Intrusion prevention systems: superior security

Tom Rowan, security consultant, Magirus

Tom Rowan

Today, most networks are protected by firewall technology. There are numerous types of firewall, but essentially they all work in the same way: allow in the authorised traffic, filter the rest. The majority of purebred firewalls do not apply any further filtering on the traffic beyond IP and service port source or destination values. Originally, network security seemed to be as simple as blocking IP addresses and filtering ports.

However, complex modern attacks use any number of attack vectors, including denial of service, protocol implementation flaws, buffer overflows, application development errors, and social engineering techniques such as phishing. These attacks all make use of otherwise legitimate connections through firewalls. To protect against these attacks it is necessary to look deeper into the traffic streams to gain application awareness. Protecting against denial of service attacks launched from one of the many botnets available for hire is also far beyond the remit of a traditional firewall.

As research organisations and security vendors are constantly explaining, a large percentage of attacks originate from inside the network perimeter. Figures between 70% and 95% have been quoted over recent years.[1], [2] Networks are still often built according to out-of-date best practices, dictating a well protected exterior shell with an open, acquiescent interior. This means that the majority of network traffic does not pass across a firewall; even advanced 'deep packet inspection' and 'application aware' firewalls cannot check traffic that does not traverse their interfaces.

The first breed of application aware systems known as intrusion detection systems (IDS) appeared in the mid 1990s. The majority were based on signatures which aimed to match malicious traffic patterns. When a specific pattern was found, the network administrator could be alerted to the presence of malicious traffic on their networks. Initially, these systems reported specious traffic, rather than blocking it.

Many argued that monitoring without taking action was akin to shutting the door once the horse has bolted. Despite these claims, the technology served to increase awareness of complex network security issues amongst the network management community. Like Antony Van Leeuwenhoek's microscope had done for bacteria, IDS showed the characteristics of network attacks as never before[3].

"In about 1998, detection became prevention, and products started to emerge that blocked attacks."

Detection became prevention

The main limitation of IDS was soon removed. In about 1998, detection became prevention, and products started to emerge that blocked attacks. Some existing products were enhanced with blocking capability, while whole new offerings also appeared to take advantage of this new market. Various blocking strategies were employed depending on whether the device was designed to sit 'in line' or stand alone on the network.

In-line intrusion prevention systems are placed so that network traffic must pass through them. When the IPS decides to stop traffic, this has the advantage that blocking actions will be completely effective. However, this approach requires that the network design must force traffic through the device in order to maximise coverage. Placement of the IPS becomes crucial to its effectiveness. Consideration must also be given to the behaviour of the network should the IPS device fail. Many include a 'fail open' relay which turns the device into a piece of wire once power is removed.

Stand-alone systems are arranged so that they gain access to traffic streams from a switch span port or by using a network tap – a piece of hardware which allows the diversion and duplication of traffic at wire speed. There are two strategies employed to give the IPS blocking powers. The IPS can send TCP RST (reset) messages that cause open connections to end suddenly. Or credentials can be supplied which empower the IPS to control firewalls and modify router or switch access control lists to dynamically block traffic – scary stuff!

A third alternative is to provision the IPS a software-only system, which is installed on each host to be protected. In this case, the IPS can block malicious traffic directly. The disadvantage is simply one of implementation. The software must be deployed to multiple hosts, will use up resources when running, may not be compatible with all operating systems, and will require maintenance and upgrades.

