



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials (Security 401)"
at <http://www.giac.org/registration/gsec>

Security in Practice- Reducing the Effort
GSEC Practical Assignment V1.4b Option 1

By Leon Pholi
April 2003

© SANS Institute 2003, Author retains full rights.

Contents

| | | |
|----------|--|-----------|
| 1 | Abstract | 3 |
| 2 | Introduction | 3 |
| 2.1 | Practical Security | 3 |
| 2.2 | The Effort Factor | 4 |
| 3 | Ten steps of network and information security | 5 |
| 3.1 | Policies and Education | 5 |
| 3.1.1 | Policy Creation | 6 |
| 3.1.2 | User Education | 8 |
| 3.2 | Perimeter Security | 8 |
| 3.2.1 | Physical Security | 9 |
| 3.2.2 | Router | 10 |
| 3.2.3 | Firewall | 11 |
| 3.3 | Password Control | 12 |
| 3.4 | Services Minimisation | 14 |
| 3.5 | Patch Management | 16 |
| 3.6 | Antivirus Solutions | 18 |
| 3.6.1 | Product Selection | 19 |
| 3.7 | Access Control | 20 |
| 3.7.1 | Permissions | 20 |
| 3.7.2 | Templates | 21 |
| 3.7.3 | Remote Access | 22 |
| 3.7.4 | Encryption | 22 |
| 3.8 | Secure Communications and Detection | 23 |
| 3.8.1 | Communications | 23 |
| 3.8.2 | Intrusion Detection | 25 |
| 3.9 | Unauthorised Network Equipment | 26 |
| 3.9.1 | Policy and Detection | 26 |
| 3.9.2 | Modems | 27 |
| 3.9.3 | WLAN | 27 |
| 3.9.4 | Other Equipment | 28 |
| 3.10 | Maintaining Security | 28 |
| 3.10.1 | Security Training | 28 |
| 3.10.2 | Backups | 29 |
| 3.10.3 | Auditing and Log Analysis | 30 |
| 4 | Conclusion | 32 |
| 5 | References | 32 |

1 Abstract

Information security is known to be at least important, if not critical, to most business and personal needs. This paper covers the ten most vital steps in attempting to achieve a good base level of security, which can then be built upon. The focus of these is on reducing the effort in order to ensure they are completed to at least a minimum degree.

The intended target audience is Network/Systems/Security administrators who need a reference guide on the fundamental steps in securing a network, why each step is important, and how to reduce the effort whilst doing it.

2 Introduction

There are many reasons why companies and individuals do not carry out even the most fundamental aspects of securing their networks, but one of the most common would have to be "It's too hard", followed closely by "I don't have time" and "Although there is plenty of information, I just wouldn't know where to begin". This paper attempts to address these issues in a couple of ways. Firstly it includes the 10 steps that have the biggest impact in securing a network, remembering the Information Security 80/20 rule, which is 80% of exploit risks can be effectively reduced using 20% of the recommended security procedures. Secondly, it addresses the need to make these tasks easier, through both technological and procedural means. This includes additional steps that, although they may only make an incremental difference to the level of security, are still preferable than not doing anything at all due to giving in to one of the excuses mentioned. Thirdly, each step is given an Effort Factor, which is a rating between 0 and 1, of what level of effort the author considers is required to complete the step to a minimum level, with 1 being the maximum effort. The effort factor can then be used as a guide when creating project timelines for the implementation of the steps in this paper.

2.1 Practical Security

The SANS institute has published "The 7 Top Management Errors that Lead to Computer Security Vulnerabilities" (see <http://www.sans.org/resources/errors.php>). Number one on this list is to "Assign untrained people to maintain security and provide neither the training nor the time to make it possible to do the job" [1]. While increasingly the commonly held perception is that network and information security is important for any company, the reality is sufficient time and resources are rarely given to this critical area. Usually there is no single qualified person who is given the resources, authority, and ultimate responsibility for information security. While this would be ideal, and is beginning to happen, overall the norm is for these roles to be divided amongst staff of varying levels within the company, such as the network administrators and the CIO (Chief

Information Officer). The result is these duties are being expected to be done along side many other tasks.

It is little wonder that the individuals who end up having to ensure integrity of resources they are in charge of, end up letting major security system design and implementation floors regularly slip through. They simply do not have the time or training to do this additional role. However since it still needs to be done, the lack of time and training should to be made up for by having quick access to supporting resources and finding ways of minimising the required work, or else it is likely little or nothing for security will be done, in order to accomplish perceived priorities.

The information security 80/20 rule is an implementation of the broader concept discovered by Italian economist Vilfredo Pareto, known as the Pareto Principle (explained at <http://www.8020info.com/principle.html>). When applied to information security, it implies that the majority of the risk that a company is exposed to can be substantially reduced by implementing the few most important procedures [2]. By focusing the security time and resources on these areas, maximum impact is achieved in a short amount of time, while the effort involved is simultaneously reduced as a positive by-product of this approach.

The exact number and types of procedures that are required to follow this principle are a matter of conjecture, and to a degree a matter of personal opinion. For example security services company Symantec promote using just three specific security controls in order to achieve the 80/20 goal; remove unneeded services, keep patches current, and enforce strong passwords (see http://securityresponse.symantec.com/avcenter/security/Content/security_articles/fundamentals.of.info.security.html) [3]. This may or may not be proven to be all that is required, however to ensure the majority of networks will be secure to a level equal or greater than the Pareto Principle, there are at least ten areas that should be looked at with high priority, all of which are generally accepted as important steps in information security by security experts.

2.2 The Effort Factor

As part of each step, a very short summary of the points to which attention should be focussed is included, under the heading "Minimum Requirements for Step". For each of these a rating is given under the heading "Effort Factor". This rating is a guide to indicate the potential time, knowledge, planning and implementation difficulty that is required. On the scale, 0 indicates no effort is required, while 1 indicates it is a very difficult task to achieve.

Note that the effort factor is purely based on the personal experience of the author, and should not be taken as a definitive level, but rather a guide for use when the topics are unfamiliar. The actual effort involved will have many dependencies, including technical environment, network size, and experience

level of the security practitioner. That said, the Effort Factor is intended to add value and clarity to the reference, in order to simplify and bring perspective to what at times can be an oversupply of information from the Internet community.

3 Ten steps of network and information security

The ten steps which will vastly improve the security of information on a network are as follows.

3.1 Policies and Education

Most companies have many different policies for all sorts of areas, and as such it is easy to view policies as 'bureaucratic red tape', just making it difficult for implementers to get the job done. There are also policies that can be applied to information security, which can vary in number, size and scope depending on company size. However regardless of the size of the company or any other factor, it is absolutely essential that they all have a security policy.

A security policy provides the basic framework which all other security practices revolve around. This is the most vital step in securing a network, as without a good & relevant security policy, even if there has been adequate security procedures initially put into place, they will quickly fail if they are not based on rules and guidelines to define what they are meant to achieve. It is also a good idea to create an initial security policy before applying other security recommendations to ensure these are in line with the policy, or else audit the current environment after major security policy changes to ensure compliance.

RFC 2196 Site Security Handbook, <http://ietf.org/rfc/rfc2196.txt?number=2196> makes the following statement about security policies:

A security policy is a formal statement of the rules by which people who are given access to an organization's technology and information assets must abide.

The main purpose of a security policy is to inform users, staff and managers of their obligatory requirements for protecting technology and information assets. The policy should specify the mechanisms through which these requirements can be met. Another purpose is to provide a baseline from which to acquire, configure and audit computer systems and networks for compliance with the policy. Therefore an attempt to use a set of security tools in the absence at least an implied security policy is meaningless. [4]

3.1.1 Policy Creation

The first step in writing a security policy is to identify the risks. This simply means taking the time to access what exactly it is that needs protecting, and defining where the risks are coming from. For example, the desktop machines for the users may contain company data. If the users are allowed external access to the Internet and email, they could introduce a virus onto the desktop machine causing data loss and a virus outbreak. They could also bring an infected disk from home, causing the same outcome. The security policy needs to cover all identified risks, specify how to minimise them, and how to deal with possible results, including roles & responsibilities for those involved.

It is important to note that the security policy should cover any laws that are applicable to the information security of the company. There are some industries that have strict government controls over the handling of information, which should be considered throughout the process. Another consideration is any other guidelines that are considered appropriate to follow. ISO 17799 is a comprehensive international standard on information security that is proving increasingly popular; the standard is downloadable after buying online [5]. For more information see <http://www.iso17799software.com/>. Compliance to this standard is not a simple task, which is one of the strengths of the standard, ensuring that those who do have clearly proven to be secure to a specific level. For the sake of keeping it simple, it is best to implement initial policy & procedures based on good resources, after which fine tune using quality standards such as these.

