Introduction to

Information Security Concepts

*A Guide for Administrators and Home Users on the design
and implementation of security for your network.*

By: Robert H. Williams III
CompTIA A+\Net+\Security+, Microsoft MCSA:Security\MCSE:Security, Cisco CCNA

Version 1.7.12.7

Introduction...

Information Security (IS) is one of the most misunderstood things within the *Information Technology* (IT) world right now. You often hear people talking about how X is more secure than Y because of some misguided logic, or how you should run weekly virus scans to protect your system as if that was all security was about, but few people seem to really grasp what's going on. And this is a shame, too, cause in truth the basic principals of security are just so simple anyone could learn them. This paper is designed to help you with just that. It will cover what security is all about, the steps you should take to protect yourself, and how to recover if there is a security issue. This paper is designed to be usable by someone of any skill level, from home users to those who already have a CISSP security specialization. To help with this, I've placed different markings throughout this paper to designate the the audience:

Top Secret - This marking denotes a more advanced level of information. Most of the information within these parts will not needed by home users. It may cover network configurations of hardware the average home user doesn't have, or it may just go farther in depth than most people would care to.

Home User - This marking tells you that the information here is aimed more for the home user. It will generally follow the "Top Secret" tags, and offer a simpler recap of the same information. Bundled tags like this well be denoted with a (1) and a (2) to let you know they are connected. If you are an advanced reader but do not fully grasp something covered in the "Top Secret" tags, check to see if there's a "Home User" tag following it. If this paper seems to be a bit too much for you, scan it for the "Home User" tags, and read those first.

Everyone - The "Everyone" tag is, as you can guess, information that applies to everyone. This includes what *Information Security* is, the goals presented within IS, and the basic steps needed to protect your data.

Advanced **-** The "Advanced" tag does not show up on it's own, but denotes when one of the other three tags is a bit more in-depth than normal. An "Everyone - Advanced" tag, for example, is aimed to be understood by both home users and IT professionals, but goes into a bit more depth than they would really need. "Advanced" topics are not harder to understand, they are just sections that contain information that isn't needed to be known, but could help with your understanding. You may wish to skip these sections during your first read of this paper.

Besides these three tags cover what groups of people should be reading what sections, this paper is split between two basic sections, design and implementation. Design should be the first step in any security plan, however, for home users, there isn't all that much of a design required. Since the home computer contains all of your data, all home users need similar security designs. It's only when you start getting into more complex networks that design is really needed. But the design section also covers concepts that home users should read about to help them understand exactly what their actions are doing.

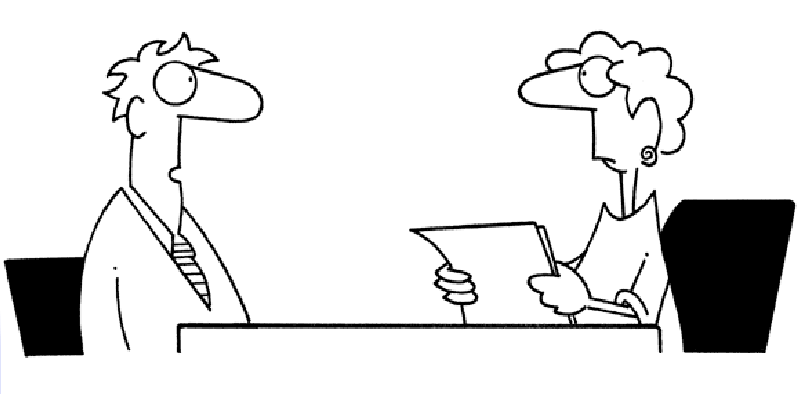

"For the sake of information security, everything on my résumé is false."

Design Topics

What security is about, how you should plan your security, and why you need it.

What is *Information Security (IS)* about? *Everyone*

Information Security has three primary goals, known as the security triad:

- ***C***onfidentiality – Making sure that those who should not see your information, can not see it.
- ***I***ntegrity – Making sure the information has not been changed from how it was intended to be.
- ***A***vailability – Making sure that the information is available for use when you need it.

