



SANS Institute

Information Security Reading Room

Protecting Your Internal Systems from a Compromised Host

Michael Nancarrow

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.

Protecting your Internal Systems from a Compromised Host

Written by Michael Nancarrow
GIAC ID michael220

Based on GSEC Assignment version 1.3, December 2001

The following Registered Trademarks and Companies appear in this Document:

Redhat
CIS
SAINT
Nmap
SARA
SATAN
Tripwire
Solaris
HP-UX
CISCO
IdeaData
DigiToll
Microsoft
SANS
CERT
AusCERT
HP-UX

© SANS Institute 2002, Author retains full rights.

Protecting your Internal Systems from a Compromised Host

Introduction

When the horse has bolted.

The concept for this paper came from a recent incident when one of our customer machines was compromised. It is designed to cover some additional aspects of systems security and design, which I believe have been ignored to some extent in the Security Essentials material and most systems admin courses.

At some stage you must concede that a system will be compromised and as such being located in a trusted or semi-trusted position on the network an effort must be made to minimise the impact and also identify the problem as soon as possible. The content of this paper has been kept brief and covered areas that have not really been emphasised enough and fall victim to lazy system management. A number of other areas like systems resource monitoring, systems file access and protection, and user management issues are generally well covered in standard system administration courses and guides and are not mentioned in the paper.

The paper is structured to provide overview of some areas to address, brief recommendations, some hints from experience and some reference sites to check out. There are 3 distinct areas the first being the system, the second dangerous information on the system and the third monitoring of the systems.

Do you stand a chance ?

The attackers knowledge of your system/s and skill-set can vary wildly. The knowledge of your systems may be as little as the ip address and the method of hacking up to complete technical details obtained from months of analysis or insider information (eg disgruntled employee, job trainee, short/long term consultant, etc). The simplicity of running tools like nmap, SAINT, SARA, SATAN, to name a few has lessened the skill set required for a hacker such that the majority, commonly known as script-kiddies, are dependent on them. Use of tools like hfnetchk (from Microsoft), CIS and Typhon can also provide the attacker information as they can analyse the system itself to provide information on other system weaknesses that provide a better entry point. If the network manager uses these tools as well, instead of relying on his variable skill level then he will be able to raise the level of skill and effort required for hacking his systems. As in the Defence in Depth model of security you do not put your reliance in one piece of software or hardware but use all tools available to increase the level of skill required. Hopefully this will be enough.

Steps to take on the systems.

An attacker having established a beachhead will try to expand his influence over the machine and also try and prevent himself getting kicked off by the owner of the machine or possibly even another hacker. Depending on the method of entry he will

probably try to import tools onto the machine to provide another mechanism of entering the machine or another more secretive backdoor, that only he may use. The attacker may bring tools able to escalate his level of privilege to allow greater influence on the system and probably bring tools that will allow exploration of the surrounding network. All these can be guarded against to some level, some by simple monitoring of the systems and others by knowing your system and what its requirements really are.

Guard against back-grading of software.

On Microsoft and Unix platforms various software vulnerabilities have been discovered over many years and are still being found. Each one of these has resulted in patched files within the operating systems.

Microsoft, though a slow starter, has provided a number of mechanisms to allow the user to know his current patch level, hfnetchk and Windows Update are some. The nice part of these tools is that they tell you what needs to be upgraded. This not only gives you ideas of what is required, but it gives you a visual reference as to what software is currently not patched on your system. If a hacker where to remove a patch the regular running of these programs may provide a clue, apart from being good housekeeping. The System File Checker tool, a Microsoft program which checks the legitimacy of files, is available for use on platforms including 2000 and XP for verification of versions. It is not a virus scanner but may be useful for a quick automated check of the system (Microsoft articles Q310747 and Q222471).

Unix variants unfortunately are not as helpful regarding the maintenance of patches. Each vendor has different mechanisms for informing users of patch availability and some vendors like Redhat offer subscription services to provide updates to their versions. The most commonly used are notices through Bugtraq, CERT, newsgroups, email distribution lists from the vendor or the sudden arrival of a pack of cd's on your desk. Careful record keeping of the RPM's installed (linux) , the PHN's from HP-UX and the equivalent from other Unices is highly recommended with auditing checks performed at regular intervals.

