

# The Philosophy of Security



## Table of Contents:

- Introduction ..... 1
- Category Mistake ..... 2
- Ockham’s Razor ..... 3
- Ockham’s Razor Misapplied ..... 3
- First Cause and Trust Anchors..... 5
- Greedy Reductionism ..... 8
- The Verification Problem ..... 9
- Confessions of an Unethical Hacker – Part 1 ..... 11
- Confessions of an Unethical Hacker – Part 2 ..... 11
- Confessions of an Unethical Hacker – Part 3 ..... 12
- People and Technology: An Analysis for Part 1 ..... 12
- People and Technology: An Analysis for Part 2 ..... 14
- People and Technology: An Analysis for Part 3 ..... 16
- How Security Technology Can Help People ..... 16
- How People Can Hurt Security Technology ..... 17
- Summary ..... 20

## Introduction

Many security whitepapers begin with an in-depth analysis of an algorithm or they begin by showing how easy it is to exploit various vulnerabilities. The intention is to scare you into performing the steps outlined by the whitepaper or buy the technology the whitepaper promotes. We are not going to do that here. This introduction to security endeavors to step back and look at security more generally and apply some basic philosophical concepts to help understand security in a more meaningful way. Essentially, we are going to use Holism and apply it to security. What is Holism?

*Holism - In the philosophy of the social sciences, the view that denies that all large-scale social events and conditions are ultimately explicable in terms of the individuals who participated in, enjoyed, or suffered them. Methodological holism maintains that at least some social phenomena must be studied at their own autonomous, macroscopic level of analysis, that at least some social “wholes” are not*

*reducible to or completely explicable in terms of individuals' behaviour (see emergence). Semantic holism denies the claim that all meaningful statements about large-scale social phenomena (e.g., "The industrial revolution resulted in urbanization") can be translated without residue into statements about the actions, attitudes, relations, and circumstances of individuals.* – Encyclopedia Britannica Online

What we will find out is that anytime security is viewed as something other than a holistic enterprise, mistakes can undermine overall security. In short, when we treat security as a holistic enterprise, we find the following:

- People are the problem
- People are the solution
- Security technology can help people make good decisions about security
- Security technology can help when people do not make good decisions about security
- Decisions made by people can render security technology ineffective

A character in a famous movie had the words: "Those who build on people build on mud" right before he met his demise. He was wrong because he underestimated the intense loyalty that a person can feel towards another person. Returning to security, we can paraphrase a more correct saying: "Those who deploy security technology without regard to people builds on mud".

Actually, talking about a specific security technology under the umbrella of the label "Security" is a type of mistake. Let's look at what is called a category mistake.

## Category Mistake

The philosopher Gilbert Ryle formally introduced the concept of applying a macro term to a micro entity as a type of mistake – specifically, the category mistake. A common example of a category mistake is when a tour of a university is given to a new student. The tour guide takes the new student around the various buildings – the "school of engineering", the library, and so on. After the tour is over, the new student says something to the effect of "that was all very nice, but where is the university?" The new student has made a category mistake – they assumed the university was a building (micro) rather than a series of buildings under a common goal or theme (macro).

A similar example can be made with automobiles. Let's assume that you are an automobile mechanic and that you have completely taken apart your car in your workshop. You tell your three-year-old son to come look at Daddy's automobile. After viewing the driveline, then engine, the wheels, and all the various parts of the automobile, your son asks: "But Daddy, where is your automobile?" Your son has made a category mistake.

Security analysts and consultants often make the exact same mistake without realizing it. Continuing with our automobile example, instead of labeling the automobile parts by their common names, let's label them SSL/TLS, Web Services, AES, and so on. A security consultant/developer/analyst making a category mistake will often stop at SSL/TLS and claim that they have found security. This behavior is equivalent to holding up a driveline of an automobile and claiming to have found the automobile.

Everyone reading should repeat the following to themselves:

- Security is not a cryptographic algorithm
- Security is not a network protocol
- Security is not encryption

These are all category mistakes. Security is a holistic enterprise involving people, processes, technology, and how they all interact. Sometimes that is hard to understand and can also be a bit intimidating. With such a definition, how do you know where to start? For example, if you were the

owner of a new business and were concerned about how to be profitable and be secure, everything that you've read so far may not help. So, let's start by making a category mistake. What? Why would we want to do that? Because this category mistake we are about to make will actually help us on the road to developing a more sensible way of talking about security: Security is about people.

In 2006, 42,642 people were killed in fatal automobile accidents in the United States (FARS, <http://www-fars.nhtsa.dot.gov/Main/index.aspx> ). From 1994 to 2006, the rate of traffic fatalities is between 40,716 and 43,510 people, per year (Ibid). Many automakers invest heavily in safety features for their vehicles and these features have saved many lives. However, one can also see that a great deal is missed by assuming that a vehicle's safety features are the only thing important when it comes to being safe on the roads. Far more important are the people on the roads, the training they've had, the decisions they make, and the environment they are operating in. The same is true regarding security. While some may object that security doesn't have much to do with such a gruesome statistic, on the contrary, many of the same technologies used to buy a book or music over the Internet are used by hospitals, police departments, fire departments, and power grids. In short, the very infrastructures that people rely on to help them and keep them safe use the same technologies that make the news for being hacked. Not a comforting thought.

Viewing security as a holistic enterprise is a bit complex and is can be intimidating. Usually, when presented with complexity, people try to simplify it. Whether they know it or not, they are often using a form of Ockham's Razor.

## Ockham's Razor

Ockham's Razor is a common sense principle that basically says the following: If you are trying to explain or predict the behavior of something, use the theory with the least amount of assumptions, everything else being equal. This principle lends itself well to security considerations as it tends to show how flexibility and complexity can be viewed as untested assumptions. For instance, there are a wide variety of ways to secure a communication session. For a given level of security that is desired, these various ways can be compared in terms of their flexibility and complexity. By viewing flexibility and complexity as untested assumptions, Ockham's razor can be applied to eliminate those methods with more untested assumptions than other methods, all else being equal.

## Ockham's Razor Misapplied

There was a popular comic strip in the US called "Calvin and Hobbes" drawn by Bill Watterson. Calvin, a boy of about six years old, would often ask questions that his dad could not or would not honestly answer. Rather than explaining, his dad would invent answers. For example, he told his son that the wind blew because trees were sneezing, or that the sun set in Arizona near Flagstaff, or that the world really existed in black and white until it turned into color in the 1930s. After these explanations, Calvin would breathlessly tell his mom that someday he wished he could be as smart as his dad. When coming to learn a new topic like security, everyone should have the inquisitiveness (but not necessarily the innocence) of Calvin. Unfortunately, after hearing more than a few security consultants and analysts talking over the years, one could come to the conclusion that they were heavily influenced by Calvin's Dad.

Explaining that the wind blows because trees are sneezing is a very simple explanation and would seem to fit Ockham's razor rather nicely, as compared to mathematical weather models. Unfortunately, it doesn't do very well when it comes to predicting behavior, or at least having some good probabilities about future behavior, which is an important part of security. One could argue that the weather man isn't a good model for predicting weather either, but they probably do a better job than postulating that trees are sneezing, so we'll assume that the weather man is better.

