



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Threat Analysis of Allowing Employee Internet Access

G7799 Gold Certification

Author: Mason Pokladnik, mason@schwanda.cc

Adviser: Stephen Northcutt

Accepted:

Table of Contents

| | |
|--|----|
| 1. Introduction..... | 3 |
| 2. Threats..... | 3 |
| 2.1. Increasing sophistication of malware..... | 4 |
| 2.2. Profit motivation of attackers | 5 |
| 2.3. Detailed attack methods | 5 |
| 2.3.1. Email based attacks | 5 |
| 2.3.2. Web based attacks | 7 |
| 2.3.3. Non web based protocols – P2P, IM, Remote access..... | 11 |
| 2.4. Legal issues | 12 |
| 3. Threat management..... | 12 |
| 3.1. Good controls..... | 13 |
| 3.1.1. Antivirus/Antispyware (AV) | 13 |
| 3.1.2. Intrusion Detection Systems (IDS) | 15 |
| 3.1.3. Gateway protections | 15 |
| 3.1.4. User education | 19 |
| 3.2. Better controls..... | 20 |
| 3.2.1. Intrusion Prevention Systems (IPS)..... | 20 |
| 3.2.2. Patch management | 22 |
| 3.2.3. Defense in depth | 23 |
| 3.2.4. Security as a software selection criteria | 25 |
| 3.3. Best controls..... | 27 |
| 3.3.1. Network architecture | 27 |
| 3.3.2. Darknets/honeypots..... | 28 |
| 3.3.3. Integrity monitoring | 29 |
| 3.3.4. Disposable computers | 30 |
| 4. Some hypothetical numbers | 30 |
| 5. Checking with the 17799 standard..... | 33 |
| 6. Conclusion – Who should get Internet access?..... | 34 |
| 7. Parting Thoughts | 36 |

1. Introduction

The ISO 17799/27001 standard provides a good minimum description of what organizations should be doing to protect themselves, but it should not be the sole focus of your security and audit control design. A better approach is to allow your information-security management-system subcommittees or technical specialists to analyze the threats your organization is likely to face. Then, design your controls around those threats, balancing the cost to mitigate a threat versus the cost of a threat occurring in your environment. Finally, after you have analyzed the threats, you can double check your policies and procedures against a regulatory or management framework, such as ISO17799, SOX, GLBA, HIPPA or PCI.

In order to demonstrate the threat analysis process, I will use the business necessity of allowing employees access to the Internet. Since many employees need access to email and the Internet to complete their jobs, this is a threat that most organizations will have to address.

During the analysis, we will cover four distinct steps that make up the major sections of the paper. You can use these sections along with the Table of Contents to locate a topic.

Step 1 - Section 2 – We will discuss several classes of threats that come along with employees using the Internet.

Step 2 - Section 3 – We will discuss some of the controls an organization can use to mitigate the threats from Section 2. This section is organized into good, better and best categories based on how effective the controls are at dealing with emerging threats.

Step 3 - Section 4 – We will look at the cost of implementing some of the controls from Section 3 in a hypothetical network and attempt to decide if the spending is justified, and if we are spending our budgets and time in the right places.

Step 4 - Section 5 – We will check our analysis against the relevant sections in the ISO17799:2005 standard to see if the threat analysis is comprehensive enough. Since ISO17799 suggests best practices for a wide array of business functions from antivirus to contracts, it may highlight threats we did not consider.

Step 5 – Section 6 – We will attempt to apply the information in your environment and look at the appropriate type of Internet access for a user.

Please note that a prerequisite step to our analysis is to inventory and place a value on the assets you are trying to protect. This prerequisite step is beyond the scope of this paper, but time spent acquiring that information, as well as understanding how your business works will allow you to apply the threat information from this paper more specifically in your environment.

2. Threats

First, let us take a quick look at the overall threat. The original protocols and systems that made up the Internet were designed to run on an open network. That network was resilient against the threats of backhoes and nuclear warheads, but not designed to keep out an attacker. Over the last decade, the attackers have transformed. Originally, the attackers were people who were interested in bragging to their friends about their conquests. Now, the attackers include criminal enterprises, like the infamous

Russian Business Network¹, who are interested in making a profit. While there are plenty of opportunists – affectionately referred to as “script kiddies” - who must use tools created by others, there now are violent organized crime groups who are combining the anonymity the Internet provides with the difficulty of prosecuting international crimes to create a unique money making opportunity. One estimate is that there are billions in potential profits from being able to control other peoples systems: then exploiting those systems or the information on them.²

The classic method of attackers is to scan the Internet for vulnerable systems and then attempt to compromise them. Now attackers have a complimentary attack vector. Why should an attacker go out and look for vulnerable hosts when the hosts can come to the attacker? An attacker can place exploits on the Internet and attract people to them using Spam, buying advertisements on web pages, and other methods. This way the computers come to the attacker without interference from firewalls or network scanning.

2.1. Increasing sophistication of malware

Malware is a term with many different definitions and names like spyware, viruses, adware, trojans, etc. I use it here as an all-encompassing term for software that may spy on you, allow others to control your computer, display advertisements or perform any number of other activities that fall into the category of your computers working for others. As with most software, the original incarnations were relatively simple programs compared to today. Early malware (viruses) would attempt to replicate itself and was occasionally destructive. The authors of this software may have written them as pranks or as a way to gain notoriety among their fellow virus writers. As more and more computers arrived on the Internet, especially those with high-speed connections, people began to notice that the same old crimes committed in the real world could also be perpetrated online. Over time, the attacks have become much more sophisticated in nature. For example, recent malware usually actively attempts to protect itself from being removed. Early simple tricks included multiple processes that would restart one another if one died and stopping antivirus services. Today’s malware can subvert your computer in ways that make it nearly impossible to remove without reinstalling the operating system. Some malware authors have even adopted reusable code to aid in quicker development. They also use tiered applications to both distribute the load of spreading their creations and to provide redundancy so that when incident handlers do manage to shut down a website, their networks can continue to operate. Many organizations cannot replicate these skills with their own internal development efforts! Obviously, something has changed to bring all of this worldwide talent into developing this complicated software. As you already know, one of the single greatest motivators is money.

¹ <http://www.washingtonpost.com/wp-dyn/content/article/2007/10/12/AR2007101202461.html>, Visited 12/3/07

² <http://www.informationweek.com/shared/printableArticle.jhtml?articleID=201203030>, Visited 10/25/2007

2.2. Profit motivation of attackers

Both the con artist and criminal organizations have adapted to the Internet. The pyramid scheme, get-rich-quick scams, and extortion have online counterparts and the bad guys seem to be dreaming up new schemes all of the time. As an organization, you could be subject to a distributed denial of service attack against your website in an attempt to extort money from you, directly, or to cause your customers to go somewhere else. Your employees could be subject to identity theft by using an infected computer or visiting websites that steal passwords from web browsers. The “crown jewels” of your company may even be under attack. If you have a product or a process that other people want, *you are going to be attacked*. The FBI has been telling anyone who will listen through its Counter Intelligence Domain program³ that foreign intelligence agents are actively attempting to get at your intellectual property. You may not just be subjected to commercial espionage either. There is considerable evidence that foreign governments have already probed U.S. government agencies and contractors for weaknesses. If you have something they are interested in stealing, you could be next. Just search for the codeword Titan Rain, to see the types of directed attacks that were happening a few years ago. The attacks today are likely even more subtle such as the subversion of certain Greek cell phone switches during the run up to the Athens Olympics.⁴ Certain Chinese IP addresses have been implicated in recent attacks against companies in the UK⁵, the German government⁶ and the US Department of Defense⁷. You can be certain that other people are doing their own cost/benefit analysis and figuring out it is much cheaper to hijack your intellectual property than it is to recreate it.

2.3. Detailed attack methods

In the following sections, I will cover some of the most prevalent and recent attack vectors. By the end of the analysis, you should be able to evaluate how capable your controls are at preventing or detecting these classes of attacks and determine if improvements in your defenses are necessary. Since the threats are always changing, you should periodically assess your controls in this fashion, or they will become increasingly less effective over time.

2.3.1. Email based attacks

Nearly everyone in corporate America has an email address. It is simply one of the easiest ways to disseminate information in a timely fashion and allow quick intercompany communications. A brief look at most companies’ websites will give an attacker a list of email addresses that they can use to launch

³ <http://www.fbi.gov/hq/ci/domain.htm>, Visited 10/25/2007

⁴ <http://www.spectrum.ieee.org/print/5280>, Visited 12/6/2007

⁵ http://business.timesonline.co.uk/tol/business/industry_sectors/technology/article2980250.ece, Visited 12/7/2007

⁶ <http://www.spiegel.de/international/world/0,1518,502169,00.html>, Visited 8/27/2007

⁷ <http://www.ft.com/cms/s/0/9dba9ba2-5a3b-11dc-9bcd-0000779fd2ac.html>, Visited 9/5/07

phishing attacks, send an infected attachment, or link to a drive-by download site. Just because your CEO's email address may not be posted on your website does not mean they will not be targeted. Most companies follow a standard format for email addresses such as firstname.lastname@company.com. Combine that information with your website, public filings with the SEC, press releases, newsgroup postings, etc. and you can easily guess legitimate email addresses, including those of senior managers.

