



# 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 in deploying IDS network sensors in high availability switched network

Sylvain Proulx

GSEC 1.4 option 2

April 14, 2002

### Introduction:

During the last couple of years, a lot of energy was concentrated to build a high availability network. The main reason was the revenue and the image of the enterprise. Firewall in cluster, redundant routers with HSRP and redundant switches. The only security defense was the use of firewalls. The upper management's concern was to have the network up and running to make revenue. Time has changed; a lot of publicity has been done about big corporation that was subject to Internet attacks. Like many others companies security has been put in management's agenda. They decided to deploy IDS in our network to improve the level of security.

The purpose of this paper is to describe the steps I have taken to implement the IDS network sensors in our high availability switched network. The paper will focus on how I deployed ISS Real Secure IDS network sensors in a switched network with firewall in cluster with redundant switches and routers.

### Before:

I work in the network engineering department for a wireless carrier (wireless = cellular phone) that provides data services and Internet access to the mobiles. Part of my job is network design and network security. Refer to the figure 1; I will describe the network that I have to protect. First our Internet connection terminates on a pair of routers with HSRP that provides redundancy. The Internet routers are connected to a pair of firewalls in cluster through a pair of switches. This is our first line of protection. We have Checkpoint firewalls running on Sun servers with Stonebeat full cluster software that provides high availability. So comes our DMZ where resides around one hundred servers that provide different function to our customers. As an example we have WAP servers, SMTP, AAA, NTP, HTTP servers etc. Most of the servers are UNIX based we also have few servers that are Windows NT or 2000. Our customer's access point is on the first line of firewall. From there they can get to the Internet or simply use our data services. There is another line of firewall that delimits the DMZ from the corporate network. All our databases sit on a segment attached to the second line of firewall. As many other companies, we did not have an incident response plan neither an incident response team.

As we can see the only defense that we had was firewalls. It is a good start but can not protect from everything. Firewalls are not bulletproof, they can not block all attacks. Depending of the policy on your firewall, whole might exist in the

firewall that let attacker go through your DMZ or internal network. As an example the Slammer worm attack that start in mid January hit computers behind firewall. How many of us had to react fast because of the fast propagation of this worm. Firewalls themselves can be port scanned or even attacked. Techniques are well described how to port scan a firewall to determine the type and version of it. One of them is using Nmap and scans the ports use by different firewalls, you have to know which port to look for. Checkpoint use 256-258, Microsoft Proxy use 1080 and 1745.

```
Starting nmap V. 3.00 ( www.insecure.org/nmap )
Host (myfirewall) appears to be up ... good.
Initiating NULL Scan against (myfirewall)
The NULL Scan took 1 second to scan 3 ports.
Adding open port 256/tcp
Adding open port 258/tcp
Adding open port 257/tcp
Interesting ports on (myfirewall):
Port      State  Service
256/tcp   open   FW1-secureremote
257/tcp   open   FW1-mc-fwmodule
258/tcp   open   Fw1-mc-gui
Nmap run completed -- 1 IP address (1 host up) scanned in 19 seconds
```

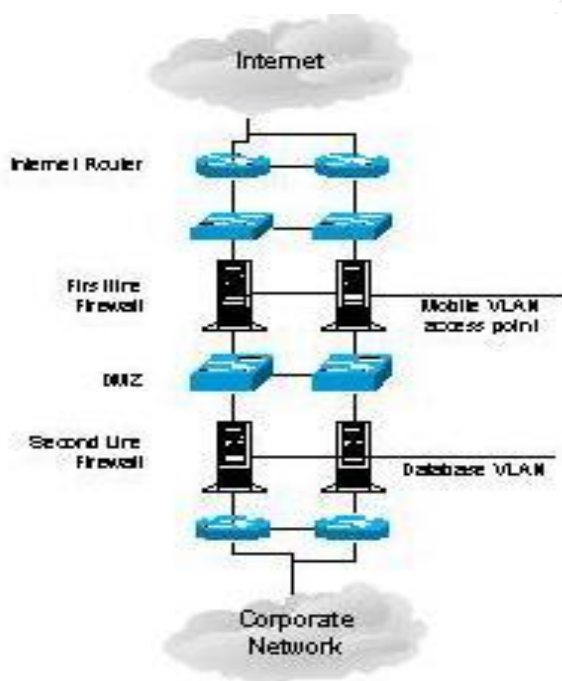


Figure 1

With over 5000 employees in the internal network and over a million cellular subscribers that have to use the servers in our DMZ and access the Internet, the management decided to deploy other measures to defend our network. First they hired a Security manager. He convinced the upper management to deploy a

network sensors based IDS to enhance the security of our organization. It is here the adventure started with the IDS.