To move to a more complicated security example, let's see how a couple of simple mistakes can lead to a misapplication of Ockham's Razor.

Example\_User is a user in the EXAMPLE Domain. This person has two accounts on the Internet for books and for jewelry, 4 email accounts, and is also an Enterprise Administrator for the Example Domain.

Internet Book Store

Login: [login1@freemail.com](mailto:login1@freemail.com)  
Password: 1ReMM&2ndDEVICE#

Internet Jewelry Store

Login: [myUserName@anotherfreemail.com](mailto:myUserName@anotherfreemail.com)  
Password: A\*isBourne\$YETI!

Corporate User Login

Login: Example\_User  
Password: \$M0neyThat'sWhatIWant!  
Domain: EXAMPLE  
Email: [example\\_user@example.corp](mailto:example_user@example.corp)

Corporate Enterprise Admin Login

Login: Example\_EA  
Password: WOW!I'mAnEntAdminForExample!!!  
Domain: EXAMPLE  
Email: [example\\_EA@example.corp](mailto:example_EA@example.corp)

Intranet Web Server

Login: Example\_User  
Password: \$M0neyThat'sWhatIWant!  
Domain: EXAMPLE

All of these passwords and logins are too much for Example User to remember. Example User believes that writing a password down is a horrible breach of security, so Example User decides to do some research into the Internet Book Store and the Internet Jewelry Store and found out the following:

- The servers used to store account information are located in a highly secure building – more security than his company's buildings
- The servers used to handle account information meet higher security standards than his company's servers
- The servers that handle a user logging-in utilize a higher security cipher suite than his company's servers
- The servers reside in a location that is required by law to report any breach in privacy of information. His company was not under the same obligation for its own employees.

Based upon this information, Example User decides to apply Ockham's Razor and this results in:

Internet Book Store

Login: [example\\_EA@example.corp](mailto:example_EA@example.corp)  
Password: WOW!I'mAnEntAdminForExample!!!

Internet Jewelry Store

Login: [example\\_EA@example.corp](mailto:example_EA@example.corp)  
Password: WOW!I'mAnEntAdminForExample!!!

Corporate User Login

Login: Example\_User  
Password: WOW!I'mAnEntAdminForExample!!!  
Domain: EXAMPLE  
Email: [example\\_user@example.corp](mailto:example_user@example.corp)

Corporate Enterprise Admin Login

Login: Example\_EA  
Password: WOW!I'mAnEntAdminForExample!!!

Domain: EXAMPLE  
Email: [example\\_EA@example.corp](mailto:example_EA@example.corp)  
Intranet Web Server  
Login: Example\_User  
Password: WOW! I'm An Ent Admin For Example!!!  
Domain: EXAMPLE

Is this a misapplication of Ockham's Razor?

Analysis: Here we have an interesting scenario. Based upon the research that Example User has performed, one may be confused about whether Example User has done anything wrong. What would happen if a "hacker" broke into the database of the Internet Book Store and came across the email and password for Example User? Example User has revealed critical information to the "hacker" (i.e., for those readers who are unaware, an Enterprise Administrator of an Active Directory environment is an extremely powerful user with many privileges). If it took a month for the Internet Book Store to realize their database had been compromised, the damage that could occur against the EXAMPLE Company could be extensive. Compare that to the original way Example User had the usernames/passwords configured – no company information is revealed should a "hacker" retrieve this information – in other words, "all else is not equal". In short, Example User needs to go back to the first approach.

The first approach doesn't solve the problem that Example User was running into though – too many usernames/passwords to remember. How does Example User solve that? Well, first memorize the Enterprise Administrator login and give it a strong username/password that doesn't reveal anything. Next, write down the usernames and passwords for Example User's personal accounts (e.g., Internet Book Store) and keep them with the same security that Example User provides to credit cards, driver's license, and other personal information – whether that is at work or at home. What? Write them down? Isn't that horrible security procedure? It depends. We are memorizing the critical account (Enterprise Admin) and writing down the passwords for personal accounts that probably use credit cards with fraud protection anyway. Simply protect them with the same care as your credit cards and you should be fine. Alternatively, a file can be created with the passwords and then the file would be encrypted with a pass-phrase. This procedure allows for the passwords to be managed and stored on the computers where the user will be doing the work from.

As we continue to promote security as a holistic enterprise, we've seen some category mistakes that people make and we've seen a person performing incorrect application of Ockham's Razor. Another thing that tends to undermine security as a holistic enterprise is the things that need to be done before security can even begin – things we will call trust anchors.

## First Cause and Trust Anchors

Trust anchors are those things that need to be setup before security can even begin. Many companies promoting a specific security technology often do not talk about trust anchors because they usually require separate out-of-band configuration – this tends to dirty up the ease-of-use message. Another way of explaining trust anchors is through a philosophical concept called First Cause.

Imagine a line of upright dominoes that is so long it would take many lifetimes to watch them fall down from beginning to end. If you were born and lived during the middle of this process, you may begin to wonder what caused the dominoes to start falling in the first place. Essentially, something had to kick-start the process. This idea can be referred to as First Cause or the Unmoved Mover of the very first domino. You can see a similar line of thought in modern science as the Big Bang Theory.

Security has similar questions, but usually they are about trust. For instance, it is all very well and good to talk about a security solution using SSL/TLS, Web Services, Signed XML Documents, Kerberos Tickets, and so on. Ultimately, there is a point where the "rubber meets the road" so to

speak and you often have to dig to get that information out. Here is an example of a [security developer \(SD\)](#) and a street wise [potential customer \(PC\)](#) having a conversation about their remote device management software and its advertised security:

SD: We have an incredible remote device management solution that is completely secure and no one anywhere has anything like it

PC: What security does it use?

SD: Web Services on top of SSL/TLS

PC: How does the device know that it is talking to your management station?

SD: We use SSL

PC: How does the device know that it is talking to your management station?

SD: Um... We use Digital Certificates?

PC: Ah! So my device needs a trusted CA certificate, trusted access to a real time clock, trusted access to a Domain Name Server, and trusted access to a Lightweight Directory Access Protocol Server or Hyper-Text Transmission Protocol server for the Certificate Revocation List or trusted access to an Online Certificate Status Protocol server.

SD: Um...

PC: Well, I'm assuming the device needs to verify that the management station's certificate is valid. I mean it has to make sure the certificate hasn't expired, it has to make sure that the management station's name and network address match, it has to make sure that the certificate hasn't been revoked, it has to make sure that the certificate is being used according to its certificate purpose and so on. The device does do this doesn't it?

SD: Um... Yes

PC: How do these things get configured on the device?

SD: Oh, that's easy – the management station does it automatically!

PC: Don't we have a chicken-egg problem here? I mean how does the device know that the management station is really the management station if the management station has to configure the things that would prove to the device that it is the management station?

SD: Um... I believe you can configure them manually as well.

