

Westline Athena Aegis IPS 510L V2.1

Technical Evaluation

An NSS Group Report



First published May 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>

©1991-2005 The NSS Group

All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the authors. This report shall be treated at all times as a confidential and proprietary report for internal use only.

Please note that access to or use of this Report is conditioned on the following:

1. The information in this Report is subject to change by The NSS Group without notice.
2. The information in this Report is believed by The NSS Group to be accurate and reliable, but is not guaranteed. All use of and reliance on this Report are at your sole risk. The NSS Group is not liable or responsible for any damages, losses or expenses arising from any error or omission in this Report.
3. NO WARRANTIES, EXPRESS OR IMPLIED ARE GIVEN BY THE NSS GROUP. ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT ARE DISCLAIMED AND EXCLUDED BY THE NSS GROUP. IN NO EVENT SHALL THE NSS GROUP BE LIABLE FOR ANY CONSEQUENTIAL, INCIDENTAL OR INDIRECT DAMAGES, OR FOR ANY LOSS OF PROFIT, REVENUE, DATA, COMPUTER PROGRAMS, OR OTHER ASSETS, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.
4. This Report does not constitute an endorsement, recommendation or guarantee of any of the products (hardware or software) tested or the hardware and software used in testing the products. The testing does not guarantee that there are no errors or defects in the products, or that the products will meet your expectations, requirements, needs or specifications, or that they will operate without interruption.
5. This Report does not imply any endorsement, sponsorship, affiliation or verification by or with any companies mentioned in this report.
6. All trademarks, service marks, and trade names used in this Report are the trademarks, service marks, and trade names of their respective owners, and no endorsement of, sponsorship of, affiliation with, or involvement in, any of the testing, this Report or The NSS Group is implied, nor should it be inferred.

TABLE OF CONTENTS

INTRODUCTION	1
Intrusion Prevention Systems (IPS)	1
Host IPS (HIPS).....	2
Network IPS (NIPS).....	2
Rate-Based IPS (Attack Mitigator)	3
Detection Methods.....	3
Pattern Matching	4
Stateful Pattern Matching	4
Protocol Decode	5
Heuristic Analysis	7
Anomaly Analysis	7
Which Detection Method Is The Best?	7
Implementation Challenges.....	8
Requirements for effective prevention.....	9
The NSS Intrusion Prevention Group Test.....	10
Performance	11
Security Effectiveness	14
Usability	16
WESTLINE ATHENA AEGIS IPS 510L V2.1	17
Executive Summary.....	17
Architecture.....	17
Athena Aegis IPS Appliance.....	17
Intrusion Management Centre (IMC).....	18
JConsole.....	19
Performance	19
Security Effectiveness	20
Usability	21
Installation.....	21
Configuration	22
Policy Management.....	25
Alert Handling	30
Reporting and Analysis.....	33
Verdict.....	35
Contact Details	37
APPENDIX A – TEST RESULTS.....	38
The Test Environment	38
Section 1 – Detection Engine	38
Section 2 – Evasion	40
Section 3 – Stateful Operation.....	42
Section 4 – Detection/Blocking Performance Under Load	44
Section 5 – Latency & User Response Times.....	48
Section 6 – Stability & Reliability	50
Section 7 – Management and Configuration.....	50
Westline Athena Aegis IPS 510L V2.1 Test Results	52
Section 1 - Detection Engine	52
Section 2 - IPS Evasion.....	52
Section 3 - Stateful Operation	54
Section 4 - Detection/Blocking Performance Under Load.....	54
Section 5 - Latency & User Response Times	55
Section 6 - Stability & Reliability	55
Section 7 - Management Interface	55

TABLE OF FIGURES

Figure 1 - Aegis IPS: Architecture.....	18
Figure 2 - Aegis IPS: Viewing alert details in the multi-pane JConsole.....	23
Figure 3 - Aegis IPS: Monitoring sites within JConsole.....	24
Figure 4 - Aegis IPS: Viewing/Editing Attack Objects	26
Figure 5 - Aegis IPS: Policy editor	27
Figure 6 - Aegis IPS: Editing Rules.....	28
Figure 7 - Aegis IPS: Installing Policies	29
Figure 8 - Aegis IPS: Alert Summary	30
Figure 9 - Aegis IPS: Alert Monitor.....	31
Figure 10 - Aegis IPS: Alert Analysis	32
Figure 11 - Aegis IPS: Specifying report criteria	33
Figure 12 - Aegis IPS: Typical report.....	34

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 reboot. 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:

$$\text{throughput} = \text{window} / \text{RTT}$$

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*.

WESTLINE ATHENA AEGIS IPS 510L V2.1

Executive Summary

Westline offers a range of Athena Aegis appliances from dual 10/100Mbps to 5 x Gigabit interfaces. Although the IPS 510L tested is Gigabit-capable (only one network segment can be monitored in-line, despite sporting two fibre and two copper Gigabit interfaces), Westline rates it at only 200Mbps when all signatures are enabled. This could make the cost per megabit/cost per port higher than it should be for a device of this type.

Overall, the performance of Athena Aegis IPS510L is very good providing the limit of 200Mbps is not exceeded. Blocking rates were adequate out of the box, and improved to over 90 per cent following an update, and resistance to the most common evasion techniques was excellent. Throughput and latency were excellent under all normal network loads (up to 200Mbps) and across all packet sizes.

We also found the IPS510L to be very stable and reliable under extended attack, although resistance to high-levels of DOS/DDOS attacks (such as SYN floods) was poor. We would recommend deploying this device behind an additional layer of protection, such as an attack mitigator or firewall with DOS/DDOS mitigation capabilities.

The device can be managed via the Java-base JConsole connecting to the Intrusion Management Centre (IMC). The IMC can be located either on a dedicated management server, or on the sensor, or both at the same time, and the management system can be deployed as a two-tier or three-tier or hybrid (mixture of the two) system. Alert handling is adequate and reporting is fairly basic.

Architecture

There are three main components to the Athena Aegis IPS system:

- [Athena Aegis IPS Appliance](#)
- [Intrusion Management Centre \(IMC\)](#)
- [JConsole management client](#)

Athena Aegis IPS Appliance

Athena Aegis IPS is offered in a range of appliances supporting various bandwidth requirements and physical port arrangements, each of which can be configured in transparent or gateway (routed) mode. The latter mode of operation makes use of the integrated stateful firewall, and allows the Aegis IPS to be installed as the main perimeter security device if required (though we would not recommend this with the current version due to lack of protection against high-volume DOS/DDOS attacks).

The product submitted for testing was the Aegis IPS 510L, a 1U rack-mount appliance with two fibre and two copper Gigabit monitoring ports.

Despite this physical port arrangement, only one port pair can be used to monitor a single in-line segment at a time - the remaining ports will remain idle. In addition, despite being Gigabit-capable, Westline rates the device at just 200Mbps with typical network traffic when all signatures are enabled.

Also on the front panel is a dedicated 10/100Mbps Ethernet management port, a serial console port, port indicator LEDs, and a two-line LCD display with four buttons. In normal use, the LCD panel will show the Aegis IPS general information such as IPS version, engine status and CPU utilisation.

Via the four buttons, it can also be used to perform some basic initial configuration of the device - including setting the management IP address and netmask - as well as shutting down/restarting the sensor.

Standard hardware is used throughout the Aegis IPS product range, based on Intel hardware with 512MB to 2GB memory, depending on the model. No network processors or other hardware acceleration techniques are employed in the current version. The appliance operating system is custom written, which has allowed Westline to take full advantage of multi-processor systems to provide maximum packet processing performance.

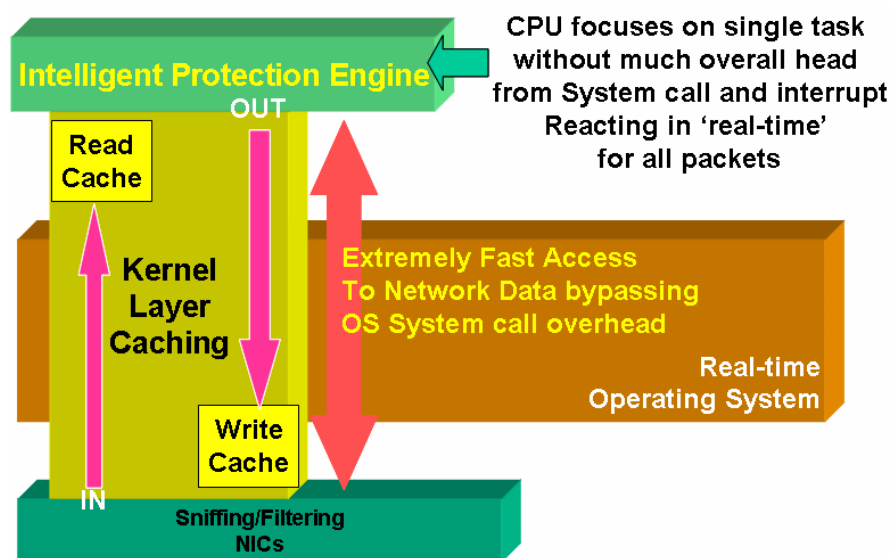


Figure 1 - Aegis IPS: Architecture

In order to make use of SMP architectures, the IPS engine (originally based on open source software) is multithreaded, dedicating one CPU to capturing packets, and another for processing and sending/dropping them. Unlike typical Linux-based operating systems, this all happens at the kernel level to maximise performance and minimise latency when under load. When the engine is slightly overloaded, the cache can help minimise packet drops.