Programs that store a record of the characteristics of a systems files, including size, modification dates, check-sums of the contents are required to guarantee no modifications have occurred if the hacker is skillful. Tripwire and similar programs are possibly the only means of understanding if a patch has been removed or modified as patch histories can be modified. After each upgrade/patch new reports will need to be generated with new keys for files. (ref. www.tripwire.com)

System Hardening

System "hardening" is the reduction of security holes within the basic operating system, running applications and auditing of the system such that the manager understands what potential security holes exist.

System hardening tools and manuals can have two purposes. The first is to harden the system to prevent attack. The second to provide a record of the changes you have made to the system for later reference. If you perform an upgrade then system

hardening will need to be performed again to ensure you have a secure platform. The hardening process should also be revisited at regular intervals to ensure that no changes have been made that compromise the system either by internal staff, unwittingly, or by an attacker who has had access to the system and has opened security holes for use. Dependent on the operating system tools are available that can be automated. If you receive a log of the output, this should be retained securely as a handy reference to the security compliance of the system after the hardening process.

The Cerberus Internet Scanner and Typhon, a later version of CIS, provide good reports in HTML format for Microsoft systems. The changing of some of the parameters related to security in the registry or the system auditing will appear on the reports and you can keep the reports as a baseline as to how this system compares to when you set it up. These tools also provide remote scanning of other systems, provide suggestions on what to do and Typhon has a port scanner built in.

On the Unix platform the Bastille (<http://www.bastille-linux.org/>) program has been the only hardening program I have been able to find that is attempting to not be vendor specific. Some vendors have provided either guides or specific tool kits for their product. The JASS toolkit from Sun is an example of this (<http://www.sun.com/security/jass/>).

Manuals and guides on system hardening are available from:

- www.sans.org
- www.securityfocus.com
- www.hackerwhackers.com
- www.linux-sec.net
- www.microsoft.com/security/
- www.uscert.org
- www.cert.org

all operating system vendor sites and numerous security sites located on the Internet.

Protecting the other systems from the compromised host

An attacker will try to use the resources on the system to discover as much as possible about the network structure before embarking on any discovery. Restricting the information in a system to as little as possible about the rest of the network is advisable, particularly on machines related to Internet tasks, like web-servers and other DMZ systems. This will not stop an attacker finding information but it can reduce the quality of information and require the importing of additional tools that could tip off the manager of something being wrong. Some bad habits are:

- Copying hosts files and not tailoring them to the needs of the system
- Copying Lmhosts files and not tailoring
- Maintaining system logging files on the systems and for extended periods
- Copying nfs export files to new systems
- Allowing the systems to do DNS zone transfers (`ls -d <domain-name>`)
- Making the system a DNS server unnecessarily

- Allowing netbios over TCP/IP when not required (use of nbtstat command)
- Failing to disable the IPC\$, C\$ and admin\$ shares if not required
- Copying or using hosts.equiv files
- Maintaining copies of wtmp and utmp files on systems and for long periods
- Not deleting unused or old employee accounts

All these mechanisms can provide access to IP addresses, system names and usernames resident on the network, who have accessed the system, and therefore are live hosts to target. In the case of the syslog file it can provide a list of IP addresses and the mechanism used to connect to the system so that a structured attack can be made that will not necessarily arouse suspicion eg: extract from a syslog file (names changed to protect the innocent)

```
Mar 22 13:30:08 6C:nautilus telnetd[255556]: connect from nemo.city.com
Mar 22 13:40:09 6C:nautilus telnetd[255683]: connect from kirk.city.com
Mar 22 14:20:08 6E:nautilus login:[256764 ]: ?@douglas.city.com as miken
```

Gives the system name, the type of connection made and even a user login. The attacker can now restrict his activities to a very limited window. The WTMP and/or UTMP files will give the same information interactively to the "who" or "last" command from information in these files (extract from the "who -m" command).

```
miken ttyq0    kirk.city.com. Fri Mar 22 14:17  still logged i
taylorj ttyq3  douglas.city.com. Fri Mar 22 13:27 - 14:43 (01:16)
admin  ttyq0   nemo.city.com. Fri Mar 22 13:07 - 13:39 (00:31)
hackera ttyq0  douglas.un der.com.au  Fri Mar 22 12:53 - 12:54 (00:00)
note: ( I have checked this on Linux,(RedHat), AIX, IRIX, HPUX and Solaris)
```

The worst part about the above is that I did not require root to get this information.