Next, the security policy needs to include exceptions to the standard situations. This is necessary because if a scenario is not covered by the policy, it would likely mean the simplest solution will be found rather than the most secure. An example is if there is a business critical file that needs transferring to another company that the security lockdowns are preventing, ways to circumvent the security may be attempted. Rather than have this happen, the policy should state what to do in such situations, where possible specific to each area in the policy. It should also have a general rule that may include making sure ad-hoc changes are approved, recorded and later analysed for finding ways to remove the need for this to be repeated.

Another part to include in each section of the security policy is clearly defined roles and responsibilities. The users, administrators, management, contractors, and anyone else must be aware what is required of them in relation to information security. Some users, for example, are simply unaware that it is unsafe to access their personal email account and run an executable file sent to them from some unknown person. By describing what is expected, as well as what to do, such as whom to notify if the above example happened and what needs to be done by the person notified, reduces the chances of human error or oversight.

Finally it is quite important that the security policy is written in conjunction with management. Without management buy-in, there will be little impact made through the security policy. Appropriate management representatives should be adding input as to their expectations during the process, so the result is the right combination of applying maximum appropriate security considering business priorities and processes. There is also a need to involve management in creating the policy because at the end of the day they will have to sign off approval for implementation. It must be someone from outside the IT department that signs off on the security policy, as its impact is company wide.

The security policy needs to be broad in terms of its intended audience. It should be clear for all levels in the company to understand, not technology specific, but cover everything relevant to information security. Because of this, the security policy should not require continual updating. From the broad nature of the security policy, security procedures can be developed. These can detail the specifics of how the policies are implemented, down to day to day tasks and technology solutions used to achieve the security policy goals. The procedures that are developed from the policies can be regularly updated as technology and operations change, and as such only need signing off from within the IT department.

The effort involved in creating a security policy may seem to be quite a lot, as there are many considerations, but there are ways to reduce the effort required. Most companies that do not have a formal security policy do have procedures that are part of the daily operations. If these procedures are not written, there are usually at least informal implied procedures that exist. Once identified, these procedures can be used in reverse to formulate an appropriate security policy. By discovering what the purpose of security related procedures are, when they are done and who has responsibility for them, an inferred policy statement can be made.

For example if there are methods currently employed for virus protection such as antivirus software on the desktops which are routinely updated and recorded, a policy statement indicating this must be done can quickly be deduced. It would still be necessary when complete to go through the entire policy to make sure it is thorough enough, especially considering the procedures used to at least in part create it may not be all that's needed, or may not cover the complete requirements for each issue. This is where the aforementioned steps such as identifying risks becomes necessary to prevent policy holes.

Another way to reduce the effort in creating the policy is to use policy templates and published examples. SANS have available issue specific security policy examples and templates as guides (see <http://www.sans.org/resources/policies/>) [6], as well as tips and information sharing as one of the topics in its reading room (see the Security Policy topic,

<http://www.sans.org/rr/policy>) [7]. Others such as Cisco have also published examples (<http://www.cisco.com/warp/public/126/secpol.html>) [8].

3.1.2 User Education

Additionally to security policies & procedures, an essential part of information security before getting to technical aspects is education. A recent survey by the Computing Technology Industry Association (CompTIA) located at http://www.comptia.org/pressroom/get_news_item.asp?id=207, indicated that “where agencies and companies have looked primarily to technology for network safety, in over 63 percent of identified security breaches, human error looks to be a major, underlying factor” [9]. Following on from this the survey also found “80 percent of respondents saying that a lack of IT security knowledge, training or failure to follow security procedures were the root causes of human error” [9]

Education for all users is absolutely critical to reduce the flaws in security that can result from human errors and omissions, no matter how good the technology based solutions used are. While it would be great to send the entire company on long intense courses, there are easier ways to raise user awareness. As stated earlier, part of the role for the security policy is to do just that, inform and raise awareness. Therefore once the policy has been approved, it should be added to be read as part of the employee induction procedures. Furthermore, the security policy needs to be easily available and accessible for staff to be able to refer to at any time. For example it can be published on the Intranet, as long as all staff then can access it.

Finally it would be ideal and very beneficial for management to organise short occasional security awareness meetings, such as bi-annually, where staff in small groups can be reminded of best practices, be informed of any changes, and ask questions in an open and relaxed forum. This would not only vastly raise awareness but also make users and management more comfortable with the fact they have security responsibility in their personal role to some extent.

| Minimum Requirements for Step | Effort Factor |
|---|---------------|
| Security policy creation using existing procedures | 0.7 |
| Educate users using security policy & make policy readily available | 0.1 |

3.2 Perimeter Security

Perimeter security in this context covers both the physical and data perimeter boundaries of a computer network. In order to protect the perimeter, a number of different areas need to be looked at.

3.2.1 Physical Security

The first aspect to consider for perimeter security is physical security. This has often been the easiest way for data to fall into the wrong hands, and is regularly overlooked. Even if you follow all other security practices to the letter, neglecting physical security can inevitably lead to a compromise of data. Physical security should be a part of the security policy, as it is a fundamental element of information security.

If unwanted individuals can gain access to the physical information on the network, many of the other security mechanisms will be useless. The incorrect underlying security principles for some networks can be the internal network is considered 'safe', while only on public networks such as the DMZ and the internet is the data in any real danger. However even if this were true, with physical access that assumption is gone and potential danger is exponentially larger. Physical access means the possibility of being able to bypassing security mechanisms is almost definite. For example a server may be password protected and logged off, however someone with an operating system on a disk can boot to that and access all data on the machine, bypassing existing access controls. If they cannot do that they could always remove the hard disk and take it with them, meaning access would not be a matter of if, but when.

The first step in ensuring physical security is defining restricted areas and non restricted areas. Having done the identifying risks part of security policy creation will make this step slightly easier. Having restricted areas enables definition of what is to be protected simpler. An example of this would be a server room or data centre. Assess what else needs physical protection other than these restricted areas, such as backup tapes, network hubs & cabling etc. Next, consider appropriate method of securing these areas. Options to consider are many, including security guards, various kinds of locks, biometrics, inspections, cameras, sensors, and many others.

While this may seem a lot to have to decide, at a minimum all important equipment should be added to server rooms and cabling cabinets which should be locked with keys given only to those that need them, and employees should have ID passes including visitors for easy identification. Monitoring of the designated security areas should be carried out also, if not by security cameras then an audit trail of who has access to these and when needs to be undertaken, either through technical or procedural means. Secure off-site storage of backup tapes needs to be considered such as in a fire/electromagnetic proof safe.

Other options include preventing boot to floppy in the BIOS, and locks for the PC case, and cable locks which are particularly useful for laptops since they are much more easily stolen. There are almost an endless number of ways for preventing physical access to assets, however the more layers that can be added to the physical security the better. The effort can be reduced by simply starting with the most critical such as those listed in the previous paragraph and continuing to add layers when time and budgets avail.

3.2.2 Router

Next is data perimeter security. The data network perimeter is usually made up of combinations of firewalls and routers. When a packet arrives into the network usually the first device it needs to travel through is a router at the perimeter. While the primary function of a router is not security, and it may seem OK to leave this to other devices such as firewalls, defence in depth is the goal, that is the more layers of defence there are the more protected the information is.

Both a router and a firewall usually have the ability to perform ingress and egress filtering on the packets that traverse them. Ingress filtering refers to packets coming into the network, and egress filtering refers to packets leaving the network. Of these ingress, or what is allowed into the network, is usually the most critical.

Routers have the ability to do this through access control lists (ACL). There are two common types of ACL, standard and extended. It is through using an extended ACL that allows ingress or egress filtering based on source or destination, as well as port or other header information. It can also be used to configure services running on the router. Using these features it is possible to set quite advanced criteria on what traffic is allowed to pass through the router and how the router will behave. For example Oracle recommends (at http://www.oracle.com/ecostructure/blueprint_rec/network_security_analysis_and_recommendations.htm) stopping the finger, bootp and other unneeded services as well as preventing IP redirects [10]. Others such as a checklist published on the SANS site (see <http://www.sans.org/score/checklists/CiscoChecklist.doc>) suggest in greater detail many services & ports that should be blocked like SNMP [11]. One of the ten steps in this guide is indeed stopping unnecessary services, and the external router is a good place to start.

However blocking many critical ports on an existing network may have unknown consequences. Therefore to reduce the effort on this one, start by only allowing the right IP ranges on each interface to prevent IP spoofing, block those basic services that are obviously and clearly not needed (including such things as preventing IP redirects etc), and then clamp down further after the firewall rules have been defined, by which stage there should be a clearer understanding of network services needed to be running.