Once the attackers have identified a target, they can proceed in several ways. The first email borne attacks came against the mail transfer agent (MTA) programs that are responsible for delivering emails to the destination computer. Sendmail was one of the first and most popular MTAs in the early Internet, and in 1988, an early bug⁸ in a debug routine was used to help the Morris worm⁹ propagate throughout a significant portion of the Internet. One would hope in the almost 20 years since then that this type of issue had been eradicated; however, even modern MTAs like Microsoft Exchange have recently had remotely exploitable bugs.¹⁰

Other attacks focus on email clients, such as Microsoft Outlook or Lotus Notes. Just like any application that must parse input before processing it, an email client can have vulnerabilities too. In programming you assume that the more complicated a program is, and the more features it supports, the more likely it is that something will go wrong. As we added in images, encoded attachments and other complications on top of the text of email messages, we created new opportunities for mistakes. In 2003, just viewing a crafted html formatted email was enough to allow you to run any command you wanted on a vulnerable computer running Outlook.¹¹

In 1999, Melissa (the first successful attachment based worm) hit the Internet and raised the consciousness of many IT departments about the threat presented by attachments. Before then, some organizations realized that certain types of attachments such as executable files did not need to be sent by email. After Melissa, organizations had to adapt to a new threat where even "data" files could carry malicious content. As a result, now all attachments need to be scanned. In most cases, executable attachments can be removed from emails, but images, office documents, pdf files, zip files and other attachments could all be holding malicious content or intellectual property, and email gateways are the natural place to look for it.

Using social engineering techniques in email messages has always been a successful attack method. Recent attempts have included sending targeted emails to executives claiming to be from the Better Business Bureau or IRS requesting the recipient download a form describing a complaint filed against the company.¹² The Storm worm fooled millions of people by asking them to view a fake e-card. In the case

⁸ <http://www.securityfocus.com/bid/1/exploit>, Visited 12/7/2007

⁹ http://en.wikipedia.org/wiki/Morris_worm, Visited 12/7/2007

¹⁰ <http://www.securityfocus.com/bid/23809>, Visited 9/15/2007

¹¹ <http://www.securityfocus.com/bid/6923>, Visited 9/15/2007

¹² <http://isc.sans.org/diary.html?storyid=2979>, Visited 9/15/2007

of the Storm worm, just clicking on the link was enough to launch several exploit attempts against your web browser even if your user was wise enough not to download and run an exe file from the site.¹³ Since the bad guys have access to Spam filters and antivirus programs too, they will always be able to tweak their messages so that some get through Spam filters. Therefore, filtering will always be a step behind a social engineering attack.

2.3.2. Web based attacks

The web browser is a complex application that can be extended through entire programming languages such as Javascript, Java applets, Flash applications and ActiveX controls. While these technologies make for a richer web experience for the end user, each of them is a new attack vector in a program running inside your firewall. You may have heard several commentators expressing their views that the network perimeter is no longer capable of preventing attackers from getting inside.¹⁴ Filtering network traffic cannot protect against all attacks, and laptops on the road usually connect directly to the Internet. In either case, once a web page and its associated content has made it to the client machine, whatever controls you have set up on that machine are your last line of defense.

Attackers are constantly advancing the technologies behind the web-based attack. Some groups have gone as far as selling prebuilt frameworks for turning a compromised web server into an attack platform. One package for sale, called Mpack, was recently installed on web servers owned by an Italian web hosting company. Once installed, the attackers were able to choose from a buffet of prebuilt exploits using a web based administrative interface. Mpack then appended the attack script to every page sent out by the web servers, so anyone who visited one of the thousands of sites hosted on those computers came under attack.

The first web-based attacks went after the browser itself, and all browsers have had their share of remotely exploitable bugs. As the browsers have become more secure – although how much more secure is up for debate – the attacks have begun to focus on the helper programs and add-ons that people install to support new content. In the next several sections we will cover some of the most often attacked helper programs and browser features.

2.3.2.1. Javascript

While many helper programs have to be consciously downloaded and added to your browser, almost all modern browsers ship with an implementation of Javascript. Javascript and AJAX (asynchronous Javascript and XML) are credited with allowing the creation of today's highly interactive, web-based applications such as Google Maps. These "Web 2.0" applications demonstrate just how powerful even an interpreted language can be. That same power can be used to change context menus in your browser, resize windows, and open new content such as an IFrame with an exploit created specifically for your browser and operating system. Attackers can use dynamic code obfuscation routines to hide

¹³ <http://isc.sans.org/diary.html?storyid=3063>, Visited 9/15/2007

¹⁴ <http://www.microsoft.com/technet/community/columns/secmgmt/sm0907.msp>, Visited 12/7/2007

the real intention of their scripts from users and intrusion detection sensors. By altering the code nearly every time someone visits a page, attackers make it extremely difficult to write a signature to detect malicious content. Your web browser, on the other hand, happily runs through the routines necessary to expose the real code and run it.

Some really creative people are exploring what you can do with this technology. Recently, a tool named Jikto was released that could be used to turn your web browser into a port scanner for a remote attacker. While the code for Jikto was not originally going to be made public, it was accidentally released¹⁵ and can be found on the Internet. Since you can tunnel IP packets in nearly anything, – even DNS – it is just a matter of time until someone writes a script to provide full network access by creating a tunnel through your web browser to your network.¹⁶ In such a case, if you could not detect and eliminate this traffic, you might as well not even bother having a firewall.

2.3.2.2. Flash

The flash plug-in powers multimedia applications such as streaming video on youtube.com and online games. It also has its own scripting language that is currently being abused to redirect web browsers to malicious web sites. Attackers are buying advertisements on websites and then, inside their flash ads, redirecting users to malicious sites like winantispyware .com (link purposely broken) which use Javascript to launch exploits against your machine. If vulnerable, your machine becomes infected with a wide array of malware.

Another “feature” of the flash plug-in is its ability to access your webcam and microphone. This is disabled by default, but if you want to double check your settings, right click on any flash item in your browser and choose settings.

2.3.2.3. Java

The Java programming language was created at Sun in order to provide a write once/run anywhere environment for programmers. Java programs – known as applets - need a runtime component that is specific to each operating system it runs on. The runtime is supposed to execute applets in an isolated environment, called a sandbox, where it can do little harm to the system. The sandbox concept has no middle ground. You are either in it or you are not, and all a user has to do is click on one dialog box to give an applet full access to the system. That level of access on a Windows system also gives it access to the registry, which is where Java security settings, Internet Explorer security settings and many others are stored. Letting any malicious Java program out of the sandbox means you are allowing it to make the decisions. It could configure your registry to allow any website of its choice, or any publisher of its choice to run applets with full access to your system instead of inside the restricted environment.

¹⁵ <http://portal.spidynamics.com/blogs/spilabs/archive/2007/04/02/Jikto-in-the-wild.aspx>, Visited 12/7/2007

¹⁶ http://www.doxpara.com/slides/DMK_BO2K7_Web.ppt, Visited 12/7/2007

Then again, a more subtle attack may not even require leaving the sandbox. As the “pure evil java popup” demonstrates, you could replace a person’s browser window with your own re-creation.¹⁷ If the attacker knew enough about the target, they could collect login or personal information, and Java’s cross platform design means it works on Windows, Mac, Linux and many other places with a browser and Java runtime. At the time this paper was written, this is still a current vulnerability. This particular threat is mitigated by the fact that the user still has to navigate to a malicious webpage to be affected.

2.3.2.4. ActiveX

Since we started putting computers on the Internet, there have been many poor design decisions that led to security problems. The Internet Explorer/ActiveX combination has to be among the worst of them. In an attempt to offer an easier development environment and less rigid plug-in architecture than the Netscape browser, Microsoft unwittingly opened a flood of potential problems. Now, instead of having to worry only about browser vulnerabilities and vulnerabilities for the plug-ins installed on purpose, you now have to worry about every piece of software you install potentially adding a new ActiveX control to your system and marking it safe for scripting. Acrobat Reader, Winzip, FTP programs, Video codecs, and nearly all audio players are just some of the programs currently exploited to install unwanted software on Windows systems. This means your patch management process has to mature beyond turning on automatic updates and walking away. The problem is further exacerbated by the use of Internet Explorer as a rendering engine by other programs. Programs that may have been designed only to run on a local computer are installing ActiveX controls that are now potentially accessible to the Internet. Any ActiveX control registered on the system could possibly be used to attack your system. Look under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Internet Explorer\ActiveX Compatibility in the registry of a Windows computer to get an idea of the scope of the issue. Microsoft has tried to limit this with the default settings of Internet Explorer to only those classes marked “safe for scripting.”¹⁸ However, there have been numerous bugs within Internet Explorer allowing attackers to bypass the security settings or change what security zone they are running under, thus giving them access to all controls installed on the system.

Since Internet Explorer has had the largest market share of any browser for the last several years, many websites have tweaked themselves or implemented features that require you to install an ActiveX control, which makes using an alternative web browser difficult for some users. This means we will continually deal with the consequences of a design from a much less hostile Internet, which valued code reusability over security.

¹⁷ <http://www.hackademix.net/2007/08/07/java-evil-popups/>, Visited 10/5/2007

¹⁸ <http://support.microsoft.com/kb/240797>, Visited 9/14/2007

2.3.2.5. Other add-on programs

The programs discussed so far are just some of the most common add-ons attacked. Recent exploits for Adobe Acrobat,¹⁹ Real Player,²⁰ Quicktime,²¹ etc. highlight that we are deploying new vulnerabilities faster than patches. We may have firewalled ourselves off from the Internet and prevented many direct attacks, but any application that interacts with the Internet can become another way into the “soft chewy center” of most networks.

2.3.2.6. Interface design

The consistent weak link in the security chain does seem to be the user. We humans do not make the best security decisions under easy circumstances.²² So, when a web browser presents a question like “There is a problem with the site’s security certificate. Do you want to continue anyway?” Why are we surprised when they make the wrong decision and end up giving away their credit card information?

