Cisco IPS-4255 V5.0(3)

Technical Evaluation

An NSS Group Report

First published July 2005 (Version 1.0)

Published by The NSS Group
Security Testing Laboratories
Mas la Carrière, Route de Ganges
30440 Sumène, France

Tel : +33 (0)4 67 81 49 11
E-mail : info@nss.co.uk
Internet : http://www.nss.co.uk

TABLE OF CONTENTS

TABLE OF FIGURES

The NSS Group

The NSS Group is the world's foremost independent security testing facility.

With British headquarters, and security and network infrastructure testing facilities in the South of France, The NSS Group offers a range of specialist IT, networking and security-related services to vendors and end-user organisations world-wide.

The NSS Group's Security Testing Laboratories are available to vendors and end-users for fully independent testing of networking, communications and security hardware and software.

The NSS Group also operates certification schemes for vendors and certification bodies, and currently provides evaluation and certification of a wide range of security products, including IDS/IPS appliances, firewalls, VPNs, Web Application firewalls, multi-function security appliances, cryptographic devices and PKI products.

Output from the labs, including detailed research reports, articles and white papers on the latest network and security technologies, are made available on the NSS web site at http://www.nss.co.uk.

The NSS Group awards are recognised world-wide as being the most desirable and essential when it comes to security products. Vendors consider the awards to be a crucial step in any security-related marketing campaign, whilst feedback from readers of the reports indicates that participation in an NSS Group test and/or one of the **NSS Approved** awards is a prerequisite for any security product in order to be considered for purchase.

Foreword

Following the huge success of the first comprehensive *Intrusion Prevention System* (IPS) test of its kind, The NSS Group is pleased to present the results of its third IPS Group Test, the largest so far, which includes a number of new products not included in the first two reports.

As with the first two Editions, this exhaustive review will give readers a complete perspective of the capabilities, maturity and suitability for immediate deployment of each of the products tested. The NSS Group established this test as IPS products are being actively deployed as a new layer in defence-in-depth security architectures.

The NSS IPS Group Test evaluates the performance, reliability, security effectiveness, and usability of Network IPS products. The test consists of seven sections within three primary areas: *performance and reliability, security accuracy,* and *usability*.

Overall, the brand new test suite contains over **800 individual tests**, many of which are run multiple times, to provide the most thorough and complete evaluation of IPS products available anywhere today. The NSS Group has developed advanced testing methodologies for both *Rate-Based IPS* and *Content-Based IPS* products, since these devices are often very different in operation, although all products tested in this edition of the report are content-based.

It is worth pointing out that not every product submitted for testing receives an *NSS Approved* award. Standards are very high, and only those appearing in this report have received ***NSS Approved*** awards. For this latest edition, **ten** vendors submitted a total of **twelve** products for testing, and **eight** of these passed our stringent testing to receive **NSS Approved**. It is heartening to note that this is a much-improved success ratio over Edition 2.

We believe that our IPS test methodologies - which have been updated again for this test - will become the *de facto* standard for testing in-line Intrusion Prevention/Attack Mitigation devices, and the *NSS Approved* logo an essential item on the list of requirements when purchasing these products.

We also believe that this report is essential reading for anyone considering deploying Intrusion Prevention Systems in their networks, either in a test or live situation, and we hope that you find it both informative and useful in making your purchasing decisions. The latest **IPS Group Test** report can be viewed on-line at *www.nss.co.uk/ips*

Bob Walder

INTRODUCTION

In a survey commissioned by VanDyke Software, some 66 per cent of the companies who responded said that they perceive system penetration to be the largest threat to their enterprises.

The survey revealed that the top eight threats experienced by those surveyed were *viruses* (78 per cent of respondents), *system penetration* (50 per cent), *DoS* (40 per cent), *insider abuse* (29 per cent), *spoofing* (28 per cent), *data/network sabotage* (20 per cent), and *unauthorised insider access* (16 per cent).

Although 86 per cent of respondents use firewalls (a disturbingly **low** figure in this day and age, to be honest!), it is apparent that firewalls are not always effective against many intrusion attempts. The average firewall is designed to deny clearly suspicious traffic - such as an attempt to telnet to a device when corporate security policy forbids telnet access completely - but is also designed to allow some traffic through - Web traffic to an internal Web server, for example.

The problem is, that many exploits attempt to take advantage of weaknesses in the very protocols that **are** allowed through our perimeter firewalls, and once the Web server has been compromised, this can often be used as a springboard to launch additional attacks on other internal servers. Once a “rootkit” or “back door” has been installed on a server, the hacker has ensured that he will have unfettered access to that machine at any point in the future.

Firewalls are also typically employed only at the network perimeter. However, many attacks, intentional or otherwise, are launched from within an organisation. Virtual private networks, laptops, and wireless networks all provide access to the internal network that often bypasses the firewall. Intrusion detection systems may be effective at detecting suspicious activity, but do not provide *protection* against attacks. Recent worms such as Slammer and Blaster have such fast propagation speeds that by the time an alert is generated, the damage is done and spreading fast.

Intrusion Prevention Systems (IPS)

The inadequacies inherent in current defences has driven the development of a new breed of security products known as *Intrusion Prevention Systems* (IPS). This is a term which has provoked some controversy in the industry since some firewall and IDS vendors think it has been “hijacked” and used as a marketing term rather than as a description for any kind of new technology.

Whilst it is true that firewalls, routers, IDS devices and even AV gateways all have intrusion prevention technology included in some form, we believe that there are sufficient grounds to create a new market sector for true *Intrusion Prevention Systems*.

These systems are proactive defence mechanisms designed to detect malicious packets within normal network traffic (something that the current breed of firewalls do not actually do, for example) and stop intrusions dead, blocking the offending traffic automatically before it does any damage rather than simply raising an alert as, or after, the malicious payload has been delivered.

Within the IPS market place, there are two main categories of product: *Host IPS* and *Network IPS*, with the latter being further sub-divided into *Content-Based* and *Rate-Based* (or *Attack Mitigation*) systems.

Host IPS (HIPS)

As with Host IDS systems, the Host IPS relies on agents installed directly on the system being protected. It binds closely with the operating system kernel and services, monitoring and intercepting system calls to the kernel or APIs in order to prevent attacks as well as log them.

It may also monitor data streams and the environment specific to a particular application (file locations and Registry settings for a Web server, for example) in order to protect that application from generic attacks for which no “signature” yet exists.

One potential disadvantage with this approach is that, given the necessarily tight integration with the host operating system, future OS upgrades could cause problems.

Since a Host IPS agent intercepts all requests to the system it protects, it has certain prerequisites - it must be very reliable, must not negatively impact performance, and must not block legitimate traffic. Any HIPS that does not meet these minimum requirements should never be installed in a host, no matter how effectively it blocks attacks.

Network IPS (NIPS)

The Network IPS combines features of a standard IDS, an IPS and a firewall, and is sometimes known as an *In-line IDS* or *Gateway IDS (GIDS)*. The next-generation firewall - the *deep inspection firewall* - also exhibits a similar feature set, though we do not believe that the deep inspection firewall is ready for mainstream deployment just yet.

As with a typical firewall, the NIPS has at least two network interfaces, one designated as *internal* and one as *external*. As packets appear at either interface they are passed to the detection engine, at which point the IPS device functions much as any IDS would in determining whether or not the packet being examined poses a threat.

However, if it should detect malicious traffic, in addition to raising an alert, it will discard the packet(s) and mark that flow as bad. As the remaining packets that make up that particular TCP session arrive at the IPS device, they are discarded immediately.

Legitimate packets are passed through to the second interface and on to their intended destination. A useful side effect of some NIPS products is that as a matter of course - in fact as part of the initial detection process - they will provide “*packet scrubbing*” functionality to remove protocol inconsistencies resulting from varying interpretations of the TCP/IP specification (or intentional packet manipulation).

Thus any fragmented packets, out-of-order packets, or packets with overlapping IP fragments will be re-ordered and “cleaned up” before being passed to the destination host, and illegal packets can be dropped completely.

One thing to watch out for - don't let the "reactive" IDS vendors kid you into believing that they have *intrusion prevention* capabilities just because they can send TCP reset commands or re-configure a firewall when they detect an attack (a worrying piece of FUD that we have noticed in some IDS marketing literature recently).

The problem here is that unless the attacker is operating on a 2400 baud modem, the likelihood is that by the time the IDS has detected the offending packet, raised an alert, and transmitted the TCP Resets - and especially by the time the two ends of the connection have received the Reset packets and acted on them (or the firewall or router has had time to activate new rules to block the remainder of the flow) - the payload of the exploit has long since been delivered….. *game over*! Our guess is that there are not many crackers using 2400 baud modems these days….

A true IPS device, however, is sitting in-line - **all** the packets have to pass through it. Therefore, as soon as a suspicious packet has been detected - and **before** it is passed to the internal interface and on to the protected network, it can be dropped. Not only that, but now that flow has been flagged as suspicious, **all** subsequent packets that are part of that session can also be dropped with very little additional processing. Oh, and for good measure, some products are also capable of sending *TCP Resets* or *ICMP Unreachable* messages to the attacking host.

Rate-Based IPS (Attack Mitigator)

Most NIPS products are basically IDS engines that operate in-line, and are thus dependent on protocol analysis or signature matching to recognise malicious content within individual packets (or across groups of packets). These can be classed as *Content-Based IPS* systems.

There is, however, a second breed of Network IPS that ignores packet content almost completely, instead monitoring for anomalies in network traffic that might characterise a flood attempt, scan attempt, and so on. These devices are capable of monitoring traffic flows in order to determine what is considered "normal", and applying various techniques to determine when that traffic deviates from normal. This is not always as simple as watching for high-volumes of a specific type of traffic in a short space of time, since they must also be capable of detecting "stealth "attacks, such as low-rate connection floods and slow port scan attempts.

Since these devices are concerned more with anomalies in traffic flow than packet contents, they are classed as *Rate-Based IPS* systems - and are also known as *Attack Mitigators*, as they are so effective against DOS and DDOS attacks.

Detection Methods

At one time, most Network IDS/IPS products based their alerts purely on pattern matching packet contents against a database of known signatures. Then came a new breed of offerings that approached the problem in a completely different way - by doing a full protocol analysis on the data stream. Others began to use heuristics or anomaly-based analysis to determine when an attempted attack had taken place.

Today, most IDS/IPS employ a mixture of these detection methods in a single product, though some will be more biased towards one method than another.

According to Cisco, there are five main methods of attack identification (*source: Cisco Systems, The Science of Intrusion Detection System Attack Identification*):

Pattern Matching

Pattern matching in its most basic form is concerned with the identification of a fixed sequence of bytes in a single packet. In addition to the tell-tale byte sequence, most IPS will also match various combinations of the source and destination IP address or network, source and destination port or service, and the protocol. It is also often possible to tune the signature further by specifying a start and end point for inspection within the packet, or a particular combination of TCP flags.

The more specific these parameters can be, the less inspection needs to be carried out against each packet on the wire. However, this approach can make it more difficult for systems to deal with protocols that do not live on well defined ports and, in particular, Trojans, and their associated traffic, which can usually be moved at will.

Although it is often quite simple to define a signature for a particular exploit, basic pattern matching can often be too specific, sometimes requiring multiple signatures to be defined for minor variations in exploits. They are also prone to false positives, since legitimate traffic can often contain the relatively small set of criteria supposedly used to determine when an attack is taking place.

This method is usually limited to inspection of a single packet and, therefore, does not apply well to the stream-based nature of network traffic such as HTTP sessions. This limitation gives rise to easily implemented evasion techniques.

Stateful Pattern Matching

Stateful pattern matching offers a slightly more sophisticated approach, since it takes the context of the established session into account, rather than basing its analysis on a single packet.

Stateful IPS products must consider arrival order of packets in a TCP stream and should handle matching patterns across packet boundaries. Thus, if the exploit string to be matched is *foobar*, and the exploit is split across two packets, with *foo* in one and *bar* in another, the simple packet matching IPS will miss the attack, since it will not be able to match the complete string. The stateful IPS, however, will maintain the session context and reassemble the traffic stream, once again making the complete string available to the detection engine.

This requires more resources than simple pattern matching, since the IPS now has to allocate large amounts of memory and processing power to track a potentially large number of open sessions for as long as possible. This approach does make IPS evasion that much more difficult, though far from impossible.

Direction of traffic is also important here, both in terms of quality of detection and performance.

Client-to-server traffic inspection is the process of applying detection mechanisms to the “request side” portion of a communication - for example, in HTTP this could be the “GET” request coming from a client.

Client-to-server traffic inspection is typically activated to protect all traffic whether internally or externally generated. As the size of the traffic in terms of byte count is relatively small, the processing load placed on the IPS will be lower.

Server-to-client traffic inspection is the process of finding an attack in the "response side" portion of a communication - for example, in HTTP the server-to-client traffic could be the web page and content returned from the server as a result of a "GET" request. Server-to-client traffic, as in this example, is often much larger than the client-to-server traffic in terms of byte count. As a result, the processing load that is placed on an IPS is greater for server-to-client traffic.

Some vendors do not implement server-to-client signatures at all. Often this is for performance reasons, but sometimes it is a design decision by those vendors who also offer HIPS products, which are often better placed to detect the types of exploits executed by malicious response traffic as opposed to request traffic. Some vendors do include server-to-client signatures, but recommend they are disabled when performance is paramount. Bi-directional detection can have a significant impact on performance in some cases - those products which can handle this situation with zero or minimal impact on performance are worth closer inspection (although this level of performance often comes with a higher price tag).

It should be noted that there are situations where disabling server-to-client signatures is reasonably safe, and - happily - these are usually the situations where the highest levels of performance are demanded. Typically, this would be where an IPS is deployed within the network perimeter, where it is unlikely that purely internal HTTP response traffic is likely to be malicious. Perimeter defences would normally be deployed with both client-to-server and server-to-client signatures enabled, but perimeter devices rarely have the same performance requirements as internal ones.

Protocol Decode

Protocol decode IPS take a radically different approach to simple pattern matching IPS products - though sometimes not quite as radically different as the marketing folks would have you believe. With this technique, the IPS detection engine performs a full protocol analysis, decoding and processing the packet contents in the same way that the target client or server application would. It also tends to be stateful.

Although this may seem like using a sledgehammer to crack a nut, it does have the advantage of highlighting anomalies in packet contents much more quickly than doing an exhaustive search of a signature database. It also has the advantage of greater flexibility in capturing attacks that would be very difficult - if not impossible - to catch using pure pattern-matching techniques, as well as new variations of old attacks. These are attacks which - although changing only slightly from variant to variant - would normally require a new signature in the database for the "traditional" IPS architecture, but which would be detected automatically by a complete protocol analysis.

One of the first things the protocol decode engine does is to apply rules defined by the appropriate RFCs to look for violations. This can help to detect certain anomalies such as binary data in an HTTP request, or a suspiciously long piece of data where it should not be - a sign of a possible buffer overflow attempt.

One simple example of how this might work concerns searching Telnet login strings for one of the many well-known login names that rootkits tend to leave behind on the system. A pattern matching system might scan *all* Telnet traffic for *all* these patterns, in which case the more patterns you add, the slower it becomes (not *always* the case, but a reasonable assumption for the purposes of this example).

In contrast, a protocol analysis system will decode the Telnet protocol and extract the login name. It can then perform an efficient search in a binary-search tree or a hash table for just the login name, which should scale much better as new signatures are added.

In theory, therefore, protocol decoding should offer more efficient processing of traffic and improved scalability as more signatures are added, compared to a pure pattern matching solution. In reality, pattern matching solutions rarely opt for a “brute force” approach (there are some extremely intelligent and efficient pattern matching mechanisms available), and so the differences are not always as marked as the marketing people would like us to believe.

Note also, that pattern matching and protocol decoding are not mutually exclusive, as some would lead you to believe. A protocol analysis IPS can only go so far with its protocol decodes before it too will be forced to perform some kind of pattern matching, albeit against a theoretically smaller subset of “signatures”.

One major downside, of course, is that if a completely new type of exploit does surface, it is likely that the developer will have to create new protocol decode code to handle it, whereas the pattern matching approach can allow the administrator to develop a custom signature much more quickly on site.

Protocol decoding does offer a number of advantages, however. It minimises the chance for false positives if the protocol is well defined and enforced (although false positives can be higher if the RFC is ambiguous), and can also be more broad and general to allow the IPS to detect minor variations of an exploit without having to implement separate signatures.

You may see this technique referred to in several different ways:

- *Protocol decode*
- *Protocol Anomaly Detection*
- *Protocol validation*

Each of these terms, if strictly applied, could use a slightly different approach to the problem. For example, we would expect a *protocol decode* engine to perform the sort of additional pattern matching and length checking mentioned above on the field contents in order to detect specific exploits or buffer overflows.

Pure *protocol validation* or *Protocol Anomaly Detection* engines, however, might go no further than decoding just enough to be able to determine if the packet follows the RFC to the letter. If not, they will raise an alert - but in allowing a packet to pass, they cannot be sure that the contents will not contain a means of exploit that just happens to conform with the RFC.

Beware the marketing hype in this particular area – no matter what architecture is used, the performance figures and detection rates in a live deployment will speak for themselves.

Heuristic Analysis

Heuristic-based signatures use some kind of algorithmic logic on which to base their alarm decisions. These algorithms are often statistical evaluations of the type of traffic being presented.

A good example of this type of signature is one that would be used to detect a port sweep. This signature looks for the presence of a threshold number of unique ports being touched on a particular machine. The signature may further restrict itself through the specification of the types of packets that it is interested in (that is, SYN packets). Additionally, there may be a requirement that all the probes must originate from a single source, and even that valid SYN ACK packets must be seen to be returned by the host being probed.

Signatures of this type will react differently on different networks, and can be a significant source of false positives if not tuned correctly, requiring some threshold manipulations to make them conform to the utilisation patterns on the network they are monitoring. This type of signature may be used to look for very complex relationships as well as the simple statistical example given.

Anomaly Analysis

The final approach is to forget about trying to identify the attacks directly, and concentrate instead on ignoring everything that is considered “normal”. This is known as “*anomaly-based*” IPS, and the basic principle is that, having identified what could be considered “normal” traffic on a network, then anything that falls outside those bounds could be considered an “intrusion” - or at the very least, something worthy of note. This is generally better suited to passive IDS rather than in-line IPS devices, given its propensity for false positives.

The primary strength of anomaly detection is its ability to recognise previously unseen attacks, since it is no longer concerned with knowing what an attack looks like - merely with knowing what does not constitute normal traffic. Its drawbacks, of course, include the necessity of training the system to separate noise from natural changes in normal network traffic (the installation of a new - perfectly legitimate - application somewhere on the network, for example).

Changes in standard operations may cause false alarms while intrusive activities that appear to be normal may cause missed detections. It is also difficult for these systems to name types of attacks, and this technology has a long way to go before it could be considered ready for “prime time”.

Which Detection Method Is The Best?

Which detection method to choose is a difficult question, and in all honesty, it is not one with which most of those evaluating these products should concern themselves.

Adequate performance to handle the traffic to which the sensor will be exposed, accuracy of alerts, low incidence of false positives, and centralised management and reporting/analysis tools are far more important than how the packets are processed.

In some instances, the lines blur between methodologies to the point where they become almost indistinguishable.

For example, most protocol decode analysis engines alert the user to the presence of protocol violations that are not directly related to any known attack but are “anomalous” (for example, length-based buffer overflow detection). Therefore, in this instance the engine has attributes of an anomaly-based system.