As you can see, the security triad can be remembered as the letters *CIA*. These principals are simplistic when broken down, but when you think about it more in depth, all steps taken within security are to help complete one or more of these three security goals.

When most people think about Information Security, they will generally only think of the first item, *Confidentiality*, and for good reason, since that's all the media seems to think security is about. Confidentiality is also, ironically, the one of the three goals you most often do not need. A public web-site does not want to be confidential, it would defeat the point of being public.

In order to promote Confidentiality, you have several tools at your disposal, depending on the nature of the information. Encryption is the most commonly thought of method used to promote Confidentiality, but other methods include Access Control Lists (ACLs) that keep people from having access to information, using smart cards plus pin numbers to prevent unauthorized people into your building and looking around, or even explaining to your employees what information about the company they can and can not disclose over the phone.

Integrity is the part of the triad that affects the most people in the IT world, but few seem to notice it, and fewer still think of it as a security issue. The files on your operating system must maintain a high level of integrity, but worms, viruses and trojans are a major issue in IT, and can also be a way that an attacker can get information out of your network, or inject his own information into it. And integrity is not just about malicious parties, it also covers items such as disk errors, or accidental changes made to files by unauthorized users. Access control lists (ACLs), physical security, and regular backups all fall under integrity (And sometimes confidentiality and availability. One fix can solve multiple problems).

Availability is the part of the triad most administrators have to worry about at work, and with good reason. It's the most common, and most visible, part of the security triad, and it is part of the job duties of just about every administrator, even non-security based ones. It's mostly about system uptime for them, but it can also cover subjects such as accidentally denying a user access to a resource they should have, having a user locked out of the front door because the biometrics does not recognize his fingerprints (False negative), or even major issues such as natural disasters, and how the company should recover in case of one.

66%

The percent of people who have seen co-workers store password information at work, such as on Post-It notes.

Source: www.rsa.com

How do I protect my information? *Everyone* - *Advanced*

Now that you know the goals of security, you may ask: "how do I apply them?" Well, first, you must decide what needs protected. In other words, you need audit all of your assets, from information stored on servers to physical items such as staplers, if your duties call for it. This topic will be covered with more depth later, but right now, we'll keep it simple. Since most people reading this are applying the principals here just to information security, we will first focus on information classifications. There are many different ways of classifying information, but many of them follow the same basic principals. According to Microsoft's view of information, there are four types of information:

- ***P***ublic
- ***I***nternal
- ***C***onfidential
- *S*ecret

While it may not be as cool as remembering ***CIA***, the word ***PICS*** should help you remember these four data types. But remember, while Microsoft and others use these classifications of data, not all groups follow this as a standard. In other words, it's just not as wide spread as the talk about the CIA model, and some companies may use their own models.

Depending on the type of data, security is compromised just by exposing the information to others. With other types of data, however, damage is only done if the data was altered or unavailable. Here is a more in-depth explanation of the four major data types:

Public data is designed to be shown, so there is no reason to protect it from being seen, and thus confidentiality is not a concern. If Public data is changed or destroyed, however, you lose something you can remember by the letters PTR, or ***P***oin***T***e***R***: Prestige, Trust, and Revenue. Public data needs to be accessible, but only a few users or machines should be able to change it. Examples of Public data for businesses may be information on your company web site (this does not include any configuration files used by the website however), or any documentation sent to all consumers of your product or services. For home users it may be your personal homepage, or something akin to a myspace page. While it would do no harm for this data to be seen by others, if this data was changed in transit, the results could be disastrous. And funny. Must mostly disastrous.

