated-errata_05-01-2010.pdf. This edition includes all updates as of May 1, 2010. We reference the specific sections from this source in our text.

[17] Bosworth et. al., *CSH,* "Chapter 44.2.1: Security Policy Guidelines." Dr. Kabay presents a thorough overview of a security policy's place in the organization. Truly, without a security policy "it is impossible to demonstrate due diligence in the protection of corporate assets" (pg. 1148).

[18] Headquarters, Department of the Army, "Information Assurance," *Information Management (Army Regulation 25-2)*, Washington, DC: March 23, 2009. Retrieved August 7, 2010, from the US Army Web site: http://www.army.mil/usapa/epubs/pdf/r25_2.pdf. Specifically, we reference the reader to chapter 4, section 16 "Protection Requirements" that specifically addresses security requirements for all information systems equipment and media processing, handling, or storing classified data.

[19] Lewis Page, "MoD: We lost 87 classified USB sticks since 2003," *TheRegister.com*, July 18, 2008. Published online at http://www.theregister.co.uk/2008/07/18/mod_secret_usb_sticks/. In this article, the British Ministry of Defense (MOD) reported that they had lost 87 classified USB sticks since 2003. The losses were due either to theft or to simple misplacement, and the Ministry reported that "nothing of significance" was actually lost; the affected classification range on the lost data included "Restricted" (the lowest form of classification) to "Secret" (the next-to-highest classification). The reasoning: The Ministry acknowledged that data is routinely over-classified with the result that "Secret" information is generally of little to no significance. However, the actual number of lost USB sticks "continue to be adjusted" so that the actual data loss is not known. The data classification issues raised by this event themselves would make the basis for an excellent paper.

[20] Shon Harris, "Information Security and Risk Management," *CISSP All-in-One (AIO)*, 4th ed., (New York: McGraw-Hill, 2007), pg. 57. Harris defines the three types of controls (Administrative, Technical, and Physical).

[21] Mex Cooper, "Hot mail: stolen email passwords appear online," *The Sydney Morning Herald*, October 6, 2009. Published online at http://www.smh.com.au/technology/security/hot-mail-stolen-email-passwords-appear-online-20091006-gjq9.html (accessed: July 25, 2010).

[22] Harris, *AIO*, pg. 265. "Phishing" is a type of social engineering attack (one that relies on a victim's innate sense of trust).

[23] Charlotte Gill and Jonathon Weinberg, "Online black market where stolen bank passwords sell for as little as 40p," *The London Daily Mail*, April 14, 2009. Published online at http://www.dailymail.co.uk/news/article-1169777/Online-black-market-stolen-bank-passwords-sell-little-40p.html (accessed: August 7, 2010).

[24] Harris, *AIO*, pg. 245. Harris uses the term "keystroke monitoring" but the meaning is the same.

[25] Robin Wauters, "Another Security Tip For Twitter: Don't Use "Password" As Your Server Password," *TechCrunch.com*, http://techcrunch.com/2009/07/15/another-security-tip-for-twitter-dont-use-password-as-your-password/ (accessed: August 8, 2010).

[26] Yan Han, "On the Clouds: A New Way of Computing,"*Information Technology and Libraries* (June, 2010), pp. 87 – 92. In this article, Han reports on his findings while implementing a cloud computing environment for the University of Arizona at Tucson.

[27] Bosworth et. al., *CSH,* "Chapter 28.3: Identification and Authentication," pg. 787. This section provides a stinging indictment on passwords and lists nine password-usage risks that work directly against the goal of overall security.

[28] Shaun Ivory, "Whisper32 Password Manager," http://www.ivory.org/oldwebsite/whisper.html (accessed: August 16, 2010). This older (2004) password manager allows the storage and retrieval of an unlimited number of passwords. However, numerous free password manager applications exist; the reader can perform a Web search for "free password managers" to scan through the thousands of available applications.

[29] Peter R. Stephenson, ed. Information Security Essentials: Section 1 (Auerbach Publishing, ISBN 978-1-4398-0030-0, 2009), pg. 308. Retrieved on June 6, 2010 from https://norwich.angellearning.com/AngelUploads/Content/MSIA_2_0/_assoc/msia_s1/msia_s1_readings/ise_section_1_et.pdf. This reading references a sample Internet Usage Policy, and an employee's password

responsibilities certainly apply here. The larger goal here concentrates on creating and applying meaningful policies to ensure that the overall security policy for an organization has the best chance of succeeding.

