



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Maintaining a secure network connected to the Internet is becoming more difficult as time goes on. New viruses are released daily, higher machine speeds and more sophisticated and automated tools mean that hackers can scan and attack wide sections of the Internet at a time. More consumer broadband connections mean that there is an increasing source of relatively vulnerable high performance computers available to be harvested as “bots” and used as proxies for illicit activity. The increasing complexity of current operating systems also means more vulnerabilities available for exploitation.

Faced with this increasing threat from the Internet, virtually all connected organizations are now actively engaged in some level of IT security activities. According to the 2003 FBI/CSI Computer Crime and Security Survey, 99% of organizations use virus protection, and 98% use firewalls.¹ The use of more sophisticated tools, such as Network Intrusion Detections (NID) Systems increased from 60% in 2002 to 73% in 2003² as organizations look for additional assurance their systems have not been compromised. At present, large corporations tend to be the biggest users of these technologies³, but the CSI survey indicates that many medium businesses are also adopting this technology.

Following the example of firewalls and virus protection, where large corporations first adopted the technology and it then “trickled down” to smaller organizations, one would expect to see an increasing percentage of small organizations implementing NID Systems. Here, however, the situation may be different. Small businesses operate with different efficiencies, different resource distributions, and different investment criteria than their larger cousins, and the question of whether or not it makes sense to implement a Network Intrusion Detection system for these organizations is not necessarily settled. This paper examines the issues involved in the deployment of a NIDS from the perspective of the small organization.

For the purposes of this discussion, a small business will be defined as having between 20-99 employees and is not run out of a home. There are over 670,000 of these organizations in the United States, and they employ more than 20 million people.⁴ These organizations almost universally have one or more computer available for business purposes⁵, and over 85% have Internet access.⁶ In addition to small retail or wholesale businesses, organizations in this size range include engineering consulting firms, health care practices, law offices and other organizations where technology may be a central component of their operations without being the focus of their activities. However, their networks will almost certainly contain their “crown jewels”, whether they are proprietary or client confidential plans, patient information, or supplier and customer information. Despite their size, these organizations have significant security concerns and a legitimate interest in keeping their business information and internal networks secure.

The Network Intrusion System discussed here be primarily of the turnkey device or single box model. A device of this type usually sits at a network boundary, typically at the internal/external network boundary, and usually gathering data from inside or outside the firewall or both. A device of this type usually has one or more sensors, which are essentially NIC cards, then some kind of processing, storage and reporting component. Conceptually, these systems operate in one of two ways: Misuse Detection and Anomaly Detection.⁷ In Misuse detection, the system compares the gathered signals to a database of signatures of malicious activity types, and any signal that matches one of the signatures is flagged as an alert. In anomaly detection, a baseline profile of typical, legitimate network traffic is developed and established. The network traffic from the sensors is then compared to this baseline, and any identified departures are reported as alerts.

In practice, most commercial devices tend to be primarily signature based like virus detection systems so they need periodic updates of these signature to detect the most current threats. In addition to generating alerts, most types also store at least the header information for any packets generating alerts for further analysis by an intrusion detection analyst. Another feature, called Active Response, that many NID systems offer is the ability to automatically react to detected alerts. The responses try to block or disconnect the connection generating the alerts to quickly and automatically protect the network from the threat.

So, what are the reasons a small business would consider deploying a NIDS? Small businesses are more dependent on computers and the Internet than before and the threat to business computers is changing - the old wisdom that most attacks come from the inside is wrong. Most attacks now come from the Internet, and the threat from the Internet is increasing every year. The number of attacks in the first half of 2003 is 19% higher than the same period in 2002.⁸ Further, as large and medium businesses implement more sophisticated Internet defenses, it may have the effect of focusing attention on smaller businesses as hackers look for targets with a higher probabilities of success. Clearly, as small businesses use the Internet more and the threat from Internet attacks increases, the risk increases. To help them mitigate this risk, they will find much of the attention of influential people and organizations in the IT industry is focused on deploying IDS systems.