Figure 1 – Alert message from Internet Explorer 6 regarding an SSL certificate



Many of the programs we use are not helping the users to make better decisions. Nearly every user when presented with a cryptic error and the question “Do you want to continue?” is going to press yes. This type of problem is typical in host-based intrusion detection systems and many other security

¹⁹ <http://isc.sans.org/diary.html?storyid=3537>, Visited 9/15/2007

²⁰ <http://isc.sans.org/diary.html?storyid=3528>, Visited 9/15/2007

²¹ <http://isc.sans.org/diary.html?storyid=3618>, Visited 9/15/2007

²² “Passwords revealed by sweet deal”, <http://news.bbc.co.uk/1/hi/technology/3639679.stm>, Visited 9/14/2007

products as well. In many cases, administrators may disable security features in products rather than trust an end user to make a decision when confronted with a security related question.

2.3.2.7. Cross site scripting (XSS)

Far from the final threat, but the last one in this section, is cross-site scripting (XSS.) It is unfortunate that actually fixing XSS bugs is often out of the control of the end users of the websites. However, the consequences of those bugs are felt by the end user, such as when clicking on a link can cause you to turn over your login session to another site. If that site happens to be your bank, CRM system or other high value site, that is a high price to pay for something that is beyond your control to remediate.

2.3.3. Non web based protocols – P2P, IM, Remote access

While email was the killer application of the past decade, a completely new generation is entering the workforce. These young adults have never lived in a time without cell phones or the Internet. Their comfort level with new forms of communication such as instant messaging and social networking websites requires companies that want to attract new talent to support these technologies. Legitimate programs like Skype have impressive NAT/Firewall traversal technologies built-in allowing them to pass through proxies with ease, and the malware authors are emulating them. You can no longer assume that an outgoing request to a web server, p2p network or any protocol is not a bot phoning home for commands. The use of deep packet inspection technologies, that are aware of these new protocols, is necessary to know what is really going on in your network.

The two largest categories of non web/email traffic right now are peer to peer networking (P2P) and instant messaging (IM.) Both types of programs have had highly publicized, remotely exploitable security problems – especially the instant messenger programs. Both also provide yet another way for malicious files to enter and sensitive information to leave your company through their file transfer capabilities. P2P programs are especially bad since users will install them planning to download free music or other content, and unknowingly, may share the content of their system with the rest of the P2P network. In one recent incident, a user inadvertently shared out a network diagram of one of the Pentagon's classified networks.²³ Instant messaging can also cause headaches for companies who, due to industry regulations, have to retain a copy of all of their communications. If you are not blocking or providing corporate versions of these applications, your users may have already installed them.

Growing in popularity is the remote access program. Citrix regularly runs nationwide ads for their product GotoMyPC with the promise of saving users the drive into the office to finish that last little bit of work over the weekend. An encrypted remote control utility controlled by a third party sounds almost identical in description to an attack tool to me.

There are sure to be other applications to come. The social networking phenomenon exemplified by Myspace, Facebook, and LinkedIn continues to grow, and people continue to post a gold mine of information useful to social engineers - whether they are stalking information or people. Our networks

²³ <http://computerworld.com/action/article.do?command=viewArticleBasic&articleId=9027949>, Visited 9/20/2007

are also beginning to possess some intelligence about the traffic going through them. Whether it is providing low latency to voice traffic or preventing the outbreak of a worm, we need better visibility into what is going over the wire to keep business assets performing for the business. Otherwise, you are almost guaranteed to be providing free hosting and bandwidth for sending Spam and launching denial of service attacks.

2.4. Legal issues

Some businesses have learned the hard way that technology and the “anonymity” it seems to provide sometimes can bring out the worst in people. In a company, of any size, there will probably be someone browsing the Internet for pornography or other content inappropriate for the workplace.²⁴ Employees might also use the email system to harass and intimidate coworkers. While the seedier side of the Internet is infamously used as an attack point for installing malicious software, the potential legal liabilities of allowing this kind of activity to go on could fund many security projects.

When you are discussing this issue with your lawyers and marketing people, ask them questions such as what it could cost your brand to be labeled yet another company that lost people’s social security numbers or credit cards due to not detecting and removing malicious software on your network? You should also inquire about the potential liability of allowing an attacker to use your network as a staging point to launch attacks on other companies.

3. Threat management

The preceding section on threats is not intended to be a comprehensive list. However, if you are able to deal with these classes of threats and identify new ones as they arise, you are much more likely to retain control of your network. In this section, we will discuss ways you can meet these threats, so that you can evaluate your existing controls and some new controls to see if you have reduced the risk associated with employee Internet access to an acceptable level.

The defenses against these new threats have had to evolve at a rapid rate from simple firewalls to network and host based programs constantly scanning for known and unknown threats. Some defenses, such as antivirus, are nearly universally employed in organizations. Others can vary wildly in the number of organizations using them, effectiveness and costs, so it can be hard to determine which controls are the best for your environment.

In this section, as I discuss the various controls, I will attempt to rate them on a few factors:

- Effectiveness – how well a class of controls addresses certain attacks
- Cost – When possible, an estimate of per computer cost for acquisition may be included

²⁴ <http://www.websense.com/global/en/PressRoom/PressReleases/PressReleaseDetail/?Release=050509928>,
Visited 12/7/2007

- Tradeoffs – Security is often a tradeoff between safety and functionality. Some controls may only be appropriate for high criticality assets.

3.1. Good controls

I will start with controls considered the bare minimum that most environments should be using. The lack of these controls would probably be noted in an audit, but their mere presence is not an indicator of an effective control either. If you notice that a basic control, such as an antivirus scanner, no longer prevents threats as effectively as it did in the past, you must be willing to replace it. You might choose another antivirus vendor or choose to implement a new control to help compensate for any weaknesses you have discovered in your antivirus software.

3.1.1. Antivirus/Antispyware (AV)

Antivirus has the distinction of being one of the original security controls deployed in most companies. Until a few years ago, the signature based technology used by both antivirus and antispyware products had been successful in dealing with the majority of threats encountered by the average Internet user. However, starting in the early part of this decade, the motivation of the average attacker changed. When mass mailing email worms, like Melissa, crippled email systems around the world, its author was still looking for notoriety and press attention. Today's malware author would prefer to go undetected for as long as possible so that they can continue to use your computer resources at will to deliver Spam, collect personal information and host websites to extend the number of computers under their control.

A classic military study is warhead versus armor. If one side builds a better tank or buries their command post underground, the other side just builds a bigger bomb to defeat the new defenses. AV/antispyware vendors have been trying to fight malware using signature-based technology that looks for the known hallmarks of a bad file. The attackers have responded by designing systems to modify their applications in ways that make it difficult for traditional signature based scanning to keep up. In some cases, every visit to a malicious webpage can create a custom script obfuscation designed to thwart scanners looking for known issues. When a drive-by download site automatically starts invisibly installing programs to your computer, they often send the malware in two or more different stages. The first stage that installs itself can be quite simple and changed often to avoid detection. Then the initial program downloads additional malware to enable the full functionality, thus giving the attacker remote control of your computer.

In order to respond to these threats, vendors are beginning to implement behavioral (or heuristic) modules that not only look for known bad software, but look for actions which "good" programs should not be performing in order to respond to new threats for which signatures do not exist. Unfortunately, that is not easy because many activities that a malicious program may want to perform are the same ones that many legitimate programs will perform when installing or running. Attempting to block all programs from performing behaviors - such as configuring the registry to start the program automatically, or going to a website and attempting to download an update - would break the functionality of many programs. It would probably prevent your antivirus from updating itself. Therefore, there is a risk of misidentifying a legitimate program and preventing it from running properly.

Several vendors have begun to implement behavioral technology in their products, but many of them have it disabled, or at least have severely limited the actions they will block by default, because they know it will generate help desk calls. The security consulting firm, Intelguardians, created a research tool called SPYCAR – a reference to the standard AV test file called EICAR – for an *Information Security* magazine article they were writing on this subject.²⁵ The program performed several of these gray area actions to see if antispymware vendors would protect you against an unknown threat, and the results were not encouraging. Many vendors either could not block the actions or had that ability disabled by default. The saddest reaction to SPYCAR – which is not malicious and comes with full uninstall support for any changes it makes to a system – was from a few companies, who instead of addressing the questions about the abilities of their product, just marked the SPYCAR executable as spyware, so that it would be blocked by signature instead of its behavior. This type of response from a vendor does nothing to address the actual risk faced by organizations. In emails from Tom Liston, one of SPYCAR's authors, he indicates that a future version of SPYCAR may need to incorporate the same anti-detection technologies as real malware in order to remain a valid testing tool in the face of vendors that would rather give people a false sense of security.

Should you find yourself looking for a new vendor, ask a few questions about how they are responding to these new threats. Do they recognize and consider software protection packers such as Themida and UPX? Are they still improving their signature-based technology with rapid releases for new malware and support for new attack vectors like obfuscated scripts on web pages? Make sure you test behavioral analysis in your environment to see what will need to be tuned and how many false positives will be found.

Antivirus/spyware products are one of your last lines of defense, especially for mobile users. If your vendor has not been updating their product to deal with new threats, or the behavioral component causes too many false positives when enabled, this may be an area in which you could make a substantial improvement by changing products. You should also pay close attention to what the products actually do, and not just the marketing hype, when evaluating them. Some vendors may claim they have a host-based intrusion prevention system when what they really have is a firewall and/or signature based application blacklist.

Effectiveness – varies greatly by vendor

Cost – Pricing starts at \$25 for new licenses, per machine, plus an additional annual subscription fee. Some vendors have moved to subscription only models.

Tradeoffs – Signature based engines are a mature, well understood technology. The new behavioral features could block legitimate applications and create an increase in help desk activity, so testing is important.

