



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Who's Afraid of the IPS?

Mike Jawetz

GSEC Practical Assignment Version 1.4b (amended August 29, 2002)

Abstract

Network security spans multiple technologies that focus on specific security requirements. Firewalls concentrate on protecting the perimeter, antivirus technology searches for viral or mobile code attacks. Intrusion detection systems look for patterns of traffic that could lead to system failures or hijacks and vulnerability assessment look for system weaknesses. Each discipline develops “best of breed” products that are used in conjunction with common best practices that recommend deployments that include all of these devices to provide a sound defense in depth strategy. As these systems get better, there is a desire to move from a defensive mode into a more proactive prevention mode that could improve the return on investment for these devices.

An Intrusion Prevention System (IPS) is the logical next step in the evolution of network and host security. These systems provide a combination of firewall, intrusion detection, antivirus protection and vulnerability assessment into a single device that can be deployed in a number of different ways within network architectures. Each of the underlying technologies offers significant value on its own and can be obtained from various vendors, each with their own ideas for a “best of breed” implementation. The question is, by taking concepts from each of these components and encapsulating them into a single device, is a user being served up a “silver bullet” that has the power and opportunity to do more harm than good? Customer support groups report multiple instances of unhappy customers suffering from configuration errors, device failures and poorly thought out deployments that have resulted in lost productivity all in the name of network security. The worry is that these devices will block legitimate traffic, cause network latency and present a single point of failure¹. Should access to network services be controlled by technologies that rely on detection and response methodologies that are notoriously imperfect? Until each technology can demonstrate complete mastery over its domain, whether it is in blocking malicious traffic, or virus protection, or intrusion detection or vulnerability assessments, we may not want to dedicate all of these functions to a single device in hope of finally solving a significant component of network security. This paper examines the different technologies incorporated into IPS devices and presents the argument that we may not be ready to deploy intrusion prevention systems as a single, inline device in the hope that it will provide network intrusion protection and prevention.

Reliable prevention depends on accurate information from multiple sources

A “defense in depth” methodology recommends deploying multiple detection and recognition technologies. These devices must be flexible enough to accommodate the unpredictable evolution of attacks, changing network and host characteristics, and business impacts. They must also work together in concert to be effective in protecting the enterprise. So what kinds of protection do these discrete systems provide and where do they fall short?

Firewalls – Protection at the perimeter

There are several different kinds of firewalls that are used to provide perimeter protection. They include a full application inspection (or proxy) firewall, a stateful inspection firewall and a packet filter firewall. Depending on the type of network traffic and configurations deployed, each design has its own inherent advantages and disadvantages.

Proxy Firewalls

A full application inspection (or proxy) firewall is the most secure and flexible of the traditional firewall types. A full application inspection firewall provides protection at the application layer, thus protecting your network from the most common types of attacks. Since it understands the application portion of the data packet, you can design very granular security policies. For example, an administrator can create a rule allowing file transfer protocol (FTP) PUT commands, but denying all other FTP requests. In addition, a full application inspection firewall supports a wide range of authentication mechanism. Therefore, not only can your security policy be more granular, but it can also be based on users rather than easily spoofed IP addresses.

Another benefit of a full application inspection firewall is that for every connection request, it creates another new connection request on the behalf of the requester. By creating a brand new connection the firewall is essentially protecting the back-end devices from network-based attacks that rely on improperly constructed data packets or fragments.

These are the most secured of all the traditional firewall technologies; flexible and granular security policies; easier to maintain; more network address translation (NAT), and logging features; user authentication provides a higher degree of protection than IP-address based policies.

The drawbacks are that they are more resource intensive than other firewall types and only major protocols are supported via native proxies. Because of the intense inspection they are doing, performance can be slow relative to the other types described next.

Stateful Inspection Firewalls

A stateful inspection firewall (or circuit-level) can maintain “state” information about a connection. By maintaining a state table of the connection, it reduces the number of

rules that are necessary to maintain versus a packet filter firewall by 50%. Secondary features include the ability to NAT traffic and perform detail logging. A stateful inspection firewall normally works at the network and session layer exclusively.

