



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>

Case Study – Technical Refresh of a Commercial Intrusion Detection System
and the not so Witty Worm
Doug McLaren
GSEC Practical 1.4b option 2
April 19, 2004

Abstract:

This paper describes the more interesting parts of a project to rationalize the Intrusion Detection system (IDS) on an E-Commerce Gateway implemented in the heady days of the Dot.com boom. Those heady days have been replaced with a very cost conscious overhead cutting hangover for those Dot.com's that survived. Rather than describe the upgrade of the Internet Security Systems (ISS) Products from Real Secure 6.5 to 7.0 and the upgrade of Real Secure Work Group Manager 6.5 to Site Protector 2 Service Pack 3, both of which are far better documented on the ISS web site ([1] ISS), this paper describes the issues and opportunities that this upgrade allowed the small security team I am part of to encounter and resolve. This will be much more relevant to a larger section of the Security Community and will demonstrate the many different security related subjects as that a security specialist needs to be familiar with.

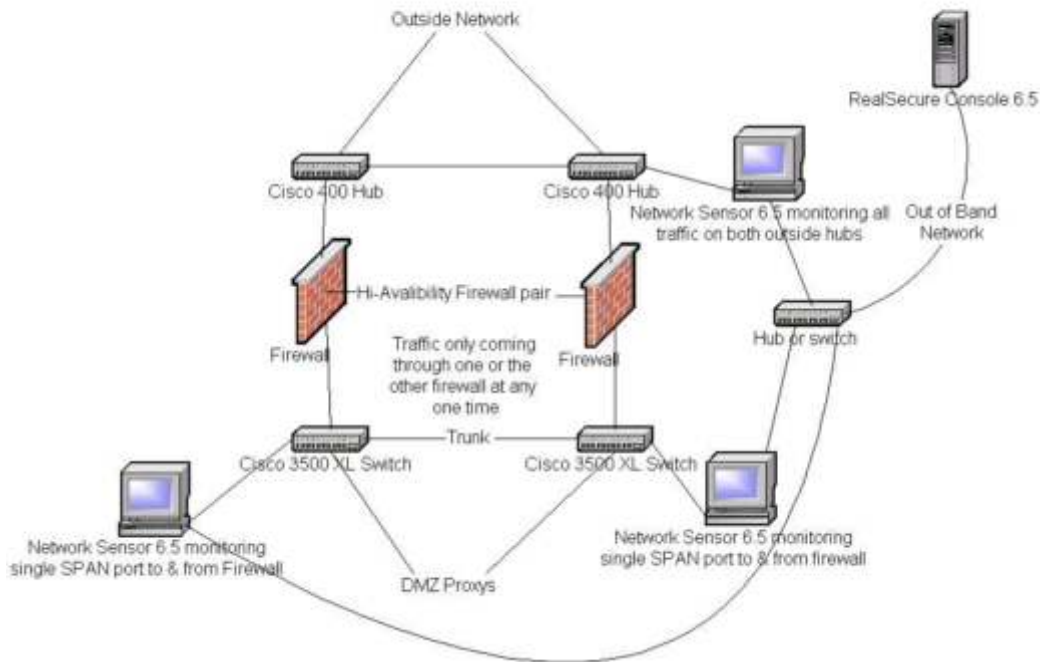
The system detailed in this paper has grown to comprise a large number of network based IDS (Real Secure Network Engines), host based IDS (Real Secure Server Sensors) and IDS central management systems (Real Secure Work Group Managers). Due to the budgetary constraints it was difficult to justify the upgrading of a large number of sensors spread over a number of different systems that were slowly becoming obsolete. It describes the way that improvements were made to systems without spending vast sums of money on the latest very expensive new products that the IT industry wishes to sell us and that we wish to buy. This paper describes the following Issues with IDS's and Cisco Switch SPAN/Monitoring, Support of Operating systems on Hardware and getting Windows2000 server to run foreground tasks more efficiently, A number of important functions of the Site Protector and Network Engine Version 7.0 systems which are not so well documented have also been added. As a Postscript the effects of the real attack of the Witty worm on a production system are described and its effect. There are many resources describing the Witty Worm in minute detail, I will not regurgitate them in this case study as only those facts pertaining to this case study have been described. A list of the security improvement that needs to be implemented to stop such worms in future has also been added.

Before:

The system being described in this paper is an e-commerce gateway implemented four years ago during the Dot.com boom. When the system was implemented there was plenty of money to implement fairly expensive systems. This is a scenario seen in many commercial systems, quotes such as "We shall build our systems from the best of breed products", "we do not allow open source products on our system (no Linux or Snort)", "Commercial products only", trip off the lips of senior managers. This has led to the enviable situation of having five separate environments, two Development, Test, Production and Pre-Production (a complete copy of production) and also the situation that each new addition to the production system needs to be added and justified also to the other four systems as well. In the example environment this had led to a large implementation of ISS (Internet Security Systems) products, including the real secure product suite, including real secure Network engines (a network-based IDS), Real secure Server Sensors (a host based IDS) and Real Secure Workgroup Manager (a Central Management platform). The Systems were based on Windows NT 4 workstations and servers and only where it was commercially necessary, were systems based on Windows 2000 allowed to be implemented. This had left the intrusion detection system based on Windows NT4 and the ISS Real secure 6.5 Product suite (which could be run on NT4), which ISS were still producing new Vulnerability signatures for (even if ISS were now putting most of its development and support effort into the new Real secure 7.0 Product Suite, (which the network engines only ran on Windows 2000).