As we have already mentioned, most protocol analysis systems are also reduced to performing some form of pattern-matching process following the protocol decode. Likewise, even the most basic pattern-matching systems perform some form of protocol analysis - even if it is only for a limited range of protocols. In truth, almost all Network IPS systems are already adopting a hybrid architecture.

By and large, therefore, the *pattern-matching vs. protocol decode* debate is one of religion - something for the marketing departments to shout about. Why should the average user care what happens under the hood as long as the product does what it claims to do - detect and prevent intrusions?

Implementation Challenges

There are a number of challenges to the implementation of an IPS device that do not have to be faced when deploying passive-mode IDS products. These challenges all stem from the fact that the IPS device is designed to work in-line, presenting a potential choke point and single point of failure.

If a passive IDS fails, the worst that can happen is that some attempted attacks may go undetected. If an in-line device fails, however, it can seriously impact the performance of the network.

Perhaps latency rises to unacceptable values, or perhaps the device fails closed, in which case you have a self-inflicted Denial of Service condition on your hands. On the bright side, there will be no attacks getting through! But that is of little consolation if none of your customers can reach your e-commerce site.

Even if the IPS device does not fail altogether, it still has the potential to act as a bottleneck, increasing latency and reducing throughput as it struggles to keep up with up to a Gigabit or more of network traffic. Devices using off-the-shelf hardware will certainly struggle to keep up with a heavily loaded Gigabit network, especially if there is a substantial signature set loaded, and this could be a major concern for both the network administrator - who could see his carefully crafted network response times go through the roof when a poorly designed IPS device is placed in-line - as well as the security administrator, who will have to fight tooth and nail to have the network administrator allow him to place this unknown quantity amongst his high performance routers and switches.

As an integral element of the network fabric, the Network IPS device must perform much like a network switch. It must meet stringent network performance and reliability requirements as a prerequisite to deployment, since very few customers are willing to sacrifice network performance and reliability for security. A NIPS that slows down traffic, stops good traffic, or crashes the network is of little use.

Dropped packets are also an issue, since if even one of those dropped packets is one of those used in the exploit data stream it is possible that the entire exploit could be missed.

Most high-end IPS vendors will get around this problem by using custom hardware, populated with advanced FPGAs and ASICs - indeed, it is necessary to design the product to operate as much as a switch as an intrusion detection and prevention device.

It is very difficult for any security administrator to be able to characterise the traffic on his network with a high degree of accuracy. What is the average bandwidth? What are the peaks? Is the traffic mainly one protocol or a mix? What is the average packet size and level of new connections established every second - both critical parameters that can have detrimental effects on some IDS/IPS engines? If your IPS hardware is operating “on the edge”, all of these are questions that need to be answered as accurately as possible in order to prevent performance degradation.

Another potential problem is the good old *false positive*. The bane of the security administrator’s life (apart from the script kiddie, of course!), the false positive rears its ugly head when an exploit signature is not crafted carefully enough, such that legitimate traffic can cause it to fire accidentally. Whilst merely annoying in a passive IDS device, consuming time and effort on the part of the security administrator, the results can be far more serious and far reaching in an in-line IPS appliance.

Once again, the result is a self-inflicted Denial of Service condition, as the IPS device first drops the “offending” packet, and then potentially blocks the entire data flow from the suspected hacker. If the traffic that triggered the false positive alert was part of a customer order, you can bet that the customer will not wait around for long as his entire session is torn down and all subsequent attempts to reconnect to your e-commerce site (if he decides to bother retrying at all, that is) are blocked by the well-meaning IPS.

Another potential problem with any Gigabit IPS/IDS product is, by its very nature and capabilities, the amount of alert data it is likely to generate. On such a busy network, how many alerts will be generated in one working day? Or even one hour? Even with relatively low alert rates of ten per second, you are talking about 36,000 alerts every hour. That is 864,000 alerts each and every day. The ability to tune the signature set accurately is essential in order to keep the number of alerts to an absolute minimum. Once the alerts have been raised, however, it then becomes essential to be able to process them effectively. Advanced alert handling and forensic analysis capabilities - including detailed exploit information and the ability to examine packet contents and data streams - can make or break a Gigabit IDS/IPS product.

Of course, one point in favour of IPS when compared with IDS is that because it is designed to prevent the attacks rather than just detect and log them, the burden of examining and investigating the alerts - and especially the problem of rectifying damage done by successful exploits - is reduced considerably.

Requirements for effective prevention

Having pointed out the potential pitfalls facing anyone deploying these devices, what features are we looking for that will help us to avoid such problems?

- **In-line operation -** only by operating in-line can an IPS device perform true protection, discarding all suspect packets immediately and blocking the remainder of that flow

- **Reliability and availability -** should an in-line device fail, it has the potential to close a vital network path and thus, once again, cause a DoS condition. An extremely low failure rate is thus very important in order to maximise up-time, and if the worst should happen, the device should provide the option to fail open or support fail-over to another sensor operating in a fail-over group (see below). In addition, to reduce downtime for signature and protocol coverage updates, an IPS must support the ability to receive these updates without requiring a device re-boot. When operating inline, sensors rebooting across the enterprise effectively translate into network downtime for the duration of the reboot
- **Resilience -** as mentioned above, the very minimum that an IPS device should offer in the way of High Availability is to fail open in the case of system failure or power loss (some environments may prefer this default condition to be “fail closed” as with a typical firewall, however - the most flexible products will allow this to be user-configurable). Active-Active stateful fail-over with cooperating in-line sensors in a fail-over group will ensure that the IPS device does not become a single point of failure in a critical network deployment
- **Low latency -** when a device is placed in-line, it is essential that its impact on overall network performance is minimal. Packets should be processed quickly enough such that the overall latency of the device is as close as possible to that offered by a layer 2/3 device such as a switch, and no more than a typical layer 4 device such as a firewall or load-balancer.
- **High performance -** packet processing rates must be at the rated speed of the device under real-life traffic conditions, and the device must meet the stated performance with all signatures enabled. Headroom should be built into the performance capabilities to enable the device to handle any increases in size of signature packs that may occur over the next three years. Ideally, the detection engine should be designed in such a way that the number “signatures” (or “checks”) loaded does not affect the overall performance of the device.
- **Unquestionable detection accuracy -** it is imperative that the quality of the signatures is beyond question, since false positives can lead to a Denial of Service condition. The user MUST be able to trust that the IDS is blocking only the user selected malicious traffic. New signatures should be made available on a regular basis, and applying them should be quick (applied to all sensors in one operation via a central console) and seamless (no sensor reboot required)
- **Fine-grained granularity and control -** fine grained granularity is required in terms of deciding exactly which malicious traffic is blocked. The ability to specify traffic to be blocked by attack, by policy, or right down to individual host level is vital. In addition, it may be necessary to only alert on suspicious traffic for further analysis and investigation
- **Advanced alert handling and forensic analysis capabilities -** once the alerts have been raised at the sensor and passed to a central console, someone has to examine them, correlate them where necessary, investigate them, and eventually decide on an action. The capabilities offered by the console in terms of alert viewing (real time and historic) and reporting are key in determining the effectiveness of the IPS product.

The NSS Intrusion Prevention Group Test

The NSS Group conducted the first comprehensive IPS test of its kind, now updated in this Edition.

This exhaustive review will give readers a complete perspective of the capabilities, maturity and suitability of the products tested for their particular needs.

As part of its extensive IPS/Attack Mitigator test methodologies (see section on *Testing Methodology* later in this report for detailed methodologies, updated for this latest test) The NSS Group subjects each product to a brutal battery of tests that verify the stability and performance of each IPS tested, determine the accuracy of its security coverage, and ensure that the device will not block legitimate traffic.

If a particular IPS has been designated as *NSS Approved*, customers can be confident that the device will not significantly impact network/host performance, cause network/host crashes, or otherwise block legitimate traffic.

To assess the complex matrix of IPS/Attack Mitigator performance and security requirements, the NSS Group has developed a specialised lab environment that is able to exercise every facet of an IPS product. The test suite contains over 800 individual tests that evaluate IPS products in three main areas: *performance and reliability, security accuracy,* and *usability*.

This thorough review should give readers a complete perspective of the capabilities, maturity and suitability of the products tested for their particular needs.

Performance

Any IPS is expected to be reliable (not crash), to never block legitimate traffic, and to not unduly affect network or host system performance.

The latency and throughput of a Network IPS (NIPS) or Attack Mitigation device must be on a par with other equipment in the network on which it is deployed, and in this respect, an in-line NIPS must strive to perform much more like a switch than a typical passive security device, especially when it is necessary to install more than one NIPS in the same data path.

Detection/Blocking Performance Under Load

This group of tests verifies that the IPS does not adversely impact legitimate traffic, even when new TCP connections are being created rapidly. We also verify that the sensor is capable of detecting and blocking exploits when subjected to increasing loads of background traffic up to the maximum bandwidth supported as claimed by the vendor. An IPS that misses attacks under load can be evaded. An IPS that adversely affects legitimate background traffic will not stay in-line for long.

A fixed number of exploits are launched with zero background traffic to ensure the sensor is capable of detecting our baseline attacks. Once that has been established, increasing levels of varying types of background traffic are generated **through** the IPS device in order to determine the point at which the sensor begins to miss attacks.

All tests are repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic (or up to the maximum rated throughput of the device in 25 per cent increments should this be less than 1Gbps). The test is conducted with UDP, HTTP, and mixed-protocol traffic and includes packet rates up to 453,000 packets per second and connection rates up to 20,000 connections per second.

Latency & User Response Times

In any network environment latency is important. Latency may impose an upper bound on throughput and it also has an impact on interactive applications, thus affecting user response time. As such, it is important to understand the impact of latency introduced by a NIPS and to determine the maximum acceptable delay, which will be different for each network.

There is a direct relationship between latency introduced by a networking device and the maximum throughput allowed by that device on a single TCP connection. There is a critical value for the *round trip time* (RTT) of a packet in each network, and if the latency is below this critical value, TCP throughput will be unaffected - instead, it is the line speed of the underlying network which becomes the bottleneck. Above this critical value, however, TCP throughput is negatively impacted. To be specific, the maximum throughput achievable for any given TCP connection in a zero loss network is expressed as:

throughput = window / RTT

where *window* is the maximum TCP window size (64 Kbytes by default) and RTT is the round trip time in the network.

This equation tells us that the throughput of a TCP connection is inversely proportional to network latency (note that this is TCP throughput for *one* connection - the aggregate bandwidth is not affected by latency). In other words, if you double latency, you halve throughput.

Consider adding a NIPS in an internal Gigabit network where the RTT is 200 microseconds. The critical value for RTT in a Gigabit network is 500 microseconds (below which it may no longer be possible to achieve 1Gbps of throughput), which means the NIPS can add a maximum of 300 microseconds to the RTT without affecting the network. In this particular case, therefore, for an internal, high speed deployment, the administrator may determine that his chosen IPS device needs to be capable of sub-300 microsecond latency under normal traffic loads.

Of course, the latency of an IPS device may vary significantly based on packet size, complexity of the protocol, presence of attack traffic, or simply the makeup of the normal traffic passing through it. For example, Gigabit segments, will rarely carry only a single TCP connection. Rather, a saturated Gigabit segment could be supporting hundreds, if not thousands of TCP connections, and this multiplexing eases the impact of latency on the overall throughput on the segment.

Although each of these connections carries only a fraction of the total throughput, a few connections tend to dominate. The maximum latency for a NIPS is then determined by the utilisation of the fastest connection. For example, in a Gigabit Ethernet segment carrying 10,000 TCP connections the fastest connection might have a throughput of 250Mbps. In this case, the critical value for round trip latency is as high as 2 milliseconds.

Assuming the latency without the NIPS is 300 microseconds, an administrator may therefore determine that his chosen NIPS device must be capable of 1700 microsecond round trip latency (850 microseconds in each direction).

Such critical value calculations are important when TCP connections achieve maximum throughput, which is true for large data transfers.

For smaller data transfers, and non-TCP applications like NFS, latency has a more direct impact on user experience - response time is directly proportional to latency. That is, *doubling latency doubles response time.* In these situations, the latency of the network in which a NIPS is deployed determines the acceptable latency of the NIPS.

Consider deploying a hypothetical NIPS with 1 millisecond one-way latency in the following scenarios:

- In internal corporate LANs, the round trip latency could be in the 200-300 microsecond range. Deploying our hypothetical NIPS would increase the maximum round trip latency to 2.3 milliseconds, an increase of just over 700 per cent. The time to copy a large group of files, for example, would increase by a factor of seven.
- In inter-campus corporate networks connected over a MAN, the latency could be in the 500-1000 microsecond range (or less). Deploying our hypothetical NIPS would increase the maximum round trip latency to 3 milliseconds, a minimum increase of 300 per cent. The time to copy a large group of files, for example, would increase by at least factor of three.
- Internet facing connections experience round-trip latency from 10-100 milliseconds. Deploying our hypothetical NIPS would increase the round trip latency by 1-10 per cent, which would have only a minor impact on the user experience.

The latency of the NIPS must therefore be evaluated in the context of the network in which it is deployed. For example, to protect networks that are accessed over the public Internet, one-way NIPS latencies in the 1-2 millisecond range would be acceptable. Whereas for NIPS deployments on MAN/WAN links, NIPS latencies of well under 1 millisecond would be essential. And as we have already mentioned, for deployments on internal networks where latencies are a few hundred microseconds, NIPS latencies of less than 300 microseconds would be more appropriate.

Network administrators have laboured long and hard to reduce latency within the corporate network to an absolute minimum. Core network devices such as switches are frequently chosen as much on their performance - packet loss and latency under all load conditions - as any other feature. Given that Network IPS devices are operating in-line, it is not surprising that they will be evaluated in a similar way.

For this reason, part of The NSS Group methodology uses very similar testing techniques to those we would normally employ when testing switches (in order to determine *packet latency*), in **addition** to measuring *application latency*. This group of tests determine the effect the IPS sensor has on the traffic passing through it under various load conditions. High packet latency will lower TCP throughput. High application latency will create a negative user experience.

Bi-directional network latency of a range of differently-sized UDP packets is measured under three test conditions: with no load, with 500 Mbps of HTTP traffic (or half the rated load of the device if this is less than 1Gbps), and while the device is under a heavy SYN flood attack (up to 10 per cent of the rated throughput of the sensor).

Spirent Avalanche and Reflector devices are also used to generate HTTP sessions through the device in order to gauge how any increases in latency will impact the user experience in terms of failed connections and increased Web response times.

This “*application latency*” is measured both with no background load and while the device is under attack.

Stability & Reliability

These tests verify the stability of the IPS device under various extreme conditions. Long-term stability is critical for an in-line IPS device, where failure can produce network outages.

In the first part of this test, we expose the external interface of the sensor to a constant stream of attacks over an extended period of time. The device is configured to block and alert, and thus this test provides an indication of the effectiveness of both the blocking and alert handling mechanisms. A continuous stream of exploits mixed with some legitimate sessions is transmitted through the sensor at a maximum rate of 90 per cent of the claimed throughput of the device for eight hours with no additional background traffic.

The device is expected to remain operational and stable throughout this test, blocking 100 per cent of recognisable exploits, raising an alert for each, and passing 100 per cent of legitimate traffic. If any recognisable exploits are passed - caused by either the volume of traffic or the IPS device failing open for any reason - this will result in a FAIL. If any legitimate traffic is blocked - caused by either the volume of traffic or the IPS device failing closed for any reason - this will also result in a FAIL.

In the second part of the test we stress the protocol stack of the device under test by exposing it to malformed traffic from the ISIC test tool for eight hours. The device is expected to remain operational and capable of detecting and blocking exploits throughout the test to attain a PASS.

We scan the management interface for open ports and active services and report on known vulnerabilities. We also stress the protocol stack of the management interface of the NIPS by exposing it to malformed traffic from the ISIC test tool. The device is expected to remain (a) operational and capable of detecting and blocking exploits, and (b) capable of communicating in both directions with the management server/console throughout the test to attain a PASS. We also note whether the sensor detects the ISIC attacks even though targeted at the management port.

Security Effectiveness

Detection Accuracy & Breadth

This group of tests verifies that the NIPS will not block legitimate traffic (*Accuracy*) and is capable of detecting and blocking a wide range of common exploits (*Breadth*). Although *breadth* is extremely important, *accuracy* is critical because a NIPS that blocks legitimate traffic will not remain in-line for long.

We have a number of trace files of normal traffic with “suspicious” content, together with several “neutered” exploits that have been rendered completely ineffective. The IPS attains a “PASS” for each test case if it does **not** raise an alert and does **not** block the traffic. Whilst it is not possible to validate completely the entire signature set of any IPS, this test demonstrates how accurately the IPS detects and blocks a wide range of common exploits, port scans, and Denial of Service attempts.

This test is repeated twice: the first run with blocking disabled on the IPS in order to determine which attacks are detected and how accurately they are detected (*Attack Recognition Rating*); the second run with blocking enabled in order to determine which attacks are blocked successfully regardless of how they are detected or what alerts are raised (*Attack Blocking Rating*).

Following the initial test run, each vendor is provided with a list of CVE references of the attacks missed and is allowed 48 hours to produce an updated signature set. This updated signature set must be released to the general public as a standard signature/product update before the report is published - this ensures that vendors do not attempt to code signatures just for this test.

Naturally, Rate-Based IPS devices will not respond to the same attack traffic as Content-Based devices, and so for those the Detection Accuracy tests involve detecting and mitigating a wide range of rate-based attacks such as port scans, SYN floods, connection floods, and so on. We note which of these are mitigated completely, which are mitigated partially, and which require the use of built-in firewall capabilities.

Resistance To Evasion Techniques

These tests verify that the IPS is capable of detecting and blocking basic exploits when subjected to varying common evasion techniques. An IPS that cannot detect attacks subjected to these “script kiddie” evasion techniques is easily bypassed.

The tests consist of four parts (only the third is applicable to Rate-Based devices):

- ***Baselines -*** *This establishes that the IPS is capable of detecting and blocking a number of common basic attacks (our baseline suite) in their normal state, with no evasion techniques applied.*
- ***Packet Fragmentation and Stream Segmentation -*** *The baseline HTTP attacks are repeated, running them through fragroute using 19 evasion techniques.*
- ***URL Obfuscation -*** *The baseline HTTP attacks are repeated, this time applying 9 URL obfuscation techniques made popular by the Whisker Web server vulnerability scanner.*
- ***Miscellaneous Evasion Techniques -*** *Certain baseline attacks are repeated, and are subjected to 7 protocol- or exploit-specific evasion techniques, including altering default ports, inserting spaces in FTP command lines, inserting non-text Telnet opcodes in FTP data streams, and RPC record fragging.*

For each of the evasion techniques, we note if (i) the attempted attack is blocked successfully (the primary aim of any IPS device), (ii) the attempted attack is detected and an alert raised in **any** form, and (iii) if the exploit is successfully “decoded” to provide an accurate alert relating to the original exploit, rather than alerting purely on anomalous traffic detected as a result of the evasion technique itself.

Stateful Operation

If the IPS is tracking TCP session state, then it has the potential to introduce denial of service when the session table becomes full (too many connections) or if it can’t keep up with the creation of new sessions (too many connections per second).

As with latency and bandwidth, the number of connections supported by the IPS and its connection per second rate should be matched to the network.

For example, a fully saturated Gigabit Ethernet link can handle 22,000 5KByte transfers per second. Assuming each connection lasts 20 seconds, the IPS should be able to handle 448,000 simultaneous connections. These numbers scale proportionately for slower networks. Any IPS that doesn't offer these capabilities will impact performance of Web or e-commerce servers.

The aim of this section is to be able to determine whether the IPS is capable of monitoring stateful sessions established through the device at various traffic loads without either losing state or incorrectly inferring state.

