V-Secure V100 V7.0

Technical Evaluation

An NSS Group Report

First published January 2005 (Version 1.0)

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

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

TABLE OF CONTENTS

TABLE OF FIGURES

The NSS Group

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

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

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

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

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

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

Foreword

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

As with Edition 1, 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.

It is interesting to note that between publishing Edition 1 and Edition 2 the analyst groups who were previously so sure that IDS was dead and IPS stillborn have now come around to our way of thinking - while the so-called "*deep inspection firewalls*" are not ready for prime-time deployments, security administrators need to make the best use of the technology that **is** available, and for now that means a combination of firewalls, in-line intrusion prevention devices, and intrusion detection systems. They are likely to be in use for quite some time to come, too!

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. This edition also sees the introduction of a new *Rate-Based IPS* methodology to complement our exiting *Content-Based IPS* methodology used in Edition 1. This has allowed us to more accurately test Rate-Based/Attack Mitigation products, and two devices were tested against this new methodology in the latest report (one of them actually tested against **both** methodologies – a first).

It is worth pointing out that not every product submitted for testing receives an *NSS Approved* award. Standards are very high, and out of nine products signed up for this group test initially, only the five included in the final Edition 2 report have received ***NSS Approved*** awards.

We believe that our IPS test methodologies - which have been updated 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 **IPS Group Test (Edition 2)** 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 the 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.

Implementation Challenges

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

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

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

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

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

Dropped packets are also an issue, since if even one of those dropped packets is one of those used in the exploit data stream it is possible that the entire exploit could be missed. Most high-end IPS vendors will get around this problem by using custom hardware, populated with advanced FPGAs and ASICs - indeed, it is necessary to design the product to operate as much as a switch as an intrusion detection and prevention device.

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

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

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

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

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

Requirements for effective prevention

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

- **In-line operation -** only by operating in-line can an IPS device perform true protection, discarding all suspect packets immediately and blocking the remainder of that flow
- **Reliability and availability -** should an in-line device fail, it has the potential to close a vital network path and thus, once again, cause a DoS condition. An extremely low failure rate is thus very important in order to maximise up-time, and if the worst should happen, the device should provide the option to fail open or support fail-over to another sensor operating in a fail-over group (see below). In addition, to reduce downtime for signature and protocol coverage updates, an IPS must support the ability to receive these updates without requiring a device re-boot. When operating inline, sensors rebooting across the enterprise effectively translate into network downtime for the duration of the reboot
- **Resilience -** as mentioned above, the very minimum that an IPS device should offer in the way of High Availability is to fail open in the case of system failure or power loss (some environments may prefer this default condition to be “fail closed” as with a typical firewall, however - the most flexible products will allow this to be user-configurable). Active-Active stateful fail-over with cooperating in-line sensors in a fail-over group will ensure that the IPS device does not become a single point of failure in a critical network deployment
- **Low latency -** when a device is placed in-line, it is essential that its impact on overall network performance is minimal. Packets should be processed quickly enough such that the overall latency of the device is as close as possible to that offered by a layer 2/3 device such as a switch, and no more than a typical layer 4 device such as a firewall or load-balancer.
- **High performance -** packet processing rates must be at the rated speed of the device under real-life traffic conditions, and the device must meet the stated performance with all signatures enabled.

Headroom should be built into the performance capabilities to enable the device to handle any increases in size of signature packs that may occur over the next three years. Ideally, the detection engine should be designed in such a way that the number “signatures” (or “checks”) loaded does not affect the overall performance of the device.

- **Unquestionable detection accuracy -** it is imperative that the quality of the signatures is beyond question, since false positives can lead to a Denial of Service condition. The user MUST be able to trust that the IDS is blocking only the user selected malicious traffic. New signatures should be made available on a regular basis, and applying them should be quick (applied to all sensors in one operation via a central console) and seamless (no sensor reboot required)
- **Fine-grained granularity and control -** fine grained granularity is required in terms of deciding exactly which malicious traffic is blocked. The ability to specify traffic to be blocked by attack, by policy, or right down to individual host level is vital. In addition, it may be necessary to only alert on suspicious traffic for further analysis and investigation
- **Advanced alert handling and forensic analysis capabilities -** once the alerts have been raised at the sensor and passed to a central console, someone has to examine them, correlate them where necessary, investigate them, and eventually decide on an action. The capabilities offered by the console in terms of alert viewing (real time and historic) and reporting are key in determining the effectiveness of the IPS product.

The NSS Intrusion Prevention Group Test

The NSS Group conducted the first comprehensive IPS test of its kind, now updated in this Edition. This exhaustive review will give readers a complete perspective of the capabilities, maturity and suitability of the products tested for their particular needs.

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

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

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

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

Performance

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

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

Detection/Blocking Performance Under Load

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

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

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

Latency & User Response Times

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

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

throughput = window / RTT

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

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

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

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

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

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

Such critical value calculations are important when TCP connections achieve maximum throughput, which is true for large data transfers. For smaller data transfers, and non-TCP applications like NFS, latency has a more direct impact on user experience - response time is directly proportional to latency. That is, *doubling latency doubles response time.* In these situations, the latency of the network in which a NIPS is deployed determines the acceptable latency of the NIPS.

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

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

- Internet facing connections experience round-trip latency from 10-100 milliseconds. Deploying our hypothetical NIPS would increase the round trip latency by 1-10 per cent, which would have only a minor impact on the user experience.

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

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

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

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

Spirent Avalanche and Reflector devices are also used to generate HTTP sessions through the device in order to gauge how any increases in latency will impact the user experience in terms of failed connections and increased Web response times. This "*application latency*" is measured both with no background load and while the device is under attack.

Stability & Reliability

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

In the first part of this test, we expose the external interface of the sensor to a constant stream of attacks over an extended period of time. The device is configured to block and alert, and thus this test provides an indication 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*.

V-SECURE V-100 V7.0

Executive Summary

The V-Secure IPS series is a family of attack detection and mitigation appliances designed to detect and prevent rate-based attacks across multiple network segments at speeds of over 100Mbps. A range of models are available covering 10Mbps, 100Mbps and fractional Gigabit networks.

The V-100 under test here is currently the mid-range 100Mbps model. A dedicated 1U appliance designed to monitor a single network segment at 100Mbps speeds, the device sports two copper 10/100Mbps ports for in-line operation, and a single 10/100Mbps port for management. Under normal network conditions, we can verify the 100Mbps rating of this device.

With a range of innovative technologies under the hood, we found the V-100's detection and mitigation capabilities to be excellent. We also found the V-100 to be very stable and reliable, coping with our extensive reliability tests with ease and without succumbing to most common evasion techniques.

The GUI is comprehensive and user-friendly, and contains a range of excellent monitoring and reporting capabilities. An "enterprise" version is also available for managing multiple devices.

Architecture

The V-Secure offering supports a three-tier architecture, consisting of the *V-Secure Management Studio* and the *V-Secure IPS* appliance.

V-Secure Management Studio

V-Secure Management Studio enables management of one or more V-Secure IPS devices (one at a time) from a central graphical console. The components of V-Secure Management Studio are:

- ***V-Secure Management Server -*** *Provides central management functions for V-Secure IPS, log file collection, and client authorisation for SuperVisor clients.*
- ***SuperVisor Client -*** *Enables management of the security features of V-Secure IPS and provides a user interface for individual V-Secure IPS configuration (no correlation or multi-device management capabilities).*
- ***V-Secure IPS Reporter -*** *Provides detailed reporting services based on V-Secure IPS log analysis.*

The V-Secure IPS can also be managed using the following additional products:

- ***NetVisor Management Console -*** *This is a Java-based client which connects to the V-Secure Management Server to provide centralised management of multiple V-Secure appliances across the network. Unlike the SuperVisor client, NetVisor is an enterprise-level product designed to provide comprehensive management, alerting, monitoring and reporting across multiple devices simultaneously.*

It provides a high-level topological graphical display of the network, with alerts correlated across multiple devices and physical locations and highlighted on a graphical map. The administrator can drill down into countries, sites, and individual devices to determine the cause of the alerts. The reporting capability is also far more advanced than the HTML-based V-Secure IPS Reporter module, providing much more comprehensive filtering and reporting capabilities across multiple devices rather than having to select one device at a time.

- ***CLI (command line interface) -** is used for secure direct set-up and appliance management. This interface communicates directly with the appliance's embedded configuration mechanisms.*

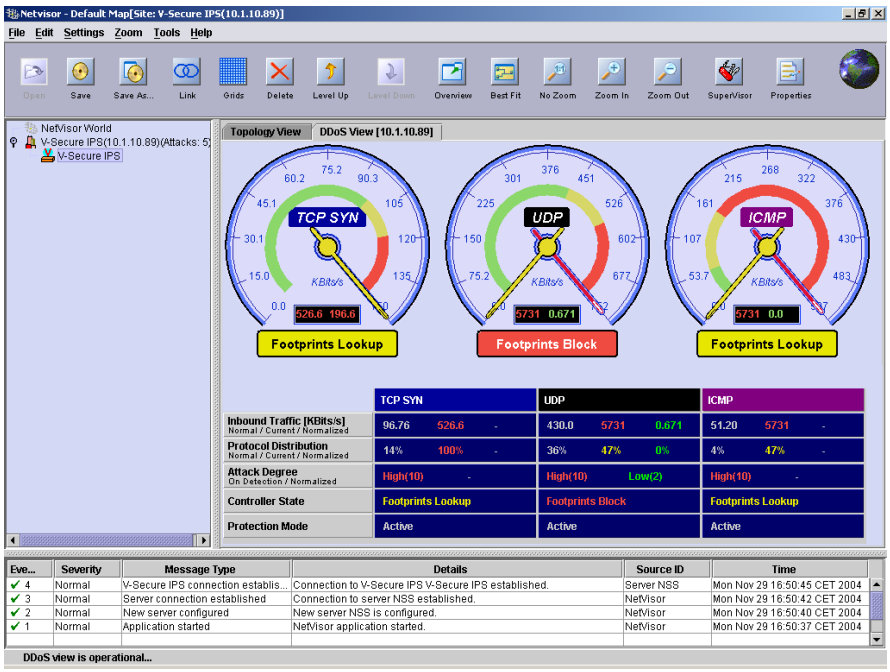

Figure 1 - V-Secure: NetVisor Dashboard

V-Secure IPS Appliance

Three appliances are available, all 1U rack-mount devices running on a 750MHz IBM PowerPC platform:

- ***V-Secure V-10 -** This is a 10Mbps device, with 128MB RAM, two copper 10Mbps detection ports and a single 10/100Mbps management port.*
- ***V-Secure V-100 -** This is a 100Mbps device, with 256MB RAM (max 1GB), two copper 10/100Mbps detection ports and a single 10/100Mbps management port.*
- ***V-Secure V-1000 -** This is a "fractional Gigabit" device (rated in the region of 250Mbps), with 256MB RAM (max 1GB), two copper 10/100/1000Mbps detection ports (which can be replaced by fibre ports at extra cost) and a single 10/100Mbps management port.*

The unit submitted for testing was the V-100, to be tested as a 100Mbps device.

This is a 1U device with two copper 10/100Mbps ports at the rear of the device, and one 10/100 management port at the front.

The V-100 is fitted with 256MB of RAM as standard (expandable to 1GB). Coloured status LEDs on the front panel provide an immediate indication of when traffic is being passed or blocked, and a mini-D serial port connector provides access to the Command Line Interface (CLI).

Figure 2 - V-Secure: V-100 appliance

Since this is a standard PowerPC platform (running VXworks), there is no built-in redundancy or high-availability features in terms of multiple fans or power supplies. The custom dual-port network card used for the detection interfaces does, however, provide a fail-open bypass capability to allow traffic to flow in case of hardware failure. This can be configured to fail closed if required

V-Secure also provides (at additional cost) an Active-Passive High Availability (HA) mechanism. This enables two redundant V-Secure IPS devices to provide mutual backup in case of application, hardware or link failure. The High Availability solution includes "*link-down*" trigger configuration options such as configuration of link down time-out and idle line time-out. Configuration files and the adapted network base-lines are automatically synchronised between the two devices, thus enabling the administrator to manage the two redundant devices through one console.

At the time of writing, V-Secure is working on an Active-Active HA capability that will be available in a future release.

