

NetScreen-IDP 500 V3.0

Technical Evaluation

An NSS Group Report



First published January 2004 (Version 1.0)
Revised March 2004 (Version 1.1)

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

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

©1991-2004 The NSS Group

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

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

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

The NSS Group Limited is registered in England & Wales, Reg No. 3233843
Registered Office: Montagu House, 81 High Street, Huntingdon, Cambs, PE29 3NY, England
Tel +44 (0)7005 802 953

TABLE OF CONTENTS

INTRODUCTION	1
Intrusion Prevention Systems (IPS)	1
Host IPS (HIPS).....	2
Network IPS (NIPS).....	2
Implementation Challenges	3
Requirements for effective prevention.....	4
The NSS Intrusion Prevention Group Test.....	6
Performance	6
Security Effectiveness	9
Usability	11
NETSCREEN-IDP 500 V3.0.....	12
Executive Summary.....	12
Architecture.....	12
IDP Sensor	13
Detection Engine	14
High Availability	15
IDP Management Server.....	16
User Interface (UI).....	16
Performance	16
Security Effectiveness	18
Usability	19
Installation.....	19
Configuration	20
Policy Management.....	22
Alert Handling	28
Reporting and Analysis.....	31
Verdict.....	34
Contact Details	36
APPENDIX A – TEST RESULTS.....	37
The Test Environment	37
Section 1 – Detection Engine	37
Section 2 – IPS Evasion	39
Section 3 – Stateful Operation.....	41
Section 4 – Detection/Blocking Performance Under Load	42
Section 5 – Latency & User Response Times.....	47
Section 6 – Stability & Reliability	48
Section 7 – Management and Configuration	49
NetScreen-IDP 500 V3.0 Test Results	50
Section 1 - Detection Engine	50
Section 2 - IPS Evasion	50
Section 3 - Stateful Operation	51
Section 4 - Detection/Blocking Performance Under Load.....	52
Section 5 - Latency & User Response Times	53
Section 6 - Stability & Reliability	53
Section 7 - Management Interface	53

TABLE OF FIGURES

Figure 1 - NetScreen-IDP: Architecture	12
Figure 2 - NetScreen-IDP: Sensor rear panel	13
Figure 3 - NetScreen-IDP: The UI	19
Figure 4 - NetScreen-IDP: The Dashboard display	20
Figure 5 - NetScreen-IDP: Attack Objects	21
Figure 6 - NetScreen-IDP: Configuring dynamic groups	22
Figure 7 - NetScreen-IDP: Policy Editor	23
Figure 8 - NetScreen-IDP: Defining custom signatures	24
Figure 9 - NetScreen-IDP: Notification options	26
Figure 10 - NetScreen-IDP: Sensor settings	27
Figure 11 - NetScreen-IDP: Log Viewer	28
Figure 12 - NetScreen-IDP: Tuning the Policy via the Log Viewer	29
Figure 13 - NetScreen-IDP: Quick Reports from the Log Viewer	30
Figure 14 - NetScreen-IDP: Top 10 Attacks report	31
Figure 15 - NetScreen-IDP: Log Investigator	32
Figure 16 - NetScreen-IDP: Enterprise Security Profiler (ESP)	33
Figure 17 - NetScreen-IDP: Signature update process	36

The NSS Group

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

Based in the UK with separate 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 throughout Europe and the United States.

The Group consists of two wholly-owned subsidiaries :

- *NSS Network Testing Laboratories*
- *Network Security Services*

NSS Network Testing Laboratories are available to vendors and end-users for fully independent testing of networking, communications and security hardware and software.

NSS Network Testing Laboratories also operates certification schemes for vendors and certification bodies, and currently provides certification of firewalls, VPN's, crypto products 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>

Network Security Services provides a range of security-related services to vendors and end-users including security policy definition, IDS, firewall and VPN implementation, network security auditing and analysis, and penetration testing.



Foreword

The NSS Group is pleased to present the results of the first comprehensive *Intrusion Prevention System (IPS)* test of its kind.

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.

Contrary to recent analyst claims, we do not believe that “IDS is dead” or that “IPS is stillborn”. So-called “*deep inspection firewalls*” **may** be where the industry is heading in the long term, but they are simply not ready for prime-time deployments at this point in time. Until they are, 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.

Please note that we fully recognise that each of the above product-types could be considered as “intrusion prevention” systems in some respect, along with Anti Virus gateways, desktop firewalls, and other products designed to “prevent” malicious activity on a network or individual host. However, we also believe that the marketing terms for each of these products are well established, and that the grounds for creating a new market segment - referred to as *Intrusion Prevention Systems (IPS)* - specifically for those products which are evaluated as part of this report is valid. We have defined what **we** consider to be IPS (both host- and network-based) in the introductory text to this report.

You might not like the marketing hype, but the time for quibbling over the terminology is over - it is now time to get down to the serious issue of evaluating the technology behind it.

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

Overall, the suite contains over **750 individual tests**, many of which are run multiple times, to provide the most thorough and complete evaluation of IPS products available anywhere today.

We believe that our IPS test methodology will become the *de facto* standard for testing in-line intrusion prevention 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.

Bob Walder

INTRODUCTION

In a recent 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*.

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 a malicious packet, in addition to raising an alert, it will discard the packet 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.

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 has conducted the first comprehensive IPS test of its kind. 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 test methodology (see section on *Testing Methodology* later in this report for detailed methodology) 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 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 750 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) 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 1.48 million packets per second and connection rates up to 20,000 connections per second.

Latency & User Response Times

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

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

To be specific, the maximum throughput achievable for any given TCP connection in a zero loss network is expressed as:

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

where *window* is the maximum TCP window size (64 Kbytes by default) and RTT is the round trip time in the network. This equation tells us that the throughput of a TCP connection is inversely proportional to network latency (note that this is TCP throughput for *one* connection - the aggregate bandwidth is not affected by latency). In other words, if you double latency, you halve throughput.

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

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

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

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

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

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

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

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

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

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

Bi-directional network latency of 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.

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:

- **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*.

NETSCREEN-IDP 500 V3.0

Executive Summary

The NetScreen Intrusion Detection and Prevention (IDP) system is a turnkey appliance-based system which uses as many as eight detection methods to detect malicious network traffic.

The IDP-500 is capable of operating in both in-line mode (as an *Intrusion Prevention System*) or as a passive *Intrusion Detection System* attached to a SPAN or mirror port on a switch.

The IDP-500 is designed for 500Mbps networks and can handle its maximum rated bandwidth under most normal traffic conditions likely to be encountered on a sub-Gigabit network. Latency, too, is within acceptable limits under normal conditions, increasing significantly only when the device is under severe DOS/DDOS attack.

We found the IDP-500 to be to be very stable and reliable, coping with our extensive reliability tests with ease and without blocking any legitimate traffic or succumbing to common evasion techniques.

The management system has been well designed to handle management and configuration of large numbers sensors across the enterprise. Policy definition and deployment is flexible and powerful, and the alert handling and reporting/forensic capabilities are extensive.

Architecture

IDP uses a three-tier architecture that consists of the *IDP Sensor*, the *IDP Management Server*, and the *IDP User Interface*.

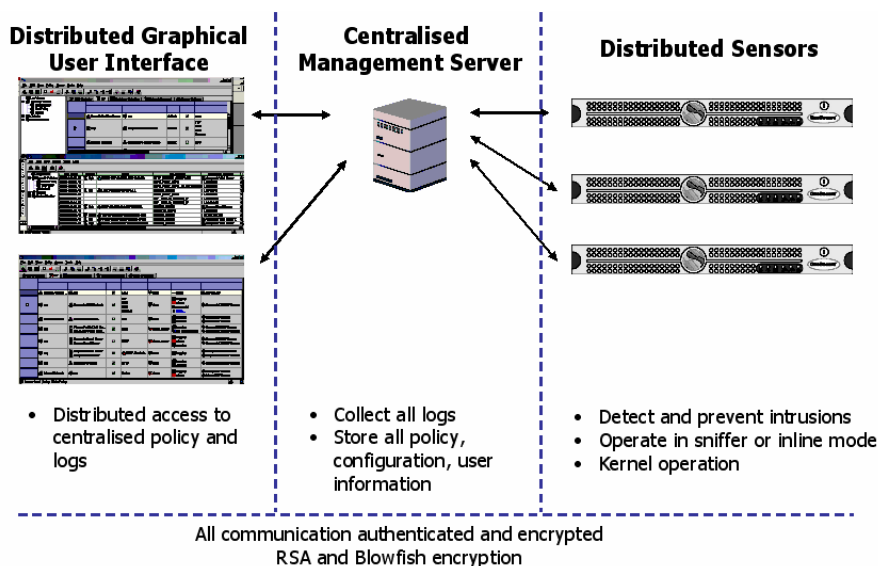


Figure 1 - NetScreen-IDP: Architecture

- **IDP Sensor** - monitors the network on which the IDP system is installed. The Sensor is a hardware appliance that runs the Linux-based IDP Sensor software.

- **IDP Management Server** - stores and manages all Attack Objects (including attack signatures and protocol anomalies), log information, rulebases, and Security Policies. Multiple Sensors can be managed by a single Management Server.
- **User Interface (UI)** - a graphical interface for interacting with the IDP system. The UI is used to access remotely and manipulate the information stored on the Management Server.

IDP Sensor

Based on the Dell PowerEdge 1750 1U rack-mount hardware platform with a single Xeon processor and 4GB RAM, the IDP 500 submitted for this test sports two fibre Gigabit interfaces for packet forwarding and two copper Gigabit interfaces for management and high availability. This device is rated at 500Mbps, but the range also includes the IDP-10 (20Mbps) and the IDP-100 (200Mbps). In the IDP-100 and IDP-500 models, the dual fibre ports can be replaced by eight 10/100Mbps ports allowing up to four 10/100Mbps segments to be inspected in-line at the same time.

The Sensor's primary task is to detect suspicious and anomalous network traffic based on specific rules defined in IDP rulebases. If the Sensor is running in-line, it can also take a predefined action against malicious traffic.

IDP Sensors can be deployed as *active gateways* or *passive sniffers*. A *passive sniffer* Sensor connects to a switch or hub in promiscuous mode and sniffs the network traffic as it passes by. The Sensor monitors network traffic, records security events, and can create alarms for attacks. However, because a sniffer Sensor cannot take direct action against the attack, it cannot prevent it.




Figure 2 - NetScreen-IDP: Sensor rear panel

An *active gateway* Sensor sits between the network and a firewall, or a network and a DMZ, and takes an active role in protecting the network. When it detects intrusions or attacks defined by Security Policies, the Sensor can drop, block, or ignore the suspicious connection, or drop only the suspicious packets. Because active gateway Sensors are installed in-line they can take direct action against the attack, thus preventing it.

The Sensor can be deployed in router, bridge, or proxy ARP modes, providing a high degree of flexibility in deployment options. The IDP appliance supports the 802.1Q VLAN protocol at the kernel level and can generate 802.1Q frames in all in-line IDP deployment modes. It also allows the creation of multiple virtual routers on a single IDP Sensor. A new Layer 2 transparent install mode has also been added to this release. This allows IDP to be installed like a "bump in the wire" without any changes to networking architecture. Another benefit of this install mode is that IDP can now pass non-IP traffic. For many this is the simplest in-line deployment option, and this is how it was deployed in our tests.

Another new feature in this release is the ability to use a non-forwarding interface to send TCP reset packets. This is useful when IDP is attached to span ports of switches that do not allow incoming traffic through the span port.

The Sensor stores logs on the local hard drive initially, and then transmits the logs to the Management Server. If communication between the Sensor and Management Server fails, the Sensor attempts to restore communication and stores new logs on the local file system until available disk space is consumed. If no disk space remains and communication has not been restored, the Sensor first deletes sysinfo logs (debug and error messages) followed by packet captures to free up disk space. If communications cannot be restored before all local disk space is consumed, the Sensor ceases to detect new attacks.

Detection Engine

IDP operates as an in-line, active device that looks at all traffic and determines what constitutes an intrusion. If instructed by the Security Policy, IDP can deny traffic associated with the intrusion by dropping the connection or associated packets (when operating in-line). Rules are created in the IDP rulebases to specify what actions IDP takes when a particular attack is detected. Detection methods (rulebases) include:

- **Stateful Signature** - IDP includes several hundred signature Attack Objects that use stateful signatures to detect known attacks. The administrator can view, create, edit, or delete signature Attack Objects using the Signature Editor in the UI.
- **Protocol Anomaly** - Protocol anomalies are deviations from the standard protocol. Most legitimate traffic adheres to the published specifications (RFC) for protocols, but traffic that doesn't produces an anomaly. The IDP system uses the RFC to create protocol anomaly Attack Objects for deviations from the protocol's published specification. The latest release has added a further eight protocols, bringing the total number of decoded protocols to almost 60.
- **Backdoor** - Backdoor detection uses network traffic patterns and heuristics of packet transmissions to detect interactive traffic, a common sign of an attacker using a Trojan or backdoor. Unlike anti-virus software, which scans for **known** backdoor files or executables on the host system, IDP detects the interactive traffic that is produced when backdoors are used, but without requiring a specific signature. This method ensures that IDP can detect all backdoors, both known **and** unknown, even if the data is encrypted. Normal bi-directional services (such as FTP) are excluded from the default policies.
- **Traffic Anomaly** - A traffic anomaly is a pattern that indicates abnormal network activity such as port scans and other distributed network attacks. The IDP system counts the number of ports scanned in a specified time period and uses this traffic flow analysis to identify scans.
- **IP Spoofing** - IP spoofing occurs when a packet uses a fake IP address, and the IDP system detects IP spoofing by comparing the IP addresses of packets to the IP addresses of devices on the protected network. An IP address is considered spoofed if an incoming packet uses an IP address that belongs to a Network Object on the internal network, or if an outgoing packet uses an IP address that does not belong to a Network Object on the internal network

- **SYN Flood Detection** - IDP detects and prevents SYN floods by ensuring that the TCP handshake is performed correctly. In sniffer mode, the default SYN flood rule tells IDP to passively monitor the transfer of packets between the client host and the server using a timer to ensure that connections are established promptly. The timer IDP uses for the connection establishment is shorter than the timer the server uses for the connection queue. If the client host does not send an ACK packet to the server, as would be the case during a SYN flood attack, the IDP connection timer expires and the default SYN-Protector rule is triggered.
- **Layer 2 and Denial of Service Detection** - Attackers can manipulate Layer 2 protocols to perform ARP attacks (such as ARP cache poisoning) and other MAC attacks. The IDP system detects Layer 2 attacks by defining implied rules on the IDP Sensor. Implied rules include ARP table restrictions, fragment-handling, connection timeouts, byte/length thresholds for data packets, and other mechanisms.
- **Network Honeypot** - Whilst not strictly speaking a detection method, the IDP Network Honeypot impersonates open ports on existing servers in order to lure attackers into attempting exploits for the purpose of evidence gathering.