An IPS that does not maintain TCP session state can flood the management console with false-positive alerts. Although this should not directly impact the IPS blocking function, it can make it very hard to perform forensic analysis of the attacks. In addition, if the default condition of the sensor is to block all traffic for which it does not believe there is a current connection in place, then an inability to maintain state under extreme conditions could result in the sensor blocking legitimate traffic by mistake.

In the first part of this test, we transmit a number of packets taken from capture files of valid exploits, but without first establishing a valid session with the target server. In order to receive a "PASS" in this test, no alerts should be raised for any of the actual exploits. However, each packet should be blocked if possible since it represents a "broken" or "incomplete" session.

In part two, we test whether the sensor is capable of preserving state across increasing numbers of open connections, as well as continuing to detect and block new exploits while not blocking legitimate traffic when the state tables are filled. Various numbers of TCP sessions from 10,000 to 1,000,000 (one million) are tested.

This test is run in both the out-of-box configuration and then repeated after applying any tuning recommended by the vendor (if applicable) to increase the size of the state tables.

Usability

After quantitatively evaluating the network performance and security effectiveness of the IPS, we qualitatively evaluate the features and usability of the product.

This evaluation provides the reader with valuable insight into product features, how easy it is to install the IPS and perform common, day-to-day operations with the management console. Areas evaluated include *installation, configuration, policy editing, alert handling,* and *reporting and analysis*.

CISCO IPS-4255 V5.0(3)

Executive Summary

Cisco offers a family of appliances for passive IDS and in-line IPS deployments, designed to detect and prevent attacks across multiple network segments at up to 500Mbps in-line.

The Cisco IPS-4255 under test here is currently the top of the range IPS offering - a dedicated 1U appliance designed to monitor and protect multiple network segments at speeds up to 500Mbps in-line. The device sports four copper 10/100/1000Mbps ports which can be configured in various combinations of in-line and passive modes. An additional copper 10/100Mbps port is used for dedicated management.

Overall, the performance of the IPS-4255 was acceptable as a 500Mbps device, easily handling the rated 500Mbps of traffic under normal network conditions. Attack recognition and blocking ability were excellent out of the box, and improved to a perfect 100 per cent following an update, whilst resistance to all common evasion techniques tested was also perfect.

Latency was acceptable for the rated throughput, allowing the IPS-4255 to be deployed anywhere in the network - either at the perimeter or in the core. Latency increased significantly under load, however.

We found the IPS-4255 to be very stable and reliable under extended attack, and the handling of high-levels of DOS/DDOS attacks (such as SYN floods) was excellent, although such attacks are not currently mitigated for the target host.

A range of management options are offered out of the box, from an extensive Command Line Interface (CLI), which will be familiar to anyone who has ever managed Cisco networking or security devices, to a comprehensive centralised management system capable of handling multiple devices deployed throughout a large enterprise.

All of these are included in the base price of the product (although the high-end CiscoWorks VMS is restricted to a small number of devices), which is nice.

Policy management and deployment capabilities are second to none, and alert handling is capable. In the past, Cisco IDS/IPS products have been lacking in reporting tools, but with the acquisition of Protego Networks this gap has been plugged.

Between the CLI, IDM, MC and now MARS, Cisco offers a broad range of management, monitoring and reporting options that will suit most administrators.

Architecture

Cisco provides a range of IPS appliances and management solutions which can be deployed in either a basic two-tier or more advanced three-tier management architecture.

The main components of the Cisco solution are as follows:

Cisco IPS 4200 Series sensor appliances

These deliver intrusion prevention via dedicated, purpose-built devices that protect multiple network segments through the use of up to eight interfaces (via optional four port modules), and support dual operation simultaneously, in both promiscuous and prevention modes. Cisco IDS/IPS appliances provide a range of performance as follows:

- ***Cisco IDS 4215 -*** *80Mbps passive mode, 65Mbps in-line*
- ***Cisco IPS 4240 -*** *250Mbps (passive mode and in-line)*
- ***Cisco IPS 4255 -*** *600Mbps passive mode, 500Mbps in-line*
- ***Cisco IDS 4250-XL -*** *1000Mbps passive mode, 800Mbps in-line*

In addition, IDS/IPS modules are available for Cisco Catalyst 6500 Series switches. Each appliance runs a heavily modified version of Red Hat Linux, hardened and complete with a familiar Cisco Command Line Interface (CLI) and customised packet drivers and packet capture libraries written to cope with extreme traffic loads.

The device under test was the Cisco IPS-4255, a 1U rack-mount appliance designed for up to 500Mbps in-line. The 4255 sports four copper 10/100/1000Mbps ports which can be configured in various combinations of in-line and passive modes - two in-line pairs, one in-line pair and two passive mode IDS, or four passive mode IDS. An additional copper 10/100Mbps port is used for dedicated management.

A software bypass capability is built in to the sensor allowing it to fail open or closed, as required by the administrator, should the sensor application fail or resources become exhausted. Optional modules provide hardware bypass capabilities, allowing the entire device to fail open under power loss or hardware failure.

There is a single power supply, and neither power supply nor fans are hot-swappable. No other High Availability (HA) features are available.

Meta Event Generator (MEG)

Cisco IPS incorporates sensor-level event correlation that gives security administrators an automated method for enhancing the confidence level of the classification of malicious activity detected by the sensor. This provides a mechanism that allows for corresponding actions to contain worm and virus injection vectors, as well as worm propagation. This is accomplished through the following techniques:

- *Correlation of alarms pertaining to worms that exploit multiple vulnerabilities*
- *Meta event generation for sequences of actions leading up to worm infestation*
- *Automated elevation of severity ratings when groups of events signify worm/virus activity*
- *Enhancement of alarm fidelity through simultaneous triggers based on hybrid detection algorithms*

Nimda is a prime example of a worm that exploited multiple vulnerabilities during its propagation across networks. Typically, the various alarms that pertain to each of these exploits will trigger within a short time interval.

Using MEG, the user can specify logic that will consolidate all events pertaining to a certain worm into a single meta event, called "Nimda", for example (see below).

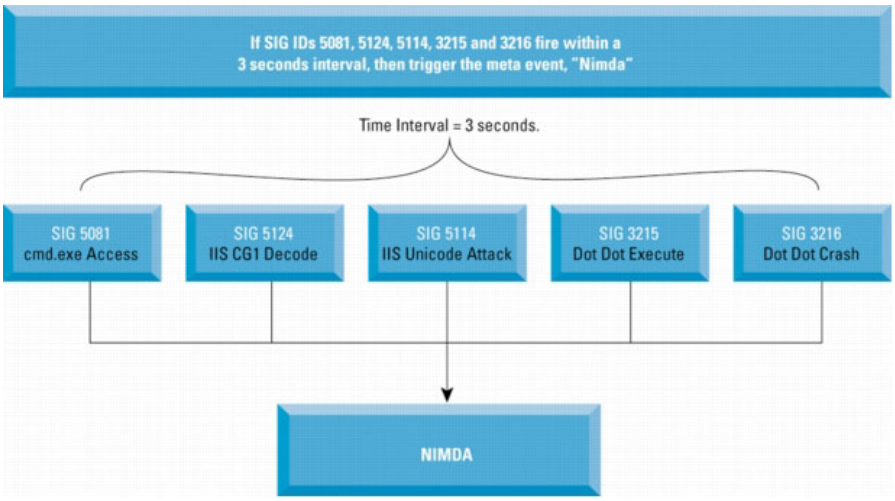

Figure 1 - Cisco IPS: Correlating related events using MEG

In doing so, the user can also specify a time interval during which these disparate events must be detected in order for the correlation algorithm to trigger the actual meta event.

Through the use of similar logical parameters, MEG can also be customised to deliver protection from environment-specific threats that may not be related to a universally known worm activity.

Although the user-defined nature of this feature makes it very flexible and powerful, Cisco recognises that the knowledge to create such meta events is lacking in may organisations. Thus, signature updates delivered by Cisco that pertain to multifaceted worms such as Nimda will also deliver the associated meta event. Along with this meta event, the user will be given information that indicates the individual signatures that form the meta event.

Historical trend analyses performed to characterise the lifecycle of worms often reveal a certain sequence of actions that are detected just prior to penetration. These actions occur in the “probing phase”, when a chain of reconnaissance activities is performed against the target network. MEG allows the user to define the precursors to worm penetration by specifying a logical algorithm that triggers when a particular sequence of events occur.

For example, if a certain number of hosts are pinged, followed by port scans on a defined set of ports, followed by a buffer overflow targeting hosts on a particular range of IP addresses, then trigger a single meta event “X”. In this case, the resulting meta event will attain a higher fidelity rating by virtue of the correlation that was performed. Additionally, this meta event can be assigned an automated response action that will stop the worm that has been detected.

As worms propagate through the network, they typically generate multiple IPS events of varying degrees of severity. When there is no relationship established between such disparate events they could be assigned low severity ratings since, by themselves, they do not pose a significant threat. However, when these events are considered in the context of a sequence of related events, they could collectively indicate worm or virus activity.

Cisco's Meta-Event Generator links these seemingly unrelated lower severity alarms into a high severity, high risk event, enabling the administrator more confidently to drop the associated packets (see below).

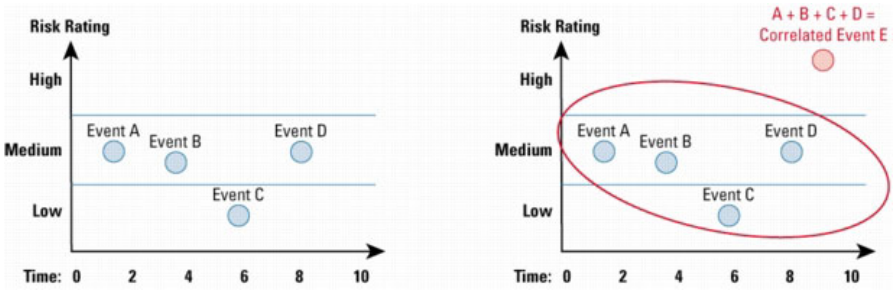

Figure 2 - Cisco IPS: MEG can correlate multiple low level events to create a high level alert

Lastly, MEG can be used to correlate events that are generated through the use of the hybrid detection techniques available in Cisco IPS sensor software.

For example, if a denial of service (DoS) activity is detected through the triggering of a traffic anomaly algorithm and a classical “flood” type of signature, MEG can be used to corroborate one event with the other, thereby delivering a single meta event that indicates a higher likelihood that the DoS activity has actually occurred. As always, the most appropriate response actions could then be configured to mitigate the DoS condition.

The effectiveness of the IPS is enhanced when such correlation algorithms are embedded into the sensor, as opposed to performing them purely at the monitoring console. When event correlation is performed at the sensor level, the sensor can take automated response actions that are based on the correlated incident rather than the individual - possibly low-level - events.

Command Line Interface (CLI)

A full-featured Cisco IOS Software-like CLI that provides device configuration over a Secure Shell (SSH) connection or via direct keyboard/video connection to the sensor.

IPS Device Manager (IDM)

The *IPS Device Manager* is a web-based Java application that resides on the sensor and is accessed via a secure, encrypted TLS link. Standard Netscape and Internet Explorer Web browsers are used to connect directly to the sensor to perform various management and monitoring tasks.

IDM is designed to manage a single device at a time in a two-tier management configuration, and it includes basic alert monitoring capabilities.

CiscoWorks VMS

A multi-device configuration and alarm management tool offering a unified, view of all security events across the enterprise.

With the CiscoWorks VMS solution, events from all types of security devices, including firewalls, VPNs, and IPS, can be viewed from a single console in a browser-based GUI.

Multiple security devices can be configured and managed, making it easier to manage security across the enterprise.

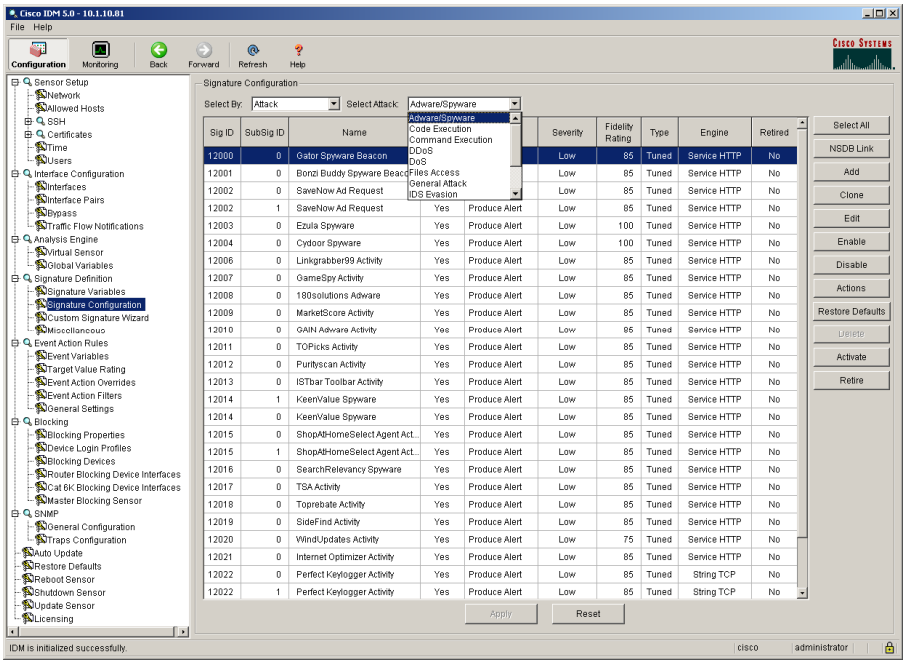

Figure 3 - Cisco IPS: IPS Device Manager (IDM)

CiscoWorks VMS is designed as a three-tier management system and requires a dedicated management server. It provides advanced policy management and deployment capabilities, role-based access, Security Monitor for alert handling, basic correlation capabilities via user-defined alert rules, basic reporting, and automated updates.

Performance

The aim of this section is to verify that the sensor is capable of detecting and blocking exploits when subjected to increasing loads of background traffic up to the maximum bandwidth supported as claimed by the vendor.

For each type of background traffic, we also determine the maximum load the IPS can sustain before it begins to drop packets/miss alerts. It is worth noting that devices which demonstrate 100 per cent blocking but less than 100 per cent detection in these tests will be prone to blocking **legitimate** traffic under similar loads.

The Cisco IPS-4255 was tested up to 500Mbps, the rated throughput of the device. Performance at almost all levels of our load tests was good, with 100 per cent of all attacks being detected and blocked under most load conditions.

The one exception was Test 4.2.4 where we determined that there was a limit of approximately 8,000 HTTP connections per second (equating to approximately 400Mbps) meaning the test could not be completed. Beyond this limit, legitimate connections began to fail, but all attack traffic was still blocked successfully.

However, we would happily confirm Cisco's 500Mbps rating for this device under normal network conditions.

Basic latency figures were good for a device of this type at most traffic loads and packet sizes, ranging from 112µs with 125Mbps of 256 byte packets, to 209µs with 500Mbps of 1000 byte packets. Behaviour through most of the tests with no background traffic was very predictable, with relatively small increases in latency as traffic levels increased. The one exception was at 500Mbps of 256 byte packets, where the device began to drop packets.

Placing the device under a half load of 250Mbps of HTTP traffic we noted significant increases of almost 300 per cent with 256 byte packets (112µs to 447µs), almost 200 per cent with 550 byte packets (131µs to 385µs), and 134 per cent with 1000 byte packets (169µs to 397µs).

However, we would still consider these results to be acceptable for a 500Mbps device. HTTP response times were good, and in most deployments the IPS-4255 could be situated anywhere on a network with no more than 500Mbps of traffic, either internally or at the perimeter.

SYN Flood protection is implemented for the sensor, and 50Mbps of SYN flood traffic had no significant effect on the IPS-4255. Latency increased by only a few microseconds across all packet sizes and HTTP response times were barely affected. Note, however, that the SYN Flood was **not** mitigated at all for the target server (this feature is planned for a future release) although appropriate alerts were raised.

The IPS-4255 performed consistently and completely reliably throughout our tests. Under eight hours of extended attack (comprising millions of exploits mixed with genuine traffic) it continued to block 100 per cent of attack traffic, whilst passing 100 per cent of legitimate traffic.

Exposing the sensor interface to ISIC-generated traffic had no adverse effect, and the device continued to detect and block all other exploits throughout and following the ISIC attack.

Please refer to the *Testing Methodology* section for full details of the methodology used and performance results.

Security Effectiveness

We installed one sensor with the latest signature pack, and configured it with all attack signatures enabled, plus some key audit/information-only signatures. All "*retired*" signatures were un-retired and enabled.

Out of the box, blocking performance was excellent at 89 per cent, and was improved to a perfect 100 per cent following the application of a signature update after 48 hours.

Detection/recognition rate was slightly lower (95 per cent following update) due to the fact that some protocol anomalies (i.e. overlapping fragments, invalid TCP options, etc.) are always blocked without alerting by the normaliser engine. It is possible to configure normaliser signatures to produce alerts if required.

We noted a minimum of "noise", with few test cases raising multiple alerts for a single exploit, and the accuracy of the exploit descriptions was high.

Performance in our “false negative” tests was very good out of the box, and there is every indication that, wherever possible, signatures are written for the underlying vulnerability rather than specific exploits.

A major concern in deploying an IPS is the blocking of legitimate traffic, and the IPS-4255 did block two of our false positive test cases out of the box. This was rectified following the signature update, but we continued to note numerous port-based backdoor alerts during our FTP tests.

The Cisco IPS arrives with a sensible default policy with PASS and BLOCK actions set for appropriate signatures (i.e. where the confidence level of a signature is high then the action will generally be set to BLOCK). It should be possible to deploy this product successfully using the default settings in most organisations.

Resistance to known evasion techniques was excellent, with the IPS-4255 achieving a clean sweep across the board in all our evasion tests. *Fragroute, Whisker, ADMmutate* and even *RPC record fragging* all failed to trick the IPS-4255 into ignoring valid attacks.

Not only were the fragmented and obfuscated attacks blocked successfully, but all but two of them were decoded accurately as well when not blocking. Once blocking was enabled, certain of the more complex fragmentation and segmentation evasion techniques raised alerts from the normaliser only.

Out of the box, the IPS-4255 handled 500,000 open connections (the claimed maximum) and no tuning parameters are available to the administrator to adjust this. Default operation of the device is to age out old connections when the state tables are full or resources are low, and this behaviour is not configurable.

Stateless “exploits” are not alerted upon (this is correct behaviour in order to be resistant to *Stick* and *Snot* tools) and mid-flows are blocked by default. It is possible to configure the device to allow mid-flows if required.

Please refer to the *Testing Methodology* section for full details of the methodology used and performance results.

Usability

This part of the test procedure consists of a subjective evaluation of the features and capabilities of the product, and covers *installation, configuration, policy editing, alert handling,* and *reporting and analysis*.

Installation

Installation tasks are minimal thanks to the turnkey appliance approach. All that is required is to log in at the sensor console and work through a text-based installation routine in order to set up the initial communication parameters of the sensor to allow it to communicate with the management console. Cisco has done a good job of implementing a complete Cisco CLI on the sensor platform so that the sensor has a familiar look and feel to anyone with Cisco experience.

Whichever management console is being used - Cisco *IPS Device Manager* (IDM) or *CiscoWorks VMS* - then all that is required is to define the communication parameters for each sensor to begin managing and monitoring via the GUI.

All documentation is supplied on CD-ROM and is also available on-line. The quality of the documentation is very good.