PC: Oh – that means I'll have to have a trusted administrator configure them with a trusted laptop on a trusted network. I guess we can do that. My device setup is outsourced, but none of these settings really undermines my network security, so I don't mind providing them to my outsourcer. So, the device has determined it is talking to a trusted management station, how does the management station know that it is talking to a trusted device?

SD: We use a proprietary Web Service and keep our Web Services Device Language secret.

PC: Well, that is okay I guess, assuming no one ever figures it out and posts it to the Internet. How do you prevent from even establishing a connection to an untrusted device?

SD: We use SSL.

PC: Yes, we established that. Are you requiring the device to have a digital certificate?

SD: Oh yes!

PC: Ah! So my management server needs a trusted CA certificate, trusted access to a real time clock, trusted access to a Domain Name Server, and trusted access to a Lightweight Directory Access Protocol Server or Hyper-Text Transmission Protocol server for the Certificate Revocation List or trusted access to an Online Certificate Status Protocol server.

SD: Um...

PC: Well, I'm assuming the management station needs to verify that the device's certificate is valid. I mean it has to make sure the certificate hasn't expired, it has to make sure that the device name and IP address match, it has to make sure that the certificate hasn't been revoked, it has to make sure that the certificate is being used according to its certificate purpose and so on. The management station does do this doesn't it?

SD: Um... Yes.

PC: How does the device get a digital certificate?

SD: Oh, that's easy – the management station does it automatically!

PC: Don't we have a chicken-egg problem here? I mean how does the management station know that the device is really the device if the management station has to configure the things on the device that would prove to the management station that it is a trusted device?

SD: Um... I believe you can configure the digital certificate manually as well.

PC: Oh – that means I'll have to have the outsourcer do more configuring. Unfortunately, to assign the device a certificate, I'll have to give my outsourcer access to my Certificate Authority – a definite “no-no”. I'll just have to wait until the device is on my network to assign a trusted certificate.

SD: Um... Okay.

PC: Okay, so we've established a secure SSL connection which has authenticated the device and the management station to each other, how does the web service determine what to do next?

SD: We use user authentication. We have Single Sign On capability. You send us your domain credentials, we validate them and determine what group you belong to and then grant you rights off of that group.

PC: What?

SD: Yes – that way you don't have to remember multiple usernames and passwords. It works just like logging into the domain.

PC: Um – I don't think so. The only two things that know my username password are myself and the Key Distribution Center that is part of the Domain Controller in my Active Directory environment. When I'm authenticating myself, I'm sending over Kerberos Tickets, not my username/password pair. Why on earth would I want to send your device my domain credentials?

SD: Um – for ease of use?

PC: Does your web service support Kerberos tickets to authenticate a user over the SSL channel?

SD: Um – no.

PC: Well, unless my domain credentials are converted into some form of security token before being sent to your device, I'm really not interested. Do you have any alternatives?

SD: Well, we support Role based authentication where an Administrator can specify a username, password, and role.

PC: Perfect. How do the Administrator credentials get configured?

SD: Well, we have defaults for the Administration credentials. You could have your outsourcer configure them too.

PC: Give my outsourcer my device's administration credentials?

SD: ahhhhhhhhhhhhhh!! (Runs screaming from the room)

In short, trust anchors are those things that need to be in place before security can even begin. As you can see, having trust anchors for security can really impact things like ease-of-use and ease-of-configuration. It is very important to understand what needs to be configured in order to establish these trust anchors for the security of a given solution. Also, not only what needs to be configured, but also, who is going to be configuring these items on the device in question. What are some of the trust anchors in the previous solution?

- A secure Public Key Infrastructure (PKI). Easily the most overlooked and hardest part of using SSL with digital certificates. Many corporations who have implemented a PKI have a team of experts that do nothing but manage the PKI and maintain its security. It is non-trivial to do correctly.
- The configurations on both the device and management structure needed to support digital certificates (e.g., the trusted CA certificate).
- The implementation of SSL – is it implemented correctly on the management station and device (e.g., a well tested and supported version of OpenSSL for instance?).
- The implementation of the application that is using SSL – is it using SSL correctly, is the proper SSL version being used, insecure cipher suites eliminated, enforced CRLs, correct time, and so on.
- The configuration of administration credentials on the device.

All of these things need to happen before secure device management can even begin! Hence, why we call them trust anchors. Note that we can also ask the same trust questions about SSL – after all, why should you trust the SSL protocol? Ultimately, it will come down to the type of answers you are satisfied with. Let's examine SSL.

- Used in the industry several years and has gone through 4 different revisions – SSLv1.0, SSLv2.0, SSLv3.0, and TLS 1.0/1.1
- Standardized by the Internet Engineering Task Force
- Widely deployed via OpenSSL and has been reasonably analyzed.
- Supports open encryption and hashing algorithms such as AES and Triple DES.

These seem reasonable answers, but we will talk about this more in the section called The Verification Problem. Back to our potential customer (PC) and security developer exchange (SD), you can see,

just saying “We use SSL” as our Security Developer did is not enough of an answer to really explain anything, much less justify the security claim being made. With our view of security as a holistic enterprise, we can see the people questions – “who configures what settings, where does this configuration take place, when does this configuration need to be done, how is this configuration performed, and what knowledge do I need to give them in order for them to be successful at the tasks they are assigned to do” are very important security questions to answer in our example (Note: this doesn’t mean that cryptography is unimportant)

We found our trust anchors using a methodology known as reductionism. We needed to eliminate some variables and focus in on the things that need to be established that our security protocol for device management has to rely on. However, reductionism can be tricky – there are good forms and bad forms.

## Greedy Reductionism

If you’ve made it this far, you’ve hopefully realized that Security, in its general form, is best analyzed as a holistic enterprise – basically a complex system worth more than the sum of the parts. However, when a certain part of security must be analyzed, some sort of elimination of variables and focusing in on relevant but simpler aspects of a complex security system is required – this would be following a methodology which we will call reductionism. Reductionism is extremely useful for developing explanations and predictions for complicated systems. For us, we are using reductionism as a technique by focusing on a specific relative part of a system that is of interest to us.

As an example, let’s look at the things that someone needs to do to keep their automobile in good shape. They could spend all their energy learning everything about that automobile – essentially all the knowledge that the team of engineers that designed it had and then develop a service plan. An alternative is to make the assumption that things that break down are usually the moving parts. Instead of studying the entire automobile, we can now simply study the moving parts and develop a service plan around that. This would be an example of using reductionism as a technique to help simplify problems (of course, they could simply read their owner’s manual maintenance schedule as well).

However, reductionism can be misused and when it is misused, we will call it greedy reductionism, using a term from a famous philosopher (Dennett). Here is where simplifying things too much results in a type of category mistake. For instance, in the previous example, saying an automobile is “just the sum of its moving parts” would be an example of Greedy Reductionism.