Security Modules

Because of constantly changing network traffic patterns, an effective IPS needs to quickly adapt to its surroundings without human intervention.

Detection mechanisms used by the IPS must be capable of distinguishing between normal and abnormal behaviour, even though the differences between them may be subtle. In case the IPS misidentifies traffic, it must incorporate a self-correcting mechanism in order to eliminate false positives.

The V-Secure IPS uses a number of techniques - represented by different security modules - in order to achieve this:

State Machine Module

The *State Machine Module* follows every individual user connection and verifies that each connection behaves according to accepted protocol standards. This is a deterministic process that detects and prevents protocol anomalies and organises the necessary data used in the behavioural analysis performed by other modules.

Signal Processing Module

The behaviour properties of a single TCP connection can change dramatically during a session and still be considered normal. As a result, the difference between normal and abnormal TCP connection behaviour is particularly hard to evaluate with a reasonable level of certainty.

TCP typically delivers data at a predictable rate (packets or bytes per second) according to the physical line bandwidth and the properties of the TCP stack. However, TCP can also deliver data at non-predictable rates due to legitimate communications problems, such as router failures, bandwidth peaks, host failures, and so on.

In addition, some applications make efficient use of TCP as the transport layer for transaction control, whilst others do not. Therefore, some applications can keep an idle TCP connection open for a long period of time, while others will terminate an idle connection after a short delay (the latter represents more efficient operation).

The *Signal Processing Module* and *Spectrum Analyser Module* provide a practical solution for reliably detecting the misuse of TCP resources and malicious attacks that can be attributed to TCP misuse.

The *Signal Processing Module* aggregates time measurement data collected from the State Machine and transforms them into the *frequency domain*. The result is a spectrum of frequency groups divided by their *intensities*, which represent the number of users that belong to each one of the frequency groups. A *Spectrum Analyser Module* then examines spectral development over time, and produces traffic behaviour parameters. These behaviour parameters are then sent to a fuzzy logic decision engine, which assesses how best to repel an attack.

Spectrum Analyser Module

The Spectrum Analyser divides and compares (in real-time) user behaviour groups in relation to protocol state machines. One way to detect abnormal network activities, for example, is to analyse a particular behaviour group compared to other behaviour groups according to the appropriate network characteristics. In a sense, the Spectrum Analyser mimics the thought process of a network security expert, except that it does the work in real-time.

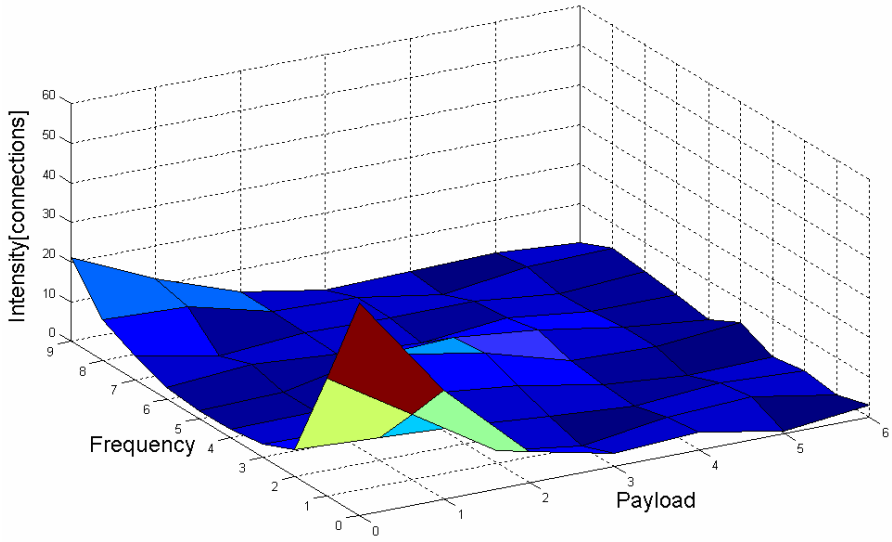

Figure 3 - V-Secure: The Spectrum Analyser

The Spectrum Analyser, is particularly astute at detecting and preventing the misuse of CPU and memory resources on networked servers.

These misuse activities are often the result of denial of service attacks such as connection floods, SYN floods, and known and unknown application floods.

The surface shown in *Figure 3* represents an aggregation of time measurements (TCP properties) at a certain moment in time. The *Frequency* axis represents the TCP packet's rate (Hz), whilst the *Payload* axis represents the TCP's average load (bytes/packet).

Together, these two axes characterise the behaviour of connections. The Intensity axis represents the number of connections belonging to each square (as defined by the frequency and the payload groups), and the intensity of colour represents the total number of connections. The frequency and payload of each connection are thus captured and presented to the Spectrum Analyser Module.

V-Secure claims that values of frequency and payload provide a good approximation of TCP connection behaviour for a given application, at a given point in time. These measurements, as visualised in the graph, present a comprehensive view and a method by which each group's intensity can be compared to others. Looking at the graph, we can see abnormal behaviour in part of the figure (the colour intensities of the abnormal squares).

Specifically, it is immediately apparent that the abnormal square is the one whose position is at the lower left (origin), because when compared to every other square, its intensity is much stronger than the others. In this case, the deep red colour represents a situation where there are large numbers of TCP connections with long idle times and a small average packet size. In comparison with other colours and values in the graph, it is apparent that this is abnormal behaviour which is worthy of additional investigation. In this case all the resources are located in a particular square which represents misuse of memory resources.

By comparing intensities in this way, it is possible to distinguish the difference between normal and abnormal behaviour.

Fuzzy Logic Module

The *Fuzzy Logic Module* is the decision engine of the system. It collects parameters from all the other modules in the system and assigns them an anomaly weight according to an adaptive *fuzzy membership function*. It then correlates these parameter weights and produces real-time decisions represented by a "*degree of attack*" value. Based on these degree of attack figures, the system is then able to introduce counter-measures to repel a perceived threat.

The use of "fuzziness" to represent features (e.g. TCP's connection properties) helps smooth the abrupt separation of normality and abnormality and provides a measure for the degree of normality or abnormality for a particular parameter. The fuzzy logic algorithm can process a large number of parameters, decide about their degree of anomaly, correlate between them and produce conclusions or decisions in real-time. The fuzzy logic algorithm provides an excellent balance between significance and precision.

It also overcomes traffic analysis difficulties that Internet communications usually present, providing a relatively simple way to draw definite conclusions from vague, ambiguous or imprecise information. Difficulties such as incomplete knowledge on the one hand and noisy signals on the other (something that usually happens when dealing with Internet traffic) are handled well by the fuzzy logic algorithm. V-Secure claims that fuzzy logic consumes less CPU and memory resources than other “traditional” analysis and approximation methods.

In the figure below, the XY plane shows the fuzzy inputs, and the Z-axis represents the degree of attack. The fuzzy logic module can determine when attack conditions exist by monitoring for unusual peaks and/or troughs in the decision surface model.

The Fuzzy Logic Module includes adaptive capabilities, meaning the sensitivity of the module is continuously being tuned in order to match the characteristics of the protected network. The adaptive algorithms include IIR (*Infinite Impulse Response*) filters that continually average traffic parameters and shape the fuzzy logic membership functions accordingly.

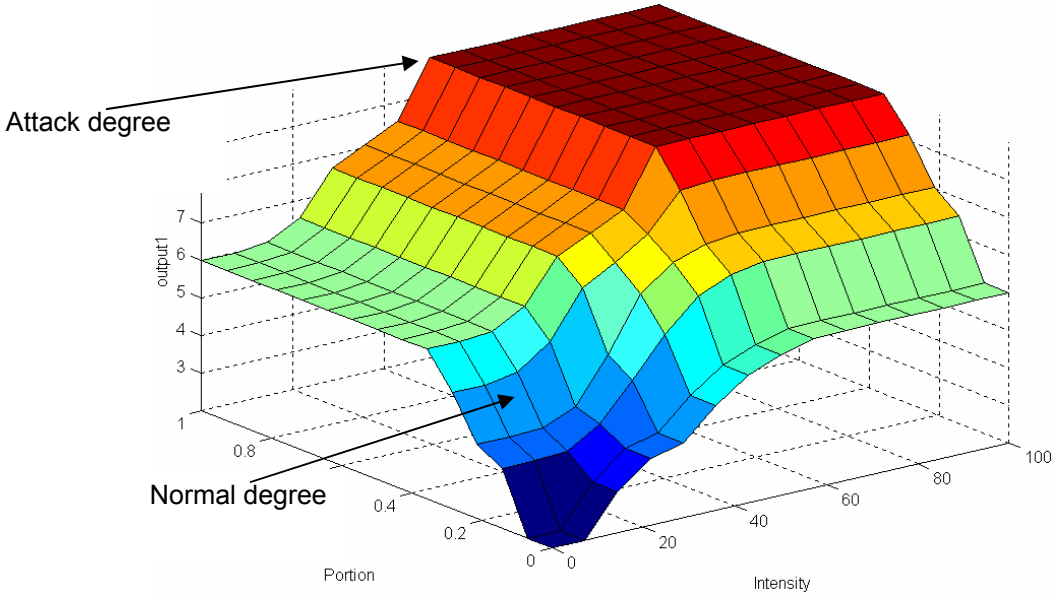

Figure 4 - V-Secure: Fuzzy Logic Decision Surface

The system learns normal traffic parameters such as:

- *Bandwidth parameters*
- *Protocol type distribution*
- *State machine behaviour parameters, such as transition frequencies between the protocol states both in the transmission and the application layers*

These capabilities allow the V-Secure IPS to establish normal behaviour baselines according to the date and the time of day.

The Fuzzy Logic Module can thus be thought of as an adaptive “expert system” that does not require human intervention to configure rules or thresholds.

A system that relies upon manually-tuned thresholds and rules produces wildly disparate detection quality, depending mostly on the individual skill level of the system administrator, as well as daily (or hourly) changes in what can be perceived as “normal” network traffic.

Attack Detection Module & Footprint Detection Module

V-Secure IPS incorporates traffic filters that are based upon a *Trap Buffer Algorithm*. The algorithm is responsible for characterising detected attacks, and creating dynamic filters in response. The resulting filters mitigate detected attacks in a way that does not adversely impact legitimate users.

In addition to the ability to block packets according to IP addresses and ports, V-Secure IPS can filter packets according to just about every parameter in the protocol header. The system can also terminate TCP connections.

When the *Attack Detection Module* determines that an attack is occurring, the *Footprint Detection Module* characterises the attack. This is achieved by using statistical analysis to develop one or more “*footprints”* of packets participating in the attack, such as values of one or more packet header fields, or, in some cases, information from the packet payload (such as a UDP DNS query string). The *Intrusion Response Module* filters incoming traffic participating in the attack, using the footprints developed by the footprint detection module

The system is adaptive, automatically reacting to changes in characteristics of an attack during the attack's life cycle by changing the footprint data dynamically.

Closed Feedback Module

The *Closed Feedback Module* is responsible for activating the system’s preventive counter-measures. It synchronises between attack detection and prevention sub-systems in order to pinpoint the optimal counter-measure that will mitigate an attack without affecting legitimate traffic.

The *Closed Feedback Module* achieves this by repeatedly checking the results of its actions, and continually refining the filters and comparing with the results of the previous filter to ensure that the attack continues to be mitigated effectively. If the action was successful in reducing the attack, then the system will continue to further refine the filter. The moment a new filter configuration increases the degree of attack or is ineffective in reducing the attack, the system rolls back to the previous good filter, or searches for a more appropriate response.

This means that it can take a variable amount of time to mitigate a new attack depending on the overall traffic composition on the network versus the composition of the attack, and the mitigation is not instantaneous as the optimal footprint is being determined (which may briefly allow attack packets through as the filter is refined completely). However, during the period when the device is optimising the footprint, it automatically throttles the bandwidth available to the suspicious traffic. Also, once the final filter is in place, the V-Secure device sends TCP reset packets to the victim to eliminate all of the malicious connections. Thus, the impact on legitimate users throughout the mitigation process is kept to an absolute minimum.

The closed feedback operation is also performed very quickly to minimise the duration that a legitimate address could be blocked unnecessarily (generally less than 1 second). As soon as the attack stops, the system ceases all counter-measures immediately.