At present, it would be difficult to read about Information Technology (IT) or IT security without encountering a wide array of advice in print and online recommending or assuming your organization has deployed a NIDS. It easy and perhaps necessary to be influenced by these sources because they are a valuable source of information and analysis not otherwise available to a small IT department. People running or working in small business IT departments do wide variety of things in a day; everything from thinking about the future development of the network to replacing a bad fan in a workstation.

They don't have the time to research every new idea for running their networks, and they usually don't have a test lab. So they depend on published information to help guide policy and make decisions.

In the case of NIDS, the advice is universally in favor of deployment. Starting the parade is the National Institute of Standards and Technology (NIST) "...Intrusion detection systems have become a necessary addition to the security infrastructure of most organizations."⁹ In *Network Intrusion Detection (Third Ed.)*, in the Organizational Issues section, pg. 321 "The seven most important things to do if Security matters:" Item number seven is "Implement intrusion detection and incident response."¹⁰ The IT press and the vendors of these systems of course recommend deployment, especially of their products, but even taking that into consideration, there was enough coverage of IDS systems to attain "buzzword" status in the 2003 CSI Computer Crime and Security Survey.¹¹ The lone voice of dissent was Gartner Consulting's 2003 press release¹² that was critical of the current state of IDS technology. The firestorm of responses to this article affirmed the status of IDS systems as a darling of the IT security community. Furthermore, this article was critical only of the state of current NIDS technology and called for building additional capability, including much of NID functionality, into firewalls. In essence, this article was a statement in favor of deploying better NID-like technology rather than giving up the idea entirely. Given this, it appears that virtually the entire industry supports deployment of NIDS systems or at least something like it.

So we know they're recommended by knowledgeable people in the industry, but what do they actually do to improve the security of organizations? In a nutshell, they are the eyes of a network. As defined above, NIDS systems capture and analyze traffic across some network boundary. Assuming a typical Internet/network boundary that a small organization might have with a firewall, sensors will typically be placed on the WAN (outside) and the LAN (inside) side of the firewall. These will log data on every signal back to the monitoring station.

With the sensors placed at these points, it becomes possible to observe, analyze and document traffic traveling into and out of the network. With sensors in these positions, a number of analyses become possible:

- The data from the outside sensor can be analyzed to provide information on the type, frequency, source and the target of reconnaissance scans and attacks. This information can then be used to identify specific scans, specific attacks, specific targets, and to an extent specific sources of malicious signals coming at the internal network.
- The NIDS will show breaches of the firewall. The classic sign of this is a questionable signal showing up both in the outside and inside sensors. When this happens, and there is no established session from within the LAN, it's time to have a look at the firewall rules to see why this is happening.

- An NIDS can identify compromised machines inside the network. For example, if it is configured to capture SMTP outbound traffic, any virus like Sober.G that carries it's own SMTP engine will show up immediately as a problem, when huge numbers of SMTP outbound messages show up on the alert listing. Similarly, a machine with the MS Blaster worm would have also shown up in the NIDS logs with masses of port 1433 traffic. An NIDS is perhaps the most effective tool available to quickly identify and react to these situations.
- It is the only way an analyst can identify attacks and scans that don't match a predefined signature. By analyzing the logs of traffic, usually on the outside interface, it is possible to identify patterns showing new scans and attacks that are not captured by the NIDS signature library.
- It can provide records of network traffic for forensic analysis. When the worst happens and a network compromise occurs, the NIDS logs may be able to provide information how the compromise was engineered, and what activities were conducted during the time the network was compromised.

All of these above analyses are different parts of the same idea. As the "eyes" of the network, it makes observation and recording of network traffic possible. If analysis resources are added, it makes it possible to answer many questions about the signal environment outside the firewall, the effectiveness of the firewall, and the kinds and volume of traffic flowing through the network.