Sometimes security products are marketed with Greedy Reductionism in mind. For example, let’s assume that a company marketed an encrypted hard disk for a printer or multi-function device (MFP). The marketing department for the encrypted hard disk claims that buying this product results in “peace of mind” for your printed and imaged documents because no one will be able to recover your documents using forensics. Unfortunately, this marketing strategy is using greedy reductionism. Let’s look at an actual path of a confidential document stored on an intranet web server:

- A user brings up a confidential document from an internal web server. This user has a meeting and would like everyone to have a printed copy, so the user prints multiple copies. The internal web server obviously has a copy of the document on its hard drive and any backup tapes or DVDs.
- Unless the web browser was using some form of transmission security (e.g., IPsec, HTTPS, etc...), the document probably went over the company’s local network in the ‘clear’ and could be captured. Even if a secure transmission was used, if the trust anchors were not setup correctly, then other attacks can be used to read the document without the knowledge of the server or client.

- If HTTP was used (a popular protocol) to read the document, a proxy server could be involved and there is probably a cached copy of the document in the proxy server's RAM and potentially on the proxy server's hard disk
- There is probably a "deleted" copy of the document on the user's hard drive that was used to render the document in the browser (i.e., a temporary file). *Note: "deleted" is used in quotes to indicate that a normal user believes the file has been deleted, but the file can be recovered via specialty software or forensics.*
- There is probably a "deleted" copy of the spooled print file on the user's hard drive. If network print spoolers (Windows, NetWare, UNIX/LINUX, and so on) were used instead of direct printing, the document was probably sent in the clear to the network print spooler and a copy exists on the network print spooler's hard drive.
- When the user or a print spooler sends the document to the actual network printer, unless the machine was printing using IPsec or another security technology to the actual printer, the print image of the file was probably sent over the local network in the clear.
- There is probably a copy of the raster image on the printer's hard drive.
- If the user forgot a printout (e.g., due to paper jam, too many copies, delayed print job, etc...), there is a paper copy available at the printer. If there was a paper jam, there may be partial copies in the recycle bin after the jam was cleared.
- The user decides that an outsourcer under a trusted non-disclosure agreement needs a copy of the document as well and emails one of the printouts directly to them from an MFP. Unless it is the same machine as was used to print the document, there is probably another copy on the MFP's hard drive.
- The document was probably sent in the clear over email, available to be sniffed.
- The document may in fact be stored by email servers along the way and perhaps "deleted" as well. *Note: These electronic copies are available on servers that are probably not covered by your security policy!*
- There is probably a "deleted" copy of the PDF on the outsourcer's hard drive when it was viewed via email.
- There is probably a "deleted" copy of the spool file on the outsourcer's hard drive when it was printed. In addition, if an intermediate print spooler is used, there is a "deleted" copy on that hard drive.
- The document was probably sent to the outsourcer's printer in the clear and could be sniffed.
- The outsourcer's printer probably has a "deleted" copy of the raster image on its own hard drive.
- If the outsourcer forgot to pick up the printout, there is a copy by their printer. Any problems with the print job, there are probably partial copies in the recycle bin.
- The outsourcer probably saves the PDF file. If it was an internal server, there is probably a copy on its hard drive and potentially any backup tapes or DVDs.
- After the meeting is over, a user inadvertently places the document in a normal paper recycle bin rather than the confidential document bin.

Greedy reductionism will often result in a false sense of security by making security seem easy and not looking at the big picture. Looking at security holistically, one can see that while buying an encrypted hard disk for a printer/MFP may be a good step in certain circumstances, there are also many other ways to obtain these documents as well. If your documents are important enough to buy an encrypted hard disk for your printer, then the security around all the other ways of obtaining the document probably should be evaluated too. For the sake of argument, let's assume that all the previous ways of capturing a document were locked down and a customer purchased an encrypting hard drive for their printer. All is well right? Well, now we can then begin down the road of The Verification Problem.

## The Verification Problem

Let's work through a simple example.

Our imaginary customer is evaluating encrypting hard drives for his printers in the finance department. The customer is confident all other ways of accessing these sensitive documents have been closed and is now trying to close the final way – a forensic analysis of a printer’s hard drive by a hacker that is able to get his hands on one. The customer’s main worry is that the electronics recycling firm being used by him is actually owned by a larger corporation that is a fierce competitor to the customer’s own products, which may lead to some conflicts of interest and hacking opportunities.

The customer purchases four different encrypting drives from different manufactures and places each one in a different printer. After a few days of use, a question occurs to the customer: “How do I know these devices are actually encrypting data? How do I know that they aren’t just regular hard drives with a high price?” The customer decided to run his own tests. He sent each printer the same file – a 500 page ASCII text document filled with the letters of the English Alphabet (e.g., “ABCDE...”). He then removed each drive and placed them one at a time in a free drive slot in his own computer. He then ran some tests.

Hard Drive A: The document appeared to be encrypted, but the meta-data about the document (e.g., author, title, date, and so on) was used for reporting and was not encrypted.

Hard Drive B: All the data was encrypted using AES-256. Unfortunately, the key was simply a SHA-256 (Secure Hash Algorithm with 256 bits of message digest) hash of the hard drive serial number.

Hard Drive C: All the data was encrypted using AES-256. Unfortunately, the key was simply the first 256 bits of the actual data of the document being sent.

Hard Drive D: All the data was encrypted using AES-256 and the customer wasn’t able to find the actual key value. Looking at the manual for the drive, the manufacturer indicated that a random number was used as the key and was unique to each drive.

The customer had a good friend who was also a very good hacker. He gave Drive D to his friend and asked him to find out the contents. In about an hour, the friend returned with the document that was printed. The customer was dismayed. It seems that the company that made Drive D did indeed store a random number for each drive, but they kept track of the actual value and correlated it with the serial number. A disgruntled employee of the company had posted this serial number-to-key database to the underground hacking community.

The customer was upset at what he saw as horrible implementations of security. He immediately went and looked at the manufacturer’s warranty statements. Dismayed, he saw that the hard drives themselves were under standard warranties, but the encryption function was under a “use at your own risk” warranty. Unbelievable! The customer didn’t have a legal recourse, so he thought he would do as much as he could to educate the public about these questionable products, hoping that consumer pressure would result in better products. On his blog, he posted his test results in great detail. He was immediately taken to court by the four encrypting drive manufacturers for violating the Digital Millennium Copyright Act (DMCA) and taken to jail.

What can one do when it comes to verification of a security product claims? That is a very good question. We probably need to start with what standards the product is compliant with, such as Common Criteria Certifications (CCC) and Federal Information Processing Standards (FIPS) as a way of “limiting the field” of products so to speak. It is much better to start with products that have compliance with some security standards; we just need to make sure that we do not stop there. We can then look at whether the product has passed any independent third party testing (by someone you trust of course!). We may want to do our own testing as well (just be careful what you do with your findings). Much like scientific theories, there can never be complete verification of the functions of a security product. Remember, at one time everyone believed the world was flat. Based upon the information they had at the time, that was a reasonable belief to have. As time moved on and more things were discovered, maintaining the world was flat was no longer reasonable. If our best theories can be proven false at any time, including the theories used to develop security products, how can product verification be done? This is the “Verification Problem”. We attempt to combat The Verification Problem with Testability and Falsification. In short, some things that are important are the following:

- Are the claims made by the security product testable? Who tested them? Is there on-going testing? Are the testing results public? Have the claims been independently verified?