3.2.3 Firewall

The next network device an arriving packet will usually go through is a firewall. A firewall has much more powerful and granular tool for this as security is its primary role. One thing must be said to companies currently without a firewall- get one! It is that simple. A firewall is the gateway to the internal network infrastructure, without one there is no gate.

That said, there is one situation equally as bad as having no firewall, and that is having a poorly or non configured firewall. Continuing the analogy there is little advantage having a large, thick steal gate if it is left open. Therefore even more importantly than a router there needs to be a clear set of rules defined on the firewall for what is allowed into and out of the network. In addition, there will be different rules needed to be set for the public servers in the DMZ to the private internal servers, although this can be achieved through using multiple tiers of firewalls as well. Also the right kind of firewall needs to be purchased depending on the role. For example the firewall acquired may be packet filtering, stateful so that decisions made are related to the connection and not just the individual packet, or application which allows advanced logging and authentication options amongst others. Questions should also be asked about the potential number of connections needed, and load balancing & redundancy requirements.

Although there are general guidelines available for configuring firewalls, such as one published by the National Institute of Standards and Technology in the US (<http://csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf>) [12], mostly the exact configuration needed will vary from one environment to another. Therefore to reduce the effort while preventing too many unexpected results, start by only doing ingress filtering. Egress filtering is a valuable step to be done so this should still be done afterwards. Then assess what services need to be connected to in the DMZ, allow only these, and allow only returning traffic into the private internal network, i.e. traffic originated from the internal network.

Finally and most importantly of all set a default deny policy, where if a packet does not meet this criteria it is dropped and logged. The logs will then show any connections to services that are failing, and if necessary these can be explicitly allowed to the specific device. Regular analysis of logs for both routers and firewalls even post set up will ensure configurations are correct and suspicious activity dealt with.

A basic set up such as this should not be terribly involving, although appropriate personnel would have to be notified before proceeding, but would make a monumental difference to the level of protection of the information assets.

| Minimum Requirements for Step | Effort Factor |
|---|---------------|
| Physical locking away, identification, basic monitoring etc | 0.6 |
| Router basic configuration & services lockdown | 0.3 |
| Firewall ingress filtering & default deny | 0.4 |

3.3 Password Control

Password protection is probably the most commonly used information protection mechanism, and as such it is easy to over simplify its implementation. The way passwords are chosen and changed should be scrutinised, as poorly chosen passwords that have long life cycles are little better than no passwords at all. Simple passwords and even simple passwords with extra characters can usually be 'cracked' in seconds by programs downloadable from the Internet.

Symantec makes the following recommendations on password selection and change (see http://securityresponse.symantec.com/avcenter/security/Content/security_articles/fundamentals.of.info.security.html) [3]:

1. Minimum password length should be 8. Administrator passwords should be fifteen characters or greater.
2. Use a password max age of 60 or 90 days. Any less may cause users to write down passwords.
3. Use a password min age of not less that 14 days. Do not want users to cycle back to their "favorite" password.
4. Keep password history at least 10.
5. Use a password filter so that users are forced to use a combination of alphanumeric and non-alphanumeric characters.
6. Audit for empty and weak passwords using a password strength analysis tool. Tool should use password dictionaries from the world of sports, Star Wars, Star Trek, Disney, J.R.R. Tolkien, Monty Python, etc.
7. Remove all default accounts from applications and devices.
8. Rename well known account names such as administrator, sys, system, if possible.
9. Remove inactive accounts.

Other good points are ensure passwords are stored with non-reversible encryption such as shadow passwords in Unix, use at least one letter (preferably both upper & lower case), one number and one symbol in each password, and use one time passwords where possible. Finally enforce the settings chosen and audit accounts for logging and analysis purposes.

Password management is quite likely the most critical aspect of user security education. It would be ideal if on hearing the need for strong frequent

passwords, they would be done without prompting. Human nature dictates, however, that the easiest path will be chosen. Therefore a two forked approach is required, education and enforcement. Education is used to raise awareness and inform staff on security concerns, such as not leaving passwords written on a Post-It note stuck on the monitor, or sharing their password with other staff. Enforcement is acknowledging human behaviour and ensuring passwords are controlled by rules.

Password policy settings for enforcing strong passwords can be done through most operating systems with the aid of password filters. Examples of filters are passwd+ for Unix, passfilt.dll for Windows NT 4, and the Local Security Policy settings in Windows 2000\XP. Password enforcement on a large network may seem a daunting task, but thankfully there are tools to reduce the effort required. Rather than changing policy settings such as minimum password length on each individual machine, use tools which centralise the management of passwords.

An example of powerful password management functionality is using the capabilities of Active Directory in a Microsoft Windows 2000/2003 domain. Using the features of Active Directory Group Policy makes it a relatively simple task to set a password policy on all the machines in the domain, and make changes any time. Microsoft offer plenty of support for how to configure group policy settings, best practice methods, and impact such as for Windows 2000 password policy settings. "Securing the Domain Infrastructure" at <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/pr odtech/Windows/SecWin2k/05secdom.asp> contains examples of some of these settings [13]. Other operating system and third party options also exist for security management centralisation, all with the purpose of making password control easier, some of which can cater for applications other than the OS as well.

Apart from enforcing passwords, the passwords should be audited to ensure compliance with the password policy decided upon. Again, this could be a difficult task in all but the smallest environments. However there are many tools that exist which make this process much more manageable. Microsoft offer the Microsoft Baseline Security Analyser (MBSA) as a free download at <http://www.microsoft.com/technet/treeview/default.asp?url=/TechNet/Security/tools/tools/MBSAHome.ASP>, which offers a simple interface and is easy to use [14]. Another highly recommended way of auditing passwords is to use a password cracking program on each system, after gaining written permission from relevant personnel. These are a very good way of finding out the strength of passwords actually in place. There are many open source and corporate products which allow thorough password auditing.

Finally accounts on all systems should have logging turned on for auditing purposes. These logs may be very useful down the track for analysing inconsistencies that might indicate suspicious circumstances, such as repeated account log on failures and using a non standard account for

accessing information. As with password policies these can usually be configured with ease even in the Enterprise using options such as Microsoft's Windows 2000 Group Policy. For example see the Windows 2000 Server Baseline Policy section of <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/pr odtech/Windows/SecWin2k/06basewn.asp> [15].

| Minimum Requirements for Step | Effort Factor |
|---|---------------|
| Operating System password enforcement and auditing using tools for centralisation | 0.2 |

3.4 Services Minimisation

In order for a machine to be compromised, an attacker must take advantage of a service running on that machine. When the attacker discovers and has access to a service with a known vulnerability, the compromise is already well under way. There is a well known saying along the lines that the only secure computer is disconnected, powered off, and buried deep underground in concrete. While this is obviously not possible in the real world, the closest that can be achieved to this is by only running the absolute bare minimum services.

If you don't need it, turn it off. For example the SNMP vulnerabilities that have surfaced would have been a far smaller threat if devices that didn't need this protocol had the associated service switched off. An administrator can sleep far easier at night knowing they are safe from vulnerabilities no matter what the patch level is since effected services are not even available for compromise. A positive by-product of minimising services is that it also makes the next step easier, there is less to have to patch if only minimal services are running.

Symantec offer the following advice on service reduction (see <http://securityresponse.symantec.com/avcenter/security/Content/security.articles/fundamentals.of.info.security.html>) [3]:

1. Define the role of the information system. Avoid using systems in more than one role, e.g., a web server should not also be the ftp server. A single role system makes it easier to define which services should be running and which services should not.
2. Determine which services are needed on the system (legitimate business need) and remove all others. If it is a public facing web server then obviously some web server components should be running, but the vast majority of other services installed by default should not. To discover services that handle network connections use netstat to view open ports and then systematically eliminate all services which are not legitimately needed for the system's purpose.

3. Determine which features within the service should be enabled. In many instances it is important to disable features within a service when there is no legitimate need such as ISAPI extension mapping in IIS. Also there are many sample script files that are installed by default by enterprise applications in many instances there samples do not undergo the same level of quality assurance and can be used to compromise the service. Remove sample scripts.
4. Public facing systems such as web servers, DNS, email servers, etc., should have priority in this first step. Internal servers should then have next priority when defining role and removing unneeded services.

Another point to add to these is service swapping. Although a service may be needed, there may well be another more secure method of doing the same thing. A good example of this is the telnet service, used most commonly for command line connectivity to Unix based machines. This service is inherently very insecure due to all traffic using it being sent as clear text. In most instances it is quite simple to change to a more secure alternative, such as SSH, which encrypts all transmitted traffic. The telnet service can then be disabled and the more secure alternative used in its place.