The remaining operational questions concern the usefulness and effectiveness of NIDS systems. How well do they do what they're supposed to do? The current conventional wisdom is that IDS systems generate a lot of false positives, and they aren't particularly accurate. Older tests of IDS systems have generally supported this view. A 1999 study of two products, Internet Security Systems (ISS) RealSecure 3.0 and Cisco NetRanger 2.1.2 conducted by IBM Zurich, found that one product detected 30 out of 42 attacks (71%), and the other a rather mediocre 18 of 30 (56%) of attacks that these products were supposed to detect according to the documentation.¹³ One product generated 8000 alerts in a month of real world testing, of which at least half were false positives.¹⁴

More recent studies have shown better results. In a 2002 NSS IDS test, six IDS devices were tested that detected between 61.5% and 86.2% of 109 total attacks.¹⁵ Information on the false positive rate of these devices is not available. In a 2003 release of similar tests (the NSS Group Ltd. group test edition 4)¹⁶ four 100Mbps maximum IDS systems from Cisco, Internet Security Systems (ISS), NFR Security Inc., and the open source Snort system were extensively tested. This time, considerably better results were obtained. Each of the four NIDS systems was tested against 82 simulated attack types, both with default settings, and with custom signatures (and also made available to the public). To measure resistance to false positive alerts, NSS subjected each system to 14 tests using legitimate traffic with various alert triggering characteristics. The following table summarizes the these two metrics for the four IDS systems.

IDS System	Default Detection Tests Passed	Default Detection Percentage Passed	Custom Detection Tests Passed	Custom Detection Percentage Passed	False Positive Tests Passed	False Positive Percent Passed
Cisco Secure IDS	78/82	95%	78/82	95%	12/14	86%
ISS Proventia A201	72/82	87%	82/82	100%	14/14	100%
NFR NID-310 v3.2.1	52/82	63%	74/82	90%	13/14	93%
Snort 2.0	63/82	77%	N/A	N/A	10/14	71%

Source: The NSS Group Ltd.¹⁷

As can be seen from this table, even compared to the next most recent NSS study, the systems have improved, capturing 63-95% of attack types versus 61.5-86.2%. Comparing the false positive data is not as clear, but the test results indicated this is a priority for improvement. Based on this upward trend, it does seem that at least these four NIDS system manufacturers/developers are significantly improving their products year by year, and that at least the accuracy portion of the conventional wisdom is probably due for a review.

Even with these improvements, there remain two significant shortcomings about NIDS systems that need to be kept in mind when considering what they do. The first is that all the tests above are conducted with known attacks and with signatures designed to detect these attacks. These test do not measure system performance when confronted with an unknown attack, especially if that attack is unlike anything stored in the attack signature library. The second point is that although these four systems (and probably other manufacturers as well) appear to be making progress reducing false positives, it's not clear that these have been reduced to an acceptable level.

Regarding the first point, the way virtually all available NID systems work is similar to virus detection systems, essentially pattern matching network signals to pre-defined attack signatures stored in a library. This means that, like virus detection systems, they are vulnerable to "zero day" effects where the attack is alive and proliferating before the corresponding attack signature is released and downloaded. In this situation the NIDS system may be unable to detect it unless there are people monitoring the system and analyzing the IDS data as it is logged. Even then, the detection process will be considerably longer than via the signature library.

This is a rather significant point. Information from the 2003 Symantec Internet security report shows that the newest vulnerabilities are much more likely to be attacked than old vulnerabilities. More than a third of new attacks target vulnerabilities less than six months old, and nearly two thirds target vulnerabilities less than a year old.¹⁸ In this dynamic situation, a signature-based NIDS is going to require significant analysis resources to detect most threats.

The second point is that even though the false positive rate has probably been substantially reduced, the large volumes of traffic that can flow across even a small business sized Internet connection makes small false positive rates troublesome. If we assume a system is operating with average packet sizes on a 20% utilized T-1 line it will be transferring about 353,000 packets an hour.¹⁹ Even if the NIDS systems are 99.999% correct (less than one false positive per hundred thousand packets) this will result in 3.5 false positives per hour. This could be easily handled by IDS monitoring staff, but it would be a problem if the IDS system was an automatic responding system. As mentioned above, these systems will react in some way when it decides a packet or connection is an intrusion attempt. Usually these systems react by attempting to disconnect the session or by blocking or shunning the detected connection or IP address for a specified period of time.

