



SANS Institute Information Security Reading Room

Security from Scratch ... How to Achieve It

Alan Davies

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Security from Scratch ... How to Achieve It

Alan Davies

September 4, 2001

If you find yourself in a situation where you're working for a company that has put together an IT infrastructure and the only real concerns have been functionality and performance, then this document is aimed as a guideline for starting off an information security culture. This will be achieved through policies and the use of various tools to analyse your systems and network – the end result should be a series of reports you can present on the current state of security in your company and a roadmap built on that to improve it based on a risk analysis.

First of all, the goals of securing a network can be summarised as ensuring the following: [1]

- Confidentiality
- Integrity
- Availability

Since there is no one technology or process that can be implemented in the name of total security, the aim is to develop a defence in depth strategy. One of the great things about this is that it doesn't require a huge budget – lots of small hurdles are generally much more effective than one big expensive one. High-end solutions can be added on later as needs outgrow the initial implementation.

To simplify the process of putting in place some basic and generally low cost security, I will break it down into 8 areas. When starting off, you will most likely be faced with limited or non-existent budget for tools. You will also likely have few resources to get the work done, keep it up to date and keep it adequately monitored. Hopefully this document will provide the proof of concept required to start the ball rolling and create an IT security department – or at least an initial function similar to this if in a small company. There is a 'Software Used' section at the end of this document, which lists all tools mentioned throughout, and the URL you can get them from.

Know your system

In order to protect your network and systems, you must know them. Without a good understanding of how your network works and what is connected to it, you stand no chance of securing it. If they are not already in existence – make some network schematics mapping out the internal LAN and WAN structure, clearly showing any outside links (be they to private networks which are not owned by the company or to the internet). You should also keep documentation on every server in your company – this should state the functions of that server (i.e. the services that it provides), who needs to access it and who can grant access to it.

You need to pay particular attention where your network connects to the internet. If your ISP manages it for you, what are they doing to protect you? Can you verify they are actually doing it? What other services do they offer – what other services

do their competitors offer? If you manage your own internet access, you need to know what is currently done to secure it. As firewalls are such a huge topic, I would highly recommend reading through William Davis's excellent paper "Firewalls: What I Wish I'd Known When I Was Getting Started". [2]

If you're using RAS, how secure is it? What level of authentication are you using? Maybe you're using VPN's instead (or indeed as an office to office link rather than to support remote users). If so, the same questions apply.

You should also inspect the traffic on your network so you know what is actually going on day to day – both from a capacity management and an intrusion detection point of view. This will reveal what protocols are in use, what systems are being accessed and who is accessing them. It will give you a benchmark for the normal performance of your network that can then be used later to help identify abnormal behaviour. There are many products out there, both dedicated hardware and software solutions, which can create reports for you with this information. I have personally found the OneTouch Series II (and its reporter software) from Fluke Networks particularly good as a hardware device and LANScan Professional from LANScan Software Limited to be good as a software option.

Risk Assessment

The Ohio State University Computer and Information Science department states "Risk analysis involves determining what you need to protect, what you need to protect it from, and how to protect it. It is the process of examining all of your risks, then ranking those risks by level of severity. This process involves making cost-effective decisions on what you want to protect." [3]

The schematics produced in the previous step will identify most of the assets you need to protect. Identifying the threats can be achieved by examining what will threaten the integrity, confidentiality or availability of those assets. From a technical point of view there are many tools that can scan systems to see what vulnerabilities they are exposed to. You need to decide how much of a threat each of these vulnerabilities is. From a non-technical point of view, you should look more towards threats relating to the physical access of systems and the quality/reliability of them.