So how can the effort to do this be reduced? Stopping services on many machines could be very difficult on a large scale, particularly as this could have impact on the functionality of critical applications. There is also the tendency to leave the services running when administrators are unsure of their purpose. This is where defining the role of the operating system (point 1) helps. By only having one role per server, it is far easier rather than worrying about the cumulative effects of changes.

It is also a good idea to create test servers configured the same way as production, then minimise the services including the ones still unsure about even after research, until the right balance between security and functionality is achieved. Test the changes for each setup before them rolling out in the order noted in point four, reducing the size of the task into more manageable goals. Once a particular setup is determined, document it as a Standard Operating Environment (SOE) so that it can be rolled out to machines with the same configuration.

There are tools for some applications which can aide in making determining which features within the service should be enabled (point 3). These tools often make the task a lot easier. An example of this is Microsoft's IIS lockdown utility, which is a free download located at <http://www.microsoft.com/windows2000/downloads/recommended/iislockdown/default.asp> [16]. Using a simple wizard it is possible to customise the service as required and remove the unneeded components.

Once the services on the machines have been minimised, they should also be routinely audited to check all is as expected. Again doing so can be made simpler using the various tools available. Microsoft's MBSA [14] mentioned

earlier is one such tool, as well as others such as the GFI LanGuard Network Security Scanner (<http://www.gfi.com/lannetscan/index.htm>) [17], the Nessus project from the open source community (information on features found at <http://www.nessus.org/features.html>) [18], and many more.

| Minimum Requirements for Step | Effort Factor |
|---|---------------|
| Service reduction creating SOEs and rolling out in stages | 0.4 |

3.5 Patch Management

When an attacker attempts to compromise a machine or network, they are usually looking for one of two things, either an open door such as an unprotected file share, or a known vulnerability. A vulnerability can be any number of flaws in an application and can usually be exploited in multiple ways. Examples are escalation of privileges meaning the attacker gains more rights on the system than they should, or triggering errors that mean the application crashes, thereby causing denial of service.

The protection against vulnerabilities is simple, patching the application. Patching applications closes holes in them reducing the possible options for compromise. Unfortunately this is a never ending battle; as soon as one vulnerability is closed another is often discovered. A system that is fully up to date with patches and service packs one day may be exposed to a major flaw the next, making it difficult for administrators to keep up.

Apart from problems with the workload, there are often other issues with patches, that is although they may close security holes, they can often cause the application to error and even fail, making patch deployment particularly on production servers risky. These two factors combined can make it tempting for administrators to avoid patching systems altogether. This line of reasoning is very dangerous, as unpatched systems on the network not only risk compromise for themselves, but they can then be used as staging points to attack other systems on the same network or other networks.

The risk for administrators of damage and loss of control related to a compromise will normally always outweigh the potential risk of patch installation, although each patch should be assessed to check this is indeed the case. Consider this- what is the most valuable asset to an organisation? Usually the answer is the same; it's the information and intellectual property that the company has that sets it apart from the competition. Anything that is a threat to the core of the business needs to be given the highest priority, which means timely and consistent installation of relevant patches.

Symantec gives the following advice on patch management (see http://securityresponse.symantec.com/avcenter/security/Content/security_articles/fundamentals.of.info.security.html) [3]:

1. Identify available patches from vendor sites. This should be automated if at all possible or a patch tool should be used.
2. Identify systems that are not running the latest patch. A patch tool can identify systems without current patch levels.
3. Download and test patches on test systems. Vendors run patches through QA but the release cycle is often much shorter than the long regression testing employed during a regular release cycle.
4. Deploy patches to systems. Public and internal servers receive first priority.
5. Monitor systems. Determine if the service behavior has changed.

Point 3 is critical and addresses the concerns about the impact of patch deployments. They should always be tested first on non production systems set up identically to their live counterparts, to ensure the systems function as expected. When given enough time and testing as previously defined in security policies and procedures, they can then be rolled out to the production servers in order of priority, and then checked again for correct functionality as described in points 4 & 5.

Reducing the effort involved in following the patch management best practices can be achieved in a number of ways. Firstly there needs to be ways of identifying patches when they are available. Patching systems is always a race between the good guys deploying the patches to systems before the bad guys can discover a system is vulnerable and take advantage of it. Depending on the nature of the exploit this could be as short as a matter of hours or less, therefore timely notification is important.

Notification can be achieved through subscription services to the product vendors, newsletters from reputable security web sites, and appropriate newsgroups can also be used. Make sure one way or the other all operating systems and applications are covered so an administrator will always be informed of an update. Because this is a manual approach always consider an automated method of notification wherever possible. Microsoft Windows Update is a free online service that allows this, located at <http://windowsupdate.microsoft.com> [19]. Using this service an Automatic Update utility can also be downloaded to make patching easier. As long as the system has connectivity the internet, relevant patches will automatically be flagged, and can even be configured to be automatically downloaded and installed.

In order to identify systems that need patching (point 2), there are a number of tools available to make this easier. As well as Windows Update, Microsoft have the Microsoft Baseline Security Analyser (MBSA) [14] mentioned previously which allows remote checking of the patch levels of Windows systems, as well as other security holes. The MBSA is an improved version of

the Hfnetchk utility that Microsoft had released earlier (see <http://www.microsoft.com/technet/security/tools/tools/hfnetchk.asp> [20], as Hfnetchk is a command line version. Other tools exist which are able to indicate the patch levels for various operating systems also. The Center for Information Security (see <http://www.cisecurity.org>) [21] offer security tools which not only check patch levels for various OS's, but check compliance with security lockdown benchmarks of various levels, as created by the cooperative effort from some of the leading information security authorities. It is recommended to spend time going through these recommendations and considering applying them after thorough local testing. Tools such as these and others should be used to routinely audit all machines on the network to check they are all patched to the appropriate level.

Once the systems to be patched have been identified, and tested on non production machines, they need to be deployed beginning with the internet exposed and critical internal servers. Microsoft Windows Update service [19] again can aide in this, as can others such as the Red Hat Linux Up2date subscription based service (see <http://www.redhat.com/docs/manuals/RHNetwork/ref-guide/up2date.html>) [22] which works somewhat similar to the Microsoft version. Increasingly there are options such as these within quality non OS applications to allow checking the patch level and updating the software far easier, and should be a consideration when purchasing any new application.

To reduce the effort further, and gain more control of patch management, Microsoft offer another free application called Software Update Services (SUS- see <http://www.microsoft.com/windows2000/windowsupdate/sus/>) [23]. SUS is basically an internal version of Windows Update, giving administrator's control over which Microsoft patches are approved for installation and allowing automated patch rollout. Microsoft Systems Management Server (SMS- see <http://www.microsoft.com/smsserver/default.asp>) [24] is a product that offers a more advanced version of MS SUS features amongst other advantages, such as reporting options.

Patching many and varied systems and applications can seem like a daunting task, but with the right tools and management techniques the effort involved can be substantially reduced in this fundamental security practice.

| Minimum Requirements for Step | Effort Factor |
|---|---------------|
| Patch management using notification services and automated rollouts | 0.3 |

3.6 Antivirus Solutions

Viruses, Worms, Trojans etc have become regular vocabulary in today's inter-connected world. The media will occasionally grab one particular threat and

suddenly public awareness of this threat skyrockets, and then gradually dissipates. The reality is the risk of infection from malicious code is constant, even old viruses can cause havoc on unprotected systems. New viruses most often are a modified version of a previous one, although sometimes a virus is released that takes advantage of new techniques, designed to increase the rate of infection and avoid detection.

The damage caused by these pests is varied, and can even go unnoticed. Many are harmless and simply self propagate as proof of concept code. Some are hoaxes and don't actually exist. Some however do cause major damage to the systems they infect, resulting in considerable time and resources to clean up after infection has been detected. One way viruses spread is through human interaction, so the Security Policy and user education should cover what to look for, and how to deal with this threat. Another way is through automatic detection and exploitation of a known vulnerability, which is yet another reason why services should be minimised and patches up to date. This includes investigating software application version improvements for virus prevention, for example more recent versions of MS Outlook automatically block attachments with dangerous extensions. Additionally some such as worms tend to be self propagating, and even mutating code, making them very difficult to detect and stop using conventional methods.

3.6.1 Product Selection

Fortunately there exists an entire market of products whose sole role is to make this task easier. There is a diverse range of solutions available to combat these threats, so the biggest hurdle is usually which one to select and how to manage the products successfully. Firstly, decide where to target viruses effectively. The answer to this is usually threefold, one is at the perimeter such as the firewall, the second is critical infrastructure applications which carry risk, such as email servers, and third on the machines at risk themselves, that is all servers and desktops including laptops and machines used for work at home.