During testing, we found that the speed with which attacks were mitigated was generally very good, and was remarkably consistent across all levels of traffic (even with very slow port scans and connection floods) - something which is not true of devices which rely on fixed thresholds to determine what is and is not “anomalous” traffic. The Trap Buffer approach would appear to produce the most optimal response for exploits in the general family of DDoS attacks.

Another feedback methodology employed by the system is a *dynamic blocking period*. When the system detects an attack, it initiates a very short blocking period, during which the system traces the blocked user and observes their behaviour. If subsequent activities represent legitimate network usage (such as application recovery from dropped packets), then the system immediately reduces the blocking duration to zero and releases the user.

If the user’s abnormal activities persist, then the system automatically increases the blocking duration to repel the attack. The dynamic blocking process is performed very quickly, so that normal traffic is not impacted.

In order to minimise false positives, V-Secure combines deterministic rules with heuristic and adaptive rules. The decision engine correlates between deterministic events, such as those coming from the *State Machine Module*, which will include such events as a session’s compliance with protocol standards, and traffic behaviour parameter values being generated by the *Spectrum Analyser*.

Performance

The aim of this section is to verify that the sensor is capable of detecting and mitigating attacks 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 V-Secure V-100 was tested up to 100Mbps, 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 mitigated under all load conditions. We would happily confirm V-Secure’s 100Mbps rating for this device. The V-100 handled 115,000 simultaneously open connections.

Basic latency figures were excellent at all traffic loads and with all packet sizes, ranging from 52µs with 25Mbps of 256 byte packets, to 72µs with 100Mbps of 1000 byte packets. Behaviour throughout the tests with no background traffic was very even and predictable, remaining at well under 100µs under all load conditions and with all packet sizes.

Placing the device under increasing loads of HTTP traffic also had minimal effect on the latency figures, which ranged from 57µs with 25Mbps of 256 byte packets, to 95µs with 100Mbps of 1000 byte packets. All of these figures are well inside the limits we would expect to see for a 100Mbps device, meaning the V-100 could be situated anywhere on a 100Mbps network, either internally or at the perimeter.

Naturally, performance when under attack is critical for an attack mitigation device, and the V-100 did not disappoint overall. Mitigation was handled well at almost all levels, but the V-100 does have a limitation on handling packets per second of around 70,000pps. Below this level, however, latency was excellent when under attack, ranging from just 72µs at 20Mbps of SYN flood traffic to 250µs at 60Mbps.

The V-100 performed consistently and completely reliably throughout our tests. Under eight hours of extended attack (comprising millions of sessions of legitimate traffic interspersed with some attacks) it continued to mitigate 100 per cent of attack traffic, whilst passing 99.992 per cent of legitimate traffic (blocking an average of 80 out of 1,000,000 legitimate sessions per run).

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

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

Security Effectiveness

We installed one sensor and enabled all the protection mechanisms apart from those which were not applicable in our test environment (*IANA protection, Outgoing NAT IP Scan,* and *Access Lists Policy*). We then configured the expected bandwidth settings and ran through all of our load generation tests, allowing the V-100 to learn from our “normal” traffic.

Attack detection/mitigation was excellent, with the V-100 detecting and successfully mitigating all but one of our attacks. Our Back Orifice controller was able to contact the Back Orifice Server installed on the protected network, receiving a large number of packets in response. These outbound packets were detected as an *Outbound UDP Flood*, but the V-100 was unable to mitigate it successfully.

Performance in the high volume detection/mitigation test was impeccable across the board, with perfect detection and mitigation at all load levels.

A major concern in deploying an in-line device is the blocking of legitimate traffic. Once we had configured the *Trusted Groups* and *Approved* the *Protected Hosts* the V-100 ran through every single one of our tests without raising a single false positive alert.

Resistance to our evasion attempts proved very good, with the V-100 successfully detecting all of the fragmented and slow attacks we ran. It would appear to be very difficult - perhaps impossible - to evade this device by simply slowing down port scans and connection floods thanks to the fuzzy logic mechanism employed to compare “*normal*” vs. “*abnormal*” traffic.

When employing URL obfuscation techniques to attempt Web vulnerability scans, the V-100 successfully detected and blocked all except the Whisker splice. V-Secure is working on this at the time of writing.

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 V-Secure IPS is straightforward. The *Management Server* and *SuperVisor Client* are both native Windows applications, and install quickly and easily from CD.

Initial configuration of the V-Secure IPS appliance - to set the IP address, net mask, management IP address and management server secret key (to ensure secure communications between sensor and management server) is all performed via the Command Line Interface (CLI) over a serial connection.

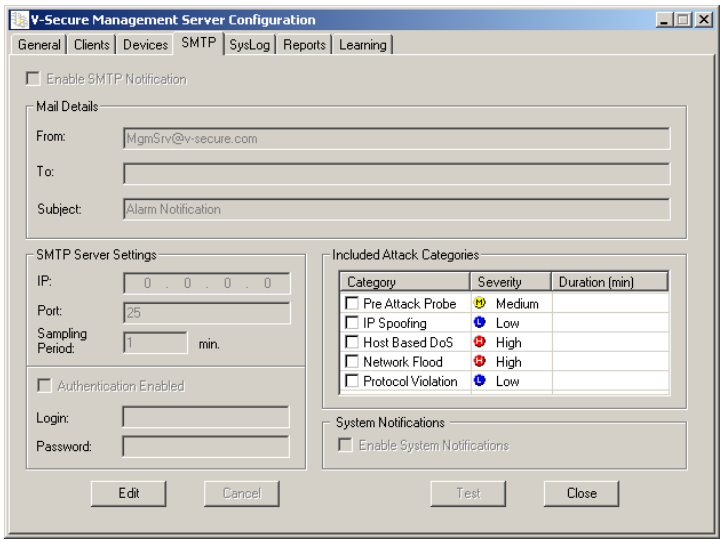

Figure 5 - V-Secure: Configuring SMTP alerts in the Management Server

Because this is a three-tier management system, it is then necessary to configure the Management Server itself. This is mainly a matter of adding details of authorised clients (i.e. those hosts which will run the SuperVisor Client software) and those appliances which will be managed and monitored by the server. A maximum of 25 appliances can be monitored from a single server (there is no restriction on the number of clients), and only those appliances and clients defined here will be able to communicate with the Management Server.

If required, this is also the place to configure SMTP notifications, *syslog* alerts, and system-wide reporting.

It is possible to restrict SMTP and syslog alerts to specific attack categories of required - for example, send alerts to syslog for floods, DOS attacks, protocol violations and IP spoofing attacks, but only send e-mail alerts for floods and DOS attacks. All attacks are reported to the SuperVisor Client by default.

Without enabling reporting at this stage, the Management Server will not create log files for processing by the V-Secure Reports Generator. Retention periods can be set for long-term and short-term logs, and scheduled report periods and recipients can be defined if required.

The documentation appears to be very good, with clear explanations, screen shots, and step-by-step instructions on how to manage the V-Secure IPS system.

User guides are available (PDF only) for the *V-Secure IPS Management Server* (with SuperVisor Client), the *CLI* and the *NetVisor Client*.

Configuration

On starting the SuperVisor Client, the administrator is prompted for a password only. This is because V-Secure does not enforce any form of role-based access to the system, and does not even have the concept of user accounts. Instead, each authorised client PC is given access to the system based on the SuperVisor permissions defined in the Management Server:

- *None*
- *Monitoring only*
- *Monitoring and run-time control*
- *Monitoring and configuration*

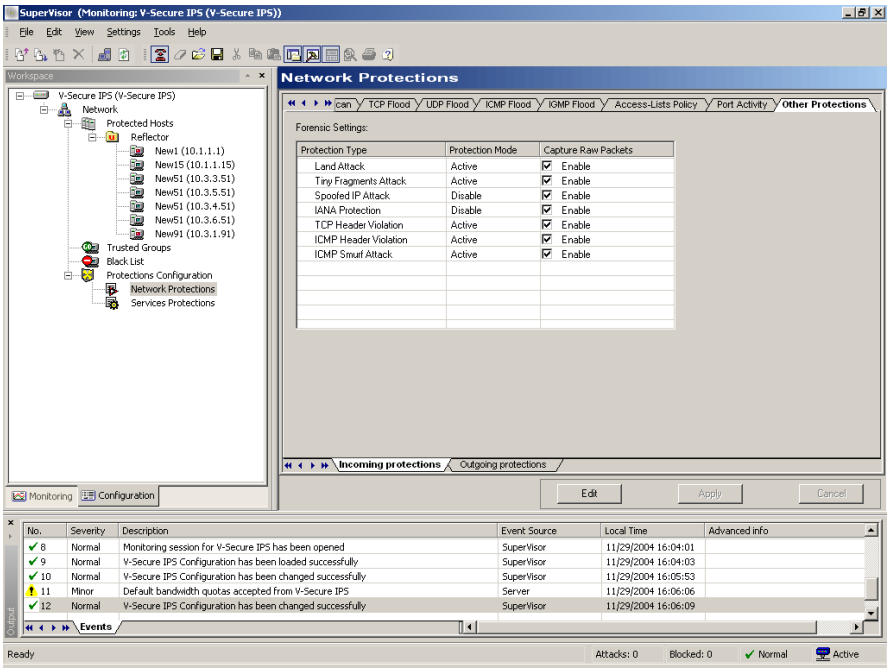

Figure 6 - V-Secure: Configuring Network Protection in SuperVisor Client

Note that it is also possible to restrict access to specific attack categories for configuration purposes:

- *Pre-attack Probes*
- *IP Spoofing*
- *Host-Based DOS*
- *Network Flood*
- *Protocol Violation*

So although it is possible to restrict a client to a monitoring only role rather than a full administrator, or to being able to configure only specific types of attack response, the granularity is only as far as the host PC. We would prefer to see this extended to provide true role-based user accounts offering different levels of access per-user from the same PC.

After successful authentication to the Management Server, the administrator is provided with a list of available appliances which are managed by that server. Selecting from the list attaches to that V-Secure appliance with the appropriate permissions as defined for the client.

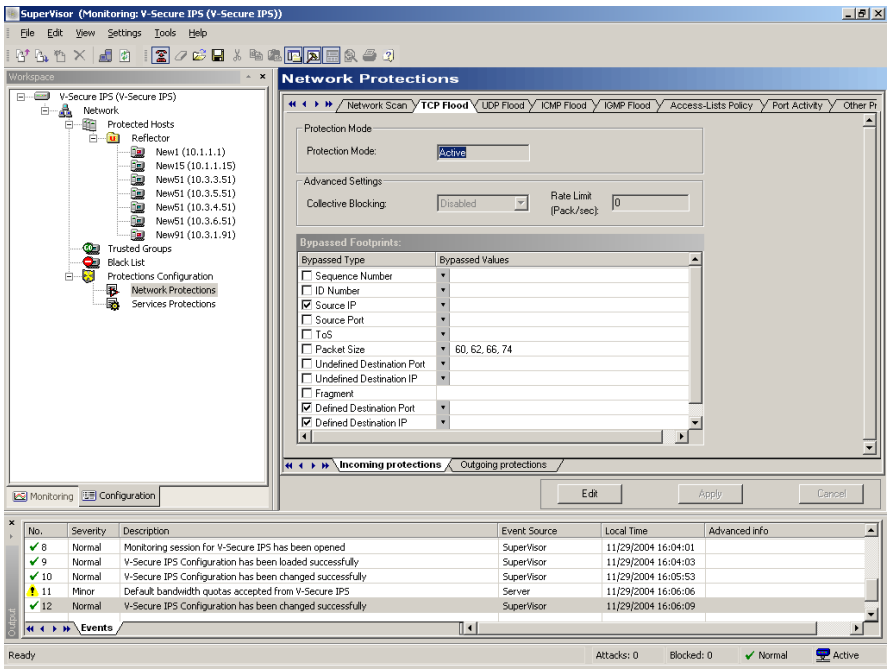

Figure 7 - V-Secure: Configuring TCP Flood parameters

Thus, the SuperVisor Client is only capable of managing, configuring and monitoring a single appliance at a time. Whilst this is clearly not a scalable solution in large deployments, it is adequate for smaller deployments of one or two devices. Where large numbers of appliances need to be managed from a single console, the Java-based *NetVisor Client* is available at additional cost. This provides the ability to deploy policies to, and view alerts and reports from, multiple clients simultaneously.