On paper, this is a great thing. It would be especially attractive for a small business or organization because it implies that the box can be deployed and subsequently won't need a lot of management beyond making sure it's up and running and updating its attack signature files every now and then. Unfortunately this is probably exactly how not to deploy an IDS because sooner or later (probably sooner) the machine is going to start cutting off or blocking legitimate connections. This frustrating and irritating to users and management alike and it will result in either loosening the rule set or getting rid of the device. Either response means alerts that should get investigated won't be and the investment in the system is diminished or lost.

In contrast to the above scenario, virtually all sources indicate that the effective way to operate an NID system is to have dedicated staff analyzing the logs on a daily basis. If the NID system is to be the eyes of the network, its most important function is producing log files of traffic at various sensors. Daily analysis of the logs is important to avoid having already large volumes of log data get completely out of hand, and to be able to quickly react in the event something does come up. Following up on a potential security breach is vastly more useful and less costly to contain if the alert is minutes old rather than days old. If it's minutes old, you can probably prevent most of the damage on the spot. If it's weeks old, then the situation becomes an expensive and difficult forensic investigation.

The quality of the analysis is also important to NID, so while it's tempting to think that the logs can be examined by a network or system administrator "in their spare time", it's unlikely that this approach will be successful. Interpreting the logs is a specialized skill independent of other network skills qualified IT staff may have. IDS monitoring skills are acquired as much by experience as by training. They will stay sharp through continuous use and decline otherwise. An old high school math teacher of mine used to say "Math is not a spectator sport". Intrusion detection is the same way: it requires constant practice and exercise to develop and stay sharp.

It is also unlikely that one could properly process a week's worth of log data in addition to all the other activities a Network admin must attend to.

For many of the same reasons, an even better approach would be to have continuous IDS monitoring and analysis. Having people watching what's happening 7x24 has several advantages over having dedicated monitoring during business hours. Continuous monitoring means data will be processed and examined in real time rather than processing 168 hours of data in a 40 hour time span. It avoids the backlog of data accumulated over the weekend that needs to be evaluated before the current day's activity can be examined. It also means the response to events that happen to occur on a weekend will be real time and not several days old by the time its checked. Similarly, continuous monitoring avoids potentially having a sixteen hour delay detecting a successful attack if the logs are only monitored during business hours. In addition, continuous monitoring and the required larger staff this will take also generates an additional benefit by having a team of analysts available. The combination of skills, experience and viewpoint will increase the effectiveness of IDS deployment far beyond what is gained by the increased time monitoring the console.

Unfortunately, for a small business, employing a team of analysts working on a single component of a (hopefully) robust security infrastructure is almost certainly not a realistic deployment scenario. For an organization with fewer than 99 people. Adding even one person is a substantial investment, requiring careful thought and consideration. Adding a staff of at least 5 IT people, the minimum necessary to provide 7x24 monitoring, borders on the absurd. Fortunately, there are security service providers (SSPs) who do this for a living. These organizations generally have a staff of analysts, dedicated equipment, and the efficiencies of a larger IDS installation to provide 7x24 monitoring and response. These organizations can likely provide this service for considerably less money than the cost of hiring and managing a team of analysts.

The natural question here is, of course, "and how much would each of those wonderful ideas cost?" And the answer is that for a small business to install and run an NID system in-house, it's very expensive. In *Network Intrusion Detection, Third Edition*, Novak and Northcutt present a scenario that would spend \$15, 000 on hardware alone.²⁰ Another study, *Justifying the Expense of IDS, Part One: Calculating ROI for IDS*, by Kinn and Timm estimated that an in-house network IDS system would cost \$10,000, and a management station would cost about \$5000 for a total technology cost of \$15,000, identical to Northcutt and Novak's figures.²¹

Both of these estimates are consistent with costs of the four IDS devices reviewed by NSS and discussed above. These costs are shown in the small table below for the appliance with a single sensor comprising a minimal configuration most suitable for a small organization.