During:

Real Secure IDS is a signature base system that looks for attack signatures that meet a specific patterns indicating suspicious activity. There are three forms of IDS; Network based, Host based and Stack based. Network based IDS uses raw network packets as data source. Typically this kind of IDS uses a network adaptor in promiscuous mode that listens and analyses all traffic. Host based IDS use software that continuously monitor specific logs on a specific host. Stack based IDS look for all incoming and outgoing packet as they traverse the TCP/IP stack to the OSI layers. Each form has their pros and cons. Network based IDS allow strategic deployment and require less management than deploying host based sensors. Deploying network based IDS in a switched may be challenging. Host based IDS uses logs containing events that occurred. This type of detection is more accurate and less prone to false positive.

As our first experience with an IDS and having around 100 servers in our DMZ and database LAN we decided to deploy network based sensors at strategic access point. The decision was based on few criteria; time to be in service, cost and ease of integration. It was faster and cheaper to deploy few network sensors than a hundred host based sensors. Since we are on a switched network we knew that choice was a bit challenging. We decided to deploy 4 network sensors in our network (refer to figure 2). Sensor 1 monitors the traffic going to/from the internet. Sensor 2 monitors all traffic in the DMZ. Sensor 3 monitors the traffic to the database LAN and finally the sensor 4 monitors the traffic going to/from our internal network.

© SANS Institute

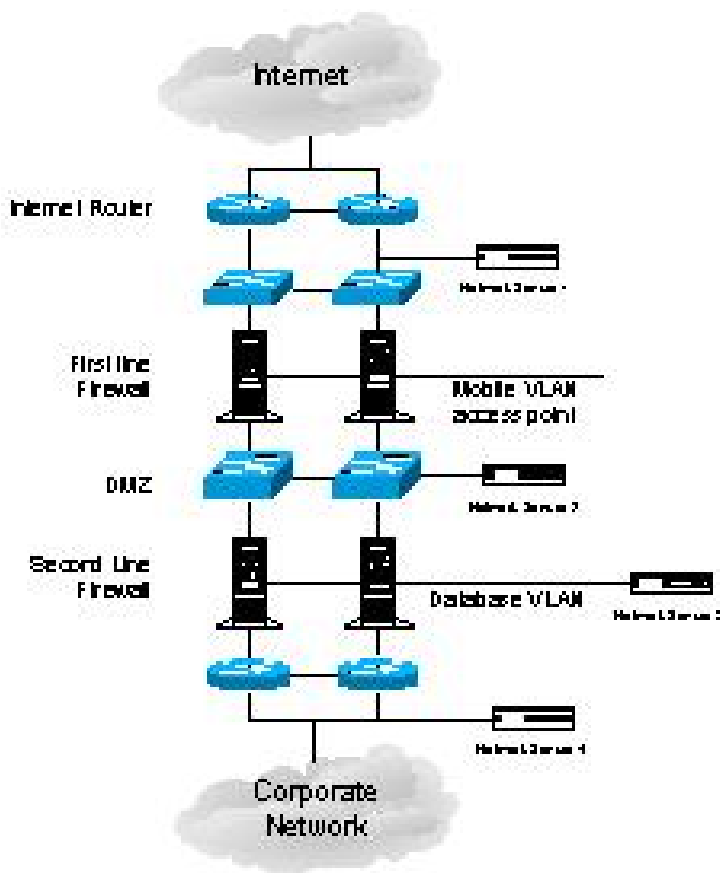


Figure 2

Real Secure IDS is managed through the WorkGroup Manager and has three components; console, event collector and database. All components may be installed on one computer or may be distributed on a unique computer. For scalability and performance reasons we choose to install each components on a different computer. The minimum system requirement for any component is an Intel Pentium II 400 MHz with Microsoft Windows 2000 server SP-1-3 or NT 4.0 SP4-6a, 450MB hard disk and 256MB memory. We installed the console and the database on two different Compaq Proliant DL380, one CPU 1.4 GHz Pentium III with 512 cache, 1GB memory, and two 36MB hard disk in RAID 0. The event collector resides on a Compaq DL320, one CPU 1.0 GHz Pentium III, 1 GB memory, two 36MB hard disk in RAID 1. We decided to use the Microsoft SQL 2000 as our database since the MSDE is limited to 2 GB of data. Real Secure IDS network sensor support different OS, Windows, Solaris 2.6 & 2.7, Red Hat Linux 7.3 kernel 2.4.18-10 SMP. We installed our four network sensors on four Sun Netra T1 DC200 with Solaris version 2.8. Each network sensor comes with 2 built in NIC. All the computers are on the same LAN and are protected by a firewall from the internal network. We wanted to protect the data collected by the