Configuration

For remote management functions, Cisco bundles the *IPS Device Manager* (IDM) with the sensor for basic management and monitoring. In addition, Cisco diehards will be pleased to learn that there is a full-featured Cisco CLI environment available when logging into the sensor locally - everything that can be configured via the GUI utilities can also be configured via the command line.

However, for enterprise deployments with more than a couple of sensors, the recommended management solution is CiscoWorks VMS. This is a Java-based “umbrella” management system for a range of Cisco devices such as firewalls, VPN and, of course, IDS/IPS sensors. For those organisations which find the included IDM to be too limited, then CiscoWorks provides enterprise-level IPS management, monitoring and alerting capabilities, and it will be the focus of this evaluation.

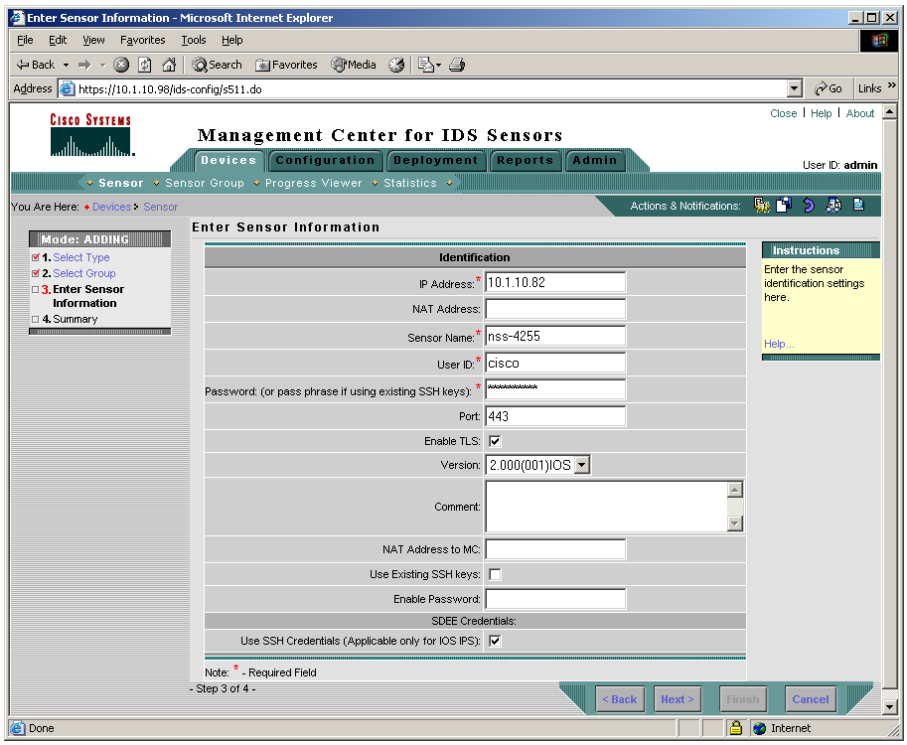

Figure 4 - Cisco IPS: Defining sensor devices in CiscoWorks VMS

Login to the CiscoWorks console is controlled by a user name and password combination. User accounts are created within CiscoWorks (or can be managed by ACS) and each account can have different levels of access to the system, providing granularity of administrative function. When using ACS, it is also possible to control authorisation on a per-device or per-group basis.

Once logged on, a hierarchical menu down the left of the screen provides access to CiscoWorks’ administrative functions (security settings, logging options, server configuration, user account maintenance, and so on), and the *Management and Monitoring Centre*.

A number of diverse modules can be added to this section to enable the administrator to manage a wide range of devices - such as firewalls, IPS, IDS, routers, etc. - from a single point.

The system provided for evaluation contained a *Management Centre for IDS Sensors* plug-in for the *Management Centre*, and a *Security Monitor* plug-in for the *Monitoring Centre*. These provide similar functionality to IDM, but offer more flexibility in some of the configuration functions, together with the ability to manage multiple sensors across the corporate network. They also handle configuration and monitoring tasks for both IDS and IPS devices.

The layout of the Management Centre (MC) is very straightforward, consisting of a number of tabs along the top of the screen providing access to *Devices, Configuration, Deployment, Reports* and *Admin* functions. Selecting any of these will bring up additional tabs below for access to different functions, and a very useful “you are here” indicator provides constant notification of which options have been selected to get to the screen currently displayed.

Sensors can be logically grouped together - perhaps according to location or function - and a sliding *Object Selector* tab to the left of the screen allows the administrator to select either an individual sensor or an entire group, to which all subsequent administrative operations will be applied. This can slide out of the way to reclaim screen real-estate, and the current scope is always shown on subsequent configuration screens.

Many parameters can be set globally, or at group level, and can be made mandatory at either of those levels, thus preventing them from being overridden on the individual sensors. Where the configuration is not enforced in this way, the administrator can specify settings at the global level which can then be overridden as required at lower levels in the device hierarchy.

Clearly, not all admin functions are available at global or group level, since some are tied closely to the operation of the individual sensor.

At the top right of the MC GUI is a small *Actions & Notifications* toolbar, containing buttons to generate and deploy new configurations, save pending changes, undo pending changes, view current background operations (useful for monitoring deployment tasks), and configure automatic signature updates.

Policy Management

The key to policy management in VMS is the ability to create *groups* and *subgroups* below the *global* root node in the Management Centre (MC)device tree. Groups can be created to mirror any policy structure which could, for example, be based on location (i.e. country or department) or function (i.e. Web server, FTP server, private or public-facing).

All, or some, settings can be inherited from the nodes above the one being edited, and settings applied to any node in the tree are automatically applied to all sensors within that node during the deployment operation. Changing the policy applied to any sensor is simply a matter of dragging and dropping it between nodes in the device hierarchy. This is a very flexible and intuitive method of policy management, and it works well in this implementation.

A “policy” is made up of a number of configuration parameters (accessed via the tree menu to the left of the screen once the scope has been selected) and, of course, a number of signatures.

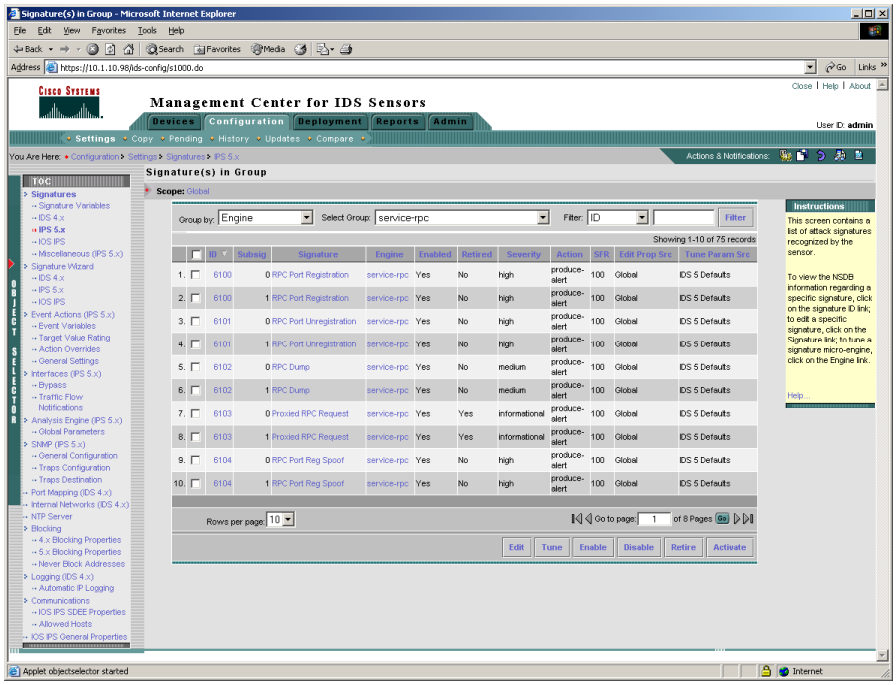

Figure 5 - Cisco IPS: Policy definition

By default, all built-in signatures are displayed in a single huge list when the GUI is first accessed. A drop-down menu provides access to several ways of grouping the signatures:

- *Signature ID*
- *Engine*
- *L2/3/4 Protocol*
- *Attack*
- *OS*
- *Release*
- *Service*

Selecting any one of these from the first drop-down menu populates a second drop-down menu allowing further granular selection of related groups of signatures. For example, if *Engine* is selected in the first menu, the administrator can select on which engine (roughly equivalent to a protocol decoder in most cases) to focus on in the second menu (*Normaliser, HTTP, RPC, SQL, SMB*, and so on). If *OS* is selected in the first menu, the administrator gets to select from a list of available *operating systems* for which signatures are available in the second menu.

Signatures can be edited, allowing the administrator to change the active flag, retired flag (retired signatures are removed completely from the sensor reducing the memory footprint), severity level, fidelity rating, actions, mandatory flag and override flag. The following actions are available:

- ***Deny attacker -*** *block all subsequent packets from source IP*
- ***Deny packet -*** *drop attack packet*
- ***Deny connection -*** *drop attack packet and all subsequent packets on that connection*

- ***Log attacker packets*** *- log inbound packets only*
- ***Log victim packets*** *- log outbound packets only*
- ***Log pair packets*** *- log packets in both directions*
- ***Modify packet*** *- rewrite and pass packet, removing obfuscation/evasion attempts or protocol anomalies*
- ***Produce alert*** *- send short alert to console*
- ***Produce verbose alert*** *- send complete alert to console*
- ***Request block connection*** *- use external device (i.e. firewall or router) to block connection*
- ***Request block host*** *- use external device (i.e. firewall or router) to block subsequent packets from source IP*
- ***Request SNMP trap*** *- send trap to external SNMP console*
- ***Reset TCP connection*** *- send TCP resets to source and target hosts*

Rather than opt for a fixed “block or allow” methodology, Cisco IPS uses dynamic *Risk Ratings* that are assigned to alerts generated from IPS sensors. The intent of this Risk Rating is to provide the user with an indication of the relative risk of the traffic or offending host continuing to access the user's network.

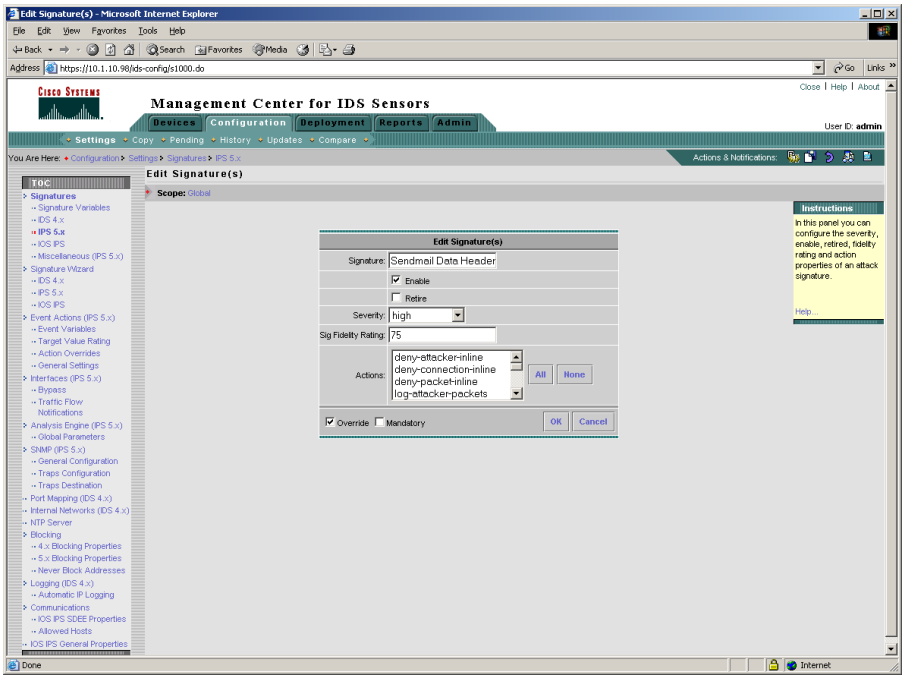

Figure 6 - Cisco IPS: Editing signatures

This rating can be used either to illuminate the events that require immediate administrator attention in the classic IDS promiscuous mode. It can also provide a means for developing risk-oriented event action policies when the sensor is deployed in the in-line IPS mode.

The Risk Rating is presented as an integer value in the range from 0 to 100 - the higher the value, the greater the security risk of the trigger event for the associated alert. The Risk Rating is a calculated number that has four primary components:

- ***Alert Severity Rating (ASR)*** *- A user-modifiable weighted value - Informational, Low, Medium, or High - that characterises the damage potential of the suspect traffic.*

- ***Signature Fidelity Rating (SFR) -*** *A user-modifiable weighted value that characterises the fidelity of the signature. Several factors affect the fidelity of a signature: whether it is relevant for a particular OS, service, application, or patch level, or whether it is prone to false positives, for example.*
- ***Attack Relevancy Rating (ARR) -*** *An internal weighted value that characterises any additional knowledge that the sensor may have about the target of the event. This information is used to clarify some of the uncertainties related to signature fidelity. As the sensor builds a more definitive picture of the target host, the risk associated with the event can be better defined.*
- ***Target Value Rating (TVR) -*** *A user-defined value that represents the user's perceived value of the target host. This allows the user to increase the risk of an event associated with a critical system and to de-emphasize the risk of an event on a low-value target.*

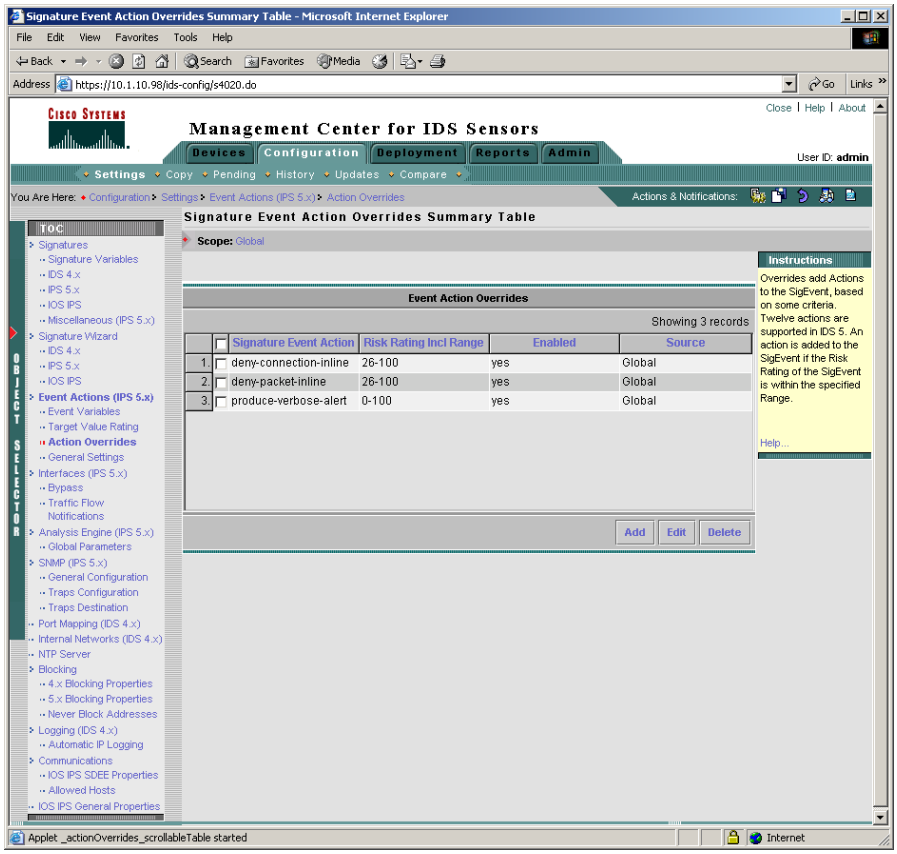

Figure 7 - Cisco IPS: Action Event Overrides

One of the most powerful features of the Cisco IPS system is the *Event Action Override* which allows the administrator to override default actions based on the calculated *Risk Rating*. The TVR in particular can have a significant effect on the overall Risk Rating based on the perceived (administrator-defined) value of the servers being protected. The effect could be a raising or lowering of the default Risk Rating as calculate by Cisco IPS.

Rather than be forced to make mass alterations to the default signature action settings, however, the administrator can apply global overrides on a per-action basis, based on the source of the event (individual sensors, sensor groups or sub-groups) and the Risk Rating (which can be expressed as a single value or a range).

Thus, for example, the administrator could decide that all events with a calculated risk rating of greater than 90 should be blocked, no matter what default action has been configured. Conversely, he could decide that all those events with Risk Rating below 10 are not even worthy of an alert.

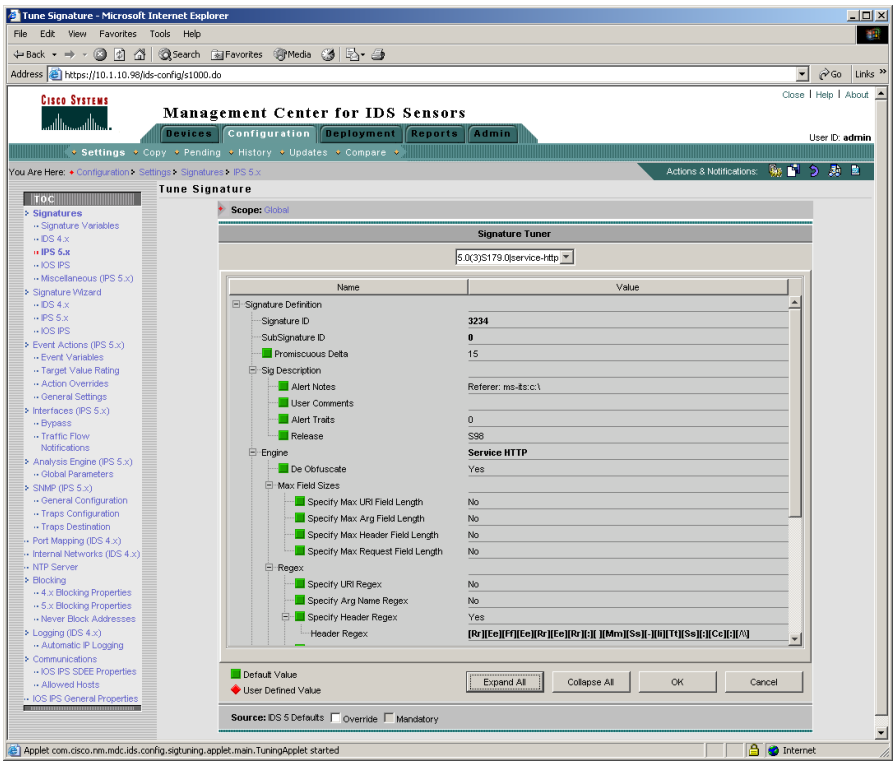

Figure 8 - Cisco IPS: Tuning signatures

This is an extremely powerful and flexible system - the nice thing about Cisco's implementation is that it does not rely purely on its *Signature Fidelity Ratings* and calculated Risk Ratings, but also allows the administrator to have input by specifying *TVRs* for critical servers.

Groups of signatures can be edited be clicking on the check boxes alongside each one, or by checking the "*select all*" box. Search criteria can also be used to search for specific groups of signatures within the current focus making it possible to search for all signatures whose description contains the phrase "*Back Orifice*", or all signatures that have a *severity* setting of "*High*". These too can be edited individually or *en masse* if required. All in all, the signature search, selection and editing capabilities are very powerful and very easy to use.

Signatures can also be *tuned*, which goes further than the basic *edit* capability, providing the means to access and modify all of the detailed signature parameters if required.