There is no redundancy built in to the device (certain appliances feature dual power supplies), and nor is there a High Availability solution on offer at present. Westline does, however, offer an optional copper bypass device to allow traffic to pass unimpeded should the Aegis IPS appliance fail.

Intrusion Management Centre (IMC)

The *Intrusion Management Centre* (IMC) is the server component which provides alerting, reporting and device management capabilities.

A unique feature of Aegis IPS is that the IMC software can be built-in to the sensor or external. Each sensor comes with a built-in management server which is enabled by default, allowing one or more consoles to connect directly to the sensor out of the box, thus operating in a two-tier mode.

In addition, identical IMC software can be installed on an external server running Red Hat Linux with a MySQL database for centralised provisioning, consolidated reporting and multi-sensor management from a single console.

This configuration is known as hybrid mode, since both two-tier and three-tier modes are supported simultaneously, allowing the console software to connect either to the IMC server, or to the sensor directly, or both as required. In this mode, alerts are stored locally on the sensor as well as centrally on the IMC server.

The built-in IMC software is identical to the external version in all respects apart from the lack of any scheduling capability. Each sensor comes with this software on the installation CD and a free single-sensor management license is provided. Licenses to manage multiple sensors can be purchased as required.

The Aegis IPS can also be configured with a pure 3-tier architecture. In this mode, the internal IMC software is disabled completely and the console can no longer connect to the sensor directly, but only via the central IMC server. Whilst this can provide a slight performance boost in heavily loaded networks, the downside with this mode is that alerts are no longer stored locally on the sensor. Thus if communication between IMC server and sensor is interrupted for any reason, alerts will be lost.

Multiple read-only connections can be made to any IMC, meaning that multiple consoles can connect to an IMC server or to a sensor directly, and that multiple IMC servers can connect to each sensor. Read/write management control is achieved by an individual console acquiring a single management token per sensor.

JConsole

This is a Java-based console which provides a graphical management interface for the Aegis IPS system. JConsole can connect to the IMC software over a secure channel, either directly to the Aegis IPS sensor, or to a central IMC server.

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 Athena Aegis IPS 510L was tested up to 200Mbps, the rated speed of the device, and performance at all levels of our load tests was impeccable, with 100 per cent of all attacks being detected and blocked under all load conditions. We would happily confirm Westline's 200Mbps rating for this device.

The basic latency figures of the Athena Aegis IPS 510L were acceptable across the board under all traffic loads. They ranged from 315µs with 50Mbps of 256 byte packets, to 416µs with 200Mbps of 1000 byte packets. Behaviour throughout the tests with no background traffic was consistent and predictable, with small increases as additional network load was applied from 50Mbps to 200Mbps.

Placing the device under a half load of 100Mbps of HTTP traffic actually had the effect of reducing latency further, falling from 315µs to 156µs with 256 byte packets, 343µs to 187µs with 550 byte packets, and from 403µs to 244µs with 1000 byte packets. The effect of this is that latency under normal traffic loads is excellent for a device of this type - HTTP response times, for example, were also excellent.

The device did have problems with SYN flood traffic, however, and would perform best when situated behind an attack mitigation device or a firewall with SYN flood mitigation capabilities. Latency when under 20Mbps of SYN flood traffic was excessive, as was packet loss. The overall effect was an increase in response times and a high number of failed transactions.

This is an issue which Westline recognises and is working on at the time of writing. Future releases will include per-IP based rate limiting capabilities to help mitigate such attacks. For now, our recommendation would be to deploy the device as mentioned above, with additional protection in front.

Athena Aegis 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 99.9975 per cent of legitimate traffic. This is considered acceptable for a device of this type.

Exposing the sensor interface to ISIC-generated traffic had no long-term adverse effect on the device, although communications between the management console and the management server/sensor were interrupted during the 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 updates, using a slightly modified version of the *Detect_All* policy, disabling *Low* severity informational signatures only. All exploit signatures are enabled in this policy.

Out of the box, signature recognition was adequate at 75 per cent, and was improved to 88 per cent following the application of a signature update after 48 hours. Blocking performance was actually four per cent higher throughout - giving a much more creditable final result of 92 per cent - due to various types of malformed/invalid packets (such as those used in our DOS test suite) being blocked silently.

Performance in our “false negative” tests was adequate out of the box, and improved slightly following the signature update. However, there were still four misses out of the 14 test cases following the signature update, indicating that many signatures are written for specific exploits rather than for the underlying vulnerability.

Having said that, we did notice that where a deliberate attempt was not made to evade the device, Athena Aegis IPS did extremely well at detecting most of the variants that we include alongside our basic test cases (i.e. if five different exploits are known for a vulnerability, the signature is generally written well enough to detect all of them).

A major concern in deploying an IPS is the blocking of legitimate traffic. The device failed only one test case in our false positive test suite, and this was rectified following the signature update. No other false positives were noted during stress testing with either *Avalanche* traffic or real-world traffic mixes.

Resistance to known evasion techniques was very good, with the Athena Aegis IPS achieving a clean sweep across the board in most of our evasion tests. *Fragroute* and *Whisker* both failed to deceive the device into ignoring valid attacks, and all of the attempts were decoded accurately.

Of the miscellaneous evasion techniques, non-text Telnet opcodes in FTP data streams and RPC fragmentation both proved troublesome. Full RPC and FTP protocol parsers are under development for a future release.

Out of the box, Westline claims that Athena Aegis IPS can handle just over 250,000 open connections, and this was verified in our tests. The default operation of the device is to reject new connections when the state tables are full or resources are low, and this meant that once we exceeded the connection limit the device continued to maintain state on existing connections (thus detecting our half-open exploit), but inevitably, as a consequence, began to block legitimate traffic. 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 (a mid-flow violation alert is raised, if configured). It is not possible to configure the device to allow mid-flows.

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 of the Aegis IPS appliance is very straightforward, even though it is not possible to perform the entire initial configuration via the front-panel LCD display. Instead, the serial console port is used, which brings up a menu offering the following options:

- *Change IPS system mode (transparent or gateway/routed mode)*
- *System reboot*
- *System shutdown*
- *Change IP address of administrative interface*
- *Sensor information*
- *Diagnostic information*

- *Restore default configuration*
- *Change time zone of the system*
- *Set up default gateway*
- *Change time of the day*
- *Display current time*
- *Change sensor to use external IMC*
- *Change sensor to use built-in IMC*
- *Configure third-party firewall*
- *Set up DNS server*
- *Set up DMZ and external interfaces (routed/gateway mode only)*
- *Set up routes*
- *Set up allowed IPs for management*
- *Change monitoring port pair (copper or fibre)*

Note that there is no option to set a password or secret key for sensor-to-IMC/console communications. Instead, all sensors have a default password which is used to establish initial connection to the IMC or console. We consider this to be too insecure, even though it is necessary to specifically define which hosts are allowed to connect to the sensor for management (thus making it harder for a rogue console to gain access to a sensor). The user name and password can be changed via the GUI once initial connection has been established.

If an external IMC is required, this can be installed on a Red Hat Linux platform which must be supplied by the end user (Westline sells the IMC as a software-only product). Although the procedure is documented in the manual, we felt it was not quite automated enough, requiring several commands to be issued at the command line (including separate and unassisted installation of the MySQL database) rather than wrapping everything in a single, automatic installation routine.

Installation of the JConsole software is a very simple process, especially since there is no requirement for a separate Java Runtime Environment (JRE) to be installed in order for it to operate.

Documentation is adequate, consisting of separate PDF-based user guides for the sensor and IMC.

Configuration

The *JConsole* management GUI consists of three panes, any of which can be disabled to create more screen real estate.

The pane on the left side of the screen is called the *Function Window*, and contains a hierarchical menu of “sites” (Westline’s term for devices - whether IMC servers or sensors) and management options within those sites. The right-hand pane is the *Management Area*, which is used to configure whichever option is currently selected in the Function Window.

As entries are selected in the Function Window, a new tab is opened in the Management Area, and any number of these tabs can remain open allowing the administrator to switch quickly between them. Below these, running the full width of the screen, is the *Message Window*, which displays messages relating to sensors and IMC servers in real time.

A *site* is actually an instance of the IMC, and thus can either be a single sensor, or a central management server which controls multiple sensors. A site can be added to a console by anyone who knows the IP address, user name and password of that device, and whose own IP address has been defined as an allowed management console at the sensor. There is no other control, meaning that it is possible for two different administrators to define the same sensor within their own console and configure it independently of each other.

The only protection against this situation is that, although multiple read-only access is allowed in all cases using only the site password, write access is only granted to the person who has acquired the “management token” for that site. This is accomplished via a menu option within the console (it is released in the same way, or by exiting the console), and prevents two people from actually updating a sensor at the same time. However there is no protection from two separate administrators updating the same sensor with different policies at different times - a potentially dangerous situation from a security standpoint, made worse by the fact that there is no audit log available listing what changes were made to which sensors and policies.