- What is the company's response if any of the claims are falsified? Are there legal obligations for customer notifications? For product replacement? For liability?
- Are there clear indications the product is working and doing its job properly? Are their indications when the product is not doing its job properly? Are there diagnostics that can be run to test the product out periodically?

Okay - why are we talking about something so specific when this whitepaper is about security as a holistic enterprise? To ensure that everyone understands that security technology has to deal with the Verification Problem in much the same way as scientific theories do. There may come a day when an announcement is made that the security technology you rely on isn't as secure as you originally thought when you deployed it (i.e., "What do you mean that the world isn't really flat?"). Such an announcement may result in a "cold prickly" feeling rather than a "warm fuzzy" feeling, especially if you relied solely on that technology without regard to the people around it.

The good news is that more than likely your security won't be compromised by the techniques listed in this section. The bad news is there are much easier ways of compromising your security. The really bad news is when Security is not viewed as a Holistic Enterprise, these ways are almost trivial. Let's look at a few exploits of an imaginary unethical hacker.

## Confessions of an Unethical Hacker – Part 1

It was hard for the last few weeks to wake up on a Friday morning and hit the bar, but that is where a person that I will call X and his peers came in after work. They worked the late shift as a clean up crew for Company Y – Sunday night from 11pm to 4am through Thursday night. Friday morning, they always stopped in for a few drinks. I had got to know X and decided the time was right to show him my fake business card – "Certified Ethical Hacker and Licensed Penetration Tester". That always got a laugh. You see, I told X, the company that employs you hired me to break into their network. If I can, they'll give me a bonus. I'm willing to give you that bonus if you help me. If you are caught, it is okay as I'll simply say that you work for me, and they've promised me that nothing will happen – after all, they are paying me to do this. X seemed skeptical, but after I told him how much the bonus was and showed him my fake contract, he was all for it. It is really simple, I told X, just go by each printer and MFP they have, get the documents that are in the "to be picked up" pile – you know, the documents that people have printed and have forgotten to pick up, place them in an MFP, send them to this email address, and then put them back where they were before. That is it – you don't need to take anything or even do anything illegal. Do this once a week, preferably on Friday, for a month and the bonus is yours! I even showed him a video on my laptop of exactly what he needed to do on the control panel of the MFP – basically put papers in the scanner, press the "email" button, type in the email address, and then hit "start". A month later, I had a lot of information for that company's competitor – quite amazing what employees print out and don't ever pick up.

## Confessions of an Unethical Hacker – Part 2

I love Halloween. Company Y has a few buildings, a few hundred people, and they always have a Halloween get together where everyone dresses up. The day is pretty easy – not much real work getting done – and the vast majority of people don't actually know each other. People bring their kids in, have some fun, play some games, and rarely are ever at their cubes. I always show up a bit early dressed up like the Headless Horseman – we'll, with a pumpkin as a head – since I don't want anyone to know who I am. I'm carrying a lot of trays filled with cookies – not because I'm a sweet guy, but because I need someone to open the door for me. I have an employee badge – not a real one, but a fake one. It doesn't work on the card control, but I have it hanging around my neck anyway. They are so easy to fake with modern color printers and most employees will leave them on the dashboard of their car while at the gas station or local grocery store – so I know just what they look like. Just have the "badge" hanging around your neck and have your hands full and the door will get opened for you. Everyone is so helpful. Once inside, I just walk around – check out the organizational charts posted everywhere and find where the managers are sitting. I plant a few

keystroke loggers – I’m getting pretty good at it – in and out of their cubical really fast. They aren’t ever in the cubicles – they have celebrations to go to! Then I wander around to all the buildings and eat all day on the trays of food people have out. Most people leave early on Halloween – got to take their kids Trick-or-Treating or be at home to hand out candy. That’s when I go back and collect those keystroke loggers and head home. Well, not home really – I just need to find that insecure wireless network in a suburb and use it to access the Company Y’s VPN – their VPN endpoints are in DNS which makes it easy. They use SSL, but only do server authentication. Their firewall has a cut-through-proxy feature that allows them to enter their username and password, and I have plenty of those.

## Confessions of an Unethical Hacker – Part 3

X was the head of the finance department and lived in the hills, at least according to the white pages. The bad news for him was that he lived in an area with no broadband connectivity. I expected to see him a lot at the company’s main site, about an hour from his house. But, after a few days of watching, I only spotted him once. Looking at the yellow pages, I saw that Company Y had a remote office about 20 minutes from his house. I decided to investigate. Sure enough, it looked like X was stopping in there and doing his work and only coming into the main site once a week or so. I could see his office from a local café which had free Internet access. Looking at the wireless access point in the cafe, I could see that the café was on a cable broadband modem. Teasing a tech-savvy clerk a bit about cable modems and load sharing, she responded that DSL wasn’t out there yet and cable was their only option.

I looked a bit silly in overalls, with my name tag “Jon”, and my toolbox, but I figured it would be effective. At lunchtime on a day that X was at the main site, I stopped by after disconnecting the outside cable line. “Networking problems – dispatch told me to check it out – luckily I was right next door”. Cool! “Can it get to your networking equipment?” – Yep – right over here. In a small wiring closet, I connected my access point to a mirrored port on the switch I configured. I verified I could connect (securely – I don’t want anyone else to do that!) and went back outside and connected the cable. Everything was fine and I was a genius, at least according to an employee that was working over lunch on a critical issue. I had fixed it so fast they didn’t even have to report the problem to their IT department! Yea! Back at the café, I connected my laptop wirelessly to the access point I placed on their network and verified I could capture packets. I’ll be doing the same thing tomorrow when X shows up.

## People and Technology: An Analysis for Part 1

Did our imaginary unethical hacker seem to possess a lot of technical knowledge? Not really. He was an exploiter of people and used that to gain unauthorized access. Once access was gained, there was simply no security to block him.

We started our discussion of Security as a Holistic Enterprise by knowingly making a category mistake. We said Security is about people. In fact, to repeat from the introduction:

- People are the problem
- People are the solution
- Security technology can help people make good decisions about security
- Security technology can help when people do not make good decisions about security
- Decisions made by people can render security technology ineffective

Let’s look at our imaginary unethical hacker’s first confession. This confession had an unauthorized person digitally sending documents to a competitor. Someone technology focused may say: “We require domain credentials to be entered in order for digital sending to take place. Problem solved!” What would someone people focused say? Let’s start with some observations about people printing in the workplace:

- People print documents and then get distracted – a phone call, a meeting, and so on and forget to pick up those documents.

- In many businesses, there is a good distinction between super secret documents and documents that are not confidential. Unfortunately, most documents fall into the grey area in between. In fact, without proper identification, there may be a debate between two peers on whether a document is confidential or not.
- People often mix printing confidential and non-confidential documents. For instance, printing the latest Dilbert cartoon to post at your cubical or banners for the holiday party. This intermingling of business confidential documents and non-confidential documents often result in the business confidential documents being mistakenly treated as non-confidential.
- In fact, usually the due diligence that a business would like to see performed for its business confidential documents is often performed for an employee's personal data instead. For instance, the activity known as Print & Sprint is more likely performed when an employee is printing their stock share plan performance summary than with a confidential internal reference specification.
- Many individuals with a variety of different levels of access to confidential documents often use the same printers to print them out. An intern from college doing research and printing out publicly available documents as compared to a chief technology officer printing out the latest prototype design of a new product.
- Many companies encourage environmentally conscious behavior – often placing recycle bins directly next to printers. Often, partial documents that were part of a paper jam are often placed in the recycle bin. Sometimes, documents that haven't been picked up and are taking up space are placed there. If these aren't recycling bins, they are usually trash bins. The confidential bins are usually further down the hall.