Microsoft systems can even provide information remotely due to the IPC\$ feature. The following outputs are from the Cerberus Internet Scanner against the localhost (127.0.0.1) the first with IPC\$ share available the second with the share disabled.

IPC\$

Account Name :jmichael

The jmichael account is a normal USER, and the password was changed 38 days ago. This account has been used 0 times to logon. Comment : User Comment : Full name :John Michael

Account Name :manager

The manager account is an ADMINISTRATOR, and the password was changed 43 days ago. This account has been used 2 times to logon. This account is the renamed original default Administrator account. Comment Chief Information Officer :Built-in account for administering the computer/domain User Comment : Full name :

NO IPC\$

No NetBIOS Session Service

This information provided gives the username, a measure of how often the account is used, personal name details and group membership. This is all great information for somebody that doesn't know your real name, position and wants to get in. Both Unix and Microsoft provide easily accessible information on when an account was last

02/06/03 Page 5

used. A great place for a hacker to start on cracking a password without anybody complaining.

Steps to take on the Network

The monitoring of network resources and utilisation are paramount in the protection of your other systems once one system is compromised. A number of tools are available to look at network traffic and to analyse what is happening, unfortunately the sheer speed of some networks means that these systems can easily get overloaded and care must be taken in their placement and maintenance.

For ease of explanation I have broken this area up into two sections which I describe as passive and pro-active monitoring. The definition of “passive” monitoring used is the non-generation of traffic to monitored systems, “pro-active” is defined as the active generation of traffic to monitored systems.

Both Passive and pro-active means are available to collect data in a number of key areas, which can allow the network manager to identify an abuse of network and systems resources. Data that should be collected on network traffic and compared at regular intervals includes:

- Aggregated current network load
- Typical network load patterns
- Network traffic load per system
- Acceptable performance levels (latency, eg ping timeouts and late collisions)
- Traffic profiles of protocols
- Top talking and listening systems (top network users)
- Top conversations (not necessarily the same as talking and listening)
- Malformed traffic

Passive monitoring of traffic

Like the monitoring of an electricity network or water supply, knowledge of what is happening is vital and sensors provide critical information on without actively contributing to a reduction of resources.

All the above parameters can be monitored to provide baseline statistics and noticeable discrepancies or additions can be a sign of system compromise and attacker probing. Generally this is known as “sniffing” the network. This type of network analysis requires the positioning of specialised devices in critical positions where traffic congregates into single pipes or where a specialised feature of a network devices, eg port/VLAN spanning or mirroring, allows the aggregation of traffic to feed into a single analysis interface. Generally two types of probes are used though both essentially are doing the same job which is to capture packets from the network and analyse them.

The first are Intrusion Detection Sensors or IDS's. These devices are becoming quite popular after September 11 and the hysteria that management now feels about a threat that already existed. These devices are really just packet sniffing engines that direct the traffic through advanced filters to isolate packets potentially capable of

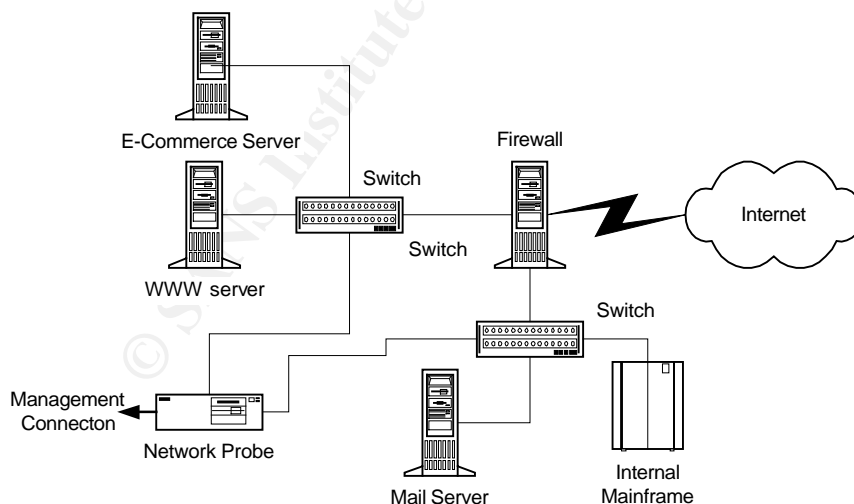
corrupting a listening executable on a system for malicious purposes, eg a system crash or execute code of choice by the attacker.