The latest V3.0 release allows administrators to use compound signatures, where multiple individual signatures (both stateful signatures and protocol anomalies) can be combined into a single compound signature to detect complex attacks in a single session. The language for creating custom signatures has also been improved to include Perl style regular expressions to make it easier to create signatures that can be used to detect an underlying vulnerability rather than a specific exploit based on that vulnerability. In an attempt to further increase the accuracy of attack detection, V3.0 has increased the number of stateful signature fields to over 500.

High Availability

When operating in-line, it is essential that traffic is not interrupted by device failures. It is therefore possible to deploy the IDP system in a *high availability* (HA) configuration to provide failure protection, either in a standalone configuration or using third-party hardware. In high availability, two to sixteen Sensors are joined together in a HA cluster that provides failure protection and/or load balancing. In *load balancing* mode, all Sensors in the cluster share network traffic equally, and if a Sensor fails, network traffic is redirected to the other Sensors in the cluster. With *hot standby*, a primary Sensor handles all network traffic while a secondary Sensor stands by. If the primary Sensor fails, network traffic is redirected to the secondary Sensor.

State information is shared between Sensors in the cluster using state-sync interfaces and dedicated network interface cards. High availability configurations are available for transparent, bridge, router, and proxy-ARP deployment modes.

Also available when operating in bridge mode (IDP 10 and IDP 100 only) is the *IDP Bypass Unit* - a fail-open network device for the IDP system. If traffic flow through the IDP appliance is disrupted, the Bypass Unit can automatically reroute traffic.

IDP Management Server

The NetScreen-IDP *Management Server* consists of software that manages system resources and data that drives IDP functionality. The Management Server software can be installed on the IDP appliance for evaluation purposes or for single Sensor deployments on small networks (IDP 100 only). In all other situations, the Management Server must be installed on a separate computer running Red Hat Linux 7.2 or Red Hat Enterprise Linux Version 3 or Solaris Versions 8/9 (this is the recommended configuration due to performance considerations).

The Management Server centralises the logging, reporting, data, and Security Policy management for the IDP system. Logs from multiple Sensors are consolidated into a single, central repository, and the Management Server also provides a single point of management, policy definition and deployment. All objects, Security Policies, and logs are stored in databases on the Management Server and are administered using the User Interface by one or more administrators.

Management Server communication with the other two tiers of the IDP system (the Sensor and User Interface) is encrypted and authenticated to provide an additional level of security.

The Management Server gathers log records for security events using a proprietary communication mechanism utilising a modified version of UDP which NetScreen claims is more reliable than ordinary UDP, requires less CPU and memory resources from servers, and reduces the number of acknowledgement packets on the network. Alerts are stored in a proprietary database which NetScreen also claims offers high levels of performance (over 50,000 logs per second).

When IDP detects a DoS attack, the Sensor increases its transmission rate and injects redundant packets to increase the odds of transmitting the logs to the central location.

User Interface (UI)

The User Interface (UI) is software that provides a graphical environment for centrally managing IDP. The UI is a java-based software application that can be installed on multiple computers across the network.

Although the UI supports multiple users, only one user at a time can take control of the Management Server - this eliminates concerns about synchronisation or data loss. Each administrator can configure the UI with his own preferences - IDP stores user preferences and custom Log Viewer views in the central database so that they remain consistent when accessed from different client machines.

Performance

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

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

It is worth noting that devices which demonstrate 100 per cent blocking but less than 100 per cent detection in these tests will be prone to blocking **legitimate** traffic under similar loads.

The NetScreen-IDP 500 was tested as a 500Mbps device and performed well in our performance tests, demonstrating that it could handle most traffic loads up to the rated speed of 500Mbps, successfully detecting and blocking 100 per cent of all our exploits.

Under certain extreme traffic loads - unlikely to be found in most "normal" real-world networks, especially sub-Gigabit networks for which this device has been designed - there was some packet loss or increase in latency leading to failed HTTP connections. However, even when the device was placed under such extreme loads and started to drop packets it continued to block 100 per cent of all exploits successfully.

Providing traffic loads are kept within the bounds indicated in these tests (i.e. in particular "normal" traffic with less than 7,500 HTTP connections per second and 450,000 packets per second), then we would have no hesitation in rating this device at 500Mbps.

Being a software-based IPS product running on an off-the-shelf hardware platform, latency is higher than you would expect to see in a pure networking device such as a switch. However, under most normal traffic conditions, latency is well within acceptable limits for a 500Mbps network.

Behaviour throughout the tests was very predictable, increasing as both average packet size and traffic load increased and ranging from 224µs with 125Mbps of 64 byte packets to 341µs with 500Mbps of 1514 byte packets. The only situation where latency increased significantly beyond the norm was when we loaded the device with more than 450,000 64-byte packets per second - not a normal occurrence in a sub-Gigabit network.

When the device was loaded with 250Mbps of background traffic (half the rated speed of the product), there was a slight increase in latency of 64 byte packets, and a slight reduction in latency of 440 byte and 1514 byte packets. This demonstrates that high loads of normal network traffic would not have a serious negative impact on overall latency figures.

With the SYN-Protector enabled in "Relay" mode, the degradation when under a 50Mbps SYN flood attack was much more significant. Although the SYN flood was effectively mitigated, the latency was increased to the point where we noted a significant number of failed connections in our HTTP traffic. The IDP-500 is better as a straight IPS than as a DOS/DDOS mitigation device, and NetScreen has better DOS/DDOS mitigation capabilities built into its firewall product line. To balance this, however, it should be noted that with the SYN-Protector disabled, **all** of our HTTP connections failed when under DOS attack, proving the effectiveness of the SYN-Protector despite the latency issues.

The NetScreen IDP-500 performed consistently and reliably throughout our tests. Under extended attack it continued to pass legitimate traffic whilst blocking attack traffic in a consistent manner.

Exposing both the sensor and management interfaces to an extended run of ISIC-generated traffic had no effect on either its performance or stability and the device continued to detect and block all other exploits throughout and following the ISIC attacks.

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

Security Effectiveness

We installed one NetScreen-IDP 500 sensor with the latest signature set. The Policy was created from scratch, switching on all Attack Objects from *Critical* down to *Low* (omitting only *Informational* signatures, which are more suited to auditing purposes), as well as turning on *Backdoor Protection* and the *SYN-Protector*. We also enabled fragmented packet detection and flow error detection in order to identify certain DOS and ICMP attacks.

Signature recognition (with blocking disabled) was excellent out of the box at 96 per cent and the blocking performance was identical. The new compound signature capability in particular has enabled far wider coverage than was possible with previous versions. This figure was increased to 100 per cent following a signature pack update which NetScreen provided within 48 hours.

Most of our “false negative” (modified exploit) cases were detected correctly out of the box, demonstrating that the NetScreen signatures are generally designed to detect the underlying vulnerability rather than a specific exploit. Once again, following the signature pack update, the IDP-500 detected 100 per cent of these test cases. Resistance to false positives was also generally very good, and a clean sweep was achieved in our tests following a signature update.

The NetScreen IDP-500 performed impeccably in all of our evasion tests, being one of the few products to collect a clean sheet across the board in this area. *Fragroute*, *Whisker*, *ADMmutate* and even *RPC record fragging* all failed to trick the IDP-500 into ignoring valid attacks.

Note that not only were the fragmented and obfuscated attacks blocked successfully, but almost every one of them - a single *fragroute* attack being the sole exception - was decoded accurately as well.

The IDP-500 was not tricked into alerting on our stateless exploits, indicating that it would be resistant to TCP-based exploits launched via replay tools such as *Stick* and *Snot*. Blocking of mid-flow streams, whilst disabled out of the box, is possible by enabling a single signature if required.

Out of the box, the device maintained state on up to 100,000 open connections, successfully detecting our half-open exploit as it was completed. It also continued to detect and block new exploits as we maintained 100,000 open connections, and no legitimate traffic was blocked during this stage of these tests.

With a simple change to the flow-related parameters in the GUI, it was possible to maintain state across 500,000 open connections (though NetScreen actually recommends 220,000 as the maximum on this device).

If the configured maximum number of sessions is exceeded, the sensor will refuse new connections, meaning that legitimate traffic is blocked and new attacks are not alerted. However, state is maintained completely on existing flows and no attack traffic is ever allowed through the device.

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 is very straightforward. The Management Server can be installed on any *secure and trusted* Red Hat Linux 7.2, Enterprise Version 3, or Solaris 7/8 computer. It is installed via a simple shell script that walks the administrator through all the steps necessary to install the software, and this process takes no more than five minutes.

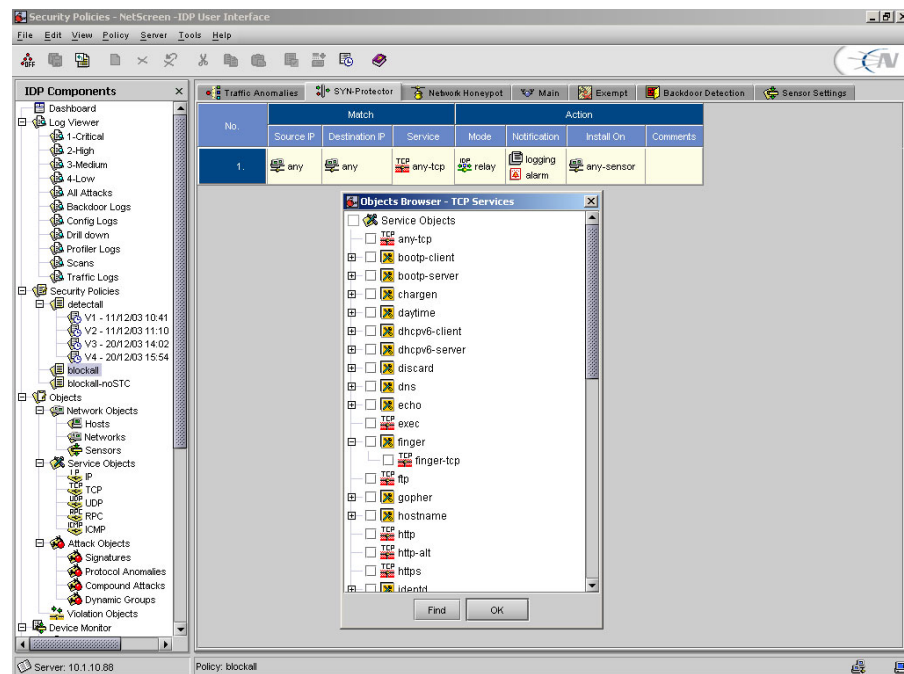


Figure 3 - NetScreen-IDP: The UI

Next, the Sensor is connected to the network to be protected, and one of the interfaces connected to the management network. Using the *Appliance Configuration Manager (ACM)*, a Web-based tool, a Wizard guides the administrator through the six-section configuration process of configuring the Sensor software. During this process, key parameters such as Sensor name, passwords, deployment mode, network interface settings, SSH settings, date and time are configured.

The administrator is also prompted to enter a One Time Password (OTP) and is provided with a unique identification number (VIN) for the Sensor. This is used later when adding Sensor devices in the UI.

Finally, the NetScreen User Interface (UI) is installed. Versions are provided for Windows and Red Hat Linux. Once network components have been added (starting with the Sensors themselves) and a security policy defined and deployed (see following section), the IDP 500 is ready to go.

Configuration

The rule-based, centralised management system makes it easy to configure, update, and maintain the IDP system from a single administration point if required. An unlimited number of IDP Sensors can be managed with a single, logical Security Policy thanks to the ability to apply individual rules to specific sensors within the same policy. Once the administrator has defined an enterprise-wide Security Policy it can be distributed at the push of a button. All relevant configuration and signature information is automatically sent to the IDP Sensors in a secure manner (authenticated and encrypted).

