Network Security at the Institute Level

Christian Schwingenschlögl and Andreas Pilz

Institute of Communication Networks, TU München

1 Motivation

One of the worst failures in an institutes network is the (often complete) ignorance of security issues. In our experience, this is mainly due to the following reasons:

- Nobody is really responsible for security management.
- A wrong sense for security, i.e. someone did a security check years ago and there was no related activity afterwards.
- Securing the Network is often considered a "non-productive" task. Unfortunately, this attitude is not likely to be changed before a major security incident occurs. But as everyone knows - the alarm systems installed directly after a burglary are the most effective.
- The institutes network is considered as "not interesting" for potential attackers.

The Sans institute also provides a more general list (1) about the most common mistakes people make that lead to security breaches. To consider the last point given above - in our experience there are a lot of reasons for potential attackers to focus on an institutes network. In the later chapters we talk about different kinds of sensitive data - one possible motivation for external persons to get access to the network. As the machines of an institute usually also have a very good connection to the internet they are a optimal platform for attacks on other targets. In this case the institute and the university can get considerable legal troubles. Imagine your institute as the last link in an attack on a big company - if you are very lucky you can find the traces on your machines which prove your innocence. But usually the luck is ending at this point.

2 Inventory

In this section we will talk about the problems a "typical" institute is facing when considering security.

As stated above - the most effective alarm system is the one installed directly after the burglary. So we start with the assumption that some of the institutes machines have been cracked.

2.1 Possibly Compromised Systems

Usually, if your systems have been cracked, the intruder installs so called "root toolkits" on your system. Some system-binaries are replaced by the toolkit and you only see what the intruder wants you to see. The problem here: If your system has been cracked and the

intruder has root-permission (only a small step from a cracked account) you will not detect the crack via your logfiles.

The only way to be sure your system has not been compromised is the usage of tools like tripwire (2). This tool is basically a file integrity checker. It passes over the file system and generates a cryptographic hash for each file. It should be run at installation time of your system, the produced hashes should also be stored on removable media and locked away. If you suspect the system to be compromised, integrity checkers like tripwire (or similar tools like e.g. L5 or SPI) are the only chance to prevent your harddisk from being nuked. If there are no checksums from installation time the only way is to format and re-install all suspicious machines. Even if they have not been compromised - without the checksums there is no way to detect it for sure.

2.2 No emergency plan

In the typical institute, there is usually one "guru" who does the system administration (very often in addition to his normal work). If we are in a technical institute, the chances are high that some members of the technical staff also know some basics about system administration. However, even in this "best case" the number of security experts at an institute is usually < 1. To have an emergency plan in this "best case" is a very good thing. Lets define the "worst case" as the guru on vacation and the rest of the staff without any considerable technical background. Without any emergency plan you hopefully have up-to-date backups stored in some really secure place.

But what should be written in your emergency-plan? In the "worst case" defined above, the emergency plan should usually consist of some very high-level descriptions. If the staff can detect an attack in progress they can at least disconnect important server-machines from the network and prevent them from serious damage. In all other cases the emergency-plan can be made more and more sophisticated. Important things to know are e.g. the root password (if the administrator is out of range), where the connections to the internet are located, the important server machines and the location of the log-files. It is also important to know the location of the backups and the mechanism to install them.

Besides a lot of valuable information about firewalls, (3) also gives the two probably most important rules for responding to security incidents:

- Rule1: Don’t Panic!
- Rule2: Document!

Having an emergency plan, at least the risk of having destroyed valuable data by your own staff is greatly reduced. If everything goes well, a coordinated approach can help a lot to gather valuable informations about the attack - a necessary base for any countermeasures.

2.3 Underestimation of risks

At the institute level, we have often seen a huge underestimation of risks. This is mainly due to the facts we have pointed out in the previous sections.

Our experience shows the institutes networks not directly at the electronic front line (our assumption is that in this case most institutes won't be the owner of their networks any more), however, a look in the log-files usually shows considerable interest from external people.