The second are general LAN Probes, which are essentially the same devices as above but the back-end software filters are designed to analyse the traffic for protocol utilisation, system conversation, network traffic load patterns and some times for packet analysis. These are commonly used for network baselines and capacity planning with the occasional problem solving. These devices are becoming more important in that they now have the speed and capacity to analyse traffic and also provide alarming mechanisms based on traffic threshold through the use of back-end databases which they use to store either the packets structure or a summarised version of the traffic.

The Diagram below shows the possible position of the “sniffing” device relative to the traffic flows through the firewall. The sniffer, being a passive device, would have the capability to not be network active, which has some serious advantages in regards to security of the network and the data it collects. The device, since it has no need to transmit traffic of it’s own, can be isolated from the network and therefore immune to any attack on the network or potential compromise by attackers. Depending on how many interfaces the device has you could increase security by:

- 1/ **Not** giving the device an IP address (as the Interface is only listening, works with SNORT, TCPDUMP and any Linux based tool I have tried so far).
- 2/ Place the device in IP network not routed to the main network.
- 3/ If possible cut the transmit wires in the connecting cable (device can not respond).
- 4/ Turn off Broadcast and Multicasts on the connecting switch port (ie does not respond to ARP)

Diagram 1 Passive System Monitor position



Depending on the capability of the switches it would be possible to mirror all traffic destined to and arriving from the Internet both for the DMZ and the internal network to the monitoring ports on the network probe.

An interesting point to this is the current development of some products that store a summary of the packet information that flows through collector devices and feeds this in to back-end databases like MySQL and Oracle. This simplification of data structure whilst reducing the forensic value has major advantages in the speed of response of such a database and the possibility of setting alarm thresholds on the protocols that are exceptionally out of range for an individual system. One product which I have experimented with from IdeaData called DigiToll while essentially meant to be used as a Billing database for network utilisation is so effective that it allowed the alarming of thresholds for a change in the standard FTP Traffic over a 1 hour period. This allowed the identification of the WAREZ site within our customer DMZ and the limitation of damage to themselves and our network. Such alarms could be established for other probes such as unusual amounts ICMP or ARP traffic across a VLAN, unusual protocols from a stable system or even port scanning probes or abnormal port utilisation like 6667 trojans. Similar capabilities are available on other IDS and Network probe based systems like CISCO Secure Policy Manager and HPOpenview platforms though in lesser detail and at greater cost as they require specialised probes. Theoretically this capability should be available on the SNORT IDS system as well when the packets are logged within a MySQL database, however I have concerns about the amount of data stored from a high speed network and the complexity of the queries required.

An interesting device on the market is the Packetshaper from Packeteer. This Device actually sits on the path of the traffic to monitor the throughput and can actively identify new protocol streams and conversation flows.. Whilst not storing the packets the device simplifies the traffic flow into a database structure which enables active querying and allows identification of per protocol per system analysis. This device is actually targeted at edge of network and gateway interfaces. Whether it is worth the risk for another point of failure for simple network monitoring I have my doubts.

Ideally both the functions of an IDS and a Network Monitor should be able to be combined in the one platform, however the networks that usually require these functions are so active that it is probably a bit early in database technology to combine the functions. What is more interesting is the distributed nature of these devices and the possibility of running a combined function appliance or probe that feeds to separate back-end systems. This should not be difficult to implement as the Unix Kernel will apparently fork data streams of network traffic to two programs.

Passive mode monitoring, though it can be more expensive due to the need for dedicated equipment, offers a more secure and potentially extremely effective means of alarming on network changes.