IDS. All the OS on each computers has been secure according to the SANS consensus guide Securing Windows 2000: Step-by-Step and Solaris Security: Step-by-Step.

In order to help automate the install process of all the IDS components, it is recommended to install the software on all the components (network sensors, collectors, consoles and database) before starting the configuration. The sequence should be:

1. Install the Asset and Enterprise databases
2. Install the network sensors
3. Install the Event collectors
4. Install the Console or WorkGroup manager

The database server will contain both the Enterprise and Asset databases. These two databases are installed from the Workgroup Manager Install program and are required before installing any other components. To install only the database part we used the 'Custom Installation'. This part of the install process will create 2 new tables in the SQL database instance: the Asset database (RSAsset60) and the Enterprise database (ISSED). Both can be viewed from the SQL enterprise manager application which is the database management tool. The first step to install the network sensors is to de-compress the TAR file and install the network sensor software. Two options are available: manual install or script install. We use the script install. The install process creates the following start up script: `/etc/init.d/realsecure`. To automate the configuration of the Stealth interface, we added the configuration command in the script in the start section; `/usr/sbin/ifconfig eri1 plumb -arp up #interface 1 in stealth mode`. The installation of the Event Collector and the Console on their specific server was completed using the WorkGroup Manager install program with the CUSTOM option to install only the component needed. The communications between every component are all encrypted using a pre-shared key that is built during installation on every server. The key also controls if a server is allowed to talk to another. For example, the Key on the Event Collector must be copied onto each probe it needs to connect to. Before starting to configure the IDS system, crypto Keys need to be distributed between servers. Although there is a menu in the console to perform this task, it is easier for the deployment of the first console to manually copy the Keys that need to be copied. The key is generated when the software is install on a system and is different with every install.

On Console server, the key is located in:

*C:\Program Files\ISS\Realsecure 6.0 console\Keys\CerticomNRA*

On Event Collectors, the key is located in:

*C:\Program Files\ISS\Realsecure 6.0 Event Collector\Keys\CerticomNRA*

On Unix servers, the key is located in:

*/opt/ISS/issSensors/network\_sensor\_1/Keys/CerticomNRA*

The Network Sensors need to have the key of every Consoles and every Event Collectors talking to it. The Event Console needs to have the key of every Consoles talking to it.

After the software installation the next step is to connect the promiscuous interface of each network sensors to the dedicated LAN to be monitored. Here comes the challenge. There are three main methods to tap into a switched network; Taps, hubs and spanning ports also call port mirroring. Each of the methods has its advantages and disadvantages. Before choosing the best method we had to understand the network topology and how traffic was going through the network.

The network is build on redundant Cisco 3512 switches for the Internet access and redundant Cisco 6509 for the DMZ, VLAN Database and the corporate network refers to fig. 3.