2.4 No general security policy

Let's have a look on the environment we are typically facing in an institute. Usually we have to deal with

- often changing users (e.g. students in lab-courses)
- partially old machines with old software (i.e. the old (often self-made) software needs the old OS
- students doing system administration

If the institute has some technical background, also the following points are worth mentioning:

- users with widely varying technical backgrounds
- users are often installing software themselves
- very heterogeneous platforms
- often "experimental" software is used
- a highly dynamic network structure (i.e. often the operating system and security relevant software is changed, installed or removed without central organisation).

All these points underline the need for a well defined security-policy. This policy shold contain the necessary organisational points to keep the institutes network under control, especially things like:

- current network architecture
- which machines are in the network?
- which operating system and which version is running on which machines?
- who is responsible to install the various patches becoming available for the different platforms?
- common policy for students (e.g. do they get local root access to machines of the institute? Removal of old accounts, initial passwords, password-handling in lab-courses, ...)

2.5 Network structure - Security

Usually the network of research institutes is very heterogeneous:

- A lof of different computers and operation systems are used.
- For research reasons there may be a lot of services necessary (audio and video conferencing).
- It may be necessary, that users of special services need more privileges, e.g. root-privilegs.

– At a institute there is a very high fluctuation of staff, e.g. students working on their diploma thesis. Some of them need root-accounts on their local machines. Therefore they are able to sniff passwords or mount any other home directory, if the noroot-flag of the home directory stored on the NFS server is not set.

Therefore providing security in a institute network is a very complex task.

2.6 Passwords

A lot of security mechanisms are based on passwords, e.g. local login or login via telnet, ssh, ftp. If someone is able to steal or sniff your password, he or she can pass these security mechanisms and is able to abuse your account.

Therefore no default-passwords, passwords everyone knows or weak passwords should be used. Accounts without any password are completly forbidden.

As we mentioned above, it may be necessary, that students need root-passwords. It is recommended to change these root-passwords after the student has finished his or her diploma thesis.

Moreover it might be useful, to create a General Password Policy for the network, in order to provide a uniform password handling.

2.7 Sensitive Data

There are usually a lot of sensitive data stored within a network of a institute, e.g. e-mail of the professor or of the collaborates, unpublished results of current research work, password files, ...

As especially student-hackers are often interested in very sensitive files containing examination papers or examination results, these files should be created, edited or stored on computers, which are not connected to any network. It is recommended to even encrypt these files after editing (e.g. PGP). Moreover examination papers and results should not be printed on a network printer, because the print job itself can be easily copied or redirected to another printer.

2.8 WLAN

Not too long time ago, the wireless LAN was a somewhat exotic hardware. Today, we do not yet find it everywhere, but it is becoming standard equipment in a lot of institutes. The WLAN equipment is easy to install, easy to use and - if security is not considered - does widely open the electronic equivalent of your institutes front door. Two popular examples of security holes created by WLAN-systems (or, better, the complete ignorance of the related security-questions) can be found in (4) and (5). Article (4) describes the hack of WLAN-equipped roboters at the EXPO 2000 in Hannover, article (5) describes a incidentially discovered security-hole in an universities network.

The main problems with WLAN are

- the potential intruder does not need to get access to your rooms
- in contrast to the wired ethernet, eavesdropping is possible without expensive hardware
- anyone within the range of your WLAN can access your network (if DHCP is enabled, the intruder is superuser on a machine within your network)

Most of these troubles are easy to circumwent: To deal with the eavesdropping problem, the WLAN equipment has to support link encryption (sounds trivial but is often overlooked: the link encryption has to be activated). To circumwent the third problem stated above, it is recommended to register the MAC addresses of users who want to use your DHCP service. So DHCP can be restricted to known users.