Additional information can be found at the following websites:

www.ideadata.com.au

www.cisco.com/univercd/cc/td/doc/pcat/sqplmn.htm

www.snort.org

www.hp.com and www.agilent.com (for HP Openview and NetMetrix using remote LAN Probes)

www.packeteer.com

Pro-active network monitoring

Generates traffic on the network by either actively querying or testing the network devices, which are involved in either generating traffic or controlling the traffic on the network. This requires that the querying device have an IP address and be actively known on the network through ARP and other responses.

Similar to passive mode this breaks down into two areas. Cooperative querying of network devices and systems for information on throughput, current state of links and other information through the use of installed agents. The second being Un-Cooperative querying or active probing of systems to test for problems with the device and confirm the current status of open or closed services.

Cooperative Querying of Network Devices and Systems

This type of monitoring is the most common used due to it's cost effectiveness and the support given by the majority of devices and systems for the Simple Network Management Protocol (SNMP). A number of management software platforms are freely available for this type of monitoring at various degrees of sophistication and alarming capability. The SNMP protocol also allows the network device or system to send messages to network platforms if an alarm condition exists on the device (eg link down, unauthorised MAC address, system log full, etc).

The problem with this type of querying is that you are distributing the processing of frames and traffic to active devices in the network whose prime directive is the processing of data as fast as possible or forwarding the data on to target systems. Recently this has become less of an issue with the large increases in processing power and the use of more ASIC based technology to reduce the utilisation of the main processor.

Unfortunately you are at the mercy of the vendor in regards to the amount of information available and some devices are limited to the overall packet and octet throughput. Some vendors are becoming more sophisticated in this area as seen by CISCO's Netflow technology, which is gradually seeping down their network devices. Netflow provides application and protocol level information, which would allow greater capability in monitoring for excessive protocol utilisation, ICMP polls and other signatures of network compromise.

This type of monitoring does open up the possibility of compromise for both the system being queried and the querying system due to additional services being open. This has recently been of significance due to the poor security practices in the processes of some companies coding and poor implementation of the SNMP agents and collectors by vendors. The Computer Emergency Response Team released "CERT[®] Advisory CA-2002-03 Multiple Vulnerabilities in Many Implementations of the Simple Network Management Protocol (SNMP)" showing how badly implemented SNMP had been by the vendors and that the expediency of cross-compiling code from different platforms had not be done with thorough examination of the code base. Some Unix Vendors actually have SNMP running on their platforms when first installed with no community string set so that any rogue user could get into the

system if either the system manager is unaware of this or fails to close the hole, HP-UX and Sun Solaris definitely do this.


There are a number of precautions that can be taken in regards to ensuring minimal issues occur, which have dealt with to a large extent lately, but a quick list would be:

- Change the SNMP community String for Read and Write
- Make sure it is not easy to guess
- If possible make sure the agent has a fixed management IP address
- Don't let it through the firewall
- Turn it off unless required
- Ensure it is off after a fresh install, or at least configure to other than defaults.
- SNMP activity logging on the devices be recorded if feasible
- The highest version of SNMP be used were possible (currently version 3)
- Ensure you are running the latest version of firmware available for the device
- Ensure that you are running the latest version of SNMP management software
- If version 3 capable devices configure the security features particularly username/password

A large amount of information on this can be located on most vendor sites and of course the CERT advisory at www.cert.org/advisories/CA-2002-03.html

Diagram 2 shows a basic situation where only one management station is required to collect data from all systems it's position within the secure internal network will require opening SNMP holes in the firewall to allow complete monitoring, not an ideal situation

Diagram 2. SNMP Management Station



These tools can be extremely effective in the establishment of network baselines for comparison and can particularly highlight system issues over extended time periods. Their interaction with the various types system and devices provides immediate feedback on a number of easily readable graphical interfaces, dependent on the tool interface and are extremely easy to deploy.

Implementation of SNMP systems is recommended but a high level of caution will be required and it is strongly recommended that you check your vendor site for the latest information and upgrades. Use of the data that is collected could point out unexpected activity and most management stations will allow you to program some alarms on threshold events which can be forwarded to email, pager or SMS alerting systems.