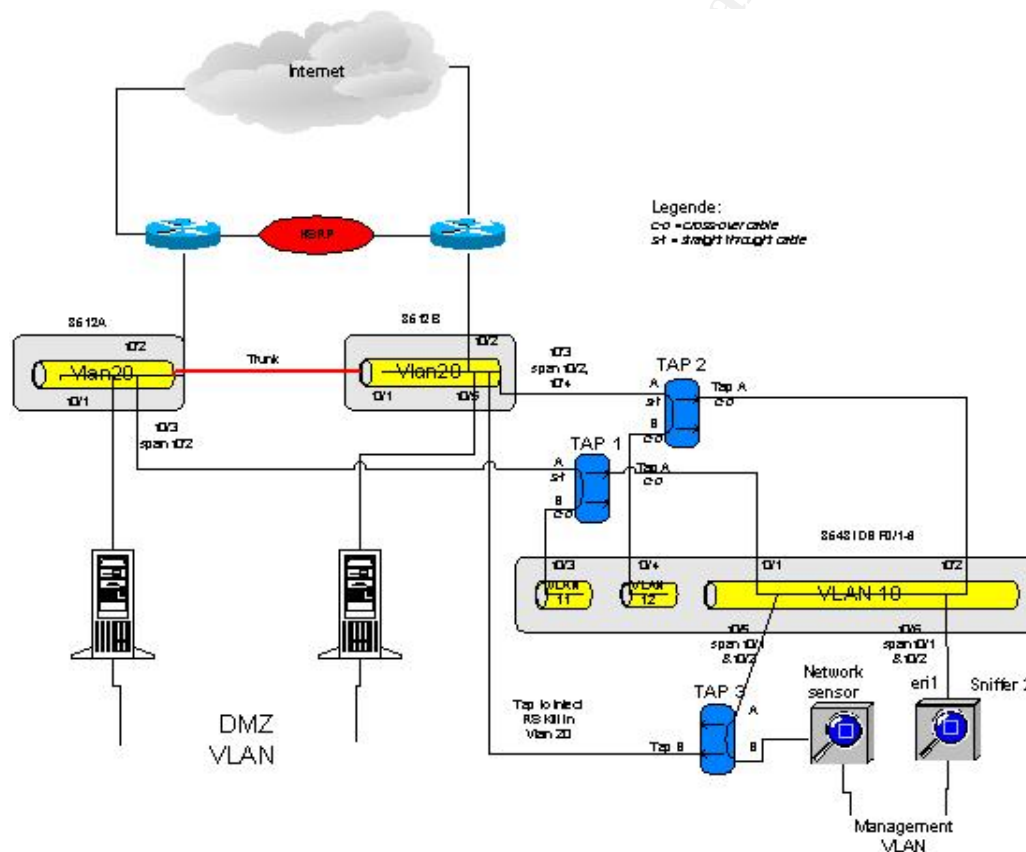


Figure 3

In our architecture all firewalls are in cluster using Stonebeat software from Stonesoft. The Stonesoft clustering technology share a common unicast IP address between clustered nodes per VLAN and we used a shared multicast MAC address on each nodes per VLAN. A multicast MAC address has the least significant bit of the most significant byte as 1, meaning that the first byte is an odd number; ex; 01:00:00:00:00:00 and 09:00:00:00:02:03 are multicast MAC addresses. By default a Cisco switch does not forward any packet with a MAC multicast. We have to statically define on which port a multicast MAC addresses will be forward to, see the example below.

```
mac-address-table static 0900.0000.0002 FastEthernet0/1 FastEthernet0/12 vlan 20
mac-address-table static 0900.0000.0002 FastEthernet0/2 FastEthernet0/1 FastEthernet0/12
vlan 20
mac-address-table static 0900.0000.0002 FastEthernet0/12 FastEthernet0/1 vlan 20
```

Routers facing the Internet are redundant using HSRP, refers to fig. 3. We have a pair of router using HSRP; Hot Standby Router Protocol. . In HSRP there is only one router active and the other one is in standby. Each router is connected to a different Ethernet switch. A trunk port is set-up between the two switches. When a packet from the Internet wants to gets to our DMZ it is routed by one of the one router and are forwarded to both firewalls at the same time. When a packet is leaving our DMZ to the Internet the packet is forwarded by one firewall at a time. So if we want to capture all the incoming packets without duplicate packets we have to monitor only the router ports on switch A & B.

We decided to use a combination of TAP and span ports, we also added another switch dedicated to the network sensor, refer to figure 3. TAP are fault tolerant device and does not impact the traffic flow. The TAP protects the network sensor by prohibiting people to establish a direct connection to it. It prevents the bridge loop that may happen by connecting two switches together. On the Cisco 3512, the SPAN port copy all the traffic both direction (egress & ingress) from one port to the span port. The limitation of SPAN on the Cisco switches 3500XL series is only limited by the number of ports on the switch. SPAN does not copy corrupted packets, CRC errors etc... This design permits the use of the functionality of injecting RSKILL. In table 1 there is a copy the port configuration of both switches and the one for the IDS. On the IDS switch we remove the spanning tree on each VLAN. The spanning tree algorithm that was running in each VLAN was blocking one of the two port were the TAP was connected. This is the command to configure a span port on a Cisco 3512;

```
3512A#config t
3512A(config)# int fa0/3
3512A (config)# port monitor FastEthernet0/1
3512A (config)# port monitor fastEthernet 0/2
```