These firewalls tend to be very fast. They are also easier to maintain than a packet filtering firewall.

On the down side, the rule sets are order dependent and require specialized tools to make sure an administrator doesn't inadvertently overwrite a DENY rule with an ALLOW rule. It is possible that a stateful inspection firewall will leave a network susceptible to application level attacks, such as buffer overflows and improper application methods. Since it does not understand the application protocol, it could leak internal information out through the application payload. This implementation also provides no authentication.

Packet Filtering Firewalls

A packet filtering firewall simply allows or denies data packets based on the protocol type, and the IP address source, and destination. All decisions are made at the network layer.

The advantage of this type of firewall is that it is very fast. They are often used in conjunction with other types of firewalls where speed is extremely important and throughput cannot be impeded. Once the majority of packets have been filtered out, the more resource intensive firewall architectures can be put to better use.

These firewalls tend to be management intensive; rule sets are order dependent and require two rules to define by-directional traffic flow. They are not very secure since decisions are based on IP headers only, which are very easy to spoof and bypass security policies. A packet filter firewall does not provide protection against the most common application attacks.

Vulnerability scanners – Put potential vulnerabilities into context

Administrators need an up-to-date picture of what's running on their network, where the holes are, and what's patched and what's not. Vulnerability scanners are designed to reveal where the deficiencies are and what needs to be done to correct them². When looking at the functionality and benefits the criteria are:

1. **Mapping.** Many network managers aren't even sure exactly what's running where. Sometimes, surprises spring up even when they think they know what's out there. Network mapping is a critical first step in any network security project.
2. **Vulnerability analysis.** Most of these scanners come with more than 1,000 tests to find software and configuration problems. But do those tests work? How many vulnerabilities will slip in under the radar?

3. **Data management.** Large networks generate hundreds, if not thousands, of records of network map and vulnerability information. How do these tools let network managers sort, sift and report on all that data--and how hard they have to work to do it?
4. **Performance.** These tools aren't designed to run in real-time, but normally you'd want a scan of even a large network to complete in less than a day. It's critical that the benefits these tools provide outweigh the potential loss if systems crash as a result of using them.

These tools search for misconfigured application servers, such as web servers; and network components, such as switches and routers, that are vulnerable to known problems. They look for out-of-date applications, especially those with known problems. And they often search for applications that are enabled by default--but perhaps shouldn't be, such as RPC services on UNIX hosts or the UDP ECHO program on Windows NT/2000. Vulnerability scanners are also security oriented, so they often look for "information leakage" from systems through domain name servers (DNS) and other avenues, including the simple network management protocol (SNMP) and the Windows registry.

Most vulnerability scanners take a three-phase approach to testing: Given a network range by the security manager, the vulnerability scanner attempts to determine which IP addresses are in use. This phase usually includes tools such as *ping*.

The vulnerability scanner attempts to determine which applications and services are running on these systems, and their configurations. The vulnerability scanner uses a variety of techniques, ranging from simply trying to connect (a port scan) to gathering actual socket information out of SNMP.

The tool employs a long series of tests to find out if each system is susceptible to a particular known bug or problem. Smarter products iterate between phases two and three, learning more and using that information to launch additional tests. Others have ways of pruning their decision tree to save time and minimize the risk of overloading the target systems.

There are many variations within these three phases. Some products try to brute-force guess passwords on accounts. Others assume a "friendly" environment and connect to servers with administrative access to look for problems at the system level. Some are more devious, and will try to evade a network IDS.

IDS – Detection, analysis and response for hosts and networks

Intrusion detection systems (IDS) are often lumped into two categories: host based and network based. A host based IDS (HIDS) system is comprised of agents loaded on servers, desktops and/or network elements such as routers that monitor various system components. The information is sent to a management station that is both the repository and management point for multiple agents. A console communicates with the management station to examine the data and configure the agents. Each

agent can examine such system components as file integrity, host activity and user authentication. A network based IDS (NIDS) can also be a three tier (or in some cases a two-tier) architecture. Sensors monitor subnets, VLANs or individual servers looking for known signatures, protocol or behavior anomalies and denial of service attempts. HIDS and NIDS can work together to provide visibility into the network that are unique perspectives from one another. While firewalls are designed to block specific types of traffic based on ports or addresses, a productive network must be allowed to pass traffic (port 80 for web, port 53 for DNS and port 21 for FTP for instance) through the perimeter. A typical network based intrusion detection system provides deep packet and (sometimes) protocol inspection, event analysis and correlation and automated response policies using combinations of signature and anomaly detection methods.