The layout of the SuperVisor Client user interface follows the tried-and-tested three pane format:

- ***Workspace Area –*** *This is the left hand pane which contains a hierarchical display of the configuration or monitoring options for the device being managed*
- ***Parameter Area –*** *This is the right hand pane, which contains the available parameters for the currently selected configuration option, or the output for the currently selected monitoring option.*

Where multiple configuration options are available (such as Inbound *and* Outbound *settings, for example), these are presented as tabbed windows, making it easy to switch between them.*

- ***Output Area –*** *This pane runs along the bottom of the screen and contains text-based system function information such as when a monitoring session has been opened to a device, or when a configuration has been uploaded successfully (or failed to upload).*

Two toolbars are available above the panes, and panes can be activated or deactivated (hidden) at the click of a mouse. Overall, although a lot of information is presented on-screen at times, the user interface is very easy to navigate and clever use of tabbed windows makes it possible to present large numbers of configuration options without overwhelming the administrator.

V-Secure IPS can be configured to run in one of three different operating modes:

- ***Transparent Mode –*** *V-Secure IPS behaves like a layer 2 bridge device. Traffic is passed straight through the appliance without any security testing. Transparent Mode is primarily used for device testing.*
- ***Detection Mode –*** *V-Secure IPS performs configured security tests but does not enforce the configured policies. The administrator will see the alerts and any dynamic changes to the V-Secure configuration (such as creation of new service/host combinations, etc.), but no mitigation will be performed on the live traffic. Detection Mode can be used to test the effectiveness of a particular security policy configuration without the risk of adversely affecting legitimate traffic.*
- ***Active Mode –*** *V-Secure IPS performs configured security tests and enforces the configured security policies, mitigating traffic where necessary. Active Mode is used under normal operating conditions to protect the network.*

Policy Management

Configuring a security policy within V-Secure is reasonably straightforward if you do not wish to tweak the default settings (not recommended until you are familiar with the system). For most environments, it is enough to configure only the *Network* settings, which consist of *Global Settings* and *Network Protections*.

The *Global Settings* allow the administrator to specify available bandwidth for inbound and outbound traffic (this is used to set “quotas” for different protocols as a starting point to determine when flood attacks are in progress), the IP address range of the protected network (to help prevent spoofing attacks), and the *learning period* (the statistical sampling period - day, week or month - used by V-Secure IPS to establish baseline traffic characteristics). An additional check box allows the administrator to reset the baseline learned statistics at any point, forcing the system to re-learn the baselines following major changes to the characteristics of the protected network traffic.

The *Network Protection* settings allows the administrator to specify which protection mechanisms are to be activated, deactivated, or operated in detection-only mode for both inbound and outbound traffic.

These protection mechanisms cover a wide-range of common rate-based attacks, such as:

- **Network Flood Attacks –** The Network Flood attack creates large volumes of irrelevant traffic in order to fill available network bandwidth, denying use of network resources to legitimate users. V-Secure can be configured to detect and block network flood attacks by defining attack footprints, which are selected fields in the packet header or payload. V-Secure automatically detects the footprints using techniques described in the *Architecture* section of this report, and is capable of generating dynamic filters to protect against the attack. The system can then be manually configured to bypass certain footprint types, which prevents traffic from being blocked based on the value of the bypassed footprint. V-Secure IPS generates filters based on the narrowest possible footprint initially, building broader filters as it becomes necessary to protect against an attack. When the configured footprint-based filters do not protect the network sufficiently, the system can employ *Collective Blocking* as a last resort, where all inbound packets of the same type are blocked regardless of their footprint. Network Flood attack types include:
    - TCP SYN Flood
    - UDP Flood
    - UDP Flood with ICMP Back Scattering
    - ICMP Flood
    - IGMP Flood
- **Network Scan Attacks –** Network Scans are generally performed before an attack and are used to determine vulnerable services and hosts in the network and report the information back to the attacker. V-Secure IPS recognises these scans and blocks the pattern of the scan, both inbound scans from the outside world and outbound scans from inside the protected network. For outbound scans, V-Secure recognises scans initiated from both the public IP addresses of devices such as Web, Mail, and FTP servers, or from private IP addresses of internal user workstations connecting to the outside network through a NAT/proxy. Different sensitivities can be set in order to detect even very low-rate scans, and V-Secure can be set to block completely the IP address of the offender or block it only for those services not defined within SuperVisor (which ensures legitimate services are not blocked to legitimate hosts when their IP address has been spoofed by an attacker). Where total blocking is not desired, the suspicious connection can be rate limited instead. Network Scan types include:
    - TCP Host Scan
    - TCP Port Scan
    - UDP Host Scan
    - UDP Port Scan
    - ICMP Sweep
    - Stealth Scan
- **TCP SYN Attacks –** TCP SYN attacks deny user access to a host by using all available host resources to handle incomplete TCP connection requests. These can be thought of as low-rate SYN floods, and they are often detected before the network-wide SYN Flood protection mechanism kicks in.
- **Abnormal Port Activity Attacks –** Port Activity protection is designed to protect the network from self-propagating worm attacks which use random or pseudo-random propagation methods.

These methods allow the worms act against the protected network in such a way that no source IP can be identified. By tracking unusual port activity, V-Secure can detect and block these attacks, which appear as persistent attempts to access illegal ports from distributed IP source addresses.

- **Connection Flood Attacks –** Connection Flood attacks consume server resources by creating a large number of TCP connections with the server and filling those connections with unnecessary traffic.
- **Web Crack Attacks –** Web Crack attacks target Web servers by scanning for vulnerable files and scripts and attempting to gain access to authentication data (password files, etc.) to access the Web server

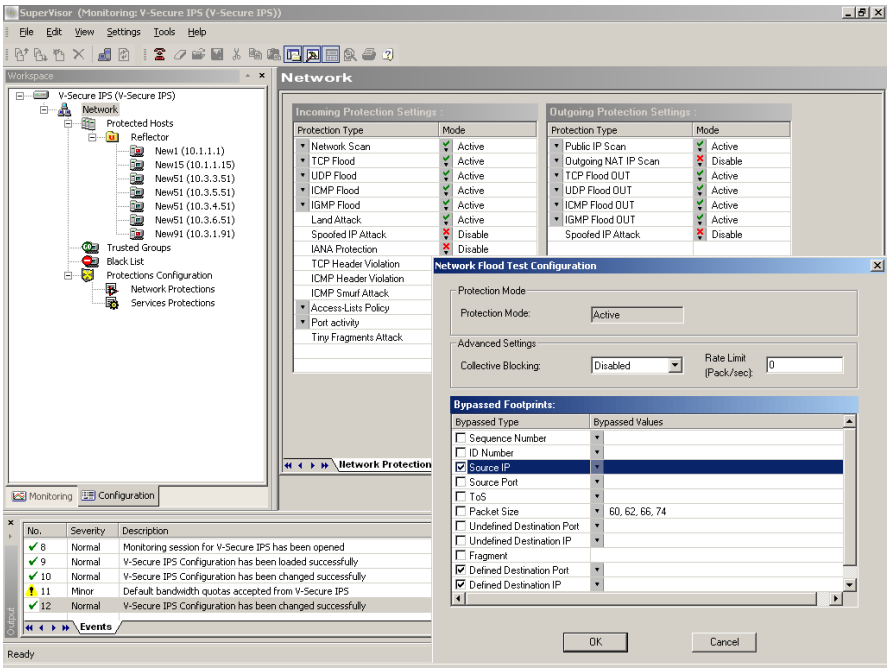

Figure 8 - V-Secure: Configuring Network Global Settings

- **Spoofed IP Attacks –** Spoofed IP attacks attempt to gain access to the network by sending a packet with a source IP address located on the destination network. The network assumes this is a trusted address and allows the packet into the network
- **Outbound Spoof Attacks –** Outbound packets generated by DDoS tools, or “zombies,” with spoofed source IP addresses. These attacks generate a large amount of irrelevant traffic, denying network capacity to legitimate traffic.
- **Land Attacks –** A variation on the Spoofed IP attack, the Land attack attempts to gain access to the network by sending a packet with identical source and destination IP addresses. The addresses are members of the destination network.
- **IANA Attacks –** IANA attacks attempt to gain access to the network by spoofing IANA reserved private IP address space. The source of these attacks is difficult to determine, because IANA addresses are not routable over the Internet and cannot be traced
- **TCP Header Violation Attack –** Attacks that try to exploit existing vulnerabilities in some of the TCP/IP stack implementations through corrupted TCP packets (i.e., non-standard TCP header parameters).

- **ICMP Header Violation –** Attacks that try to exploit existing vulnerabilities in some of the IP stack implementations through corrupted ICMP packets (i.e., unknown ICMP message types, etc.)
- **ICMP Smurf Attacks –** Attacks that create network congestion by sending large numbers of ICMP echo reply packets
- **Terminal Telnet Attacks –** Many times hackers use a terminal telnet application in order to gain access to and exploit public applications such as Web, FTP, etc.
- **Tiny Fragments Attacks –** Tiny IP fragments carrying malicious content in the application layer. These packets are used to evade attempts to detect exploitation of vulnerable network applications

Some of these protection methods (i.e. *Web Crack*, amongst others) offer the ability to capture raw packets from the suspicious session which can later be viewed using Ethereal. Many of these provide a number of individual configuration parameters, mainly based around certain traffic characteristics which are to be ignored when determining attack conditions. As mentioned before, unless you are confident in amending these at the outset it is best to leave the system to run through its learning phase (in detection mode, initially, to prevent false positive conditions from disrupting legitimate traffic) and then tune these settings directly from the alerts raised by the system.

This is a much more efficient way of configuring the system than attempting to wade through all the possible tuning parameters before deploying the first security policy.

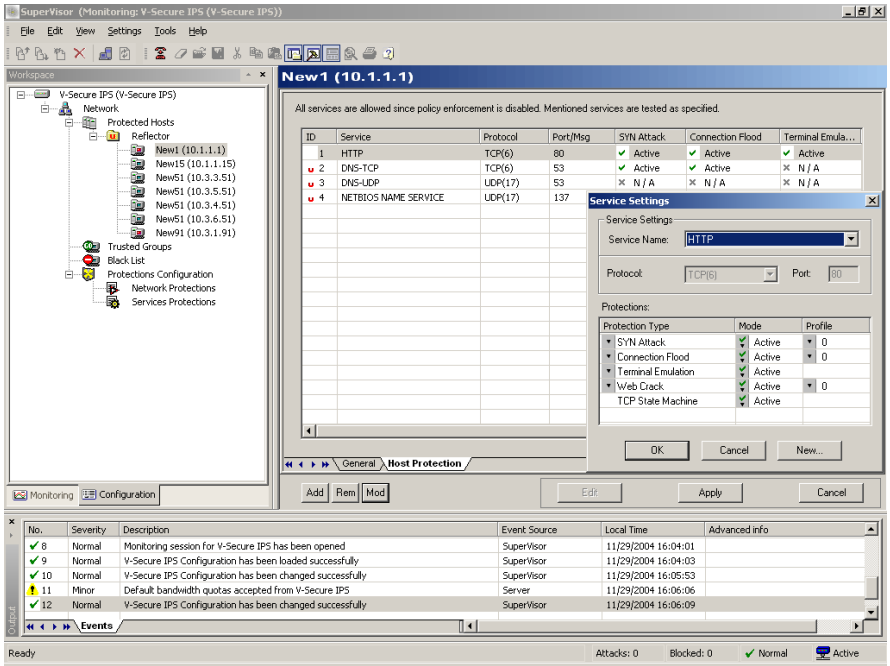

Figure 9 - V-Secure: Configuring Protected Hosts

Having set the *Global Settings* and *Network Protection* settings – a five minute task – the policy can be uploaded to the appliance, and this very basic policy is good enough to provide generic protection immediately. Note that there is no provision for creating, saving or deploying multiple policies in SuperVisor, since it is restricted to managing a single device at a time.

If it is known at the outset what services are running on which protected servers, it is possible to define these ahead of time and specify the protection modes to be applied to each server. A wide range of services are defined in the system, and it is possible to define new ones by entering the port and protocol over which the service runs.

For each service, a number of server-specific (as opposed to network-wide) protection mechanisms can be activated. For example, by listing HTTP as a service on a particular host, that Web server can be protected from general TCP protocol violations, SYN Attacks, connection floods, terminal telnet attacks, and Web cracking attempts.