As can be seen from the table below, the average cost is \$12,665 for the three turnkey appliances. Adding the odds and ends required to actually get the machine into the rack and into production would probably add a few hundred to a thousand dollars, so a round \$15,000 provides a reasonable estimate for the hardware cost of these three devices.

NIDS Vendor and model	Cost
Cisco IDS-4235	\$12,500
ISS Proventia A201	\$9,995
NFR NID-310	\$15,500
Average	\$12,665

Source: The NSS Group Ltd.²²

The remaining system, Snort, is an open source product and so the software is free, but the machine to run it on won't be. Running Snort with some kind of front end (say Hogwash or ACID) to organize the data requires a substantial machine. It will need a decent processor, two reasonable NICs, at least a gigabyte of memory and 30-40GB of disk space for the database, and it should be a mirror or other fault tolerant configuration. It will also need a securely configured operating system. All of this needs to be carefully built, burned in, tested, hardened, and then validated. Add all this up and the machine will probably cost around \$7,500 ready to run.

This is about half the cost of the turnkey devices, but Snort has no simple mechanism for retrieving and installing updates or any other maintenance tasks, so there would be higher costs associated with updating and running Snort. Since this is not a comparison of these four products, the device estimated costs presented in the two sources (Northcutt and Novak, Kinn and Timm) of the device will be used (\$15,000).

In any case, the small differences in hardware costs pale in comparison to the much higher costs for analytical support. Northcutt and Novak's estimates the cost would be about \$85,000 per year for one analyst²³, which yields a total cost of about \$100,000 to get started with a minimum system monitoring the system during business hours only. To provide 7x24 monitoring on the same system would be five times that or \$425,000 per year. This should be considered a minimum estimate because it doesn't consider the extra management and administration costs that might be required by the addition of five analysts or the cost of maintaining the system.

Similarly, Kinn and Timm estimated a slightly lower recurring cost of \$75,000 for the analyst, but added a 15% of the technology cost per year to maintain the system.²⁴ Excluding additional management and administration, the initial year's investment would then be \$90,000, and \$77,250 a year thereafter. For the 7x24 continuous monitoring scenario, the costs go up to \$390,000 for the first year and \$377,250 after that.

On the other hand, Kinn and Timm presented costs for a Security Service Providers, which this study called MSSPs (Managed Security Service Providers) that were surprisingly different. In this case the cost of the technology remained the same at \$15,000, but the management of the NIDS would cost only \$24,000 per year.²⁵

In this scenario, again excluding management and administration costs, the initial year's total (hardware and management) costs would then be \$39,000, and thereafter the annual cost would be \$26,250 and this was for a 7x24 level of coverage. The following table summarizes the costs for each of these solutions and puts them side by side for easy comparison.

Source and type of coverage	Northcutt and Novak Business hours only	Northcutt and Novak Continuous Coverage (7x24)	Justifying the Expense of IDS Business hours only	Justifying the Expense of IDS Continuous Coverage (7x24)	Justifying the Expense of IDS MSSP Continuous Coverage (7x24)
Technology Cost	\$15,000	\$15,000	\$15,000	\$15,000	\$15,000
Labor Cost	\$85,000	\$425,000	\$75,000	\$375,000	\$24,000
Total First year	\$100,000	\$435,000	\$90,000	\$390,000	\$39,000
Total subsequent year	\$85,000	\$425,000	\$77,250	\$377,250	\$26,250

Sources: Network Intrusion Detection (3rd Edition) and Justifying the Expense of IDS, Part One: An Overview of ROIs in the IDS

From a cost perspective, this is a no-brainer. There is simply no doubt that the MSSP or Security Service Provider solution provides the most coverage at the best cost. Not only were the costs for the SSP continuous coverage solution were vastly better than the in-house cost estimates for continuous coverage, but they were only a third of the cost of the lowest estimated in-house costs for business hours only coverage. This result appears consistent through the sources of the numbers.