Most of the signatures written by Cisco are open for inspection and modification, and they can be cloned to provide the starting point for custom signatures. In addition, custom signatures can be created from scratch using the *Signature Wizard*, which will walk the administrator through this sometimes complicated process. All of the capabilities of all the protocol decoders are available for custom signatures, in addition to basic pattern matching and regular expression functions.

Whenever changes are made to a sensor configuration, the details are stored on the *Pending* tab awaiting deployment. The old two-stage process of *generating* configuration files and then *deploying* them is still available for those who want complete control, but the process has now been simplified, requiring only a click of the *Generate & Deploy* button on the main screen. Changes are deployed to all applicable sensors automatically (i.e. all those sensors in the group or sub-group to which the changes have been applied). Pending changes can be discarded using the *Discard* button or by deleting them from the *Pending* tab - this is a feature which we consider essential, but which is often missing from competing products.

A full change history is recorded, with every configuration file that has ever been applied saved as an audit trail. It is also possible to roll back a change to a previous configuration if required. Change management and policy deployment features in VMS are second to none.

Alert Handling

Alert Handling within CiscoWorks VMS is handled by the *Security Monitor* plug-in, which is capable enough as a basic day-to-day alerting tool and which can handle events from multiple devices.

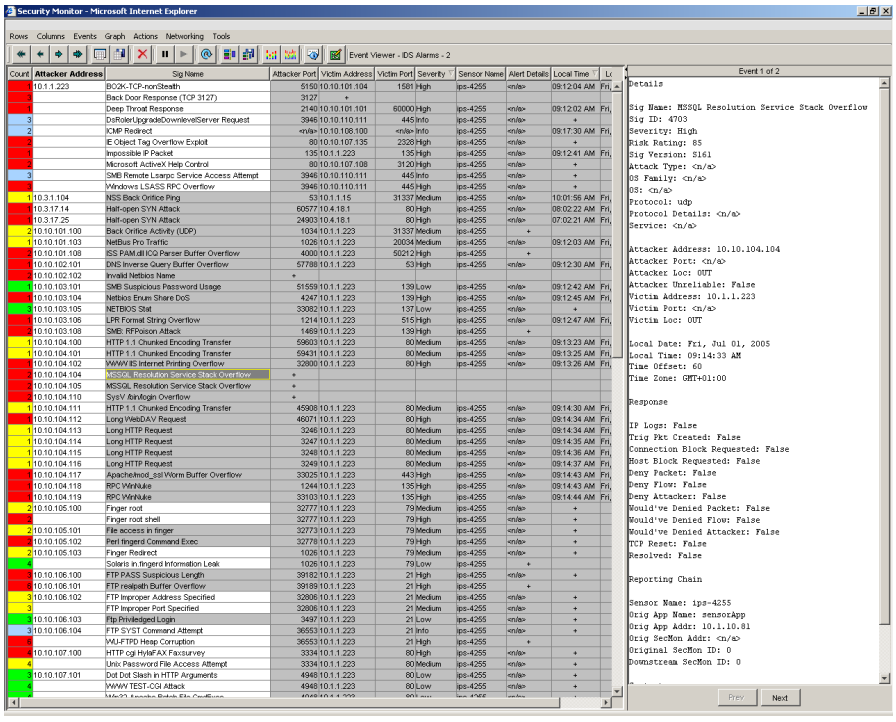

Figure 9 - Cisco IPS: Security Monitor

Once the devices to be monitored have been defined to the Security Monitor in the *Device* tab, the *Monitor* tab can be used to monitor both the status of each device, and the events that are raised.

When launching the *Event Viewer*, the administrator is presented with an initial screen of options allowing him to select the *Event Type* to view (Event Viewer can handle alerts from more than just IPS devices), the *Column Set* (the column layout: *By Victim, By Attacker, By Signature, All, Default* or *Last Saved*), event *start* and *stop* times, and the *Filter* to apply.

Filters can be created based on a wide range of criteria contained within the event (such as source IP, destination IP, severity, and so on) and these can be redefined and applied on the fly during a monitoring session. In common with the Column Set, the current criteria can be saved to re-use as the "*Last Saved*" set next time the viewer is launched. It would be nice if it were possible to save multiple named Column Sets and Filters, but for now only the last one of each saved is available for re-use.

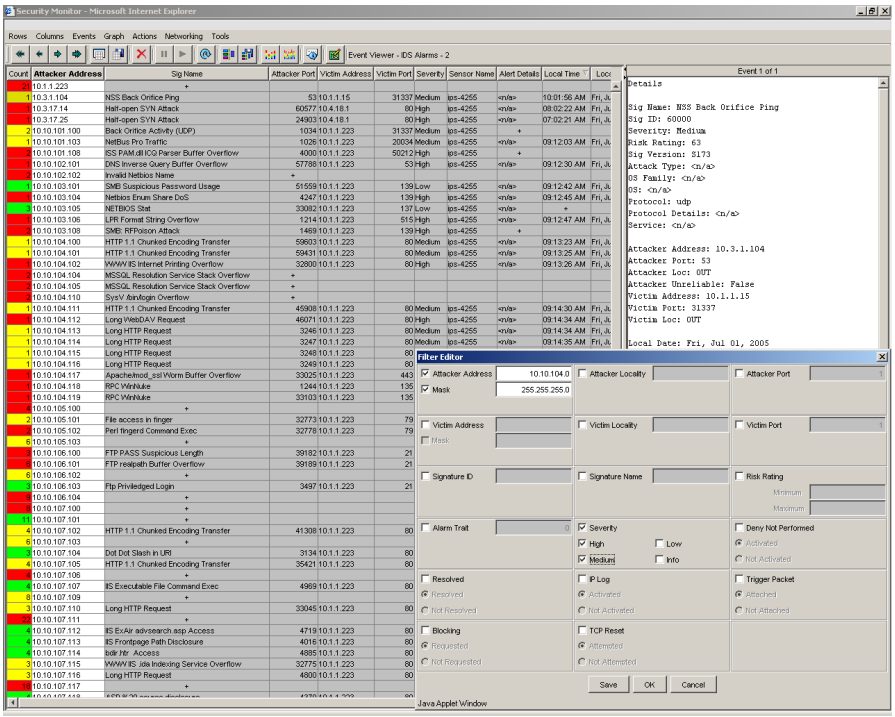

Figure 10 - Cisco IPS: Applying Security Monitor filters

The default Column Set is comprehensive, including the date and time, signature name, ID, severity, source IP/port, destination IP/port, sensor ID, and a separate pane on the right of the screen providing complete event details for any selected alert (including context data where applicable). Additional columns can be added (or existing columns deleted) as required.

Multiple Event Monitors can be launched simultaneously, each providing a different view of current events, and it is possible to drag and drop columns of data to re-order and resort the entries. For example, dragging *Source IP* to the left of the screen sorts and groups (and counts) by source IP. Dragging *Severity* to the left of the screen creates a new view sorted and grouped by the severity code.

The entire contents of the database can be expanded or collapsed using menu options to the left of the screen, providing a quick way to switch between detail and summary views. It is possible to view the full context buffer for an alert, and the viewer is capable of showing context data for an entire group of alerts in the same window.

On selecting any of the alerts, menu options provide the capability to view the trigger packet in an Ethereal-like display, block the host that perpetrated the exploit, or block the entire network from which it came. It is not possible, however, to annotate events, or mark them as *resolved* or for *further investigation*.

Although it may appear simple at first glance, and although it is sometimes not the most intuitive interface when it comes to resorting, expanding and collapsing the entries, the Event Viewer is actually a very flexible and powerful tool.

One final option worth mentioning is the *Event Rule* capability. Event Rules allow the administrator to specify related criteria, thresholds and actions in order to provide some very basic correlation. For example, a rule could be created to watch for a particular type of port scan from three sensors. If the scan is detected on all three sensors within a given time frame then an alert is generated, an e-mail is sent to the administrator and a script is run to reconfigure the firewall.

Whilst the Security Monitor is a perfectly capable tool for monitoring and investigating events in real-time, the customisation, filtering, correlation and forensic analysis capabilities are basic. With the acquisition of Protego Networks, Cisco can now offer another tool which can take alert handling to a higher level - *Cisco Security Monitoring, Analysis & Response System* (Cisco Security MARS).

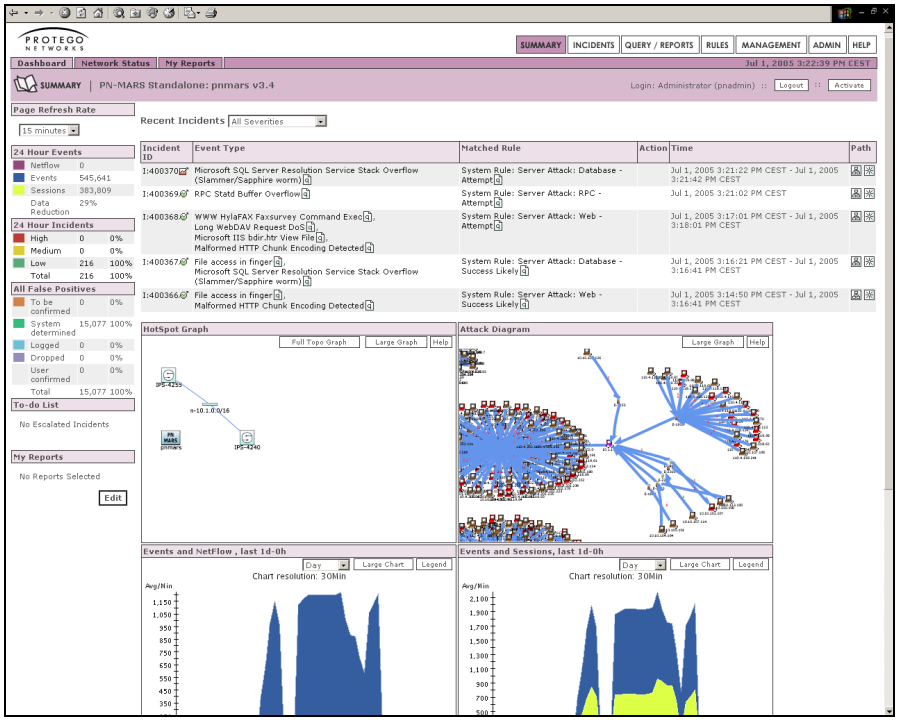

Figure 11 - Cisco IPS: MARS Summary screen

Intended mainly as an in-depth reporting, analysis and correlation system, its main features will be described in the following section. However, it is worth mentioning at this point that the myriad custom reports that can be defined within MARS can be configured to run in real-time, thus providing alert monitoring capabilities. In addition, there are a number of excellent graphical summary displays that provide good drill-down capabilities.

Although MARS can provide additional detail, and more flexible ways of drilling down to that detail, than Security Monitor, Security Monitor has a more lightweight feel to it. Some administrators might prefer Security Monitor for the basic, rapid-fire, day-to-day stuff, and MARS for the more in-depth analysis after the fact.

However, it should also be pointed out that MARS' ability to infer connections between multiple events and raise a single incident as a result can actually result in less information for the administrator to process, thus making him more efficient in tracking down and eliminating the most serious problems.

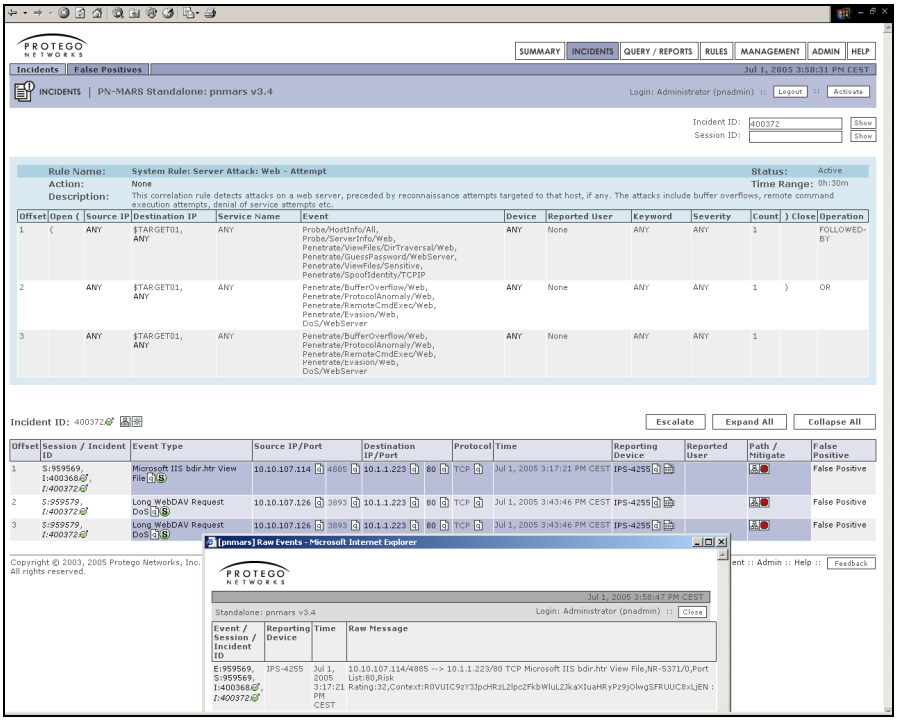

Figure 12 - Cisco IPS: MARS Incident display

Incidents are displayed in real-time in the *Incidents* display, with each entry in the display containing the incident name, the rule matched, the action taken, and the date and time. Hyperlinks are provided on almost every data item, allowing the administrator quickly to:

- *View the individual incident - the display consists of the rule matched and every raw event that was matched against that rule*
- *View the path taken by the attacker through the network (MARS is aware of the complete network topology),*
- *Access a graphical display of the attack vectors (pictorial representations of the source and target hosts, intermediate network devices, and which individual attacks were launched between each pair of hosts)*
- *Launch a query searching for all incidents with matching Rule*
- *Launch a query for all incidents with matching Event Types*

We found the correlation capabilities to be very impressive, and the mix of real-time graphical displays with rapid-fire hyperlinked queries to be very usable and powerful.

Reporting and Analysis

Security Monitor provides a number of basic summary reports:

- *24 Hour Alarm Metrics*
- *30 Day Alarm Metrics*

- *30 Day Details: Alarm Destinations*
- *30 Day Details: Alarm Source/Destination pairs*
- *30 Day Details: Alarm Sources*
- *30 Day Details: Alarms*
- *30 Day Details: Alarms by Hour/Day*
- *30 Day Details: Top 50 Alarms*
- *30 Day Details: Top 50 Alarm Destinations*
- *30 Day Details: Top 50 Alarm Source/Destination pairs*
- *30 Day Details: Top 50 Alarm Sources*
- *30 Day Alarm Summary*
- *Detailed Alarms By Sensor*

None of these is customisable in any way. On choosing a report, the administrator is presented with an extensive set of filtering options including severity, date, source, destination, signature and category. The report can then be scheduled to run immediately or at a later time. If scheduled for a later time it can also be made to repeat at regular intervals, and the finished report can be e-mailed to one or more recipients.

The completed report is viewed on-screen in a browser window and can be printed out using the standard browser print function. No other print or export options are offered.

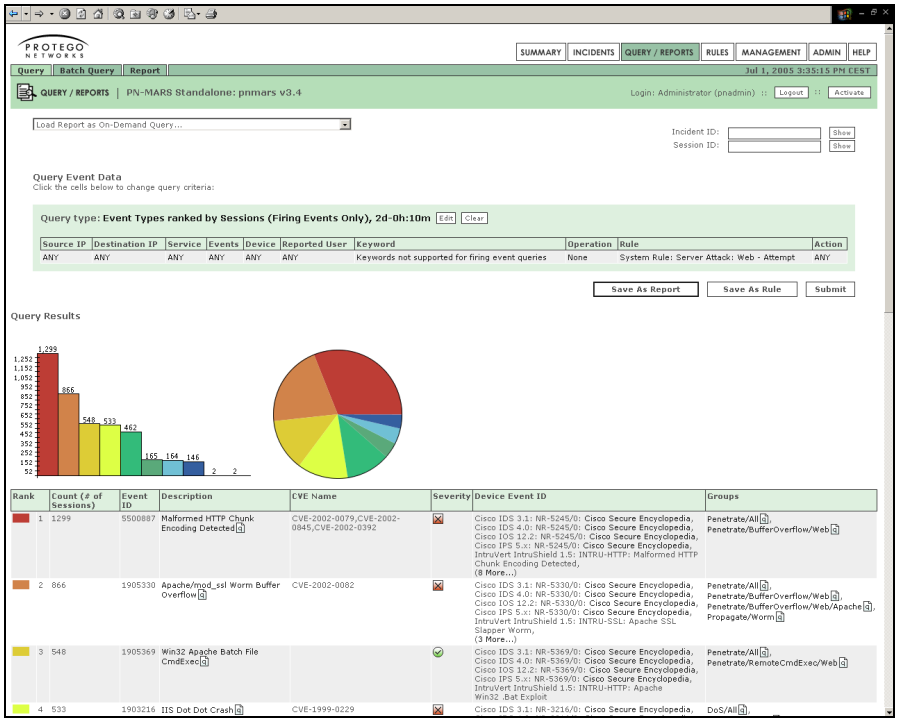

Figure 13 - Cisco IPS: MARS reports

The summary reporting is adequate, but is very high level and not suitable for forensic analysis. More detailed forensic reporting and correlation across multiple devices (both Cisco and third party, and not restricted to IDS/IPS) is offered via the aforementioned MARS system. This is provided as a turnkey appliance, and a range of appliances is available depending on size of operation and amount of data to be analysed.

MARS is first and foremost a data correlation system. At the heart of the product is a huge database of rules which describe events and conditions that must occur in order for incidents to be created.

Devices are configured within MARS, following which MARS will poll those devices at regular intervals for raw event data, storing that data in its own local database. This data - which can be from IPS/IDS sensors, firewalls, routers, switches and other devices (both Cisco and third party) - is then matched against the rules database which is capable of inferring connections between events occurring on different devices and in different time-frames.

MARS offers over 80 pre-defined reports, all of which can be modified via the flexible and intuitive report generator. Not all of these are IDS/IPS related, of course - many of them will relate to other devices and some are designed to satisfy operational requirements and assist in regulatory compliance efforts including Sarbox, GLBA, HIPPA, FISMA, and Basel II.

The report generator can modify any of the standard reports or can be used to generate new reports for action and remediation plans, incident and network activity, security posture and audit, and departmental reports - all in data, trend and chart formats.

Selection criteria are extensive, including source and destination IP, service, device, event type, users, keywords, rules and actions. Any number of criteria can be combined with Boolean operators to create extremely complex reports, and reports can be based on either high-level correlated incidents or the low-level individual raw events. The system also provides for batch and e-mail reporting.

In addition to the query/reporting capability, MARS also provides a high-level summary screen showing the most recent events in a list at the top of the screen (with the same rapid hyperlink access to drill-down data mentioned in the Alert Handling section), followed by a graphical attack view (shows the known network topology overlaid with attacks), and graphs of top destinations, top events and sessions, and top events and netflows.

In past evaluations of Cisco software we have bemoaned the lack of reporting capabilities. It is nice to see that gap plugged so effectively in this latest release - MARS is an extremely powerful and useful product.

Verdict

Performance

The Cisco IPS-4255 was tested up to 500Mbps, the rated throughput of the device. Performance at almost all levels of our load tests was good, with 100 per cent of attacks being detected and blocked under all but the most extreme (10,000 connections per second) load conditions.

Even when pushed beyond its limits (8,000cps), the IPS-4255 continued to block all malicious traffic successfully. We would happily confirm Cisco’s 500Mbps rating for this device under normal network conditions.