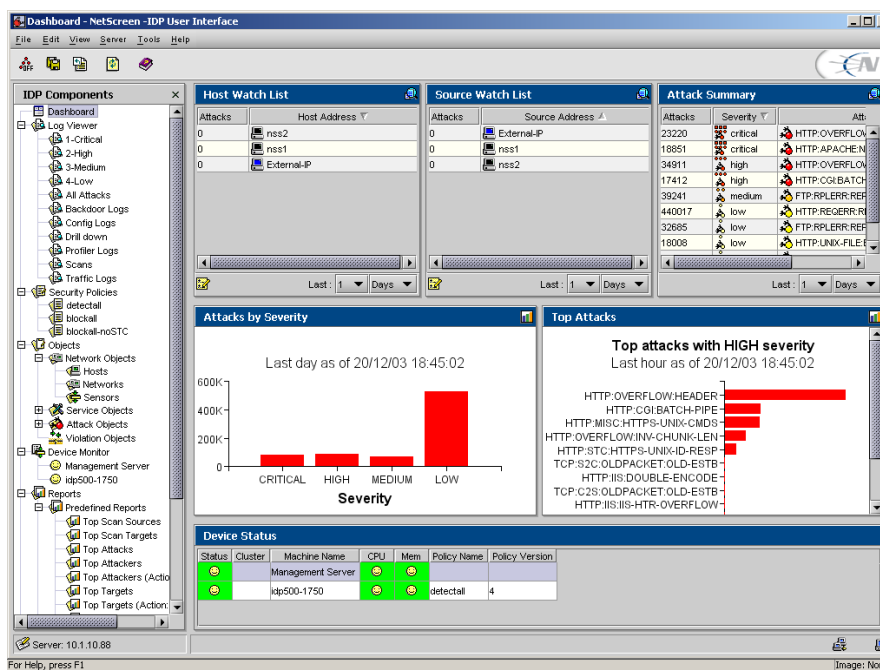


Figure 4 - NetScreen-IDP: The Dashboard display

An unlimited number of remote User Interface installations (running on either Windows- or Linux-based systems) can access and manage all aspects of the IDP system. The UI contains eight main components:

- [Dashboard](#)
- [Log Viewer](#)
- [Security Policy Editor](#)
- [Object Editor](#)
- [Device Monitor](#)
- [Reports](#)
- [Log Investigator](#)
- [Profiler](#)

Each of the above options is accessed via a hierarchical menu tree in a pane on the left of the UI (some of the above options have several sub-levels within the menu tree), whilst data is displayed in a second pane on the right. The Dashboard component is the default component shown in the main display area when the administrator first logs in.

This provides a one-screen high-level view which displays the vital statistics of the network and the IDP system (though it can be extremely slow to refresh when the number of alerts is excessive).

The data that appears in the Dashboard is real time, and refreshes periodically from the IDP Management Server. The administrator can customise the Dashboard to display only the information that is most important to him, including:

- **Target Host Watch List** - to monitor target hosts and/or networks of particular interest
- **Source Watch List** - to monitor activity from particular source addresses
- **Attack Summary** - list of the most recent/common attacks
- **Reports** - Any one of the built-in reports can be regularly refreshed on the Dashboard
- **Device Status** - To monitor the status of the Management Servers and all Sensors (availability, CPU utilisation, disk space, etc.). Thresholds can be set to flag devices as requiring attention if critical parameters are exceeded (i.e. CPU utilisation pegged above 90%, or disk space below a certain value).

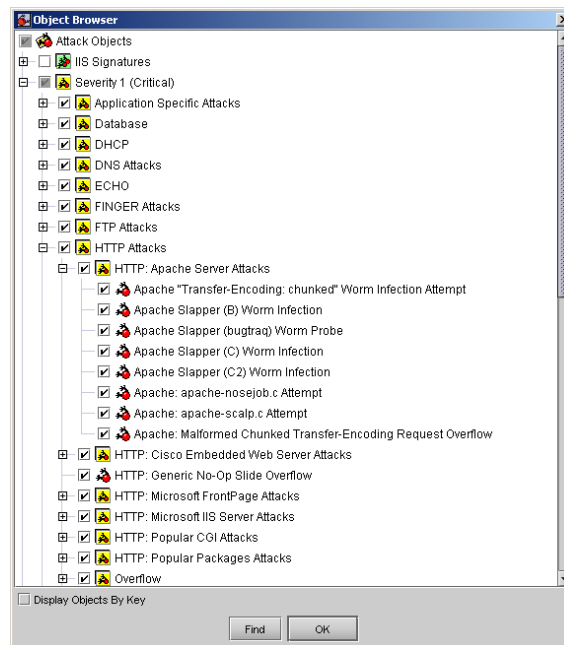


Figure 5 - NetScreen-IDP: Attack Objects

The first step in making the IDP system functional is to create *Network Objects* for each of the IDP Sensors. A Network Object is the IDP system representation of the components on the network, including:

- **Network Objects** - representing network components (individual hosts, networks, servers, Sensors).
- **Service Objects** - representing services running on the network (such as FTP, HTTP, and Telnet)
- **Attack Objects** - representing attack signatures and protocol anomalies, used to create rules for Security Policies.
- **Violation Objects** - allows the administrator to define which hosts and networks are allowed communicate, and on which ports. The Profiler will then report on any violations to these rules, enabling the enforcement of a corporate policy which goes beyond pure security.

Network Objects are the basic building blocks used in Security Policy rules to specify both the network components and the Sensor(s) which will protect those components.

Attack Objects are used in IDP rules to detect attack signatures and protocol anomalies that could compromise the network being protected. IDP includes a database of several hundred Attack Objects that can be updated weekly - a simple update procedure allows the administrator to connect to the NetScreen Web site or to upload Objects from a file which has been previously downloaded and stored locally.

The update procedure provides a detailed report of which Objects are new, which are updated, and which will be deleted (a nice touch, and one we would like to see more often in other products) - custom amendments can be preserved if required. The administrator has the option to select all the suggested changes or only some of them as required - this is a very well thought out update procedure.

Policy Management

Attack Objects are categorised and grouped according to severity first (*Critical, High, Medium, Low* or *Informational*), and target application or platform second. This is useful in that it allows the administrator to create a policy which, for example, monitors only for *Critical* exploits.

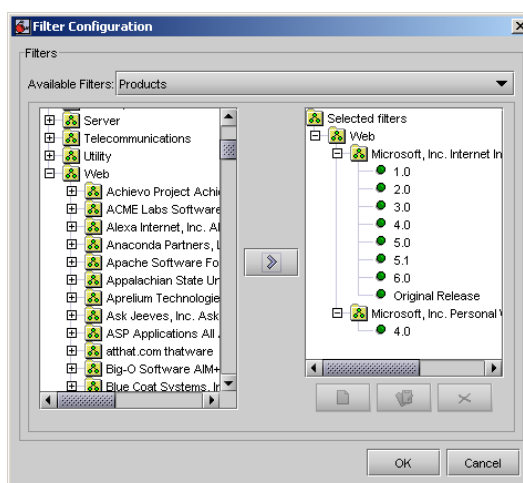


Figure 6 - NetScreen-IDP: Configuring dynamic groups

The latest software release also adds the concept of *dynamic groups* of signatures, which makes it possible to categories signatures by various properties including vendor (i.e. all Microsoft applications), application (i.e. all IIS or Apache signatures), protocol (i.e. SMB only), direction (i.e. client to server), severity (i.e. critical only) and potential for false positives (i.e. never).

These filters can be combined in many different ways, allowing the administrator to build very specific groups (i.e. all critical SMB signatures that are client-to-server and have a low possibility for false positives) or wide-ranging ones (i.e. all Microsoft applications). Whenever signature updates are downloaded, the appropriate signatures are added to the correct dynamic groups automatically, ensuring, for example, that your IIS policy automatically includes all the latest IIS signatures with no intervention required. This is an excellent feature.

The Object Editor provides the means to view and edit pre-defined Objects, as well as create new ones (including custom signatures).

When defining Security Policies in the IDP UI, they take the form of multiple *rules* laid out in a format that will be very familiar to most firewall administrators. Rules are organised into *rulebases*, collections of rules that use the same detection method to detect and prevent network intrusions.

There are seven IDP rulebases which combine to form a Security Policy:

- [Traffic Anomaly](#)
- [SYN-Protector](#)
- [Network Honeypot](#)
- [Main \(Attack Objects\)](#)
- [Backdoor Detection](#)
- [Sensor Settings \(sensor-specific parameters\)](#)
- [Exempt \(to allow certain traffic to be ignored\)](#)

Each rulebase uses a different detection mechanism to protect against intrusions (see *Architecture* section for details).

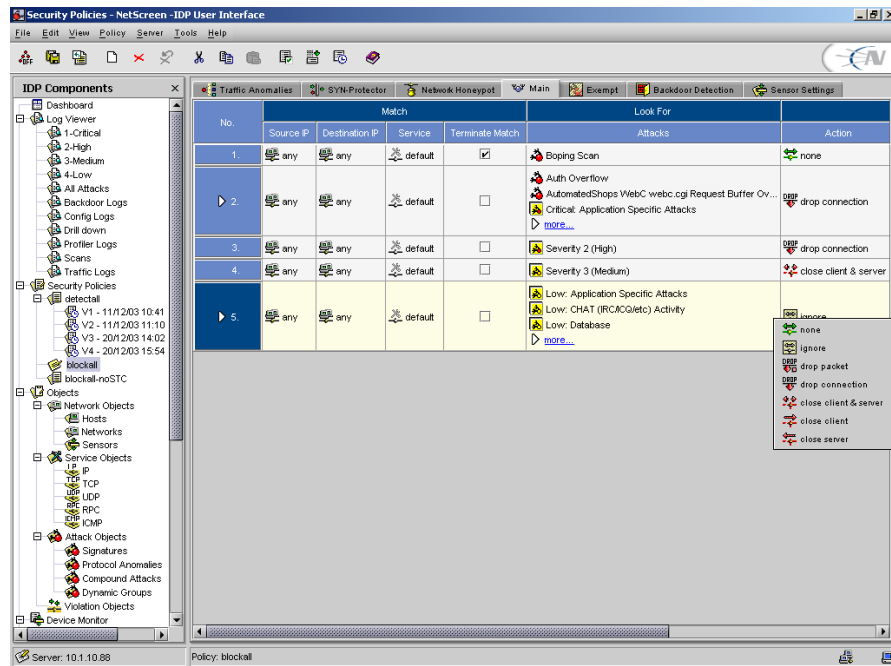


Figure 7 - NetScreen-IDP: Policy Editor

It is possible to have any number of Security Policies stored in the IDP system (three sample Policies are provided as examples), but each IDP Sensor can use only one Security Policy at a time. Installing a Security Policy on an IDP Sensor gives the Sensor a set of instructions about the actions the Sensor should take in response to specific network situations. The same Security Policy can be installed on multiple IDP Sensors, or unique Security Policies can be installed on each Sensor in the network.

When creating custom rules, the administrator defines the patterns or circumstances to be matched against each packet (including protocol, source and destination IP addresses, source and destination ports, flow direction, context within packet, packet contents, and so on), and the actions taken when a match is found.

Each Attack Object has five tabs of information which can be accessed via the Signature Editor:

- The **Basic tab** specifies the attack properties and the signature (attack pattern) that IDP uses to match attacks. This tab also includes the context in which the attack pattern occurs.
- The **Binding tab** specifies the service that the attack uses to attack the network.
- The **IP tab** specifies the packet fields, values, and options in a malicious packet.
- The **Protocols tab** specifies data length, flags, and other detailed data for the protocol header of a malicious packet.
- The **References tab** specifies known network security references for this attack.
- The **Extended Information tab** that provides detailed information from the TruSecure IntelliShield database

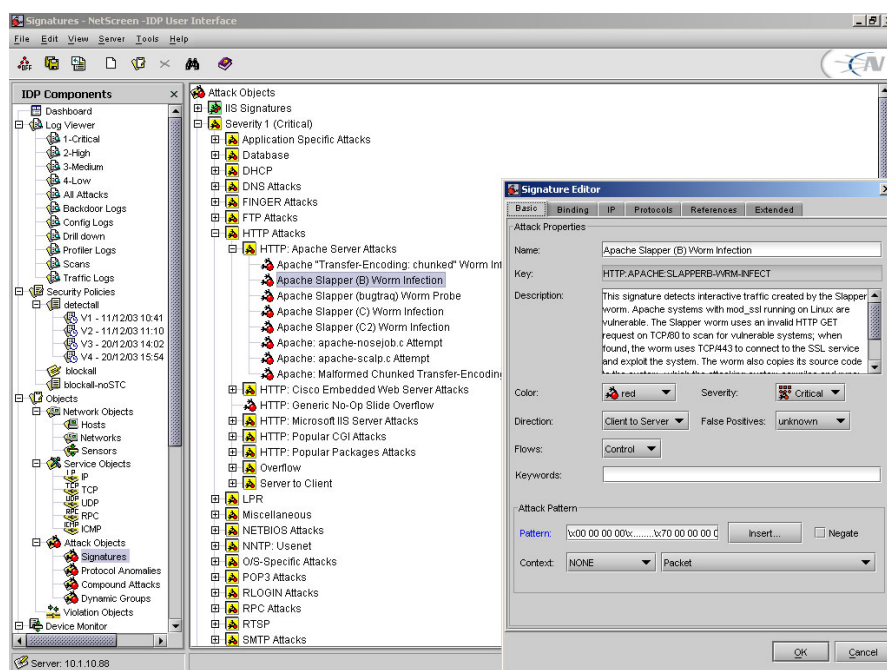


Figure 8 - NetScreen-IDP: Defining custom signatures

The creation of new pattern-matching signatures is relatively straightforward, although a knowledge of regular expressions would be a requirement in most cases for specification of the exploit patterns to be matched. Protocol anomaly signatures can neither be edited nor created from scratch, since these require extensive coding and can thus only be added or modified by NetScreen developers.

However, the latest release does add the ability to define *Compound Signatures*, which allow multiple individual signatures (both stateful and protocol anomaly) to be combined into a single signature to enable detection of complex exploits. This is an extremely powerful and flexible capability, providing for the creation of some very complex and accurate custom signatures.