²⁵ http://searchsecurity.techtarget.com/magazinePrintFriendly/0,296905,sid14_gci1257060,00.html, Visited 12/11/2007

3.1.2. Intrusion Detection Systems (IDS)

As with antivirus, intrusion detection systems try to locate malicious traffic through signature and behavioral analysis. Just like antivirus, the bad guys have seen our defenses and come up with some creative workarounds. IDS is a maturing technology that was once declared “dead” in a Gartner research report as it was assumed that IDS systems would be superseded by intrusion prevention systems (IPS.) However, there is still a role for IDS systems given their ability to spend more time analyzing traffic and monitor links passively. The greatest downside of IDS systems is the possibility of information overload. An IDS must be tuned carefully for your network, or the number of alerts it raises may be impossible to follow-up on.

When budgeting for an IDS, be careful to include the human element in your equations. Tracking down IDS alerts takes time and training. After all, if you are not going to check out the alerts, then why do you have an IDS in the first place? This has led many organizations to outsource their IDS systems to third parties. Just remember you do not relieve your company of the responsibility to protect its assets by doing so; you are just managing the vendor instead of the system itself. Make sure you are receiving the level of protection you were expecting and paid for.

Some ways that IDSs can distinguish themselves include resistance to evasion techniques, low false positive rates, and the ability to help analysts quickly answer whether or not an incident is real. Some systems capture the actual packets that generated an alert, which adds to the acquisition costs, but can reduce the ongoing management costs. Your IDS architecture should be designed with more than just external attacks in mind. It should also alert based on signs of internal infections such as clients attacking other systems or running unauthorized services like email and web servers.

Effectiveness – Good alerting system, if you respond to the alerts in a timely fashion.

Cost – Snort is free but requires a lot of knowledge. Commercial systems are priced according to the throughput of the links they are monitoring. Outsourced providers can bundle systems and management, so they may be cheaper if you consider the cost of employee time.

Tradeoffs – Intrusion detection systems are notoriously labor intensive. Outsourcing your IDS may remove visibility into how well the systems are working. Whether internally managed or outsourced, you should audit your IDSs periodically to make sure they are being updated and alerts are being handled.

3.1.3. Gateway protections

Connections from one network to another, or from one level of sensitivity to another, have always been a perfect location to scan for and remove traffic that does not need to move between networks. The original perimeter protection device is the firewall. While the firewall is very good at its job of filtering out traffic, the average firewall does not normally look at the content of the traffic for performance reasons. As firewalls and patching have made it harder to attack the underlying operating system, attackers have moved up the OSI stack and started attacking applications directly. Since this type of

attack traffic is not readily detectable by looking at the header sections of a packet, it has led to the creation of several tools that are intended to scan the content of network traffic before it ever reaches its destination.

3.1.3.1. Proxies

The first applications to enter the content scanning area were proxies. A proxy handles outgoing connections at the application protocol level such as HTTP and FTP. Instead of a program, such as a web browser, contacting a site directly, all traffic is routed through the proxy, which is configured to be the only device allowed to make outgoing connections from the network. The proxy can then perform several useful tasks such as caching frequently accessed content, authenticating and logging all outgoing connections, and can even deny or filter traffic. Proxies make it much easier to manage your outgoing connections as clients can be denied direct access to the Internet. This allows you to take a default-deny stance and configure your routers where only a few computers may access the Internet directly. Unfortunately, they also have several drawbacks. Since a proxy must be able to handle all of the traffic for the devices behind them, the more functions they are asked to perform, and the more devices you put behind them, the more likely they are to become a network bottleneck. Proxies also must understand the protocols that you want to send through them. If your proxy is unable to handle instant messaging or some other business application your company needs, then you end up creating exceptions that bypass your proxy. Finally, like most protections, the bad guys have just adapted their programs to work through most proxies. So the level of protection provided against malware that uses Internet Explorer to phone home after you have typed in your password is somewhat limited. If you do use a proxy, one of the greatest benefits you may see is the logging of outgoing connections. Doing some simple analysis of sites visited and user-agent strings can allow you to identify infections on internal machines.

Effectiveness – Many proxies can block known bad sites and allow you to use third party tools to scan content for viruses. Throughput requirements, rapid malware repacking and script obfuscation all reduce a proxy's ability to filter out malicious code. Some of the best benefits may come from manual or automated analysis of log files.

Cost – There are free proxies available such as Squid and even free URL block lists like Dansguardian. A commercial product like Microsoft ISA server is licensed per processor starting around \$5000 for the first processor. Additional protections such as virus scanning will cost extra.

Tradeoffs – Proxies can be limited in the protocols they can handle, and you may keep finding the need for more proxies to handle instant messaging or video conferencing. As new applications arrive on the Internet, proxy servers may handle them poorly or not at all.

3.1.3.2. Email gateways

A very common perimeter device in networks is the email gateway. Sendmail and Exchange are two of the most common mail transfer agents (MTAs) on the Internet, and each has had its share of remotely exploitable vulnerabilities. A few examples include vulnerabilities that allow you to send commands to

be executed via email or having a system email you its password file so you can crack it without ever accessing the system.²⁶ Many organizations have decided that having a less complicated MTA program in their DMZ networks could help reduce their exposure to such problems.

The gateway is an ideal place to apply filtering to reduce the load on internal system by filtering out Spam, viruses and other email borne nastiness. New technologies that prevent connections from known Spammers, and quarantine the Spam that is received, can save real money by preventing the 75% or more²⁷ of emails thought to be Spam from making it to your internal systems. Once removed, those emails never need to be backed up or stored for compliance purposes, and users do not have to spend time cleaning out the junk to find legitimate business communications. Email gateways also have become the place where you can make a final check for your intellectual property leaving the building or encrypt content to limit its distribution to those authorized to see it.

Companies that do not want to manage email gateways internally, can also turn to the many service providers in the industry. Companies like Postini, Messaglabs and Frontbridge can filter email in their data centers around the world and then forward only the reduced email load to your organization.

Try to look for programs or hosted services that have a flexible rules engine that can quickly adapt to new threats. Over the past few years variants have tried techniques such as password protected zip files and placing their message in images and PDF files to fool detection engines. Unlike many areas of IT, the larger vendors in the email protection business will probably be some of the quickest to respond to new Spam variations. Since they can sample a larger percentage of the email going around, they will typically detect and respond to mass mailing worms faster.

Effectiveness – Even a simple gateway that scans for viruses, blocks executable extensions and removes some Spam will likely be paid back very quickly in the time it saves users and reduced need for helpdesk and incident response. Achieving extremely high Spam removal rates (greater than 90-95%) can be difficult since you will start to quarantine more emails that are legitimate.

Cost – Protection on individual machines will cost around \$25 per machine. Gateway devices are priced by the mailbox and start at \$15 to \$30 per mailbox. Hosted services are usually in the same range.

Tradeoffs – Whether you manage your own systems or outsource, filtering emails has become a business necessity. As long as your system does not block too many legitimate emails, few people will complain about having to deal with less Spam.

3.1.3.3. Web/protocol filtering gateways

Web/protocol filtering gateways come in several different implementations, but most work like an intrusion detection system with an active response component. The gateway device monitors the traffic

²⁶ http://en.wikipedia.org/wiki/Sendmail#History_of_Vulnerabilities, Visited 11/10/07

²⁷ http://www.postini.com/news_events/pr/pr020105.php, Visited 11/10/07

at a choke point before it leaves the organization's network and heads out to the Internet, and when it sees a policy violation it responds by blocking the traffic. This can take place by resetting the TCP connection or responding back to the client request faster than the intended server can (OSI layer 3 monitoring), dropping the packet (OSI layer 2 monitoring), or by reconfiguring a firewall to filter out the traffic. Also included in this category are specialized OSI layer 2 devices that perform like an intrusion prevention system. They can scan traffic for malicious code in web pages, and other protocols, and drop any packets where malicious content is detected.

Many of these devices started out as web filtering devices that would block access to certain categories of sites such as pornography and hate sites, which have caused legal problems in the past. This type of technology was already readily available in proxy type systems, but can be easier to implement as a filtering device since you do not have to reconfigure any clients. The more complex programs in this category can monitor many different types of protocols beyond HTTP and FTP including instant messaging, peer-to-peer, remote access, pop3/imap, and streaming audio and video. New support can be added as quickly as the protocols can be analyzed.

Unlike a proxy, a filter does not need to be involved in the data flow between the client and server unless it detects a problem and reacts to block the connection in some way. This brings with it some good and bad consequences. While it can be a benefit that new applications will usually just work without the need for setting up a new server, the way you might need to with a proxy, it can also be a problem, as your ability to log and manage a new protocol is dependent on your filter's ability to decode it. This may leave your filter blind to the existence of some traffic that it cannot interpret. Keep in mind that filtering is a default-allow technology, with explicit denial of known bad traffic, while proxies are a default-deny technology with explicit allowance of supported traffic. A basic principle of secure design is that default-deny will be the safer choice, but as always, you must balance functionality, security level, and cost.

Effectiveness – Varies by vendor. Products with better coverage of protocols and constantly updated databases will protect better. There are known evasion techniques for bypassing both URL filtering and virus gateway scanning including the Google translation function for the former and script obfuscation and multistage infections for the latter. Filters may be particularly vulnerable to new protocols, so the detection and logging of unknown traffic is a good feature to look for.

Cost – A 500 user appliance starts at \$4000 from Barracuda. The more threats you want to protect against, like protocol filtering for IM and P2P, the higher the costs will be. Hosted content filtering starts at \$20 per user annually.

Tradeoffs – While it is easier to monitor and block a new protocol than to write a proxy for it, the default-allow nature of filtering technologies always leaves them responding to something new. What you gain with filtering is a more functional network that can respond to changes without needing to wait for a new proxy server to be written. What you lose is the slightly more secure proxy-based model where protocols are enabled as the need arises thus giving you more control over what traffic is transiting your network. Since a filtered network has more ways to communicate with the Internet, it

will face a wider array of threats, the same way a computer running unnecessary services can be attacked in more ways than a computer using a hardened configuration.

