

Halmstad University

School of IDE

Network Security, 7.5 hp

Network Design and Computer Management

Intrusion Detection Systems

Spring term 2008

Supervisor: Yan Wang

By:

Muhammad Sarmad Munir

Saad-ul-Hassan

Introduction: Critical Issues

- Prevention
 - Avoiding possibility of break-ins, exploitation and breaches of privacy
 - Verifying identity
 - Placing access controls on data
 - Encrypting communication
 - Some software rules for keeping out hackers
- Intrusion Detection/Security Surveillance
 - How do we know if our security was successful or not?
 - How can we determine whether security was breached?

Intrusion Detection Systems (IDS)

- As with cryptography,
 - IDS are the latest craze in computer security.
 - IDS is over-hyped but nonetheless interesting.
- IDS is the last thing you should think of implementing at site.
 - No sense in having a security camera to watch thieves rob your bank, unless you lock the door.
 - Otherwise this is just a kind of voyeurism.
 - IDS will probably not save you from any advanced attacks, because the sophisticated intruder knows how to get around the alarms, or fool them.
 - Probably the best one can hope for with an IDS is to gauge an impression of how many potential intruders are trying the lock.

What is an Intrusion?

- Can be difficult to define.
 - Someone trying to login at root once?
 - Someone trying to login at root fifty times?
 - Port scanning, SATAN or ISS scan?
 - Someone trying a known security hole?
- The aim of an intrusion detection system is to detect break-ins in progress so that something can be done about them.
- Obviously the first thing to worry about is how difficult it is to break in in the first place.
- If we have done the job of securing data well enough, why are we worried that anyone will be able to get in?

Problems with IDS

- Intrusion detection is a form of fault-diagnosis.
- Faults (in a security system) are not supposed to happen, but the fact is that they do happen.

- As with all fault diagnosis systems, IDS give the wrong answers from time to time.
- Because it is so difficult to define what intrusion actually means in a generic sense, intrusion detection systems tend to err on the side of caution and report many *false positives*, i.e. false alarms.

More Problems with IDS

- Very difficult problem to do in real time.
- What does real-time mean?
- Some attacks are stealthy and occur over many hours or days.
- How can we make a prompt notification about such attempts?
- The intrusion detection will have to be fast to detect quick break-ins, but have a long memory in order to see slow ones
 - Like the thief digging a tunnel into the bank with a tea-spoon.
- How will we be alerted or notified about intrusions?
 - By alarm on screen? By E-mail or pager alert?
 - What if attacker first knocks out E-mail or the pager link?

Drawing the Line - Privacy

- User privacy is a major problem.
- If an intrusion detection system examines everything going on within the system, looking for suspicious behaviour, is that an intrusion of privacy?
- What if humans never see the data, but only the warnings?
- Where do we draw the line between justified and unjustified surveillance?
 - Law enforcement agencies have been arguing for years!
 - E.g. CARNIVORE

Intrusion Detection – Practical Approach

- In 1986 Dorothy Denning published a paper "An intrusion detection model" which has been the basis of much of current thinking on the subject.
- Basically one audits system activities, and network communications traffic, looking for suspicious signatures. What is a signature?
 - File existence/checksum violations (Viruses, Trojan horses)
 - File permission violations
 - Illegal processes/missing processes
 - Packet sniffers
 - Port scanners (nmap)
 - Covert communications (eggdrop, irc-bots)
 - Suspicious traffic
 - Tampering with a honey-pot

Packet Sniffers

- Packet sniffing is the act of capturing packets of data flowing across a computer network.
 - S/W or device used to do this is called a packet sniffer.

- Has legitimate uses to monitor network performance or troubleshoot problems with network communications.
- Also widely used by hackers and crackers to gather information about networks they intend to break into.
- Possible to capture data like passwords, IP addresses, protocols being used on the network and other information that will help the attacker infiltrate the network.

Detecting Packet Sniffers

- Detecting rogue packet sniffers on a network is not easy.
- By its very nature the packet sniffer is passive. It simply captures the packets that are traveling to the network interface it is monitoring. That means there is generally no signature or erroneous traffic to look for that would identify a machine running a packet sniffer.
- There are ways to identify network interfaces on your network that are running in promiscuous mode and this can be used as a means for locating rogue packet sniffers.
- If interested, look at the [Ethereal](#) program. Learn what types of information can be discerned from the captured data and how you can use it.
- Be aware that users on your network may be running packet sniffers, either experimenting out of curiosity or with malicious intent.