2.9 Intrusion Detection

One possible motivation to break in an institutes network is eavesdropping. As in this case no malfunctions of the systems are expected, without effective intrusion detection the intruder will stay unnoticed for a long time. One present-day example of this is the microsoft break-in. According to (6), the intruders stayed unnoticed for about 3 months.

A common situation in institutes is the existence of various different log-files nobody really cares about. In this case, the intruder will stay unnoticed for very long times and only be discovered incidentially. A big step towards a systematic intrusion detection is the central collection and a periodic inspection of all the logfiles. As one of the first things expected from an intruder is the alteration of the system logs to blur or remove his traces, an efficient and secure mechanism for the storage and backup of the logfiles is essential.

Moreover, there are intrusion detection systems available for all common platforms from numerous sources. In (7), Stephen Northcutt gives an overview about some intrusion detection solutions. This overview gives valuable information about the basic functionality of different itrusion detection solutions. However, the product cycle in intrusion detection is very fast, so it is recommended to get an up-to-date overview before making a buying decision.

It is important not to understand an intrusion detection system as a fully automated watchdog for your network. In (7), a false positive fate of about 90 percent is given for most of the intrusion detection solutions. So it currently provides some help to the administrator resp. security analyst, not more and not less.

2.10 Webserver

Regarding the excessive access from the outside, the webserver is a very special machine in the institutes network. Due to his exposed position and the usually numerous functions it provides, the webserver should be somehow separated from the institutes core network. This can be achieved using e.g. a demilitarized zone (3).

2.11 Different Subnets

Often different subnets are used in an institute. Here the main problem is with the configuration of a firewall. On the one hand, access from unauthorized people should be made as hard as possible. On the other hand, it is not desirable to administrate all the necessary services in each subnet. To prevent the firewall from being looped, the implementation of a secure, encrypted tunnel can be used between the two subnets to create a VPN. There are various implementations available for this purpose, from the commercial side as well as free software, e.g. freeswan (8). We are planning to implement a secure tunnel in the next weeks and also some related performance evaluation.

2.12 Special Case: Network Research

If your institute has something to do with computer science or network research you will discover very soon that this is a special case. Without control, such institutes tend to behave like vespiaries, unfortunately much less well-fortified. As a lot of people are (more or less) familiar with system administration in such institutes, the risk is high that anyone installs anything on any platform without prior notice. To prevent a total disaster in terms of security, either the network access from the outside has to be very restricted (usually not very popular in this environment) or a good security policy has to be in place. It is absolutely necessary to know who is allowed to do which things. Also a common policy regarding e.g. diploma students, which have access to the institutes machines, can only be enforced this way.

2.13 Multimedia Services

As teleteaching and videoconferencing is emerging, more and more "exotic" hardware can be found in the institute. As of the variety of available hardware, we can not go into much detail on this point, however, we recommend the "black box" approach to detect possible unsecured entry-points in the institutes network: if the hardware is possible to be connected to the telephone network and it also has a connection to the local network it is worth to have a detailed look on it.

3 Solution - Architecture

3.1 Network Architecture

A lot of security problems can be solved by a well thought out network architecture.

Firewall (packet filter)

A firewall is not just a packet filter. It is a device or mechanism, which is supposed to keep the bad guys out of your network. A good firewall succeeds in keeping the bad guys out, while still letting you cleanly use your network. If it doesn't, it's a bad firewall.

An example for a firewall is a mechanism for filtering IP packets according to a filtering policy. Therefore the traffic from outside into the local network and vice versa can be restricted. Your local network can be split into two zones, too:

– demilitarized zone DMZ: The services provided located in this zone are available for anyone. These services are for example the web- and ftp-server.
– local internal network: The computers and network components of the internal network are “located” in this zone. The firewall restricts access to these components according to the filter policy.