Next, find a product or products that meet these needs, as well as ensuring the quality of the software is high beyond just marketing hype. One thing to consider is despite promising to detect 'x' number of viruses, does it include 100% of the "Wild List"? The Wild List, located at <http://www.wildlist.org/>, covers the viruses currently circulating and infecting computers, and as such is more relevant than pure virus numbers [25].

As explained at <http://frontpage.kmoraine.com/antivirus.asp>, a site that offers some advice on Antivirus solutions:

In 1992 [ICSA Labs](#) (then the NCSA and now a division of TruSecure) established a certification process that provides a consistent and accurate means of comparing antivirus detection rates. This process favors anti-virus software that can detect viruses in the wild, requiring 100% detection

of viruses on the WildList. Other viruses are considered to be less important.

The ICSA Labs testing criteria are well designed and the testing process is thorough and performed by professional virus researchers. Look for the ICSA Labs Certified logo on antivirus software products and check the latest test results at <http://www.icsalabs.com>. [26]

The product(s) chosen should have the ability and be configured for both on-access scans as well as routine full scans. Relying on users to remember to scan their systems is rarely a good idea, so make sure it can be configured for routine full system scans such as weekly, and also on access for checking files are clean automatically immediately before they are accessed.

Finally it is vitally important that any antivirus solution chosen is easy to update and manage. Antivirus software contains a list of fingerprints for detecting the various kinds of viruses, including boot sector, programs, macro, and others, and as new viruses are released this quickly becomes outdated. Therefore any product used needs to have easy ways to ensure the list for detecting the latest viruses is regularly up to date. Having out of date virus software can be little better protection than none at all, so these should have scheduled automatic updating which is routinely audited to check compliance.

Enterprise level antivirus management software is used to make the updating and auditing process easier by centralising control. Having access to fast timely reports means there is little excuse for all systems not to have an antivirus solution installed and up to date. Both McAfee (<http://www.mcafee2b.com/>) [27] and Symantec (see <http://enterprisesecurity.symantec.com/>) [28], as well as others, offer enterprise level management solutions to make it relatively simple to achieve this goal.

Ensuring all systems are protected from malicious code is a relatively simple task, as long as the time is taken to make sure the methods are thorough, up to date and well managed.

| Minimum Requirements for Step | Effort Factor |
|---|---------------|
| Antivirus solutions installed, automatic updates and centralised management | 0.2 |

3.7 Access Control

3.7.1 Permissions

Once access to a system has been achieved, whether legitimate or not, what is available to that user depends on the account used to connect and what that account has been configured to be allowed to do. Most modern operating systems allow quite granular control of rights through share and file permissions. If the operating system does not allow file level permissions to be applied, it should really be upgraded to one that does. An example is Windows 95/98 do not allow file permissions since they run FAT or FAT32 file systems. Windows 2000/XP does allow folder/file permissions through NTFS, as long as this option has been selected at install or converted to NTFS later.

All file system permissions applied should use the principle of least privilege, that is only the basic permissions needed should be given. This is particularly relevant for shares, as these are doorways to the file system of the machine. Therefore never give a share full access to everyone permissions unless absolutely needed. The system can have hidden shares by default; it may be a good idea to remove these as well if they are not needed. This is also a good time to check the strength of the passwords on the system as described in step 3.3. There are programs available that are able to search for weak or unprotected shares on a network at the click of a button, so don't leave doors wide open.

Apart from manual user connections, checking permissions are also important for isolating what a service can do. If even after minimising services one is compromised, the service itself can only do as much as the account it runs under is allowed. It is for this and other reasons it is important to lock down the permissions on critical operating system folders and files. Guides are available with recommendations of how to configure permissions on these files, such as in Microsoft's Securing Windows 2000 Server Guide (<http://www.microsoft.com/technet/security/prodtech/windows/secwin2k/a0603.asp>) [29], as well as non-vendor resources such as the Linux Benchmarks offered by the Center for Information Security on their web site at <http://www.cisecurity.org> [21].

3.7.2 Templates

Since rolling out many different permission changes on a large scale is difficult and prone to error, ways have been developed to reduce the effort involved. Such tools have been developed by Microsoft, collectively called the Security Configuration Tool Set, which can be investigated by going to <http://www.microsoft.com/windows2000/techinfo/howitworks/security/sctoolset.asp> [30]. These comprise of multiple MMC snap-ins that enable a template to be created per SOE, which after testing can simply be imported into similar environments or used as part of an auditing process to check compliance. Auditing of permissions using the preferred tool such as a vulnerability assessment tool is important to do regularly, as permissions can quickly become out of date creating potential holes in the security infrastructure.

The templates can be used to enforce many different security settings in Windows as defined in Security Policies/Procedures, such as services to allow, user rights, registry settings, and of course permission changes to system files. A template can be built from scratch, a default template used, or one can be downloaded and configured for the environment, such as those offered by the Center for Information Security [21] for exactly this purpose. Always ensure however any changes are thoroughly tested on non production machines before rolling out.

3.7.3 Remote Access

Apart from local and server permissions, Remote Access Service (RAS) permissions need to be considered. Usually remote access is given by some sort of remote access server. Authentication services are provided by this server and then access is given to the network. There are many different protocols available for remote access authentication, such as Password Authentication Protocol, and Challenge Handshake Authentication Protocol, amongst others. These various options should be researched into with high consideration given to the strongest authentication option possible used. The defaults are often only available for backwards compatibility with older operating systems so more recent protocols considered more secure are a far better choice.

Many vendors provide software in order to improve manageability of multiple remote access servers and to centralise authentication control. Two types of software which offer this functionality are RADIUS and TACACS+. It is advisable to consider an implementation of one of these types of remote access authentication controls to simplify managing this security step. Integration into the network operating system such as Active Directory also improves manageability and can immediately tie in with authorised access levels on the network.

3.7.4 Encryption

Access can also be controlled using encryption techniques. Sensitive data can be encrypted on a machine ensuring only those with the correct decryption key can access it, adding another layer to the security. The Security Policy should cover when to use encryption, such as for highly sensitive email, and Security Procedures should define how it is to be done including data recovery considerations. There are many different products on the market which offer encryption services, one of the most popular being Pretty Good Privacy (PGP- see <http://www.pgp.com/>) [31]. Despite its name, PGP offers excellent encryption features and so far has stood the test of time.

At a minimum, important documents on laptops and mobile PC's should be encrypted. This is because they are outside the controlled environment created on the corporate network, so mobile computers don't always have the

same level of protection. Combined with the fact laptops and the like often carry highly sensitive executive level information, their protection becomes critical. To make this a bit easier Windows 2000 and later contain inbuilt encryption capabilities, using the Encryption File System (EFS). One advantage of this is the encryption and decryption process can be transparent to the user, it is simply a check box in the advanced options of the folder where sensitive documents are contained.

Another way to make EFS even easier to use is by configuring the Data Recovery Agent through Active Directory. By doing this the account which can recover data when the original keys are lost can be the same, reducing the workload required for backing up Data Recovery Agent keys for each laptop. The capabilities of encryption should be combined with permissions and enforced through policies to ensure sensitive data in all locations have sufficient layers of protection.

Access controls can require multiple methods to ensure protection, but combining these options will enable the level of protection to be sufficient.

| Minimum Requirements for Step | Effort Factor |
|---|---------------|
| Permissions with templates, RAS centralised control, inbuilt encryption options | 0.3 |

3.8 Secure Communications and Detection

While following the other steps in this guide may give some assurances that the data stored locally on machines is reasonably secure, data in transit is also something that needs attention. If important information is sent around as clear text, it is only a matter of time before someone with a simple packet sniffer comes along and reads the data. Additionally if someone is in the process of sending communications for the purpose of infiltrating security mechanisms, there needs to be a way to discover this and alert appropriate personnel.

3.8.1 Communications

As with other steps there are different ways to achieve these goals, only some of which will be covered here. One emerging standard for secure communications is IPSec. Using IPSec, it can be assured the source is valid & hasn't changed, and also depending on the implementation the data can also be encrypted end to end preventing the data from being read in transit. IPSec is an integral part of IPv6, the next generation of Internet Protocol. IPSec can be used alone or in conjunction with other technologies, so is supported by many products including some operating systems, enabling taking advantage of its capabilities easier.

The most common use of IPsec is in relation to Virtual Private Networks (VPN). A VPN can be used to create a secure tunnel between two locations, usually used when the data needs to travel through a public network such as the internet. A VPN is most often used for client machines connecting to a corporate network for remote access, or for secure communication between two businesses avoiding the expense of other alternatives such as leased lines. Both these scenarios are important times to take advantage of technologies such as a VPN, as without doing so all sorts of company information assets may be exposed.