If the administrator prefers not to define these ahead of time, the system will detect the host/service combinations as unknown services and present them in the GUI to be approved or rejected, at which time the required protection mechanisms can be specified. Where the same protection mechanisms are to be enabled for several servers (for example, all the Web servers on a DMZ), then *Protection Groups* can be defined. The protection mechanisms defined for the Group are applied automatically to all the servers within it.

Other items which can be defined at the outset – if known – to reduce tuning time later, are *Trusted Groups* and *Black List* hosts. *Trusted Groups* (IP addresses which are known to be trustworthy and thus will never be considered an attacker) and *Black List* hosts (IP addresses which are known to be untrustworthy, and who are not permitted to communicate with protected networks or hosts over any protocol under any circumstances). Black Lists can be created for both inbound and outbound traffic.

Alert Handling

V-Secure's monitoring capabilities enable the administrator to analyse both current and historical attacks using the SuperVisor Client. V-Secure IPS normally generates countermeasures automatically against inbound attacks, but it is also possible to launch *Run-Time Commands* to override the standard countermeasures, and counter any attack immediately based on the most current data.

Available Run-Time Commands include the ability to add a footprint or an IP address to a bypass list (or remove), insert to black list, or perform a WHOIS lookup.

Each alert appears in the *Active Attacks* tab of the *Monitoring* window, showing the attack type, degree of attack (*low, medium* or *high*), start time, end time, source IP, and status (*currently active* or *closed*). Selecting any alert brings up additional information below depending on whether the *Anomaly* view or *Attack Footprint* view option has been selected.

The *Anomaly* view displays information about overall traffic and protocol mix and shows anomalies in traffic patterns that indicate an attack condition.

The *Attack Footprint* view displays all of the footprint data which has been gathered to identify the attack. If the administrator is happy with how the attack is being mitigated, then there is nothing for him to do.

However, every entry in the footprint table shows detailed information about each field in the TCP header, UDP header or ICMP header of each packet that was determined to form part of the attack, and the data which is actively being used to create the mitigation filter (the "*footprint*") is highlighted in red.

Each item has a checkbox against it to allow the administrator to fine tune the footprint by deselecting (bypassing) one or more items of footprint data. For example, if the source IP is being used to filter packets and the administrator feels this is too broad (perhaps because it could lead to a legitimate host being blocked if IP addresses are being spoofed) then he can select the source IP address for bypassing.

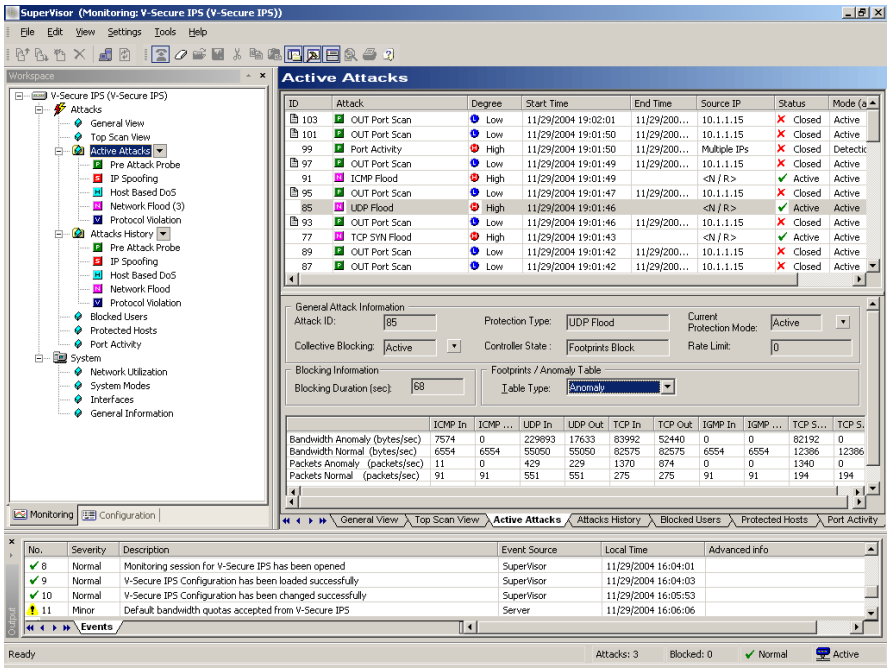

Figure 10 - V-Secure: Active Attacks (Anomaly view)

The effect is immediate, and it is possible to see the process by which V-Secure re-evaluates the footprint as it takes account of the fact that it can no longer use that particular piece of data. The fact that the source IP address should not be used to footprint that particular kind of attack is recorded in the detailed configuration settings for that attack type so it will not be used in future. It is possible to further refine this manually, or to remove the bypass should it prove to be an error, by accessing the detailed configuration parameters directly in the *Configuration* tree in the *Workspace Area*.

Once an attack has ceased, it is marked as *closed* in the *Active Attacks* window, and after a period of time, all closed attacks are moved to the *Attacks History* window, so they no longer clutter up the current view. If individual hosts are blocked for perpetrating an attack, they are displayed in the *Blocked Users* tab. As with Active Attacks, once the blocking period on a host has expired, it is marked as *closed*, though it remains in the Blocked Users tab for analysis purposes.

It is possible to configure V-Secure to recognise new hosts and services on the protected network when a host on the public network attempts to contact it.

When V-Secure detects a new host, or an existing host with a new service, on the protected network, the host appears in the *Protected Hosts* section of the *Configuration* tree in the *Workspace Area* with an exclamation point against it. In addition, a new host appears in the *Protected Hosts* tab in the *Monitoring* window - also with an exclamation point against it - and a new service appears in the *Host Protection* tab for that particular host.

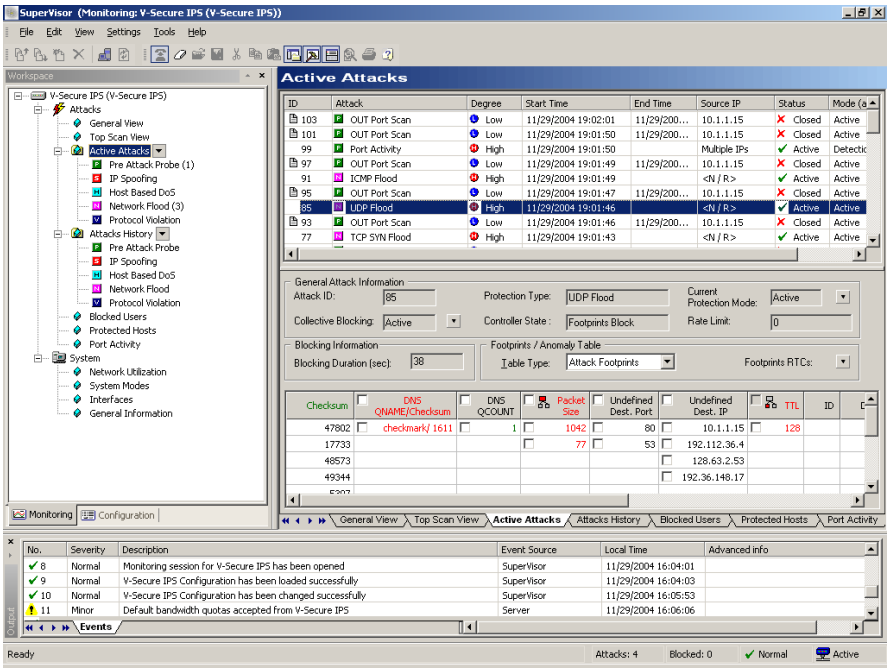

Figure 11 - V-Secure: Active Attacks (Attack Footprint view)

This provides notification for the administrator that someone has tried to access a new or unrecognised service on a host on the protected network. One of two things can happen at this point. The administrator can mark the combination as bad, thus ensuring that communications to that host and port are blocked. Or it can be marked as *Approved*, at which point it becomes a *Protected Host*, allowing the administrator to define which host-specific protection mechanisms should be activated (protocol violations, SYN Attacks, connection floods, terminal telnet attacks, and Web cracking attempts, for example).

Reporting and Analysis

As mentioned in the previous section, V-Secure's monitoring capabilities enable the administrator to analyse both current and historical attacks using the SuperVisor client.

V-Secure IPS monitors the following activities:

- *Overall attack analysis*
- *Top Scan analysis*
- *Active attacks*
- *Historical attack activity*
- *Blocked users*
- *Protected hosts*
- *Port Activity analysis*
- *Network utilisation*

- *System modes*
- *Network interfaces*
- *System information*

Some of these - like *Active Attacks, Attacks History, Blocked Users* and *Protected Hosts* - are text-based screens which we have covered in the previous section. Others - such as the *General Attack View, Top Scan View, Port Activity* and *Network Utilisation* - are all graphical screens which provide useful insights into the current operational status of the V-Secure appliance and the traffic being monitored.

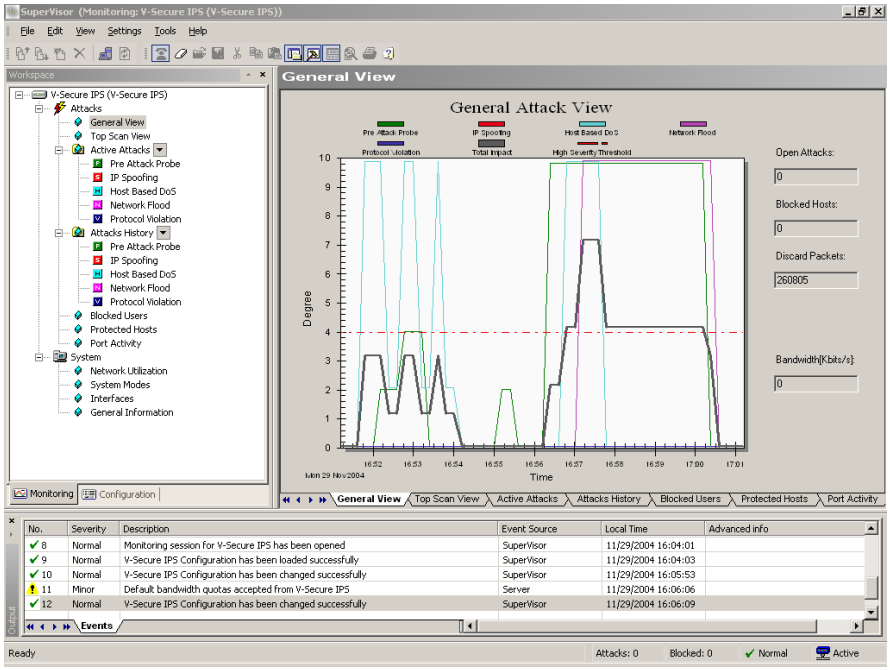

Figure 12 - V-Secure: Monitoring Attacks

The *V-Secure Reporter* is a web-based application residing on the Management Server that can be accessed by using a standard Web browser through a simple username/password authentication procedure. Thus, it is not necessary to provide access to the SuperVisor Client in order to allow an administrator or other user to view report data.

The reports are somewhat basic, with no drill-down capability, and with the rather more serious restriction (in multi-device deployments) that it is only possible to create reports for one device at a time. However, given that this is a rate-based IPS system rather than an IDS or content-based IPS, there is not much in the way of detailed forensic analysis required, so the basic Reporter is probably adequate for most situations. For those who need something more extensive, plus the ability to consolidate reports across multiple devices, the NetVisor Client provides this at additional cost.

A simple query screen is available to specify the type of report to run, attack category to focus on, host address, resolution (day, week, month), start date and V-Secure appliance ID. It is not possible to specify a start and end time - you must specify the start date and then accept one of the pre-defined reporting periods.