Port Scanning

- Common way for hackers to gather information about a network, is to perform a port scan.
- Port scanner is simply a program which attempts to establish a network connection to every single port number 1,2,3,4,5....5000,... on every host on the network.
- By seeing what kind of response it gets, the program is able to guess what network services are running on which hosts.
- This gives hackers a good idea about what services can be exploited.
- Sometimes, also possible to see version numbers of software and thereby identify any servers which have known security holes.

Port Scanners

- Port scanners are freely available.
 - For example, nmap
- Because early port scanners could be detected as they activated little used ports, or all ports in a sequence, newer port scanners have so-called stealth scan functionality
 - A stealth scan occurs over a long period of time, and randomizes the order of the ports scanned

Port Scanner Output

- A port scan looks like this:

```
host% nmap 192.0.2.*
```

Starting nmap V. 2.07 by Fyodor (fyodor@dhp.com, www.insecure.org/nmap/)
Host (192.0.2.0) seems to be a subnet broadcast address (returned 23 extra pings). Skipping host

Strange error from connect (101):No ports open for host (192.0.2.0)

Interesting ports on cadeler30-gw.uninett.no (192.0.2.1):

Port	State	Protocol	Service
23	open	tcp	telnet
79	open	tcp	finger
6001	open	tcp	xwin

Interesting ports on sigmund.iu.hio.no (192.0.2.2):

Port	State	Protocol	Service
9	open	tcp	discard
13	open	tcp	daytime
37	open	tcp	time
79	open	tcp	finger
80	open	tcp	http
111	open	tcp	sunrpc
487	open	tcp	saft
515	open	tcp	printer
706	open	tcp	unknown
751	open	tcp	kerberos_master
1024	open	tcp	unknown
1025	open	tcp	listen
2003	open	tcp	cfingerd
2049	open	tcp	nfs
5308	open	tcp	cfengine
6000	open	tcp	xterm

Honey Pots and Tar Pits

- A *honeypot* is a trap set to detect or deflect attempts at unauthorized use of information systems.
- A *Sticky Honeypot* or *Tar Pit*, an internet-attached server that acts as a decoy, luring in potential hackers and responding in a way that causes their machine to get "stuck", sometimes for a very long time.
- Generally consists of a computer, data or a network site that appears to be part of a network but which is actually isolated and protected, and which seems to contain information that would be of value to attackers. A honeypot that masquerades as an open proxy is known as a sugarcane.
- A honeypot is valuable as a surveillance and early-warning tool. While often a computer, a honeypot can take on other forms, such as files or data records, or even unused IP address space. Honeypots should have no production value and hence should not see any legitimate traffic or activity. Whatever they capture can then be surmised as malicious or unauthorized.

Honey Pot Method

- **Honeypots** can carry risks, and must be handled carefully. If not properly walled off, can be used to break into system.
- *Honeypots* are designed to mimic systems that an intruder would like to break into but limit the intruder from having access to an entire network. If a honeypot is successful, the intruder will have no idea that s/he is being tricked and monitored. Most honeypots are installed inside firewalls so that they can better be controlled, though it is possible to install them outside of firewalls. A firewall in a honeypot works in the opposite way that a normal firewall works: instead of restricting what comes into a system from the Internet, the honeypot firewall allows all traffic to come in from the Internet and restricts what the system sends back out.

Honey Pots Opportunity

- By luring a hacker into a system, a honeypot serves several purposes:
 - The administrator can watch the hacker exploit the vulnerabilities of the system, thereby learning where the system has weaknesses that need to be redesigned.
 - The hacker can be caught and stopped while trying to obtain root access to the system.
 - By studying the activities of hackers, designers can better create more secure systems that are potentially invulnerable to future hackers.

Honey Pots Goals

- Learn how intruders probe and attempt to gain access to your systems. The general idea is that since a record of the intruder's activities is kept, you can gain insight into attack methodologies to better protect your real production systems.
- Gather forensic information required to aid in the apprehension or prosecution of intruders. This is the sort of information often needed to provide law enforcement officials with the details needed to prosecute. More important, when you decide you're going to build a honeypot you must first realize that you're playing with fire and can easily get burned.

KF Sensor

- KFSensor - Honey Pot server
- Attract and detect hackers and worms by simulating vulnerable system services and trojans.
- By acting as a decoy server it can divert attacks from critical systems and provide a higher level of information than can be achieved by using firewalls and NIDS alone.
- Known attacks patterns are identified by the signature IDS engine and the signature details are available in the events view. (Output shown below)

The screenshot shows a window titled "Event - 1020" with a close button in the top right corner. Below the title bar are four tabs: "Summary", "Details", "Signature", and "Data". The "Signature" tab is currently selected. The window is divided into two main sections: "Event" and "Signature".