Some Interesting links for software and information are:

www.mrtg.org (Freeware SNMP query tool and grapher)
www.cisco.com/go/netflow
www.hp.com and www.agilent.com (for HP Openview and NetMetrix)
www.compaq.com/support/files/networking/software/Compaq_Network_Management_Software_V226.html (Free SNMP management software)

Un-Cooperative Querying of Network Devices and Systems

This is a nice way of putting trying to crack the systems. Most network managers who have played with security will run a security scanner once and then say all is fine and leave it alone for a few months. The problem is that he will probably run it during business hours, look at the results on the screen fail to take notes or save the output and report to management that the systems are secure. The establishment of a security baseline of the systems is just as important as the network traffic baseline for the detection of issues and holes. The same as System Hardening is required after every upgrade, any changes to devices which are network connected should be audited from a remote perspective the same as if you were the attacker and wanted to crack the system.

There are a number of well known tools that can be scheduled to perform this testing from the command line on Unix and Microsoft platforms. These tools should be run at regular but differing time intervals to ensure that:

- Users don't understand when it runs and turn off services during it's running
- All systems are eventually scanned, even those turned off sometimes
- A baseline is established as to what services are available during different times of the day

Care must be taken in regards to the running of these scans particularly against older operating systems, which have been known to crash due to the older applications having resource holes.

The tools themselves have differing outputs that vary from simple command line response, like nmap, to HTML pages and text file generation like the Security Auditor's Research Assistant (SARA) tool. The Unix tools are just as effective against Windows platforms and some have been ported to windows, nmapnt for example. All the tools I have examined have a simplified syntax, which will also allow for the testing of complete subnets as can be seen from the sample provided.

The following is a cut down version of a nmapnt scan performed against a Windows 95 system on my home network.

```
Command line > Nmapnt -sS -v -O 10.1.1.0/29
Host me (10.1.1.2) appears to be up ... good.
Initiating SYN half-open stealth scan against me (10.1.1.2)
Adding TCP port 5000 (state open).
Adding TCP port 139 (state open).
Adding TCP port 1025 (state open).
The SYN scan took 1 second to scan 1523 ports.
For OSScan assuming that port 139 is open and port 1 is closed and neither are firewalled
Interesting ports on me (10.1.1.2):
(The 1520 ports scanned but not shown below are in state: closed)
Port      State  Service
139/tcp   open   netbios-ssn
1025/tcp  open   listen
5000/tcp  open   fics
TCP Sequence Prediction: Class=trivial time dependency
Difficult y=2 (Trivial joke)
Remote operating system guess: Windows NT4 / Win95 / Win98
```

The sample is here just to illustrate the simplicity of running the tools freely available and the amount of information that they can discern. The nmap program has identified which ports are open from a basic scan and made an accurate guess as to the base operating system. It required little preliminary work including the installation of winpcap on the system and unzipping of the nmapnt zip file.

Automation of this would be extremely easy the difficulty being the comparison of results from previous scans. Nmap has a number of options, which allow the results to be reformatted and other programs are available which provide a friendlier environment for some users. Due to the need to test remote locals and because of isolation by Firewalls a client server application is available as well called nessus. This is more advanced than nmap and would allow a centralised client platform to manage and generate specific segment reports from sensors installed on devices in the LANs.

In the event of a compromised system being used on your network the reports from these tools will provide a forensic time-line from which you may be able to:

- Discover where the initial infection occurred
- Provide an external viewpoint of the system compromise
- Show the level of security on the system was acceptable (maybe job saving)
- Prove the level of risk for each system
- Establish a clear audit trail for later analysis
- Hopefully catch the infection as well if a systems security level alters

The use of network scanning as an occasional action item before an auditor comes can not be acceptable practice. A single system compromise potentially can cripple a company financially and its reputation.

The following links are for “free” software tools that have been mentioned above, I seriously recommend that you give them a try.

www-arc.com/sara/
www.nmap.org
www.eeye.com/html/Research/Tools/nmapnt.html
www.wwdsi.com/saint/
www.nessus.org
<http://www.cisecurity.org/> (site for router security tool, cisco)

Are my systems really behaving?