The following reports are available from V-Secure IPS Reporter:

- ***General Summary -*** *provides information on attack activity against the protected network for the selected time interval*
- ***General Attack -*** *provides information on specific attack type activity against the protected network for the selected time interval*

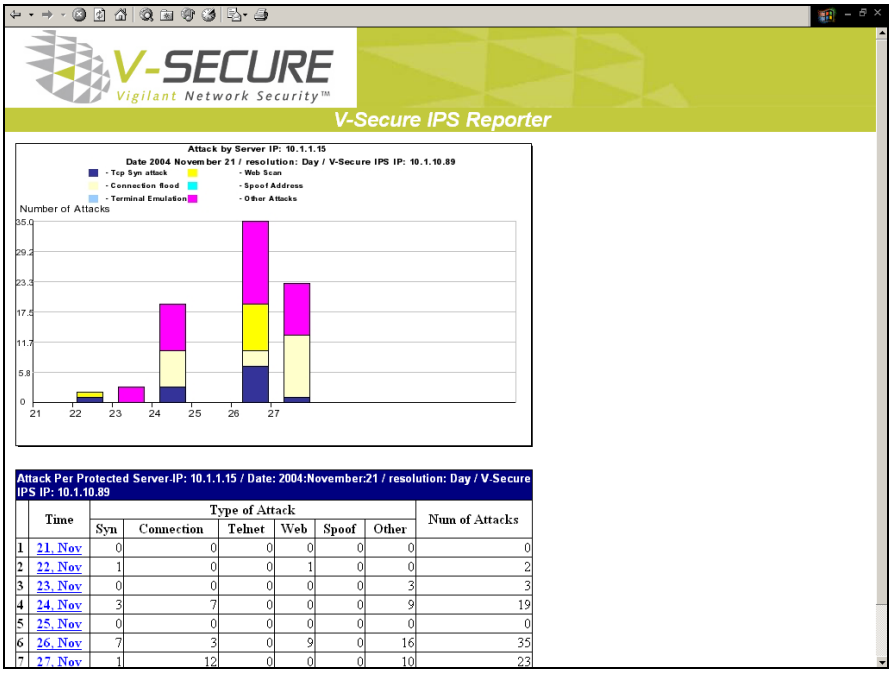

Attack Per Protected Server IP: 10.1.1.15 / Date: 2004:November:21 / resolution: Day / V-Secure IPS IP: 10.1.10.89

	Time	Type of Attack						Num of Attacks
		Syn	Connection	Telnet	Web	Spoof	Other	
1	21, Nov	0	0	0	0	0	0	0
2	22, Nov	1	0	0	1	0	0	2
3	23, Nov	0	0	0	0	0	3	3
4	24, Nov	3	7	0	0	0	9	19
5	25, Nov	0	0	0	0	0	0	0
6	26, Nov	7	3	0	9	0	16	35
7	27, Nov	1	12	0	0	0	10	23

Figure 13 - V-Secure: Typical report

- ***Host Attack -*** *provides information on attacks against a specific host over a selected time interval (as both a graph and a table(*
- ***Top Attacker per Host -*** *provides information on the source IP address of the ten most frequent attackers against V-Secure IPS (as both a graph and a table)*
- ***Top Blocked Attackers -*** *Similar to the Top Blocked Attackers report, this report provides information on the source address of the ten most frequent attackers against a specific host on the protected network (as both a graph or a table)*
- ***Network Utilisation Statistics -*** *provides information on network traffic over time. This report displays data on overall traffic, protocol mix, and packet discards*
- ***Event View -*** *provides highly detailed information directly from the Management Server's log files on every incoming attack during the selected time interval. The Management Server log files update automatically from all managed V-Secure IPS devices every night at midnight. Attack status is current as of the most recent log file update.*

Verdict

Performance

The V-Secure V-100 was tested up to 100Mbps, 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 mitigated under all load conditions. We would happily confirm V-Secure's 100Mbps rating for this device.

Basic latency figures were excellent at all traffic loads and with all packet sizes, ranging from 52µs with 25Mbps of 256 byte packets, to 72µs with 100Mbps of 1000 byte packets. Behaviour throughout the tests with no background traffic was very even and predictable, remaining at well under 100µs under all load conditions and with all packet sizes.

Placing the device under increasing loads of HTTP traffic also had minimal effect on the latency figures, which ranged from 57µs with 25Mbps of 256 byte packets, to 95µs with 100Mbps of 1000 byte packets.

All of these figures are well inside the limits we would expect to see for a 100Mbps device, meaning the V-100 could be situated anywhere on a 100Mbps network, either internally or at the perimeter.

Naturally, performance when under attack is critical for an attack mitigation device, and the V-100 did not disappoint overall. Mitigation was handled well at almost all levels, and latency was excellent when under attack with all but the heaviest loads, ranging from just 72µs at 20Mbps of SYN flood traffic to 250µs at 60Mbps. Performance under the heaviest loads was hampered by the overall packet processing limit of the device, which is around 70,000 packets per second.

While the overall packet processing limit of the device seems low for a 100Mbps device, it should be noted that this limitation is on traffic passing through the appliance - higher loads of mitigated traffic can be handled, since they are rejected at the external interface rather than being passed through the appliance.

The V-100 performed consistently and completely reliably throughout our tests. Under eight hours of extended attack (comprising millions of sessions of legitimate traffic interspersed with some attacks) it continued to mitigate 100 per cent of attack traffic, whilst passing 99.992 per cent of legitimate traffic. We would prefer to see 100 per cent of legitimate traffic passed, but this is a very small percentage of failures under an extreme load of continuous attack traffic.

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

Security Effectiveness

Attack detection/mitigation was excellent, with the V-100 detecting and successfully mitigating all but one of our attacks. Our Back Orifice controller was able to contact the Back Orifice Server installed on the protected network, receiving a large number of packets in response. These outbound packets were detected as an *Outbound UDP Flood*, but the V-100 was unable to mitigate it successfully. This should be fixed in future releases.

Performance in the high volume detection/mitigation test was impeccable across the board, with perfect detection and mitigation at all load levels.

A major concern in deploying an in-line device is the blocking of legitimate traffic. Once we had configured the *Trusted Groups* and *Approved* the *Protected Hosts* the V-100 ran through every single one of our tests without raising a single false positive alert.

Resistance to our evasion attempts proved very good, with the V-100 successfully detecting all of the fragmented and slow attacks we ran, and all but one of the Whisker evasion techniques. It would appear to be very difficult - perhaps impossible - to evade this device by simply slowing down port scans and connection floods thanks to the fuzzy logic mechanism employed to compare "*normal*" vs. "*abnormal*" traffic.

Usability

Whilst recognising that there is no halting the tide of Java-based applications, there is no denying that a native Windows application is very much slicker, faster, and easier to use. So it was with the SuperVisor Client.

Initial configuration and deployment is very straightforward, especially if you configure the basic network-level global protection in *Detection* mode and allow the V-Secure device to tell you which hosts and services it finds for you to configure further. If you have too many protected hosts (in the hundreds), however, the GUI has some issues handling them en masse (i.e. when creating a protected group) and the overall performance (latency) of the system suffers as well. Most organisations are likely to have tens rather than hundreds of specifically protected servers behind each individual V-Secure appliance, however, so this should not be too much of a problem.

Reporting is very basic, with the Web-based *Reporter* utility offering high-level reports with no drill-down capability and no extensive query or report scheduling capabilities. Nevertheless, monitoring capabilities are excellent, with a wide range of monitoring screens providing good insight into the device behaviour and underlying network traffic. For a rate-based system, this level of monitoring is more important than detailed reporting in our opinion, and V-Secure scores highly here.

All in all, SuperVisor offers a very intuitive and usable means of managing, monitoring and configuring a single V-Secure device. For those who require an enterprise-capable solution with the ability to manage multiple-devices effectively, V-Secure offers an additional-cost option in the form of NetVisor.

Contact Details

Company name: V-Secure Technologies Inc.

Internet: www.v-secure.com

Address:
Park 80 West, Plaza II
Suite 200
Saddle Brook, NJ 07663
USA

Tel: +1 (201) 291 2845

Fax: +1 (201) 2912742

E-mail: info@v-secure.com

APPENDIX A – TEST RESULTS

The aim of this procedure is to provide a thorough test of all the main components of an in-line rate-based IPS/Attack Mitigation 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 in front of 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 external 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 rate-based exploits accurately, whilst remaining resistant to false positives. During the attacks, the victim is expected to remain available and responsive.

All tests in this section are completed with **no background network load**, and only live exploits/attack tools are used (no replay traffic).

Whilst the various replay tools are generally useful for testing signature based IPS/IDS systems, their use in conjunction with rate-based attack mitigators is unpredictable, and thus the use of live tools is preferred in this case.

Test 1.1 - Attack Detection/mitigation

Whilst it is not possible to validate completely the entire detection / prevention range of any sensor, this test attempts to demonstrate how accurately the sensor detects and blocks a wide range of common rate-based attacks, port scans, and Denial of Service attempts.

The sensor is installed and all possible detection modes are activated. The vendor is permitted to tune the product (or to configure the device to learn automatically) in order to match the expected loads of attack and background traffic - just as they would for a normal customer. All attacks are run with no load on the network and no IP fragmentation.

Our attack suite covers the following areas:

- ***Test 1.1.1 -*** *SYN Flood*
- ***Test 1.1.2 -*** *TCP SYN Attack (low-rate SYN Flood)*
- ***Test 1.1.3 -*** *ICMP Flood*
- ***Test 1.1.4 -*** *Distributed Denial Of Service (DDOS)attack*
- ***Test 1.1.5 -*** *UDP Flood*
- ***Test 1.1.6 -*** *IGMP Flood*
- ***Test 1.1.7 -*** *Connection Flood (fast)*
- ***Test 1.1.8 -*** *Connection Flood (slow)*
- ***Test 1.1.9 -*** *Random protocol violations (invalid packets)*
- ***Test 1.1.10 -*** *Trojan response (external host attempts connection to internal Trojan and receives response)*
- ***Test 1.1.11 -*** *ICMP Sweep (inbound)*
- ***Test 1.1.12 -*** *ICMP Sweep (outbound)*
- ***Test 1.1.13 -*** *SQL Slammer*
- ***Test 1.1.14 -*** *Spoofed IP attack*
- ***Test 1.1.15 -*** *Web vulnerability scan*
- ***Test 1.1.16 -*** *Port Scan (full TCP connect)*
- ***Test 1.1.17 -*** *Stealth Port Scan*
- ***Test 1.1.18 -*** *FIN Port Scan*
- ***Test 1.1.19 -*** *UDP Port Scan*
- ***Test 1.1.20 -*** *Null Port Scan*
- ***Test 1.1.21 -*** *Xmas Port Scan*
- ***Test 1.1.22 -*** *IP Protocol Port Scan*
- ***Test 1.1.23 -*** *ACK Port Scan*
- ***Test 1.1.24 -*** *Window Port Scan*

We expect all the attacks to be reported in as straightforward and clear a manner as possible, and alerts to be raised in a timely manner.

It is necessary to recognise that different devices detect and mitigate rate-based attacks in different ways. For example, where SYN proxies are utilised, a flood attack could be mitigated instantly with no SYNs reaching the victim, whereas if thresholds are used, some attack packets will inevitably reach the victim before the attack can be mitigated.

Thus, our criteria for determining whether or not an attack has been **successfully** mitigated is as follows:

- The victim remains alive and responsive (i.e. returning Web requests in a timely manner) throughout the attack
- It is possible to make valid requests to the victim from external hosts **and** (in certain circumstances) from the *apparent* attacking host, and receive responses in a timely manner
- The attack is detected and mitigated within a reasonable time frame (i.e. it is not allowed to have a detrimental effect on the victim before it is mitigated)
- Once the attack has been detected, no further attack traffic from the attacking host is allowed through for the duration of the test.

Test 1.2 - High Volume Attack Detection/Mitigation

Whereas the previous tests determine the device's ability to detect and mitigate a wide range of attacks under normal conditions (using the live attack tools), the level of flooding is generally fairly low.

This test generates a subset of the previous attacks at very high volumes (up to 80 per cent of the rated bandwidth) to determine if the rate of attack has any effect on the device's ability to detect and mitigate it.

The following attacks are repeated at various levels (10%, 20%, 40% and 80% of the rated bandwidth of the device under test) and we test for successful mitigation (as defined in Test 1.1) in each case.

- ***Test 1.2.1 -*** *SYN Flood (DOS from single source IP)*
- ***Test 1.2.2 -*** *SYN Flood (DDOS from multiple source IPs)*
- ***Test 1.2.3 -*** *Smurf*
- ***Test 1.2.4 -*** *Teardrop*
- ***Test 1.3.5 -*** *ICMP Flood*
- ***Test 1.2.6 -*** *UDP Flood*

Test 1.3 - 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 in-line devices.