Let us fast forward to today's Commercial environment. The latest quote from my Manager is "where can we reduce overheads" and "Projects only happen if there is a real business need". Due to budgetary constraints a project to update the hundreds of Windows NT 4 servers and workstations had been shelved a number of times and looked like being shelved again. With Managers being asked to cut overheads this was not a good time to have a system slowly become obsolete. Then came the catalyst for the major upgrade, ISS released an End of Life Notice for Real Secure Network Sensor 6.5 running on NT4 which was calculated to lose Support and have the last XPU (X-Press Updates) upgrade around September 2003, of which the e-commerce gateway had 22 on all the different environments. The announcement of the last XPU upgrade meant no new Signatures would be developed, so we would not have visibility of any new vulnerability. The replacement software was the newer version network engine 7.0 (which we were licensed to use), which unfortunately did not run on Windows NT 4. I took this news to the security manager who tasked me with writing a proposal for running a project to resolve this issue with a warning that almost all project work was at a stand still and only cost justified projects were likely to get the go ahead.

Figure 1 Typical IDS design before Technical Refresh Project



I raised a scope change to kick off the project. This was the first process in the path to justifying the work that would be required and the scope of the work involved. While the paper work was being completed I carried out a review of our present system. What improvements could be carried out for little or no capital cost? Which area's of the system were old and had been superseded by newer products and importantly what issues were outstanding with the present system and this new proposal. These points are summarized below.

- At the time the project was being implemented there were major security concerns over how secure VLAN 's were, so a decision was made to not rely on VLAN's and have many smaller low cost Switches, these were implemented in pairs for System resilience. This lead to two Network sensors being required for each point that required monitoring. (See Figure 1). The issue that needed to be resolved was the high cost of implementing a network IDS for every network switch. It had been noticeable that since cost had become an issue, the number of new projects having intrusion detection systems being specified and then being dropped from the proposals due to the high cost of implementing them was very high, as there was not much change from \$12000 for each network sensor and two were needed (one for each switch) at a choke point. See Figure 1. Each of these two Network engines in a pair, take it in turns to monitor the traffic passing through the firewalls. This is seen to be a waste of a very expensive resource.

- The management platform for the real secure Sensors and Engines was Real Secure Workgroup Manager 6.5 (WG M). This would need to be upgraded to support the new sensors, would an upgraded WGM 6.7 or a new management platform called Site Protector 2.0 (SP) provide a better management platform. ISS provide both WGM and SP free with Real Secure Sensors.
- There were still a small number of Version 5.0 Server sensors which would not work with either of the two new management platforms, so would need to be upgraded before the new management platform was in place.
- All the Network sensors were at the periphery of the network in front of the Web and Mail proxies, so it was not possible to see if attacks on the web and mail proxies had been successful or not, without looking at the proxy logs, which it was not possible to do quickly as a call needed to be raised for another team to do this. If the issues with the switches could be resolved it could be possible to re-deploy sensors to cover this.
- Which operating system platform would be best for the Version 7.0 Network Sensors?
- The present network engine hardware was Compaq rack mounted server platforms on which only Windows 2000 Compaq supported server. This was support by ISS but worked much faster on Windows 2000 Professional, the workstation version.
- Removing the second server in an NIDS pair will reduce the options when there is a hardware failure of the single remaining NIDS system. Can this risk be mitigated?
- Some Traffic is encrypted where the Network IDS is monitoring at the System boundary and it is only possible to see and analyze the un-encrypted traffic behind the proxies.
- Internet facing Network Engines are vulnerable to miss-configuration as they rely on only stealth mode interfaces being connected directly to the Internet. One mistake and it can create a backdoor into your system, could a second way of protecting the stealth interface be found and implemented.

During:

This gave me some questions that needed answering before a proposal could be put forward. I requested from my line manager that four proof-of-concept projects were kicked off. Begged and Borrowed equipment allowed all four to go ahead

- Which operating system should be used for running the network sensor 7.0 application? Assuming there was likely to be no new money for hardware. This was carried out by myself with help from colleagues with better Linux skills
- Was it possible to aggregate the traffic from two switches onto one Network Sensor? Carried out by myself with some help from the networks team
- Which of two possible management platforms should we upgrade too. Carried out by myself
- Could a one-way diode such as a TAP be implemented inexpensively, to protect the network engine as a defense in depth measure? Carried out by myself.

The Operating System Proof of concept project

This Project was managed and over half the work done by myself, with my colleagues doing the technical Linux install and testing, the main criteria for this proof of concept was to identify acceptable Operating systems that Network Sensor 7.0 was supported on. After looking on the HP web site for the Compaq DL320 servers for the support matrix ([2] HP) and after looking through the very good online documentation on the ISS web site ([1] ISS), two viable upgrade Paths were identified, Windows2000 server (The Windows Support team would not support Compaq unsupported Windows2000 Professional) and Red Hat Linux 7.3 Workstation.