Antivirus – Detecting and removing malicious code

Whenever a user opens, moves, reads, or mails a file, the AntiVirus mechanism on the system scans it for known viruses. Every email that goes through a mail server has its attachments scanned by that mail server's scanner. All servers are automatically scanned overnight for infected files. These scanners are good only if they're kept up-to-date with the latest virus definitions.

Most antivirus software works by looking for known, defined viruses that have already been found and dissected by security companies. Virus-scanning software uses definition files and updates to detect new viruses.³

These tools are typically composed of three components, a scanning application, a scanning engine and virus definitions:

Scanning application - Contains the user interface and configuration options, identifies which files are to be scanned, and handles any communication with the user or administrator, such as alerts.

Scanning engine - The heart of the antivirus product, which conducts the actual scan of a file, detects strings or strange behaviors (through heuristics), and performs any necessary repair or cleaning operations.

Virus definitions - Contains a database of known virus signatures as well as instructions on how the engine should clean or repair infected files.

In conjunction with these antivirus tools, companies try to educate their users not to open attachments from people they don't recognize, as this is the single most common way for viruses to spread within an enterprise.

These security technologies are not infallible

Intrusion prevention combines components of all the technologies discussed here. Each has its own challenges to conquer. There are issues with all of them that encourage continued development. How does each address such problems as incomplete attack coverage? What about inaccurate detection? What about the performance challenge each faces at its own point of presence? Each technology is briefly examined here to make the case that a true IPS solution, used in the manner in which it was designed, may not be able to satisfy the lengthy list of requirements necessary to justify its position in the network security infrastructure at this time.

Firewalls – Protecting the perimeter

Firewalls provide port, protocol and direction blocking. One of their greatest strengths is also one of their greatest weaknesses; a misconfigured firewall can block legitimate traffic or let malicious traffic in. Just having a firewall is not enough; knowing how to configure it properly, maintain it over time and monitor the information all become critical as part of the deployment ⁴.

Typically there is no deep packet inspection as the primary function is to filter out the known traffic to be blocked. Speed is of the essence, so minimizing laborious inspection functions is often the preferred configuration for these devices.

Vulnerability Assessment Tools – Finding the soft spot

So the goal is to test known weaknesses and misconfigurations against network assets. Is the tool aware of system changes, updates and additions? What about mobile systems? What is the underlying business impact to correct or prevent vulnerability? Vulnerability assessment tools can present information that misrepresent the network when devices are added, removed or reconfigured. Wrong information can cause incorrect preventative responses. For example: patching systems that don't need it or alerting on vulnerabilities on OS'es that are not present in the enterprise. This can take time and introduce unknown side effects. Checking for vulnerabilities can have negative effects on systems under test.

Antivirus tools – Only as good as the last update

Antivirus software is constantly changing to keep up with the high number of new threats that are identified daily. Detection and removal must be performed at both the gateway and at the server/host level. Viruses can and are easily deposited on the network through some internal path (a floppy disk for example) and are not detected until servers have been infected.

Complicating matters is that different AV solutions use different approaches ⁵. Some antivirus companies will detect "families" of viruses by using a strong heuristics base. These vendors may be able to detect and catch new viruses without specific samples. Other companies prefer the scientific method of exact identification, which

enables the AV program to be more precise, more rigorous. Heuristics has the drawbacks of increased overhead and more false positives, while exact identification will allow minor variations to slip through the net, not to mention zero-day viruses. Since our interest is in protecting corporate resources it is incumbent upon us to seek the best possible solution.

Intrusion Detection Systems – Multiple methods, but how accurate are they?