Basic latency figures were very good with no traffic, and still within acceptable limits for a device of this type with a half load of 250Mbps of HTTP traffic. HTTP response times were also good, and in most deployments the IPS-4255 could be situated anywhere on a network with no more than 500Mbps of traffic, either internally or at the perimeter.

SYN Flood protection is implemented for the sensor, and 50Mbps of SYN flood traffic had no significant effect on the IPS-4255.

Latency increased by only a few microseconds across all packet sizes and HTTP response times were barely affected. Unfortunately, the SYN Flood was **not** mitigated at all for the target server (this feature is planned for a future release) although appropriate alerts were raised.

The IPS-4255 performed consistently and completely reliably throughout our tests, blocking 100 per cent of attack traffic whilst passing 100 per cent of legitimate traffic. Exposing the sensor interface to ISIC-generated traffic had no adverse effect.

Security Effectiveness

Out of the box, blocking performance was excellent at 89 per cent, and was improved to a perfect 100 per cent following the application of a signature update after 48 hours. Detection/recognition rate was slightly lower (95 per cent following update) due to the fact that some protocol anomalies are always blocked without alerting by the normaliser engine, but this is not important for an in-line IPS device.

We noted a minimum of “noise”, with few test cases raising multiple alerts for a single exploit, and the accuracy of the exploit descriptions was high. Performance in our “false negative” tests was very good out of the box, and there is every indication that, wherever possible, signatures are written for the underlying vulnerability rather than specific exploits.

A major concern in deploying an IPS is the blocking of legitimate traffic, and the IPS-4255 did block two of our false positive test cases out of the box. This was rectified following the signature update, but we continued to note numerous port-based backdoor alerts during our FTP tests, which we considered poor (these can easily be disabled, of course).

The *Meta Event Generator* (MEG) appeared to work well in correlating low-level attacks to infer a connection and raise higher-level alerts for worm activity, etc. This is a very useful feature, though it will be interesting to see how MARS handles custom meta events when it attempts to apply its own correlation rules.

It was also good to see entire categories of signatures for new threats such as Spyware/Adware, Voice Over IP (VOIP), covert channel activity, and even viruses (via a new partnership with Trend Microsystems).

Resistance to known evasion techniques was excellent, with the IPS-4255 achieving a clean sweep across the board in all our evasion tests. *Fragroute, Whisker, ADMmutate* and even *RPC record fragging* all failed to trick the IPS-4255 into ignoring valid attacks. Not only were the fragmented and obfuscated attacks blocked successfully, but all but two of them were decoded accurately when not blocking.

Out of the box, the IPS-4255 handled 500,000 open connections (the claimed maximum) and no tuning parameters are available to the administrator to adjust this.

Default operation of the device is to age out old connections when the state tables are full or resources are low, and this behaviour is not configurable. It would be nice to see a choice offered to the administrator of whether to age out old connections or block new ones.

Usability

The Cisco product offers a range of management options from an extensive Command Line Interface (CLI), which will be familiar to anyone who has ever managed Cisco networking or security devices, to a comprehensive centralised management system capable of handling multiple devices deployed throughout a large enterprise. All of these are included in the base price of the product (although the high-end CiscoWorks VMS is restricted to a small number of devices), which is nice.

IDM is the direct device management software produced by the sensor team which is designed to allow the administrator to manage a single device at a time (or a small number of devices). This provides a basic, two-tier management infrastructure which does not require a separate management server. IDM provides everything needed to deploy and manage small numbers of sensors, including a very basic alert monitoring capability. Its biggest advantage is that it provides a very simple, straightforward means to manage a single device, and is often more up to date than VMS, which is produced by a separate development team and sometimes lags behind sensor development.

CiscoWorks VMS is where it steps up a gear into true centralised management and monitoring for multiple sensors, though this requires a separate management server host to implement the three-tier architecture. Signature editing, tuning and creation are well catered for, and the interface makes it simple to search for and make changes to large groups of signatures in one hit. Management of policies/groups and deployment across multiple sensors is made as straightforward as possible - the “scope” selection tool makes it possible to select single sensors, groups of sensors, or all sensors in order to make bulk changes.

VMS also includes a basic event monitoring tool (Security Monitor) which is perfect for day-to-day monitoring tasks. Data can be grouped and sorted in a variety of ways by simply dragging and dropping columns, though this is the limit of its analysis capabilities.

In order to go further, Cisco now offers its latest acquisition from Protego Networks - *Cisco Security Monitoring, Analysis & Response System* (Cisco Security MARS) - as an extra cost option. Provided as a range of turnkey appliances, this is an extremely comprehensive correlation, monitoring and reporting tool that pulls together alerts from a wide range of Cisco and third party security and networking devices, stores them in a central database, and correlates them into incidents. On top of this is a very powerful custom report generator which enables the administrator to create real-time and historical reports on the data, presented in almost any way imaginable.

In the past, Cisco IDS/IPS products have been lacking in reporting tools - with the acquisition of Protego Networks, however this gap has been plugged. Between the CLI, IDM, MC and now MARS, Cisco offers a broad range of management, monitoring and reporting options that will suit most administrators.

The CLI provides a very rapid means of single-device configuration without having to handle relatively sluggish Java GUIs. The IDM also provides a relatively rapid means of configuring single devices without having to deploy a complex three-tier management infrastructure. For those who have many sensors to manage, CiscoWorks VMS provides a comprehensive three-tier management system, and the built-in Security Monitor is perfect for basic, day-to-day monitoring tasks.

Finally, MARS offers comprehensive monitoring and reporting capabilities across not just every IPS device in the organisation, but every other networking and security device too.

There is, however, one fly in the ointment. Although VMS has seen considerable improvements since we first looked at Cisco products in our labs, and now appears to be a very flexible and powerful product, it continues to suffer from the fact that its development lags behind that of the sensor and its associated IDM. This means that users who are committed to VMS are often left with no choice but to use IDM immediately following a new sensor release whilst they wait for the VMS developers to incorporate new sensor features.

In addition, in this particular test, although the sensor release was on its third iteration - and both it and IDM seemed very stable - the VMS release was brand new and showed signs of a hurried QA process. There were several cosmetic and functional bugs which manifested themselves during testing, but which should be fixed by the time this report is released.

Luckily, none of these was enough to prevent the product from performing its allotted task as an IPS device, since it is perfectly possible to deploy and manage it via the IDM.

Contact Details

Company name: Cisco Systems, Inc.

Internet: www.cisco.com

Address:
170 West Tasman Dr
San Jose
CA 95134-1706
USA

Tel: +1 408 526 7660

Fax: +1 408 527 0883

APPENDIX A – TEST RESULTS

The aim of this procedure is to provide a thorough test of all the main components of an in-line Intrusion Prevention System (IPS) device in a controlled and repeatable manner and in the most “real world” environment that can be simulated in a test lab.

The Test Environment

The network is 100/1000Mbit Ethernet with CAT 5e cabling and Cisco 6500-Series switches (these have a mix of fibre and copper Gigabit interfaces). All devices are expected to be provided as appliances - if software-only, the supplier pre-installs the software on the recommended hardware platform. The sensor is configured as a perimeter device during testing (i.e. as if installed behind the main Internet gateway/firewall). There is no firewall protecting the target subnet.

Traffic generation equipment - such as the machines generating exploits, Spirent Avalanche and Spirent Smartbits *transmit* port - is connected to the “external” network, whilst the “receiving” equipment - such as the “target” hosts for the exploits, Spirent Reflector and Spirent Smartbits *receive* port - is connected to the internal network. The device under test is connected between two “gateway” switches - one at the edge of the external network, and one at the edge of the internal network.

All “normal” network traffic, background load traffic and exploit traffic will therefore be transmitted **through** the device under test, from external to internal. The same traffic is mirrored to a single SPAN port of the external gateway switch, to which an Adtech network monitoring device is connected. The Adtech AX/4000 monitors the same mirrored traffic to ensure that the total amount of traffic never exceeds 1Gbps (which would invalidate the test run).

The management interface is used to connect the appliance to the management console on a private subnet. This ensures that the sensor and console can communicate even when the target subnet is subjected to heavy loads, in addition to preventing attacks on the console itself.

Section 1 – Detection Engine

The aim of this section is to verify that the sensor is capable of detecting and blocking a wide range of common exploits accurately, whilst remaining resistant to false positives. All tests in this section are completed with **no background network load**. The latest signature pack is acquired from the vendor, and sensors are deployed with **all** available attack signatures enabled (some audit/informational signatures may be disabled).

Test 1.1 - Attack Recognition

Whilst it is not possible to validate completely the entire signature set of any sensor, this test attempts to demonstrate how accurately the sensor detects and blocks a wide range of common exploits, port scans, and Denial of Service attempts. These are updated/changed for every new test, and all exploits are run with no load on the network and no IP fragmentation.

Our attack suite contains over 100 basic exploits (plus variants) covering the following areas:

- ***Test 1.1.1 -*** *Backdoors (standard ports and random ports)*
- ***Test 1.1.2 -*** *DNS/WINS*
- ***Test 1.1.3 -*** *DOS*
- ***Test 1.1.4 -*** *False negatives (common exploits which have been modified to remove or alter obvious "triggers" - this ensures that the signatures are coded for the underlying vulnerability rather than a particular exploit)*
- ***Test 1.1.5 -*** *Finger*
- ***Test 1.1.6 -*** *FTP*
- ***Test 1.1.7 -*** *HTTP*
- ***Test 1.1.8 -*** *ICMP (including unsolicited ICMP response)*
- ***Test 1.1.9 -*** *Reconnaissance*
- ***Test 1.1.10 -*** *RPC*
- ***Test 1.1.11 -*** *SSH*
- ***Test 1.1.12 -*** *Telnet*
- ***Test 1.1.13 -*** *Database*
- ***Test 1.1.14 -*** *Mail*
- ***Test 1.1.15 -*** *Voice*

A wide range of vulnerable target operating systems and applications are used, and the majority of the attacks are successful, gaining root shell or administrator privileges on the target machine.

We expect all the attacks to be reported in as straightforward and clear a manner as possible (i.e. an "RDS MDAC attack" should be reported as such, rather than a "Generic IIS Attack"). Wherever possible, attacks should be identified by their assigned CVE reference. It will also be noted when a response to an exploit is considered too "noisy", generating multiple similar or identical alerts for the same attack. Finally, we will note whether the device blocks the attack packet only or the entire "suspicious" TCP session.

This test is repeated twice: the first run with blocking disabled on the sensor (monitor mode only) in order to determine which attacks are detected and how accurately they are detected (*Attack Recognition Rating*); the second run with blocking enabled in order to determine which attacks are blocked successfully regardless of how they are detected or what alerts are raised (*Attack Blocking Rating*)

The "**default**" *Attack Recognition Rating-Detect Only* (ARRD) and *Attack Recognition Rating-Block* (ARRB) are each expressed as a percentage of detected/blocked exploits against total number of exploits launched with the default signature set as received by NSS. This demonstrates how effective the sensor can be when simply deploying the default configuration.

Following the initial test run, each vendor is provided with a list of CVE references of the attacks missed, and is then allowed 48 hours to produce an updated signature set. This updated signature set **must** be released to the general public as a standard signature/product update before the report is published - this ensures that vendors do not attempt to code signatures just for this test.

The sensor is then exposed to a second round of identical tests and the "**custom**" ARRD/ARRB is determined. This demonstrates how effective the vendor is at responding to a requirement for new or updated signatures.

Both the *default* and *custom* ARRD/ARRB figures are reported.

Test 1.2 - Resistance To False Positives

The aim of this test is to demonstrate how likely it is that a sensor raises a false positive alert - particularly critical for IPS devices.

We have a number of trace files of normal traffic with "suspicious" content, together with several "neutered" exploits which have been rendered completely ineffective. If a signature has been coded for a specific piece of exploit code rather then the underlying vulnerability, or if it relies purely on pattern matching, some of these false alarms could be alerted upon.

The product attains a "PASS" for each test case if it does **not** raise an alert and does **not** block the traffic. Raising an alert on any of these test cases is considered a "FAIL", since none of the "exploits" used in this test represents a genuine threat. A "FAIL" would thus indicate the chance that the sensor could block legitimate traffic inadvertently.

- ***Test 1.2.1 -*** *False positives*

Section 2 – Evasion

The aim of this section is to verify that the sensor is capable of detecting and blocking basic exploits when subjected to varying common evasion techniques.

Test 2.1 - Baselines

The aim of this test is to establish that the sensor is capable of detecting and blocking a number of common basic attacks (our baseline suite) in their normal state, with no evasion techniques applied. Note that common/older attacks have been chosen deliberately for this particular test to ensure that ALL products tested have signatures in place for the evasion tests.

- ***Test 2.1.1 -*** *Baseline attack replay*

Test 2.2 - Packet Fragmentation and Stream Segmentation

The baseline HTTP attacks are repeated, running them through fragroute using various evasion techniques, including:

- ***Test 2.2.1 -*** *IP fragmentation - ordered 8 byte fragments*
- ***Test 2.2.2 -*** *IP fragmentation - ordered 24 byte fragments*
- ***Test 2.2.3 -*** *IP fragmentation - out of order 8 byte fragments*
- ***Test 2.2.4 -*** *IP fragmentation - ordered 8 byte fragments, duplicate last packet*
- ***Test 2.2.5 -*** *IP fragmentation - out of order 8 byte fragments, duplicate last packet*
- ***Test 2.2.6 -*** *IP fragmentation - ordered 8 byte fragments, reorder fragments in reverse*

- ***Test 2.2.7 -*** *IP fragmentation - ordered 16 byte fragments, fragment overlap (favour new)*
- ***Test 2.2.8 -*** *IP fragmentation - ordered 16 byte fragments, fragment overlap (favour old)*
- ***Test 2.2.9 -*** *TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with invalid TCP checksums*
- ***Test 2.2.10 -*** *TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with null TCP control flags*
- ***Test 2.2.11 -*** *TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with requests to resync sequence numbers mid-stream*
- ***Test 2.2.12 -*** *TCP segmentation - ordered 1 byte segments, duplicate last packet*
- ***Test 2.2.13 -*** *TCP segmentation - ordered 2 byte segments, segment overlap (favour new)*
- ***Test 2.2.14 -*** *TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with out-of-window sequence numbers*
- ***Test 2.2.15 -*** *TCP segmentation - out of order 1 byte segments*
- ***Test 2.2.16 -*** *TCP segmentation - out of order 1 byte segments, interleaved duplicate segments with faked retransmits*
- ***Test 2.2.17 -*** *TCP segmentation - ordered 1 byte segments, segment overlap (favour new)*
- ***Test 2.2.18 -*** *TCP segmentation - out of order 1 byte segments, PAWS elimination (interleaved dup segs with older TCP timestamp options)*
- ***Test 2.2.19 -*** *IP fragmentation - out of order 8 byte fragments, interleaved duplicate packets scheduled for later delivery*
- ***Test 2.2.20 -*** *TCP segmentation - ordered 16 byte segments, segment overlap (favour new (Unix))*

For each of the evasion techniques, we note if (i) the attempted attack is blocked successfully (the primary aim of any IPS device), (ii) the attempted attack is detected and an alert raised in **any** form, and (iii) if the exploit is successfully “decoded” to provide an accurate alert relating to the original exploit, rather than alerting purely on anomalous traffic detected as a result of the evasion technique itself.

Test 2.3 - URL Obfuscation

The baseline HTTP attacks are repeated, this time applying various URL obfuscation techniques made popular by the Whisker Web server vulnerability scanner, including:

- ***Test 2.3.1 -*** *URL encoding*
- ***Test 2.3.2 -*** */./ directory insertion*
- ***Test 2.3.3 -*** *Premature URL ending*
- ***Test 2.3.4 -*** *Long URL*
- ***Test 2.3.5 -*** *Fake parameter*
- ***Test 2.3.6 -*** *TAB separation*
- ***Test 2.3.7 -*** *Case sensitivity*
- ***Test 2.3.8 -*** *Windows \ delimiter*
- ***Test 2.3.9 -*** *Session splicing*

For each of the evasion techniques, we note if (i) the attempted attack is blocked successfully, (ii) the attempted attack is detected and an alert raised in **any** form, and (iii) if the exploit is successfully "decoded" to provide an accurate alert relating to the original exploit, rather than alerting purely on anomalous traffic detected as a result of the evasion technique itself.

Test 2.4 - Miscellaneous Evasion Techniques

Certain baseline attacks are repeated, and are subjected to various protocol- or exploit-specific evasion techniques, including:

- ***Test 2.4.1 -*** *Altering default ports/passwords for backdoors*
- ***Test 2.4.2 -*** *Inserting spaces in FTP command lines*
- ***Test 2.4.3 -*** *Inserting non-text Telnet opcodes in FTP data stream*
- ***Test 2.4.4 -*** *Polymorphic mutation (ADMmutate)*
- ***Test 2.4.5 -*** *Altering protocol and RPC PROC numbers*
- ***Test 2.4.6 -*** *RPC record fragging (MS-RPC and Sun)*
- ***Test 2.4.7 -*** *HTTP exploits to non-standard port*

For each of the evasion techniques, we note if (i) the attempted attack is blocked successfully, (ii) the attempted attack is detected and an alert raised in **any** form, and (iii) if the exploit is successfully "decoded" to provide an accurate alert relating to the original exploit, rather than alerting purely on anomalous traffic detected as a result of the evasion technique itself.

Section 3 – Stateful Operation

The aim of this section is to be able to determine whether the sensor is capable of monitoring stateful sessions established through the device at various traffic loads without either losing state or incorrectly inferring state.

Test 3.1 - Stateless Attack Replay (Mid-Flows)

This test determines whether the sensor is resistant to stateless attack flooding tools - these utilities are used to generate large numbers of false alerts on the protected subnet using valid source and destination addresses and a range of protocols.

The main characteristic of many flooding tools is the fact that they generate single packets containing "trigger" patterns without first attempting to establish a connection with the target server. Whilst this can be effective in raising alerts with some stateless protocols such as UDP and ICMP, they should never be capable of raising an alert for exploits based on stateful protocols such as FTP and HTTP.

In this test, we transmit a number of packets taken from capture files of valid exploits, but without first establishing a valid session with the target server. We also remove the session tear down and acknowledgement packets so that the sensor can not "infer" that a valid connection was made.

In order to receive a "PASS" in this test, no alerts should be raised for any of the actual exploits (although "mid-flow" alerts are permitted).

However, each packet should be blocked if possible since it represents a "broken" or "incomplete" session.

- ***Test 3.1.1 -*** *Stateless attack replay*

Test 3.2 - Simultaneous Open Connections (default settings)

This test determines whether the sensor is capable of preserving state across increasing numbers of open connections, as well as continuing to detect and block new exploits when the state tables are filled. It also attempts to determine whether or not the sensor will block legitimate traffic once state tables are filled. This test is run using the default sensor settings (no tuning of sensor parameters).

A legitimate HTTP session is opened and the first packet of a two-packet exploit is transmitted. The Spirent Avalanche (on the "external" interface of the sensor) then opens various numbers of TCP sessions from 10,000 to 1,000,000 (one million) with the Spirent Reflector (on the "internal" interface of the sensor). The initial HTTP session is then completed with the second half of the exploit and the session is closed. If the sensor is still maintaining state on the first session established, the exploit will be recorded. If the state tables have been exhausted, the exploit string will be seen as a non-stateful attack, and will thus be ignored.

Both halves of the exploit are required to trigger an alert - a product will fail the test if it fails to generate an alert after the second packet is transmitted, or if it raises an alert on either half of the exploit on its own.