3.1.4. User education

Trusted experts in the information security industry like Marcus Ranum and Bruce Schneier are beginning to doubt the effectiveness of training users as a security control.²⁸ The inherent thing that makes a computer useful to us in the first place, its programmability, is often used against the user to trick them into giving up control. However, until IT departments make great strides in controlling what programs can or cannot run on a system, and software and interface designers make changes to help users make better choices, the user will still be a large part of securing systems and information.

Since it is difficult even to define a successful user education program, deciding what to spend on one can be a difficult task as well. The best programs will focus on the motivations for doing the right or wrong things. Positive motivating factors may include forcing users to pass a quick five question security quiz before being allowed Internet access for a 24 hour period. Negative factors will likely include the popular “you could be fired” motivation. Recommendations also include limiting the amount of information presented at any one time, and pointing out when information is also applicable to home computing such as protecting kids online and preventing identity theft. A little bit of information repeated often is going to be remembered much better than a two hour presentation once a year.

Training should be tailored to certain job roles.²⁹ Software developers, helpdesk staff, and information owners are just some of the groups that have a need for training specific to their job role over and above typical end user training. Your content should be tailored to current pain points. If your technology is doing a good job of preventing a threat, then focus on other items where users are still making bad decisions that are costing the company money. I also recommend individual training for anyone repeatedly infected by downloading shareware or caught handling sensitive information inappropriately.

Effectiveness – This is difficult to assess and is the cumulative effect of many factors. Generally speaking, if your program is tailored to the audience and delivered in smaller chunks, more often, throughout the year, then you have a better chance of the information sinking in. Creative delivery by good speakers, entertaining stories in newsletters or posters will also aid in retention. A boring half-day training seminar every couple of years will have a much lower chance of changing habits.

Cost – You must count the cost of taking employees away from doing their jobs if performing formal training. Using shorter techniques such as lunchtime training or the previously mentioned quizzes prior

²⁸ http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1256941,00.html, Visited 11/10/2005

²⁹ http://www.computer.org/portal/site/security/menuitem.6f7b2414551cb84651286b108bcd45f3/index.jsp?&pName=security_level1_article&TheCat=1001&path=security/2006/v4n5&file=bsi.xml, Visited 11/25/2007

to Internet access may reduce costs. Formal training at both new employee orientation and change of job role may still be necessary.

Tradeoffs – A poorly implemented program will waste money, people's time, and not increase security over the long term. As a security awareness program is easily done poorly, and difficult to do well, it is easy to see why user education has such a bad reputation.

3.2. Better controls

Controls in the "better" category begin to go beyond the traditional security controls in most companies. They may cost more money in acquisition or ongoing costs than the "good" controls, but when fully implemented, they can provide a higher level of protection for assets that are more valuable. Systems that handle valuable information like human resources data, customer lists and marketing plans may be targeted for commercial espionage. Mobile workers like sales people and executives come with their own set of unique challenges since they connect to many untrusted networks and may receive updates and patches less frequently. Assets with more sensitive information or that operate in a higher risk environment may need an equally higher level of protection.

3.2.1. Intrusion Prevention Systems (IPS)

A step beyond intrusion detection is intrusion prevention. This can take place at the network or host level and is fast becoming one of the more effective tools for preventing and detecting infections in a timely manner. Using signature or behavioral techniques an IPS can prevent many infections from ever happening, but they are hardly a one-stop solution. Each type has limitations on the activities it can monitor and an incorrect configuration can break applications. For this reason, IPSs tend to come with many of their features disabled and extensive tuning is required before you can use the full capabilities of the product.

3.2.1.1. Network-based IPS (NIPS)

NIPS devices have the benefit of protecting multiple devices and the ability to watch traffic for patterns that might indicate signs of an infection. This allows a NIPS to begin to block traffic using signatures that identify known bad traffic or use behavioral methods to alert or block traffic that is out of the ordinary. The problem is, in order to do the latter, your NIPS device has to know what good traffic looks like. NIPS devices must also process traffic in an extremely short amount of time. A detection only device can detect and alert in seconds or hours after the traffic has passed by the device, while a prevention device must make the decision on whether or not to drop a packet in milliseconds, or it will become a network bottleneck. The greater the amount of traffic to be monitored, the more specialized hardware and extensive tuning is needed to achieve acceptable performance. This issue is exacerbated by the recent merging of security devices into a single platform now being sold under the category of unified threat management (UTM.) For a brief presentation on some of the benefits and risks of UTM devices, please see the following presentation by Manuel Santander and me on the subject.

http://www.sans.edu/resources/student_projects/200709_003.doc

A NIPS device possibly may be evaded through a variety of techniques that exploit their low latency requirements, so a NIPS device may not eliminate the need for a NIDS on your network. An incorrectly configured NIPS can cause hard to diagnose connection and throughput problems or even bring your network to a halt. Take special care when updates are applied to the NIPS rules or major changes to applications take place - especially if you are using a behavioral detection engine to prevent a self-inflicted denial of service by your own equipment.

Effectiveness – Highly effective at preventing known network based attacks. The ability to block unknown attacks will vary based on the detection engine and what rules you have enabled. The more behavioral/anomaly prevention rules you have running, the more likely you will start seeing false positives and begin to block legitimate network traffic.

Cost – Pricing is based on capacity. Units to protect a T1 (1.544 Mbit/sec) connection can run on commodity class hardware and will cost from \$3000 and higher depending on additional features. UTM devices will cost more but the combined price may be less than buying pieces individually. Units to protect gigabit/sec speeds are considerably more expensive.

Tradeoffs – NIPS like NIDS takes ongoing manpower to tune and check alerts. The more hosts you need to protect, the more tuning will be required. Limiting the types of systems and applications protected by a NIPS will help by allowing unneeded rules to be disabled.

3.2.1.2. Host-based IPS/IDS (HIPS)

While NIPS/NIDS provide coverage for multiple devices, they suffer from two major shortcomings. First, no NIPS or NIDS device will ever be powerful enough to examine all of the programs and scripts that pass by it on the network for evaluation of whether the code is malicious or not; nor can it protect against threats it never sees, such as when your executives check their email and surf the web from a coffee house over an unencrypted wireless access point. Second, a NIPS/NIDS cannot see actions that only happen on a client device such as programs loaded via USB drives or CDs. In order to address these issues, a new class of protection on the host itself was created that allows for the monitoring of suspicious activity and blocking of unauthorized programs and network connections. Host-based IPSs can protect critical assets like laptops that are outside of the traditional perimeter protections and even help verify that those machines are safe before allowing them to connect back to the corporate network.

HIPSs can monitor nearly any activity that takes place on a system using both signature and behavioral technologies, but not all HIPSs are created equally. Since there is no standard definition of what a HIPS is, you may find vendors have renamed an existing product such as a client firewall or antivirus product and started calling it a HIPS. You will need to ask many questions about application whitelisting/blacklisting, system call monitoring, buffer overflow detection and other issues to see whether the system you are looking at provides any additional protections above and beyond a traditional Antivirus/spyware program.

Effectiveness – In a more static environment where applications change infrequently, a HIPS can provide very effective protection - especially if you can define all allowed applications in advance and deny any others from running. In environments where applications change often, you can prevent known bad applications from running, but malicious code is constantly changing. Either new programs are created or old ones can be repacked. To stop them a product will need more advanced tools including buffer overflow detection and the ability to deny applications using criteria such as uncommon packing techniques and attempts to make changes to sensitive parts of the operating system.

Cost – Either as an add-on to your current antivirus/spyware or a separate product, you can expect to pay \$25 or more per seat for new purchases. You may pay less if you are migrating from a competing product.

Tradeoffs – HIPS suffer from the same false positive issues as NIPS, especially when using behavioral detection. If you configure the product to prompt users when a suspicious activity takes place, make sure the dialogs offer the correct choice as a default and expect them to allow the action anyway. If you configure the product to deny applications silently, expect to receive helpdesk calls on why certain applications do not work.

3.2.2. Patch management

Microsoft, Redhat and others thankfully have made operating system level patching much easier over the past few years with products like Windows Software Update Services. The bad guys, of course, responded in two ways. First, they started analyzing the patches to see what was being fixed, and used that information to create exploits for unpatched systems. Second, they began looking for other targets such as applications where vendors may provide either less comprehensive or no automatic patching system at all. Microsoft's Automatic Update technology, while very helpful, has really just highlighted the fact that you need a comprehensive vulnerability management strategy for all of your operating systems and applications. Please do not forget that routers, switches, printers, and anything else with a network port on it have operating systems; and those operating systems all have vulnerabilities in them, whether or not they have been discovered and publicized.

Since we are addressing the threats associated with accessing the Internet, I will only address patching client systems. Whether you use a standalone product, or one integrated into an overall client management product that provides software delivery and asset management functions, you need a tool that can identify known vulnerabilities on systems. Then, once a vulnerability has been approved for remediation, the system should let you know whether a patch has been successfully applied or not. Since no vendor can possibly maintain a vulnerability information database for all existing software, your product must also be able to support custom detection and patches.

The combination of the complexity of modern software and the overall lack of securing development skills leads us to the inevitability of vulnerabilities. They already exist whether or not anyone has discovered them, and so the patches are going to keep coming.

Effectiveness – While you can modify an exploit in an infinite number of ways to bypass AV/antispayware/IPS/etc. removing the vulnerability means all of that work by the bad guys hopefully does not matter. There have been plenty of cases where the bad guys understood the vulnerability better than the people patching it, so we also have to be prepared to patch the same problem multiple times.

Cost – If you already use centralized software deployment, you may already have a basic patching capability. However, in many environments, the need for reporting and paying a vendor to research patches instead of doing it internally may require purchasing an add-on module or standalone product. If you are in a Windows environment, WSUS and Automatic Updates are free for modern versions of Windows. Products such as HFNetchk Pro cost about \$25 per seat.