IDP is one of the more open systems currently available, providing the means to copy and edit the built-in Attack signatures. All of the critical fields from the protocol decoders are exposed to the UI, and can be used in regular expressions to match for suspicious traffic.

Note that this is more than simply “pattern matching”, where a signature has to look for a particular pattern in a lengthy data stream. Instead, it is possible for an individual field - say the FTP user name - to be isolated and used in complex expressions. These can then be combined together (perhaps looking for a particular user name followed by a password greater than a specified length) into the aforementioned Compound Signatures.

The IDP contains a default set of actions that the Sensor can perform against network traffic that matches the specified pattern or circumstance when operating in-line. These actions include :

- **None** - takes no action against the connection.
- **Ignore** - ignores the remainder of a connection after a Attack Object in a Security Policy rulebase is matched.
- **Drop packet** - drops a matching packet before it can reach its destination but does not close the connection. Used to drop packets for attacks in traffic that is prone to spoofing, such as UDP traffic, where the dropping of a connection for such traffic could result in a denial of service that prevents subsequent receipt of traffic from a legitimate source IP address.
- **Drop connection** - drops the connection without sending a RST packet to the sender, preventing the traffic from reaching its destination. Used to drop connections for traffic that is not prone to spoofing.
- **Close client and server** - closes the connection and sends a RST packet to both the client and the server. If the IDP Sensor is in sniffer mode, IDP sends a RST packet to both the client and server but does NOT close the connection.
- **Close client** - closes the connection to the client, but not to the server.
- **Close server** - closes the connection to the server, but not to the client.

If the current network traffic matches a rule, IDP can perform an IP action against future network traffic that uses the same IP address. IP actions are similar to other actions in that they tell IDP to drop or close the connection, but they also allow the Sensor to block the attacker for a specified amount of time.

A number of notification options are also available whenever an attempted exploit is detected:

- **Log** - Generate a log record for security events
- **Alarm** - If alarm is selected and the rule is matched, IDP places an alarm flag in the alarm column of the Log Viewer for the matching log record (used primarily for rapid filtering of alerts with higher than normal levels of importance)
- **Tracking Session Data** - Session data includes the number of packets passed, the number of bytes, and the elapsed time. Detailed session data can be tracked in order to enter the information into a log or traffic analysis tool, monitor user activity on the network, determine the length of sessions, or monitor the number of bytes transferred
- **SNMP Trap** - Sends an SNMP trap to the SNMP Manager.

- **Syslog Message** - Sends an entry to a syslog server.
- **E-mail** - Sends an e-mail to designated recipients.
- **Run Script** - Run custom scripts that take information about a security event and process it in a specific way.
- **Logging Packets** - It is possible to record the individual packets in the network traffic that matched a rule by capturing the packet data for the attack. By default, only the trigger packet is captured, though it is also possible to specify a number of packets to be captured both before and after the trigger - a nice feature.

Action and Alert options can only be set at the individual rule level. Given that a single rule could contain a single Attack Object or every single Attack Object available, it could be said that this is as flexible as it needs to be.

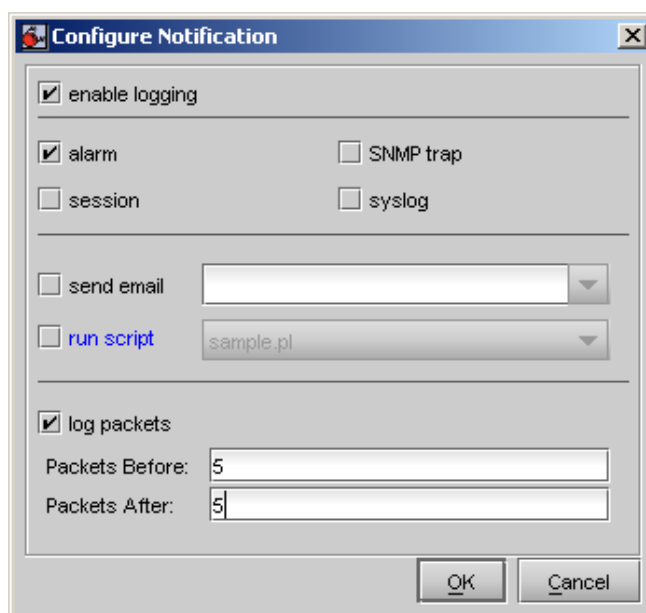


Figure 9 - NetScreen-IDP: Notification options

However, in order to improve readability, most Security Policies would consist of multiple rules. The ability to set Action and Alert options on a global Policy-wide basis would thus be a welcome addition, saving the administrator the effort of having to change these options for every single rule in a Policy if, for example, he wanted to turn on or off packet logging for all Objects.

When editing Policies, it is often useful to be able to search for existing signatures that fit specific criteria. The IDP UI contains a search facility which allows the administrator to search for one or more signatures based on various criteria including name, key words, description, CVE reference, BugTraq ID, false positive rating or target application/platform. This is available at both the Attack Object level and within the Policy Editor (the latter being a very welcome new capability in the V3.0 release).

It does, however, only find signatures one at a time, requiring you to keep pressing "next" in order to find all matches for a given search. It also allows you to search on only a single parameter at a time. It would be nice if it allowed the specification of multiple parameters in a single search and returned multiple signatures in the result set, perhaps then allowing the application of bulk edits to the entire group.

The *Exempt* rulebase allows the administrator to specify that certain legitimate traffic be ignored for the purposes of matching for exploits. For example, vulnerability scanners can trigger many signatures, and so it would be possible to define rules that allowed a specific machine to scan one or more servers on the corporate network without hindrance. These rules can be defined manually, or by selecting an alert and clicking on the “Exempt” menu option. The latter method is ideal for rapidly fine tuning a security policy based on false alarms which may have been raised.

Finally, it is worth noting that all of the main Sensor operational parameters - such as flow table sizes, flow timeouts, fragmentation reassembly parameters, maximum flows, thresholds, etc. - are exposed in the *Settings* tab in the Policy Editor. Thus it is possible to alter default values to accommodate environments where, for example, it is necessary to monitor more than the default 100,000 open connections. The default values are adequate for most circumstances, but it is nice to see parameters such as these made available for editing in such a straightforward manner.

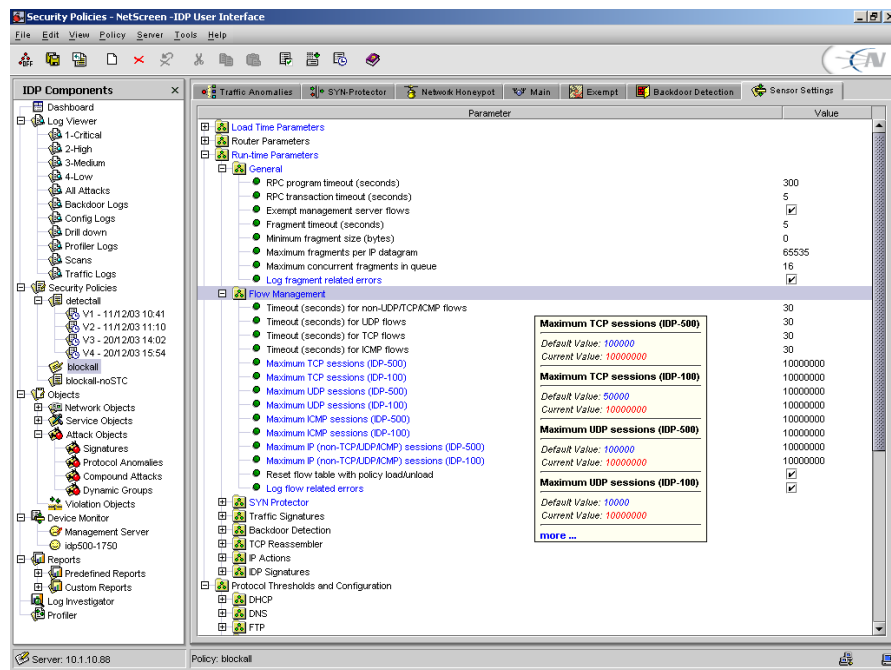


Figure 10 - NetScreen-IDP: Sensor settings

Once a Security Policy has been created and verified, it can be deployed to one or more Sensors - a list of Sensors appears during deployment allowing the administrator to select all Sensors (the default), groups of Sensors or an individual Sensor. One useful addition would be the ability to logically group Sensors together in order to select a specific collection of Sensors (those protecting DMZ networks, for example) with a single check box.

When creating or editing individual rules within a Policy, it is possible to specify whether that rule applies to all Sensors or one or more individual devices. When the Security Policy is subsequently deployed, the rules become active only on the Sensors selected in the rulebase. This provides the option to have a single global Security Policy, with individual rules activated or deactivated on a per-Sensor basis. Alternatively, it is also possible to create multiple Security Policies to be applied to different Sensors.

This is an extremely flexible method of handling Policy creation and deployment, and is a feature we would like to see implemented in more IDS products.

A Sensor can only use one Security Policy at a time, and when a new Security Policy is installed, it overwrites any existing policies on the Sensor. Each time Policy is deployed, however, the IDP saves a version of the Policy currently installed. A version is a copy of the original Security Policy, and contains information about when the policy was installed on the Sensors and which user performed the install.

It is possible to view, reinstall, and delete previous versions of a Security Policy. However, because Policy versions are designed to provide an audit trail of the rules that were operating on a Sensor at a given time, it is not possible to edit previously installed Security Policy versions.

The administrator can print and export (as PDF or postscript) all rulebases in a Security Policy.

Alert Handling

In the Log Viewer, the administrator can view log records that IDP Sensors generate based on criteria defined in the Security Policies.

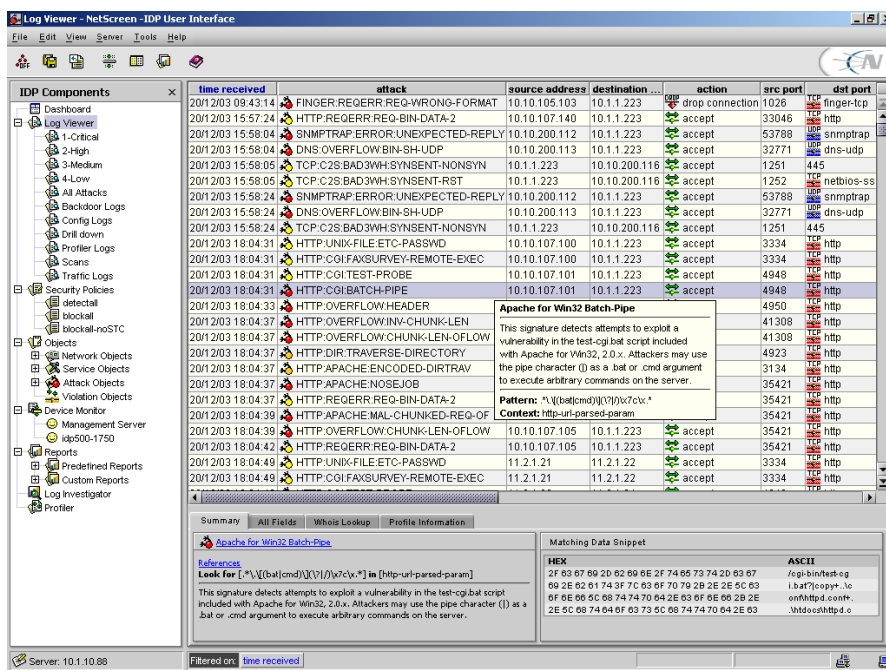


Figure 11 - NetScreen-IDP: Log Viewer

The Log Viewer displays log records in table format and can be modified using filters or column settings. When viewing log entries it is important to be able quickly to access all information relevant to that particular alert.

Thus, IDP provides one-click navigation between the Log Viewer (shows incidents), the Session Viewer (shows associated packet contents in detail), the Signature Editor (shows why the incident was triggered), the Security Policy Editor (shows which Policy triggered the event), and the Enterprise Security Profiler (ESP) (provides an on-demand view of network activity for detailed analysis).

The *Signature Editor* also contains a *References* tab that provides links to outside information about the attack, such as the attack CVE number and description, as well as an *Extended Information* tab that provides detailed vulnerability and exploit information from the TruSecure IntelliShield database.

The ability to access directly the Security Policy that triggered a particular alert is extremely useful when attempting to refine Policies to reduce the incidence of false positives. If an alert is identified as a false positive, the administrator can access the Policy with a single click, disable or tune that Attack Object, and re-deploy the Policy immediately.