The process of setting up a VPN could be quite complex depending on how it is handled and which product is used, although there are easier ways of doing it. One way to setup a LAN to LAN VPN is utilising the inbuilt capabilities of existing firewalls or routers. A firewall to firewall or router to router VPN may mean no additional hardware resources are required and it is just a matter of configuring them correctly. Additionally there may be inbuilt capabilities within the operating system to set up a server for client remote access through VPN. These make the setting up of a VPN from both the server and client ends very simple, for example using the inbuilt VPN capabilities of Windows 2000 Server.

As explained on the Microsoft web site at <http://www.microsoft.com/windows2000/server/evaluation/business/remote.asp> in relation to this:

For basic networks, the New Connection Wizard walks you through the set-up of the remote access server for both direct-dial connections and for VPN. Plug and Play modem configuration makes it easy to install a modem in the server. The wizard asks a few simple questions to determine if you want dial-up, VPN, or both types of access, and asks you which network interface to allow the connections to come in on. Next, it presents you with a simple list of known users; just check the box on the users you want to allow remote access for, and you're essentially finished. [32]

It is important to note some of the negatives with using a VPN. Firstly it can't be used with Network Address Translation (NAT). Secondly encrypting the data can mean not only the bad guys will be prevented from reading it, firewalls, IDS software and virus scanners also cannot analyse this data to perform their functions. Thirdly using a VPN is implying trust with the party at the other end; therefore compromise of the remote end can quickly lead to compromise of the corporate network. This is one reason for example, why home users connecting to the corporate office need to be included in antivirus policies. However as long as these negatives are taken to consideration, VPN technology has an important role in securing the communications on a network.

Another method for securing communications, particularly in regards to secure web site transactions, is using Secure Sockets Layer (SSL). SSL uses certificates to prove the authenticity of a web site, and ensures the traffic between the client's computer and the web site is encrypted. SSL capabilities are included in most modern web servers and are usually fairly easy to set up. A certificate can be purchased from a trusted supplier on the Internet such as Verisign (see <http://www.verisign.com/products/site/index.html> for more information) [33], and when enough proof of identity is given a certificate can then be used to verify the web site. Although it doesn't have to be done this way and an internal Public Key Infrastructure (PKI) and certificate management services can be set up, it is far simpler to use the resources of a third party. Then it is just a matter of importing the certificate to the site and configuring the site to use it. Other than SSL and VPN, other forms of encrypted communications can also be utilised, such as PGP for securing email as mentioned in the previous step. As before, policies need to define what should be used and when to use them.

3.8.2 Intrusion Detection

Not only is it important to secure the network infrastructure, and secure communications, but also to be alerted when an intrusion has been detected. Intrusion Detection Systems (IDS) can be divided into network-based and host-based, depending on their location and role. Without an IDS there is often little way of knowing a compromise has even taken place. In any survey on security intrusions the most worrying statistic is the "I don't know" group. It isn't possible to clean up the mess of an intrusion and prevent it from happening again if it is not even known about, and it is very difficult to deal with even if known if the extent of the damage is unclear.

Enter Intrusion Detection Systems, whose job it is to send alerts triggered by suspicious activity and record impact and changes. Both network and host based IDS play an important role in this, and should be used together. A very popular network based IDS is Snort (see <http://www.snort.org/>), a free highly configurable open source application with powerful IDS features [34]. There are also open source versions for host based IDS options such as the Linux version of Tripwire, available at <http://www.tripwire.com/products/linux/> [35].

Using a product such as Tripwire makes it quite quick to discover a compromise has occurred and what has changed as a result, including aiding in discovery of a root kit being installed. A root kit is when operating system files are replaced with Trojan versions, and as the operating system still behaves as normal can be very difficult to detect. Of course there are commercial versions of both kinds of IDS products that should also be investigated to discover the most suitable options for the environment.

Because these products are so powerful, they can be difficult to configure when inexperienced. One way that these products are made easier to use is when they are integrated into existing services. For example Microsoft's

Internet Security and Acceleration Server (ISA), includes a network IDS (see <http://www.microsoft.com/isaserver/evaluation/features/security/intrusion.detection.asp>) [36], that requires little configuration on top of the firewall set up and is based on the technology of Internet Security Systems (www.iss.net). Capabilities for a simple host based system can already exist on the machine, for example using TCPWrappers and Syslog capabilities on Unix/Linux machines. Adding simple log monitoring software which has event triggering can turn the machine in question to a reasonably functioning host IDS without too much effort or cost.

Securing communication and setting up intrusion detection isn't a one step process, but with a little consideration can quickly become a powerful highly integrated part of the network security arsenal.

| Minimum Requirements for Step | Effort Factor |
|---|---------------|
| Communication with IPSec, VPN, SSL etc | 0.5 |
| Network & host IDS using integrated solutions and intelligent third party decisions | 0.3 |

3.9 Unauthorised Network Equipment

The topic of Unauthorised Network Equipment may not always be in security checklists, but this growing problem should be included as an important security consideration. Essentially detecting unauthorised equipment means finding and dealing with "back doors" to the now quite secure network that has been configured following the previous steps. Ensuring there are many layers of security from the perimeter in is all very good and well until someone comes and installs a modem to an ISP from their machine, effectively bypassing much of this protection. Even worse than this scenario is not knowing about it for months or even years, causing untold amounts of damage or at least risk to the core of the network.

3.9.1 Policy and Detection

The first requirement in tackling this important issue is the same as in each of the previous steps; they must be covered in Security Policies. What is acceptable and not acceptable for installing on the network, who is authorised to approve equipment installation and who must perform the installations should be clearly spelled out in policies and procedures. This information needs to be included in user education as described in step 1 so that no users of the network have the reason or excuse that they weren't aware.

Once appropriate equipment and procedures have been defined, routine audits need to be undertaken to ensure the entire network is complying fully

with this policy. Examples of equipment that should be included in unauthorised equipment audits are modems, wireless LANs (WLAN), Remote Access Servers (RAS), Virtual Private Network (VPN) Servers, and any other equipment that enables unapproved access to the corporate network.

3.9.2 Modems

Once again using software tools can make the job of equipment detection easier. For modems, there are a couple of ways to enable detection. One is to use vulnerability scanning software that has this capability, remembering to always gain written permission first. A vulnerability scanner that also provides reports on hardware installed on a remote system is a good way to cover more needs with a single tool and allow different types of audits to be combined into one, reducing the effort & resources required. Therefore considerations such as this should be given before the purchase of a vulnerability scanner.

A second method for modem detection is to use a tool for War Dialing. War Dialing is the term used to describe automated phone dialing and recording software that attempts to discover potential modem targets. It is always a good idea to do this for all company phone lines before someone with bad intentions does the same, as long as permission is given to do so. There are a number of resources available on how this can be done, such as "War Dialing Your Company- A HowTo" [37], a GSEC Research Project available at http://www.giac.org/practical/GSEC/Dave_Owens_GSEC.pdf, and "Modem Scanning" [38], a publication from the Computer Security Resource Center (CSRC) located at <http://csrc.nist.gov/fasp/FASPDocs/security-controls/USAIDModemScanBSP.htm>.

3.9.3 WLAN

Another increasingly popular method of network access is using a WLAN. Setting up a WLAN can be as simple as plugging in a wireless access point, so doesn't need a great deal of technical competence. Unfortunately this means the convenience of wireless can be setup by someone with no security considerations whatsoever. Apart from the obvious situation of setting up unauthorised wireless access with pre-conceived plans to do the wrong thing, even if it is done with the best of intentions having wireless access is like waving a red flag to a bull. It is then simply a matter of someone walking or driving past with a basic hardware set up to detect and abuse an opportunity to use company resources for free Internet access or worse.

As with modems, it is up to network/security administrators to find and shut down rogue wireless networks before anyone else can find them. A popular tool to aide in doing this is software called Netstumbler. According to their web site at <http://www.netstumbler.com/>, "once the downloadable software has been installed on a user's PC, it automatically sniffs out available Wi-Fi

network 'hotspots.' " [39]. Using software such as this with a basic hardware setup like a laptop, wireless card and antenna, the administrator can physically check the corporate offices and surrounds in order to discover these networks.

3.9.4 Other Equipment

Finally RAS, VPN Servers and the like need to be identified and noted as to whether they are authorised to be configured in this role or not. A negative impact of having simple wizards to set up VPN servers as described in step 8 is that once again the staff are capable of doing so without considering the impact on security. All servers with this potential should be checked as part of the auditing process. The easiest way to include these checks is to do them as part of other audit procedures, such as when scanning hosts for running services. Remote access related services should then be given high priority for further investigation when analysing the results.