The obvious upgrade path was to Windows2000 server but there would be an increased overhead for administration and support required from the Windows Server team and a lot of extra unwanted software that would be difficult to remove and harden. Also it was found that the Network Sensor 7.0 application worked much faster on Windows 2000 Professional (than Windows 2000 server). This was due to the shorter quantum value for Windows Professional, which gives a more responsive foreground task rather than the longer quantum useful for efficient file sharing services. The issue was best described in Microsoft Knowledge Base article 259025 ([3] Microsoft), which describes the issue and Ludens, Douglas Article "Optimizing Windows 2000 4" ([4] Ludens), which describes the solution. This describes the 3 registry key changes that are required to allow the Windows 2000 server install, that was supported on the hardware to work in a way which allows the Network Sensor

Application to work efficiently as a foreground Application by shortening the quantum value (e.g. the time slice given to each thread before a context switch can occur).

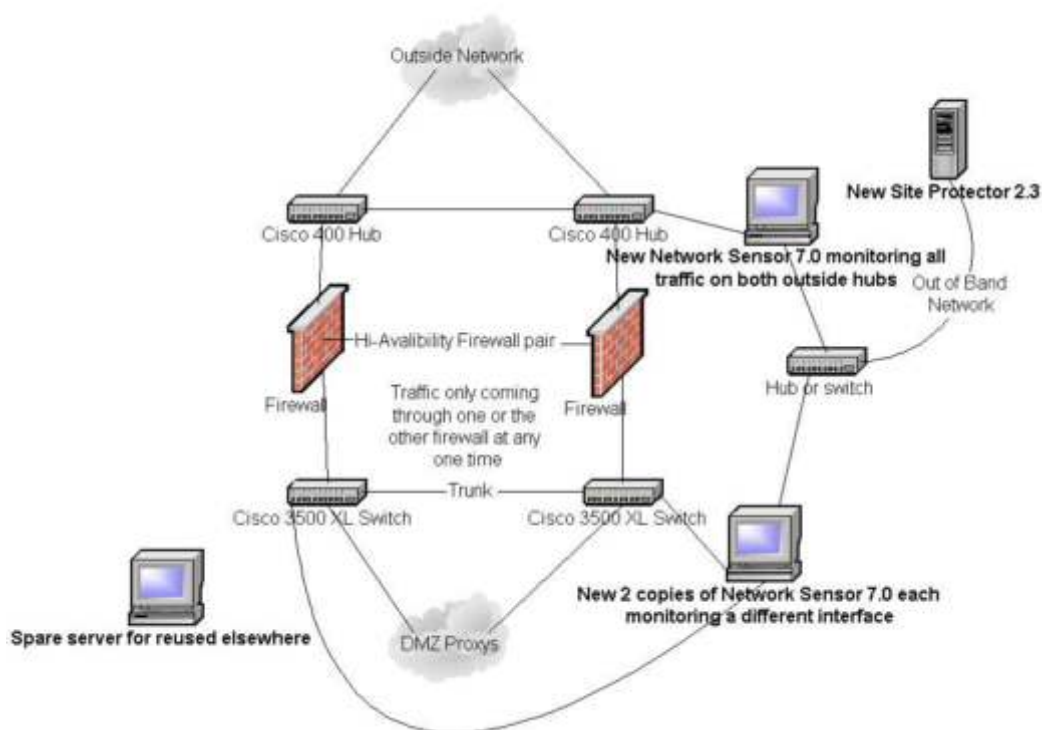
Both these platforms allow the installation of 2 copies of the Network engine software, so allowing us to replace 2 of our present Servers with one new system (where the engines would not be heavily utilized). So each server would be used 100 % of the time rather than at present a server pair each used 50 % of the time. So the Proposal was to replace our present 22 systems with a smaller number of single and double interface systems. This would allow a number of Servers to be removed or reused so helping to justify the project with reduced overheads without impacting the functionality and allowing better flexibility.

The Unix team was happy to do the Red Hat Linux Workstation Support (third line Support) and Security was happy to do second line Support for Linux. The fact that it was possible to install a cut down version of Linux with little more than the TCP/IP protocol stack and kernel reduced both the security risks and hardware requirements, both being important.

It is believed at present the only new hardware required over our present system is an extra network card in each server. There are no new software purchase or support costs (excluding Windows2000 server license or Red Hat Linux Workstation license for which a reduced number will be required) and we asked our ISS software suppliers to confirm the licensing for multiple interface network engines, their answer not surprisingly was that a license was required for each interface so no software savings could be expected (although it was found that the Site Protector license process counts 2 network engines on one server, as 1 license. When this was pointed out to our supplies their response was this was a feature to allow more flexibility to their clients not a change in license policy). One advantage that the Linux build had, was when reconnecting the control channel after a failure of the Site Protector console, the Windows 2000 network sensors would not reconnect as they thought that they already had a connection open. This problem was never seen with the Linux builds.

The outcome of this Proof of concept was Red Hat Linux was much preferred from a technical and support point of view, but both solutions provided an adequate platform for a single or dual network intrusion detection system. At this point the Management decided to implement the "safe" Microsoft solution. This was good news as the management chose the solution they felt safe with and more importantly provided the money for this much-needed project to go ahead.