Log records contain a great deal of data about the activity on the network, but it is not always necessary nor desirable to see all the information at one time. The administrator can use the Log Viewer to manipulate, analyse, and export the information contained in the log records to display only the information that is of immediate interest. For example, using the Log Viewer it is possible to:

- View complete, summarised, or detailed information for each log record
- Show, hide, or move columns, and filter data by column headings
- Create a new instance of the Log Viewer that shows only custom filters and column settings - these can be saved as custom views
- Set flags on log records to indicate a specific priority or action
- Launch the Packet Viewer
- Print or export log record data

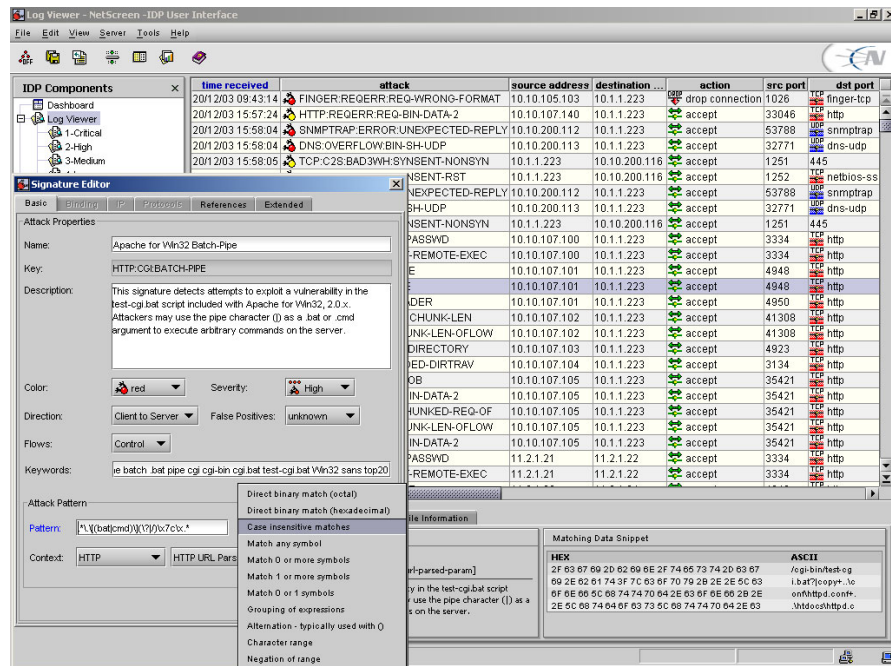


Figure 12 - NetScreen-IDP: Tuning the Policy via the Log Viewer

Although the main part of the Log Viewer screen provides a tabular view of the log entries, selecting individual entries populates four tabs at the bottom of the screen which provide more detail on the alert in question.

The “All Fields” tab duplicates all the data from the log record in a condensed format removing the need to scroll across and thus making it easier to read. The “Summary” tab provides a detailed description of the security event (including extensive material from the TruSecure database), and includes links to external references (such as the CVE database).

The remaining two tabs at the foot of the screen provide an instant whois lookup for the selected IP address, and access to detailed host, application and networking information collected by the Enterprise Security Profiler (ESP) module (covered in more detail in the Reporting and Analysis section).

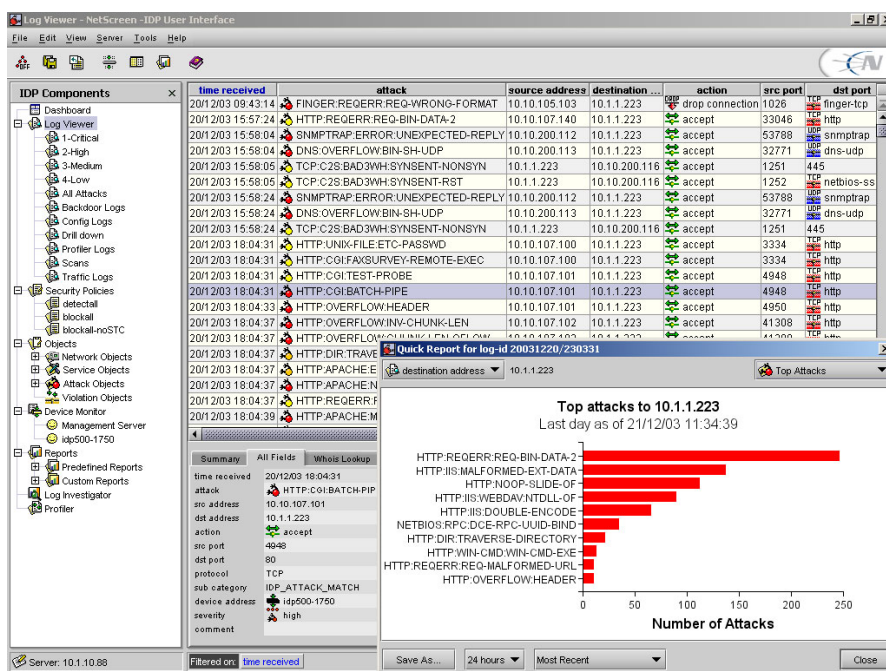


Figure 13 - NetScreen-IDP: Quick Reports from the Log Viewer

Right-clicking on any log entry brings up a sub-menu which allows the administrator immediately to:

- Set filters based on that entry - this includes or excludes all other similar entries in order to provide a more focussed list of alert entries
- Search for similar alerts - this finds similar entries one after the other each time the “next” button is pressed, rather than grouping them together for display as would happen when setting a filter
- Display packet data, security policy or attack/signature data
- Run Quick Reports - this produces rapid reports based on search criteria populated automatically from the selected alert. For example, by right clicking on the destination address a Quick Report can be launched which instantly provides the recent Critical Severity attacks to that particular address. This can be further configured (i.e. to display Top 10 attacks to that address) and then saved as Custom Reports to be re-run on demand.
- Flag the alert.

A number of flags can be set to aid other administrators, including the priority of an alert, and whether it: is a false positive; has been assigned to an investigator; should be followed up; is pending; is closed.

There is no way to correlate or group related alerts together, however, either automatically or manually.

It is possible to create multiple unique, customised views of the log data within the Log Viewer by:

- *Changing the column settings (reorder columns)*
- *Filtering by column headings (hide columns)*
- *Filtering by cell content (display all records with a particular source IP address, or ignore all records targeted at a particular host)*
- *Collapsing and/or hiding similar log records to aid readability*

Any number of such customised views can be saved along with all applicable filters enabling them to be recalled and used repeatedly. It is also possible to export log records for use in other applications. However, to export to XML, CSV, SNMP, SMTP, Syslog, script, or Postgres database format, it is necessary to use the command line interface (although a UI is also available for export to a Postgres database)

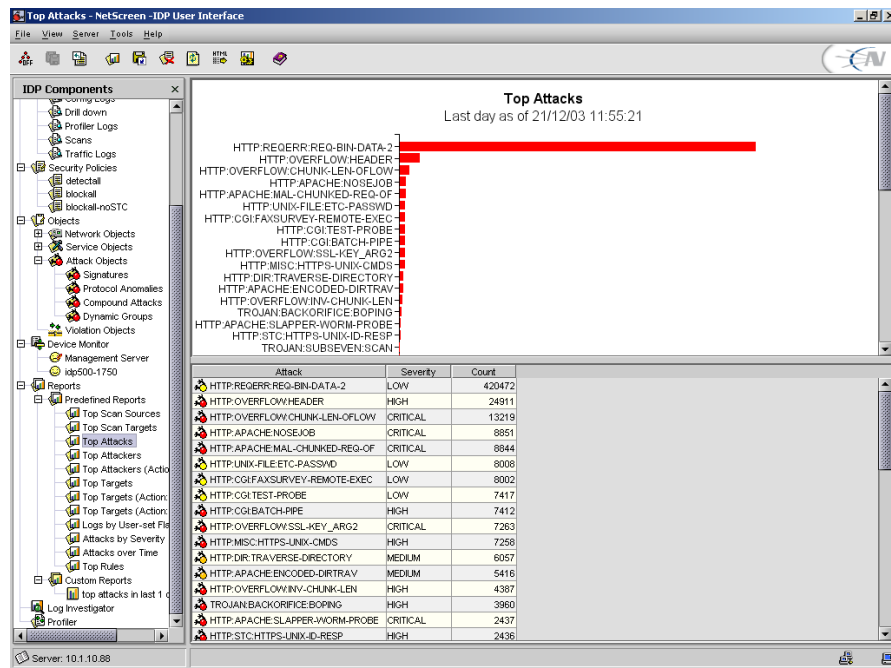


Figure 14 - NetScreen-IDP: Top 10 Attacks report

New log files are created each day automatically. Log records which are no longer needed can be purged from the Management Server database. The administrator can also set the IDP system to automatically purge logs when the Management Server disk space becomes low. Log entries are lost once purged - there is no archive option.

Reporting and Analysis

There are a number of pre-defined graphical reports included in the UI which are based on log records that summarise threats and traffic behaviour over time.

These include:

- [Top Targets](#)
- [Top Targets \(where Action was Drop or Close\)](#)
- [Top Targets \(where Action was None\)](#)
- [Top Scan Sources](#)
- [Top Scan Targets](#)
- [Top Attacks](#)
- [Top Attackers](#)
- [Top Attackers \(where Action was Drop or Close\)](#)
- [Logs by User-set Flag](#)
- [Attacks by Severity](#)
- [Attacks Over Time](#)
- [Top Rules](#)

Reports present log record data using Log Viewer filters. After selecting a report it is possible to set individual options for its appearance, including selecting a date range and chart style. It is also possible to view the log records on which the report table entries are based (by right-clicking on the data to view and selecting the “View in Log Viewer” option), allowing the administrator to “drill down” quickly from a high level graphical summary to a detailed view.

Custom reports are a new feature in the latest release, and provide the administrator with the means to specify his own selection criteria, report content (from a list of data collected for each alert), time period and report style (count or time based, bar or pie chart).

These can be named and saved in the *Custom Report* section of the main menu tree. In addition, “Quick Reports” can be created rapidly from the *Alert Viewer* based on criteria taken directly from the alert being viewed at the time, and these too can be saved in the *Custom Reports* section for later use.

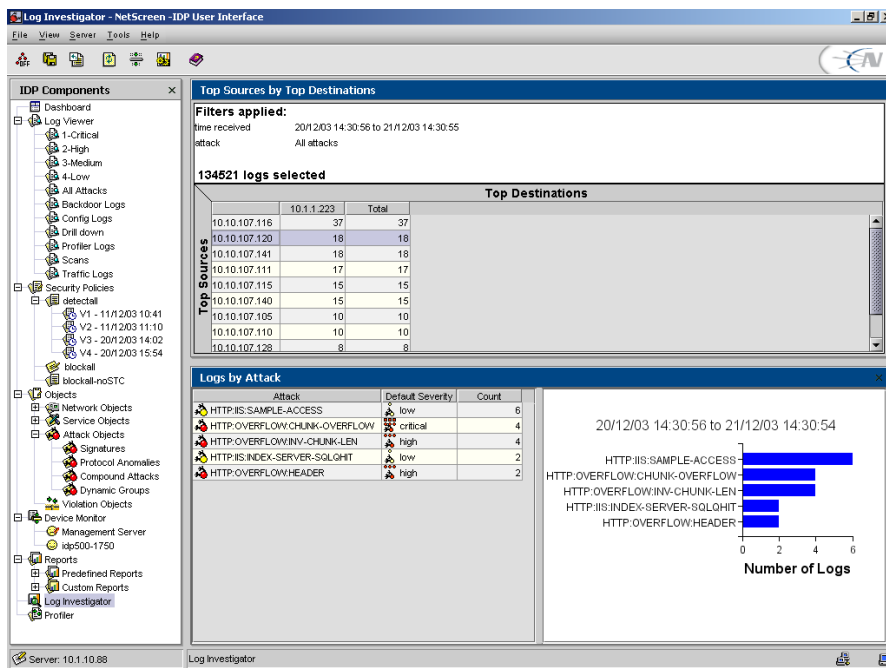


Figure 15 - NetScreen-IDP: Log Investigator

The *Log Investigator* is designed along the lines of a dynamic “reporting spreadsheet” and is intended to provide the administrator with the means to “slice and dice” log entries in various ways. It allows him to select different data sets via the use of filters and include/exclude options to incrementally refine the search for a specific range of entries.

The Log Investigator uses a default filter that selects only log records with the value of *ATTACK* in the category column of the Log Viewer. The administrator can choose which criteria he wishes to see on the left and right axes of the table, and then set additional filters to narrow the focus of his analysis of the data. For example, he might choose to filter by the *flag* column and have the Log Investigator only consider log records with *ATTACK* in the category column and a *HIGH* setting in the flag column.

For individual table cells, or entire rows and columns, the administrator can “zoom in” to focus on the destination ports, attacks, or time. The Log Viewer then appears in a separate window, displaying the log records which produced the entries in the Log Investigator. The filters that produced the Log Investigator entries appear on the status bar.

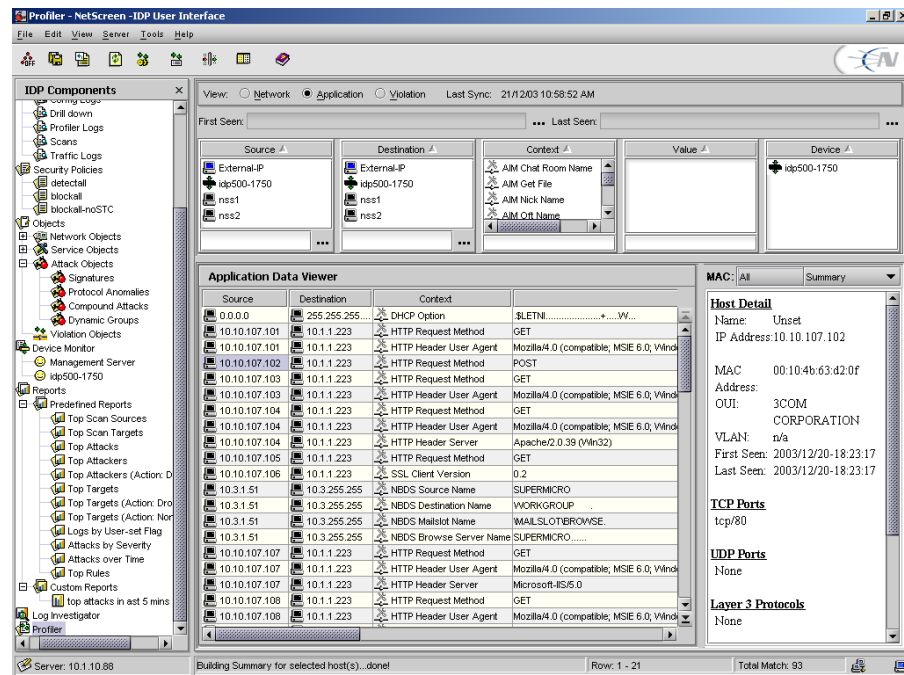


Figure 16 - NetScreen-IDP: Enterprise Security Profiler (ESP)

The final forensic and reporting tool is new to the latest V3.0 release - the *Enterprise Security Profiler* (ESP). This is intended to help administrators monitor their network security posture by collecting forensic data on network traffic and storing that data to form a *network profile*. Data can then be used to identify changes in network behaviour, which in turn can be translated into policy modifications. In addition to being a proactive analysis and planning tool, data collected by ESP provides administrators with a useful post attack analysis tool.