Signatures vs rules

Whichever the deployment model, there are a number of ways in which an IPS

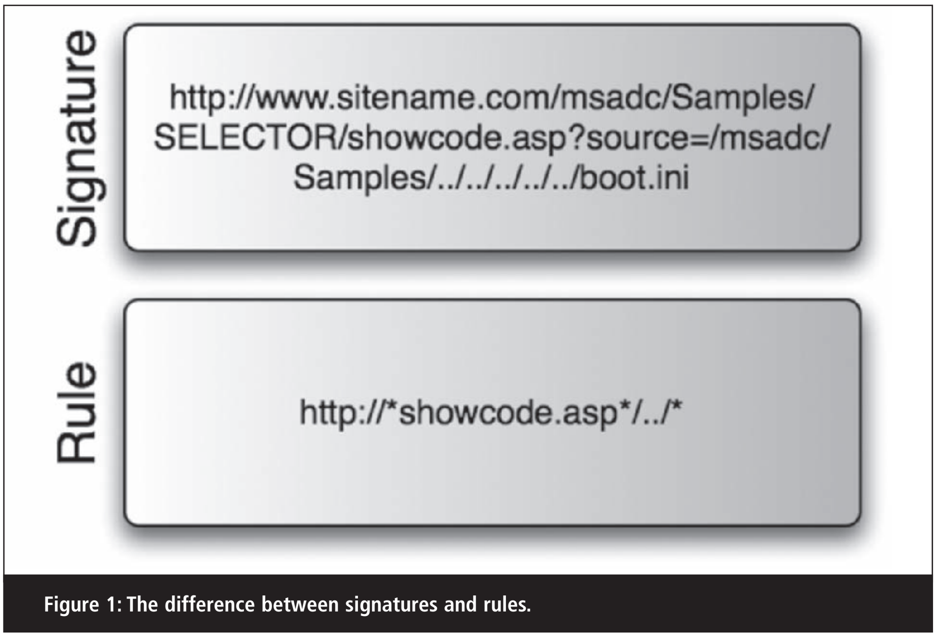
Signature
http://www.sitename.com/msadc/Samples/SELECTOR/showcode.asp?source=/msadc/Samples/../../../../../boot.ini
Rule
http://*showcode.asp*/../*

Figure 1: The difference between signatures and rules.

can detect the presence of unwanted network traffic. These may be grouped into three main categories: Signatures, traffic rates and anomaly detection. Signatures are the most common way to proceed. The majority, although not all, of IPS products include signatures of some kind. These are almost exactly akin to the antivirus signatures familiar to network administrators. A pattern within the network traffic is matched against the shape of a known attack. There are two subdivisions of the term: signatures and rules. These are alike, but fundamentally control the effectiveness of the deployment.

A signature is a pattern that identifies a specific known attack. A rule identifies the use of a known vulnerability. There is a subtle difference which can be understood using this analogy. Consider a glass window. This can be broken in any number of ways, including a hammer, a thrown brick or a bullet. A corresponding signature-based IPS for the window would match a person with a hammer, a brick in flight, and an inbound high-velocity round. The IPS would recognise and block each of these attacks when presented exactly. An attack that uses a fire extinguisher would not be recognised – and would consequently succeed.

> ***"A rule based on the vulnerability rather than the exploit would look for the behaviour that all possible variants of the worm would need to exhibit in order to compromise a system."***

A rule-based IPS, however, would have rules to block against the vulnerability itself. In this case, the vulnerability relates to the fragility of the window to percussive attack. This means that the rule would match any massive inbound object which has sufficient momentum to produce a fracture of the glass. In other words, anything heavy or speedy which is about to hit the window would be blocked.

In a computer network, the signature scenario might relate to the ability of the IPS to block worm variants A and B, but allow variant C straight through. A rule based on the vulnerability rather than the exploit would look for the behaviour that all possible variants of the worm would need to exhibit in order to compromise a system.

Zero day blocked?

Most tier one IPS vendors claim to block the vulnerability not the variant. Confusingly, they mostly discuss signatures rather than rules, but this is just semantic. In reality, when a major new attack evolves, initial signature updates are often issued which fix a single variant, while vendor labs frantically work on a cure to the vulnerability itself. This is released as a replacement to the original update. Interestingly, many vendors produce a specific signature for each variant anyway, allowing network administrators the luxury of a report showing exactly which variant of the attack class has been foiled.

Sometimes, these claims are taken further and vendors offer the ability to stop a zero-day attack. This is somewhat of a holy grail in the IPS world, because if this were true, there would be no need for further signature (or rule) updates. What is meant by these claims is that once a suitable rule is downloaded, any possible future attacks using a specific vulnerability will not succeed – in some senses avoiding zero day attacks, although not all future zero day attacks!

Signatures can also be provided to match any other interesting traffic rather than just attacks. For example, administrators might be interested to be alerted whenever a specific type of traffic is seen on a network: perhaps UNIX portmapper traffic (port 111) on a truly homogenous Windows network. This might not signify an attack – there may be nothing 'wrong' or malicious in the portmapper traffic itself – but may signify mis-configuration or the installation of unauthorised software.

Connection rate limiting

Another way to protect the network which does not rely on signatures is to apply connection rate limits. This method is particularly effective against denial of service (DoS) attacks. This class of attack often uses otherwise perfectly legitimate traffic to flood a network, using all available bandwidth and server resources. Of course, legitimate looking traffic will not be picked up by signatures or rules.

Limiting the number of connection requests allowed into a network will protect against flood attacks by keeping traffic volumes below bandwidth and server resource thresholds. If a host is generating superfluous volumes of traffic, this will be noticed and dropped once it reaches a defined level. However, this simplistic approach may also deny any new legitimate connections too.