Unauthorised network equipment is an area that does deserve attention so that much of the other security improvements done on a corporate network are not simply bypassed.

| Minimum Requirements for Step | Effort Factor |
|--|---------------|
| Equipment detection using scanning software and including in existing audits | 0.2 |

3.10 Maintaining Security

Unfortunately there is no point in time where it is possible, having done everything recommended to sit back and decide the information is now secure and the security related tasks are complete. In order for security to be effective, it must be looked at on an ongoing basis, using well defined routines to ensure all aspects are as up to date as possible. These routine security tasks need to be converted into procedures that are then implemented by the appropriately trained personnel.

3.10.1 Security Training

This is an important point, as having someone untrained look through the logs for a specific device may defeat the purpose of what is trying to be achieved. Indeed utilising personnel with strong skills specific to each step in the guide, or up skilling staff in the specific security related area they are working on, should be a major consideration and should be done wherever possible, as the impact and quality will then be far greater. These needs and their importance were briefly covered under education in the first recommended

step, but for emphasis education expands from being important not only for users of the network but for the security practitioners charged with carrying out recommended guidelines and enhancing the existing corporate environment. In short, all staff need to be trained in security to the level of the role security plays in their job, which affects all staff to varying degrees.

3.10.2 Backups

The security specific procedural routines that need to be done will depend on the infrastructure and policy specific to the particular company. There are however some common themes that will usually need to be done in any organisation. One of these is backups, which while they may not immediately spring to mind as particularly related to security, the use of backups and how the media is handled is something that does need to be considered from a security perspective. Backups are used for multiple purposes, the first being the obvious data recovery functionality in the event of data loss or corruption.

Additionally however, backups are used to recover from a security incident where backups can be used to return a system back to a known pre-incident state. Backups can also be used for computer forensics, where a snapshot of a system can be used of the system after the incident has occurred for later analysis and possible evidence for prosecution. The best practice recommendations for the proper handling of media from a forensics perspective are a relatively complex area so will not be attempted to be covered here, suffice to say this is a role backups play in security.

There are multiple ways to achieve a backup. The most common method of performing routine backups is by using a backup program, such as the one included in Windows 2000 systems, dump, tar or dd in Unix/Linux, or a third party choice such as Veritas BackupExec for Windows (see <http://www.veritas.com/products/category/ProductDetail.jhtml?productId=bews>) [40]. One consideration in choosing a backup program is manageability, such as options for centralisation, to reduce the effort required when there are many systems to backup. An example is Legato Networker software (www.legato.com/products/networker/), which allows management of backups for many operating systems including Windows, Macintosh, Linux and Unix [41].

Another method of performing backups is using imaging or cloning techniques, the latter being quite useful for forensics. Disk drive imaging creates a copy of all the contents of a drive and stores this in the form of a file located elsewhere. Disk cloning takes this concept a step further and copies all the data from a disk at a lower level, including all the sectors and even ones with free space. A commonly used product for these types of backups is Symantec Ghost (for more information see <http://enterprisesecurity.symantec.com/products/products.cfm?ProductID=3&EID=0>), which makes this task quite simple [42]. Using an application such as this enables before and after comparisons, due to being good for base-lining

systems for later use. Also these applications are used for replicating an identical copy of the system elsewhere such as for performing restores in the event of data recovery.

Methods of backups, frequency, and deciding on which systems require them are all important considerations, which for the latter two can simply mean as much and many as possible, as mostly this just depends on the environment and resources available. However additionally the way backup media is handled needs to be looked at and added to policies. Media handling considerations are part of routine tasks, such as tape rotations, as well as general policies, like the location and method of offsite storage. Offsite storage should be at a physical location far from where the backups were conducted, to cater for data loss due to natural disaster. How the media is to be securely transported, and how it is stored also need to be decided.

An example is to securely store magnetic media, a safe that is fireproof as well as offering electromagnetic protection should be used, with the media within clearly identified. Finally in regards to backups, trial restores using various scenarios need to be routinely performed and part of operational procedures. Without doing this, there is no way of being sure the backup and restore procedures are correct and proficient.

3.10.3 Auditing and Log Analysis

Other than backups, other important routine tasks that should be part of normal procedures are auditing and log analysis. The auditing needs of various devices and system set ups have been covered where relevant in each step of this guide; however these audits are not designed to be one off events. Auditing is intended to be ongoing regular procedure to ensure all aspects of the infrastructure are in full compliance with the security policies and procedures in place. Ways to reduce the effort involved in auditing have also been discussed previously, often using tools which simplify the process of information gathering.

However it is not enough to simply have reports if they are not scrutinised and used for security enhancements and closing known or potential holes. Therefore time and resources must be given for this process, or else all the previous effort involved to get to that point is going to waste.

In addition to assessing the reports of routine audits and taking action according to the results, time and resources also need to apply to log analysis of the various systems and devices. This can be a large undertaking, as there are often many systems of differing types, all of which can each offer logs of different types, collectively providing a large amount of information. This information is very valuable for providing incite into what is happening on the network, and can also indicate important security events, incidents that need attention, and potential security flaws to correct.

The importance of log analysis procedures needs to be stressed and not skipped, despite the challenges associated with doing this. An IDS system may be indicating all sorts of intentionally devious activity, and a firewall may be recording regular penetrations and clear miss-configurations, but if they are not discovered and dealt with by appropriate staff in a timely fashion the effectiveness of such security measures is drastically reduced.

One way timely notifications of serious events can be achieved quite easily is through configuration of alerts. Alerts can immediately notify staff of an issue to allow them to be dealt with. Products used for any security procedures that allow thorough logging and event notification, or add on tools that provide this functionality should be investigated and given priority. On top of alerts triggered by specific known events, the logs still need to be regularly reviewed for suspicious activity. This is where security management products play a part in reducing the effort in sifting through all this information. NetIQ are one of many companies to offer a solution for this. NetIQ offer a product called VigilEnt Log Analyzer to meet the need for log consolidation and management.

According to the NetIQ site at <http://www.netiq.com/products/vlm/default.asp>:

VigilEnt Log Analyzer provides a complete enterprise solution for log archival and consolidation, security event analysis and log forensics. It enables security officers and administrators to truly analyze and understand the security events from a wide variety of operating systems, firewalls, intrusion detection systems and other devices. [43]

NetIQ also offer many other security management products aimed at the enterprise level (for more information investigate <http://www.netiq.com/solutions/security/default.asp>) [44], as do many other companies that offer competitive alternatives or entirely different product ranges with the same goals. The common intention of these products is to reduce the load of security related activities, which in turn can free up the time of those responsible, giving them more time for other important tasks.

Security maintenance is something that well defined procedures and a wide variety of tools make easier and engrained into normal operations, making the impact on resources less but ensuring the ongoing security of all corporate resources are to a high standard.

| Minimum Requirements for Step | Effort Factor |
|---|---------------|
| Security staff training, backup utilities, audit tools, alerts and using log analysing software | 0.6 |

4 Conclusion

In order to achieve the goal of security for the corporate network, ensuring the confidentiality, integrity and availability of data, there are many steps that need to be performed. Ten of the more critical aspects of information security have been covered in this paper to some degree or another, offering a background and starting point to investigate the implications of security steps that should be put into practice.

An understanding of why specific tasks are recommended can emphasise the importance and relevance to the guidelines, therefore these have been covered where appropriate. The step by step guide format may aid in describing what is involved in practical security more clearly. The application of best practice techniques even by professionals are acknowledged to not always be followed through on, with a lack of knowledge and resources most often to blame. With this in mind, ways exist to make the implementation of the techniques described easier, and to reduce the effort in terms of time, resources and management related responsibility on staff.

These methods are combinations of technical and procedural activities, often achieved by introducing additional tools, all with the purpose of achieving the maximum benefits of security whilst requiring the minimal effort. If by doing so raises the bar on the security recommendations actually put into practice by any business or company, the primary goal of this paper has been accomplished.

To further aid in these intentions, an Effort Factor for areas covered in each step can make clearer the most basic application that should be done, as well as a score indicating the difficulty in relation to other components. Not to be taken at face value, each Effort Factor may give guidance to the work required and help in designing timelines and assigning priorities, particularly when the topic is unfamiliar.

Reducing the effort in practical security application is one way to increase the likelihood of steps actually being put into place, which in turn may mean higher quality security across a broader percentage of the population, thus improving the environment most must interact with everyday. By doing so would make the increasingly interconnected world in which we live of less concern and a more pleasurable experience for all.