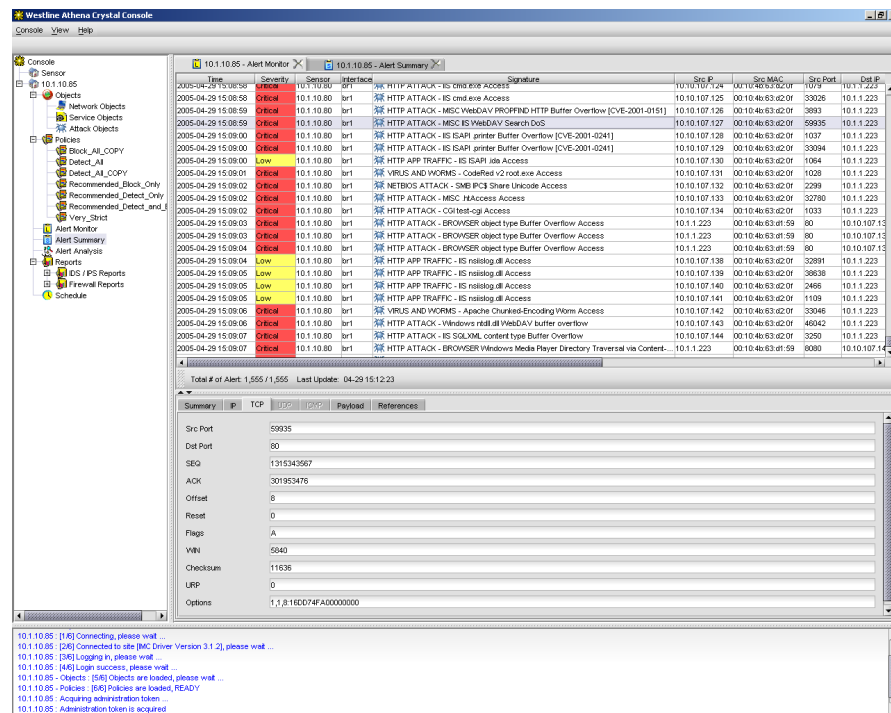


Figure 2 - Aegis IPS: Viewing alert details in the multi-pane JConsole

At the moment we feel that Westline's definition of a site is flawed. A site should be just that - a central site containing multiple devices (sensors and/or management servers) which are defined within it. Devices created within a site should be assigned to that site only, and should then be unavailable for management stand-alone or as part of a site created on another management server.

Each console is also protected by a simple password (not even a user name), which then provides unlimited access to the sites defined within that console. There is no role-based security model in place to ensure that certain administrators have access to certain sites, or are restricted to read-only access, for example.

As it stands, we feel that the Westline management model is lightweight, and suited only to small-scale deployments where there is a single administrator managing a small number of sensors. This is unfortunate, since the ability to operate the system in either two-tier, three-tier or hybrid mode makes for a very flexible and scalable system otherwise. Hopefully this will be improved in future versions, which will make the management system very much more powerful and flexible overall.

On first entering the console, the administrator is presented with a list of available sensors and/or IMC servers in the Function Window. Selecting any one of these allows him to connect to that device (more than one device can be connected from a single console) which brings up the hierarchical menu list beneath. Note that the menu list is duplicated under every connected device, and when in hybrid mode it would be possible to make changes to both the IMC server **and** the server(s) which the IMC controls.

It is up to the administrator to remember which IMC controls which sensors and ensure that changes are made in one place only. Once again, we feel that it would be better if these relationships were enforced by the IMC and console to prevent potentially conflicting changes being made in more than one place.

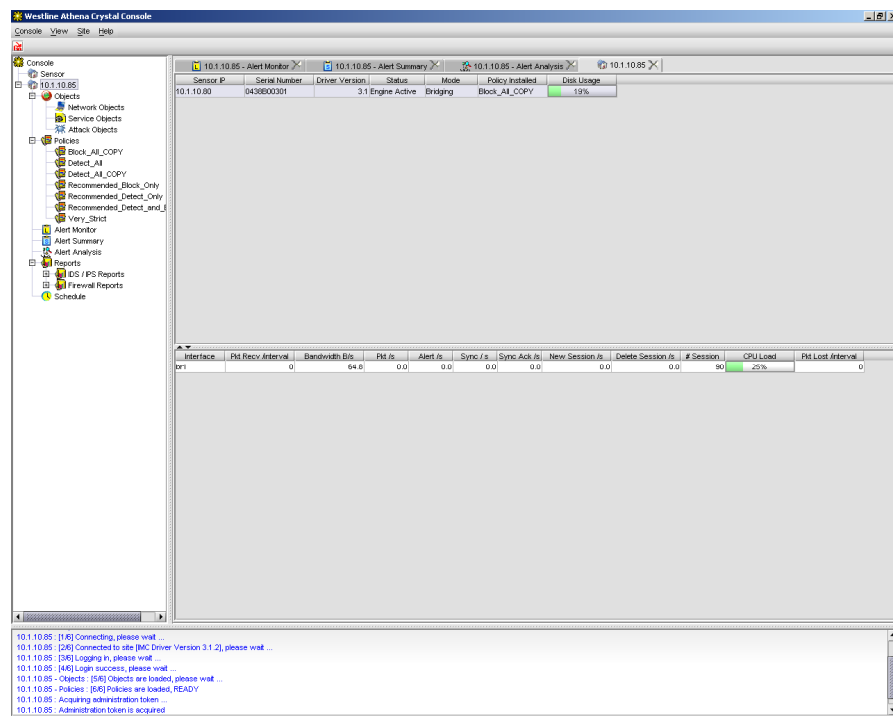


Figure 3 - Aegis IPS: Monitoring sites within JConsole

Single sensor details can be viewed by double-clicking the sensor's menu entry, and a list of sensors within a site can be viewed by double-clicking the site entry. We found the interface to be counterintuitive, on occasion - sometimes requiring a right-click, and other times requiring a double-click. Details of all available sensors in a site are listed in a table, including:

- *Sensor IP*
- *Serial number*
- *Driver version*

- *Status*
- *Mode*
- *Policy installed*
- *Disk usage*

By selecting individual sensors, it is possible to view real-time statistics, including:

- *Interface used for monitoring*
- *Packets received per interval*
- *Bandwidth bytes per second*
- *Packets processed per second*
- *Alerts generated per second*
- *TCP SYN per second*
- *TCP ACK per second*
- *New TCP sessions per second*
- *Closed TCP sessions per second*
- *Number of TCP sessions per interval*
- *CPU loading*
- *Packets lost per interval*

We found these statistics to be very accurate, and very useful for troubleshooting. It would be nice if all vendors would expose such statistics via the management GUI in this way.

Options are also available to start, stop and restart the engine, shutdown and restart the sensor, set external logging and alerting options (syslog and SNMP are supported currently), set alert options for the protocol parsers, packet decoders, TCP streaming engine, and mid-flow detection, and configure the options for SYN flood, UDP flood and ICMP flood defences.

Finally, the alert database can be backed up, restored or purged, and software updates and signature updates can be applied from local files downloaded from the Westline Web site. Signature updates can also be applied direct from the Westline Web site or from local IMC servers, and updates can be scheduled for regular, automatic application.

Policy Management

Policies are defined at site level, which means they can apply to a single sensor, or to multiple sensors controlled by single IMC.

Policies are made up of three elements:

- **Network Objects** - Entries describing individual hosts, networks (groups of hosts) and groups of networks. This allows the administrator to assign descriptive names to each of the above objects (i.e. DMZ instead of 10.10.16.0/24)
- **Service Objects** - Entries describing Ports (i.e. HTTP=TCP/80) and Firewall Services (i.e. Finger=TCP/79 and UDP/79). Although similar, Ports are used in IPS rules, whilst Firewall Services are used in firewall rules.
- **Attack Objects** - Entries describing attack signatures and groups of signatures.

When an Object window is opened in the Management Area within JConsole, it is divided into two parts: the Objects currently loaded in the IPS Engine, and the Objects currently in the configuration database (i.e. those being worked on currently in JConsole). All Objects in the Engine are read-only. As information is changed within JConsole it is changed in the configuration database only, and those Objects must then be explicitly applied to the Engine.

Attack Objects, also known as *Signature Objects* are obviously the heart of the detection and prevention system, and Aegis IPS has over 2300 signatures built-in. All signatures are assigned to an appropriate *Signature Group* to make them easier to handle. The administrator can create custom Signatures and Signature Groups as required, and existing signature groups can be duplicated. It is not possible to amend the built-in Signatures, however.