This is particularly noticeable in the case of a distributed denial of service attack

(DDoS), where the network flood originates from potentially tens of thousands of IP addresses. In a simple algorithm, there is no mechanism to determine which of these IP addresses is flooding the system and which are producing legitimate traffic. In the worst case, the IPS could actually assist in the denial of service!

The effectiveness of this mechanism can be enhanced by performing in-depth analysis of the traffic. To protect against a DDoS attack, it is necessary to detect and correlate attack packets converging from a disparate array of source IP addresses. For example, if nine hundred out of a thousand incoming connections have the same packet length and TCP checksum, it is likely that these are part of an attack; normal packets don't all look the same. These can be blocked, while allowing the other hundred connections through.

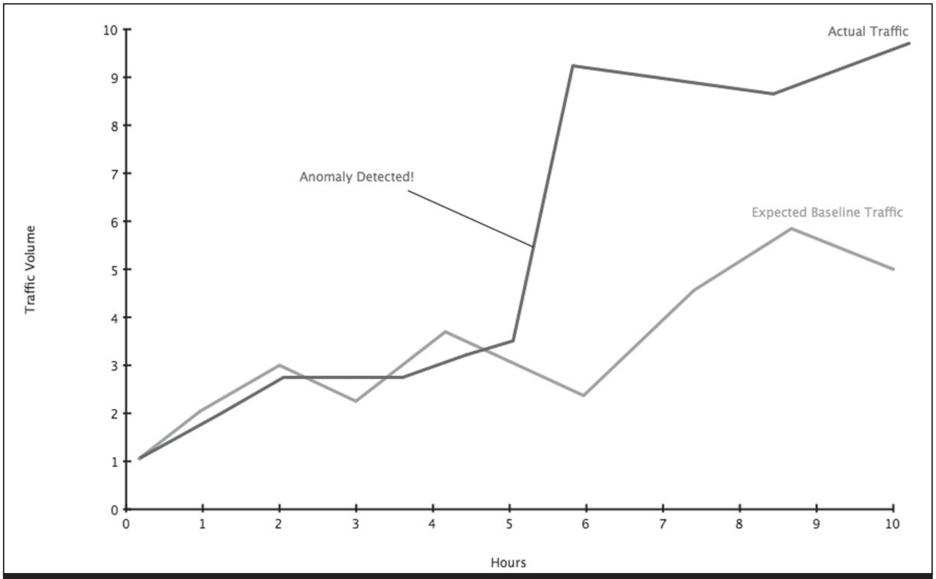

Figure 2: A timeline showing a typical anomaly.

Anomaly detection

Anomaly detection techniques start by first defining a baseline of the network. This is usually done by having the system enter a learning stage for a number of weeks. During this period, the IPS analyses the traffic streams and builds a picture of the normal behaviour of the network. This mechanism works best for green field networks where there is no risk that the network has already been compromised. How to filter out network traffic that is included into the baseline by an already present worm, spyware or other malicious activity? This is done by applying signatures to the traffic being analysed and weeding out anything that will not be acceptable once protection is fully enabled.

"It is vital that the IPS see as much of the network traffic as possible. The location of the system in the network affects this coverage."

Once the baseline has been defined, anomaly detection works by noting any deviation from the norm. As no network behaves exactly as expected at all times and yet nothing untoward is happening, it is important to be liberal about what is defined as deviation. Fuzzy logic is used to dynamically adjust the threshold beyond which activity is deemed anomalous. A traffic spike between two systems that is fairly short in duration is a very different prospect from a prolonged increase in traffic from one host destined to a large number of hosts on the internet. Clever anomaly-based systems know this and can cope with these variations.

Coverage is king

The effectiveness of the deployment of an IPS is often compromised by several factors. It is vital that the IPS see as much of the network traffic as possible. The location of the system in the network affects this coverage. Concentrating the IPS at network ingress points such as firewalls has some advantages – many attacks do come from the internet – but coverage of LAN traffic may suffer. Likewise, placing IPS on a data centre switch would protect servers, but may not pick up attacks involving the internet and workstations directly.

There are a number of IPS vendors whose hardware is capable of protecting multiple network segments at once. An eight-port IPS appliance, for example, might provide protections for four segments running in inline mode. However, the resources of the IPS system are shared amongst these four segments, meaning that it is essential to ensure that the total volume of traffic across them all will be less than the IPS is rated to handle.

