Understanding Intrusion Detection Systems

You can use firewalls, secure logins, one-shot passwords and encrypted IDs and still not know if your system has been violated. The emerging field of intrusion detection offers some surprisingly low-tech solutions to a high-tech problem.

By Clive Grace
IT Journalist

The proliferation of applications (especially Web sites) requiring some form of e-commerce, and the increasingly important role that networks play in modern business, has given new impetus to the search for developing more secure systems. Intrusion detection products are gaining widespread recognition as important tools that improve the security of a computer network. Although firewalls have traditionally been seen as the "first line of defence" against would-be attackers, intrusion detection software is rapidly gaining ground as a novel but effective approach to making your networks more secure.

Intrusion detection operates on the principle that any attempt to penetrate your systems can be detected and the operator alerted - rather than actually stopping them from happening. This is based on the assumption that it is virtually impossible to close every potential security breach; intrusion detection takes a very "real world" viewpoint, emphasising instead the need to identify attempts at breaking in and to assess the damage they have caused.

The fact is that breaking into a computer system is more often a computer user's hacker fantasy than a system administrator's reality; but, when asked, any system administrator will tell you that there's always a lingering uncertainty after all the security policies have been implemented and (hopefully) adhered to. You can never make an Internet site truly secure, just "good enough"; but how secure is "good enough" and can you detect when, or if, someone is breaking in?

The answer is usually "yes" but, whereas firewalls tend, at best, to merely bat away would-be intruders with no real cracking experience, detecting an intruder is normally best achieved by keeping up with known security breaches. You'll achieve more by finding out about new break-in methods, keeping an eye on user logs, quickly closing loopholes and keeping a firm grip on what gets installed onto your machines than by throwing money at the problem by buying "better" software packages.

The issue of rogue software installations is a very important one - it only takes one employee who decides to set up their own remote access to a desktop machine, using a product like pcANYWHERE without a password, in order to allow a back door into the system you're trying to protect. So, knowing that the risk of such an attack is real, how can you tell if this is happening on your network, and what can you do to stop it?

Identifying Attacks

First of all you need to ask yourself just how such an attack would appear from a system's or network administrator's standpoint, and that depends on what type of monitoring and logging software you're already using. Create a checklist:

- What are these monitors and logs actually pointed at?
- Where are your logs stored, and how are they stored?
- How are they protected from modification?
- Which firewall packages, products, terminal servers and internal machines are logging these events?
- How are your logs reviewed?
- Do you know what to look for in a log file?
- Do you have tools to sift the legitimate users from potential intruders?
- Is there an "alarm" mechanism installed?

Many of these factors (and there are a great many more) will determine the scope of an attack and how quickly you can detect it.

It is fair to say that most, if not all, organisations concentrate most of their intrusion detection efforts on the firewall - so much so that even seasoned system administrators can ignore internal machines and direct dial-up lines, and this is why the aforementioned pcANYWHERE scenario is all too often the reason for a security breach at an otherwise secure site.

Attackers generally don't bother with Internet connections - at least, not for their first line of attack. They prefer instead the relative stealth of an unwatched back door via dial-up connections or fax terminals. Even the most expensive and well-configured firewalls almost universally fail to respond to unusual activity coming from trusted users or external machines such as Web servers, DNS servers, administration accounts and the like.

The best advice is to thus monitor your entire network perimeter. Not just firewalls, but all your dial-up connections and your VPNs, especially those that offer direct access to file servers and so on. Don't assume the threat is always from outside, either; malicious employees or contractors with an axe to grind can happily do as they will relatively undetected if all your scanning efforts are directed towards your firewall. When monitoring your network perimeter it is often sensible to enforce some sort of security monitoring on internal machines. This is where intrusion detection software becomes useful.

Acting On A Breach

It is very difficult to decide how to go about securing a large network - not to mention influencing an even larger number of users to practice good network security. The usual top-level reaction to the presence of an intruder is to shut down all Internet connections pending the purchase of better security software. This is normally not the best policy. Despite the problem, the solution is not usually an overly technical one.

Start by taking a long hard look at how the intruder broke in and see if there are additional holes that you have not yet plugged, based on their method of entry. Examine what preventative methods you already have in place, and concentrate on isolating overall weaknesses so you can build a stronger foundation that can realistically improve security without wasting your resources on ineffective security measures.