Tradeoffs – While vendors may do their best to test patches before they are released, they can never replicate all possible environments. Therefore, you will want to make sure you test patches in a controlled environment to make sure the fix is better than the vulnerability before deploying across your enterprise.

3.2.3. Defense in depth

Defense in depth is not a control in the typical sense, as you cannot just go to a website and buy it. Rather, it is a way of selecting and implementing controls in a way that they offset each other's weaknesses and of focusing your efforts on the assets that need the most protection. Most of the controls discussed so far are implemented in a way that they protect nearly all machines equally. However, not all of your assets are equally valuable, so it makes sense that some of them are going to need additional protection. By looking at assets from a few different perspectives, you can begin to decide where to apply your time and budget in the areas that will make the most difference. The remaining controls discussed in this paper may only be applicable to a subset of systems on your network. A few tips on locating systems that may need additional protection based on Stephen Northcutt's defense in depth weblog entries follow:

- Information centric view - The information stored or processed on the system is critical to the ongoing operations of your company.
- Threat vector view - The system is exposed to a greater number of likely threats than others (e.g. a salesperson's laptop is often accessing the Internet without the protections covering computers on the corporate network.) This is the view we have been using throughout the paper.
- Attack surface view - Systems may be poorly configured or have users that are easily tricked into performing actions that give away information or cause other losses.³⁰

³⁰ <http://www.sans.edu/resources/securitylab/76/>, Visited 12/10/2007

Now that you know where to look, start by evaluating your existing controls in a holistic manner to see what improvements can be made using what you already have. As an example, signature based antivirus technology is a purely reactive technology. When an attacker releases a new program, or a repacked version of an existing one, they check it against antivirus programs to see if it is detected. Can your other controls help mitigate this?

An example, circa 2005 – A user's home computer was unusable due to spyware infections, and so the user took home a floating corporate laptop. The next day a virus detection message on that laptop initiated an incident response where we found out the user's spouse had access to the user's login information and the spouse had triggered the incident. Several policy and technical controls had to fail for the spouse to get to that point including:

| <u>Control</u> | <u>End user circumvention</u> |
|--|--|
| Policy against allowing another person to use your account | Gave username to spouse because home computer was broken |
| Policy against checking external email enforced by filtering software that blocked access to web based email, pop3, and imap | Spouse checked web based email account on a small ISP that was not in the filtering block list |
| End user training on not opening attachments you are not expecting | Employee's spouse did not attend the training |
| Anti-virus | We got lucky here |

In 2008, it is much less likely that antivirus, using signature only technology would stop today's rapidly evolving malware, and so our IT department came to the conclusion that additional end point security technology that could help with unknown threats was going to be needed. We also added this incident to our awareness training - without giving out the name of the person involved - and at least the embarrassment factor will help that one user remember the policy now. While this paper is only discussing the threat of employees having Internet access, also remember that your controls face a wide array of threats; and so when selecting and deploying them, try to implement controls that help you address multiple threats at the same time.

Defense in depth also prompts an organization to review the configurations of existing systems, applications and controls to see if a change would lead to a more secure environment. What you are trying to do is reduce the attack surface of the systems, thereby reducing the number of ways a program can be attacked or a user can make a mistake. Many operating systems and applications come with default settings, which allow much more functionality than needed to complete the job for which they were acquired. By hardening configurations and disabling unneeded features, you can eliminate access to vulnerabilities which, as we have already discussed, are there waiting to be discovered in all modern software. When analyzing your controls you may find you are not using all the features of your products. Perhaps parts of the technology were not mature when a product was initially deployed, and

they may have improved overtime, or new features have been added of which you were not aware. By evaluating your configurations, you should be able to improve your security posture without spending a great deal of money. After that step, focus the purchases you make so that they focus on your more important assets, then when one control fails to protect you, another may.

Effectiveness – Usually very effective as a layered security approach can offset weaknesses in an individual product.

Cost – When hardening existing systems, there are no acquisition costs, only implementation. Security is a highly competitive marketplace and in many segments, like AV, where there is no hardware component, you may receive a large discount when changing to a competing vendor. Some vendors may go as far as to give you the product for the cost of ongoing maintenance fees. With a little research, you could end up with a superior product with no additional spending over your current level.

Tradeoffs – The biggest tradeoff for a defense in depth strategy is time. You have to spend time to keep up with new threats, evaluate configurations and test alternative products. All of these steps take time and are difficult to accomplish for an IT department that is in constant firefighting mode.

3.2.4. Security as a software selection criteria

Many programs are not developed with security in mind: they are not developed using secure coding techniques; they come with too many features enabled; they often have poor architectural or interface design issues that create unnecessary vulnerabilities, and cause users to make poor decisions when faced with a choice. What if, as part of your software acquisition strategy, you added the inherent security of the product as selection criteria? Could that help to reduce some of the money you were spending on controls to protect software that did not consider security issues when it was written?

A few relevant examples for your consideration:

- Comparison 1 - Suppose you were in the market for a new web browser, and you were looking from a security first, functionality second perspective. You have short listed Internet Explorer 6 and Firefox version 2 with the noscript³¹ plug-in. Both browsers can show fully featured websites with one difference. If the user is using Firefox+noscript, all active content is blocked until explicitly enabled. Many websites will run fine with active content disabled, and for those whose active content you actually want to see, enabling it is only a mouse click away. Combine that with the fact that most browser based exploits require scripting to run, and you have just removed one of today's most effective attack vectors. Unless there is another business requirement, such as a site that only supported ActiveX, Firefox starts to look very attractive.
- Comparison 2 – Windows XP with Internet Explorer 6 versus Windows Vista with Internet Explorer 7. Service Pack 2 for XP made some very useful security improvements including enabling the built in firewall by default, and the ability to prevent some buffer overflow attacks

³¹ <http://noscript.net/>, Visited 11/27/07

using data execution prevention (DEP). Microsoft has included several new technologies to aid Internet users in Vista. It has further expanded the DEP support started in XP and added address space layout randomizations³² for some windows components to make it harder to guess the return pointer during a buffer overflow attack (although the ani exploit³³ from January 2007 shows it has not been implemented everywhere. ^{***}) They have also implemented the concept of least privilege with the new user account control system. While it takes a little getting used to giving consent before performing administrative level tasks, the benefit is that even if you are logged in with administrative level permissions, most programs that you run will still execute in the context of a normal user. This means that even if you were to become infected by browsing the web, the damage should be limited, and the likelihood of installing system-modifying technologies, such as a rootkit, would be reduced. Internet Explorer 7 also includes helpful interface improvements including prominent warnings when the information for an SSL protected website cannot be verified properly with a trusted certificate authority, and a warning system for known phishing sites. Compare this to Internet Explorer 6 where if there was an issue with an SSL certificate you would be prompted with a dialog box that indicates there is a problem, but you can go to the site anyway by just clicking yes.

As of this writing (pre Vista service pack 1) there are still concerns with the stability and speed of Vista especially on older hardware that may outweigh the security concerns, but it shows that Microsoft is thinking about these issues across its products. Outlook 2007 has received similar treatment by using Word 2007 to render HTML formatted emails instead of Internet Explorer. The Word HTML renderer is a much less complicated piece of code that does not support all of the HTML standards or ActiveX, thus greatly reducing the attack surface of the product.

While it is unlikely security will be the highest weighted criteria in our selection process, it is well worth discussing with your vendors. Vendors will not make security a priority in their development process until buyers start demanding security, as the U.S. government recently did in requiring all programs purchased to run under one of the NIST approved, hardened configurations of Vista³⁴ called EC and SSLE.³⁵

³² <http://www.microsoft.com/technet/technetmag/issues/2007/04/VistaKernel/>, Visited 12/4/2007

³³ <http://nvd.nist.gov/nvd.cfm?cvename=CVE-2007-1765>, Visited 12/4/2007

³⁴ <http://www.whitehouse.gov/omb/memoranda/fy2007/m07-11.pdf>, Visited 11/25/2007

³⁵ http://csrc.nist.gov/itsec/guidance_vista.html, Visited 11/25/2007

*** The ani exploit was a simple stack overflow in a Windows dynamic library that had been around since at least Windows 95. It was quickly determined that the exploit could even run on Vista the newest and most "secure" version of Windows. Due to worries about false positives, the updated compiler - in use since Windows XP SP2 - did not flag the code as a vulnerability; and when Microsoft fixed the exact same type of problem in the exact same code two years earlier, no one checked to see if there might be similar issue elsewhere in the code.

<http://www.determina.com/security.research/vulnerabilities/ani-header.html>

<http://www.informationweek.com/windows/showArticle.jhtml?articleID=199202711>

Effectiveness – Only time will tell, but as more companies begin to demand security, it is likely the vendors will follow. Making better purchase decisions now may not payback immediately, but it will reduce the ongoing cost of securing applications over time.

Cost – Once again, the major cost is time. The products must be researched, RFPs and contracts have to be written to take security into account and vendor claims have to be verified. You can try to shift some of the cost onto the vendors by requiring that security scans of their products be included with their proposals using tools such as web application scanners and vulnerability scanners.

Tradeoffs – Since it can be difficult to show the benefits of security in financial terms, it often is discounted in the financial analysis of a potential system. Security is not often the first concern of people who just want their business problem solved, so the acquisition process will have to be carefully monitored to make sure that it is given consideration.

3.3. Best controls

There may be a subset of computers in your organization that are so sensitive they must be protected with extraordinary measures. The consequences of losing a computer that helps run a nuclear reactor, the power grid, life support systems in a hospital, a pharmaceutical company's latest research or the formula for Coca-cola are potentially disastrous for a company or even life threatening. For these computers, you will need to make the best effort possible to prevent a threat, but failing that, you need to detect when an incident occurs and trigger a response quickly.