The Internet community is becoming less tolerant to the interference of hackers, spammers, script-kiddies and other unethical users of the Internet. A quick scan of e-mail lists, show a more militant attitude and aggressive frame of mind by some people to take a more pro-active role in Internet security. This has resulted in some benefits as most of these people are not the type to grab a rifle or send an email bomb or virus and they have turned their efforts to more technically orientated solutions. Some solutions that have appeared include:

Live remote probing of your computer or firewall

A number of sites will now scan your firewall systems or computer free remotely. There are some arguments about whether this is a good or bad idea to provide this so easily, but since they are already there, better to use them than not.

www.hackerwhacker.com
scan.sygatetech.com
www.dslreports.com/tools
www.grc.com (shields up)
<http://www.linux-sec.net/> (one minute Audit, left hand side screen)

I have used all these sites to test my systems on occasions just to feel a bit more secure both at home and work.

Coordinated Internet wide analysis of attacks.

The main problem with Internet traffic is that it is nearly impossible to analyse all the traffic going across multiple backbones belonging to different organisations and companies. These Internet Service Providers can not stop every piece of traffic and analyse it to see if it is legitimate or an attack against someone. Some commentators will also say that they are unwilling to, as ultimately it could reduce traffic, increase overhead costs and slow performance against other ISPs who don't monitor. Therefore a few Internet organisations have been created by some activists to try and collect information on IP addresses that harass others whilst on the Internet. These organisation collect data from firewall logs and generate reports on who is currently the top attacker, the current top method of attack and attempt to contact the ISP's involved to get some action from them. It is worthwhile to check out these sites as your IP address may be there or one of your systems may be vulnerable to the current top attack.

www.dshield.org (look at are you cracked ?)
www.mynetwatchman.com (look at hot incidents)

www.spamcop.net (allows you to report and see reports of sites none to send SPAM e-mail.)

www.incidents.org (general information on Global situation)

Advertisement “ Be Pro-Active Join DSHIELD”

Summary

As stated in the introduction this paper is based on an actual incident. It is essentially a back-track and covering of items that I had in place that I needed to resolve the issue of a WAREZ site on a customer machine. The customer was not terribly security conscious, though after two weeks without a site a mood change is in the air. After the incident I actually had a bit of time to think about the other machines that were my responsibility in the customer area and what pain I would have to go through if one of them was compromised. The process I went through on these machines closely resembles the layout of the paper where checking of patches and software upgrades was not assumed and confirmation was required. The paper installation logs stated that “system hardening had occurred” with no specification or records kept as to what was found, what tools were used and what procedures followed. System log files were ok, except that old files were stored on the machine giving a history of everything done on it. Whilst no Name servers were installed, except for external DNS, hosts files listed some of the internal systems and all the Windows systems responded to IPC\$ probes.

Like a house when you lock it up before you go out you first hide the family jewels, close and lock the windows and then finally lock your door on the way out. Hopefully you have found the information and references useful.

References

The following books and sites have provided information and opinions, which I gratefully acknowledge.

The CERT Guide to System and Network Security Practices by Julia H. Allen (particularly Chapter 2), publishers Addison Wesley, May 2001

Unix System Security by Rik Farrow, Chapters 5 and 6, publishers Addison Wesley, DEC 1990

Network Intrusion Detection, An Analyst’s Handbook by Stephen Northcutt and Judy Novak, Chapters 9 and 10, publishers New Riders, Sept 2000

Designing Network Security by Merike Kaeo, publishers Cisco Press, 1999

Hackers Beware by Eric Cole, publishers New Riders, Aug 2001

The following sites have been used for software tools, documentation are articles

Provision of Nmapnt and documentation

www.eeye.com/html/Research/Tools/nmapnt.html

Information on the Digitoll product

www.ideadata.com.au

Information on Intrusion detection and application

www.snort.org by Marty Roesch

CERT advisory on SNMP

www.cert.org/advisories/CA-2002-03.html

Reference site used to obtain Typhon software and documentation

www.net-security.org/various/software/999866655,6627,windows.shtml

© SANS Institute 2002, Author retains full rights.