[30] Bosworth et. al., *CSH,* "Chapter 28.2: Identification and Authentication," pg. 784.

[31] Harris, *AIO*, pg. 190. A *passphrase* is a sequence of characters that is longer than a password (thus a "phrase") an, in some cases, takes the place of a password during an authentication process.

[32] Bosworth et. al., *CSH*, "Chapter 29.6.2: Biometrics Authentication," pp. 818 - 820. The author covers several problems with biometrics integration including cost and cultural issues.

[33] Ibid., "Chapter 29.5: Biometric Authentication," pg. 817.

[34] Ibid.., "Chapter 28.3.10: Authentication Using Recognition of Symbols," pg. 794.

[35] Ibid., "Chapter 7.5.3: Encryption," pg. 237.

[36] Ibid., "Chapter 37.2: PKI and Certificate Authorities," pg. 1012. The classic problem with public key exchange goes like this: Alice and Bob wish to exchange public keys to have a private (encrypted) conversation. Consider if Chuck intercepts the transmission ("man-in-the-middle") – when Alice requests Bob's public key, Chuck actually provides his (and performs the same malicious operation for Bob). The net result: Chuck can understand the entire conversation between Alice and Bob despite the fact that all network messages are fully encrypted. To avoid this problem, Alice must be assured that she is actually receiving Bob's public key. PKIs exist to solve this problem. When Alice receives Bob's purported public key, she can verify that it really did come from Bob by verifying the key with the Certificate Authority component of the PKI.

[37] Pfleeger, *Security in Computing*, pp. 666 – 684. Starting with Diffie-Hellman in 1976, public key encryption systems have gone through several evolutions. Of particular interest is the possible use of quantum cryptography (see *CSH* pg. 243, Pfleeger pg. 683), which attempts to solve the problem of generating a truly random set of numbers for the strongest possible encryption.

[38] SSL Certification Request Identity Verification typically has multiple levels of support depending on the trust level that clients receiving that certificate can have in its authenticity. In our paper, we reference the Certificate Authority (CA) GoDaddy.com but the points we raise here apply to all commercial CAs. All CAs offer multiple levels of identity verification. GoDaddy.com in particular provides three levels (see https://www.godaddy.com/gdshop/ssl/turbo.asp for more information; accessed August 17, 2010). The lowest level (Standard SSL) is suitable for individuals and GoDaddy.com warrants the certificate for losses of up US$10,000.00 (a very low amount). For the highest levels of identity verification, GoDaddy.com provides up to US$250,000.00 of warranty coverage, but requires the requesting organization to prove their existence and their overall business legitimacy.

[39] Gabriel López Millán, Manuel Gil Pérez, Gregorio Martínez Pérez, and Antonio F. Gómez Skarmeta, "PKI-based trust management in inter-domain scenarios," *Computers & Security* 29 (2010), pp. 280-281. Available from the Journal homepage at http://www.elsevier.com/locate/cose. Our notes on PKI infrastructure are based based upon Section 3.1: "PKI Service Requirements" from this source.

[40] Skone, "Server Potentially Compromised – c99madshell," *serverfault*, January 20, 2010, http://serverfault.com/questions/104721/server-potentially-compromised-c99madshell (accessed: August 20,

2010). This post describes a real-world event where a Web server hosting a legacy site was hacked and the private key potentially compromised.

[41] Trustees of the Leland Stanford Junior University, "OU Design Suggestions," *Stanford Windows Infrastructure*, http://windows.stanford.edu/docs/OUDesign.htm (accessed: August 3, 2010).

[42] Bill Freehling, "'Flash crash' result of a cyber attack?," *Fredericksburg.com*, June 13, 2010. Published online at http://fredericksburg.com/News/FLS/2010/062010/06132010/553954#axzz0uGYjGVF1 (accessed: August 19, 2010).

[43] Jake Tapper, Huma Khan, Alice Gomstyn, and Daniel Arnall, "Wall Street Nervous After 'Flash Crash' Thursday," *Good Morning America*, May 7, 2010.