3.3.1. Network architecture

One of the most obvious controls would be not to allow any Internet access at all; and, if that is an option, it is probably the one you should take. Even if you take this approach, you may find a targeted attack jumping from your internal network to your protected machines so a control you need to consider is separating these assets into their own network. VLANs can be used to accomplish this, but there are attacks that can allow machines to circumvent VLANs, and configuration mistakes can allow machines to impersonate a switch and join several VLANs. The best approach would be to use a physically separated network running on separate hardware. This network needs to be firewalled and monitored at all connection points, and all ingress and egress rules should be only allowing needed traffic to pass. The tighter the rules are, the better. Many companies will want to monitor their SCADA systems from the corporate network, but there should be no access from your DMZ networks or the Internet unless it is a response to an outgoing connection. If you are going to allow any outgoing connections, they should be explicitly defined, as well. This egress monitoring is not a typical configuration in most networks, but is critical to controlling a highly protected asset. Should a machine need an update for antivirus or time synchronization, these are services you should be able to replicate on your internal network. Users of these computers should not be able to browse the Internet or check email. Move those services to a technology such as a terminal server and only allow the connections to it or provide a separate machine for those functions on the normal company network. For companies thinking they can use virtual machines to save money on hardware, just remember a virtual machine is

not a firewall; and it is not recommended that you place virtual machines with different security levels on the same hardware.

Network admission control (NAC) technologies should be able to help in this area. By checking a computer's security state and user's credentials before allowing access to the network, you can have more confidence that systems are receiving role-based access to the network. This will prevent something as simple as plugging a computer into a different wall jack from changing the protection given to a system.

Effectiveness – Allowing access from the secure network to only what is necessary to complete the job is a very effective control. In this case, you must place security over convenience to keep the machines performing their primary functions. Watch out for people trying to add to the list of what is “necessary” and verify it before any firewall rule changes are made.

Cost – Network segmentation can be one of the most expensive controls as it may require creating a separate infrastructure parallel to the existing network, but if you ever need to disconnect the networks from one another to keep a worm, which is flooding your corporate network, from disabling the secure network it is required.

Tradeoffs – The functionality of these machines should be greatly reduced requiring people to use more than one machine in most cases.

3.3.2. Darknets/honeypots

If you cannot prevent an infection from taking place, your next level of detective controls has to alert the appropriate people to initiate a timely response. Tools like an IDS can raise alerts on suspicious events in your normal traffic. A new breed of tools such as honeypots and darknets monitor for activity where there should not be any, and can raise an alert for scanning activity and even insider abuses.

Honeypots have been around for several years now, and people are beginning to branch out from their original intended purpose of observing hacking activity to see what they can learn from it. Some companies are placing honeypot machines within their own network just to see who is poking around. These machines may look just like any other server or client on the network, and they may have what looks like interesting information placed on them so that a log can be kept of who is trying to look at it. In a secure network, you could place a honeypot machine, and use it to log all connection attempts, file access or any other potential useful information. Since your honeypot machine should not normally be accessed, an alert should be raised if it is. That way the source of the alert, whether it be a real problem or not, can be tracked down. Like most intrusion detection technologies, there will be a certain amount of tuning that you will need to do to reduce false positives from “normal” network traffic, but after the tuning process, any new activity should be worth investigating.

Darknets are address ranges on your network that do not have any hosts on them. They can be small ranges inside legitimate address ranges or completely separate networks that should have no activity at all. Either way, at the router level you should be able to redirect any traffic headed to or from a darknet

range to a packet collection device for analysis. Darknets may turn up all sorts of information such as incorrect network configurations and activity of which you were previously unaware. Once you have resolved those issues, any new activity should raise an alert that unwanted activity is taking place such as an unexpected firewall rule change.

Effectiveness – The main purpose of a honeypot or darknet is to detect activity that should not be happening. Therefore, if your network is configured correctly and your alerts are tuned properly, they can be very effective as all warnings are events that need to be looked into. Reaching that point requires a lot of time and expertise.

Cost – Typically, either control could be supported with old hardware or existing network equipment. The only capital outlay would be for storage space for logging and analysis. There is a large time component involved for those responsible for monitoring and tracking down incidents.

Tradeoffs – Both tools can give you improved visibility into what is going on in your network, but can be complicated to implement and maintain. Make sure you are getting the most out of your more common security controls before trying to start into a project like a honeypot. There are also some potential legal issues associated with capturing this kind of traffic. Make sure you understand these issues and how they change from one jurisdiction to another.

3.3.3. Integrity monitoring

A common problem with control systems is that the vendors will certify an underlying operating system and then refuse to allow any changes to it. For example, they may not allow you to add security software or install patches. These types of machines are definite candidates for placement on an isolated network as discussed previously - even if they are just on a VLAN because their sensitivity level does not justify being placed on a physically separate network. These machines probably match two criteria. First, they are easily attackable due to the lack of protection and updates; and second, the applications and underlying operating systems probably do not change very often making these computers a good candidate for system integrity monitoring.

One of the original file integrity tools was Tripwire. In concept, it is simple, just cryptographically hash every file on the system and then store the results. If, in any subsequent test, the file hash does not match the one in the database, then you know the file has changed. In practice, the major complication is that files are always changing on modern computers, and even limiting the scope of the monitoring to application and operating system directories will still raise too many alerts to be useful on most systems. If your systems are relatively static, then you may be able to use integrity checking to alert you to changes in critical files.

Some newer tools such as Osiris, Samhain and the commercial version of Tripwire can monitor the OS kernel for modifications and report to a central server making database management much easier.³⁶

³⁶ <http://www.la-samhna.de/library/scanners.html>, Visited 11/27/2007

Effectiveness – In the right environment, integrity monitoring is an extremely useful control. The biggest problems are keeping on top of approved changes and hashing the correct content. Organizations that already have a strong change management process in place will be able to monitor much more content and even use integrity monitoring to alert on deviations from approved configurations.

Cost – Free open source alternatives are available including Osiris and Samhain. Commercial tools such as Tripwire have expanded to monitor more items such as databases and provide better reporting on change management issues. These expanded features can cost hundreds or thousands of dollars per machine depending on what you want to monitor.

Tradeoffs – Hashing systems with frequent changes will lead to many alerts that must be verified against the approved configuration. Companies without any type of standard configuration or change management process will quickly be buried under a flood of alerts. Start by hashing static content and then adding additional files as you have capacity to handle them.

3.3.4. Disposable computers

Desktop management has progressed to a very interesting point. Several vendors offer an entire toolset that can allow you to capture a user's personal settings, reinstall a computer from a disk image, reload any needed software packages and then restore the user's settings. While this capability is usually sold as a migration tool, it can also be used to make sure that machines in your environment are consistent with the latest standard configuration as approved by your change control board. As long as you keep your image and packages up to date, you can schedule computers to rebuild as often as every night - assuming you have the bandwidth – and this can become a part of your patch management process. A convenient side effect is that it will clean up any malware infection, even if a rootkit is installed.

Effectiveness – Whether you do this in a preventative fashion or in response to an incident, rebuilding a computer is often the only way to be sure you have remediated an infection. If you can fully automate the process to run during off-peak times, then the effects on productivity can be limited. Many computers also tend to slow down over time. Many users will notice an increase in performance after their computer is reimaged.

Cost – You may already own the technologies. Many companies already use tools like Altiris, SMS, CA, and Landesk to support their systems. The boot server and settings migration tools are usually a minimal incremental cost over the base product or already included.

Tradeoffs – Multicasting can reduce the bandwidth needed, but hardware and network failures can leave your users with no usable computer at all.

4. Some hypothetical numbers

In the battle for our networks, we are facing well-financed and capable enemies. In intelligence circles that would classify as a credible threat worthy of attention, but for many reasons getting budget for a security project can still be problematic. Technologists bear some of the blame as we are not known for our communication skills or understanding of the business. Most IT departments have also taken on

more change than they can handle at one time or another and managed a project or two into the ground. The goal of this section is not just to talk about the latest cool technologies to lock down your network, but to help you match appropriate protections to your assets, based on their value to the company. That value is not just the physical replacement cost alone, but also that assets' ability to help your company make money, as well as how long you can continue to operate without that asset. These additional factors may be considered its criticality to the enterprise. If when designing your controls, you can balance an asset's value, and the threat to it, with the cost of your security controls, then management may be more inclined to provide you with the money needed to secure those assets.

The following is just a glimpse of what some of these controls are costing organizations. It is outside the scope of this paper to help you define the value of your assets, but once you have undertaken that process, you should be able to begin to balance your controls against the threats your assets are expected to face. The traditional quantitative method begins with taking the cost of a single incident, known as the Single Loss Expectancy (SLE), and multiplying the SLE by the expected likelihood of an incident happening in a year, or the Annual Rate of Occurrence (ARO), to give you the Annualized Loss Expectancy (ALE).³⁷ While I am sure that everyone has those numbers handy for all of the threats to all of their assets, perhaps we should play it safe and do some qualitative analysis as well.

With qualitative techniques, you will not end up with detailed financial numbers, and the results are more subjective; but you should at least end up with knowledge about what is critical to your organization and where you should be focusing your security spending. We will assume that you have gone through your assets and business processes and determined which ones are critical to the company's ability to make money and how long you can survive without them. The next step is to detail the threats that have a high likelihood of damaging your systems. When you see a highly likely threat affecting a critical asset, you need to consider focusing additional spending on items like the better and best types of controls, covered earlier, to protect that asset.

The criticality assessment of an asset is not always intuitive. Take the comparison of the CEO's computer and a technical writer for a manufacturing firm. The CEO's computer is obviously important and the exposure of information on the computer could reduce stock values and cost the CSO their job. The technical writer regularly works on the company's intellectual property and could, at any point in time, be working on highly sensitive information and specifications. There are many companies where the loss of this information could be catastrophic allowing competitors to recreate products without any of the development costs. The quick lesson here is that you have to know your business. If you do not know what information is important in your company, you are probably not in a good position to evaluate the effectiveness of security and auditing controls.