3512A	3512B	3548IDS
-------	-------	---------

interface FastEthernet0/1 description FW-A duplex full speed 100 switchport access vlan 20	interface FastEthernet0/1 description FW-B duplex full speed 100 switchport access vlan 20	interface FastEthernet0/1 duplex full switchport access vlan 10
interface FastEthernet0/2 description Router A duplex full speed 100 switchport access vlan 20	interface FastEthernet0/2 description Router B duplex full speed 100 switchport access vlan 20	interface FastEthernet0/2 duplex full switchport access vlan 10
interface FastEthernet0/3 description IDS--port monitor duplex full speed 100 port monitor FastEthernet0/1 port monitor FastEthernet0/2 switchport access vlan 20	interface FastEthernet0/3 description IDS--port monitor duplex full speed 100 port monitor FastEthernet0/1 port monitor FastEthernet0/2 port monitor FastEthernet0/4 switchport access vlan 20	interface FastEthernet0/5 duplex full port monitor FastEthernet0/2 port monitor FastEthernet0/1 switchport access vlan 10
	interface FastEthernet0/5 description IDS_inject_RSKILL switchport access vlan 20	interface FastEthernet0/6 duplex full port monitor FastEthernet0/2 port monitor FastEthernet0/1 switchport access vlan 10

**Table 1**

For the other network sensors we used the same topology, i.e. TAP and SPAN. The only difference is the SPAN configuration on Cisco 6500 series is not the same as the 3500 series. On the 6500 series there is a limitation of the number of SPAN session; 2 RX SPAN sessions or 4 TX SPAN sessions. We had 3 VLAN to monitor so we had no choice to use TX SPAN session. This is the command to configure SPAN session on a 6500 switch. The port 3/11 is the port to be mirror and the port 3/23 is the destination port. We can set multiple ports to be mirrored, in the second line below the port 3/27, 4/2 through 4/25 and 4/40 through 4/43 will be mirror to the port 3/12. It is not recommended to SPAN multiple ports to one port because this one may be saturated and some packets will be missing.

```
set span 3/11 3/23 tx inpkts disable learning enable multicast enable create
set span 3/27,4/2-25,4/40-43 3/12 tx inpkts disable learning enable multicast enable create
```

We used the tool TCPDUMP to monitor and analyzed the traffic that was captured by the network sensor to make sure that we captured all packets and we did not receive duplicate packet. All the packets going through the firewall are sent to both firewalls since they use MAC multicast. Before we end up with the network configuration show in figure 3 we tried other configuration and found that we were not capturing all packets or we received duplicated packets in one

direction. So before we started our IDS we did want to be sure that our IDS network sensors was capturing all the right traffic.

After we deployed adequately the network sensors we went back to the console, it is there that all the management and monitoring is done. The console window has five areas; Activity tree window, Managed Asset window, and the high, medium and low alert windows. On the Activity tree window we find all the current events logged by the active sensor. We can view the events in three different way by choosing the tab Source, Destination or Events in the Activity tree. Most of the time I use the Events tab. It is in the Managed Assets windows that we apply the policies to the network sensors. A policy is a group of signatures that the sensor will respond to. RealSecure IDS has predefined network sensor policies for different environment. There are nine pre-defined policies for the network sensors. We can modify policy with the Policy Editor. We apply different policies to each network sensors. After a minute of monitoring we had hundreds of events. There was so much noise at the initial start up that we were not able to identify real malicious activity from the false positive. To minimize the number of false positive we had to deactivate some signatures and build filters. It took a couple of weeks to fine tune the policy on each network sensors.

Pre-defined policy can not be modified but can be saved on another name and the new policy can be modified. We can view the policy applied on a sensor by right clicking on the sensor in the Managed Assets window and choose the View Active Policy. There is five tabs on the network sensor policy editor window, Security Events, Connections Events, User-Defined Events, Filter, X-Press Updates. The Security Events and X-Press Updates are the two tabs to use to configure events to be monitored. Each signature is grouped. Each sensors policy signature has to be tailored. As an example the sensor that monitors the Internet do not need to activate DHCP signatures, there is no DHCP server. Another example is the signature Openview\_NNM\_Overflow that detects an excessively long line directed at the OpenView NNM alarm service version 6.1, this signature does not apply to our environment since we are running version 6.2. We disabled all signatures group that were not relevant to each sensor. We monitored again and we had to build filter to minimize the events number. Filters are used when you want to keep a signature on your policy but you want to filter specific events. We can build filters based on protocol (TCP, UDP or ICMP), source and destination IP, source and destination port. As an example we have a server running MRTG that monitor traffic by sending SNMP request to different nodes, we had to filter all the SNMP traffic from this server because it was generating events. ISS provides new signatures from time to time with the X-Press Updates. One important point is the documentation, we decided to document every signature we removed and filters we configured.