Figure 2 IDS System After Technical Refresh Project Completed



The Switch SPAN port Aggregation Proof of Concept

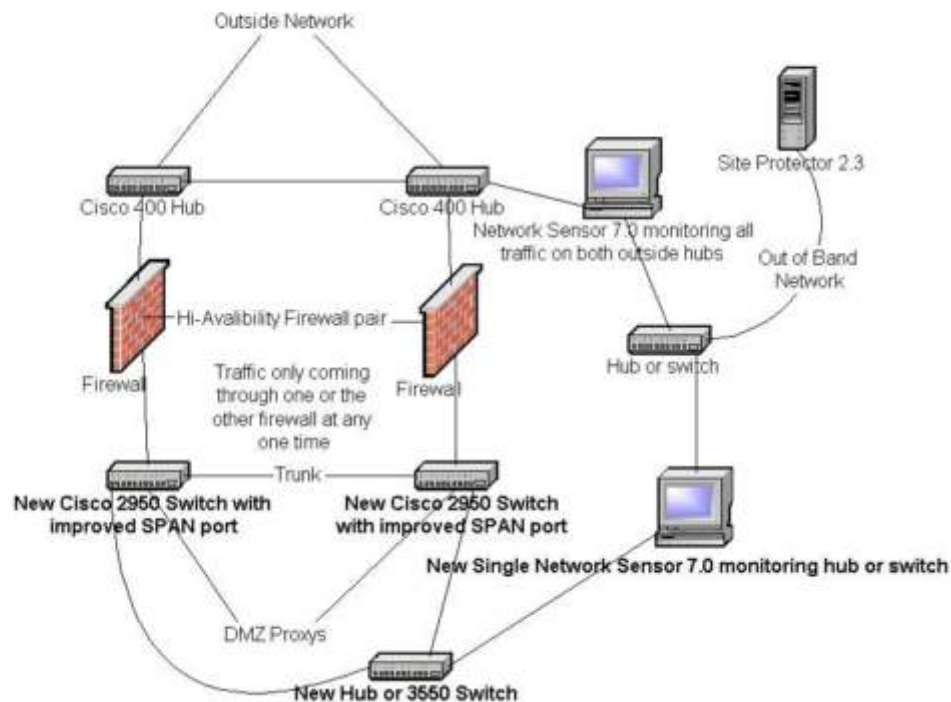
Unfortunately new hardware such as IDS TAP's and Top layer IDS load balancer Technology was excluded from this project as it would have sunk this project without trace. Another major issue that needed to be resolved was the high cost of implementing a network IDS for every network switch, especially as the networks teams liked implementing switches in pairs for fault tolerance. Since the early days of the system when the high cost of the Intrusion Detection System had not acted to deter new implementations, the new more cost conscious environment, was leading to new projects having intrusion detection systems being specified and then dropped from the proposals due to the high cost of implementing them. There was not much change from \$12000 for each network sensor and two were needed at each choke point (one for each switch) see Figure1. The Type of Switch that almost all of the sensors monitored was a Cisco 3500XL, as it turned out this switch has a very basic version of SPAN (or Port Monitoring). One idea which we were keen to try out was adding a switch or hub to connect a pair of switches together, so the traffic from the SPAN ports on two switches could be aggregated and Analyzed by one instead of two IDS sensors (A saving of \$12000 times nine or at least \$108000). I had also read about remote SPAN, so there was plenty of scope for improving our setup, I thought! The Proof of concept Test Rig soon dented my optimism, as the switches were unable to be configured to stop the spanning-tree Loop that developed (other Cisco switches had a feature called in-packets which resolves this issue).

Cisco Systems “Configuring the Catalyst Switched Port Analyzer (SPAN)” web site ([5] Cisco) gave me the bad news in great detail. A long hunt through the Cisco web pages confirmed that the fix for this issue on the 3500XL had not been released and as a newer version of Cisco’s small switch the 2950/3550 had replaced the 3500XL it was unlikely to release a fix for the 3500XL. RSPAN on the 3500XL was also tested and was found to not work. The networks team confirmed these findings and then lost all interest in assisting the security team in their quest to improve the system. The Switch battle would have to be fought another day in another project.

The GSEC paper written by Sylvain Proulx “Case study in deploying IDS network sensors in high availability switched network” ([6] Proulx) was used to confirm the solution was probably to replace the switches with newer models (either the Cisco 2950 or 3550 depending on the performance required) as switches are much less expensive than the IDS system. This has unfortunately fallen outside of the scope of this project. I managed to get a pair of 2950 and a 3500XL switch on loan as a proof of concept and was able to demonstrate the two configurations (see Figure 3). This has convinced the networks team of the feasibility of the new configuration and I have won a limited victory in getting agreement that this work can be done if a project can be funded and this work will also bring down the cost of specifying an IDS for future projects.

Figure 3

Figure 3 IDS future improvements



The Management Platform Review

After using the Real Secure Management platform called Workgroup manager 6.5 from ISS for 3 years, many of the issues and lack of functionality in this basically free product were known to us. The product was an adequate log for recent alerts but relied on good SQL skills for looking at historic data or long term trending information. Around the time this project kicked off a new management product was released called Site Protector 2.0, it was obvious that ISS was keen for its clients to move to the new Site Protector system. As providing support for two competing management systems is very expensive it was through that if the Site Protector product was stable enough and provided a better platform, then support of the old Work Group Manager would likely be dropped by ISS at some future date, this was subsequently to occur nine months later. Site Protector also has the added advantage that in future releases we can take advantage of closer integration of the ISS products Internet scanner and System Scanner with our Real Secure system. Site Protector has proved to be much easier to install and troubleshoot, the interface is much more flexible and provides a much better picture of complex related alerts and low frequency events.