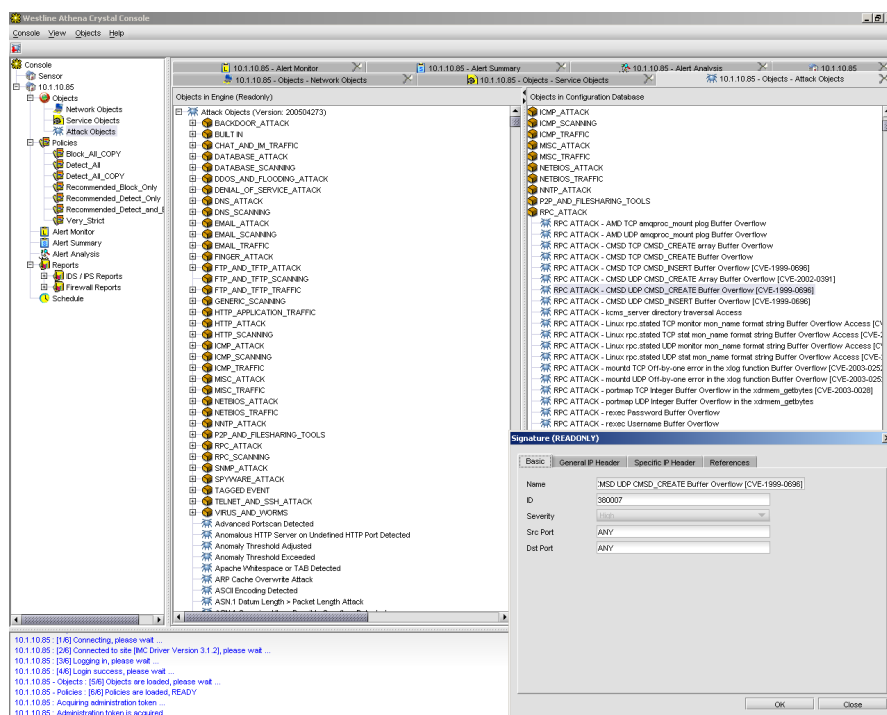


Figure 4 - Aegis IPS: Viewing/Editing Attack Objects

When creating or viewing Signatures, the Signature detail is divided into five tabs:

- **Basic** - Signature name, severity, source and destination port
- **General IP Header** - Source and destination IP, time to live, type of service, option flags, and so on
- **Specific IP Header** - Protocol-specific options, such as IP flags, ICMP type, and so on
- **Attack Pattern** - Data size, offset, depth, pattern to match, and so on. This information is viewable on custom signatures only, and unfortunately is not visible for built-in signatures. Regular expressions are not supported.
- **Reference** - CVE/Bugtraq IDs (built-in Signatures only)

Whilst it is clearly not possible to create complex signatures, this would normally be enough for most administrators.

A detection policy is a collection of IPS, Firewall and NAT rules, with associated notification and action options. There are eight built-in policies provided by the system which cannot be changed or deleted:

- **DETECT_ALL** - Alert only policy that detects all attacks, scanning and normal activities.
- **RECOMMENDED_DETECT_AND_BLOCK** - A policy that combines both alert and block actions, suitable for inspecting suspicious activities and blocking attacks. Chat programs and P2P are blocked.
- **RECOMMENDED_DETECT_ONLY** - Another version of the policy above but where all block actions are changed to alert only. Suitable for initial deployment and monitoring.
- **RECOMMENDED_BLOCK_ONLY** - A policy that focuses on blocking attacks, suitable for perimeter defence. Chat programs and P2P are blocked.
- **VERY_STRICT** - A policy that blocks all exploits and potential scanning activities. To be used with caution.

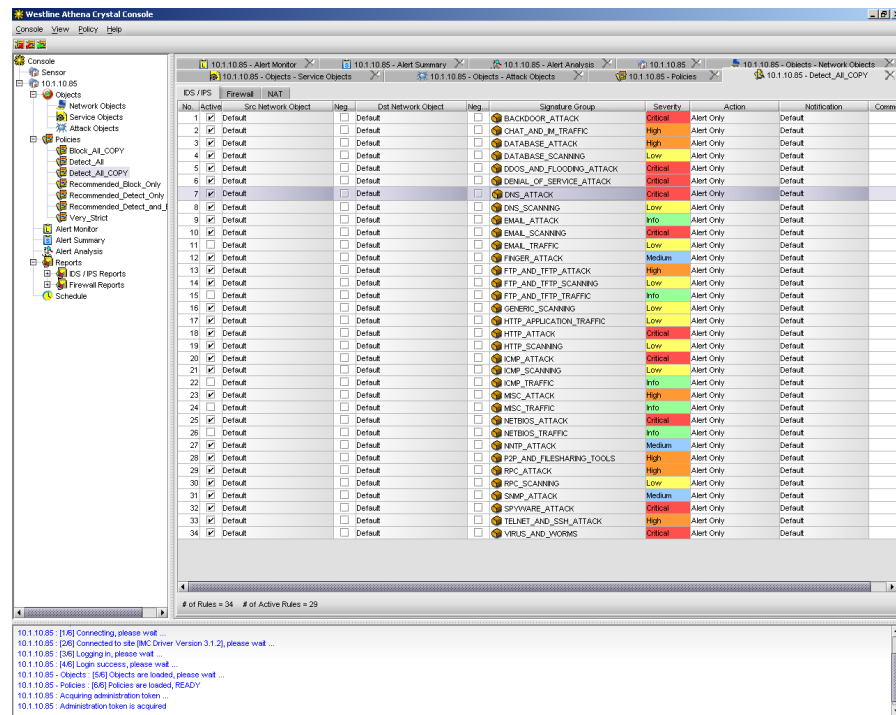


Figure 5 - Aegis IPS: Policy editor

New/custom policies can be created from scratch or based on existing policies and modified. Each policy contains multiple *Rules*, grouped by severity, each Rule containing the following components:

- **Rule number within the policy** - displayed on the alert screen when the rule is triggered for ease of identification
- **Status** - Active or disabled
- **Source Group** - Selected from the Network Objects. Can be a single host or an entire network, and it can be negated
- **Destination Group** - Selected from the Network Objects. Can be a single host or an entire network, and it can be negated
- **Signature/Signature Group** - Multiple Signatures and/or Signature Groups can be selected for each rule

- **Severity** - The default severity value as defined in each attack object can be used when raising alerts, or it can be overridden on per-Rule basis. The system allows ten custom levels besides the standard "Critical", "High", "Medium", "Low", and "Info".
- **Action** - The Action column defines the response of the sensor when a Rule is matched. The following actions are available:
 - **Default action** - drop packets in IPS (alert only in IDS), raise alert to JConsole, log packet which triggered the Rule
 - **Alert Only** - raise alert to JConsole
 - **Log TCP session traffic** - log remainder of packets for same TCP session
 - **Drop packet by IPS** - trigger packet is dropped along with the rest of the connection
 - **Close client connection** - send TCP reset to source IP
 - **Close server connection** - send TCP reset to destination IP
 - **Close connection on both sides** - send TCP reset in both directions
 - **Block by external firewall** - configure rules on third party firewalls
 - **Log packets** - Record all packets from source IP or to destination IP for specified time period after alert is raised
 - **Silent drop** - Packet is dropped (and connection is marked bad) without raising an alert
- **Notification** - The default behaviour of Aegis IPS is to log the attacking packets and send out an alert to JConsole. The following notifications are also supported:
 - No action
 - Syslog
 - SNMP Trap
- **Comment** - description of rule

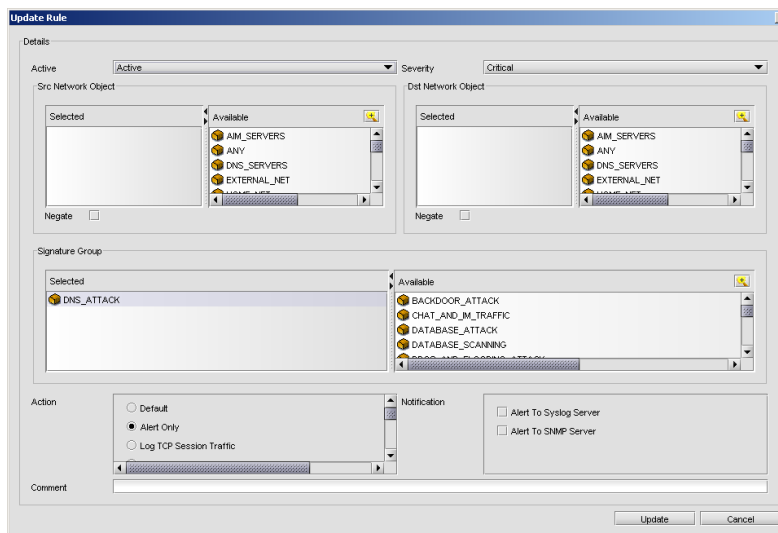


Figure 6 - Aegis IPS: Editing Rules

Firewall and NAT rules are defined in a similar way. The Firewall rules allow the administrator to block or allow packets at a lower level than the IPS Engine, perhaps simply blocking all packets from a specific source address without even inspecting them. Naturally, NAT rules come into play only when the device is deployed in routed/gateway mode.

Note that any number of Signature Groups or individual Signatures can be added to each rule, and since each rule can be restricted to specific source or destination IPS addresses/networks this allows the administrator to define very flexible and powerful policies. For example, it would be possible to drop all Web-based exploits when directed at the internal network, but alert only when directed at the DMZ.

A useful search capability is provided within the policy editor to allow the administrator to search for, say, all Signatures with "Apache" in the description when defining a Rule. Unfortunately, there is no bulk edit capability, so it is not possible to change a number of Rules - perhaps from *Active* to *Disabled* - in a single operation. This can make policy tuning quite tedious.