Internal data, also called Private data, is data that company workers generally know, but outsiders should not know. It's items such as PINs (Personal Identification Numbers) for doors if everyone shares the same pin, the location of some rooms within the building (such as server rooms or wiring cabinets), or internal procedures of the company. It's information that most company workers can find out, or may even need to know. Discovering this information is normally not a risk in itself, but it allows for better attacks. The main risk is modification, either by an outside force such as an attacker, or most cases, accidentally by an internal user (How they can accidentally change the location of a server room, I will never know, but after some stories I've seen, I'm sure someone, somewhere, has done it). Security breaches of this type of information will generally affect the operations of a business, and not much else. Most files on your OS would actually fall under this, as damage to them will only affect operations. Keep in mind, however, that internal data can also be a stepping stone to launch attacks on other, more secure, forms of data. On the flip side, removing internal data from the view of workers can cause damages to business operations, performing a form of Denial of Service (DoS) attack. For a home user, Private data could be where you store your keys, security codes for home security systems, to even less obvious items. While knowing your pet's name may not seem like any sort of security risk, if you happened to use your pet's name as a prompt in case you forgot your password, this could raise some security concerns. And if someone offered you a present, wouldn't you tell them your birthday?

Confidential data is the data used by a limited number of internal users, and should not be known to the majority of workers. This is the class Human Resources (HR) data and payroll information falls under. Read access to this data is limited to a few users, and write access is generally restricted even more. If this becomes public internally, Operations and Internal Trusts are at stake, while if reviled externally, you once again lose *PTR*, along with Operations and Internal Trusts. OS files dealing with security also fall into this area in most cases. Confidential data is just a few steps away from Secret data, and like Secret, it needs to be protected. For a home user this could be some emails you've wrote, your browser history, or a folder containing pictures and movies the rest of the household wouldn't approve of. And you know you have those.

Secret data is the data most people think of when they hear about breaches in information. This data is your trade secrets, intellectual property, and *External Secrets*, such as info held in trust for others (partner company's, or customers). Loss of this data may cause critical damage to the company, and could very well be the downfall of it. Besides the PTR loss, and maybe loss of operations, there's fines and legal actions to think of in most cases.

While this may seem like only businesses would have data that fall in these four classes, all information can be placed inside them, sometimes into more then one class. As stated before, most of the files used by your operating system would fall under Internal data. It's not something that needs to be kept secret so much as needs to be kept from being changed. Music files on your machine? They have an effect on the operation of how you run your life, and so fall under operations. Credit card information could be considered secret data as well.

How is the security of the average company setup? *Everyone - Advanced*

When you hear about all of these security breeches in the news, it's no real surprise, since most businesses have a very crude design when it comes to security of their network. Many businesses, and almost all home users, have a security model that resembles an egg shell. While the outer surface is hardened to keep people out, once there's a breach, there's nothing internal to prevent or limit access. And unlike the mighty chicken, companies have no way to make this security egg shell without holes in it. If a computer is online and networked, then any flaws in the firewall, along with any programs that access the network, form the security holes. If the computer is not online, physical access is still the hole in the shell. If it's a server in a locked room, the door itself is the hole in the shell.

The area of the shell you can attack is known as the *Attack Surface*. As you can imagine, a smaller attack surface is better, but this is only a tiny part of security. Most people and many companies think that the attack surface is all there is to security. This is putting all your egg-shelled networks in one basket. While it may be hard to get into the network, once in, everything is at the mercy of the attacker. It only takes one hole in the shell to get in, and as stated before, you can not prevent all holes in a shell.

The primary goal in IT security is to limit, not prevent, damage. While preventing damage is great, you can never prevent 100% of damage. Rather, you must try to make it harder to cause the damage, and work on lowering how often that damage can happen. While this concept is covered more in-depth later, a quick overview is as follows: all damage done can be assigned a monetary value. If your home computer was completely destroyed by fire, how much would the real damage be? While there is the cost of the computer itself, there is also the data inside the computer, damaged. Even if you recovered the physical cost of the computer, the data contained within could have more value of the computer itself. However, if you had an up-to-date backup of the data that was not destroyed in the fire, then the damage done would be the cost of a new computer, and effort required to restore said backups While you have not lowered how often the event may happen, you have lowered the damage done by it.

The egg-shell rule, while it applies to most companies, applies to home users too. Think for a moment. Chances are, you're browsing the internet right now, or will be later today, from the "internal servers" on your home network, ones with secret level data, as an administrator! And most likely you also have programs running as servers, such as yahoo messenger, with full administrative rights, and no damage control.