No wonder people would rather talk about technology solutions – people solutions are hard! There is only one problem: the technology solution of requiring domain credentials to digital send doesn't actually solve anything. First, let's argue with the technology focused solution on its own terms:

- It is never a good idea to supply your domain credentials to a computer that isn't a member of your domain (remember our Ockham's Razor example). In fact, it isn't a good idea to use your domain credentials on any computer that isn't the one you work with on a daily basis. Unfortunately, domain credentials have become the new "Driver's License" of identity in the workplace, often being used in places where they shouldn't be used.
- Many domain credentials are long, full of special characters, and are difficult to type in on the Mini-Me style of keyboards in use by most digital senders. As a result, expect a jump in fax machine usage over digital send – in short, employees finding ways of bypassing your security.

Now, let's cut to the chase:

- Problem Statement: There is an unauthorized person in physical possession of confidential documents. They can simply take them, copy them and take the copy, fax them, throw them away in a specially marked trash bag for pick up later. Requiring domain credentials to digitally send doesn't address the issue anymore than having an encrypted hard disk would. Imposing rules on employees, posting signs to pick up your documents, automatically shredding documents after 6pm and so on will not really solve the problem either. People tend to go back to their old ways pretty quickly.

A reasonably simple approach is to place printers and digital sending devices in an employee badge accessible room, with a glass door, and with a confidential bin. We can come to this conclusion not because of what we know about technology, but because of what we know about people. What are the benefits of such a solution?

- It keeps employees productive. A badge accessible room is a minor inconvenience to employees. There are no special ways to print, logins, or rules to follow (or rules to try and

bypass). Most employees walk to the coffee/tea station more times a day than to a network printer.

- It provides the ability to audit access to those devices.
- It provides the ability to control access to those devices.
- It provides a constant reminder to employees about document security.
- Most importantly, it solves the actual problem.

If you value your printed documents and there are unauthorized individuals that can easily access your printers consider treating your network printers/MFPs like you treat your internal web servers or your LAN switches, not like you treat your coffee stations.

## People and Technology: An Analysis for Part 2

Physical access security personnel have often been cut way back in our cost-cutting business climate that we operate in. In particular, the individuals that monitor incoming traffic to a business, monitor the doors in which entry can be obtained, and patrol the parking lots are seemingly on the decline. With access controls being tied into employee identification badges, a new motto is being preached: "Security is every employee's responsibility." These two things have combined to justify the reduction in physical access security personnel. In our imaginary unethical hacker's second confession, he uses physical access to a tremendous advantage and completely goes undetected by employees. If everyone is responsible, how did our unethical hacker succeed?

- It is a common mistake to think that employees at a site of more than one building actually know everyone. Many employees only know their team members or former team members really well. Members of other teams on other floors of a building or in different floors of the same building don't really know each other well. In other words, it is okay to be unrecognized.
- Halloween and Christmas tend to be times that businesses in the United States have a lot of festive things going on at work. During these times, employees tend to be more helpful and friendlier. Halloween even offers the opportunity to disguise your identity and you are usually encouraged to do so.
- Many employees are not thinking about security when they are walking into work. Instead they are talking with teammates, thinking about a problem they have to solve, thinking about things they need to do. While they may think to check for a badge, they most certainly don't examine it in any great detail.
- Usually, employee identification by other employees is primarily through visual recognition of their employee identification (e.g., badge). It is not via the following: "Let's walk 100 yards so you may place your badge on this card access control panel so I may verify that you are an employee using the security technology in your badge". Since employee to employee identification is primarily visual, many types of employee identification can be faked to appear genuine in most situations.
- At many sites, once an employee has crossed a badge control boundary that uses security technology (e.g., an employee only entrance to a building), they no longer use their badge to access anything. In short, there is only one technology barrier to overcome.

The problem we are trying to stop is what is referred to as tailgating. A successful tailgating operation by an unethical hacker can severely compromise your network and the resources on it. What our imaginary unethical hacker did was very similar to what law enforcement officials do when people are suspected of computer related crimes – they get a warrant and install keystroke loggers. Our imaginary unethical hacker had to do everything in one day. A helpful employee on a single day in the year can fully compromise your network.

What we want to do is ask ourselves a question: "What can technology do to help people make better decisions in regards to security, specifically around tailgating?" The fact of the matter is that

telling employees “Don’t do this” and having the technology deployed in such a way that it allows them to “Do that” very easily isn’t going to work. Especially if “doing that” involves helping people. Let’s go through a sample analysis assuming that employees have identification badges that have security technology for card access control:

- Identify every exit where employees are solely responsible for making decisions that could allow for tailgating. These will be our initial area of focus
- The initial sliding door or revolving door is badge controlled but can allow for more than one person to enter.
- Once inside the main door, install two employee badge controlled turnstiles, one five yards in front of the other one. They are not side by side, but instead form a line for a single line of employees to use.
- When the first one is activated by a badge, the turnstile allows for one person to enter. The first turnstile then flashes “wait” until the second turnstile is activated by the same employee badge that activated the first one. At that point, the second turnstile allows for one person to exit. Once a given employee’s badge has operated the second turnstile, it will be 15 minutes before it can operate the first turnstile again.
- Once the person has exited the second turnstile, the first turnstile will allow another employee badge immediately, as long as it wasn’t the same badge that operated the second one.
- Turnstiles are monitored by a security camera
- There is a security button by the first turnstile that can be pressed to indicate a security violation has occurred (e.g., an employee just saw someone hop over both turnstiles).

What have we done in our hypothetical analysis? Well, we’ve added some inconvenience to employees with badges, but not really a lot. It may take them a few seconds longer to enter the building. We haven’t achieved the security of Fort Knox – I mean someone can just hop over the turnstiles, but we weren’t trying to deploy the technology of Fort Knox. What we are trying to do is allow people to be people but use technology in such a way that it helps them make good security decisions. Given an employee has been educated on the dangers of tailgating, think about what they would have to do to help someone get in the building without the card access control of an employee badge working:

- “Oh, your badge isn’t working? Well I guess you can hop over the turnstiles. I won’t tell”. This response is not too plausible a response for someone educated in the dangers of tailgating.
- “Here, let me get you through the first turnstile, then I’ll throw you my employee identification and so you can get through the second turnstile and then I’ll wait 15 minutes so I can get in”. This response is not too plausible for any employee.
- “Jump on my back and we’ll go through together”. This response is not likely but it could depend on who is asking. Luckily there is a security button for other employees to press when they witness such a violation.