As you move along in your security implementation you will be able to conduct more and more accurate risk assessments. It is these reports that will get you management buy-in for your security resources. Every system and the data on it can be assigned a monetary value. More detailed information on this can be found in David Litzau's article – Risk Management: A Foundation for Information Security. [4]

Security Policy

This is one of the most important steps – creating and enforcing a security policy. Without one, you would have no way of enforcing good practice. However, without full management backing, a policy will be almost useless. It needs to be approved from the top and signed by everyone. A security policy tells people what they can and can't do with your company's technology assets. It also sets out what you do to

protect these assets (e.g. backup strategy, anti-virus deployment, incident handling, disaster recovery, etc.). It is very important to audit and review your policy once created to ensure it is still relevant and meets your company's needs.

A good policy should be short, concise and easy to digest. While you should consult with your legal and HR departments over the policy, legal talk will only confuse and make it boring to read. It is also important that the areas of the policy that address the highest risks stand out compared to other low risk items. Having a page at the front summarising the highest risks without going into detail may be beneficial. SANS have produced a document similar in purpose – Mistakes People Make that Lead to Security Breaches. [5]

Many tools can be used to monitor adherence to your policy. For example, to regularly check that your users are using strong passwords, tools such as LC3 (formally L0pht Crack) for Windows or Crack for *NIX systems can report how secure account passwords are. LC3 in particular has the ability to audit passwords on Windows NT/2000 systems and present the results *without* revealing the passwords – very useful for demonstrations and producing “safe” reports. You might also setup a packet sniffer (Iris, Dsniff or Ettercap) to check for rogue network activity such as LAN gaming. Many of these products allow you to filter the output if you're looking for something in particular. There are countless products to monitor email attachments and web surfing habits too.

An equally important step however is to make sure your users are educated on the contents of the policy. Just because they signed something, it doesn't mean they read and understood it. They may have even read and understood it, but decided to ignore it! The tone can be quite important too – it should not be vague, but should avoid using intimidating or threatening language.

It's usually a good idea to invite yourself to departmental meetings or organise separate meetings where you can talk to all of your staff about what the policy means and why it is there. The best way to get their attention is to use some “wow factor” by showing them what can happen if some policies aren't followed – cracking their passwords in a few seconds for example (making sure the results are masked!). If they now understand that there is a good reason for not using a simple password, they are more likely to follow the requirements in your policy to use a more complex one.

It can be very daunting trying to create a security policy from scratch. However, there are many sample policies that can be found on the internet for free. There are also both books and software products that have policy templates that you can customise to suit your needs. If you do a quick search on Google or any other search engine, you should get plenty of results for “sample security policy”.

Scanning Tools

In order to convince management that there is a need to spend money on security, they need to perceive a risk. The risk assessment will give them the cold facts on what needs to be done. However, sometimes a more graphic demonstration can help. If you can actually show them that it takes less than 10 seconds to crack their

password or that the internet facing equipment is being actively probed (or maybe even attacked!) then it could make your job a lot easier.

A penetration test on your network from the internet is a good idea. It is much better to see what weaknesses you have for yourself and try to fix or avoid them rather than for someone else to find them and exploit them maliciously! You can do this yourself with third-party tools (some free, some expensive) or by completely outsourcing the test (make very sure you use a reputable company). Whichever you choose, it is of the utmost importance that you get *signed* permission from upper management to run any of these tests. Not only could there be legal implications for breaking the security on certain systems or revealing sensitive passwords, but some of the tests could potentially take down the systems that they scan!

NMAP and Nessus are perhaps the best of the free tools, and indeed, the ones most likely to be run by a would-be intruder to scan your internet hosts. NMAP is an advanced port scanner, while Nessus looks for vulnerabilities in the hosts it scans. There is a Windows version of NMAP (no GUI) and Nessus (client only) but I would recommend running the Linux versions of them. One other free piece of software I like is Cerberus' Internet Scanner which, although it isn't currently as good as the likes of Nessus, makes a good addition to your arsenal of tools. Greg Saoutine of MCP Magazine provides an interesting article on the leading products in this field called "How Secure is Your Network?" [6] which includes many of the leading (and often expensive) tools for comparison.