When specifying an IPS for a network, it is necessary to estimate the total traffic volumes which will pass across the device. This is particularly important in inline mode. The throughput and latency of the network can be directly affected by a poorly performing inline IPS. Choosing an IPS which is rated above the current traffic volumes will ensure that future requirements will be easily met. Often, powerful appliances are costly, making the danger of under-specification very real.

It is best practice to aim for coverage of the entire network wherever possible. One neat way to achieve this is to use a black hole route. The majority of enterprise networks utilise RFC 1918 'private' addressing schemes[4]. Knowing that a network contains only a small subset of these IP addresses, it is possible to define a set of routes which pass all other private addresses towards the packet sniffing device (usually an IPS). This is like a default route, but only specifies private addresses which are unknown within the network; any traffic seen on these addresses is therefore immediately suspicious.

IPS bypass

Many IPS vendors include a monitor only or bypass mode. This allows the administrator to see what the IPS would have done if protection had been enabled. This allows timid administrators to become comfortable that the device is not going to make the wrong decisions on their network, resulting in blocked legitimate traffic. From experience, this is the main worry that administrators have when deploying an IPS. They are often greatly concerned about false positives, rather than whether the IPS will detect all the malicious traffic on the network.

One problem with the bypass or monitor feature is that it can be used to return the system to a pre-deployment stage whenever a problem occurs. Quite often, problems with legitimate traffic being blocked are due to poor coding in applications, leading to non-RFC compliance. Rather than the wayward code being fixed by the developers, the active protection provided by the offending signature is disabled on the IPS – all in the name of mission critical traffic. If it is mission critical, surely it is worth coding the application correctly?

Sometimes the availability of this feature means that the IPS becomes no more than a tick in an auditor's box, never being turned on for real. Like the firewall installed but implemented with an open 'any accept' policy, an IPS in this mode will allow all traffic through, malicious or otherwise. Unlike that firewall, at least the events will be logged and can be acted upon. Often, however, the improvement to network security is negligible; these are exactly the administrators who would not look at the output anyway!

Don't cry wolf!

When deploying an IPS, it is usual to have to perform some tuning for the first few months. This is necessary to normalise the traffic that the IPS blocks. The amount of tuning depends on a number of factors including the complexity of the network, the quality of vendor signatures and the effectiveness of network baselining exercises for rate limiting and traffic anomaly based systems. The task of tuning is rarely ever wholly finished, as networks change topology, purpose and content organically. The IPS must be adjusted to understand what is acceptable and what is not as the network evolves.

"As we teach children, it is not sensible to erroneously cry wolf too often."

During the initial post-installation tuning phase, an IPS can generate a large number of alerts. Unless a poor choice of default filters has been made by the by the vendor, it is not usual that it will block traffic unduly. However, it is usual to find signatures present which perform an auditing function as well as the discovery of malicious traffic. Examples might include detecting all HTTP POST operations, noting the presence of Skype or Windows Live Messenger traffic, or any number of other common traffic patterns. All of these types of signatures could generate many hundreds or thousands of alerts per day.

As we teach children, it is not sensible to erroneously cry wolf too often[5]. Unfortunately, in the tuning phase, the IPS may cry wolf many, many times, generating a huge array of alerts. These seem interesting at first, but poring over them constantly for a few days will eventually lead to apathy in even the most hardened security professional. False positives and over-alerting make the network administrator's job more onerous. This is unfortunate as most IPS products can be tuned to perform valuable blocking operations almost silently, while reporting only the most relevant information to their masters – it just takes some initial effort to get there.

One effective approach to tuning is to perform it on paper first during planning, before even installing the IPS. Having some idea what traffic is expected on a given segment will allow the tuning to take place quickly after installation; knowing for example that there is no Unix traffic on a segment, but that Windows NetBIOS files sharing is normal, leads to some instant and sweeping tuning decisions. To make best efforts at tuning, it is necessary to understand the network fully and appreciate what kind of traffic is present.

The future is bright

The addition of a security information management system (SIM) to the solution allows the correlation and management of the alerts generated. These alerts are taken from security sources including firewalls, content filters, antivirus and intrusion detection and prevention systems. The normalisation and correlation of these alerts means that, for example, a logon failure detected by a firewall and an IPS can be compared natively. This leads to reports which are more meaningful; they can include only the most important alerts.

Many large organisations now rely upon IPS to protect their networks alongside firewalls and content filtering solutions. A number of IPS vendors now market their products as providing 'network patching' as a stop gap until security patches are applied. Even the marketing of these systems makes it clear that administrators still need to apply updates to their servers eventually.