Better yet, let’s review what our helpful employee might say to our Headless Horseman coming in from the rain loaded with cookie trays on Halloween:

- “Here, let me take half of these cookies so you can use your badge to get through the turnstiles. I’ll carry the other half and meet you on the other side”, says the helpful employee.
- “My badge doesn’t work. It got caught in the car door just this morning. I guess I’ll need to go to the main lobby. Thanks anyway!”, says our frustrated headless horseman.

Our hypothetical solution is just that – hypothetical. But, it shows the kind of thought process that allows technology to be deployed in such a way that helps your security and allows people to be people.

## People and Technology: An Analysis for Part 3

In our imaginary unethical hacker's third confession, we can see he is pretty smart. He's created a problem and showed up to fix it. If you've ever seen an employee's reaction to the network going down, it is quite similar to a hungry person's reaction when their food gets stuck in a vending machine. The higher the priority of items that an employee is working on, the more stress a network outage causes. Stress usually causes poor security decisions to be made. Our unethical hacker has created a situation in which he has just been granted the authority to do just about anything with the networking equipment on that remote site. Any technology that can be deployed may prevent some attacks (e.g., 802.1X), but may not prevent others (e.g., keystroke loggers). There are too many attacks possible for someone with physical access to your networking equipment and more than likely he is not being monitored – or if he is being monitored, it is by individuals without any technical knowledge of what he is doing. Remember, the employees want to help him.

Training often needs to increase substantially for remote office employees – verification of service personnel using the yellow pages, their name, and any type of identification possible (e.g., description, badge number, and so on). There are some things we can do to help remind our employees. For instance, if the LAN equipment is locked up, rather than simply putting the key on a ring with other keys, a separate box for that key could be used – with the words “Call IT Security at 123-456-7890 before using this key” printed on the box. Signs on the locked door could also say something similar.

The bottom line is that a good unethical hacker is going to use skills that allow people to compromise technology. They can do it through induced stress or through using the helpfulness of people against them. This isn't to say that they don't use technology to exploit vulnerabilities – it is to say that some of the most devastating attacks may not involve cracking the technology at all. Putting people in a position to be successful under such conditions requires a lot of work in itself.

## How Security Technology Can Help People

After reading this far, one may get the impression that this whitepaper is anti-technology. It is not. It is striving to recognize the proper place of technology for security strategies rather than placing technology on a pedestal as a solution regardless of what people do. Let's go through an analysis of a workplace situation in which technology can help.

A small company with about 50 employees has standardized on three MFP models to handle their printing and imaging needs. To save costs, they also standardized on laptops with docking stations for personal computers. From a physical access control perspective, the company's building is badge accessed controlled and their LAN equipment and servers are in a locked room controlled by their IT department. About 15 of these employees are working on a next generation product that is critical to the success of the business. The MFPs are serviced by an outsourced company. This outsourced company keeps the MFPs up and running and deals with supplies for the next two years. The IT department believes it is a good idea to protect company's intellectual property by purchasing encrypting hard drives.

Here is a very plausible case where a company may want to deploy an encrypted hard drive. However, there is more to do.

- Make sure that they have a contract with the outsourced company in regards to legal liability concerning obtaining or distributing information. Remember, a person with authorized access to the MFP's hard disk drive also has access to the printed documents that have not been picked up, the recycle bin, and any other non-volatile storage.

- Many devices use a simultaneous combination of hard disk, flash, EEPROM, and other technologies to store a variety of different types of information. An encrypted drive may protect some information placed in non-volatile storage, but not all. These are important questions to ask the MFP manufacturer:
  - What information is stored in non-volatile storage?
  - What types of non-volatile storage is in use?
  - What information is stored where?
  - Which non-volatile storage has encryption?
  - Which encryptions meet external specification (e.g., FIPS)?
- The company should determine who manages the equipment/IT of the servers and laptops. If this is an outsourced or external company (e.g., retail service), then steps need to be taken to secure these paths to their data as well.
- The company should also evaluate their laptop security for their drives. There are fifty laptops and only three MFPs. Laptops are often removed from the building which has access control and taken to places that do not have access controls. Hence, why stolen laptops make the news quite frequently.
- The company should also evaluate how they are backing up their laptops and servers and the availability of that information (e.g., DVDs lying on a desk, and so on)

Let's look at another case, based upon the same scenario:

A small company with about fifty employees has standardized on three MFP models to handle their printing and imaging needs. To save costs, they also standardized on laptops with docking stations for personal computers. From a physical access control perspective, the company's building is badge accessed controlled and their LAN equipment and servers are in a locked room controlled by their IT department. All of their laptops and servers have encrypting storage systems and their backups are encrypted and securely stored. About 15 of these employees are working on a next generation product that is critical to the success of the business. All computers and MFPs are managed by an internal IT team staffed with employees of the company. The IT department believes it is a good idea to protect company's intellectual property by purchasing encrypting hard drives for their MFPs.

Here are some situations where having encrypted hard drives on the MFPs may help protect the company just outlined:

- Theft: The MFP itself or the hard disk drive of the MFP is stolen.
- Warranty Replacement or Upgrade: The MFP is replaced or the hard disk of the MFP is replaced due to failure or upgraded to another type.
- Selling equipment to another user/company: The MFP is sold as a used device.
- Recycling this equipment: The MFP with the hard disk is recycled.
- Throwing the equipment away: The MFP with the hard disk is thrown away in the trash.

Technology can help people when they make the wrong decisions - like forgetting to lock a door or mistakenly throwing away sensitive equipment. It can also help when in defense-in-depth situations when other security measures have been defeated – for instance a break-in which sounds an alarm but the thieves are able to escape with valuable information.

## How People Can Hurt Security Technology

Let's move away from the encrypted hard disk example to a technology like SSL. A person may object to the previous analysis and say "Look at SSL – it was a security technology and it changed the way people shop and allowed for e-commerce – people aren't really involved in SSL and therefore decisions that people make can't hurt SSL." SSL is a technology that did allow for a new consumer shopping era to be ushered in. Unfortunately, people can make decisions to undermine security