Now at this point I must state this clearly: While having a strong shell is not only desirable, but almost required for security, other security steps must also be taken. I am not saying you do not need a "shell" around your system, and I am defiantly not saying you do not need a firewall. I'm saying that this is only the beginning to security, and that other steps must be taken. These include secondary shells under the first, access controls, not having systems trust other systems (Or even allowing one subsystem to trust another sub system). And perhaps one of the biggest issues with companies, making sure security procedures are followed. While you may have designed a great security setup for all your machines, are your administrators in remote offices really following them? Do the administrators in your data room know how to react in the event of a crisis? If you walked into your server room, and just unplugged a server at random, what would happen? How quickly would your network recover? This is something you need to start thinking about. I've never met a firewall yet that would plug a server back into the wall.

How do most security breaches happen? *Everyone*

If you were expecting to see the word "hackers" then you've been watching too much television. Most security related issues are accidental issues. Blame does not always fall on a person, mind you. If the operating system crashes and is unrecoverable, it could have been caused by a hardware error, but it also could have been caused by something as simple as installing a bad bit of software. While finding the true cause will help you prevent it from happening again, tossing blame around at anything you think caused it will not. If it's not an object that you can control, then blaming it does your no good. If a hurricane wipes out your data center, do you blame the hurricane? No, but you can place blame on the location being in Florida, and to prevent this, you could move the base of operations to a new location. The problem with this, you may open it up to tornadoes or worse.

The most common security breech you hear about is cause by automated attacks, such as *Worms, Viruses,* and *Trojans*. Because these are automated attacks, they can take out a lot of machines very quickly, but they also can not adapt beyond what they were programmed to do, and can not compensate to defenses they were unprepared for. For example, if a worm was to try to log in as administrator, and you renamed the account, it would fail. However, any attacker worth his salt would be able to see the administrator account was renamed with ease. And a worm can be programed to check to see who the administrators of a machine are. When implementing security, you should keep in mind, is the effort worth the gain? And will this effort make it harder to use the system, thereby raising the operational cost?

Also keep in mind, damage is not always caused by external sources. Internally caused damage is a major issue. In terms of companies, most security breaches are internal, caused by users. This applies to things you may not take into consideration. It could be that the VP of the marketing department meant to send an email to a friend, talking about how cute the new female hire is, only to accidentally send it to the whole office she works in. It could be that the newest security patch for X program accidentally changing some system files, preventing the OS from booting. Even hardware failures, from power outages, to blown CPUs, fall under security due to the disruptions of operations.

Quick Recap *Everyone*

This section is just a quick recap of what you should now know. Security is all about the *Security Triad*, or *CIA: Confidentiality, Integrity, Availability.* When you are ready to implement a security program, you first have to decide what to protect. For home users this is simple, but for businesses, this can be a very complex process. We learned that different bits of information have different security needs, and that most security breeches are not caused by malicious attackers. We also learned that most networks have a single line of defense, and that there are ways to completely bypass this line.

Limiting the Damage *Everyone*

As I have mentioned before, the goal of IT security isn't to stop damages from happening so much as limiting the impact of said damage. In the section following this brief introduction, you will learn a formula used for *Risk Assessment*, that goes like this: *SLE x ARO = ALE*. Without going into detail, this formula states that the *Annualized Loss Expectancy,* or amount of money you will lose per year from an event, is based one two things, the *Single Loss Expectancy*, or how much damage, on average, you will lose if the event happened one time, and the *Annualized Rate of* Occurrence, or how often it happens. If a server crash caused your company an average of $1,000 every time it happened, then your SLE is 1,000. If it happened once a month, then your ARO is 12. So your ALE for this item is $12,000 per year. If you cut the SLE or the ARO in half, it would have the exact same effect, cutting your ALE in half. So in general you want to lower the cheaper of the ones to fix first if cost is an issue, or lower the easier item to fix.