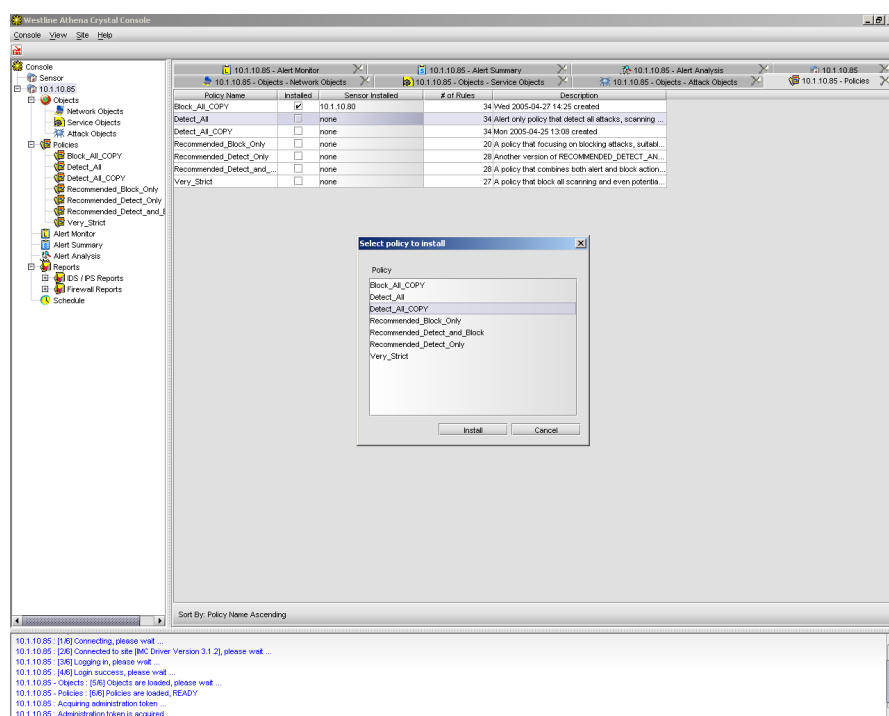


Figure 7 - Aegis IPS: Installing Policies

Once the policy has been defined, it can be deployed from JConsole to a sensor. The top level policy screen provides a list of available policies, together with a description, number of active rules, and on which sensor(s) that policy is installed. What is missing, of course, is an indication of when the policy was last updated, if it has been updated since it was installed, who updated it, and what changes were made. There is also no facility for rolling back to previous versions of a policy.

It would be nice to be able to select a policy from this list, and then choose one or more sensors to which it should be applied. Unfortunately, this is not possible. Instead, it is necessary to select a Site (device), and click on the *Install Policy* option, installing one policy to one device at a time.

This could be tedious in an environment with a large number of sensors and a small number of policies. Changing one policy could result in quite a lot of work for the administrator in deciding first where that policy is to be deployed, and then actually deploying it. This needs improvement.

Alert Handling

There are three ways to view alerts in JConsole: *Alert Monitor*, *Alert Summary* and *Alert Analysis*.

The *Alert Monitor* screen is a near real-time display of alerts as they are raised by the sensor, whilst the *Alert Summary* screen provides counts of the top 10 attack signatures, top 10 source IPs and top 10 destination IPs. The *Alert Summary* screen thus provides a rapid, high-level view of which attackers are sending which exploits to which victims, whilst the *Alert Monitor* provides the detail beneath that.

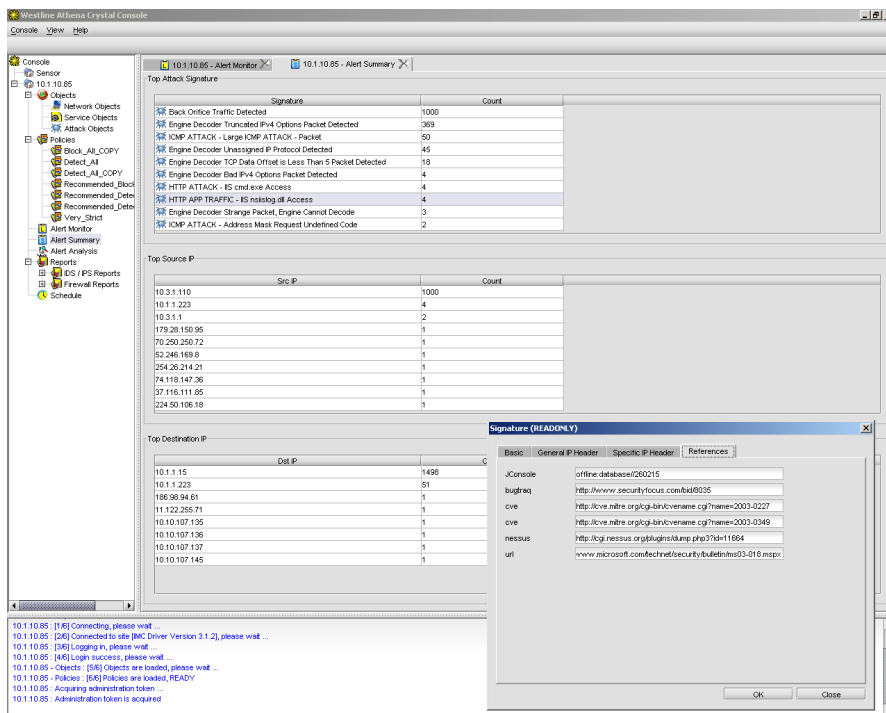


Figure 8 - Aegis IPS: Alert Summary

It is possible to drill down one level in the *Alert Summary* screen. By double-clicking on one of the attack signature entries, a separate window will appear listing the source and destination IPs involved in those particular attacks. It is a pity that it is not then possible to drill down further to see the actual alerts.

The *Alert Monitor* is divided into two windows, with a tabular list of the alerts sorted by date and time at the top, and a number of tabs containing details of the current alert selected at the bottom. The following information is presented in the tabular display:

- *Timestamp*
- *Severity*
- *Sensor ID*
- *Interface*
- *Interface in/out*
- *Signature name*
- *Source IP address*
- *Source MAC address*
- *Source port*

- Destination IP address
- Destination MAC address
- Destination port
- Protocol
- Payload (indicates if payload data was recorded)
- Event type (normal alert or packet recording)
- Rule ID within policy
- Action taken (alert, drop, etc.)
- Notification (alert, syslog, SNMP)

If alerts are arriving too quickly to view, the real-time update can be suspended at any time. It is a pity that once suspended it is not possible to re-sort on any column, add/remove columns, or right click on an entry and drill down on that data (i.e. right click on a particular source IP in order to display all alert from that IP). The *Alert Analysis* screen must be used for all drill-down operations.

The auto-update interval can be adjusted to suit the administrator, as can the maximum number of alerts displayed. Note that this display should be considered as a “window” into the MySQL database which always shows the most recent alerts - once the maximum number of alerts is exceeded the earlier ones are discarded to make way for the new ones. However, the alert details remain in the database for further reporting and analysis. This is a reasonable compromise between availability and performance.

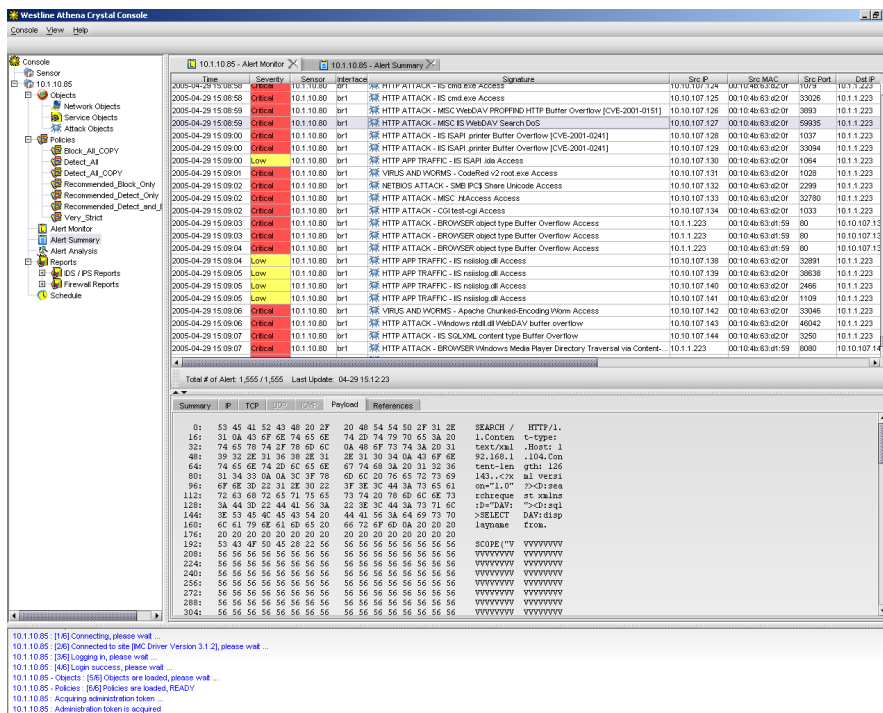


Figure 9 - Aegis IPS: Alert Monitor

Selecting an alert within the table in the top window displays details of that alert in the bottom window.

A total of seven tabs are available in this window, some of which may be greyed out depending on the protocol used by the exploit and whether or not packet data was recorded:

- **Summary** - Replicates all of the data items listed previously from the table of alerts
- **IP** - IP-specific data such as version, header length, flags, TTL, checksum, and so on
- **TCP** - TCP-specific data such as source port, destination port, sequence number, flags, and so on
- **UDP** - UDP-specific data such as source port, destination port, etc.
- **ICMP** - ICMP-specific data such as ICMP type, ICMP code
- **Payload** - Full protocol decode of the packet which triggered the alert, containing both hex and ASCII data
- **References** - List of external references such as Bugtraq ID, CVE reference, and so on.