Both Northcutt and Novak and Kinn and Timm showed very similar labor costs for the management of the solution, and these management costs are reasonable given a consideration of the labor market. Similarly, both sources showed identical costs for the device or technology that would be required, and these costs were consistent with actual prices shown for devices that would likely be implemented. The costs that are unconfirmed with other sources are the Security Service Provider solution. Given that the other labor and device costs from this source were reasonable and consistent with other sources, it's reasonable to believe that these costs are as well.

The one remaining question left is should small businesses invest in this technology?

I believe the answer is, at the current state of technology, that they should not. The benefit NID systems provide is only available with high levels of skilled management. I believe that small businesses do not have financial resources to manage NID technology.

There is little doubt that NID systems are a beneficial security tool. They are the eyes of the network that can tell you what's being probed, what's being attacked and to some extent, who is doing it. However, they do so imperfectly and at great cost.

Because this vision is imperfect, the systems are still dependant on a team of analysts to make information out of the mounds of data an NID system will produce. This means that to be effective, an NID system require significant resources to run on an ongoing basis, and those costs are too high for most small businesses.

On average the small businesses discussed here have an annual IT budget of \$1.3 million dollars²⁶ and they spend about 8% of that on security²⁷, for a total IT security budget of \$104, 000. This would barely support the minimum "in-house" solutions as the only security deployment, and even the most cost effective SSP solution would absorb a third of this budget in the first year and a quarter of it in subsequent years.

Obviously, other projects would have be abandoned to make room in the budget for NID, and Defense in Depth suggests that it is more effective to undertake smaller, more diverse projects than to put all of one's security eggs in one basket. Examples of substitutes for an NIDS might be simpler host based IDS for the critical machines, a VPN for remote access, an automated tool for patch management, and a vulnerability scanner.

Small business efficiencies may make deployment of number of individual system components preferable to a larger, centralized system. Host based IDS is a possible example. For a small business running less than 120 machines, One could afford to put fairly low maintenance host based firewalls and intrusion detection such as Symantec Client Security on every machine for about the same cost as the cost of the NID system, and it would require considerably less management and analysis.

In summary, the NID system is still a maturing technology. Its state is such that it is effective and appropriate for organizations where it is more cost effective to watch a single connection with expensive technology than it is to watch several connections with cheaper technology. It is also for organizations who have enough security infrastructure to make NIDS an additional component of defense in depth rather than replacing several other technologies. The average small business has neither of these and most likely will be better off investing in other methods of achieving Internet security objectives while NID technology evolves.