Throughout the test we load the network with a wide range of "normal" network traffic. We note how many - if any - false alarms are raised on this traffic once the device has been tuned/configured, and comment an what action is necessary to reduce or eliminate such false positive scenarios.

The product attains a "PASS" for this section if it does **not** raise an alert and does **not** block any normal traffic once the initial tuning/learning process has been completed. Raising an alert on any normal traffic once the device has been completely configured is considered a "FAIL", which would indicate the chance that the sensor could block legitimate traffic inadvertently.

It is important to note that it is impossible to state definitively whether or not a particular device is susceptible to false positives, since this depends almost entirely on the type of traffic seen in the live deployment.

- ***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 mitigating basic attacks 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.

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

Test 2.2 - Fragmentation and Timing

The SYN Stealth Port Scan is repeated, subjecting the already small packets to fragmentation and delaying the amount of time between packets in order to evade detection:

- ***Test 2.2.1 -*** *Fragmented UDP Flood (Teardrop)*
- ***Test 2.2.2 -*** *Fragmented Stealth Port Scan*
- ***Test 2.2.3 -*** *Slow Stealth Port Scan (0.4 secs between packets)*
- ***Test 2.2.4 -*** *Very Slow Stealth Port Scan (15 secs between packets)*
- ***Test 2.2.5 -*** *Slow Connection Flood (1 second between packets)*
- ***Test 2.2.6 -*** *Very Slow Connection Flood (3 seconds between packets)*

For each of the evasion techniques, we note if (i) the attempted attack is blocked successfully (the primary aim of any attack mitigation device), and (ii) if the device is capable of providing 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 Web vulnerability scans are repeated (where applicable), 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 (the primary aim of any attack mitigation device), and (ii) if the device is capable of providing 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 – Attack Mitigation Performance Under Load

The aim of this section is to verify that the sensor is capable of detecting and blocking attacks when subjected to increasing loads of background traffic up to the maximum bandwidth supported as claimed by the vendor. In other words, we expect the device to be able to successfully mitigate attacks even when heavily loaded with normal traffic.

Sensors are deployed with **all** available detection modes enabled. Each sensor is tuned or configured to learn automatically to handle the levels of traffic involved. Our “attacker” hosts launch a number of attacks at target hosts 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 level of attack and/or normal traffic seen on the wire as verification.

Having ensured that the sensor is capable of detecting our baseline attacks, increasing levels of varying types of background traffic are generated **through** the sensor.

All tests are repeated with 250Mbps, 500Mbps, 750Mbps and 1Gbps of background traffic (or up to the maximum rated throughput of the device should this be less than 1Gbps). At each level, we launch a selection of attacks from the external network and check to ensure that they are successfully detected and mitigated. We also check to ensure that the victim servers remain alive and responsive to legitimate requests from the external network throughout the tests.

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 attack packets which actually made it through.

For each type of background traffic, we also determine the maximum load the sensor can sustain before it begins to drop packets/miss alerts.

Test 3.1 - 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 transfer potential attack packets to the detection engine. It is important that the device under test is tuned accurately to ensure that this traffic is not detected as a flooding attack and is thus mitigated in error.

The range of packet sizes has been selected to mirror the maximum, minimum and average packet sizes used in our HTTP stress tests.

At each level, we launch a selection of attacks from the external network and check to ensure that they are successfully detected and mitigated. We also check to ensure that the victim servers remain alive and responsive to legitimate requests from the external network throughout the tests. Both of these checks should yield a successful result in order for the device to attain a PASS at the given load level.

- ***Test 3.1.1 - 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 3.1.2 - 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 3.1.3 - 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 achieving 100 per cent PASS 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 3.2 - HTTP "Maximum Stress" Traffic With No Transaction Delays

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 detection engine and determine how the sensor copes with detecting and mitigating attacks 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.

At each level, we launch a selection of attacks from the external network and check to ensure that they are successfully detected and mitigated. We also check to ensure that the victim servers remain alive and responsive to legitimate requests from the external network throughout the tests. Both of these checks should yield a successful result in order for the device to attain a PASS at the given load level.

- ***Test 3.2.1 -*** *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 per cent PASS rates throughout this test.*
- ***Test 3.2.2 -*** *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 3.2.3 -*** *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 3.2.4 -*** *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 3.3 - 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 3.3.1 -*** *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 3.3.2 -** 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 3.4 - Protocol Mix Traffic

Whereas 3.2 and 3.3 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 3.4.1 -** 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 3.5 - “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 3.5.1 - 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 3.5.2 - Protocol Mix (72% HTTP traffic (simulated browsing sessions as 3.5.1)) + 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 3.2 - 3.4 should be examined.

Test 3.6 - Maximum Open Connections

It is important that the device under test cannot only support a sufficiently high rate of connection set-up and tear-down, but that is can also support a sufficiently large number of simultaneous connections. This test determines its maximum capacity.

- ***Test 3.6.1 -** In this test the Spirent Avalanche is set to continually open HTTP connections with the Reflector. The Reflector is set to implement a delay on each transaction, thus ensuring that transactions remain open for a significant period of time (typical of a real world situation).*

  The Avalanche is monitored and the point where transactions begin to fail is noted. Once the exact point of failure has been determined, the test is run for several hours with that number of transactions held open throughout.

Section 4 – 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 varying levels of normal HTTP traffic and varying levels of attack traffic on the average latency and user response times.

Test 4.1 - 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 4.1.1 - 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). Note that the* ***only*** *traffic passing through the device during this test is the UDP traffic used by SmartFlow to measure latency.*
- ***Test 4.1.2 - Latency With Background Traffic Load:*** *The Avalanche and Reflector are configured to generate varying loads of background HTTP traffic through the sensor (from 25 to 100 per cent of the maximum rated bandwidth of the device under test - maximum 1Gbps - 5,000 new connections per second - average packet size 540 bytes - 225,000 packets per second). A very small bi-directional load of SmartFlow traffic at various packet sizes (256 bytes, 540 bytes and 1000 bytes) is also passed across the infrastructure switches and through the device and the packet loss and average latency are recorded at each HTTP load level.*
- ***Test 4.1.3 - Latency When Under Attack:*** *The Spirent WebSuite software is used to generate varying loads of SYN flood traffic (from a single source IP) through the sensor (from 20 to 80 per cent of the maximum rated bandwidth of the device under test - maximum 800Mbps - 1,184,000 packets per second). A very small bi-directional load of SmartFlow traffic at various packet sizes (256 bytes, 540 bytes and 1000 bytes) is also passed across the infrastructure switches and through the device and the packet loss and average latency are recorded at each attack load level. The device should be configured to detect/mitigate the attack by the most efficient method available.*

Test 4.2 - 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 4.2.1 - 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 4.2.2 - Web Response When Under Attack (10% load):*** *The Avalanche and Reflector are configured to generate HTTP traffic through the sensor as for Test 4.2.1. The Spirent WebSuite software is then used to generate SYN flood traffic (from a single source IP) through the sensor at a rate of 10 per cent of the maximum bandwidth of the device under test (maximum 100Mbps - 148,000 packets per second). Note that with the background traffic, this rest will result in a maximum load of 70 per cent of the rated bandwidth of the device under test. 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.*
- ***Test 4.2.3 - Web Response When Under Attack (20% load):*** *As for Test 4.2.2, but with a 20 per cent load of SYN flood traffic (single source IP, maximum 200Mbps - 296,000 packets per second). Note that with the background traffic, this rest will result in a maximum load of 70 per cent of the rated bandwidth of the device under test. 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.*
- ***Test 4.2.4 - Web Response When Under Attack (40% load):*** *As for Test 4.2.2, but with a 40 per cent load of SYN flood traffic (single source IP, maximum 400Mbps - 592,000 packets per second). Note that with the background traffic, this rest will result in a maximum load of 90 per cent of the rated bandwidth of the device under test. 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 5 – 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 device, where failure can produce network outages.

- ***Test 5.1.1 - Blocking Under Extended Attack:*** *For this test, we expose the external interface of the device to a constant stream of genuine traffic interspersed with occasional attack traffic for a total of 8 hours. 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 mitigate 100 per cent of attacks, raising an alert for each. Results are presented as a percentage of attacks mitigated out of the total generated through the device. If any recognisable attacks are not mitigated for any reason - caused by either the volume of traffic ***or*** *the sensor failing open - this will result in a FAIL.*

- ***Test 5.1.2 - Passing Legitimate Traffic Under Extended Attack:*** *This test is identical to 5.1.1, where we expose the external interface of the device to a constant stream of genuine traffic over an extended period of time, interspersed with occasional attack traffic. 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 percentage of legitimate traffic passed out of the total generated through the device. If an excessive level (>0.09%) of 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 5.1.3 - Resistance to ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC traffic:*** *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 a*

  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.

- ***Test 5.1.4 - Mitigation of ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC mitigation traffic:*** *As for Test 5.1.3. Results are presented as a simple PASS/FAIL - the device will receive a PASS should it prove capable of total mitigation of all the ISIC traffic.*

Section 6 – 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 6.1 - Management Port

Clearly the ability to manage the data collected by the sensor is a critical part of any attack mitigation 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 monitoring and tuning the mitigation device during an attack.

- ***Test 6.1.1 - Open ports:*** *We will scan the open ports and active services on the management interface and report on known vulnerabilities.*

- ***Test 6.1.2 - Resistance to ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC traffic:*** *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 6.1.3 - Detection of ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC traffic:*** *As for Test 6.1.2. We note whether the ISIC attacks themselves are detected by the sensor even though targeted at the management port.*

V-Secure V-100 V7.0 Test Results

Section 1 - Detection Engine

Test 1.1 – Attack Detection/Mitigation	Detected?	Mitigated?
Test 1.1.1 - SYN Flood	YES	YES
Test 1.1.2 - TCP SYN Attack (low-rate SYN flood)	YES	YES
Test 1.1.3 - ICMP Flood	YES	YES
Test 1.1.4 - Distributed Denial Of Service (DDOS) Attack	YES	YES
Test 1.1.5 - UDP Flood	YES	YES
Test 1.1.6 - IGMP Flood	YES	YES
Test 1.1.7 - Connection Flood (fast)	YES	YES
Test 1.1.8 - Connection Flood (slow)	YES	YES
Test 1.1.9 - Random protocol violations (invalid packets)	YES	YES
Test 1.1.10 - Trojan response (external host receives responses from Trojan)	YES[1]	NO[1]
Test 1.1.11 - ICMP Sweep (inbound)	YES	YES
Test 1.1.12 - ICMP Sweep (outbound)	YES	YES
Test 1.1.13 - SQL Slammer	YES	YES
Test 1.1.14 - Spoofed IP Attack	YES	YES
Test 1.1.15 - Web Vulnerability Scan	YES	YES
Test 1.1.16 - TCP Port Scan (full connect)	YES	YES
Test 1.1.17 - Stealth Port Scan	YES	YES
Test 1.1.18 - FIN Port Scan	YES	YES
Test 1.1.19 - UDP Port Scan	YES	YES
Test 1.1.20 - NULL Port Scan	YES	YES
Test 1.1.21 - Xmas Port Scan	YES	YES
Test 1.1.22 - IP Protocol Port Scan	YES	YES
Test 1.1.23 - ACK Port Scan	YES	YES
Test 1.1.24 - Window Port Scan	YES	YES
Total	**24 / 24**	**23 / 24**

Test 1.2 – High Volume Attack Detection/Mitigation		10Mbps	20Mbps	40Mbps	80Mbps
Test 1.2.1 - SYN Flood DOS (single source IP)	Mitigated	PASS	PASS	PASS	PASS
	Pass legit	PASS	PASS	PASS	PASS
Test 1.2.2 - SYN Flood DDOS (multiple source IPs)	Mitigated	PASS	PASS	PASS	PASS
	Pass legit	PASS	PASS	PASS	PASS
Test 1.2.3 - Smurf	Mitigated	PASS	PASS	PASS	PASS
	Pass legit	PASS	PASS	PASS	PASS
Test 1.2.4 - Teardrop	Mitigated	PASS	PASS	PASS	PASS
	Pass legit	PASS	PASS	PASS	PASS
Test 1.2.5 - ICMP Flood	Mitigated	PASS	PASS	PASS	PASS
	Pass legit	PASS	PASS	PASS	PASS
Test 1.2.6 - UDP Flood	Mitigated	PASS	PASS	PASS	PASS
	Pass legit	PASS	PASS	PASS	PASS