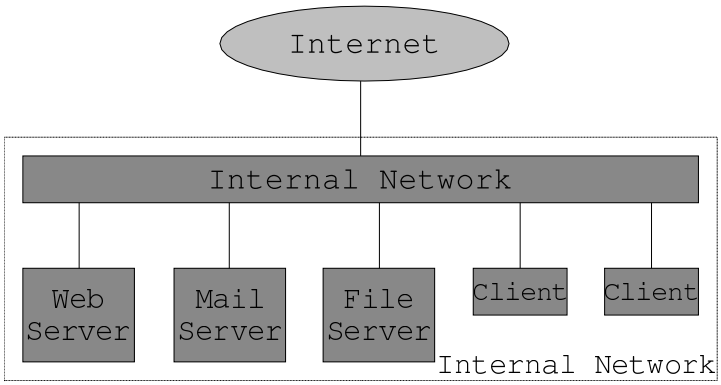

Figure 1: Network without Firewall

Masquerading

IP masquerading is a mechanism, which allows you to hide a subnet and its components of your network from the internet. The clients of the masqueraded network use IP numbers, which are designed as internal netwok numbers (e.g. 10.0.0.x or 192.168.0.x). The masquerading server restamps all outgoing packets with its own real IP address and an alternate port, and restamps all incoming packets with the masqueraded addresses and sends them over the masqueraded network to the client. As a result of this, the masqueraded IP addresses can not be addressed from anywhere except the masqueraded network itself. Therefore the masquerading mechanism is a very effective firewall.

Tunneling between the subnets

If your network consists of two separate sub-networks, than there are a lot of additional security problems. The following diagramm show such a network, where the sub-networks are located in different buildings.

The sub-networks communicate with each other using “foreign” networks. These transport networks may be malicious and observe your traffic. If you send data in plain text, than anyone connected to the transport network is able to read you data. For example, if you want to mount a directory stored on a server which is not located in your sub-network,

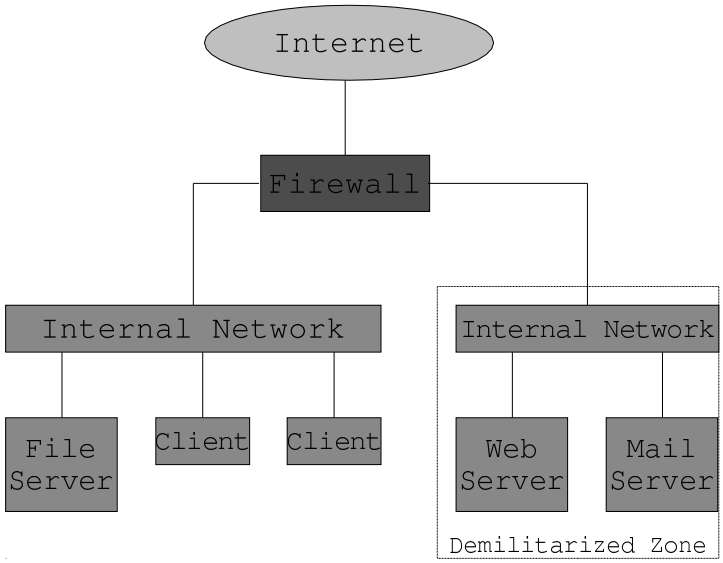

Figure 2: Network with Firewall

than your request is sent over the transport networks. As these request may contain your password in plain text, any attacker is able to extract your password.

A possible solution for this dilema is using tunneling mechanisms between subnets like IPSec. The hole traffice between the two subnets is encrypted between the security-gateways located in each subnet. Therefore it is hardly possible to sniff passwords or read confidential data, which is transporte between the two subnets. The following diagram shows an example-VPN using IPSec.

But as the encryption process is very expensive, the network performance may be reduced.