Hardening Operating Systems

The next important step is to harden both your servers and clients. A default installation of any operating system is inherently insecure. As a general guide, only enable services that are required – all others should be disabled or removed (for example a default install of Win2k Server installs IIS – not a lot of use if you're setting up a file server and a very big security hole as a default install). Also service packs and security patches should be applied as soon as possible once they are released. Please test extensively before applying to a production environment though! To keep up with the bugs and vulnerabilities that constantly appear, I would recommend subscribing to some of the free mailing lists such as the SANS security digests or Bugtraq.

There are many detailed guides which go far beyond the scope of this document for securing Windows, Macintosh and *NIX systems. The NSA has provided guides on security research for Windows 2000, Linux and Cisco routers [7] and Patrick Harris has written an article "Macintosh Internet Security Basics". [8] Also, many of the scanning tools listed in the previous step will help in making sure the changes you make actually plug security holes. Remember that, particularly with internet facing systems, this is a constant challenge. If/when you get hacked, a good security policy will be of immense help (server backups, disaster recovery plans, incident handling policy, etc.).

Intrusion Detection Systems (IDS)

An IDS is used to identify and hopefully stop unauthorised access to your systems by either internal or external sources in real time. There are two different types of IDS – host based and network based. Either can be used depending on your needs, but using both offers the best protection.

Host based IDS look at the OS to see if anything unusual happens on the system in comparison to a database of known security violations and custom policies. When something does happen it can react by logging the event, alerting someone or sometimes by stopping the action in its tracks.

Network based IDS look at the traffic on your network to see if there are any patterns in it that match any of the attack signatures in its database. As long as the database is up to date, this makes it very easy to detect any “script kiddies” who may be playing around with hacking tools. Again, the IDS can react by logging the event, alerting someone or by terminating the session.

Needless to say, it is very important that you keep your IDS database up to date. This is an area where a good vendor can make a big difference, just like with anti-virus software. Vendors such as AXENT, ISS and CyberSafe do both types of IDS. The SANS Institute web site has a truly massive amount of information on intrusion detection in their Intrusion Detection FAQ. [9] Highly recommended reading.

Auditing

Auditing is needed because no matter how much you secure your network and systems, there is always the possibility that a breach will occur. Total security is not possible – an acceptable risk level is what you look for. When a breach does occur, first of all, you need to know it happened – hopefully your IDS will have alerted you of this if it were a network attack. Many breaches go unnoticed – for example, an employee managing to crack the password of a user in the HR department and using it to read the salary database for the company. Another could be an employee/visitor/intruder sneaking into the CEO’s office during lunch, finding an unlocked workstation and reading sensitive email or documents. As long as nothing is deleted or changed it could well go totally unnoticed.

Even assuming that you have basic auditing procedures in place, a skilled hacker may be able to remove all evidence of their presence if you haven’t taken adequate countermeasures. Log all services and secure those logs. Obviously you want to have as much information about the attack as possible – if only to determine the extent of what was compromised. You may also need it as evidence to take legal action. More importantly, if you know what happened, you hopefully can stop it from happening again. Be careful though – one of the biggest challenges is managing the deluge of data this can produce!

Auditing in detail is beyond the scope of this document. Windows and *NIX systems both allow you to audit many different event types. The SANS coursework for Security Essentials covers some of this. Rajeev Gopalakrishna has written an article called “Audit Trails” [10] which describes some of the issues with auditing computer

systems. It has many good references to other books, sites, articles and reports on auditing.

Change Control

Change control plays a vital role in the security of your company's electronic assets. Over time, many systems change or fail. If you have no record as to what was changed, how can you ensure consistency with other similar systems? If a system fails and you don't have the documentation as to how it was secured and what changes were made over time, how can you hope to get it back into production quickly and safely? There is little point in having systems in your company setup completely differently because different people built them according to what they thought were sensible guidelines. If this is the case you will have no idea how secure your company is.