Test 1.3 – Resistance to False Positives	Pass/Fail
Test 1.3.1 - False positives	PASS

Section 2 - Evasion Techniques

Test 2.1 – Evasion Baselines	Detected?	Mitigated?
Test 2.1.1 - UDP Flood	YES	YES
Test 2.1.2 - TCP Port Scan	YES	YES
Test 2.1.3 - Stealth Port Scan	YES	YES
Test 2.1.4 - Connection Flood	YES	YES
Total	**4 / 4**	**4 / 4**

Test 2.2 – Fragmentation and Timing	Detected?	Mitigated?
Test 2.2.1 - Fragmented UDP Flood (Teardrop)	YES	YES
Test 2.2.2 - Fragmented Stealth Port Scan	YES	YES[2]
Test 2.2.3 - Slow Stealth Port Scan (0.4 secs between packets)	YES	YES
Test 2.2.4 - Very Slow Stealth Port Scan (15 secs between packets)	YES	YES
Test 2.2.5 - Slow Connection Flood (1 second between packets)	YES	YES
Test 2.2.6 - Slow Connection Flood (3 seconds between packets)	YES	YES
Total	**6 / 6**	**6 / 6**

Test 2.3 – URL Obfuscation	Detected?	Mitigated?
Test 2.3.1 - URL encoding	YES	YES
Test 2.3.2 - /./ directory insertion	YES	YES
Test 2.3.3 - Premature URL ending	YES	YES
Test 2.3.4 - Long URL	YES	YES
Test 2.3.5 - Fake parameter	YES	YES
Test 2.3.6 - TAB separation	YES	YES
Test 2.3.7 - Case sensitivity	YES	YES
Test 2.3.8 - Windows \ delimiter	YES	YES
Test 2.3.9 - Session splicing	NO	NO
Total	**8 / 9**	**8 / 9**

Section 3 - Detection/Mitigation Performance Under Load

Test 3.1 – UDP traffic to random valid ports		25Mbps	50Mbps	75Mbps	100Mbps	Max
Test 3.1.1 - 256 byte packet test - max 45,300pps (100Mbps)	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	
Test 3.1.2 - 550 byte packet test - max 22,000pps (100Mbps)	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	
Test 3.1.3 - 1514 byte packet test - max 12,200pps (100Mbps)	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	

Test 3.2 – HTTP "maximum stress" traffic with no transaction delays		25Mbps	50Mbps	75Mbps	100Mbps	Max
Test 3.2.1 - Max 250 connections per second - ave packet size 1000 bytes - max 12,000 packets per second	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	
Test 3.2.2 - Max 500 connections per second - ave packet size 540 bytes - max 22,500 packets per second	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	
Test 3.2.3 - Max 1000 connections per second - ave packet size 440 bytes - max 27,500 packets per second	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	
Test 3.2.4 - Max 2000 connections per second - ave packet size 360 bytes - max 32,000 packets per second	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	

Test 3.3 – HTTP "maximum stress" traffic with transaction delays		25Mbps	50Mbps	75Mbps	100Mbps	Max
Test 3.3.1 - Max 500 connections per second - ave packet size 540 bytes - max 22,500 packets per second - 10 sec delay - max 5,000 open connections	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	
Test 3.3.2 - Max 1000 connections per second - ave packet size 440 bytes - max 27,500 packets per second - 10 sec delay - max 10,000 open connections	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	

Test 3.4 – Protocol mix		25Mbps	50Mbps	75Mbps	100Mbps	Max
Test 3.4.1 - 72% HTTP (540 byte packets) + 20% FTP + 6% UDP (256 byte packets). Max 400 connections per second - ave packet size 540 bytes - max 21,500 packets per second - max 75 open connections	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	

Test 3.5 – Real World traffic		25Mbps	50Mbps	75Mbps	100Mbps	Max
Test 3.5.1 - Pure HTTP (simulated browsing session on NSS Web site). Max 470 connections per second - 2 new users per second - ave packet size 560 bytes - max 21,000 packets per second	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	
Test 3.5.2 - Protocol mix - 72% HTTP (simulated browsing sessions as 2.5.1) + 20% FTP + 6% UDP (256 byte packets). Max 370 connections per second - ave packet size 560 bytes - max 20,500 packets per second - max 150 open connections	Mitigated	PASS	PASS	PASS	PASS	100Mbps
	Pass legit	PASS	PASS	PASS	PASS	

Test 3.6 - Stateful Operation	Result
Test 3.6.1 - Maximum simultaneous open TCP connections	115,000

Section 4 - Latency & User Response Times

Test 4.1 – Latency	Packet Size	25% load	50% load	75% load	100% load
Test 4.1.1 Average latency (µs) with no background traffic (max load 100Mbps pure UDP latency measurement traffic)	256	51.81	59.72	69.34	84.22
	550	56.73	59.13	62.74	66.01
	1000	68.58	69.49	70.02	71.51
Test 4.1.2 Average latency (µs) with pure legitimate HTTP traffic (max load 100Mbps - 250 connections per second - ave packet size 540 bytes - 22,500 packets per second)	256	56.57	63.60	70.36	80.73
	550	67.40	75.66	80.96	89.07
	1000	83.18	87.47	93.68	95.34
Test 4.1.3 Average latency (µs) when under pure SYN Flood DOS attack (max load 80Mbps SYN packets - 118,400pps from one source IP)	256	71.77	99.61	184.87	N/A[3]
	550	98.40	124.06	237.85	N/A[3]
	1000	100.47	141.73	250.21	N/A[3]

Test 4.2 – User Response Times	Attempted Trans	Failed Trans	Min Page Response	Max Page Response	Ave Page Response
Test 4.2.1 - Web page response (ms) with pure HTTP background traffic (50Mbps HTTP traffic - max 250 connections per sec - ave packet size 540 bytes - max 11,250 packets per sec)	157058	0	200	206	200
Test 4.2.2 - Web page response (ms) when under **10%** attack load (50Mbps HTTP traffic - max 250 connections per sec - ave packet size 540 bytes - max 11,250 packets per sec **PLUS** 10Mbps SYN flood DOS (14,800pps from one IP)	157054	0	200	208	200
Test 4.2.3 - Web page response (ms) when under **20%** attack load (50Mbps HTTP traffic, max 250 connections per sec - ave packet size 540 bytes - max 11,250 packets per sec **PLUS** 20Mbps SYN flood DOS (29,600pps from one IP)	157053	0	200	286	201
Test 4.2.4 - Web page response (ms) when under **40%** attack load (50Mbps HTTP traffic, max 250 connections per sec - ave packet size 540 bytes - max 11,250 packets per sec **PLUS** 40Mbps SYN flood DOS (59,200pps from one IP)	157055	0	200	220	201

Section 5 - Stability & Reliability

Test ID	Result
Test 5.1.1 - Mitigation under extended attack	100%
Test 5.1.2 - Passing legitimate traffic under extended attack	99.992%
Test 5.1.3 - Resistance to ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC traffic	PASS
Test 5.1.4 - Mitigation of ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC traffic	FAIL

Section 7 - Management Interface

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

Notes:

1. Detected as “Outbound UDP Flood” but could not be mitigated successfully.
2. Some fragmented packets were detected at the “victim” host - however, the attacker **was** prevented from completing the port scan.
3. At this load there were some dropped packets, rendering latency figures inaccurate

Section 1: Detection Engine

We installed one sensor and enabled all the protection mechanisms apart from those which were not applicable in our test environment (*IANA protection, Outgoing NAT IP Scan,* and *Access Lists Policy*). We then configured the expected bandwidth settings and ran through all of our load generation tests, allowing the V-100 to learn from our "normal" traffic.

It was necessary to *Approve* entries for the *Protected Hosts* (such as the Spirent Reflector Web servers and our "victim" machines) and configure the required protection mechanisms for those hosts (once again, we configured all available protection mechanisms). It was also necessary to create *Trusted Group* entries for those hosts used in the Spirent SmartFlow latency tests, to ensure they did not trigger UDP Flood alerts. Once that had been done, the system required no further configuration throughout the tests.

Attack detection/mitigation was excellent, with the V-100 detecting and successfully mitigating all but one of our attacks.

Our Back Orifice controller was able to contact the Back Orifice Server installed on the protected network, receiving a large number of packets in response. These outbound packets were detected as an *Outbound UDP Flood*, but the V-100 was unable to mitigate it successfully (we believe this may be a bug – V-Secure is looking into it at the time of writing). The Back Orifice service was also detected as an *unknown service* and an alert raised. By using the *Access List Policy*, however, it was possible to define the ports/hosts being used by the Back Orifice programs as "off limits" and prevent access that way.

Performance in the high volume detection/mitigation test was impeccable across the board, with perfect detection and mitigation at all load levels.

A major concern in deploying an in-line device is the blocking of legitimate traffic. Once we had configured the *Trusted Groups* and *Approved* the *Protected Hosts* the V-100 ran through every single one of our tests without raising a single false positive alert.

Section 2: Evasion Techniques

Resistance to our evasion attempts proved very good, with the V-100 successfully detecting all of the fragmented and slow attacks we ran. It would appear to be very difficult - perhaps impossible - to evade this device by simply slowing down port scans and connection floods thanks to the fuzzy logic mechanism employed to compare "*normal*" vs. "*abnormal*" traffic.

When employing URL obfuscation techniques to attempt Web vulnerability scans, the V-100 successfully detected and blocked all except the Whisker splice. V-Secure is working on this at the time of writing.

Section 3: Detection/Mitigation Performance Under Load

Note that the V-Secure V-100 was tested as a 100Mbps device.
Performance at all levels of our load tests was impeccable, with 100 per cent of all attacks being detected and mitigated under all load conditions. We would happily confirm V-Secure's 100Mbps rating for this device.

Out of the box, the V-100 handled 115,000 simultaneously open connections without tuning - this is not configurable.

Section 4: Latency & User Response Times

Basic latency figures were excellent at all traffic loads and with all packet sizes, ranging from 52µs with 25Mbps of 256 byte packets, to 72µs with 100Mbps of 1000 byte packets. Behaviour throughout the tests with no background traffic was very even and predictable, remaining at well under 100µs under all load conditions and with all packet sizes.

Placing the device under an increasing load of HTTP traffic (ranging from 25Mbps to 100Mbps) also had minimal effect on the latency figures, which ranged from 57µs with 25Mbps of 256 byte packets, to 95µs with 100Mbps of 1000 byte packets.

All of these figures are well inside the limits we would expect to see for a 100Mbps device, meaning the V-100 could be situated anywhere on a 100Mbps network, either internally or at the perimeter.

Naturally, performance when under attack is critical for an attack mitigation device, and the V-100 did not disappoint overall. Mitigation was handled well at almost all levels, but the V-100 does have a limitation on handling packets per second of around 70,000pps, which we feel is a little low for a 100Mbps device. Despite the fact it can actually handle higher rates than this when mitigating attacks, we did note that it dropped packets at the 80Mbps (118,000pps) level of SYN flood attacks which caused latencies to become erratic (though not too excessive) at this level.

However, we feel that lower levels of attacks would be usual on most networks, and here the V-100 did very well, with latencies ranging from just 72µs at 20Mbps of SYN flood traffic to 250µs at 80Mbps.

Section 5: Stability & Reliability

The V-100 performed consistently and completely reliably throughout our tests. Under eight hours of extended attack (comprising millions of sessions of legitimate traffic interspersed with some attacks) it continued to mitigate 100 per cent of attack traffic, whilst passing 99.992 per cent of legitimate traffic (blocking an average of 80 out of 1,000,000 legitimate sessions per run).

Exposing the sensor interface to ISIC-generated traffic had no adverse effect, and the device continued to detect and block all other exploits throughout and following the ISIC attack. Although it mitigated the ISIC attack partially, some attack traffic did reach the intended victim.

Section 6: Management Interface

No open ports are visible to port scanners on the management interface.

The extended ISIC attack against the management interface had no effect on the appliance’s ability to detect and block attacks on the main interfaces. No alerts were raised during the attack. Console communication was interrupted during the attack, however, meaning we were unable to monitor the attack or reconfigure the device. Response returned to normal once the attack finished, and there were no residual stability problems.