Event Section:

- Sensor ID: New York
- Event ID: 1020
- Start Time: 15/07/2005 11:21:57.890
- Severity: High

Signature Section:

- ID: K FAGC174151
- Message: IIS - RBOT Worm propagation
- Reference: <http://www.keyfocus.net/kfsensor/signature/sigb> (with a "Browse" button)
- Source: KeyFocus
- Created: 03/07/2005 17:44:27.375
- Edited: 05/07/2005 23:08:08.062

At the bottom of the "Signature" section are "Edit" and "Create" buttons. At the very bottom of the window are "Next", "Previous", "Close", and "Help" buttons.

Encrypted/Fragmented Data

- Note that analyzing network traffic becomes impossible as traffic is encrypted, so this approach probably does not have much of a future.
- Today intrusion detectors have to try to reassemble fragmented packets to follow data streams long enough to analyze their content.
- Requires work at very high speed.
 - Packets get dropped.
 - The detectors can be fooled.
 - They probably have exploitable bugs.

Types of Intrusion Detection

- There are two types of intrusion detection
 - *Rule based intrusion detection*
Testing for specific occurrences
Example: seeing whether a particular private port is accessed.
 - *Statistical anomaly detection*
Looking for anything out of the ordinary
Collecting data on what 'ordinary' is

Tracking Suspicious Behaviour

- This requires a large knowledge-base of things which people have identified as suspicious already.
- Signatures are also specific to OS flavours, in practice that means you will find off-the-shelf S/W for NT and Solaris, maybe GNU/Linux.
- Databases need to be updated to detect new signatures.
- IDS can detect a few things quite well, but this approach is no more advanced than military censorship of the mail during wartime.
- How can they know about every protocol? What about meta-protocols like RPC and SMB?

Anomaly Detection

- A very interesting idea which might be used both in fault-diagnosis and security intrusion detection is the idea of *anomaly detection*.
- In anomaly detection we are looking for anything abnormal
- Could come from
 - Abnormal traffic
 - Patterns of kernel activity
 - Changes in the statistical profiles of usage

Anomaly Detection with Neural Nets

- Need some method for tracking patterns in statistical samples.
- Neural networks have been used for this, but problem here is no one really understands how neural networks work
 - NN classify information by lumping similar things into similar categories. Neural networks first have to be trained using "normal" data, then they can be switched into production mode in order to detect anomalies.
 - There is a *coarse graining* of information which means that networks throw away all but the information parameters which the network models. When a network detects a signature, there is never enough information left to be able to say why they produce the result they do! This can lead to embarrassing problems.

Problems with Anomaly Detection

- Basically anomaly detection is an unsolved problem
- There is a lot of research into it because it is very interesting and it has the potential to solve a general problem without a rule-base.
- These systems have to look at data over long periods of time and this costs a lot of data storage.

Specific Alarms

- We can rig up alarms on specific objects which reflect the way they are used and ways in which they should not be used.

- For instance, if someone edits the password file without registering first with a monitor.

Immune System

- The most efficient intrusion blocking and detection system which exists is in the human body.
 - Only one in ten million harmful organisms entering the body ever get past our initial defenses.
 - Of those, many will be destroyed by natural killer cells which react to the vandalism caused by bacteria and viruses.
 - Those which remain are targeted by a specific immune response (B and T lymphocytes) which are able to detect something like 1012 different harmful entities.
 - Even with this vast repertoire of security checks, we still get sick and need the help of doctors!

Computer Immune Systems

- The human immune system has been the inspiration for several approaches to computer security.
- Researchers in New Mexico have tried looking for suspicious activity by matching signatures of kernel system call patterns.
- Allows detection of exploits in programs which are not Trojan horses.
 - For example, when the MD5 checksum is correct, but when a program is being exploited to do something dangerous, as in a buffer overflow.

Intrusion Detection Tools

- Cfengine/Tripwire - checksums and file permissions
- [Cfengine](#) - other configuration issues
- TCP wrappers - service requests
- TCP dump/snoop - net traffic
- [Network Flight Recorder](#)
- Argus

Scary Thought

- The big problem with intrusion detection is that the intrusion detection itself system becomes a target for attack.
- Since they generally work very hard, on the edge of what is possible, it is probably not hard to fool them, or at least execute a denial of service attack on them.

Conclusions

- Though it's a new and useful technology but still it has some loop holes and more research is required
- Its good and practical approach to implement IDS in an existing network system

- Many IDS tools have been created but more evolution is expected
- Even after such extensive research in IDS, intrusion cant be prevented completely in any network system

References

- www.answers.com
- www.whatis.com
- IEEE Network May/June 1994
- <http://netsecurity.about.com/cs/hackertools/a/aa030504.htm>
- http://en.wikipedia.org/wiki/Intrusion_detection_system