The way ESP works is that it stores data already collected during the attack detection process, and makes that available via a number of database views and filters. The administrator needs to first configure the system to collect Profiler data - it does not do this by default (presumably because there will be some impact on overall performance of the sensor).

ESP collects data at both the network layer and the application layer. Network information collected includes items such as IP addresses (host connected to and from), and port number, whilst application level data includes items such as the actual application used over the network (Internet Explorer, Apache, IIS, and so on), version number, user name, and URL.

Data on individual hosts (IP address, MAC address, host name, protocols used, ports used, etc.) is stored and displayed within the *Profiler* (it has its own menu option) as well as against each alert in the Alert Viewer, where it populates its own *Profile Information* tab in the alert details window at the bottom of the main Alert Viewer screen. Once the Profile Information tab has been selected for a particular alert, the administrator can choose from a number of display options in a drop down menu, offering host detail, TCP ports used, UDP ports used, RPC services, layer 3 protocols used, and details of other machines with which the host has communicated.

The *Profiler* view - selectable via the main menu tree - provides a direct view into the Profiler database. A number of selection criteria are available to enable the administrator to examine which clients are communicating with which servers, which protocols are being used, and which applications are being run over his network.

There are three different views in the Profiler to help administrators zero in on the information they need: the *Network View* for layer 3 and 4 information, *Application View* for layer 7 data, and *Violation View* for pseudo firewall policies that help pinpoint what an administrator may not know about the network. The Network View provides data filtered and displayed on source and destination IP address, port number, MAC address, and so on.

The Application View can often be much more interesting, since it displays the context of each suspicious activity, including the server name and version, the actual command issued, the URL requested, and so on. Thus it is a simple matter to focus on FTP servers, then select all those occasions where someone tried to log in as *root* or *administrator*. Likewise, it would be a simple matter to focus on a particular version of Apache or IIS, and then see what activity has been aimed at that particular server or group of servers.

This information can then be used to set *Violation* rules, to ensure that only valid client-server communications are allowed across the network. For example, the administrator can define that FTP traffic is allowed only between one particular group of IP addresses and one particular FTP server, or that Telnet traffic is not allowed at all. This way, IDP can detect and prevent malicious traffic based purely on policy violations rather than specific exploits. The *Enterprise Security Profiler* is an extremely useful feature, and one that goes beyond pure *intrusion detection* or *prevention*, and into the realms of *policy enforcement*.

Verdict

Performance

Although equipped with Gigabit interfaces, NetScreen rates this product at a maximum of 500Mbps.

This rating was verified in our tests under almost all conditions, the only exceptions being when the network load exceeds 450,000 packets per second or 7500 HTTP connections per second. Neither of these conditions are likely to be seen on a typical sub-Gigabit network.

Latency was acceptable under all normal network conditions, though we noted that it could be an issue when the device is under a heavy DOS attack and the SYN-Protector is enabled. However, the SYN-Protector does provide some mitigation and protection against DOS and DDOS attacks, and thus the increase in latency is an acceptable trade-off in those circumstances.

In all other respects, the NetScreen IDP-500 performed consistently and reliably throughout our tests. Under extended attack it continued to pass legitimate traffic whilst blocking attack traffic in a consistent manner, and it proved impervious to extended ISIC attacks.

High Availability (HA) options are well covered with the IDP-500, with full active-active failover and even load balancing to a maximum of 16 sensors.

Security Effectiveness

Signature recognition (with blocking disabled) was excellent out of the box at 96 per cent and the blocking performance was identical. This increased to an excellent 100 per cent after application of the signature updates. Following the update, all of our “false negative” (modified exploit) cases were detected correctly, demonstrating that the NetScreen signatures are generally designed to detect the underlying vulnerability rather than a specific exploit.

We noted a general increase in accuracy and coverage with this latest release, no doubt due to the introduction of the new *Compound Signature* capabilities and a lot of good work by NetScreen in revamping the existing signature set.

Resistance to false positives was generally very good, and the NetScreen IDP-500 performed impeccably in all of our evasion tests as well as proving itself resistant to stateless flooding tools such as *Stick* and *Snot*. Out of the box, the device maintained state on up to 100,000 open connections, and this was easily increased to 500,000 with the change of a single parameter.

Usability

In terms of day to day operation, the NetScreen-IDP 500 demonstrates high levels of usability with an excellent set of centralised management tools. The three tier architecture should scale well, particularly given the claims for the proprietary inter-device communications protocol and high-performance alert logging database.

The management interface is very straightforward to use, allowing the administrator to create and modify Security Policies easily. Signature updates occur regularly and the update procedure is excellent, allowing the administrator to “vet” the changes before new signatures are downloaded.

The facility to create custom signatures is available, and is also very straightforward to use.

An extensive search facility within the Policy Editor provides the administrator with the means to effect very precise searches for specific signatures, which can then be easily copied or modified. In this respect, NetScreen IDP is one of the most open systems we have seen, and the new *Compound Signature* capability provides the means to create some very complex and accurate signatures.

Once Policies have been created or updated, a single button provides the means to deploy them to all Sensors in a single operation. One really nice feature is the automatic version control, with a complete audit trail of all previous Policies applied to each Sensor, and the ability to roll back to a previous version if required.

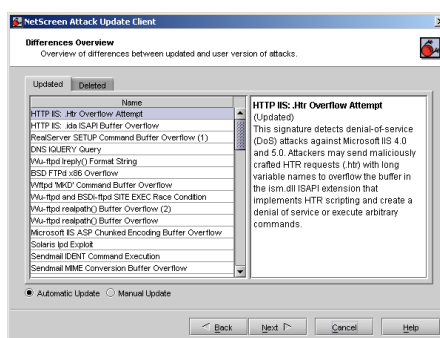


Figure 17 - NetScreen-IDP: Signature update process

The administrator is provided with a high degree of flexibility and control in how Policies are defined, stored and deployed, and we found Policy management to be very strong in this product.

Alert handling, too, was well implemented. Alerts appear quickly at the UI Console, and the descriptions are generally very accurate. The ability to apply filters and drill down to more details views within the Log Viewer and Log Investigator makes it a useful tool for the forensic investigator. New additions in the latest release include custom reports, *Quick Reports* (created from the Alert Viewer) and the excellent *Enterprise Security Profiler*. The latter provides the means for the administrator to take IDP beyond pure *intrusion detection* and *prevention*, and into the realms of *policy enforcement*, increasing its usefulness in many corporate environments.

All of these combine to make the NetScreen-IDP one of the strongest products in terms of reporting and forensic analysis.

Contact Details

Company: NetScreen Technologies

E-mail: info@netscreen.com

Internet: www.netscreen.com

Address:
805 11th Ave, Building 3
Sunnyvale, CA 94089

Tel: +1 408 543 2100

APPENDIX A – TEST RESULTS

The aim of this procedure (based on V1.0 of the NSS Group Network IPS Testing Methodology) is to provide a thorough test of all the main components of an in-line Intrusion Prevention System (IPS) device in a controlled and repeatable manner and in the most “real world” environment that can be simulated in a test lab.

The Test Environment

The network is 100/1000Mbit Ethernet with CAT 5e cabling and a mix of Allied Telesyn AT-9816GB and AT-9812T switches (these have a mix of fibre and copper Gigabit interfaces). All IPS devices are expected to be provided as appliances - if software-only, the supplier pre-installs the software on the recommended hardware platform. The IPS is configured as a perimeter device during testing (i.e. as if installed behind the main Internet gateway/firewall). There is no firewall protecting the target subnet.

Traffic generation equipment - such as the machines generating exploits, Spirent Avalanche and Spirent Smartbits *transmit* port - is connected to the “external” network, whilst the “receiving” equipment - such as the “target” hosts for the exploits, Spirent Reflector and Smartbits *receive* port - is connected to the internal network. The IPS 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 IPS appliance to the management console on a private subnet. This ensures that the sensor and console can communicate even when the target subnet is subjected to heavy loads, in addition to preventing attacks on the console itself.

Section 1 – Detection Engine

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

Test 1.1 - Attack Recognition

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

Our attack suite contains over 100 exploits covering the following areas:

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

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

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

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

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

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

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

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

Test 1.2 - Resistance To False Positives

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

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

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

- [Test 1.2.1 - False positives](#)

Section 2 – IPS Evasion

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

Test 2.1 - Baselines

The aim of this test is to establish that the sensor is capable of detecting and blocking a number of common basic attacks (our baseline suite) in their normal state, with no evasion techniques applied.

- [Test 2.1.1 - Baseline attack replay](#)

Test 2.2 - Packet Fragmentation and Stream Segmentation

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

- [Test 2.2.1 - IP fragmentation - ordered 8 byte fragments](#)
- [Test 2.2.2 - IP fragmentation - ordered 24 byte fragments](#)
- [Test 2.2.3 - IP fragmentation - out of order 8 byte fragments](#)
- [Test 2.2.4 - IP fragmentation - ordered 8 byte fragments, duplicate last packet](#)
- [Test 2.2.5 - IP fragmentation - out of order 8 byte fragments, duplicate last packet](#)
- [Test 2.2.6 - IP fragmentation - ordered 8 byte fragments, reorder fragments in reverse](#)
- [Test 2.2.7 - IP fragmentation - ordered 16 byte fragments, fragment overlap \(favour new\)](#)

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

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

Test 2.3 - URL Obfuscation

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

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

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

Test 2.4 - Miscellaneous Evasion Techniques

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

- [Test 2.4.1 - Altering default ports](#)
- [Test 2.4.2 - Inserting spaces in FTP command lines](#)
- [Test 2.4.3 - Inserting non-text Telnet opcodes in FTP data stream](#)
- [Test 2.4.4 - Polymorphic mutation \(ADMmutate\)](#)
- [Test 2.4.5 - Altering protocol and RPC PROC numbers](#)
- [Test 2.4.6 - RPC record fragging](#)

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

Section 3 – Stateful Operation

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

Test 3.1 - Stateless Attack Replay (Mid-Flows)

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

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

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

In order to receive a “PASS” in this test, no alerts should be raised for any of the actual exploits. However, each packet should be blocked if possible since it represents a “broken” or “incomplete” session.

- [Test 3.1.1 - Stateless attack replay](#)

Test 3.2 - Simultaneous Open Connections (default settings)

This test determines whether the sensor is capable of preserving state across increasing numbers of open connections, as well as continuing to detect and block new exploits when the state tables are filled.

It also attempts to determine whether or not the sensor will block legitimate traffic once state tables are filled.

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

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

At each step, we ensure that the IPS engine is still capable of detecting and blocking freshly-launched exploits once all the connections are open, as well as confirming that the device does not block legitimate traffic (perhaps as a result of state tables filling up). This test is run using the default sensor settings.

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

Test 3.3 - Simultaneous Open Connections (after tuning)

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

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

Section 4 – Detection/Blocking Performance Under Load

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

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

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

A fixed number of exploits are launched with zero background traffic to ensure the sensor is capable of detecting our baseline attacks. Once that has been established, increasing levels of varying types of background traffic are generated **through** the 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 should this be less than 1Gbps).

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

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

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

Test 4.1 - UDP Traffic To Random Valid Ports

This test uses UDP packets of varying sizes generated by a **SmartBits SMB6000** with LAN-3301A 10/100/1000Mbps **TeraMetrics** cards installed. A constant stream of the appropriate mix of packets - with variable source IP addresses and ports transmitting to a single fixed IP address/port - is transmitted through the IPS device bi-directionally (maximum of 1Gbps, or 500Mbps in each direction). 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, and the aim of this test is purely to determine the raw packet processing capability of the IPS device, and its effectiveness at passing “useless” packets quickly in order to pass potential attack packets to the detection engine.

- **Test 4.1.1 - 64 byte packets - maximum 1,480,000 packets per second:** *The “torture test” - it is unlikely that any real-life network will ever see network loads of almost 1.5 million 64-byte packets per second. This test therefore measure packet processing performance under the most extreme conditions. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic.*
- **Test 4.1.2 - 440 byte packets - maximum 260,000 packets per second:** *This test has been included to provide a comparison with our “real world” packet mixes, since the average packet size is similar. No sessions are created during this test and there is very little for the detection engine to do in the way of protocol analysis. This test provides a reasonable indication of the ability of a device to process packets from the wire on an “average” network, and we would expect all products to demonstrate good performance levels. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic.*
- **Test 4.1.3 - 1514 byte packets - maximum 81,720 packets per second:** *This test is the complete opposite of the 64 byte packet test, in that we would expect every single product to be capable of returning 100 per cent detection rates across the board when using only 1514 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 packet sizes. Repeated with 250Mbps, 500Mbps, 750Mbps and 1000Mbps of background traffic.*

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

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

The use of the Spirent Communications Gigabit **Avalanche** and **Reflector** allows us to create true “real world” traffic at speeds of up to 2.2 Gbps as a background load for our IPS tests. Our Avalanche configuration is capable of simulating over 2.5 million users, with over 2.5 million concurrent sessions, and almost 100,000 HTTP requests per second.