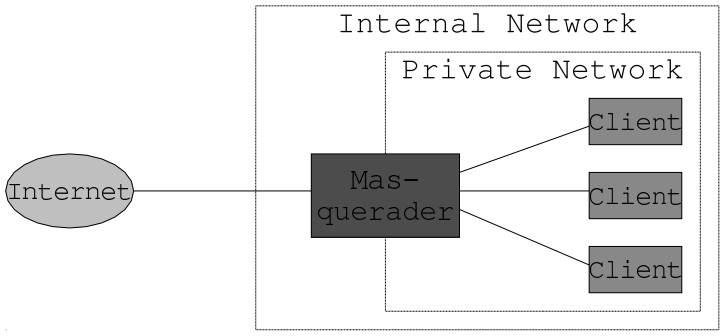

Figure 3: Masquerading

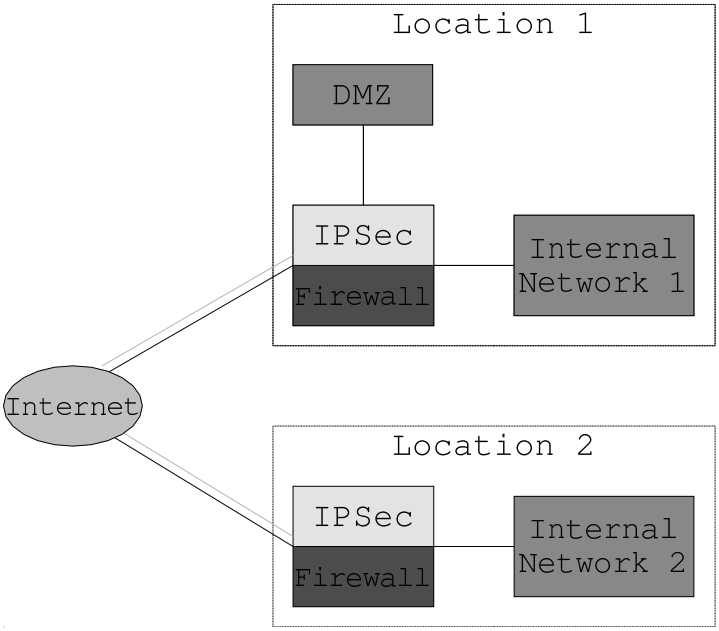

Figure 4: VPN using IPSec

4 Harden the network

In their paper (9), Dan Farmer and Wietse Venema explain an approach which is not only exciting, it also helps a lot to understand the own system and eventual strategies of the attacker. They propose to improve the security of a site by breaking into it.

4.1 Periodic System Scans

Not only that a lot of things about system administration will be learned by the cracking-attempts on the own network, it is also the approach that usually discovers the worst security holes first. We recommend two approaches here:

- Use Scanners from the Internet
- Break into your site "manually"

The first approach has the advantage that the network can be scanned really fast. Usually, a list which can be ordered by the severity of the security hole can be generated automatically. It also has the advantage that a lot of attacks work this way (at least the site is made skript-kiddy secure) - so the most common entry points in the network are secured. Popular and free available security scanners include e.g. (10), (11) and (12).

The second approach can be used to get more familiar to the own network and to learn more sophisticated techniques. As security scanners are available in new versions very often, we recommend to use them frequently. If time is left, it can be spent for manual break-in attempts.

4.2 Periodic Installation of Patches

An easy way to break into a system is the use of often well known exploits. These exploits use the unability of some programs to handle some special input or certain failure conditions. To secure your system in this direction, it is essential to

- have only really needed services running on a machine
- frequently install all relevant patches

Some security scanners describe the detected problems in detail, often with links to recommended patches to solve these problems.

4.3 Periodic Password-Crack

As in our scenario also the LAN is not fully trusted (this, however, should be the assumption for any LAN), we have to deal with the risk of "password sniffing". The encrypted passwords can be intercepted at the LAN and attacked via the well known dictionary-attack. As the original password can not be derived from the encrypted version, it is very common to use brute force for this purpose: arbitrary words are encrypted and the encrypted words are compared with the encrypted password. To protect the network against this kind of attack, it is necessary for the users to have strong passwords. One tool that can check the strength of the different user passwords is e.g. john (13). It can even be configured to check the passwords automatically and send a message to the users which don't use strong passwords.