[44] Ryan Tate, "Apple's Worst Security Breach: 114,000 iPad Owners Exposed," *ValleyWag.com*, June 9, 2010, http://gawker.com/5559346/apples-worst-security-breach-114000-ipad-owners-exposed (accessed: August 15, 2010).

[45] GoatSe Security is a somewhat mysterious group whose homepage we will not publish here.

[46] Chris Potter and Andrew Beard, "Information Security Breaches Survey 2010 (Technical Report)," for *infosecurity (Europe)* (PricewaterhouseCoopers: 2010). Available online at http://www.ukmediacentre.pwc.com/Media-Library/PwC-ISBS-report-2010-6bb.aspx (accessed: August 15, 2010).

[47] Pfleeger, *Security in Computing*, pg. 373. The seven layers of the TCP/IP model are actually defined in the International Organization of Standards' source document *ISO Publication 7498-2: 1989*.

[48] Guillermo Mario Marro, 2003, *Attacks at the Data Link Layer*, Master's thesis, University of California (Davis).

[49] Harris, *AIO*, pp. 1080 – 1087. Harris covers numerous tests here, many (such as Session Hijacking) are directly applicable to Layer 3 in the OSI protocol stack.

[50] Telnet (TCP port 23) is an older remote administration technology in which passwords are sent in clear-text (unprotected). Modern computer networks should never use Telnet and instead rely on more secure services such as Secure Shell (SSH, TCP port 22) or Virtual Private Networks (*VPN*; the TCP ports vary based on the VPN "flavor;" the UDP port for the "Internet Security Association and Key Management" Protocol or ISAKMP, is 500).

[51] Microsoft Corporation, "Response to Inaccurate Crypto-Gram Article on VeriSign Certificates," *Microsoft Technet*, http://technet.microsoft.com/en-us/library/cc751324.aspx (accessed: August 13, 2010).

[52] rsmc, "VNC Man in the Middle Exploit Code," *SecuriTeam*, http://www.securiteam.com/exploits/6S0040A6AW.html (accessed: August 17, 2010). This site contains a working C-language source file that demonstrates how to gain unauthorized access to a computer running the VNC remote administration server by intercepting communications from an authorized client. It represents a classic *authenticity* problem that protocols such as SSH and VPN go to great lengths to mitigate.

[53] In the interests of space, we cover only a subsection of network equipment here. Harris (*CISSP*) provides a more full description on pp. 536 – 548.

[54] RJ-45 is the most common type of network cable connection and uses four "twisted-pair" wires to transfer data. Numerous resources exist online describing network cabling in general and Ethernet cabling in particular (as opposed to IBM Token Ring, for example). One not-for-profit resource with a very quick introduction to this topic can be found at http://www.ertyu.org/steven_nikkel/ethernetcables.html.

[55] The IP address given (74.125.157.147) is the network address returned from www.google.com. As with most large organizations concerned with scalability, Google has many physical IP addresses all associated with the common name www.google.com. Thus, issuing a command like "ping www.google.com" (which displays the IP address for Google) can and will return different values at different times.

[56] Gio Wiederhold, "Acoustic Modem," *Computer History Exhibits Photo*, Stanford Computer Science History Display, http://infolab.stanford.edu/pub/voy/museum/pictures/display/index.htm (accessed: August 9, 2010).

[57] Licensing cost of $5,999 for Microsoft's Threat Management Gateway 2010 product verified as of August 18, 2010. Available online at http://www.microsoft.com/forefront/threat-management-gateway/en/us/pricing-licensing.aspx (accessed: August 18, 2010).

[58] "Hooah" is general slang used by Soldiers in the U.S. Army to mean anything except "no." The author has always heard it as "Heard-Understood-Acknowledged" (HUA). No authoritative definition appears to exist, see http://www.cavhooah.com/info/resources/hooah/ for one set of possible meanings.

[59] Elinor Mills, "Soldier leaked Google attack investigation details, hacker says," *cnet news*, June 12, 2010. Published online at http://news.cnet.com/8301-27080_3-20007549-245.html (accessed: August 18, 2010).

[60] WikiLeaks (http://wikileaks.org/) provides an open service where anonymous individuals may report sensitive material to the general public.