Armed with this information you can then draft up a shopping list of security methods based on a thorough risk assessment; the assessment will help you specify which products and tools are needed to enforce your new security policies. Merely buying better security software is often a waste of time and money unless you pay close consideration to how to configure and manage it. Putting it bluntly, without understanding the risk, you'll never know when your security is good enough.

Intrusion detection software is by definition a rather complex beast. Usually designed to sit between your users and the firewall (which in turn sits between your users and the Internet), intrusion detection software is often incorrectly regarded as merely being the second line of defence. With a firewall in place, a good IDS package should be able to sort out legitimate users from suspect users relatively easily.

"Misuse detection works by searching for a set of known attacks that have been stored in a system's database. The knowledge of the attacks is encoded as a set of footprints."

Detection Software Types

There are broadly two types of intrusion detection package available. The first divides the techniques of intrusion detection into two main detection methods: anomaly detection and misuse detection.

Anomaly detection packages are based on a set of statistical details that model the behaviour of your users or your host machine. The profile of a particular user, for instance, may include information such as the mean duration of telnet or ftp sessions, the number of bytes typically transmitted (in either direction); the time of day the terminals are accessed and so on. The profile of a host machine, on the other hand, might include average CPU utilisation, the average number of users logged in at specific times and from specific locations, and so on.

Anomaly detection software packages monitor the operation of your systems and constantly compare the profiles of, say, a current user session with one stored in its database. If it detects what it considers to be a large deviation from normal behaviour, it signals an alarm to the system security officer.

Misuse detection works by searching for a set of known attacks that have been stored in a system's database. The knowledge of the attacks is encoded as a set of footprints (sometimes called signatures), which are patterns that occur every time an attack takes place. Some popular misuse detection packages were developed directly as a result of Dan Farmer's SATAN (System Administration Tool for Accessing Networks) being used by hackers to determine what holes there were in systems. SATAN has a very definite footprint, and counter-measure packages like GABRIEL were soon developed to detect SATAN's very obvious network footprint as it ran through its Perl scripts to determine the weak points at a particular site.

The latest generation of misuse detection packages take this further by employing an expert system that performs pattern matching against a stored rule base. Somewhat like an antivirus package, misuse detection software requires constant updating of the rule base as new attack methods become known and new operating systems and network protocols open new holes in otherwise secure systems.

Single And Multiple Hosts

The other type of intrusion detection package is based on whether the software monitors activity on a single host or on multiple hosts interconnected by a network. Many of the earliest intrusion detection systems (many of which are still in use today) take a simple approach to examining audit data from a single machine - deriving conclusions based solely on that information. These intrusion detection packages can't detect attacks that are orchestrated from several sources, or attacks that span multiple machines in a network. Furthermore, they rely heavily on the logs provided by the underlying operating system, and this renders them fairly architecture-dependent and vulnerable to Denial of Service attacks. An intruder may manage to delay the logging mechanism, or even turn it off altogether.

Common IDS Packages

Having discussed the issues in general, we'll now discuss the advantages and drawbacks of four common intrusion detection packages.

ISS RealSecure

RealSecure by Internet Security Systems (**www.iss.net**) is probably one of the best-known intrusion detection packages on the market. When connected to a network, it listens to all traffic passing through it, searching for matches against patterns it is configured to detect. It can monitor TCP, UDP and ICMP traffic and, if a match is found, counter-measures can be implemented along with read / write server locking, IP blocking and so forth. This product has already developed a fairly large user base because it is bundled with CheckPoint Software's Firewall-1 - a very competently implemented software-based firewall package.

CyberCop

"The amount of network traffic an intrusion detection package introduces can be phenomenal, despite many ID vendors' claims."

From the Network Associates stable, CyberCop (**www.nai.com**) is very similar to ISS RealSecure in the way it detects potential break-ins, but its architecture is more distributed. CyberCop is composed of a number of sensors that are scattered among the network nodes along with a management server that collects and collates intrusion reports detected by the sensors. If an intrusion is detected, the management server presents an elaborate and concise description of the event to the security manager, who can then deal with the problem. It has an easy-to-use interface which operates under Windows NT, although a text-only console version is all that is really needed for this sort of application.