The future is always difficult to predict within the field of information security. Given that the challenge of patch management seems only to grow with time, it is likely that more focus will be given to specific network-based patching technology. This will have to move beyond the remarketing of IPS and into new ground. The technology that will improve or may even overtake IPS already exists today: inline virtual patching. This works by editing network streams to make the traffic which hits the server work as if patched; the traffic is patched, the server is not. This differs from the IPS approach by ensuring that no traffic is ever stopped – suspicious traffic is cured, not blocked.

In a modern organisation, IPS has won a strong place amongst the ubiquitous firewalls, content filters and anti virus platforms. Microscopic analysis of network traffic patterns combined with rules which detect probes against vulnerabilities mean that a mature system can give great benefit. After a planned implementation which aims to give greatest network coverage, careful tuning is usually necessary to ensure that alerts do not overwhelm administrators. Once implemented, IPS provides a value contribution to a strong network security stance.

References

1. N. Stanley, "ArcSight and Insider (or Inside?) Threat Management". October 2006 http://www.arcsight.com/articles/Bloor-Research-Insider-or-Inside-Threat.pdf
2. C. Potter et al, "DTI Information Security Breaches Survey 2006", April 2006, www.pwc.com/extweb/pwcpublications.nsf/docid/7FA80D2B30A116D7802570B9005C3D16
3. http://en.wikipedia.org/wiki/Anton_van_Leeuwenhoek
4. R. Fielding et al, "RFC 1918." Internet EFC.STD.FYI/BCP archives. June 1996. www.faqs.org/rfcs/rfc1918.html
5. http://en.wikipedia.org/wiki/Aesop's_Fables

About the author

Tom Rowan is the lead security consultant for the Magirus UK security division. He works with partners to enable them to provide best of breed security solutions to their clients. He is a 'hands on' technical consultant with over eleven years of security experience. He is equally happy installing a firewall or penetration testing a client network. In his spare time, Tom is also a RAF Reserve Officer with the Air Training Corps.

The mechanics of Vipul's Razor technology

Jamie De Guerre, chief technology officer, Cloudmark

Much has been said about the wisdom of crowds. The idea that many people can achieve results more effectively than individuals has gained credence, especially as the internet has bought those people together and allowed them to co-operate in innovative ways. These techniques apply to everything from online encyclopedias to citizen journalism – and even anti-spam technologies.

Vipul's Razor, developed by Vipul Ved Prakash in 1998 as an open source product, was the world's first collaborative spam filtration network for messaging security. Collaborative human intelligence 'identifies' a message as spam, and an automated technology verifies and prevents its proliferation[1].

Collaborative spam filtering works by allowing members of a community to identify and vote on messages. The reputation of the members is continually rated and can be coupled with an automated system of message fingerprinting. The result is a system that can be used to detect spam, email-borne viruses, and phishing threats.

Vipul's Razor provides a framework for changes to fingerprinting algorithms and adjustments to security protocols, enabling easy system adaptation to evolving or new threats. Unlike rules or heuristic-based schemes that require continuous updates and entail heavy processing, fingerprinting algorithms automatically generate lightweight fingerprints that accurately identify messaging abuses and their variants. In addition, additional algorithms to combat new classes of threats can be easily integrated without the need for changing the system architecture.

> ***"Collaborative human intelligence 'identifies' a message as spam, and an automated technology verifies and prevents its proliferation."***

Developing a global network of trusted users

Vipul's Razor began the process of developing a global network of trusted users now characterised by years of institutional anti-abuse learning. The network of trusted users, which initially numbered in the tens of thousands with the Razor community, has expanded tremendously through the rapid adoption of Cloudmark's commercial solutions by service providers, enterprises, and consumers. The network now encompasses over 180 million sources in 163 countries, including highly sophisticated reporters such as service provider abuse teams and systems administrators in addition to trusted honeypots and end users. The size and geographic scope of the network is a key factor in the collaborative process that leads to higher accuracy.

Vipul's innovation has proven that individual users can accurately distinguish between spam and legitimate email early in the lifecycle of a threat. It also began the building of a large pool or community of self-organising email readers who all receive the same unwanted messages and decide as a group by individually nominating messages as spam or not spam. Vipul's Razor further proved that this community approach, with its collaborative decision making, dramatically improved blocking accuracy.

Most importantly, this process of collective decision making reduces the training and learning curve of the spam filter, thereby reducing message misclassifications and their associated costs. One of Ved Prakash's first principles in Vipul's Razor was the preservation of legitimate communications.