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

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

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

- **Test 4.2.1** - Max 2,500 new connections per second - average packet size 1200 bytes - maximum 100,000 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 IPS sensors to achieve 100% blocking rates throughout this test.
- **Test 4.2.2** - Max 5,000 new connections per second - average packet size 540 bytes - maximum 230,000 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 IPS sensors to perform well in this test.
- **Test 4.2.3** - Max 10,000 new connections per second - average packet size 440 bytes - maximum 280,000 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 IPS sensor, and represents a very heavily used production network.
- **Test 4.2.4** - Max 20,000 new connections per second - average packet size 350 bytes - maximum 360,000 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 IPS sensor. Not many sensors will perform well at all levels of this test.

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

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

- **Test 4.3.1** - Max 5,000 new connections per second - average packet size 540 bytes - maximum 230,000 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 IPS sensors to perform well in this test.
- **Test 4.3.2** - Max 10,000 new connections per second - average packet size 440 bytes - maximum 280,000 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 IPS sensor, and represents a very heavily used production network.

Test 4.4 - Protocol Mix Traffic

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

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

- **Test 4.4.1** - 72% HTTP traffic (560 byte packets) + 20% FTP traffic + 6% UDP traffic (256 byte packets). Max 380 new connections per second - average packet size 555 bytes - maximum 22,000 packets per second - maximum 136 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 IPS sensors to perform well throughout this test.

Test 4.5 - “Real World” Traffic

This is as close as it is possible to come to a true “real world” environment under lab conditions. For this test we eliminate the Reflector device and substitute an IIS Web server installed on a dual P4 SuperMicro server with Gigabit interface. This server holds a copy of The NSS Group Web site, and is capable of handling 950Mbps 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 **25 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 much more “bursty” in nature.

- **Test 4.5.1 - Pure HTTP Traffic (simulated browsing session on NSS Web site):** Max 100 new connections per second - 25 new users per second - average packet size 1000 bytes - maximum 110,000 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 IPS sensors to perform well at all load levels in this test.
- **Test 4.5.2 - Protocol Mix (72% HTTP traffic (simulated browsing sessions as 4.5.1) + 20% FTP traffic + 6% UDP traffic (256 byte packets)):** Max 550 new connections per second - average packet size 900 bytes - maximum 130,000 packets per second - maximum 11,000 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 IPS sensors to perform well at all load levels in this test.

Note that the average packet sizes during this test do not match the various studies on Internet packet size distribution (we consider the average packet size to be in the region of 440-550 bytes).

However, given that this is a test which utilises real browsing sessions against a real server on a real network, the resulting packet distribution should be considered valid.

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

Section 5 – Latency & User Response Times

The aim of this section is to determine the effect the IPS 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.

Test 5.1 - Latency

We use Spirent SmartFlow software and The SmartBits SMB6000 with Gigabit TeraMetrics cards to create multiple traffic flows through the IPS 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 (500Mbps in each direction, 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 (64 bytes, 440 bytes and 1518 bytes) of UDP traffic with variable IP addresses and ports. At each iteration of the test, SmartFlow records the number of packets dropped, together with average and maximum latency.

- **Test 5.1.1 - Latency With No Background Traffic:** *SmartFlow traffic is passed across the infrastructure switches and through the device (the latency of the basic infrastructure is known and is constant throughout the tests). The packet loss and average latency are recorded at each packet size and each load level from 250Mbps to 1Gbps (in 250Mbps steps).*
- **Test 5.1.2 - Latency With Background Traffic Load:** *The Avalanche and Reflector are configured to generate a fixed amount of background HTTP traffic through the IPS 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 115,000 packets per second). A 250Mbps load (125Mbps bi-directional) of SmartFlow traffic at various packet sizes (64 byte, 440 byte and 1514 byte) is then passed across the infrastructure switches and through the device and the packet loss and average latency are recorded.*
- **Test 5.1.3 - Latency When Under Attack:** *The Spirent WebSuite software is used to generate a fixed load of DOS/DDOS traffic of 10 per cent of the maximum rated bandwidth of the device under test (maximum 100Mbps).*

A 250Mbps load (125Mbps bi-directional) of SmartFlow traffic at various packet sizes (64 byte, 440 byte and 1514 byte) is then passed across the infrastructure switches and through the device and the packet loss and average latency are recorded.

Test 5.2 - User Response Times

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

- **Test 5.2.1 - Web Response With No Background Traffic:** *The Avalanche and Reflector are configured to generate HTTP traffic through the IPS 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 115,000 packets per second). The minimum, maximum and average page response times and number of failed connections are recorded by Avalanche to provide an indication of the expected response times under normal traffic conditions.*
- **Test 5.2.2 - Web Response When Under Attack:** *The Avalanche and Reflector are configured to generate HTTP traffic through the IPS sensor as for Test 5.2.1. The Spirent WebSuite software is then used to generate DOS/DDOS traffic up to 10 per cent of the maximum rated bandwidth of the device under test (maximum 100Mbps). The minimum, maximum and average page response times and number of failed connections are recorded by Avalanche to provide an indication of the expected response times when the IPS device is under attack.*

Section 6 – Stability & Reliability

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

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

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

- **Test 6.1.2 - Passing Legitimate Traffic Under Extended Attack:** *This test is identical to 6.1.1, where we expose the external interface of the device to a constant stream of alerts over an extended period of time. The device is expected to remain operational and stable throughout this test, and to pass 100 per cent of legitimate traffic. Results are presented as a simple PASS/FAIL. If any legitimate traffic is blocked - caused by either the volume of traffic or the IPS device failing closed for any reason - this will result in a FAIL.*
- **Test 6.1.3 - ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC:** *This test attempts to stress the protocol stack of the device under test by exposing it to traffic from the ISIC test tool. The ISIC test tool host is connected directly to the external interface of the IPS sensor, and the ISIC target directly to the internal interface. ISIC traffic is transmitted through 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 operational and capable of detecting and blocking exploits throughout the test to attain a PASS.*

Section 7 – Management and Configuration

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

Test 7.1 - Management Port

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

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

- **Test 7.1.1 - Open ports:** *We will scan the open ports and active services on the management interface and report on known vulnerabilities.*
- **Test 7.1.2 - ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC:** *This test attempts to stress the protocol stack of the management interface of the device under test by exposing it to traffic from the ISIC test tool. The ISIC test tool host is connected directly to the management interface of the IPS sensor, and that interface is also the target. ISIC traffic is transmitted to the management interface of the IPS device (without passing through any other network equipment) and the effects noted. Traffic load is a maximum of 350Mbps and 60,000 packets per second (average packet size is 690 bytes). Results are presented as a simple PASS/FAIL - the device is expected to remain (a) operational and capable of detecting and blocking exploits, and (b) capable of communicating in both directions with the management server/console throughout the test to attain a PASS. We also note whether the ISIC attacks themselves are detected by the sensor even though targeted at the management port.*

NetScreen-IDP 500 V3.0 Test Results

Section 1 - Detection Engine

Test 1.1 – Attack Recognition	Attacks	Default ARRD	Default ARRB	Custom ARRD	Custom ARRB
Test 1.1.1 - Backdoors	5	5 ¹	5 ¹	5 ¹	5 ¹
Test 1.1.2 - DNS	2	2	2	2	2
Test 1.1.3 - DOS	12	12 ²	12 ²	12 ²	12 ²
Test 1.1.4 - False negatives (modified exploits)	13	10	10	13	13
Test 1.1.5 - Finger	4	4	4	4	4
Test 1.1.6 - FTP	5	5	5	5	5
Test 1.1.7 - HTTP	40	40	40	40	40
Test 1.1.8 - ICMP	2	2	2	2	2
Test 1.1.9 - Reconnaissance	8	7	7	8	8
Test 1.1.10 - RPC	3	3	3	3	3
Test 1.1.11 - SSH	1	1	1	1	1
Test 1.1.12 - Telnet	1	1	1	1	1
Test 1.1.13 - Database	1	1	1	1	1
Test 1.1.14 - Mail	1	1	1	1	1
Total	98	94 / 98	94 / 98	98 / 98	98 / 98

Test 1.2 – Resistance to False Positives	Pass/Fail
Test 1.2.1 - Audiogalaxy FTP traffic	PASS
Test 1.2.2 - MDAC heap overflow using GET instead of POST	PASS
Test 1.2.3 - Retrieval of Web page containing "suspicious" URLs	PASS
Test 1.2.4 - Simple SMTP QUIT command	PASS
Test 1.2.5 - Normal NetBIOS copy of "suspicious" files	PASS
Test 1.2.6 - Normal NetBIOS traffic	PASS
Test 1.2.7 - POP3 e-mail containing "suspicious" URLs	PASS
Test 1.2.8 - POP3 e-mail with "suspicious" DLL attachment	PASS
Test 1.2.9 - POP3 e-mail with "suspicious" Web page attachment	PASS
Test 1.2.10 - SMTP e-mail transfer containing "suspicious" URLs	PASS
Test 1.2.11 - SMTP e-mail transfer with "suspicious" DLL attachment	PASS
Test 1.2.12 - SMTP e-mail transfer with "suspicious" Web page attachment	PASS
Test 1.2.13 - SNMP V3 packet with invalid request ID	PASS
Test 1.2.14 - Fake DNS /bin/sh buffer overflow	PASS ³
Test 1.2.15 - Inter-firewall communication traffic (looks like FW-1 RDP header firewall bypass)	PASS
Test 1.2.16 - Fake SQL Slammer traffic	PASS
Test 1.2.17 - File copy of GIF file (contains bytes which look like MS Blaster NOP sled)	PASS
Total Passed	17 / 17

Section 2 - IPS Evasion

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

Test 2.2 – Packet Fragmentation/Stream Segmentation	Detected?	Decoded?	Blocked?
Test 2.2.1 - IP fragmentation - ordered 8 byte fragments	YES	YES	YES
Test 2.2.2 - IP fragmentation - ordered 24 byte fragments	YES	YES	YES
Test 2.2.3 - IP fragmentation - out of order 8 byte fragments	YES	YES	YES
Test 2.2.4 - IP fragmentation - ordered 8 byte fragments, duplicate last packet	YES	YES	YES

Test 2.2.5 - IP fragmentation - out of order 8 byte fragments, duplicate last packet	YES	YES	YES
Test 2.2.6 - IP fragmentation - ordered 8 byte fragments, reorder fragments in reverse	YES	YES	YES
Test 2.2.7 - IP fragmentation - ordered 16 byte fragments, fragment overlap (favour new)	YES	YES	YES
Test 2.2.8 - IP fragmentation - ordered 16 byte fragments, fragment overlap (favour old)	YES	YES	YES
Test 2.2.9 - TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with invalid TCP checksums	YES	YES	YES
Test 2.2.10 - TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with null TCP control flags	YES	NO	YES
Test 2.2.11 - TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with requests to resync sequence numbers mid-stream	YES	YES	YES
Test 2.2.12 - TCP segmentation - ordered 1 byte segments, duplicate last packet	YES	YES	YES
Test 2.2.13 - TCP segmentation - ordered 2 byte segments, segment overlap (favour new)	YES	YES	YES
Test 2.2.14 - TCP segmentation - ordered 1 byte segments, interleaved duplicate segments with out-of-window sequence numbers	YES	YES	YES
Test 2.2.15 - TCP segmentation - out of order 1 byte segments	YES	YES	YES
Test 2.2.16 - TCP segmentation - out of order 1 byte segments, interleaved duplicate segments with faked retransmits	YES	YES	YES
Test 2.2.17 - TCP segmentation - ordered 1 byte segments, segment overlap (favour new)	YES	YES	YES
Test 2.2.18 - TCP segmentation - out of order 1 byte segments, PAWS elimination (interleaved dup segments with older TCP timestamp options)	YES	YES	YES
Test 2.2.19 - IP fragmentation - out of order 8 byte fragments, interleaved duplicate packets scheduled for later delivery	YES	YES	YES
Total	19 / 19	18 / 19	19 / 19

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

Test 2.4 – Miscellaneous Obfuscation Techniques	Detected?	Decoded?	Blocked?
Test 2.4.1 - Altering default ports	YES ³	YES ³	YES ³
Test 2.4.2 - Inserting spaces in FTP command lines	YES ³	YES ³	YES ³
Test 2.4.3 - Inserting non-text Telnet opcodes in FTP data stream	YES ³	YES ³	YES ³
Test 2.4.4 - Polymorphic mutation (ADMmutate)	YES	YES	YES
Test 2.4.5 - Altering protocol and RPC PROC numbers	YES	YES	YES
Test 2.4.6 - RPC record fragging	YES	YES	YES
Total	6 / 6	6 / 6	6 / 6

Section 3 - Stateful Operation

Test 3.1 – Stateless Attack Replay	Alert?	Blocked?	Pass/Fail
Test 3.1.1 - Stateless Web exploits	NO	NO ⁴	PASS
Test 3.1.2 - Stateless FTP exploits	NO	NO ⁴	PASS

Test 3.2 – Simultaneous Open Connections (default settings)							
Number of open connections	10,000	25,000	50,000	100,000	250,000	500,000	1,000,000
Test 3.2.1 - Attack Detection	YES	YES	YES	YES	NO ⁵	NO ⁵	NO ⁵
Test 3.2.2 - Attack Blocking	YES	YES	YES	YES	YES ⁵	YES ⁵	YES ⁵
Test 3.2.3 - State Preservation	YES	YES	YES	YES	YES ⁵	YES ⁵	YES ⁵
Test 3.2.4 - Block legitimate traffic	NO	NO	NO	NO	YES ⁵	YES ⁵	YES ⁵