The downside to Site Protector was the large number of patches required (the management platform requires incremental patches but the server sensors and network engines only require the latest patch) combined with its new feature of automatically downloading of these patches off of the web site via the internet (a very good feature if your console has access to the internet which ours does not). A utility called "ManualUpgrade.exe" is provided which resolves many of the issues in this area and allows downloads from the ISS site, which can be copied to CD and installed on the systems that could not connect to the Internet, but it is still a very basic and inflexible tool. Another downside was the requirement for Microsoft IIS web server as prerequisite for installation of site protector and for the detailed help documentation. This was mostly resolved in SP 2.3 where the Help and vulnerability information is now stored on an apache web server so that the deployment manager used for the installations was the only function now requiring IIS web services which could now be disabled (and only enabled for short periods when it was required for installation needs).

As a demonstration of working within the ISS product range the following two resources that I find most useful that ISS provide are the ISS forum news group whose archives can be found at neohapsis ([7] Neohapsis) and the ISS knowledgebase ([8] ISS) See the following example and how to find and search these resources on the internet. An interesting observation of the behavior of Site Protector to Network Sensor Control traffic (Pretty constant and not related to alerts) is it is much higher than WGM 6.7 and estimates of 700 MB? a-day for each network engine has been suggested as normal? ([7] Neohapsis) Mean that monitoring a remote sensor and managing it over a WAN link will not be practical. A "LowBandWidth" parameter was added recently but was still problematical but this should be fixed in the next release

([8] ISS). Also think about out of band LAN links especially if there are a large number of sensors.

Tidy up of present System

Before any Release 7.0 Real secure Engines or Sensors can be implemented there is an amount of clear up work that needs to be done. Some old servers were on a version 5.0 OS sensor a very old product not supported by Site Protector or WGM 6.7 (due to the decision to back out the new 6.5 software due to issues, The issues have now been resolved but the upgrades had not been carried out).

The increased risk of server hardware failure effecting the NIDS monitoring of the system due to one new server replacing two old servers can be mitigated on important internet facing DMZ's (therefore high risk) by using some of the spare servers in a third row of NIDS which are monitoring behind the proxy servers, these new Sensors will also be able to analyze some of the traffic which is encrypted in front of the proxies and unencrypted behind. This extra layer of NIDS can also be tuned to look for and alarm on what would be normal Worm or Scanning traffic on the present boundary NIDS systems as the proxies at present removes 99 % of this unwanted internet background noise (See figure 2). In a future switch project (see figure 3) a number of Real secure network engine licenses can be freed up, which can be used to provide the software for this third row of NIDS systems.

The lessons learned from the pilots were passed onto the rest of the team gradually as the new systems were added to each environment until most of the issues were resolved and on Production I could just send off anyone in the security team to upgrade a part of the system with a high level of confidence, it would be completed in a consistent manner.

© SANS Institute

After:

The new management platform Site Protector has been a real success being much easier to use and as stable as the old system. The many extra features and alerts thrown up by the Network engine 7.0 have allowed use to significantly increase our knowledge of the applications hosted on the Gateway and to tune the sensors so there is a finer line between the extremes of too few alerts (missing valuable data) and too many alerts (so important information is missed or ignored). The added third layer of NIDS's more than compensates for reduced resilience.

This research for this project has furnished the team with a list of future improvement that can be planned for inclusion in future projects. Costly ones that can be implemented when the budget becomes available and configuration changes and changes, which only take time and effort. Before this project was started the team was mainly concerned with supporting the products implemented by a Central Security projects team, and implementing small upgrades projects. Unfortunately this Central Security projects team was moved on to other systems due to the lack of large expensive project work within the system and over the last 12 months the support team I am a member of has proved itself capable of filling the projects role as well as continuing to provide Security Assurance. The SANS Security Essentials course and material can be seen as an important part in creating this successful team. It's wide range of sometimes detailed subject matter gives a good basic grounding in many security related areas giving enough information in a very large number of subject areas to search the internet and other resources for more detailed information.

Projects for the future will be utilizing the newer switch Monitoring technologies discussed in this paper, Implementing IDS TAP's when the money is available, as a high priority on the external hubs to provide a defense in depth (second line of defense) to backup the stealth interface configuration which is implemented on all monitoring network IDS interfaces and includes promiscuous mode network drivers, which do not need a protocol bound to them. On the internal monitoring ports where newer switches have been implemented the switch port monitor function can be configured to provide the same function (that of a one way diode) without the need for an expensive TAP. Improving internal authentication by the use of radius server integrated with the new active directory infrastructure. The integration of other ISS products such as Internet Scanner and System Scanner so reducing the amount of consoles required. The Site Protector system will be connected through a proxy to the Internet and twice daily allowed to download ISS patches from the Internet.

The “not so” Witty Worm a postscript.

This post script is a description of a real attack on a real system and as this is a Case study an Analysis of the Witty Worm is outside the scope of this paper. If you are not familiar the Witty worm it may be useful to read a good all-round Analysis such as on the CAIDA, the Cooperative Association for Internet Data Analysis web site <http://www.caida.org/analysis/security/witty/> ([9] CAIDA) before you continue with this paper.