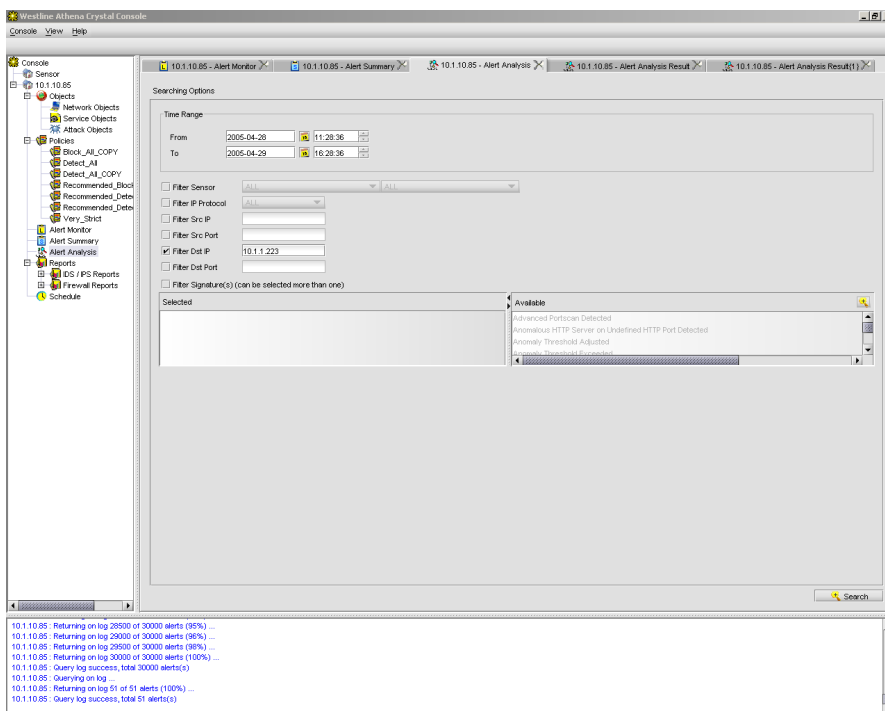


Figure 10 - Aegis IPS: Alert Analysis

More detailed historical analysis is performed via the *Alert Analysis* option. This is a cross between the *Alert Monitor* and *Reporting* functions, in that it replicates the tabular display of the *Alert Monitor*, but instead of listing current alerts in real-time, it provides access to the historical alerts stored in the MySQL database via a set of search criteria:

- **Time Range**
- **Sensor** - Single sensor, or consolidated across all sensors under the same IMC
- **Protocol**
- **Source IP**
- **Source Port**
- **Destination IP**
- **Destination Port**
- **Signature(s)** - One or more signatures can be selected

Unfortunately, it is not possible to save the search criteria for re-use.

The display is identical to that provided by the Alert Monitor, except with this option the administrator can sort on any column. Double-clicking on a Signature brings up a window containing details of that Signature, but not the pattern matched for built-in Signatures, which can make it hard to determine why a false positive alert has triggered.

It is not possible to go directly to the appropriate signature in the appropriate policy in order to edit the policy directly for tuning purposes, and nor is it possible to drill down on a particular field in order to, say, list all alerts for a particular target IP. Instead, it would be necessary to return to the selection criteria, specify the required target IP, and regenerate the query. This is cumbersome compared to some of the current competition.

Reporting and Analysis

A number of standard reports are included in the IMC covering both IDS/IPS and Firewall:

- **IDS/IPS Reports**
 - *Top Scan Sources (All/TCP/UDP/ICMP)*
 - *Top Scan Targets (All/TCP/UDP/ICMP)*
 - *Top Attacks by Severity (All/TCP/UDP/ICMP)*
 - *Top Attacks by Signature (All/TCP/UDP/ICMP)*
 - *Attacks Over Time (All/TCP/UDP/ICMP)*
 - *CPU loading*
 - *Custom Report*
- **Firewall Reports**
 - *Top Connection by Sources (All/TCP/UDP/ICMP)*
 - *Top Connection by Targets (All/TCP/UDP/ICMP)*
 - *Custom Report*

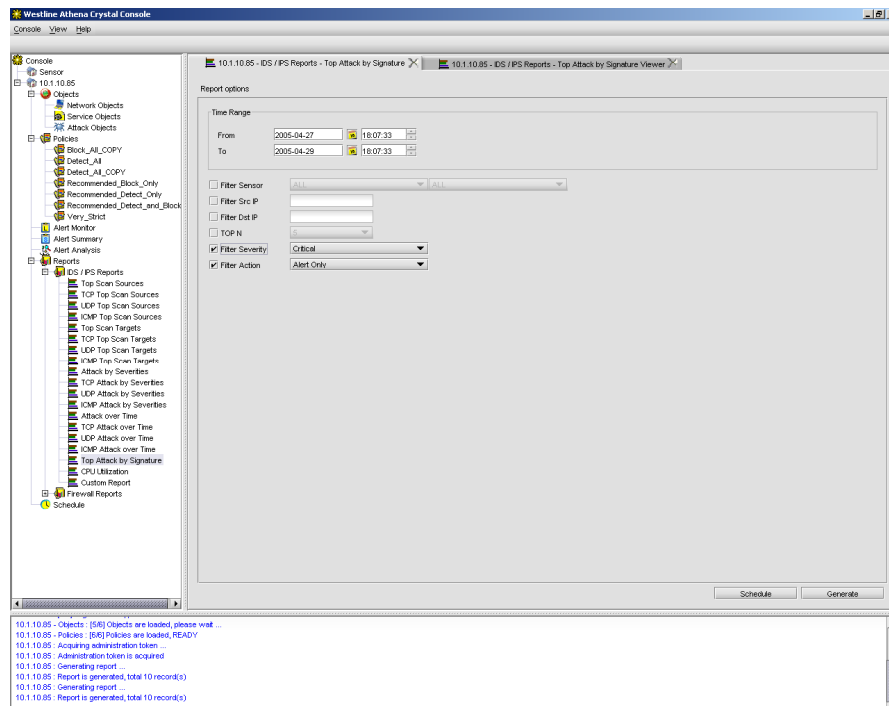


Figure 11 - Aegis IPS: Specifying report criteria

Selecting a report brings up the report option windows in the Management Area. Different reports have different selection criteria on which the results can be filtered:

- **Time Range** - The time period covered by the report. The default is one month
- **Sensor** - Reports can be for a single sensor, or consolidated across all sensors under the same IMC (the default)
- **Interface** - The default is to include all interfaces under the same sensor
- **Source IP**
- **Destination IP**
- **Source Port**
- **Destination Port**
- **Top N** - The default is top 20
- **Severity**
- **Action**
- **Signature(s)** - One or more signatures can be selected (maximum is 500)
- **Interval** - The default is one month

The *Custom Report* option provides all the same selection criteria, but allows the data columns to be specified in addition. Once the selection criteria have been established, the report can be generated immediately or scheduled for hourly, daily or weekly automated runs (though it is not possible to specify a time for the report to be run - an unfortunate omission). The finished report is presented on-screen as both a graph and text-based table.

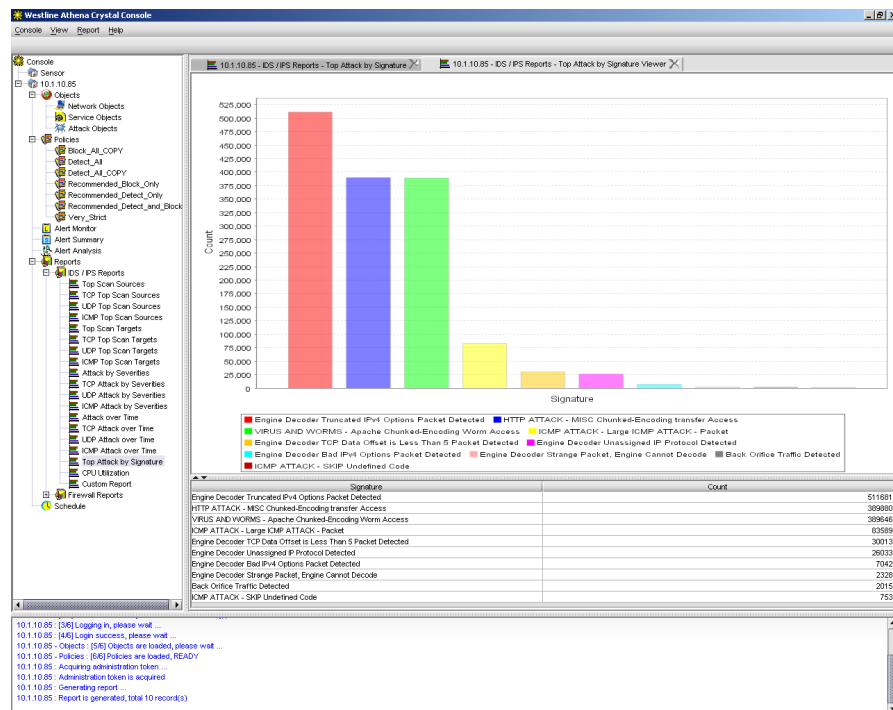


Figure 12 - Aegis IPS: Typical report

As with other options selected from the Function Area menu, each report is opened in its own tabbed window in the Management Area.

This allows the administrator to generate several reports at once and switch quickly between them. He can also customise the report format (pie chart or bar chart), colours, fonts, plot styles, and so on, before exporting or printing the report.

The report can be printed directly, saved as a PDF file, or exported to Excel, HTML, CSV or plain text formats. Scheduled reports are always saved as PDF files, and they must be manually retrieved from the IMC hard drive - it would be nice if they could be automatically e-mailed or uploaded to a Web or FTP server.