4.4 Periodic Logfile-Backup

Usually, if the intruder has the possibility to do this, the logfiles are destroyed or modified in an attack. Therefore, the periodic backup and also the safe storage of the logfiles is essential to detect an intruder and, in case, to have information about what happened.

4.5 Periodic Backup-Check

It is very tranquillizing to have periodic backups of all systems. But are you really sure everything is secured? How can you get it back? Do you need installation disks you don't have? If your systems are hacked friday night - will the backups be overwritten by the modified data?

4.6 Emergency Plan

It is necessary to create an emergency plan before being attacked, in order to be able to react to a possible hacker alarm.

There are several strategies possible:

- A very effective strategy is to immediately and completly disconnect your entire network from the outside. All damages can be minimized using this strategy. Another advantage of this procedure is, that even persons, who are less familiar with security mechanisms are able to pull the plug.

– If you want to collect additional log-files and data on the activites of your attacker, than you can redirect his request to a so-called honey-pot and disconnect him from your real network. A honey-pot is a special computer in a network with no really important information. It is used for for distracting the attacker from the real network. While the attacker is busy with the honeypot, you can easily collect further log-files and data on the attacker.

It is recommended to create an emergency plan covering several strategies, because the plan must provide procedures for collaborates partly having less experience and for very experienced users, who are able to trace back the attacker. Moreover the emergency plan should predetermine the importance of your network components. It is recommended to protect or minimize harm to your file servers in the first place. After that the single machines of each collaborate should be secured and so on. The exact order depends on your security policy and on the importance of the information stored on the machine for your productive work.

4.7 Recovery if Cracked

If your network has been cracked or if you think, that it might have been cracked, than you have to examine all networks components. If a component has been cracked, than the best thing to do is to completly begin from the scratch. The recovery procedure should be done by a trusted and reliable person. It is recommended, that the computer is finally reconnected to the network after the recovery procedure has been finished.

First of all backup all personal data, but keep in mind, that this data might have been manipulated.

Than the operating system has to be completly reinstalled after really formatting the media (HD, FD). It is recommended to install the latest patches, too.

Reconfigure the security parameters for this machine according to the security policy.

It is recommended to change all passwords, which might have been sniffed or stolen.

If you have checked and eventually reinstalled all your components, than your network has completly recoverd from the latest attack. But keep in mind: If there is only one single component still infected, than your complete network is still vulnerable.

5 Conclusion

This paper is not intended to be a complete security guide for your network. It should be a motivation and starting point for securing your network. It does also provide some ideas about possible security holes and its removal, especially for institute networks. Given the fact that for research purposes a open network architecture with little restrictions is needed, the completely secure network will stay an illusion. However, it will also stay an illusion for highly secure company networks as shown in (6).

The important thing is to find the right cost-benefit relation and to be aware of existing risks to find the appropriate level of security.

References

[1] SANS Institute: Mistake people make that lead to security breaches.
[2] Gene Kim and Gene Spafford: Tripwire.
[3] Elizabeth D. Zwicky and Simon Cooper and D. Brent Chapman: Building Internet Firewalls.
[4] Matthias Mehldau (wetter@ccc.de): WaveLAN auf der EXPO.
[5] Peter Siering: Vorsicht Nachbar!. c't. 2000, 22,
[6] BBC: Microsoft software "stolen".
[7] Stephen Northcutt: Network Intrusion Detection - An Analyst's Handbook. 1999,
[8] freeswan.
[9] Dan Farmer and Wietse Venema: Improving the security of your site by breaking into it.
[10] Worldwide Digital Security Inc.: Saint.
[11] Dan Farmer and Wietse Venema: Satan.
[12] Insecure.org: Nmap - stealth port scanner for network security auditing, general internet exploration & hacking..
[13] solar@false.com:John the ripper