Intrusion detection systems use a combination of signatures, anomaly detection and threshold counts along with statistical analysis to determine good from bad traffic. These systems must keep up with the flow of data or risk dropping packets and being unable to provide stateful inspection or recognition of fragmented attacks. IDSes continue to try and solve the problem of being less reactive and more proactive with the use of automated responses. Every product on the market deploys a slightly different detection and response methodology yet each suffers from missed attacks (false negatives), misidentified attacks (false positives) and an overall inability to proactively respond above and beyond alerting.

Signatures – Pattern matching

Pattern matching for known attacks – poorly written signature can cause false-positives. No defense for new attacks and they are maintenance intensive. They are designed to examine individual packets but many attacks are based on a series of packets or a “flow” of data.

Application layer detection – New applications introduce new vulnerabilities. This type of detection depends on applications being, “well behaved” or designed to support the RFCs as they define how a protocol works. When clever developers figure out ways to overload a data packet, they often overlook the impact the application can have when run on a secure network. Application layer detection can often trigger on legitimate use of custom applications.

Number of signatures too large to check- Performance degradation of IDS can allow attacks to evade the system. A popular example of this is an attack called, “stick”. By presenting a signature based IDS with lots of signatures that can be flagged, the system can become overwhelmed. The attacker can then slip in a “camouflage” attack that is ignored by the IDS.

Anomalies – Traffic that doesn’t quite fit

Anomalies are irregularities from expected patterns and trends. There are several different aspects of anomalous behavior that have been used to determine if network usage might be malicious.

Protocol anomalies

Many applications do not comply with published RFCs. Systems need to keep up with updates to RFCs. They must do more than simply be compliance checkers as

many well-known applications violate the RFCs. Much of the traffic seen on networks, while anomalous, is not malicious. Protocol state engines must support multiple versions of RFCs. They also need to consider the popular applications that use protocols in a way that violates the RFC but are common among enterprise traffic. This must be recognized by the system as an anomalous use of the protocol without being malicious.

Behavior anomalies

It's difficult to predict human behavior. People change. How quickly can system "learn"? What happens in the meantime to filter out the false positives? Normalized traffic can be subjective. What is normal for a significant event on the network (payday) must be flagged under "other" conditions. What are they and how are they characterized? How long will it take to characterize the behavior?

Attack profiles - Evolution and relevance of attack

Attacks can be profiled and modeled. This profile is used by security systems to help identify legitimate attacks. Profiles change, attacks evolve and attack relevance is different on every network. Attacks can be modified to evade the traditional methods of detection. As probes change the way they are launched, can it look like valid application traffic? It may. Multi-homed DNS servers are designed to provide fault-tolerance. To some intrusion detection systems the traffic they generate can look like an attack. Security technologies (firewalls, IDS, vulnerability assessment, and antivirus) need to work in conjunction with each other to recognize this as valid behavior in this configuration.

Defending against Denial of Service – Too many variables

A denial of service (DoS) or distributed denial of service (DDoS) attack is an incident in which a user or organization is deprived of the services of a resource they would normally expect to have. A common form of these attacks is a buffer overflow in which users send more traffic to a network address than the programmers who planned its data buffers anticipated someone might send. Attacking from multiple sources in concert facilitate a distributed version of these attacks. Baseline traffic patterns to characterize what constitutes a "normal" traffic flow is a method used by different devices in an attempt to understand what normal is for any given enterprise. By recognizing patterns of traffic flows, administrators can be alerted to events raised by firewalls, and host and network based IDSes.

The ability to detect attacks directly affects the ability to react appropriately and to limit the damage caused by a DoS/DDoS attack. While IDS systems have grown quite sophisticated and most products available today successfully detect most types of attacks, DoS and DDoS attacks are still difficult to detect with accuracy. The problem with DoS attacks is the sheer number of ways in which they can be executed; the increasingly sophisticated attack methods, and the growing range of systems targeted.⁶

Most of today's IDS products use a very simplistic method of detection. They compare current traffic behavior with acceptable "normal" behavior to detect DoS attacks, where normal traffic is characterized by a set of pre-programmed thresholds. These techniques establish a baseline and then look for "jumps"—situations where the volume of network traffic jumps from low to very high levels. This simplistic approach suffers from several shortcomings. First, the threshold is typically set statically and thus requires user setting for every new environment and cannot adapt to changes to the environment. Second, only a small number of thresholds can be defined because very detailed statistics for protocol breakdown are not available to users. Third, thresholds can only be applied at high aggregate levels, e.g., per subnet, due to the lack of monitoring granularity. These shortcomings can lead to false positives and false negatives in detection depending on the threshold errors. Even if detection is made correctly, a lack of granularity can limit the ability to accurately identify and block the attack traffic.