We found report generation to be slow when there were a million or more alerts in the database, and, unfortunately, it is possible to save neither the selection criteria nor the custom formatting for re-use.

Verdict

Performance

The Athena Aegis IPS 510L was tested up to 200Mbps, the rated speed of the device, and performance at all levels of our load tests was impeccable, with 100 per cent of all attacks being detected and blocked under all load conditions. We would happily confirm Westline's 200Mbps rating for this device.

The basic latency figures of the Athena Aegis IPS 510L were acceptable across the board under all traffic loads. Behaviour throughout the tests with no background traffic was consistent and predictable, with small increases as additional network load was applied from 50Mbps to 200Mbps.

Placing the device under a half load of 100Mbps of HTTP traffic actually had the effect of reducing latency further. The effect of this is that latency under normal traffic loads is excellent for a device of this type - HTTP response times, for example, were excellent.

The device did have problems with SYN flood traffic, however, and would perform best when situated behind an attack mitigation device or a firewall with SYN flood mitigation capabilities. Latency when under 20Mbps of SYN flood traffic was excessive, as was packet loss. The overall effect was an increase in response times and a high number of failed transactions.

This is an issue which Westline recognises, and is working on at the time of writing. For now, our recommendation would be to deploy the device as mentioned above, with additional protection in front.

Athena Aegis 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 99.9975 per cent of legitimate traffic. This is considered acceptable for a device of this nature.

Exposing the sensor interface to ISIC-generated traffic had no long-term adverse effect on the device, although communications between the management console and the management server/sensor were interrupted during the attack.

Security Effectiveness

We installed one sensor with the latest updates, and used a slightly modified version of the *Detect_All* policy, disabling *Low* severity audit/informational signatures only. All attack/exploit signatures are enabled in this policy.

Out of the box, signature recognition was adequate at 75 per cent, and was improved to 88 per cent following the application of a signature update after 48 hours. Blocking performance was actually four per cent higher throughout - giving a much more creditable final result of 92 per cent - due to various types of malformed/invalid packets (such as those used in our DOS test suite) being blocked silently.

Performance in our “false negative” tests was adequate out of the box, and improved slightly following the signature update. However, there were still four misses out of the 14 test cases following the signature update, indicating that many signatures are written for specific exploits rather than for the underlying vulnerability.

Having said that, we did notice that where a deliberate attempt was not made to evade the device, Athena Aegis IPS did extremely well at detecting most of the variants that we include alongside our basic test cases (i.e. if five different exploits are known for a vulnerability, the signature is generally written well enough to detect all of them).

A major concern in deploying an IPS is the blocking of legitimate traffic. The device failed only one test case in our false positive test suite, and this was rectified following the signature update. No other false positives were noted during stress testing with either Avalanche traffic or real-world traffic mixes.

Resistance to known evasion techniques was very good, with the Athena Aegis IPS achieving a clean sweep across the board in most of our evasion tests. *Fragroute* and *Whisker* both failed to deceive the device into ignoring valid attacks, and all of the attempts were decoded accurately.

Out of the box, Westline claims that Athena Aegis IPS can handle just over 250,000 open connections, and this was verified in our tests.

Usability

Westline offers a range of deployment options, from two tier (managing each sensor directly from the Console), to three-tier (using an IMC to manage multiple sensors) and even “hybrid mode” (a mix of two-tier and three-tier). This makes for a very flexible and scalable system.

However, in the current release the implementation is flawed slightly, with no role-based access, not enough security during initial connection of sensors to management servers/consoles, and not enough protection to prevent the same sensor being updated with different policies via two different management servers/consoles. Once these issues have been rectified, the Westline management model will be very impressive.

Policy management is flexible and powerful thanks to the rules-based approach which allows, in effect, different policies to be applied on a per-host or per-network basis. Some improvement is still required in the area of bulk edits and provision of audit trails for policy changes, but the current solution is very usable as it stands.

Many organisations will also find useful the ability to deploy this device in routed gateway mode complete with stateful firewall and NAT.

Alert handling is adequate, and although we would like to see more drill-down capabilities from the Alert Monitor screen, all the necessary facilities are available, including real-time alert monitoring, high-level counts and summaries, and a more detailed log analysis capability.

It would be useful to be able to go directly to the appropriate signature within a policy from an alert to disable the signature or fine tune its configuration parameters.

Reporting is fairly basic, with several built-in graphical reports and a limited custom reporting capability. The biggest omission here is the inability to save customised report layouts and selection criteria for re-use.

Contact Details

Company name: Westline Security Limited

E-mail: info@west-line.net

Internet: www.west-line.net

Address:
2/F Shui On Centre,
8 Harbour Road,
Wanchai,
Hong Kong

Tel: +852 2824 8911

Fax: +852 2824 8365

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 than 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.*

Westline Athena Aegis IPS 510L V2.1 Test Results

Section 1 - Detection Engine

Test 1.1 – Attack Recognition	Attacks	Default ARR	Default ARRB	Custom ARR	Custom ARRB
Test 1.1.1 - Backdoors	7	6	6	6	6
Test 1.1.2 - WINS/DNS	3	3	3	3	3
Test 1.1.3 - DOS	10	5 ¹	9 ¹	6 ¹	10 ¹
Test 1.1.4 - False negatives (modified exploits)	14	9	9	10	10
Test 1.1.5 - Finger	4	4	4	4	4
Test 1.1.6 - FTP	5	4	4	4	4
Test 1.1.7 - HTTP	43	30	30	41	41
Test 1.1.8 - ICMP	2	2	2	2	2
Test 1.1.9 - Reconnaissance	8	8	8	8	8
Test 1.1.10 - RPC	9	8	8	8	8
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	0	0	1	1
Total	110	83 / 110¹	87 / 110	97 / 110¹	101 / 110
		75%¹	79%	88%¹	92%

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	PASS	PASS
Test 1.2.14 - Fake DNS /bin/sh buffer overflow	PASS	PASS
Test 1.2.15 - Inter-firewall communication traffic	PASS	PASS
Test 1.2.16 - Fake SQL Slammer traffic	FAIL	PASS
Test 1.2.17 - File copy of GIF file (contains bytes which look like NOP sled)	PASS	PASS
Total Passed	16 / 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	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	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	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 ²
Test 2.2.13 - TCP segmentation - ordered 2 byte segments, segment overlap (favour new)	YES	YES	YES ²
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
Test 2.2.16 - TCP segmentation - out of order 1 byte segments, interleaved duplicate segments with faked retransmits	YES	YES	YES ²
Test 2.2.17 - TCP segmentation - ordered 1 byte segments, segment overlap (favour new)	YES	YES	YES ²
Test 2.2.18 - TCP segmentation - out of order 1 byte segments, PAWS elimination (interleaved dup segments with older TCP timestamp options)	YES	YES	YES
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 ²
Total	20 / 20	20 / 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	NO	NO	NO
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)	NO	NO	NO
Test 2.4.7 - HTTP exploits to port <> 80	NO	NO	NO
Total	4 / 7	4 / 7	4 / 7

Section 3 - Stateful Operation

Test 3.1 – Stateless Attack Replay	Alert?	Blocked?	Pass/Fail
Test 3.1.1 - Stateless Web exploits	NO ³	YES	PASS
Test 3.1.2 - Stateless FTP exploits	NO ³	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	N/A ⁴	N/A ⁴
Test 3.2.2 - Attack Blocking	PASS	PASS	PASS	PASS	PASS	N/A ⁴	N/A ⁴
Test 3.2.3 - State Preservation	PASS	PASS	PASS	PASS	PASS	N/A ⁴	N/A ⁴
Test 3.2.4 - Legitimate traffic blocking	PASS	PASS	PASS	PASS	PASS	N/A ⁴	N/A ⁴

Test 3.3 – Simultaneous Open Connections (after tuning) ⁵							
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	N/A ⁴	N/A ⁴
Test 3.3.2 - Attack Blocking	PASS	PASS	PASS	PASS	PASS	N/A ⁴	N/A ⁴
Test 3.3.3 - State Preservation	PASS	PASS	PASS	PASS	PASS	N/A ⁴	N/A ⁴
Test 3.3.4 - Legitimate traffic blocking	PASS	PASS	PASS	PASS	PASS	N/A ⁴	N/A ⁴

Section 4 - Detection/Blocking Performance Under Load

Test 4.1 – UDP traffic to random valid ports		50Mbps	100Mbps	150Mbps	200Mbps	Max
Test 4.1.1 - 256 byte packet test - max 91,000pps	Detected	100%	100%	100%	100%	200Mbps ⁵
	Blocked	100%	100%	100%	100%	
Test 4.1.2 - 550 byte packet test - max 44,000pps	Detected	100%	100%	100%	100%	200Mbps ⁵
	Blocked	100%	100%	100%	100%	
Test 4.1.3 - 1514 byte packet test - max 24,000pps	Detected	100%	100%	100%	100%	200Mbps ⁵
	Blocked	100%	100%	100%	100%	

Test 4.2 – HTTP “maximum stress” traffic with no transaction delays		50Mbps	100Mbps	150Mbps	200Mbps	Max
Test 4.2.1 - Max 500 connections per second - ave packet size 1000 bytes - max 24,000 packets per second	Detected	100%	100%	100%	100%	200Mbps
	Blocked	100%	100%	100%	100%	
Test 4.2.2 - Max 1000 connections per second - ave packet size 540 bytes - max 45,000 packets per second	Detected	100%	100%	100%	100%	200Mbps
	Blocked	100%	100%	100%	100%	
Test 4.2.3 - Max 2000 connections per second - ave packet size 440 bytes - max 55,000 packets per second	Detected	100%	100%	100%	100%	200Mbps
	Blocked	100%	100%	100%	100%	
Test 4.2.4 - Max 4000 connections per second - ave packet size 360 bytes - max 64,000 packets per second	Detected	100%	100%	100%	100%	200Mbps
	Blocked	100%	100%	100%	100%	