Even for home users, while you may not need to know this formula, it shows you that there is two parts to any security issue: How often it happens or how easy it could happen, and how much damage it can do. In the implementation section of the paper, you find a lot of information on different ways to not only prevent damage, but to limit the damage done. This leads back to the section on the security of most networks, and how they are like egg shells. Once something gets in, there's no internal checks to limit the damage done, allowing an automated attack to cause a great deal of damage, and a controlled attack, such as a hacker, will have no trouble finding the good stuff.

Glossary of Terms

0-9

5 Nines – Used mostly when talking about system uptime, 5 Nines, means the system is up and working 99.999% of the time. This generally means total downtime is less then 5 minuets per year, for a 24 hour, 7 day a week, 365 day per year time span.

A

Access Control Lists (ACLs) – An Access Control List is a listing of who, or what, has access to an object, or is denied access to an object.

Annualized Loss Expectancy (ALE) – This is how much, per year, you are losing from this event. The ALE is generated by the SLE x ARO to generate your ALE. For example, with the SLE of $4,000, and an ARO of 12, your ALE is $48,000. If the ARO was .25, then the ALE would be $1,000.

Annualized Rate Of Occurrence (ARO) – This is how often per year an event will happen. If your server goes down once for 4 hours per month, then you have an ARO of 12. If it went down for 4 hours every 4 years, then your ARO for this even is .25.

Archive Attribute – A bit inside the header of a file, on most file systems. This bit tells the computer if the file has changed since the last time it was backed up, or Archived.

Attack Surface – The area that an attack can come through, or, the visible area of a target the attacker can try to compromise. It is the sum of all the possible ways in an attacker can use to get into the the system.

Availability – The amount of time that a service is up and running correctly.

C

CIA – In security, it's also known as The Security Triad. It stands for Confidentiality, Integrity, and Availability.

Client Firewalls – Client firewalls are firewalls that run on the local machine, and are designed to protect only the local machine. They are generally less robust then Network Firewalls.

Confidential (Information) - This is the data used by a limited number of internal users, and should not be known to the majority of workers. This is the class Human Resources (HR) data and Payroll Information falls under. Read access to this data is limited to a few users, and write access is generally restricted even more. If this becomes public internally, Operations and Internal Trusts are at stake, while if reviled externally, you lose your PTR, along with Operations and Internal Trusts. OS files dealing with security also fall into this area in most cases.

Confidentiality – Insuring that data is only observed by people who should be observing it. While this is a major goal in the IT field, having complete Confidentiality can be impossible to insure at times.

D

Demilitarized Zone (DMZ) Firewalls – In firewalls, the DMZ is the area between two firewalls. The DMZ space is normally where servers that need to be accessed from the outside, such as web servers and VPNs, are located.

Denial of Service (DoS) - An attack that prevents normal usage of a service. It could be caused by using up all of your companies bandwidth from the ISP, or purposely logging into an account repeatedly with the wrong password to lock out the account.

Differential Backup – A type of backup that backs up all files with the Archive Attribute set, and does not reset the Archive Attribute.

E

External Secrets – Information held in trust for another entity. Examples of External Secrets would be customer credit card information, or trade secrets from partner companies.

F

Full Backup – A form of backup that backs up every file, regardless of the Archive Attribute. After backup is performed, the Archive Attribute is reset on all files.

I

Internal Trusts – An Internal Trust is a Trust between a company and it's employees, between two or more employees, or even between family and friends. This can be compromised in many ways, from displaying emails that could be considered offensive to some members of the trust, or reveling benefits given to one member not given to another of the same standing: example, payroll information that is often negated on the part of the employee, and isn't purely based on performance and value.

M

MCP – Standing for Microsoft Certified Professional, this is often thought of as being a certification. Being a MCP merely means you passed one Microsoft exam, it does not even have to be a technical exam, most application exams, such as Excel exams, will grant you the MCP status.