There is also the question of the kind of information that is presented to the administrator when such an attack occurs. The MS-SQL (slammer) worm that was seen earlier this year was detected by many systems that either had a signature for the known vulnerability or detected an unusual amount of traffic to port 1434, did any report the specific nature of the attack? In other words, did any IDSes report that the packet appears to attack a vulnerability identified by CAN-2002-0649⁷? Without this kind of information, how does one know how to respond? Can a policy be defined within an IPS that responds appropriately each time? This is doubtful.

Prevention cannot outweigh reliability

Everyone has heard the metaphor that one can "win the battle and lose the war". To really prevent any malicious traffic from entering your network, the single safest method that can guarantee results is to pull the wire out. This also has a significant capacity to reduce revenue generation for any company that uses web-based communications. The effort to secure the network cannot outweigh the requirement to maintain functionality, which in this case requires that we stay connected. The next less drastic step is to introduce technologies that can be defensible and ideally proactive in its attempt to protect the corporate assets. If by preventing intrusions we deny service to customers, the cost of prevention outweighs the cost of reliability. To figure out what is being spent to protect a network ask the questions; how much does it cost to buy, and how much to implement and maintain the technology? But what is the cost of the loss of reliability? By adding another bump in the cable, there is a hidden cost that must be calculated for each instance that service is terminated erroneously (and to be fair, what is the savings for preventing true loss through theft or attack)?

Passive detection systems shunt traffic stealthily

The most common configuration for a network IDS deployed today is to detect and alert. Detection is done through SPAN ports of switches, taps or hubs at critical junctions. If some kind of failure occurs within the system, no loss of network service occurs. Even if these have the ability to terminate TCP based sessions, the feature is often not used due to the inherent lack of confidence these systems exude. We would rather have an opportunity to interject some human factors before having a system send a spoofed reset packet that may or may not terminate (with extreme prejudice) a valid attack. What if the CEO was attempting to finalize some kind of transaction using an application that violated a standard network protocol? Is it worth it? "Sometimes a valid business transaction may act like an attack," says Van Nguyen, director of global security at American Presidential Lines, an ocean shipping company in Singapore. He speaks from experience. "In the past, our network-based IDS had flagged our back-up software as a legitimate attack. I definitely would not want my IDS sending TCP resets and blocking traffic automatically!"⁸. The reliability of the system needs to supercede the security feature it provides.

Exposing vulnerabilities through scanning can affect system being tested

While understanding what vulnerabilities exist in a network, keeping up with the changing environment that introduces new variables each day can be a daunting task. Systems need to be in a constant flux to keep up with the new vulnerabilities that are found in hardware platforms, operating systems, service daemons and applications. Valid tests that expose these vulnerabilities can have negative impact on the systems being tested. What if this is done during a crucial time in the business day causing the network to slow down?

Test labs that review the current breed of vulnerability scanners state that most of these products are far from being non-intrusive⁹. All caused adverse reactions on network servers.

An IPS scenario – What happens to the state of the client?

Let's say that an IPS has been installed at a network gateway. A well-formed packet arrives over port 21 requesting an FTP session. The firewall component has been configured to allow traffic on this port. The IDS component has (so far) seen no violation of the FTP protocol (RFC-959) nor has there been any matching signature detected in this session (so far). The FTP server does not currently display any known vulnerabilities that might be exploited by this file transfer (as revealed by the vulnerability scan that was recently performed). So far, productivity has not been adversely affected by the IPS system.