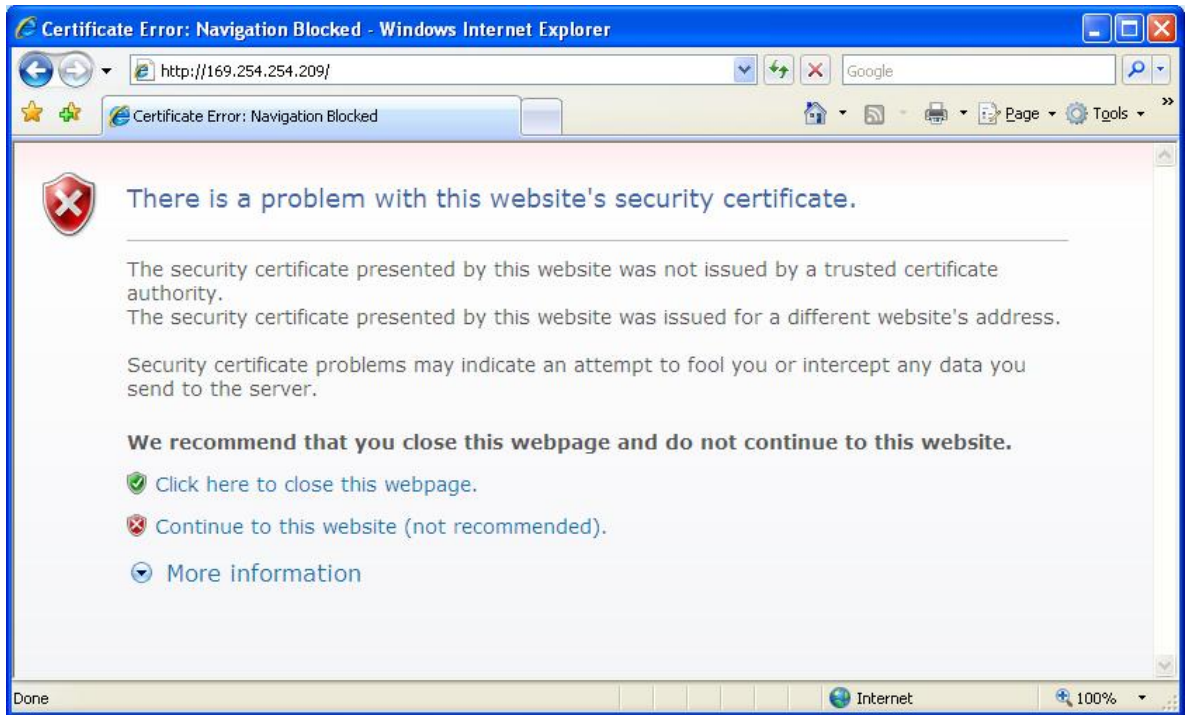
technology and Internet security around SSL is attacked in precisely those areas. As a consumer, you can ask yourself several questions that relate to SSL:

- If I only use SSL for a couple of secure shopping sites, why do I automatically trust more root CAs than I need to? Shouldn't I remove some of them?
- Why should I support SSLv2.0 if my secure shopping sites offer TLS support?
- Why don't have I CRL checking enabled?
- Can an insecure cipher suite be used in my SSL communication session?
- Have I actually clicked on the Lock Icon at the bottom to see what is being used?
- Have I actually verified the site's certificate when presented with the opportunity to do so?

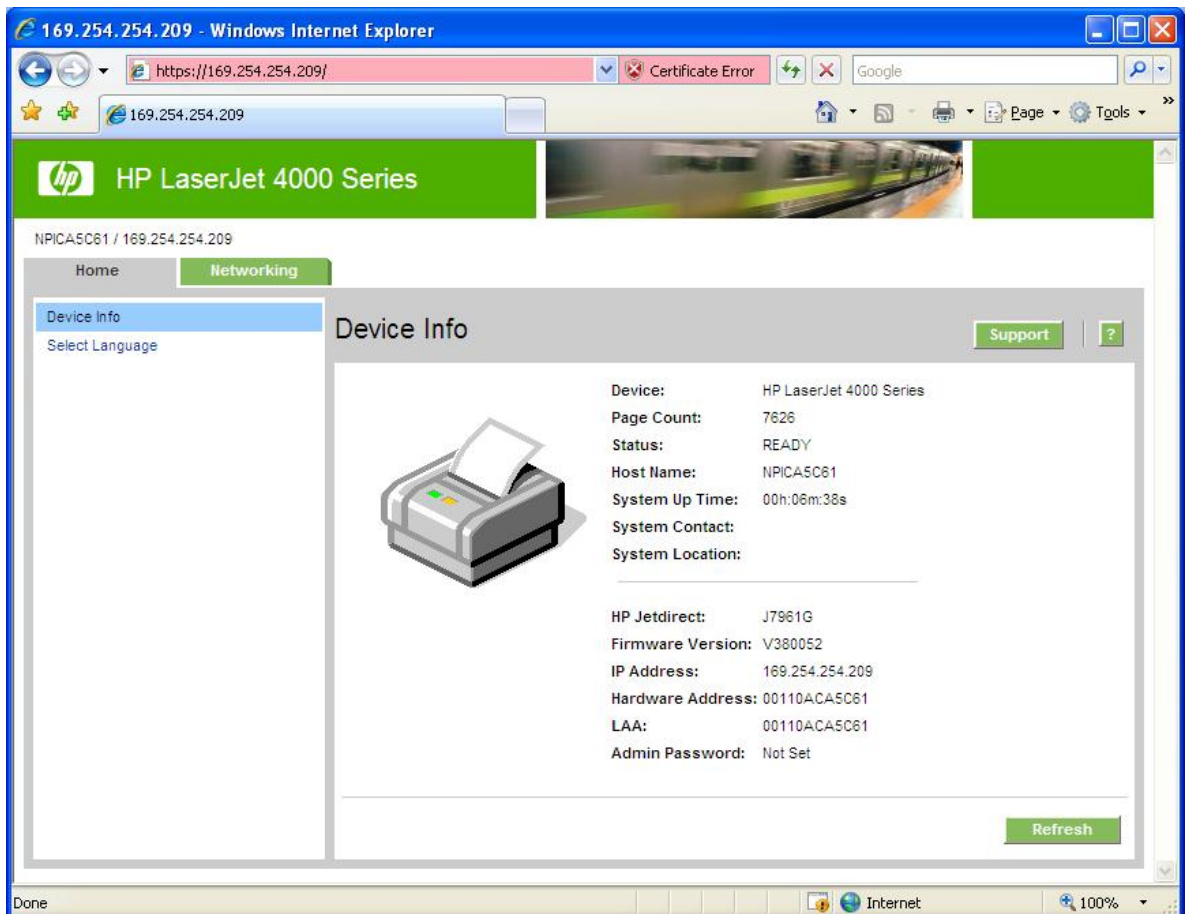
An unethical hacker could use technology to direct a user to a false web site when they are thinking they are going to a trusted website. The Internet Explorer 6 experience when an untrusted digital certificate is a pop-up dialog like this:



In many cases, a user may just click "Yes" without realizing what they are doing and then provide the unethical hacker with a lot of information – like their credit card number or sadly their domain credentials. After all, it really seems like just an annoying dialog. Luckily, the Internet Explorer 7 experience is different in a profound way. Here is an example:



This is a lot different – notice the symbols and explanatory text. The way the information is now presented, it will grab your attention. If we click the “Continue to this website (not recommended)” link, we get this:



Notice the **red** URL and the “Certificate Error” message. Why did Microsoft change the behavior so drastically? Well, because people can make decisions that hurt their security, even when they are using SSL. By moving to a different way of presenting this information to the user, they are helping the user make good decisions around security. And with that, we’ve come full circle.

## Summary

Many books have been written about security in regards to technology such as how to secure your networking equipment, how to test for vulnerabilities in technologies, how do deploy patches across the enterprise, and so on. These are all important topics and require dedicated people to implement and maintain. This whitepaper took the approach of stepping back and looking at security as a holistic enterprise and in doing so, finding more meaningful ways of using technology to help people achieve security solutions. Hope you enjoyed it!