To implement change control properly you need to find out who in your company does what within the IT department. Every role is likely to be a part of your change control procedures. This means every person in the IT department needs to be aware of them and needs to stick to them. Some examples of what might be covered by change control would be client builds for employee's workstations, server builds, router setup, etc.

Change control is a bit tedious, but well worth the effort if you have a vital internet facing server die in the middle of the night and need to not only get it back online ASAP, but make sure it is just as secure and functional as it was before it died. You might also think about keeping change control documents with your disaster recovery documents (which I'm sure you already have!) in print format and accessible externally somewhere secure. That way if your office goes up in smoke, you'll be able to find them!

Conclusions

All of the steps outlined in this document are not something you can just do and then step back and say you are secure. They are all a continuous process. As you start to implement them you will learn an incredible amount about your network and the systems connected to it. This in itself will mean you are in a far better position to know what is going on with your systems. Coupled with all the policies and monitoring/scanning tools you will be able to confidently push for changes based on an accurate risk analysis. From there it should be much easier to justify the resources you need to keep up with all the monitoring and reviewing that the security process demands. This also has the added benefit of making the internet community that little bit more secure.

References

[1] – Fried, Stephen. "Information Security: The Big Picture – Part I". SANS Institute Information Security Kickstart Highlights (2001): 1-19

- [2] – Davis, William. “Firewalls: What I Wish I'd Known When I Was Getting Started“. October 2000. URL: http://www.sans.org/infosecFAQ/start/fw_start.htm (21 August 2001).
- [3] – Fraser, Barbara. “Site Security Handbook“. September 1997. URL: <http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc2196.html> (1 August 2001).
- [4] – Litzau, David. “Risk Management: A Foundation for Information Security“. June 2001. URL: http://www.sans.org/infosecFAQ/audit/risk_manage.htm (6 August 2001).
- [5] – SANS Institute. “Mistakes People Make that Lead to Security Breaches“. URL: <http://www.sans.org/mistakes.htm> (16 July 2001).
- [6] – Saoutine, Greg. “How Secure is Your Network?”. URL: <http://mcpmag.com/Features/article.asp?EditorialsID=191> (25 August 2001).
- [7] – National Security Agency. “Operating Systems Security Research“. URL: <http://www.nsa.gov/isso/index.html> (16 August 2001).
- [8] – Harris, Patrick. “Macintosh Internet Security Basics“. September 2000. URL: http://www.sans.org/infosecFAQ/mac/mac_sec.htm (1 September 2001).
- [9] – SANS Institute. “Intrusion Detection FAQ v 1.52“. URL: http://www.sans.org/newlook/resources/IDFAQ/ID_FAQ.htm (1 September 2001).
- [10] – Gopalakrishna, Rajeev. “Audit Trails“. April 2000. URL: <http://www.cerias.purdue.edu/homes/rgk/at.html> (16 July 2001).

Software Used

Fluke Networks OneTouch Series II Reporter Software - <http://www.flukenetworks.com/uk/LAN/Handheld+Testers/One+Touch+Series+II/overview.htm>

LANScan Professional - <http://www.owensville.com/>

LC3 (formally L0pht Crack) - <http://www.atstake.com/lc3>

Crack - <http://www.users.dircon.co.uk/~crypto/> or <ftp://ftp.cerias.purdue.edu/pub/tools/unix/pwdutils/crack/>

Iris - <http://www.eeye.com/html/Products/Iris/index.html>

Dsniff - <http://www.monkey.org/~dugsong/dsniff/>

Ettercap - <http://ettercap.sourceforge.net/>

NMAP - <http://www.insecure.org/nmap/>

Nessus – <http://www.nessus.org/>

Cerberus' Internet Scanner - <http://www.cerberus-infosec.co.uk/cis.shtml>

SANS Security Digests - <http://www.sans.org/newlook/digests/>

Bugtraq - <http://www.securityfocus.com/forums/bugtraq/intro.html>