[61] Michael Muskal, "Former NSA official indicted for allegedly leaking classified information to reporter," *The Los Angeles Times (US Edition)*, April 15, 2010. Published online at http://latimesblogs.latimes.com/dcnow/2010/04/exnsa-official-indicted-for-allegedly-leaking-classified-data-to-reporter.html (accessed: August 17, 2010).

[62] Arbitron Incorporated and Edison Research, "The Infinite Dial 2010: Digital Platforms and the Future of Radio," 2010. Available via free registration from the Arbitron Web site: http://www.arbitron.com/study/digital_radio_study.asp (accessed: July 27, 2010). This study, which is full of information on the habits and trends of the American public over the last decade, provides provocative insight into the "relentless impact of the Internet." Our quote in this paper is from page 11 of the report.

[63] Department of Defense, "SSO Login Help," *Data Services Environment*, https://metadata.dod.mil/mdr/help.htm?page=sso (accessed: August 12, 2010).

[64] Millán et. al., pg. 279. Both of these statements can be found on that page.

[65] Anthony Nadalin, Marc Goodner, Martin Gudgin, Abbie Barbir, and Hans Granqvist, "WS-Trust 1.3," *OASIS*, http://docs.oasis-open.org/ws-sx/ws-trust/200512 (accessed: August 17, 2010). The link provides access to three technical documents used with WS-Trust solutions: a document with the overall goals and instructions for implementers, and two files defining data structures used by the standard as well as functions that implementers must provide, respectively.

[66] Don Smith, "The challenge of federated identity management," *Network Security* 9 (2009), pg. 8. Smith considers it to be "an annoyance to consumers to maintain their identity information securely on multiple sites" and one appreciates this massive understatement.

[67] Jordi Forné, Francisca Hinarejos, Andre´s Marín, Florina Almenárez, Javier Lopez, Jose A. Montenegro, Marc Lacoste, and Daniel Díaz, "Pervasive authentication and authorization infrastructures for mobile users," *Computers and Security* 29 (2010), pp. 502 - 503. Section 2.2 discusses evidence-based computational trust management, and the paper as a whole provides a framework for a disconnected trust environment that allows new entities to enter a local mobile *ad hoc* network (MANET).

[68] BuyerZone.com editorial staff, "Monitored Alarm Systems – Buyer's Guide," *Yahoo! Small Business*, http://smallbusiness.yahoo.com/r-article-a-40946-m-5-sc-49-monitored_alarm_systems_buyers_guide-i (accessed: August 10, 2010). Although to be taken with a grain of salt based on the non-authoritative source, the pricing range appears to be in line with our own experiences (we opted for a more complex system).

[69] Actatek, "ACTA-1K-FP – Actatek Finger Print Access," *Home Security Store*, http://www.homesecuritystore.com/p-213-acta-1k-fp-actatek-finger-print-access.aspx (accessed: August 10, 2010).

[70] Barracuda Networks, "Cloud Storage with Deduplication," *Barracuda Backup Server Series*, http://www.barracudanetworks.com/ns/products/backup_overview.php. We selected the Model 690 as a good mid-level solution that should support a small business environment.

[71] Dell Corporation, "Dell PowerConnect 6224 Ethernet Switch," *Dell Support*, http://www.dell.com/us/en/enterprise/networking/pwcnt_6224/pd.aspx?refid=pwcnt_6224&cs=555&s=biz (accessed: August 11, 2010).

[72] Microsoft Corporation, "Pricing and Licensing," *Microsoft Forefront Threat Management Gateway 2010*, http://www.microsoft.com/forefront/threat-management-gateway/en/us/pricing-licensing.aspx (accessed: August 12, 2010).

[73] Cisco Systems Corporation, "Troubleshooting, Switch Port and Interface Problems," *Cisco Technical Support*, http://www.cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a008015bfd6.shtml#pois (accessed: August 12, 2010).

[74] PCRush.com, "Precise Biometrics 250 MC Finger Print Reader," http://www.pcrush.com/product/Biometrics/532071/Precise-Biometrics-250-MC-Finger-Print-Reader (accessed: August 12, 2010).

[75] Oracle Corporation, "Oracle Access Manager Integration Guide," *Oracle 10g (10.1.4.0.1)*, http://download.oracle.com/docs/cd/B28196_01/idmanage.1014/b25347/smartcrd.htm#BJEGJGAF (accessed: August 12, 2010).