What follows is the combined cost for some of the good and better controls discussed earlier. These numbers are based on the cost of new licenses of commercial products in the category or, where indicated, a managed service since the managed service price can help us factor in the labor associated

³⁷ http://www.windowsecurity.com/articles/Risk_Assessment_and_Threat_Identification.html, Visited 11/27/2007

with a control. There are free or low cost products in many of these categories, but they may cost a company more in ongoing labor to manage the product or troubleshoot it in some cases.

Good controls

| Control | Acquisition Cost/user | Ongoing costs |
|-----------------------------|-------------------------------|---|
| Antivirus/Antispyware | \$25 | Yearly maintenance, low labor |
| Managed email gateway | None for managed service | Annual subscription \$30/mailbox |
| Managed web filtering | None for managed service | Annual subscription \$20/user |
| Security awareness training | Normally Internally Developed | \$50/hour/employee an average of 2 hours/year |
| Ongoing subtotal | | At least \$155/year |

Prices quoted from CDW 11/27/07 based on 500 users.

Better Controls

| Control | Acquisition Cost/user | Ongoing costs |
|-------------------|----------------------------------|--|
| Network-based IPS | None for managed service | \$5000/month ³⁸ |
| Host-based IPS | \$40/machine (Cisco CSA desktop) | \$10/Machine (Cisco CSA desktop SASU smartnet) |
| Patch management | None for managed service | \$25/machine HFNetChk pro |
| Ongoing subtotal | | \$45/machine/year |

NIPS managed service for the better category included NIPS, managed firewall, security log management, content filtering and vulnerability testing in the price for a hypothetical 200 person company. So the \$20 annual subscription for web/content filtering in the good controls category has been subtracted from the combined total for good and better controls.

Just considering ongoing costs for the above items, you are looking at \$155 (Good) - \$170 (Good + Better) a year, or \$77,500-\$85,000 a year for a 500 user network. When you start to compare this to the cost of cleaning up several malware infections which will cost both end user time and helpdesk/incident response staff time, you can begin to see how best practice security spending is justified. The more

³⁸ <http://www.networkcomputing.com/showArticle.jhtml?articleID=191203015&pgno=2>, Visited 11/27/2007

financial information on the cost of controls and the value of your assets you can collect, the better decisions you will make regarding what controls to invest in.

5. Checking with the 17799 standard

Now that we have completed our threat analysis, this is the proper time to go back to whatever standard you are using – in our case ISO 17799:2005 – and use it to see if you have missed any areas relevant to your analysis. You can download an audit checklist based on the 17999 standard from SANS at the following location: https://www2.sans.org/score/checklists/ISO_17799_2005.pdf

Relevant ISO 17799:2005 section number - title

6.2.1 - Identification of risks related to external parties

6.2.3 - Addressing security in third party agreements

8.2.2 - Information security awareness, education and training

10.4.1 - Controls against malicious code

10.4.2 - Controls against mobile code

10.6.1 - Network controls

10.8.4 - Electronic messaging

10.10.2 - Monitoring system use

11.4.5 - Segregation in networks

11.5.4 - Use of system utilities

11.6.2 - Sensitive system isolation

11.7.1 - Mobile computing and communications

12.1.1 - Security requirements analysis and specification

12.4.1 - Control of operational software

12.6.1 - Control of technical vulnerabilities

Whether or not you actually implement a control is a combination of threat analysis, legal requirements and negotiations with your auditor. One of the most important things is that you have evaluated the risk, and can let your auditor know the rationale behind why you did or did not implement specific controls. You may also find that certain sections of the standard may not apply to your business. If you do not sell any products over the Internet, you may not need to address the area of e-commerce. Issues of this nature are typically covered in a statement of applicability, which defines both what parts of your

business you are assessing and which parts of the standard apply. Just be prepared for your auditor to disagree with either part.

In the course of discussing controls that apply to the threat of employee Internet access, it appears that we covered all of the relevant areas except for sections 6.2.1 and 6.2.3, which deal with third party access to your network. In our case, we would need to update our threat analysis to make sure we addressed any potential issues raised by giving contractors or partners Internet access through our organization, and make sure they agreed to comply with our policies. These issues are primarily policy and contracting issues as opposed to technical controls.

It is also important to remember that while we have touched on all of these control areas, that does not mean they have been comprehensively addressed. You will need to continue analyzing other threats to your company, while using a standard such as 17799, to make sure you have addressed all the critical areas of your business. While analyzing each new threat, make sure you are using the techniques discussed in the defense in depth section to help you get the most protection from your controls against multiple threats.

6. Conclusion – Who should get Internet access?

We have covered a rather lengthy list of potential threats that come along with employee Internet access; now it is time to apply that information in your organization to determine who needs access to the Internet and how much access should they receive.

Let me propose a hypothesis. Most people in your organization do not need unfiltered access to the Internet. Only you can test that hypothesis against your organization's culture. Many companies allow reasonable personal use of corporate phone and Internet connections because it can make employees more productive.³⁹ At Google's headquarters, the company goes so far as to provide onsite: lunch and dinner, dry cleaning, a doctor, and many other services to keep employees at the office and working.⁴⁰ When looking at the issue from a recruiting perspective, companies with more open policies could be at an advantage in attracting younger, technology-savvy workers who want access to IM and flexible work arrangements.⁴¹

Problems seem to arise when employees move from reasonable personal use to wasting time and outright abuse of corporate systems. Excessive web surfing, as well as people trying to make money on the side by day trading, selling things on EBay, etc. is usually the exception rather than the rule according to the latest surveys by the Pew Internet and American Life Project.⁴² The debate over whether these activities are hurting your organization is up to management, but many agree that the

³⁹ <http://www.fedsmith.com/article/1298/>, Visited 12/31/07

⁴⁰ <http://www.google.com/support/jobs/bin/static.py?page=benefits.html>, Visited 12/31/07

⁴¹ http://apcmag.com/3113/unlock_work_internet_or_risk_losing_staff_microsoft, Visited 12/31/07

⁴² http://www.pewinternet.org/trends/Daily_Internet_Activities_8.28.07.htm, Visited 12/31/07

benefits of blocking access to pornography, hate sites and malware distribution points outweigh the costs. It is up to each organization to understand what is at risk.

I am going to attempt to classify the majority of users into one of the four following categories:

- Unrestricted access - the only people with unrestricted Internet access in our organization may only use it while performing malware research, and they are forced to use a separate network and Internet connection. Then again, if your road warriors can shut down their antivirus while surfing at the coffee shop, they are almost in the same category.
- Filtered Access - Most people in corporate America probably fall into this category, whether it is the appropriate one or not. In this category, websites and protocols which are malicious, illegal, and those that the organization has determined to be detrimental are blocked.
- Whitelist only – All Internet access is handled through proxy servers and limited to only those sites and protocols that are necessary.
- No access - All access to the Internet is blocked. In some cases, access to internal web sites and email may be allowed. In some environments, a separate reading room area may be available.

There are two approaches to categorizing an employee. In a high security environment, you should start with no access and only give employees access appropriate to the sensitivity of the information they are working with. Most environments use an alternate approach. Employees start in a less restrictive category - such as filtered access - and only if they create problems, do employers begin to restrict what they can access. The first approach is the safer one but unacceptable in many corporate and university environments. By answering the following questions, you may be able to see which approach is more appropriate to your environment.

- Are you using Internet access as a recruiting/corporate benefit?
- Is your IT staff aware of the risks and constantly updating security technologies to match new threats?
- Are you willing to pay the cost of cleaning up infections and the loss of your data?
- Would you rather manage users by exception instead of managing a whitelist of approved websites and applications?

The more of these questions to which you answer yes, the more likely you have assessed your risk. If so, you are able to legitimately start users in a filtered state and then manage by exception if a user demonstrates they are having issues. You can use quotas or whitelisting to enforce policy when necessary. The more questions you answered no - or if you are in a high security environment - the more you need to consider starting with users in the no access category, and then have people justify the access they are given. Universities are a special situation, and users may demand open access to the

Internet, but even they must follow the law. In corporate and government networks, it should be very rare that someone would have unfiltered access to the Internet.

7. Parting Thoughts

Best practices are the ones that apply to a wide array of situations, but they are not a perfect fit for every organization. Nor will they decide for you which assets need more controls and a larger share of a finite security budget. It is easy to get into a checklist mentality where you use a best practice standard or two and some auditor guidance and think you are protecting the company. I hope that you now see the benefits of starting by analyzing your assets and the threats to them. Using threat analysis, you should discover what is important and where your focus should be, instead of trying to spread your investments evenly across all assets. Then, once you know where to focus, refer to best practices to make sure that your analysis is covering all of the relevant business areas and threats.

Finally, once you think you are done, it is time to review your work again. New threats are always appearing, and there is always room for improvement. When conducting your periodic reviews, you may want to obtain some outside validation from auditors or research programs such as the annual SANS Top 20 list.⁴³ One area you can always focus on is policy. I do not know too many people who enjoy reading policies, but a well-written policy can make your users' lives much easier. If your policy can explain why it exists, and can actually be followed by the employees without preventing them from accomplishing their jobs, then you have gone a long way in helping people act in a more secure manner.

The overall goal is appropriate security. If you are constantly evaluating how you are applying controls in response to the current threats, you should be making the most effective use of your budget and personnel. If you can further show the costs of a threat to be less than the controls to prevent it, you are framing the conversation in the language that business can understand, and you will find it much easier to obtain funding for your projects.

⁴³ <http://www.sans.org/top20>