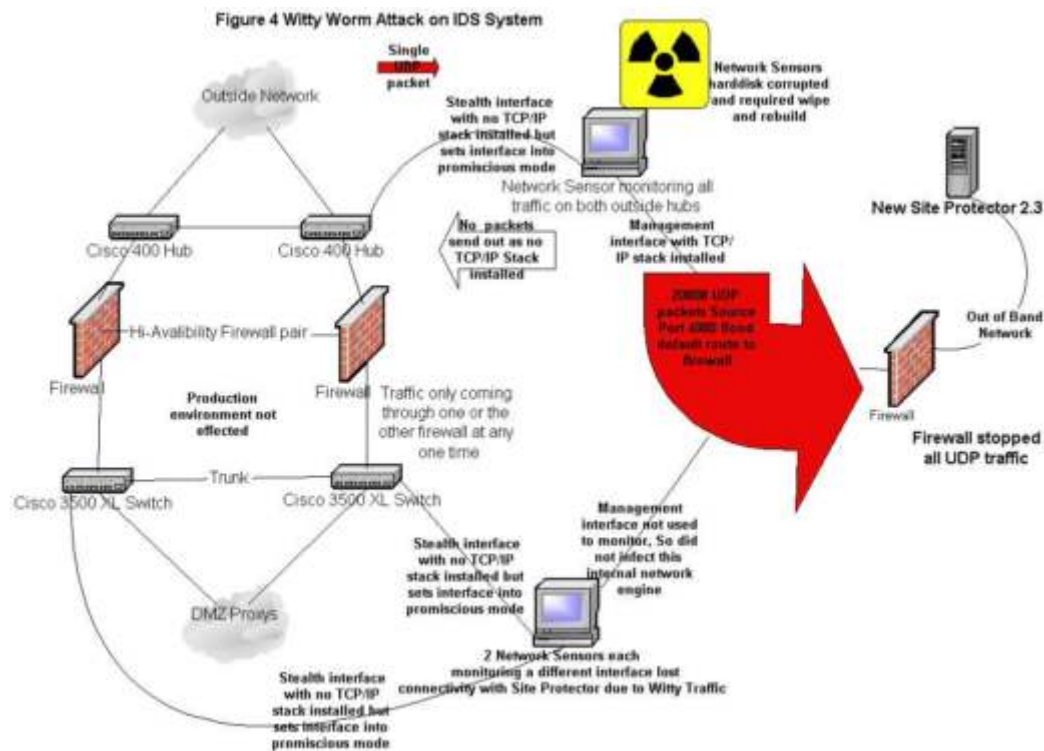
I now had an up to date intrusion detection system which was fully patched which over the next few weeks was tuned to provide all sorts of information, such as the continuous e-mails with malformed headers coming into the SMTP mail proxy's in the DMZ and not arriving at the internal mail system or the tuning of the SYNFLOOD alerts, which only occur when one of the systems internal services falls over during the day, (the kudos that can be won by getting the alert and working out which service on which server has gone down and informing the support team before the operators have alerted the correct team is great) or goes down for backup at night (I understand the operators complain we do not ask them enough questions when they phone the on-call security team between 11:45 pm and 02:15 am and get asked only three questions, was it a SYNFLOOD, was it on Source Port 25, and what time did it occur). Anyway getting back to a week ago, a few weeks after implementing the new network IDS in production, with tuning the new and interesting alerts well under way and myself nearly ready to send this paper into the post. ISS Corporation released a patch to its product range stating that this was to fix a potential memory leak in its product. The security team picked this vulnerability up on the Friday 19th March first thing in the morning after an ISS e-mail alert sent at 20:36 UTC on the 18th and discussed whether this should be patched as an emergency change (as our infrastructure could now be compromised) or push a normal change through as quickly as possible or even whether this patch should go into the normal patch cycle. All patches get tested in the many test/development/pre production environments before being released to production (unfortunately this process takes between two and five days to complete). Of course before this could be done, the change control process needed to be completed which comprises of around 12 teams needing to approve this patch (as it can be imagined this can take a long time to accomplish). As this had not been seen as an emergency as the information from ISS suggested this was only a possible exploit, which had only just been announced, and no exploit code had been released the entire process had not been completed before the attack came 36 hours after the vulnerability alert had been received. On Saturday 20th March at 13:21 UTC a single approximately 1300 byte UDP packet, source port 4000 with a randomly chosen IP address (which just happened to be one of 30 valid external IP addresses) arrived on the external sub-net just outside our external firewalls. Perhaps around 30 minutes after the first ever such carefully crafted packet was seen on the Internet.

Because during the implementation of the new IDS management platform and IDS software there was a large number of sensors not able to communicate with the management platform the alerts had been put to a low priority so the

operators would not call out the security team (as there was up to 5000 of these a day on the system. This procedure was still in place that Saturday morning so the heartbeat error of 3 sensors having difficulty communication with their console went unnoticed. Why did a small number of network Sensors (all had their management interfaces on the same out-of-band LAN) have communications difficulties you may ask and why did this stop at 09:31 GMT on Sunday Morning. Five minutes later the external Sensor crashed and lost contact for good with the console. This again went unnoticed due to the priority of the alert. We now fast forward to the first Member of the security team (not myself) to arrive on site at 7:00 UTC Monday Morning. One of the first jobs is to check the weekend logs, it was quickly spotted that 3 Sensors were not communicating with the console and one of the first actions was to reboot the 3 Sensors as this has been found to resolve most communications issues. The external sensor was found to have a faulty hard drive, as it would not reboot. The next two hours was taken up with organizing the resolution of this issue.