MCSE – Standing for Microsoft Certified Systems Engineer, this is the highest level of Microsoft certification for server and network administrators in Microsoft NT, 2000, and 2003 environments. It consists of 7 exams: An exam on desktop OS's , three exams on Server and network (Called core exams), two electives, and one network design exam. Electives generally stem between SQL exams, Exchange exams, and Security exams, but can cover a wide range of topics. The MCSE is a lifetime certification, meaning it never expires, but each updated OS has a new MCSE exam, creating people who are certified in MCSE 2000, or MCSE NT4. As of July, 2007, there are over 750,000 MCSE certificates given world wide. Keep in mind that there are less than 750,000 people with MCSEs, as many can hold two or even three difference MCSE certificates.

MCSE:Security – A specialized form of the Microsoft Certified Systems Engineer, devoted to the security aspects of Windows. To achieve MCSE:Security in Windows 2000, the user is required to take eight exams, the extra exam being the users choice of a ISA Firewall exam, or the third party CompTIA Security+ exam. The choices of exams leading up to the normal MCSE level is restricted in this path, the user must take a Security exam (70-214 in the 2000 path) in place of one of his or her electives, and must take the 70-220 exam (Again, for the 2000 path) for the design requirement. As of July, 2007, there are about 18,000 MCSE:Securities world-wide.

R

Risk Assessment - Risk Assessment is the practice of identifying risks to your business or assets, assessing the potential damage done, and recording how often said risk will occur.

S

Security+ - This is the security exam by the vendor neutral organization known as CompTIA. It can be used as an elective for the MCSE:Security specialization, taking the place of a ISA firewall exam. As of July, 2007, there are about 30,000 Security+ certifications world-wide.

Single Loss Expectancy (SLE) – This is the value of damage the average event of this type will happen. For example, if the average downtime of a server would be, say, 4 hours, and this down time causes $1,000 in loss per hour, then your SLE for this event is $4,000.

SOX - The Sarbanes-Oxley Act of 2002, also known as Public Company Accounting Reform and Investor Protection Act. This is a United States federal law, it’s primary goal is to force companies to evaluate their Internal Controls for finical reporting, and that a third party (In general, an Auditing Firm consisting of CPAs) confirm that reported practices and implemented practices do indeed match. See: http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=107_cong_bills&docid=f:h3763enr.tst.pdf for more detailed information.

V

Virus – In computers, a Virus is a program that can “infect” other programs with copies of itself. On it's own, in the purest sense, a virus can not infect other machines with itself on it's own. It must be copied via the actions of users. While it has the most publicity, a virus is still the rarest form of malware: Worms, Viruses, and Trojans.

VPN - VPN stands for Virtual Private Network. It allows two separate networks to act as one while the data from each network is transmitted over a public network, such as the Internet.

W

Worms – A worm is a program that will transmit itself to other devices on it's own, spreading across a network with no user intervention. Worms are generally single files, meaning that many times computer can only be infected by a worm once. Some single-file worms, however, use randomly generated file names, allowing for more than one concurrent instance of the worm to be running at one time.

References used:

Microsoft Course #2050A
Designing A Secure Microsoft 2000 Network

Open Sources: Voices from the Open Source Revolution: Appendix A - The Tanenbaum-Torvalds Debate
http://www.oreilly.com/catalog/opensources/book/appa.html

IPTables Connection Tracking - FTP
http://www.sns.ias.edu/~jns/wp/2006/01/24/iptables-how-does-it-work/?p=20

Linux-PAM modules etc. page
http://www.kernel.org/pub/linux/libs/pam/modules.html

Backup types
http://www.backup4all.com/backup_types.php

Applying the Principle of Least Privilege to User Accounts on Windows XP
http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/luawinxp.mspx

Using a Least-Privileged User Account
http://www.microsoft.com/technet/security/secnews/articles/lpuseacc.mspx

The Services and Service Accounts Security Planning Guide
http://www.microsoft.com/technet/security/topics/serversecurity/serviceaccount/default.mspx

MakeMeAdmin -- temporary admin for your Limited User account
http://blogs.msdn.com/aaron_margosis/archive/2004/07/24/193721.aspx

MakeMeAdmin follow-up
http://blogs.msdn.com/aaron_margosis/archive/2005/03/11/394244.aspx

An Analysis of Simile
http://www.securityfocus.com/infocus/1671