Test 4.3 – HTTP “maximum stress” traffic with transaction delays		50Mbps	100Mbps	150Mbps	200Mbps	Max
Test 4.3.1 - Max 1000 connections per second - ave packet size 540 bytes - max 45,000 packets per second - 10 sec delay - max 10,000 open connections	Detected	100%	100%	100%	100%	200Mbps
	Blocked	100%	100%	100%	100%	
Test 4.3.2 - Max 2000 connections per second - ave packet size 440 bytes - max 55,000 packets per second - 10 sec delay - max 20,000 open connections	Detected	100%	100%	100%	100%	200Mbps
	Blocked	100%	100%	100%	100%	

Test 4.4 – Protocol mix		50Mbps	100Mbps	150Mbps	200Mbps	Max
Test 4.4.1 - 72% HTTP (540 byte packets) + 20% FTP + 6% UDP (256 byte packets). Max 800 connections per second - ave packet size 540 bytes - max 43,000 packets per second - max 150 open connections	Detected	100%	100%	100%	100%	200Mbps
	Blocked	100%	100%	100%	100%	

Test 4.5 – Real World traffic		50Mbps	100Mbps	150Mbps	200Mbps	Max
Test 4.5.1 - Pure HTTP (simulated browsing session on NSS Web site). Max 950 connections per second - 6 new users per second - ave packet size 560 bytes - max 42,000 packets per second	Detected	100%	100%	100%	100%	200Mbps
	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 750 connections per second - ave packet size 560 bytes - max 41,000 packets per second - max 300 open connections	Detected	100%	100%	100%	100%	200Mbps
	Blocked	100%	100%	100%	100%	

Section 5 - Latency & User Response Times

Test 5.1 – Latency	Packet Size	50Mbps	100Mbps	150Mbps	200Mbps
Test 5.1.1 Average latency (µs) with no background traffic	256	315.45	337.83	344.80	360.32
	550	343.25	356.37	387.09	411.30
	1000	403.25	414.52	424.42	426.86
Test 5.1.2 Average latency (µs) with background traffic (100Mbps HTTP traffic, max 500 connections per second - ave packet size 540 bytes - max 22,500 packets per second)	256	156.36			
	550	187.20			
	1000	244.13			
Test 5.1.3 Average latency (µs) when under attack (20Mbps SYN flood (29600pps))	256	636.20 ⁶			
	550	652.26 ⁶			
	1000	676.66 ⁶			

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 (100Mbps HTTP traffic, max 500 connections per sec - ave packet size 540 bytes - max 22,500 packets per sec)	343103	0	201	205	202
Test 5.2.2 - Web page response (ms) when under attack (100Mbps HTTP traffic, max 500 connections per sec - ave packet size 540 bytes - max 22,500 packets per sec PLUS 20Mbps SYN flood (29600pps))	343080	106170 ⁶	202	31636 ⁶	3058 ⁶

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	99.9975%
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:

- Malformed packets are dropped silently - although not alerted upon, they are always blocked successfully
- Always blocked
- Mid-stream alert is raised - this is configurable
- Not possible to test beyond 250k open connections - connections per second too high. No configuration/tuning possible for this test
- In pure UDP tests we noted very slight packet loss - usually only a single packet over an extended run at maximum load. This is due to a bug which is being investigated at the time of writing. Throughout the tests, this had no practical effect on the device's ability to detect and block attacks with real-world traffic
- High degree of packet loss during SYN flood attack. Attack was not successfully mitigated

Section 1: Detection Engine

We installed one sensor with the latest updates, and used a slightly modified version of the *Detect_All* policy, disabling *Low* severity audit/informational signatures only. All attack/exploit signatures are enabled in this policy. No other configuration or tuning was necessary.

Out of the box, signature recognition was adequate at 75 per cent, and was improved to 88 per cent following the application of a signature update after 48 hours. Blocking performance was actually four per cent higher throughout - giving a much more creditable final result of 92 per cent - due to various types of malformed/invalid packets (such as those used in our DOS test suite) being blocked silently.

Performance in our “false negative” tests was adequate out of the box, and improved slightly following the signature update. However, there were still four misses out of the 14 test cases following the signature update, indicating that many signatures are written for specific exploits rather than for the underlying vulnerability.

Having said that, we did notice that where a deliberate attempt was not made to evade the device, Athena Aegis IPS did extremely well at detecting most of the variants that we include alongside our basic test cases (i.e. if five different exploits are known for a vulnerability, the signature is generally written well enough to detect all of them).

A major concern in deploying an IPS is the blocking of legitimate traffic. The device failed only one test case in our false positive test suite, and this was rectified following the signature update. No other false positives were noted during stress testing with either Avalanche traffic or real-world traffic mixes.

Section 2: IPS Evasion

Resistance to known evasion techniques was very good, with the Athena Aegis IPS achieving a clean sweep across the board in most of our evasion tests. *Fragroute* and *Whisker* both failed to deceive the device into ignoring valid attacks, and all of the attempts were decoded accurately.

Of the miscellaneous evasion techniques, non-text Telnet opcodes in FTP data streams and RPC fragmentation both proved troublesome (Athena Aegis IPS could handle both MS-RPC and Sun RPC fragmentation on certain signatures only, indicating that the RPC protocol decoder is incomplete). Full RPC and FTP protocol parsers are under development for a future release.

Note also that HTTP exploits can be detected on ports 80, 8080, 8180 and 8888 only.

Section 3: Stateful Operation

Out of the box, Westline claims that Athena Aegis IPS can handle just over 250,000 open connections, and this was verified in our tests. No tuning is possible to increase this, and we were unable to test much beyond this level since the high connection rates used in our 500,000 and 1 million open connections tests exceed the limit of 4,000 connections per second on this device.

The default operation of the device is to reject new connections when the state tables are full or resources are low, and this meant that once we exceeded the connection limit the device continued to maintain state on existing connections (thus detecting our half-open exploit), but inevitably, as a consequence, began to block legitimate traffic. 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 (a mid-flow violation alert is raised, if configured). It is not possible to configure the device to allow mid-flows.

Section 4: Detection/Blocking Performance Under Load

Note that the Athena Aegis IPS 510L was tested as a 200Mbps IPS device.

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

We would happily confirm Westline's 200Mbps rating for this device.

One anomaly we did note was the loss of a single UDP packet in each of our UDP performance tests. This is due to a bug which is currently being investigated, and throughout the tests it had no practical effect on the device's ability to detect and block attacks with real-world traffic.

Section 5: Latency & User Response Times

The basic latency figures of the Athena Aegis IPS 510L were acceptable across the board under all traffic loads. They ranged from 315µs with 50Mbps of 256 byte packets, to 416µs with 200Mbps of 1000 byte packets.

Behaviour throughout the tests with no background traffic was consistent and predictable, with small increases as additional network load was applied from 50Mbps to 200Mbps. Placing the device under a half load of 100Mbps of HTTP traffic actually had the effect of reducing latency further, falling from 315µs to 156µs with 256 byte packets, 343µs to 187µs with 550 byte packets, and from 403µs to 244µs with 1000 byte packets.

This is strange behaviour, and along with the single packet loss anomaly mentioned in the previous section indicates possible inefficiencies in the handling of pure UDP traffic. The effect of this, of course, is that latency under normal traffic loads is excellent for a device of this type - HTTP response times, for example, were also excellent.

The device did have problems with SYN flood traffic, however, and would perform best when situated behind an attack mitigation device or a firewall with SYN flood mitigation capabilities. Latency when under 20Mbps of SYN flood traffic was excessive, as was packet loss. The overall effect was an increase in response times and a high number of failed transactions.

This is an issue which Westline recognises and is working on at the time of writing. Future releases will include per-IP based rate limiting capabilities to help mitigate such attacks. For now, our recommendation would be to deploy the device as mentioned above, with additional protection in front.

Section 6: Stability & Reliability

Athena Aegis 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 99.9975 per cent of legitimate traffic (a total of 4 out of 200,000 legitimate sessions failed in each of the extended runs - well within acceptable limits for a device of this type).

Exposing the sensor interface to ISIC-generated traffic had no long-term adverse effect on the device, although communications between the management console and the management server/sensor were interrupted during the attack. In addition, at the height of the attack the effect was similar to that noted when subjected to a SYN flood, causing some legitimate traffic to be blocked.

A large number of ISIC-related alerts were raised during the attack and the attack was mitigated partially (many malformed ISIC packets are also dropped silently). All other malicious traffic continued to be blocked successfully during the attack, and there were no residual stability problems once the ISIC attack had been terminated.

Section 7: Management Interface

Only three ports are open by default: TCP/22348 (console to built-in management server), TCP/22349 (external management server to sensor), and TCP/9132 (out-of-band software and signature upgrade via console). Access via the management port is restricted to specified IP addresses only, making it very difficult to attack.

The extended ISIC attack against the management interface caused communications to be interrupted between the console and the management server/sensor, but these were restored successfully once the attack was terminated and there were no residual stability problems.