Bro

Whereas the previous two packages are available for very specific operating systems (especially for the burgeoning Windows NT security marketplace), Bro is a freeware package developed at the Lawrence Berkeley National Laboratory.

Its source code is freely available and the architecture on which it is based is a modular one, which means you may compile a version for a plethora of machines, networks and operating systems. Based around an event engine it makes use of a policy script interpreter, which implements your security policy via a proprietary language. You can get more information about Bro by ftp at **ftp.ee.lbl.gov/papers/bro-usenix98-revised.ps.Z**.

NID

NID is also a popular and freely-available intrusion detection package. Installed on a dedicated machine (the only software package of the four here that needs a whole machine to itself), NID monitors network traffic and scans for the presence of known attack signatures, as well as deviations from normal network behaviour. If an intrusion is detected, the security manager is notified via email, pager or a short message to the console. NID is, in comparison, a very simple package to use, but one which needs a fair bit of specialist knowledge to implement when compared against more commercial alternatives. You can get more details about NID on the Web at **ciac.llnl.gov/cstc/nid/nid.html**.

CIDF And IDWG

Despite the fact that such systems are still relatively new territory for network administrators, there are already efforts to standardise the design of intrusion detection software - especially in the way software packages from different vendors can communicate with each other. Overseen by the Common Intrusion Detection Framework (CIDF), it is an effort to develop common protocols and application programming interfaces for communication purposes. Another group, the Intrusion Detection Working Group (IDWG), is also examining methods of providing a common method of intercommunication between intrusion detection software from different vendors.

Conclusion

Intrusion detection is based on the fairly reasonable assumption that, no matter how well you configure your firewalls and "bomb-proof" your systems, you cannot be utterly certain that an intrusion will never take place or that it will never go undetected. It is therefore important to detect an intrusion attempt, preferably while it is happening, so that you can ensure that nothing has been lost and no records altered. Even if detection is made after an attack has occurred, it is still better to identify guilty parties and weak spots, and make an assessment of any damage after the event than never.

A good intrusion detection policy is often the one that is most strongly tailored to fit your network. Not all packages will be suitable for your requirements, and you will ultimately need to decide if you need all your network traffic monitored or if you can get away with only scanning key systems that initially need protection. The amount of network traffic an intrusion detection package introduces can be phenomenal (despite many ID vendors' claims), so expect your network to significantly drop in efficiency after you've implemented a tough intrusion detection policy. Bear in mind that not all systems should be treated with equal rigour, and define certain "exclusion areas" which are not scanned all the time but which can be scanned when needed - for example, immediately after your intrusion detection system has announced a potential intruder.

"Intrusion detection operates on the principle that any attempt to penetrate your systems can be detected and the operator alerted - rather than actually stopping them from happening."

Intrusion detection software tends to be something of a long-term investment. It doesn't give instant results, and there's no point where you can say "that's it; we're secure". Unlike a firewall, which offers an immediate layer of protection from the outside world, IDS software becomes more adept over time, as network profiles become known quantities. Only when systems settle down to patterns of normal usage can an intrusion detection package really be considered to be pulling its weight.

PCNA

New Reviews from Tech Support Alert

Anti-Trojan Software Reviews

A detailed review of six of the best anti trojan software programs. Two products were impressive with a clear gap between these and other contenders in their ability to detect and remove dangerous modern trojans.

Inkjet Printer Cartridge Suppliers

Everyone gets inundated by hundreds of ads for inkjet printer cartridges, all claiming to be the cheapest or best. But which vendor do you believe? Our editors decided to put them to the test by anonymously buying printer cartridges and testing them in our office inkjet printers. Many suppliers disappointed but we came up with several web sites that offer good quality cheap inkjet cartridges with impressive customer service.

Windows Backup Software

In this review we looked at 18 different backup software products for home or SOHO use. In the end we could only recommend six though only two were good enough to get our “Editor’s Choice” award

The 46 Best Freeware Programs

There are many free utilities that perform as well or better than expensive commercial products. Our Editor Ian Richards picks out his selection of the very best freeware programs and he comes up with some real gems.