Test 3.3 – Simultaneous Open Connections (after tuning)							
Number of open connections	10,000	25,000	50,000	100,000	250,000	500,000	1,000,000
Test 3.3.1 - Attack Detection	YES	YES	YES	YES	YES	YES	NO ⁵
Test 3.3.2 - Attack Blocking	YES	YES	YES	YES	YES	YES	YES ⁵
Test 3.3.3 - State Preservation	YES	YES	YES	YES	YES	YES	YES ⁵
Test 3.3.4 - Block legitimate traffic	NO	NO	NO	NO	NO	NO	YES ⁵

Section 4 - Detection/Blocking Performance Under Load

Test 4.1 – UDP traffic to random valid ports			125Mbps	250Mbps	375Mbps	500Mbps	Max
Test 4.1.1 - 64 byte packet test - max 740,000pps (500Mbps)	Detected		100%	100%	76% ⁶	55% ⁶	
	Blocked		100%	100%	100%	100%	300Mbps
Test 4.1.2 - 440 byte packet test - max 130,000pps (500Mbps)	Detected		100%	100%	100%	100%	
	Blocked		100%	100%	100%	100%	500Mbps
Test 4.1.3 - 1514 byte packet test - max 40,860pps (500Mbps)	Detected		100%	100%	100%	100%	
	Blocked		100%	100%	100%	100%	500Mbps

Test 4.2 – HTTP “maximum stress” traffic with no transaction delays			125Mbps	250Mbps	375Mbps	500Mbps	Max
Test 4.2.1 - Max 1250 connections per second - ave packet size 1200 bytes - max 50,000 packets per second	Detected		100%	100%	100%	100%	
	Blocked		100%	100%	100%	100%	500Mbps
Test 4.2.2 - Max 2500 connections per second - ave packet size 540 bytes - max 115,000 packets per second	Detected		100%	100%	100%	100%	
	Blocked		100%	100%	100%	100%	500Mbps
Test 4.2.3 - Max 5000 connections per second - ave packet size 440 bytes - max 140,000 packets per second	Detected		100%	100%	100%	100% ⁷	
	Blocked		100%	100%	100%	100% ⁷	500Mbps
Test 4.2.4 - Max 10000 connections per second - ave packet size 350 bytes - max 180,000 packets per second	Detected		100%	100%	100% ⁷	N/A ⁸	
	Blocked		100%	100%	100% ⁷	N/A ⁸	375Mbps

Test 4.3 – HTTP “maximum stress” traffic with transaction delays			125Mbps	250Mbps	375Mbps	500Mbps	Max
Test 4.3.1 - Max 2500 connections per second - ave packet size 540 bytes - max 115,000 packets per second - 10 sec delay - max 25,000 open connections	Detected		100%	100%	100%	100%	
	Blocked		100%	100%	100%	100%	500Mbps
Test 4.3.2 - Max 5000 connections per second - ave packet size 440 bytes - max 50,000 packets per second - 10 sec delay - max 50,000 open connections	Detected		100%	100%	100%	100% ⁷	
	Blocked		100%	100%	100%	100% ⁷	500Mbps

Test 4.4 – Protocol mix			125Mbps	250Mbps	375Mbps	500Mbps	Max
Test 4.4.1 - 72% HTTP (540 byte packets) + 20% FTP + 4% UDP (256 byte packets). Max 1900 connections per second - ave packet size 555 bytes - max 105,000 packets per second - max 1750 open connections	Detected		100%	100%	100%	100%	
	Blocked		100%	100%	100%	100%	500Mbps

Test 4.5 – Real World traffic			250Mbps	500Mbps	750Mbps	1Gbps	Max
Test 4.5.1 - Pure HTTP (simulated browsing session on NSS Web site). Max 50 connections per second - 13 new users per second - ave packet size 1000 bytes - max 55,000 packets per second	Detected		100%	100%	100%	100%	
	Blocked		100%	100%	100%	100%	500Mbps
Test 4.5.2 - Protocol mix - 72% HTTP (simulated browsing sessions as 2.5.1) + 20% FTP + 4% UDP (256 byte packets). Max 275 connections per second - ave packet size 900 bytes - max 67,500 packets per second	Detected		100%	100%	100%	100%	
	Blocked		100%	100%	100%	100%	500Mbps

Section 5 - Latency & User Response Times

Test 5.1 – Latency	Packet Size	125Mbps	250Mbps	375Mbps	500Mbps
Test 5.1.1 Average latency (µs) with no background traffic	64	223.659	237.992	23589.34 ⁶	23641.32 ⁶
	440	271.738	274.932	277.685	280.788
	1514	306.724	317.810	330.193	340.921
Test 5.1.2 Average latency (µs) with background traffic (250Mbps HTTP traffic, max 1250 connections per second - ave packet size 540 bytes - max 60,000 packets per second) ¹⁰	64	373.561			
	440	181.537			
	1514	206.021			
Test 5.1.3 Average latency (µs) when under attack (50Mbps SYN flood) ¹⁰	64	88530.78 ⁹			
	440	59526.55 ⁹			
	1514	55297.51 ⁹			

Test 5.2 – User Response Times	Attempted Trans	Failed Trans	Min Page Response	Max Page Response	Ave Page Response
Test 5.2.1 - Web page response (ms) with no background traffic (250Mbps HTTP traffic, max 1250 connections per sec - ave packet size 540 bytes - max 60,000 packets per sec) ¹⁰	762996	0	<1	1.025	<1
Test 5.2.2 - Web page response (ms) when under attack (250Mbps HTTP traffic, max 1250 connections per sec - ave packet size 540 bytes - max 60,000 packets per sec PLUS 50Mbps SYN flood) ⁹⁺¹⁰	762986	498476	1.051	1.960	1.616

Section 6 - Stability & Reliability

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

Section 7 - Management Interface

Test ID	Result
Test 7.1.1 - Open Ports	PASS
Test 7.1.2 - ISIC/ESIC/TCPSIC/UDPSIC/ICMPSIC	PASS
Test 7.1.3 - ISIC attacks detected against management interface?	FAIL

Notes:

1. Some Backdoor traffic was detected on “default” ports only - BackOrifice was detected on any port
2. Note that it is necessary to specifically enable fragment detection in order to detect exploits such as Teardrop, Ping of Death, etc.
3. Following signature/code update
4. It is possible to alert/block mid-flow streams by enabling a signature
5. Maximum supported sessions is 100,000 out of the box. Maximum supported sessions is 500,000 after tuning (NetScreen recommended maximum is 220,000)
6. Device is capable of processing a maximum of approx 450,000 packets per second. Above this, we noted packet loss and high latency - the device continued to block (though not detect/alert on) 100% of all attack packets
7. Approaching HTTP connections per second limit (4500cps in test 4.2.3, 7500cps in test 4.2.4). A small number of connections failed during the test, though the test remained completely stable and the device continued to detect/block 100% of exploits
8. HTTP connections per second limit reached/exceeded. It was not possible to approach 10,000cps at 500Mbps and thus not possible to run this test (though the device continued to block 100% of exploits)
9. SYN flood caused significant packet loss and dropped connections in the Avalanche HTTP traffic.
10. For tests 5.1.2, 5.1.3, 5.2.1 and 5.2.2 the background traffic is **half** the normal loads for our 1Gbps test since this device is rated at 500Mbps

Section 1: Detection Engine

We installed one NetScreen-IDP 500 sensor with the latest signature set.

The Policy was created from scratch, switching on all Attack Objects from *Critical* down to *Low* (omitting only *Informational* signatures, which are more suited to auditing purposes), as well as turning on *Backdoor Protection* and the *SYN-Protector*. We also enabled fragmented packet detection and flow error detection in order to identify certain DOS and ICMP attacks.

Signature recognition (with blocking disabled) was excellent out of the box at 96 per cent and the blocking performance was identical. The new compound signature capability in particular has enabled far wider coverage than was possible with previous versions. This figure was increased to 100 per cent following a signature pack update which NetScreen provided within 48 hours.

Most of our “false negative” (modified exploit) cases were detected correctly out of the box, demonstrating that the NetScreen signatures are generally designed to detect the underlying vulnerability rather than a specific exploit. Once again, following the signature pack update, the IDP-500 detected 100 per cent of these test cases.

With certain protocols (DNS and FTP in particular) we noted an occasional tendency for the NetScreen IDP-500 to raise several alerts for a single exploit. This is generally as a result of exploits which trigger both protocol errors and pattern matches within the same attack. Whilst we could not fault the accuracy of those alerts, it would be nice to see them correlated or aggregated somehow in the GUI, allowing the administrator to focus on a single “incident” and drill down to see the multiple alerts beneath when required.

Resistance to false positives was generally very good (a clean sweep was achieved in our tests following a signature update), though we would recommend a period of running in passive non-blocking mode to ensure that the incidence of false positives is at a minimum before switching to active blocking mode.

Section 2: IPS Evasion

The NetScreen IDP-500 performed impeccably in all of our evasion tests, being one of the few products to collect a clean sheet across the board in this area. *Fragroute*, *Whisker*, *ADMmutate* and even *RPC record fragging* all failed to trick the IDP-500 into ignoring valid attacks.

Note that not only were the fragmented and obfuscated attacks blocked successfully, but almost every one of them - a single *fragroute* attack being the sole exception - was decoded accurately as well.

Section 3: Stateful Operation

The IDP-500 was not tricked into alerting on our stateless exploits, indicating that it would be resistant to TCP-based exploits launched via replay tools such as *Stick* and *Snot*. Blocking of mid-flow streams, whilst disabled out of the box, is possible by enabling a single signature if required.

Out of the box, the device maintained state on up to 100,000 open connections, successfully detecting our half-open exploit as it was completed. It also continued to detect and block new exploits as we maintained 100,000 open connections, and no legitimate traffic was blocked during this stage of these tests.

With a simple change to the flow-related parameters in the GUI, it was possible to maintain state across 500,000 open connections (though NetScreen actually recommends 220,000 as the maximum on this device). The IDP-500 is not capable of supporting in excess of 500,000 connections, and thus it was not possible to run our 1 million connection test.

If the configured maximum number of sessions is exceeded, the sensor will refuse new connections, meaning that legitimate traffic is blocked and new attacks are not alerted. However, state is maintained completely on existing flows and no attack traffic is ever allowed through the device.

Section 4: Detection/Blocking Performance Under Load

Note that the NetScreen-IDP 500 was tested as a 500Mbps device. The IDP-500 performed well in our performance tests, demonstrating that it could handle most traffic loads up to the rated speed of 500Mbps, successfully detecting and blocking 100 per cent of all our exploits.

Under extreme traffic loads - using 64 byte UDP packets or 10,000 HTTP connections per second, for example - performance did degrade slightly as we reached some fairly hard limits - 450,000 64-byte packets per second and 7500 HTTP connections per second. Beyond these limits (unlikely to be found in most "normal" real-world networks, especially sub-Gigabit networks for which this device has been designed) there was some packet loss and/or significant increase in latency leading to failed HTTP connections. However, even when the device was placed under such extreme loads and started to drop packets it continued to block 100 per cent of all exploits successfully. Note that once the device begins to drop packets, there is the possibility that it will block legitimate traffic.

Providing traffic loads are kept within the bounds indicated in these tests (i.e. in particular "normal" traffic with less than 7,500 HTTP connections per second and 450,000 packets per second), then we would have no hesitation in rating this device at 500Mbps.

Section 5: Latency & User Response Times

Being a software-based IPS product running on an off-the-shelf hardware platform, latency is higher than you would expect to see in a pure networking device such as a switch. However, under most normal traffic conditions, latency is well within acceptable limits for a 500Mbps network.

Behaviour throughout the tests was very predictable, increasing as both average packet size and traffic load increased and ranging from 224µs with 125Mbps of 64 byte packets to 341µs with 500Mbps of 1514 byte packets. With our "average" packet size of 440 bytes in particular, latency remained low and increased in a very predictable manner from 272µs at 125Mbps to 281µs at 500Mbps. The only situation where latency increased significantly beyond the norm was when we loaded the device with more than 450,000 64-byte packets per second - not a normal occurrence in a sub-Gigabit network.

When the device was loaded with 250Mbps of background traffic (half the rated speed of the product), there was a slight increase in latency of 64 byte packets, and a slight reduction in latency of 440 byte and 1514 byte packets. This demonstrates that high loads of normal network traffic would not have a serious negative impact on overall latency figures.

With the SYN-Protector enabled in “*Relay*” mode, the degradation when under a 50Mbps SYN flood attack was much more significant. Although the SYN flood was effectively mitigated, the latency was increased to the point where we noted a significant number of failed connections in our HTTP traffic. The IDP-500 is better as a straight IPS than as a DOS/DDOS mitigation device, and NetScreen has better DOS/DDOS mitigation capabilities built into its firewall product line.

To balance this, however, it should be noted that with the SYN-Protector disabled, **all** of our HTTP connections failed when under DOS attack, proving the effectiveness of the SYN-Protector despite the latency issues.

Section 6: Stability & Reliability

The NetScreen IDP-500 performed consistently and reliably throughout our tests. Under extended attack it continued to pass legitimate traffic whilst blocking attack traffic in a consistent manner.

Exposing the sensor interface to an extended run of ISIC-generated traffic produced a significant number of useful protocol anomaly alerts, and the device continued to detect and block all other exploits throughout and following the ISIC attack.

Section 7: Management Interface

Open ports on the management interface are restricted to SSH and HTTPS in order to provide management access.

The extended ISIC attack against the management interface had no effect on the appliance or its ability to detect and block attacks and communicate with the management server. However, no alerts were raised at any time to inform the administrator that the management port was under attack.

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