-
- ¹ Richardson, Robert. "CSI/FBI 2003 Computer Crime and Security Survey". Eighth Edition. Computer Security Institute. Page 4.
- ² Ibid. page 5.
- ³ Bennett, Whitney. "Information Security." Key3Media Research. Needham MA. March 2002. Page 10.
- ⁴ "U.S. and all States, totals, 2001", United States Census Bureau.
URL: <http://www.census.gov/csd/susb/susb01.htm> (25 June 2004).
- ⁵ "Computer, Internet Use Increases at Small Businesses", 4 October 2001.
URL: http://www.clickz.com/stats/markets/smallbiz/article.php/10098_897771 (25 June 25 2004)
- ⁶ ibid.
- ⁷ "Intrusion Detection System." Small Business Computing Online Dictionary of IT Terms.
13 December 2002. URL: [HTTP://sbc.webopedia.com/term/I/intrusion_detection_system.html](http://sbc.webopedia.com/term/I/intrusion_detection_system.html). (11 May 2004) Page 1.
- ⁸ McCarthy, Linda. Ed. "Symantec Internet security report. Trends for January 1, 2003 - June 30, 2003 - Attack Trends." Symantec Corporation, Cupertino, CA, September 2003. Page 8.
- ⁹ Bace, Rebecca and Mell, Peter. "Intrusion Detection Systems." National Institute of Standards and Technology (NIST). SP800-31-2. Page 5.
- ¹⁰ Northcutt, Stephen, and Novak, Judy. "Network Intrusion Detection, Third Edition." Indianapolis: New Riders Publishing, 2002. Page 321.
- ¹¹ Richardson, Robert. "CSI/FBI 2003 Computer Crime and Security Survey". Eighth Edition. Computer Security Institute, page 7.
- ¹² "Gartner Information Security Hype Cycle Declares Intrusion Detection Systems a Market Failure-Money Slated for Intrusion Detection Should Be Invested in Firewalls." June 11, 2003.
URL: http://www.gartner.com/5_about/press_releases/pr22june2003c.jsp (25 June 2004).
- ¹³ McHugh, John, and Alan Christie and Julia Allen. "Defending Yourself: The Role of Intrusion Detection Systems." IEEE Software, September/October 2000. (2000) Page 49.
- ¹⁴ Ibid.
- ¹⁵ Kinn, David, and Timm, Kevin. "Justifying the Expense of IDS, Part Two: Calculating ROI for IDS." SecurityFocus. August 10, 2002. URL: <http://www.securityfocus.com/infocus/1621>. (27 May 2004).
- ¹⁶ "Intrusion Detection Systems-Group Test (Edition 4)." NSS Group Ltd. August 2003.
URL: <http://www.nss.co.uk/ids/edition4.htm> (14 June 2004).
- ¹⁷ Ibid. The specific URLs are: http://www.nss.co.uk/ids/edition4/cisco/cisco_results.htm,
http://www.nss.co.uk/ids/edition4/iss/iss_results.htm

http://www.nss.co.uk/ids/edition4/nfr/nfr_results.htm
http://www.nss.co.uk/ids/edition4/snort/snort_results.htm

- ¹⁸ McCarthy, Linda. Ed. "Symantec Internet security report. Trends for January 1, 2003 - June 30, 2003 - Executive Summary." Symantec Corporation. Cupertino, CA. September 2003. Page 2.
- ¹⁹ Calculation based on average packet size of 402.7 bytes. Source: "JTC 003 Mixed Packet Size Throughput." Agilent Technologies. 2002.
URL: http://advanced.comms.agilent.com/routertester/member/journal/JTC_003.html (29 June 2004)
- ²⁰ Northcutt, Stephen, and Novak, Judy. "Network Intrusion Detection, Third Edition." Indianapolis: New Riders Publishing, 2002. Page 361.
- ²¹ Kinn, David, and Timm, Kevin. "Justifying the Expense of IDS, Part One: Calculating ROI for IDS." SecurityFocus. 28 July 2002. URL: <http://www.securityfocus.com/infocus/1608>. (27 May 2004).
- ²² "Intrusion Detection Systems-Group Test (Edition 4)." NSS Group Ltd. August 2003.
URL: <http://www.nss.co.uk/ids/edition4/summary.htm>. (13 June 2004).
- ²³ Northcutt, Stephen, and Novak, Judy. "Network Intrusion Detection, Third Edition." Indianapolis: New Riders Publishing, 2002. Page 361.
- ²⁴ Kinn, David, and Timm, Kevin. "Justifying the Expense of IDS, Part One: Calculating ROI for IDS." SecurityFocus. 28 July 2002. URL: <http://www.securityfocus.com/infocus/1608>. (27 May 2004).
- ²⁵ Ibid.
- ²⁶ "Computer, Internet Use Increases at Small Businesses." 4 October 2001.
URL: http://www.clickz.com/stats/markets/smallbiz/article.php/10098_897771 (25 June 2004)
- ²⁷ "Security Spending: How much is enough?" CIO Research Reports. 12 September 12 2002.
URL: <http://www.cio.com/research/surveyreport.cfm?id=6>. (28 June 2004).

© SANS Institute 2004, Author retains full rights.