Now, a 100MB file transfer begins. As part of the prevention system, every file that is moved through this gateway must be checked for viruses. It could take many

minutes for this file to be completely transferred. The entire file must be scanned before allowing the transfer to take place. What happens to the client during this time? If a web browser is the interface being used to support the transfer, is some kind of “keep-alive” being used? This is an example of one of the ways an IPS could impede system reliability, and users productivity. If a spoofed address had attempted a denial of service attack, and the IDS component had determined that a reset packet could be sent, isn't it conceptually possible that a valid user is disconnected from the system? Security vendors are wrestling with these problems now on their own discreet devices without the overhead of trying to be good at all of these disparate technologies. What have users seen up to this point in time that would convince them to deploy an IPS in prevention mode?

Issues effecting IPS reliability

Each component of an IPS system faces various challenges to overcome the reliability issues. These problems are compounded when they are aggregated into an IPS device or appliance as a change to one component of the system may have adverse affects on the other components it cohabitates with.

False positives plague most traditional IDSes because, if improperly configured, they will register attacks as legitimate even if those attacks have no bearing on the network. For example, an IDS on a network of Apache Web servers must be told not to register attacks to Microsoft Internet Information Server, otherwise it will issue an alarm when it sees an IIS attack. Similarly, IDSes must be updated with patch information when a flaw is fixed. If it isn't updated, the IDS will set off an alarm if it registers attacks against that flaw, even if the flaw has been patched.

False positives from an IDS are irritating, because they can quickly swamp the network with nearly constant alerts. But they can be downright disastrous from an IPS tool.

What if the system experiences a hardware failure? An IPS sits in-line and if the failure mode is “open” then all traffic will be blocked until the system can be removed or replaced. As all hardware systems have a prescribed mean-time-between-failure (MTBF), it is reasonable to assume that this can happen unless there is a hot failover system standing by.

What about policy changes? Let's say that a company has decided that it now against company policy to use the network for music swapping systems (a reasonable assumption) and this occurs during an employee's time off. When they return to their job and plug in their laptop, the application formerly known as Napster causes the IPS to shutdown connectivity for a particular subnet and everyone else who has complied has now lost connectivity. All of this costs money and effects network reliability.

How about configuration changes? A good vulnerability scanner will uncover vulnerabilities that could cause system failures if attacked. Unless there is some kind of auto discovery going on, a new addition to the network may never be added to the

scanning process and thus be vulnerable to attack. The vulnerability assessment tool must be maintained as any security tool to be effective when installed and more importantly, over time.

How often or accurate are the signature updates that antivirus, firewall, IDS and vulnerability assessment tools use? Errors and omissions can cause significant problems along with the bigger problem of how timely the updates are. What if one component interprets a signature differently than another thus resulting in an unexpected response policy being enforced? How will the discrepancy be resolved (i.e. who wins)? These are battles that are being fought now within companies building best-of-breed devices. By integrating these technologies into discrete devices, the temptation to introduce concessions at each level (due to cost, space, power, etc.) becomes significant.

Defense in depth to enhance reliability

Jason Reed, at SystemExperts, says that having separate IDS and firewall devices from different vendors provides better protection because if one device fails to catch an attack, the others can. Gary Fish, chief executive at FishNet Security, suggested that some level of intrusion detection at the firewall might be good but that a best-of-breed, layered approach to security is more effective. An integrated product suite may cost less but sacrifices functionality. "It may make sense to firewall vendors to do IDS at the firewall. However, I am not sure it will ever replace conventional network and host-based IDS ¹⁰," he said.

Conclusion

Intrusion prevention may be too new and unreliable for widespread acceptance. The Gartner Group came out with a study that says by 2005, IDS systems will be obsolete and that enterprises should not bother with them due to the fact that companies will have successfully hardened their internal systems¹¹. I happen agree with Pete Lindstrom, research director for consulting firm Spire Security, who says that this will only happen, "when hackers stop putting on their thinking caps about new attack techniques, and companies stop making configuration mistakes, and the technology industry stops bringing new technology to market." Not only will IDS system not become obsolete, but Jeff Wilson, executive director of Infonetics Research, said, "though there will be continued growth in 2003, the market really takes off in 2004 due to increased global demand from customers of all sizes, and innovations in technology that make it easier to use, more accurate, and widely available" ¹². What these statistics do not reflect is that people still want to be alerted about an incident so some kind of human analysis can take place before any kind of intervention that might have significant consequences to the reliability or functionality of the network are applied. To deploy an inline intrusion prevention system is to introduce potential latencies and affect quality of service (QoS) metrics. A failure of any one of the components of the IPS whether it is the firewall, vulnerability assessment, intrusion detection, antivirus or threshold monitor could result in the

device presenting a worse case scenario where a denial of service is caused by the system itself.