At each step, we ensure that the sensor is still capable of detecting and blocking freshly-launched exploits once all the connections are open, as well as confirming that the device does not block legitimate traffic (perhaps as a result of state tables filling up). We then launch further exploits whilst the Avalanche/Reflector devices "churn" connections at the maximum level set, ensuring that the sensor is still capable of detecting and blocking freshly-launched exploits as old connections are torn down and new ones recreated constantly.

- ***Test 3.2.1 - Attack Detection:*** *This test ensures that the sensor continues to detect new exploits as the number of open sessions is increased in stages from 10,000 to 1,000,000*
- ***Test 3.2.2 - Attack Blocking:*** *This test ensures that the sensor continues to block new exploits as the number of open sessions is increased in stages from 10,000 to 1,000,000*
- ***Test 3.2.3 - State Preservation:*** *This test ensures that the sensor maintains the state of pre-existing sessions as the number of open sessions is increased in stages from 10,000 to 1,000,000*
- ***Test 3.2.4 - Legitimate Traffic Blocking:*** *This test ensures that the sensor does not begin to block legitimate traffic as the number of open sessions is increased in stages from 10,000 to 1,000,000*

Test 3.3 - Simultaneous Open Connections (after tuning)

Test 3.2 is repeated after any tuning recommended by the vendor (if applicable) to increase the size of the state tables.

- ***Test 3.3.1 - Attack Detection:*** *As Test 3.2.1 following tuning*
- ***Test 3.3.2 - Attack Blocking:*** *As Test 3.2.2 following tuning*
- ***Test 3.3.3 - State Preservation:*** *As Test 3.2.3 following tuning*
- ***Test 3.3.4 - Legitimate Traffic Blocking:*** *As Test 3.2.4 following tuning*

Section 4 – Detection/Blocking Performance Under Load

The aim of this section is to verify that the sensor is capable of detecting and blocking exploits when subjected to increasing loads of background traffic up to the maximum bandwidth supported as claimed by the vendor.

The latest signature pack is acquired from the vendor, and sensors are deployed with **all** available attack signatures enabled (some audit/informational signatures may be disabled). Each sensor is configured to **detect and block** suspicious traffic.

Our “attacker” host launches a fixed number of exploits at a target host on the subnet being protected by the device under test. The Adtech network monitor is configured to monitor the switch SPAN port consisting of normal, exploit and background traffic, and is capable of reporting the total number of exploit packets seen on the wire as verification.

A fixed number of exploits are launched with zero background traffic to ensure the sensor is capable of detecting our baseline attacks. Once that has been established, increasing levels of varying types of background traffic are generated **through** the sensor in order to determine the point at which the sensor begins to miss attacks - all tests are repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic (or up to the maximum rated throughput of the device should this be less than 1Gbps).

At all stages, the Adtech network monitor verifies both the overall traffic loading and the total number of exploits seen on the target subnet. An additional confirmation is provided by the target host which reports the number of exploits which actually made it through.

The *Attack Blocking Rate* (ABR) at each background load is expressed as a percentage of the number of exploits blocked by the sensor (when in blocking mode) against the number verified by the Adtech network monitor and target host. The *Attack Detection Rate* (ADR) at each background load is expressed as a percentage of the number of exploits detected by the sensor (with blocking mode disabled) against the number verified by the Adtech network monitor and target host.

For each type of background traffic, we also determine the maximum load the sensor can sustain before it begins to drop packets/miss alerts. It is worth noting that devices which demonstrate 100 per cent ABR (blocking) but less than 100 per cent ADR (detection) in these tests will be prone to blocking **legitimate** traffic under similar loads.

Test 4.1 - UDP Traffic To Random Valid Ports

This test uses UDP packets of varying sizes generated by a **Smartbits SMB6000** with LAN-3301A 10/100/1000Mbps **TeraMetrics** cards installed.

A constant stream of the appropriate mix of packets - with variable source IP addresses and ports transmitting to a single fixed IP address/port - is transmitted through the sensor (bi-directionally, maximum of 1Gbps).

Each packet contains dummy data, and is targeted at a valid port on a valid IP address on the target subnet. The percentage load and packets per second (pps) figures are verified by the Adtech Gigabit network monitoring tool before each test begins. Multiple tests are run and averages taken where necessary.

This traffic does not attempt to simulate any form of “real world” network condition. The aim of this test is purely to determine the raw packet processing capability of the sensor, and its effectiveness at passing “useless” packets quickly in order to pass potential attack packets to the detection engine. The range of packet sizes has been selected to mirror the maximum, minimum and average packet sizes used in our HTTP stress tests.

- ***Test 4.1.1 - 256 byte packets - maximum 453,000 packets per second:*** *This test is roughly equivalent to a 40,000 connections per second test in our HTTP stress tests (in terms of packet size and packets per second rate), and has been included to provide an indication of the packet processing performance under the most extreme conditions for most devices - it is unlikely that any real-life network will ever see network loads of over 450,000 256-byte packets per second unless under severe DOS conditions. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic.*
- ***Test 4.1.2 - 550 byte packets - maximum 220,000 packets per second:*** *This test has been included to provide a comparison with our “real world” packet mixes, since the average packet size is similar. No sessions are created during this test and there is very little for the detection engine to do in the way of protocol analysis. This test provides a reasonable indication of the ability of a device to process packets from the wire on an “average” network, and we would expect all products to demonstrate good performance levels. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic.*
- ***Test 4.1.3 - 1000 byte packets - maximum 122,000 packets per second:*** *This test is the complete opposite of the 256 byte packet test, in that we would expect every single product to be capable of returning 100 per cent detection rates across the board when using only 1000 byte packets. We have included this test mainly to demonstrate how easy it is to achieve good results using large packets – beware of test results that **only** quote performance figures using similar (or larger) packet sizes. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic.*

Test 4.2 - HTTP “Maximum Stress” Traffic With No Transaction Delays

HTTP is the most widely used protocol in most normal networks, as well as being one of the most widely exploited. The number of potential HTTP exploits for the protocol makes a pure HTTP network something of a torture test for the average sensor.

The use of multiple Spirent Communications **Avalanche 2500** and **Reflector 2500** devices allows us to create true “real world” traffic at speeds of up to 4.2 Gbps as a background load for our tests. Our Avalanche configuration is capable of simulating over 5 million users, with over 5 million concurrent sessions, and over 200,000 HTTP requests per second.

By creating genuine session-based traffic with varying session lengths, the sensor is forced to track valid sessions, thus ensuring a higher workload than for simple packet-based background traffic. This provides a test environment that is as close to “real world” as it is possible to achieve in a lab environment, whilst ensuring absolute accuracy and repeatability.

The aim of this test is to stress the HTTP detection engine and determine how the sensor copes with detecting and blocking exploits under network loads of varying average packet size and varying connections per second.

Each transaction consists of a single HTTP GET request and there are no transaction delays (i.e. the Web server responds immediately to all requests). All packets contain valid payload (a mix of binary and ASCII objects) and address data, and this test provides an excellent representation of a live network (albeit one biased towards HTTP traffic) at various network loads.

- ***Test 4.2.1 -*** *Max 2,500 new connections per second - average packet size 1000 bytes - maximum 120,000 packets per second. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic. With relatively low connection rates and large packet sizes, we expect all sensors to achieve 100% blocking rates throughout this test.*
- ***Test 4.2.2 -*** *Max 5,000 new connections per second - average packet size 540 bytes - maximum 225,000 packets per second. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic. With average connection rates average packet sizes, this is a good approximation of a real-world production network, and we expect all sensors to perform well in this test.*
- ***Test 4.2.3 -*** *Max 10,000 new connections per second - average packet size 440 bytes - maximum 275,000 packets per second. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic. With average packet sizes coupled with very high connection rates, this is a strenuous test for any sensor, and represents a very heavily used production network.*
- ***Test 4.2.4 -*** *Max 20,000 new connections per second - average packet size 360 bytes - maximum 320,000 packets per second. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic. With small packet sizes and extremely high connection rates this is an extreme test for any sensor. Not many sensors will perform well at all levels of this test.*

Test 4.3 - HTTP “Maximum Stress” Traffic With Transaction Delays

This test is identical to Test 4.2 except that we introduce a 10 second delay in the server response for each transaction. This has the effect of maintaining a high number of open connections throughout the test, thus forcing the sensor to utilise additional resources to track those connections.

- ***Test 4.3.1 -*** *Max 5,000 new connections per second - average packet size 540 bytes - maximum 225,000 packets per second - 10 second transaction delay - maximum 50,000 open connections. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic. With average connection rates average packet sizes, this is a good approximation of a real-world production network, and we expect all sensors to perform well in this test.*
- ***Test 4.3.2 -*** *Max 10,000 new connections per second - average packet size 440 bytes - maximum 275,000 packets per second - 10 second transaction delay - maximum 100,000 open connections. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic. With average packet sizes coupled with very high connection rates, this is a strenuous test for any sensor, and represents a very heavily used production network.*

Test 4.4 - Protocol Mix Traffic

Whereas 4.2 and 4.3 provide a pure HTTP environment with varying connection rates and average packet sizes, the aim of this test is to simulate more of a “real world” environment by introducing additional protocols whilst still maintaining a precisely repeatable and consistent background traffic load (something rarely seen in a real world environment).

The result is a background traffic load that, whilst less stressful than previous tests, is closer to what may be found on a heavily-utilised “normal” production network.

- ***Test 4.4.1 -*** *72% HTTP traffic (540 byte packets) + 20% FTP traffic + 6% UDP traffic (256 byte packets). Max 4000 new connections per second - average packet size 540 bytes - maximum 215,000 packets per second - maximum 750 open connections. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic. With lower connection rates, average packets sizes and a common protocol mix, this is a good approximation of a heavily-used production network, and we expect all sensors to perform well throughout this test.*

Test 4.5 - “Real World” Traffic

This is as close as it is possible to come to a true “real world” environment under lab conditions. For this test we eliminate the Reflector device and substitute an IIS Web server installed on a dual-Xeon server with Gigabit interface and 4GB RAM. This server holds a copy of The NSS Group Web site, and is capable of handling a full 1Gbps of traffic. We then capture a typical client browsing session on the NSS Group Web site, accessing a mixture of menu pages, lengthy text-based reports and multiple graphical images (screen shots) and have Avalanche replay multiple identical sessions from up to **20 new users per second**.

It should be noted that whereas the goal of the previous tests is a very predictable, consistent and repeatable background load that never varies, the nature of this test means that traffic is slightly more “bursty” in nature.

- ***Test 4.5.1 - Pure HTTP Traffic (simulated browsing session on NSS Web site):*** *Max 4700 new connections per second - 20 new users per second - average packet size 560 bytes - maximum 210,000 packets per second.*

*Repeated with 250Mbps, 500Mbps, 750Mbps and 950Mbps of background traffic. With genuine server responses to genuine **browser** sessions consisting **of multiple transactions per** session, this is a typical "real world" background load, albeit pure HTTP. Although the Web server and the network are extremely busy at the higher traffic loads, the "normal" connection rates and packet sizes should enable most sensors to perform well at all load levels in this test.*

- ***Test 4.5.2 - Protocol Mix (72% HTTP traffic (simulated browsing sessions as 4.5.1)) + 20% FTP traffic + 6% UDP traffic (256 byte packets)):** Max 3700 new connections per second - average packet size 560 bytes - maximum 205,000 packets per second - maximum 1,500 open connections.*

  *Repeated with 250Mbps, 500Mbps, 750Mbps and 950Mbps of background traffic. With genuine server responses to genuine browser sessions consisting of multiple **transactions** per session, mixed with FTP and UDP traffic, this is a typical "real world" background load. Although the Web server and the network are extremely busy at the higher traffic loads, the "normal" connection rates and packet sizes should enable most sensors to perform well at all load levels in this test.*

To gauge the effects of varying (smaller) packet sizes, connection rates and transaction delays, the results of tests 4.2 - 4.4 should be examined.

Section 5 – Latency & User Response Times

The aim of this section is to determine the effect the sensor has on the traffic passing through it under various load conditions.

Should a device impose a high degree of latency on the packets passing through it, a network or security administrator would need to think carefully about how many devices could be installed in a single data path before user response times became unacceptable or the combination of devices caused excessive timeouts. We also determine the effect of high levels of normal HTTP traffic and a basic DOS attack on the average latency and user response times.

Test 5.1 - Latency

We use Spirent SmartFlow software and The Smartbits SMB6000 with Gigabit TeraMetrics cards to create multiple traffic flows through the appliance and measure the basic throughput, packet loss, and latency through the sensor. This test - whilst not indicative of real-life network traffic - provides an indication of how much the sensor affects the traffic flow through it. This data is particularly useful for network administrators who need to gauge the effect of any form of in-line device which is likely to be placed at critical points within the corporate network.

SmartFlow runs through several iterations of the test varying the traffic load from 250Mbps to 1Gbps bi-directionally (or up to the maximum rated throughput of the device should this be less than 1Gbps) in steps of 250Mbps. This is repeated for a range of packet sizes (256 bytes, 550 bytes and 1000 bytes) of UDP traffic with variable IP addresses and ports. At each iteration of the test, SmartFlow records the number of packets dropped, together with average and maximum latency.

- ***Test 5.1.1 - Latency With No Background Traffic:*** *SmartFlow traffic is passed across the infrastructure switches and through the device (the latency of the basic infrastructure is known and is constant throughout the tests). The packet loss and average latency are recorded at each packet size and each load level from 250Mbps to 1Gbps (in 250Mbps steps).*
- ***Test 5.1.2 - Latency With Background Traffic Load:*** *The Avalanche and Reflector are configured to generate a fixed amount of background HTTP traffic through the sensor (up to 50 per cent of the maximum rated bandwidth of the device under test - maximum 500Mbps - maximum 2,500 new connections per second - average packet size 540 bytes - maximum 112,500 packets per second).*

  A 250Mbps bi-directional load of SmartFlow traffic at various packet sizes (256 bytes, 540 bytes and 1000 bytes) is then passed across the infrastructure switches and through the device and the packet loss and average latency are recorded.
- ***Test 5.1.3 - Latency When Under Attack:*** *The Spirent WebSuite software is used to generate a fixed load of DOS/DDOS traffic of 10 per cent of the maximum rated bandwidth of the device under test (maximum 100Mbps). A 250Mbps bi-directional load of SmartFlow traffic at various packet sizes (256 bytes, 540 bytes and 1000 bytes) is then passed across the infrastructure switches and through the device and the packet loss and average latency are recorded. The device should be configured to detect/block/mitigate the DOS attack by the most efficient method available.*

Test 5.2 - User Response Times

Avalanche and Reflector devices are used to generate HTTP sessions through the device in order to gauge how any increases in latency will impact the user experience in terms of failed connections and increased Web response times.

- ***Test 5.2.1 - Web Response With No Background Traffic:*** *The Avalanche and Reflector are configured to generate HTTP traffic through the sensor (up to 50 per cent of the maximum rated bandwidth of the device under test - maximum 500Mbps - maximum 2,500 new connections per second - average packet size 540 bytes - maximum 112,500 packets per second).*

  The minimum, maximum and average page response times and number of failed connections are recorded by Avalanche to provide an indication of the expected response times under normal traffic conditions.
- ***Test 5.2.2 - Web Response When Under Attack:*** *The Avalanche and Reflector are configured to generate HTTP traffic through the sensor as for Test 5.2.1. The Spirent WebSuite software is then used to generate DOS/DDOS traffic up to 10 per cent of the maximum rated bandwidth of the device under test (maximum 100Mbps).*

  The minimum, maximum and average page response times and number of failed connections are recorded by Avalanche to provide an indication of the expected response times when the device is under attack.

Section 6 – Stability & Reliability

These tests attempt to verify the stability of the device under test under various extreme conditions. Long term stability is particularly important for an in-line IPS device, where failure can produce network outages.

- ***Test 6.1.1 - Blocking Under Extended Attack:*** *For this test, we expose the external interface of the device to a constant stream of alerts over an extended period of time. The device is configured to block and alert, and thus this test provides an indication the effectiveness of both the blocking and alert handling mechanisms. A continuous stream of exploits mixed with some legitimate sessions is transmitted through the device at a maximum of 100Mbps (max 50,000 packets per second, average packet sizes in the range of 120-350 bytes) for 8 hours with no additional background traffic. This is not intended as a stress test in terms of traffic load - merely a reliability test in terms of consistency of blocking performance.*

  The device is expected to remain operational and stable throughout this test, and to block 100 per cent of recognisable exploits, raising an alert for each. Results are presented as a simple PASS/FAIL. If any recognisable exploits are passed - caused by either the volume of traffic ***or*** *the sensor failing open for any reason - this will result in a FAIL.*

- ***Test 6.1.2 - Passing Legitimate Traffic Under Extended Attack:*** *This test is identical to 6.1.1, where we expose the external interface of the device to a constant stream of alerts over an extended period of time. The device is expected to remain operational and stable throughout this test, and to pass 100 per cent of legitimate traffic. Results are presented as a simple PASS/FAIL. If any legitimate traffic is blocked - caused by either the volume of traffic* ***or*** *the sensor failing closed for any reason - this will result in a FAIL.*

- ***Test 6.1.3 - ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC:*** *This test attempts to stress the protocol stack of the device under test by exposing it to traffic from the ISIC test tool. The ISIC test tool host is connected directly to the external interface of the sensor, and the ISIC target directly to the internal interface. ISIC traffic is transmitted through the sensor (without passing through any other network equipment) and the effects noted. Traffic load is a maximum of 350Mbps and 60,000 packets per second (average packet size is 690 bytes). Results are presented as a simple PASS/FAIL - the device is expected to remain operational and capable of detecting and blocking exploits throughout the test to attain a PASS.*

Section 7 – Management and Configuration

The aim of this section is to determine the features of the management system, together with the ability of the management port on the device under test to resist attack.

Test 7.1 - Management Port

Clearly the ability to manage the alert data collected by the sensor is a critical part of any IDS/IPS system. For this reason, an attacker could decide that it is more effective to attack the management interface of the device than the detection interface.

Given access to the management network, this interface is often more visible and more easily subverted than the detection interface, and with the management interface disabled, the administrator has no means of knowing his network is under attack.

- ***Test 7.1.1 - Open ports:*** *We will scan the open ports and active services on the management interface and report on known vulnerabilities.*
- ***Test 7.1.2 - ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC:*** *This test attempts to stress the protocol stack of the management interface of the device under test by exposing it to traffic from the ISIC test tool. The ISIC test tool host is connected directly to the management interface of the IPS sensor, and that interface is also the target. ISIC traffic is transmitted to the management interface of the IPS device (without passing through any other network equipment) and the effects noted.*

  Traffic load is a maximum of 350Mbps and 60,000 packets per second (average packet size is 690 bytes). Results are presented as a simple PASS/FAIL - the device is expected to remain (a) operational and capable of detecting and blocking exploits, and (b) capable of communicating in both directions with the management server/console throughout the test to attain a PASS.