I am not one of the world's early risers so I arrived at 9:30 UTC to a huddle of Security specialists discussing the weekend's unusual occurrences. There was going to be a delay in getting the external sensor rebuilt as the Windows Server team was busy (aren't they always). So I suggested using one of the old sensors left over from the project, which already had Windows 2000 on so could be quickly installed. Next on with the important first procedure of the day "Strong black Coffee" and a scout around a few important Internet sites. Thankfully one of the first sites I checked was the SANS Internet storm center's website ([10] SANS) and specifically the diary ([10] SANS) which contained a roundup of the Witty Worms activity through the weekend. As I read down the diary, I decided to look deeper, the Worm activity which occurred on the weekend fitted some of the symptoms from my own system that weekend, so I asked one of my colleagues to check out the firewall logs specifically for any UDP source port=4000 traffic, and another to check out and find out as much about the Witty Worm as possible, while I planned what would be required to be done if my suspicions were true. By the time the firewall logs were checked and our fears were confirmed we knew enough about the worm to predict what the consequences would be and sent a colleague out to confirm everything was as expected. I confirmed that none of the other environments were connected to the Internet, and made sure that the networks team would not connect anything to the Internet without our agreement. We now knew why our external sensor had died and would not boot, we also knew that there had been no other infections. The security manager was informed as soon as the out-break was confirmed and informed the clients and by 11:15 UTC the patches had been transferred to Site Protector and were being installed on Production. By 11:35 UTC after customizing a new network sensor to replace the destroyed external sensor and installing and patching the new sensor. The external sensor was again live and by 13:00 UTC production patching was complete. Now only the other environments not connected to the Internet needed patching which took another day.

Figure 4 shows how what happened next could have been far worse, but because we had been pretty paranoid about security, the need for an out-of-band management network and the need for stealth interfaces (no TCP/IP protocol bound to monitoring network cards) for all non management network cards on the network IDS systems only one easily rebuilt system was taken down.



The lessons learned from this attack are “be suspicious of Normal occurrences such as faulty hard disks and multiple server problems when too many of them occur at the same time”. A day after the Witty Worm attack a meeting was called by the security manager to review the incident and to inform others not directly involved of what had occurred. I was asked to talk every one through this in a not too technical fashion, so the non-technical staff present could gauge the consequences. At the end of the presentation I proposed a list of improvements that would reduce the effect of a future such attacks on the system. This is when you should dust off your stack of “ready made” security improvement proposals projects (that don’t get implemented because security is not normally at the top of the managements list of priority’s) waiting to be thrust into the hands of a senior manager who wants his system to remain secure at any cost!

Infrastructure and Intrusion Detection Improvements in the light of likely future worm attacks.

A steady increase in Polymorphic worms (that have multiple methods of spreading) being released onto the Internet has made the Internet outside of

our system a very hostile place. The following suggestions have been proposed after a post mortem meeting two days after the attack to discuss the Witty Worm.

As worm technology improves so must our defenses. Else we may suffer an attack with more serious consequences in the future. This list was presented to the security manager while the attack was very much on his mind and will become the basis for a future security project

- Every few days a new group of XPU updates is released for new vulnerabilities and patches. This takes a lot of time to get on to Production and should be patched within hours so new worms and viruses can be monitored. The suggestion is that automatic downloads take place, where the Site Protector console polls the ISS web site for updates over the internet, this will be done twice daily automatic downloading new XML files which just contain lists of XPU files that need updating for Site Protector and the IDS at 05:00 and 17:00 every 12 hours (could be 05:30 and 17:30), the system then works out which of its components need patching. At present this is not possible due to firewalls and proxy servers in the way. Followed by a manual download and install of XPU's by the security person on earlies at 7:00 and later between 17:00 to 17:30). This would only download from www.iss.net from inside the secure zone and uses certificates, which is adequate security. Would need to modify a number of systems. The Downloads could be quite large so may have to restrict the hours that this could be done outside core usage hours depending on the urgency.
- The Alerting process on the system did not pickup the issues with the external network engine, this is due to the alerting being tuned to callout requirements. This could be much improved if we use 2 different methods (each which can be tuned independently). At present alerting for sensors being down either gives you no alerts or 1 alert a minute, which will occur 780 times x number of sensors in one night. This could be improved if we allowed SNMP Traps from sensors to the nearest management server, which we have tried to get implemented in the past. This requires firewall rule changes plus effort from the Enterprise management Team. Operations alerting via SNMP tuned for just the operators and email alerting with confidential information removed (as SMTP messages in new software now customizable).
- As an emergency change a number of our network engines have already been changed to stop any software communication over the TCP/IP stack of the management interface as Windows 2000 now allows any IP protocols, UDP or TCP ports to be blocked, (in our case only TCP 2998, 901 and 904, and no UDP were required) this provides some protection from a Polymorphic Witty worm type attack in the short term. As a long-term solution, an old spare router should be placed between the external network engine and the rest of the devices on the security management network in question with access lists denying

everything apart from the three TCP ports. This spare old router will only require a support cost, which is not clear at present.