What would happen if simply exploiting a flaw found in the IPS device could cause multiple points of failure? If predictions come true and this “silver bullet” of an appliance is deployed widely in an enterprise, then the margin of error will need to be infinitesimally small. The staff in the quality assurance department will certainly be earning their money.

Actual intrusion prevention happens where there is real-time response and remediation of systems while an attack is happening, or even before it happens, according to Ron Moritz, senior vice president for Computer Associates International Inc.'s eTrust solutions group. The technology to do that will take at least another three to five years to develop.¹³

With all this uncertainty, turning over the keys to this infant technology is premature. They can't be everywhere to protect everything¹⁴. And like the three little pigs, we should be wary until we have technology that is as similarly robust as a little brick house.

References

1. Messmer, Ellen “Intrusion Prevention raises hopes, concerns”, Network World, 11/04/02
<http://www.nwfusion.com/news/2002/1104prevention.html>
Accessed on: May 7, 2003
2. Snyder, Joel “How vulnerable?”, Information Security, 03/01/03
<http://www.infosecuritymag.com/2003/mar/cover.shtml>
Accessed on July 9, 2003
3. Thaddeus, Jude “Antivirus tool solves only half the problem”, Computerworld, 10/9/00,
<http://www.computerworld.com/securitytopics/security/story/0,10801,52081,00.html>
Accessed on July 15, 2003
4. Securius Newsletter , “Defending the national strategy to secure cyberspace”, Volume 3.03, 11/18/2002,
http://www.securius.com/newsletters/Defending_the_National_Strategy_to_Secure_Cyberspace.html
Accessed on: July 15, 2003
5. Bechtel, Ken “Anti-virus, defence in depth” SecurityFocus, 04/22/03,
<http://www.securityfocus.com/infocus/1687>
Accessed on: July 19, 2003

6. Gong, Dr. Fengmin Deciphering Detection Techniques: Part 3, Denial of Service Detection, IntruVert Networks, 01/2003
Accessed on: May 3, 2003
7. Common Vulnerabilities and Exposures, The MITRE Corporation,
<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CAN-2002-0649>
Accessed on August 10, 2003
8. Cummings, Joanne “From intrusion detection to intrusion prevention”,
Network World, 09/23/2002, <http://www.nwfusion.com/buzz/2002/intruder.html>
Accessed on June 21, 2003
9. Novak, Kevin “VA scanners pinpoint your weak spots”, Network Computing
06/26/2003, <http://www.nwc.com/1412/1412f2.html>
Accessed on August 10, 2003
10. Savage, Marcia “Vendors combine locks to keep out intruders”, IT Week
11/26/2002, <http://www.itweek.co.uk/Analysis/1137119>
Accessed on July 30, 2003
11. Jaques, Robert “Intrusion detection ‘a waste of space’”, IT Week 06/13/2003,
<http://www.itweek.co.uk/News/1141600>
Accessed on August 12, 2003
12. Jaques, Robert “Intrusion detection set for growth”, IT Week 06/06/2003,
<http://www.itweek.co.uk/News/1141436>
Accessed on August 12, 2003
13. Robinson, Brian “Intrusion prevention – how it works”, FCW.com 03/10/2003,
<http://www.fcw.com/supplements/homeland/2003/sup1/hom-prevent1-03-10-03.asp>
Accessed on August 12, 2003
14. Shipley, Greg “Security Watch: Don’t get bitten by the NIPS hype”, Network
Computing 06/13/2003
<http://www.networkcomputing.com/1411/1411colshipley.html> Accessed on
August 17, 2003