  Test 7.1.3 - *We note whether the ISIC attacks themselves are detected by the sensor even though targeted at the management port.*

Cisco IPS-4255 V5.0(3) Test Results

Section 1 - Detection Engine

Test 1.1 – Attack Recognition	Attacks	Default ARRD	Default ARRB	Custom ARRD	Custom ARRB
Test 1.1.1 - Backdoors	7	6	6	7	7
Test 1.1.2 - WINS/DNS	3	2	2	3	3
Test 1.1.3 - DOS	10	5	9	6	10
Test 1.1.4 - False negatives (modified exploits)	14	11	11	14	14
Test 1.1.5 - Finger	4	4	4	4	4
Test 1.1.6 - FTP	5	4	4	5	5
Test 1.1.7 - HTTP	43	41	41	43	43
Test 1.1.8 - ICMP	2	2	2	2	2
Test 1.1.9 - Reconnaissance	8	5	7	6	8
Test 1.1.10 - RPC	9	7	7	9	9
Test 1.1.11 - SSH	1	1	1	1	1
Test 1.1.12 - Telnet	1	1	1	1	1
Test 1.1.13 - Database	1	1	1	1	1
Test 1.1.14 - Mail	1	1	1	1	1
Test 1.1.15 - Voice	1	1	1	1	1
Total	**110**	**92 / 110**	**98 / 110**	**104 / 110**	**110 / 110**
		84%	**89%**	**95%**	**100%**

Test 1.2 – Resistance to False Positives	Default	Custom
Test 1.2.1 - Suspicious FTP traffic	PASS	PASS
Test 1.2.2 - HTTP “exploit” using incorrect method	PASS	PASS
Test 1.2.3 - Retrieval of Web page containing “suspicious” URLs	PASS	PASS
Test 1.2.4 - Simple SMTP QUIT command	PASS	PASS
Test 1.2.5 - Normal NetBIOS copy of “suspicious” files	PASS	PASS
Test 1.2.6 - Normal NetBIOS traffic	PASS	PASS
Test 1.2.7 - POP3 e-mail containing “suspicious” URLs	PASS	PASS
Test 1.2.8 - POP3 e-mail with “suspicious” DLL attachment	PASS	PASS
Test 1.2.9 - POP3 e-mail with “suspicious” Web page attachment	PASS	PASS
Test 1.2.10 - SMTP e-mail transfer containing “suspicious” URLs	PASS	PASS
Test 1.2.11 - SMTP e-mail transfer with “suspicious” DLL attachment	PASS	PASS
Test 1.2.12 - SMTP e-mail transfer with “suspicious” Web page attachment	PASS	PASS
Test 1.2.13 - SNMP V3 packet with invalid parameter	FAIL	PASS
Test 1.2.14 - Fake DNS /bin/sh buffer overflow	PASS	PASS
Test 1.2.15 - Inter-firewall communication traffic	FAIL	PASS
Test 1.2.16 - Fake SQL Slammer traffic	PASS	PASS
Test 1.2.17 - File copy of GIF file (contains bytes which look like NOP sled)	PASS	PASS
Total Passed	**15 / 17**	**17 / 17**

Section 2 - IPS Evasion

Test 2.1 – Evasion Baselines	Detected?	Blocked?
Test 2.1.1 - NSS Back Orifice ping	YES	YES
Test 2.1.2 - Back Orifice connection	YES	YES
Test 2.1.3 - FTP CWD root	YES	YES
Test 2.1.4 - ISAPI printer overflow	YES	YES
Test 2.1.5 - Showmount export lists	YES	YES
Test 2.1.6 - Test CGI probe (/cgi-bin/test-cgi)	YES	YES
Test 2.1.7 - PHF remote command execution	YES	YES
Total	**7 / 7**	**7 / 7**

Test 2.2 – Packet Fragmentation/Stream Segmentation	Detected?	Decoded?	Blocked?
Test 2.2.1 - IP fragmentation - ordered 8 byte fragments	YES	YES[2]	YES
Test 2.2.2 - IP fragmentation - ordered 24 byte fragments	YES	YES	YES
Test 2.2.3 - IP fragmentation - out of order 8 byte fragments	YES	YES[2]	YES
Test 2.2.4 - IP fragmentation - ordered 8 byte fragments, duplicate last packet	YES	YES	YES
Test 2.2.5 - IP fragmentation - out of order 8 byte fragments, duplicate last packet	YES	YES	YES
Test 2.2.6 - IP fragmentation - ordered 8 byte fragments, reorder fragments in reverse	YES	YES	YES
Test 2.2.7 - IP fragmentation - ordered 16 byte fragments, fragment overlap (favour new)	YES	YES[2]	YES
Test 2.2.8 - IP fragmentation - ordered 16 byte fragments, fragment overlap (favour old)	YES	YES	YES
Test 2.2.9 - TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with invalid TCP checksums	YES	YES	YES
Test 2.2.10 - TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with null TCP control flags	YES	YES	YES
Test 2.2.11 - TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with requests to resync sequence nos. mid-stream	YES	YES	YES
Test 2.2.12 - TCP segmentation - ordered 1 byte segments, duplicate last packet	YES	YES	YES
Test 2.2.13 - TCP segmentation - ordered 2 byte segments, segment overlap (favour new)	YES	YES[2]	YES
Test 2.2.14 - TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with out-of-window sequence numbers	YES	YES	YES
Test 2.2.15 - TCP segmentation - out of order 1 byte segments	YES	YES	YES[1]
Test 2.2.16 - TCP segmentation - out of order 1 byte segments, interleaved duplicate segments with faked retransmits	YES	NO	YES[1]
Test 2.2.17 - TCP segmentation - ordered 1 byte segments, segment overlap (favour new)	YES	YES[2]	YES
Test 2.2.18 - TCP segmentation - out of order 1 byte segments, PAWS elimination (interleaved dup segments with older TCP timestamp options)	YES	NO	YES[1]
Test 2.2.19 - IP fragmentation - out of order 8 byte fragments, interleaved duplicate packets scheduled for later delivery	YES	YES	YES
Test 2.2.20 - TCP segmentation - ordered 16 byte segments, segment overlap (favour new (Unix))	YES	YES	YES
Total	**20 / 20**	**18 / 20**	**20 / 20**

Test 2.3 – URL Obfuscation	Detected?	Decoded?	Blocked?
Test 2.3.1 - URL encoding	YES	YES	YES
Test 2.3.2 - /./ directory insertion	YES	YES	YES
Test 2.3.3 - Premature URL ending	YES	YES	YES
Test 2.3.4 - Long URL	YES	YES	YES
Test 2.3.5 - Fake parameter	YES	YES	YES
Test 2.3.6 - TAB separation	YES	YES	YES
Test 2.3.7 - Case sensitivity	YES	YES	YES
Test 2.3.8 - Windows \ delimiter	YES	YES	YES
Test 2.3.9 - Session splicing	YES	YES	YES
Total	**9 / 9**	**9 / 9**	**9 / 9**

Test 2.4 – Miscellaneous Obfuscation Techniques	Detected?	Decoded?	Blocked?
Test 2.4.1 - Altering default ports	YES	YES	YES
Test 2.4.2 - Inserting spaces in FTP command lines	YES	YES	YES
Test 2.4.3 - Inserting non-text Telnet opcodes in FTP data stream	YES	YES	YES
Test 2.4.4 - Polymorphic mutation (ADMmutate)	YES	YES	YES
Test 2.4.5 - Altering protocol and RPC PROC numbers	YES	YES	YES
Test 2.4.6 - RPC record fragging (MS-RPC and Sun)	YES	YES	YES
Test 2.4.7 - HTTP exploits to port <> 80	YES	YES	YES
Total	**7 / 7**	**7 / 7**	**7 / 7**

Section 3 - Stateful Operation

Test 3.1 – Stateless Attack Replay	Alert?	Blocked?	Pass/Fail
Test 3.1.1 - Stateless Web exploits	NO	YES	PASS
Test 3.1.2 - Stateless FTP exploits	NO	YES	PASS

Test 3.2 – Simultaneous Open Connections (default settings)							
Number of open connections	10,000	25,000	50,000	100,000	250,000	500,000	1,000,000
Test 3.2.1 - Attack Detection	PASS	PASS	PASS	PASS	PASS	PASS	PASS
Test 3.2.2 - Attack Blocking	PASS	PASS	PASS	PASS	PASS	PASS	PASS
Test 3.2.3 - State Preservation	PASS	PASS	PASS	PASS	PASS	PASS	FAIL
Test 3.2.4 - Legitimate traffic blocking	PASS	PASS	PASS	PASS	PASS	PASS	PASS

Test 3.3 – Simultaneous Open Connections (after tuning)							
Number of open connections	10,000	25,000	50,000	100,000	250,000	500,000	1,000,000
Test 3.3.1 - Attack Detection	PASS	PASS	PASS	PASS	PASS	PASS	PASS
Test 3.3.2 - Attack Blocking	PASS	PASS	PASS	PASS	PASS	PASS	PASS
Test 3.3.3 - State Preservation	PASS	PASS	PASS	PASS	PASS	PASS	FAIL
Test 3.3.4 - Legitimate traffic blocking	PASS	PASS	PASS	PASS	PASS	PASS	PASS

Section 4 - Detection/Blocking Performance Under Load

Test 4.1 – UDP traffic to random valid ports		125Mbps	250Mbps	375Mbps	500Mbps	Max
Test 4.1.1 - 256 byte packet test - max 226,500pps	Detected	100%	100%	100%	82%	423Mbps
	Blocked	100%	100%	100%	100%	
Test 4.1.2 - 550 byte packet test - max 110,000pps	Detected	100%	100%	100%	100%	500Mbps
	Blocked	100%	100%	100%	100%	
Test 4.1.3 - 1514 byte packet test - max 61,000pps	Detected	100%	100%	100%	100%	500Mbps
	Blocked	100%	100%	100%	100%	

Test 4.2 – HTTP "maximum stress" traffic with no transaction delays		125Mbps	250Mbps	375Mbps	500Mbps	Max
Test 4.2.1 - Max 1250 connections per second - ave packet size 1000 bytes - max 60,000 packets per second	Detected	100%	100%	100%	100%	500Mbps
	Blocked	100%	100%	100%	100%	
Test 4.2.2 - Max 2500 connections per second - ave packet size 540 bytes - max 112,500 packets per second	Detected	100%	100%	100%	100%	500Mbps
	Blocked	100%	100%	100%	100%	
Test 4.2.3 - Max 5000 connections per second - ave packet size 440 bytes - max 137,500 packets per second	Detected	100%	100%	100%	100%	500Mbps
	Blocked	100%	100%	100%	100%	
Test 4.2.4 - Max 10000 connections per second - ave packet size 360 bytes - max 160,000 packets per second	Detected	100%	100%	100%	N/A[3]	400Mbps
	Blocked	100%	100%	100%	N/A[3]	

Test 4.3 – HTTP "maximum stress" traffic with transaction delays		125Mbps	250Mbps	375Mbps	500Mbps	Max
Test 4.3.1 - Max 2500 connections per second - ave packet size 540 bytes - max 112,500 packets per second - 10 sec delay - max 25,000 open connections	Detected	100%	100%	100%	100%	500Mbps
	Blocked	100%	100%	100%	100%	
Test 4.3.2 - Max 5000 connections per second - ave packet size 440 bytes - max 137,500 packets per second - 10 sec delay - max 50,000 open connections	Detected	100%	100%	100%	100%	500Mbps
	Blocked	100%	100%	100%	100%	

Test 4.4 – Protocol mix		125Mbps	250Mbps	375Mbps	500Mbps	Max
Test 4.4.1 - 72% HTTP (540 byte packets) + 20% FTP + 6% UDP (256 byte packets). Max 2000 connections per second - ave packet size 540 bytes - max 107,500 packets per second - max 375 open connections	Detected	100%	100%	100%	100%	500Mbps
	Blocked	100%	100%	100%	100%	

Test 4.5 – Real World traffic		125Mbps	250Mbps	375Mbps	500Mbps	Max
Test 4.5.1 - Pure HTTP (simulated browsing session on NSS Web site). Max 2350 connections per second - 10 new users per second - ave packet size 560 bytes - max 105,000 packets per second	Detected	100%	100%	100%	100%	500Mbps
	Blocked	100%	100%	100%	100%	
Test 4.5.2 - Protocol mix - 72% HTTP (simulated browsing sessions as 2.5.1) + 20% FTP + 6% UDP (256 byte packets). Max 1850 connections per second - ave packet size 560 bytes - max 102,500 packets per second - max 750 open connections	Detected	100%	100%	100%	100%	500Mbps
	Blocked	100%	100%	100%	100%	

Section 5 - Latency & User Response Times

Test 5.1 – Latency	Packet Size	125Mbps	250Mbps	375Mbps	500Mbps
Test 5.1.1 Average latency (µs) with no background traffic	256	112.03	128.45	179.28	N/A[4]
	550	131.42	143.30	159.10	171.89
	1000	169.85	178.08	192.54	209.52
Test 5.1.2 Average latency (µs) with background traffic (250Mbps HTTP traffic, max 1250 connections per second - ave packet size 540 bytes - max 56,250 packets per second)	256	447.68			
	550	385.76			
	1000	397.80			
Test 5.1.3 Average latency (µs) when under attack (50Mbps SYN flood (74,000cps))	256	119.34			
	550	139.05			
	1000	177.96			

Test 5.2 – User Response Times	Attempted Trans	Failed Trans	Min Page Response	Max Page Response	Ave Page Response
Test 5.2.1 - Web page response (ms) with no background traffic (250Mbps HTTP traffic, max 1250 connections per sec - ave packet size 540 bytes - max 56,250 packets per sec)	789935	0	201	229	204
Test 5.2.2 - Web page response (ms) when under attack (250Mbps HTTP traffic, max 1250 connections per sec - ave packet size 540 bytes - max 56,250 packets per sec PLUS 50Mbps SYN flood (74,000cps))	772581	0	201	233	205

Section 6 - Stability & Reliability

Test ID	Result
Test 6.1.1 - Blocking Under Extended Attack	100%
Test 6.1.2 - Passing legitimate traffic under extended attack	100%
Test 6.1.3 - ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC	PASS

Section 7 - Management Interface

Test ID	Result
Test 7.1.1 - Open Ports	PASS
Test 7.1.2 - ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC	PASS
Test 7.1.3 - ISIC attacks detected against management interface?	NO

Notes:

1. Always blocked by normaliser (silent block or normaliser alert)
2. Only decoded successfully when not blocking - with blocking turned on, only normaliser alerts are raised
3. Not possible to run this test - maximum HTTP connection rate is approx 8,000cps
4. Not possible to run this test - dropped packets at 100% load

Section 1: Detection Engine

We installed one sensor with the latest signature pack, and configured it with all attack signatures enabled, plus some key audit/information-only signatures. All "*retired*" signatures were un-retired and enabled (this is not recommended in normal deployments, since most signatures are retired for a reason, and reinstating them simply serves to increase the memory footprint of the sensor application) .

Out of the box, blocking performance was excellent at 89 per cent, and was improved to a perfect 100 per cent following the application of a signature update after 48 hours.

Detection/recognition rate was slightly lower (95 per cent following update) due to the fact that some protocol anomalies (i.e. overlapping fragments, invalid TCP options, etc.) are always blocked without alerting by the normaliser engine. It is possible to configure normaliser signatures to produce alerts if required.

We noted a minimum of "noise", with few test cases raising multiple alerts for a single exploit, and the accuracy of the exploit descriptions was high. Performance in our "false negative" tests was very good out of the box, and there is every indication that, wherever possible, signatures are written for the underlying vulnerability rather than specific exploits.

A major concern in deploying an IPS is the blocking of legitimate traffic, and the IPS-4255 did block two of our false positive test cases out of the box. This was rectified following the signature update, but we continued to note numerous port-based backdoor alerts during our FTP tests.

The Cisco IPS arrives with a sensible default policy with PASS and BLOCK actions set for appropriate signatures (i.e. where the confidence level of a signature is high then the action will generally be set to BLOCK). It should be possible to deploy this product successfully using the default settings in most organisations.

Section 2: IPS Evasion

Our extensive testing initially uncovered a problem in the normaliser engine which caused sensor crashes, but this was fixed during the testing. Users should ensure that they have the latest engine update installed.

Resistance to known evasion techniques was excellent, with the IPS-4255 achieving a clean sweep across the board in all our evasion tests. *Fragroute, Whisker, ADMmutate* and even *RPC record fragging* all failed to trick the IPS-4255 into ignoring valid attacks.

Not only were the fragmented and obfuscated attacks blocked successfully, but all but two of them were decoded accurately as well when not blocking.

Once blocking was enabled, certain of the more complex fragmentation and segmentation evasion techniques were no longer decoded fully, but instead raised only alerts from the normaliser. Naturally, this has no effect on either blocking ability nor the administrator's ability to determine that an attack is underway.

Section 3: Stateful Operation

Out of the box, the IPS-4255 handled 500,000 open connections (the claimed maximum). No tuning parameters are available to the administrator to adjust this.

Default operation of the device is to age out old connections when the state tables are full or resources are low. This means that it is technically possible to evade the IPS-4255 once the state tables are full, since it will allow attack traffic from aged-out connections at that point. Our half-open exploit was not detected when completed, for example. This behaviour is not configurable.

Stateless “exploits” are not alerted upon (this is correct behaviour in order to be resistant to *Stick* and *Snot* tools), and mid-flows are blocked by default. It is possible to configure the device to allow mid-flows by disabling a single normaliser signature, but Cisco does not recommend this. Alerts are raised by the normaliser when mid-flows are blocked, but these can easily be disabled if required.

Section 4: Detection/Blocking Performance Under Load

Note that the Cisco IPS-4255 was tested as a 500Mbps IPS device.

Performance at almost all levels of our load tests was good, with 100 per cent of all attacks being detected and blocked under most load conditions.

The one exception was Test 4.2.4 where we determined that there was a limit of approximately 8,000 HTTP connection per second (equating to approximately 400Mbps) meaning the test could not be completed. Beyond this limit, legitimate connections began to fail, but all attack traffic was still blocked successfully.

However, we would happily confirm Cisco’s 500Mbps rating for this device under normal network conditions.

Section 5: Latency & User Response Times

Basic latency figures were good for a device of this type at most traffic loads and packet sizes, ranging from 112µs with 125Mbps of 256 byte packets, to 209µs with 500Mbps of 1000 byte packets.

Behaviour through most of the tests with no background traffic was very predictable, with relatively small increases in latency as traffic levels increased from 125Mbps to 500Mbps across each packet size. The one exception was at 500Mbps of 256 byte packets, where the device began to drop packets, making the latency figures unusable.

Placing the device under a half load of 250Mbps of HTTP traffic we noted significant increases of almost 300 per cent with 256 byte packets (112µs to 447µs), almost 200 per cent with 550 byte packets (131µs to 385µs), and 134 per cent with 1000 byte packets (169µs to 397µs).

However, we would still consider these results to be acceptable for a 500Mbps device. HTTP response times were good, and in most deployments the IPS-4255 could be situated anywhere on a network with no more than 500Mbps of traffic, either internally or at the perimeter.

SYN Flood protection is implemented for the sensor, and 50Mbps of SYN flood traffic had no significant effect on the IPS-4255. Latency increased by only a few microseconds across all packet sizes and HTTP response times were barely affected. Note, however, that the SYN Flood was **not** mitigated at all for the target server (this feature is planned for a future release) although appropriate alerts were raised.

Section 6: Stability & Reliability

The IPS-4255 performed consistently and completely reliably throughout our tests. Under eight hours of extended attack (comprising millions of exploits mixed with genuine traffic) it continued to block 100 per cent of attack traffic, whilst passing 100 per cent of legitimate traffic.

Exposing the sensor interface to ISIC-generated traffic had no adverse effect, and the device continued to detect and block all other exploits throughout and following the ISIC attack.

Section 7: Management Interface

Open ports on the management interface are restricted to TCP/22 (SSH) and TCP/443 (HTTPS).

The extended ISIC attack against the management interface had no effect on the appliance and its ability to detect and block attacks (although communications between the console and sensor were sluggish during the attack). No alerts were raised during the attack.

The sensor continued to work perfectly throughout and following the ISIC attack, and there were no residual stability problems.