5 References

- [1] SANS Institute. "The Top 7 Management Errors that Lead to Computer Security Vulnerabilities" 14 May 1999. URL: <http://www.sans.org/resources/errors.php> (30 March 2003)

- [2] 8020 Info Inc. "About 8020 Info". Section: The 80:20 Principle. URL: <http://www.8020info.com/principle.html> (30 March 2003)
- [3] Symantec. "80-20 Rule of Information Security". URL: http://securityresponse.symantec.com/avcenter/security/Content/security_articles/fundamentals.of.info.security.html (30 March 2003)
- [4] Fraser, Barbara Y. "RFC 2196 Site Security Handbook". September 1997. URL: <http://ietf.org/rfc/rfc2196.txt?number=2196> (30 March 2003)
- [5] "The ISO17799 Service & Software Directory". 27 April 2003. URL: <http://www.iso17799software.com/> (27 April 2003)
- [6] SANS Institute. "The SANS Security Policy Project". Section: Need an Example Policy or Template? URL: <http://www.sans.org/resources/policies/> (30 March 2003)
- [7] SANS Institute. "Security Policy Issues". InfoSec Reading Room. 4 December 2002. URL: <http://www.sans.org/rr/policy> (30 March 2003)
- [8] Cisco Systems. "Network Security Policy: Best Practices White Paper" 24 April 2003. URL: <http://www.cisco.com/warp/public/126/secpol.html> (27 April 2003)
- [9] CompTIA. "CompTIA Survey Reveals Human Error Most Likely Cause of IT Security Breaches". CompTIA's Press Room. 18 March 2003. URL: http://www.comptia.org/pressroom/get_news_item.asp?id=207 (30 March 2003)
- [10] Oracle Corporation. "Network Security and Analysis Recommendations". Section: IOS Router Tips. URL: http://www.oracle.com/ecostructure/blueprint_rec/network_security_analysis_and_recommendations.htm (5 April 2003)
- [11] Naidu, Krishna. "Cisco Checklist". Security Concensus Operational Readiness Evaluation (S.C.O.R.E.). URL: <http://www.sans.org/score/checklists/CiscoChecklist.doc> (5 April 2003)
- [12] Wack, John. Cutler, Ken. Pole, Jamie. "Guidelines on Firewalls and Firewall Policy". US National Institute of Standards and Technology (NIST) Computer Security Response Center (CSRC). January 2002. URL: <http://csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf> (5 April 2003)

- [13] Microsoft Corporation. "Chapter 5- Securing the Domain Infrastructure". Section: Domain Policy, Password Policy. Microsoft Solution for Securing Windows 2000 Server. 5 February 2003. URL: <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/pr odtech/Windows/SecWin2k/05secdom.asp> (6 April 2003)
- [14] Microsoft Corporation. "Microsoft Baseline Security Analyser". URL: <http://www.microsoft.com/technet/treeview/default.asp?url=/TechNet/Security/ tools/tools/MBSAHome.ASP> (6 April 2003)
- [15] Microsoft Corporation. "Chapter 6- Hardening the Base Windows 2000 Server". Section: Windows 2000 Server Baseline Policy. Microsoft Solution for Securing Windows 2000 Server. 5 February 2003. URL: <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/pr odtech/Windows/SecWin2k/06basewn.asp> (6 April 2003)
- [16] Microsoft Corporation. "IIS Lockdown Tool (version 2.1)". 14 February 2001. URL: <http://www.microsoft.com/windows2000/downloads/recommended/iislockdow n/default.asp> (6 April 2003)
- [17] GFi Software. "GFi LANGuard Network Security Scanner Overview". URL: <http://www.gfi.com/lannetscan/index.htm> (6 April 2003)
- [18] Nessus Project. "Nessus Security Scanner Features". URL: <http://www.nessus.org/features.html> (6 April 2003)
- [19] Microsoft Corporation. "Microsoft Windows Update". Version 4. URL: <http://windowsupdate.microsoft.com> (13 April 2003)
- [20] Microsoft Corporation. "Hfnetchk". URL: <http://www.microsoft.com/technet/security/tools/tools/hfnetchk.asp> (13 April 2003)
- [21] Center for Internet Security. "The Center for Internet Security". URL: <http://www.cisecurity.org> (13 April 2003)
- [22] Red Hat Inc. "Chapter 2- Red Hat Update Agent". Red Hat Network Basic User Reference Guide. URL: <http://www.redhat.com/docs/manuals/RHNetwork/ref-guide/up2date.html> (13 April 2003)
- [23] Microsoft Corporation. "Software Update Services". URL: <http://www.microsoft.com/windows2000/windowsupdate/sus/> (13 April 2003)

[24] Microsoft Corporation. "Microsoft Systems Management Server". 23 April 2003. URL: <http://www.microsoft.com/smsserver/default.asp> (27 April 2003)

[25] The Wild List Organization. "The Wild List Organization International". URL: <http://www.wildlist.org/> (18 April 2003)

[26] Internet Kmoraine. "Choosing Antivirus Software". Internet Kmoraine Support. 22 February 2003. URL: <http://frontpage.kmoraine.com/antivirus.asp> (18 April 2003)

[27] McAfee Security. "About McAfee Active Virus Defense". URL: <http://www.mcafeeb2b.com/aboutmcafeeb2b/default.asp> (18 April 2003)

[28] Symantec Corporation. "Symantec Enterprise Solutions". URL: <http://enterprisesecurity.symantec.com/> (18 April 2003)

[29] Microsoft Corporation. "Appendix C - Optional File System Permissions". Securing Windows 2000 Server. 5 February 2003. URL: <http://www.microsoft.com/technet/security/prodtech/windows/secwin2k/a0603.asp> (18 April 2003)

[30] Microsoft Corporation. "Security Configuration Tool Set". 19 April 1999. URL: <http://www.microsoft.com/windows2000/techinfo/howitworks/security/sctoolset.asp> (18 April 2003)

[31] PGP Corporation. "PGP Corporation". URL: <http://www.pgp.com/> (18 April 2003)

[32] Microsoft Corporation. "Connecting Telecommuters and Remote Employees". 29 December 2000. URL: <http://www.microsoft.com/windows2000/server/evaluation/business/remote.asp> (20 April 2003)

[33] VeriSign Inc. "SSL Certificates". URL: <http://www.verisign.com/products/site/index.html> (20 April 2003)

[34] Snort Organisation. "Snort Open Source Network Intrusion Detection System". 27 April 2003. URL: <http://www.snort.org/> (27 April 2003)

[35] Tripwire Inc. "Tripwire Open Source, Linux Edition". URL: <http://www.tripwire.com/products/linux/> (20 April 2003)

[36] Microsoft Corporation. "Integrated Intrusion Detection". Microsoft Internet Security & Acceleration Server. 03 May 2001. URL:

<http://www.microsoft.com/isaserver/evaluation/features/security/intrusion.detection.asp> (20 April 2003)

[37] Owens, Dave. "War Dialing Your Company: A Howto". 10 December 2000. URL: http://www.giac.org/practical/GSEC/Dave_Owens_GSEC.pdf (20 April 2003)

[38] Craft, James P. "Modem Scanning". Version 1.1. US National Institute of Standards and Technology (NIST) Computer Security Response Center (CSRC). 23 January 2001. URL: <http://csrc.nist.gov/fasp/FASPDocs/security-controls/USAIDModemScanBSP.htm> (20 April 2003)

[39] Netstumbler Software. "Netstumbler dot com FAQ (Frequently Asked Questions)". Version 0.3. URL: http://www.netstumbler.com/modules.php?op=modload&name=FAQ&file=index&myfaq=yes&id_cat=1&categories=Official+NetStumbler+Version+0.3+FAQ (20 April 2003)

[40] Veritas Software Corporation. "Veritas BackupExec for Windows Servers". URL: <http://www.veritas.com/products/category/ProductDetail.jhtml?productId=bews> (21 April 2003)

[41] Legato Systems Inc. "Legato Networker". URL: www.legato.com/products/networker/ (21 April 2003)

[42] Symantec Corporation. "Symantec Ghost Corporate Edition 7.5". URL: <http://enterprisesecurity.symantec.com/products/products.cfm?ProductID=3&EID=0> (21 April 2003)

[43] NetIQ Corporation. "VigilEnt Log Analyzer". URL: <http://www.netiq.com/products/vlm/default.asp> (21 April 2003)

[44] NetIQ Corporation. "Security Management and Administration". URL: <http://www.netiq.com/solutions/security/default.asp> (21 April 2003)

Other Resources-

Cole, Eric. SANS Security Essentials + CISSP CBK. SANS Institute, 2003