- Buy 2 IDS TAP's \$800 each + 2 rack mount plates for 19 inch rack \$60 each. These would not be inline with production traffic to the fire wall. One for production and one for pre production (when that gets connected to the internet intermittently). This would act as a one-way diode, so would still get hit by a future worm but no remote control would be possible as no packets could be sent outwards. It is important to do this as this hardware stops this from potentially becoming a backdoor into the system. This has been suggested by the central security project staff but has never been included in a release.
- On our system we generally do not use the UDP protocol, (except for DNS Port 53 and SNMP management traffic) so modifying the ISP Router access lists outside of our system on the Internet would stop most attacks before they reach the firewall and IDS. This would have stopped the Witty work attack reaching our external IDS, requires 3 lines of access lists. Normally Denial of service attack use UDP or ICMP (ICMP already blocked using access list).
- The Witty worm only affected ISS Software running on Windows systems not Unix or Linux. In the recent past most widespread and successful worms have attacked just windows system (I count myself a windows supporter most of the time). The suggestion is to rebuild the external network engines as Red Hat Linux platforms (as very few worms have been produced that attack the Linux Kernel or TCP/IP stack directly). Most so-called Linux Vulnerabilities are seen in the Linux Applications, very few (the minimum number) of which would get installed. Some of you that have read this paper from the beginning will now be smiling as we unofficially named this proposal the "I told you so proposal" (see page 6 for an explanation). This would also allow much quicker rebuilds, (windows server support team took two weeks to rebuild and harden the operating system. Linux is now being used within the Data center and a team has been set up to provide Linux support.

The Witty Worm came as a big shock to many of us with its plethora of Firsts, which is a very sobering to list. I have copied this quote from the CAIDA, web site, "An analysis by Shannon, Colleen and Moore, David of the spread of the Witty Internet Worm in March 2004" as it demonstrates what we can expect from a small number of future worms which may find favour with Virus Developers who do not like copying the many standard Windows exploits so much in vogue today. This worm proves we need to continually review and invent new ways of providing protection for our systems.

"While the Witty worm is only the latest in a string of self -propagating remote exploits, it distinguishes itself through several interesting features:

Witty was the first widely propagated Internet worm to carry a destructive payload.

Witty was started in an organized manner with an order of magnitude more ground-zero hosts than any previous worm.

Witty represents the shortest known interval between vulnerability disclosure and worm release -- it began to spread the day after the ISS vulnerability was publicized.

Witty spread through a host population in which every compromised host was doing something proactive to secure their computers and networks.

Witty spread through a population almost an order of magnitude smaller than that of previous worms, demonstrating the viability of worms as an automated mechanism to rapidly compromise machines on the Internet, even in niches without a software monopoly". ([9] CAIDA)

The recent "Witty Worm" released in March 2004 would have destroyed all of the external sensors and without installing and fully patching Version 7.0 I would have had to stop monitoring the internet outside the systems firewalls for a considerable length of time. If this project had not been completed in late in February 2004 and the network engines had not been upgraded to Version 7.0 and instead of having to remove all the external sensors from the system as there was no XPU patch for the old 6.5 sensors for this worm, the systems external sensors only suffered a small outage before being rebuilt and patched before being placed back on that big bad outside world we all know only too well. Without a very paranoid defense -in-depth design of DMZ environments, this worm (or a similar worm) could have infected a large part of this e-commerce system as most servers contain ISS products that at the time were un-patched. Having all the Windows 2000 servers (40 plus servers) hard disk corrupted at once would have put parts of the system offline for days while they were restored and would strained the system restore procedure in place for this system.

References:

- [1] ISS (Internet Security Systems) Online Documentation in pdf format URL:
<http://www.iss.net/support/documentation/>
- [2] HP (Hewlett Packard) "ProLiant Server Certification". URL:
<http://h10018.www1.hp.com/wwsolutions/windows/index-dl.html>
- [3] Microsoft Corporation, Microsoft Knowledge Base article 259025 available on windows.about.com URL:
<http://windows.about.com/gi/dynamic/offsite.htm?site=http://support.microsoft.com/support/kb/articles/Q259/0/25.ASP> 12/03/2003
- [4] Ludens, Douglas writing for windows.about.com "Optimizing Windows 2000" URL: <http://windows.about.com/library/weekly/aa001008c.htm>
- [5] Cisco Systems "Configuring the Catalyst Switched Port Analyzer (SPAN)" <http://www.cisco.com/warp/public/473/41.html> - topic5 01/09/03
- [6] Proulx, Sylvain. GSEC practical paper "Case study in deploying IDS network sensors in high availability switched network" URL:
http://www.giac.org/practical/GSEC/Sylvain_Proulx_GSEC.pdf
- [7] ISS forum News Group Archives. Stored on Neohapsis URL:
<http://archives.neohapsis.com/archives/iss/> e.g. Very helpful e-mail from Awan, Farrukh on the ISS forum on Mon 02/06/2003 at 15:06, subject RE: [ISSForum] SP and WAN performance
- [8] Internet Security System "Knowledgebase", where previous answers to helpdesk calls can be found URL: <http://www.iss.net/support/knowledgebase/>
And search for "LowBandWidth"
- [9] CAIDA, (the Cooperative Association for Internet Data Analysis), An analysis by Shannon, Colleen and Moore, David "the spread of the Witty Internet Worm in March 2004" URL:
<http://www.caida.org/analysis/security/witty/> March 2004
- [10] SANS (SANS Internet storm center), Handler's Diary March 22nd 2004 written by Ullrich, Johannes <http://isc.sans.org/diary.php?date=2004-03-22>, 22/03/04