After we reduced the number of events we look at what kind of response we would apply to each event. There are seven different responses for network

sensor we can apply on an event. We can use any combination of them. The DISPLAY response notifies the console of an event. The LOGDB logs the events to database in three different ways; LogWithoutRaw logs summary information about the event, LogWithRaw logs with the entire binary session, LogFiltered logs an event and remove any occurrence of the same event. The way you log will have a direct impact on your database size. Email sends an email to an administrator. RSKILL sends a TCP Reset to both source and destination of a TCP connection; it works only with protocol using TCP as SMTP, HTTP. SNMP send a trap to an IP destination. OPSEC sends a message to the firewall that will block the source of the event for a period of time. User Specified runs a user response, as an example; when an event is happening the response can be a command that will send a ping to the management IP of a network sensor that will trigger an action on the sensor. View Session records a session that we play back as an example we can record a Telnet session to see what an intruder will do on a host. The network sensor response policy can be configured through the Managed Asset window by right clicking on a network sensor and choose responses. In our deployment we used the DISPLAY and LogWithoutRaw for all of our active signatures. We configured SMNP trap to be sent to our HP Openview server for all the high priority events.

Database back-up is important since you want to keep your evidence in a safe place. We developed a back-up plan with the SQL 2000 tools. We scheduled a weekly full Backup of the ISSUED database. Also we scheduled a daily differential backup and every 2 hours we backup the transaction logs. For now all of the backups are stored locally. There is a plan to migrate the backup on another server for security reasons. ISS has a utility to purge the database issDBMS. With this utility we can schedu when to purge for the events older than so many days. This utility helps us to keep the size of the ISSUED database workable. When you want to generate reports it is less time consuming when you have a smaller database.

Reporting is another feature that we can through the console. We activate the reports windows by selecting the View – Reports menu from the console. There are pre-built reports and editable criteria that we can generate. Reports are good tool to show to the management that we are target of scans and attacks. But also it was a good tool to identify the high runner's signatures.

Before we deployed the RealSecure IDS we did not have an incident response plan. We had to develop an incident response plan. The first step was to choose the members of the team. A team was formed that included individuals from system admin, network engineering and management. Our security manager developed incident response plan to help us in deciding how to handle attacks. This plan identifies action to take during the SQL slammer worm attacks that we detected during the week-end of the 25<sup>th</sup> January. Detecting intrusion is a first step but the other one is what should we do with it.

## After

To deploy security tool in a high availability switched network we should understand its architecture and all the components that are connected to it. During the installation of such a security tool we should make sure of what we monitor. The use of a tool such as TCPDUMP helps a lot to see what we monitor. We found out that using IDS network sensors increase the level of security by seeing unwanted traffic and malicious activity on our network. Unwanted traffic is the kind of traffic that bad configured server may generated. Malicious activity is the one we should care about like SQL Slammer worm. The security of our network increase, now we can detect attacks that before we were not able to see. Since we deployed our IDS, we are now able to generate reports that help us to show to the upper management that the Internet is a jungle. Even if we feel a little more secured there is limitation to IDS network sensor. The IDS network sensors can tell that someone is trying to get into your network but can not tell you if someone really get into it. Network sensor can not detect malicious activity on secured communication as VPN, SSH or SSL. In the near future we will have a look on deploying host base sensor.

© SANS Institute 2003, Author

## References:

Joel Scambray, Stuart McClure and George Kurtz "Hacking Exposed"

<http://www.hackingexposed.com>

How To Guide-Implementing a Network Based Intrusion Detection System

[http://www.iss.net/support/product\\_utilities/realsecure\\_tech\\_center/disclaimer.php?filedown=switched.zip](http://www.iss.net/support/product_utilities/realsecure_tech_center/disclaimer.php?filedown=switched.zip)

System Requirements, Real Secure Protection System

[http://documents.iss.net/literature/RealSecure/rs\\_sysreqs.pdf](http://documents.iss.net/literature/RealSecure/rs_sysreqs.pdf)

Multicasting with Stonesoft Clustering Products

<ftp://download.stonesoft.com/web/Support/StoneBeat/Technical%20Notes/SGSBTECNMulticasting.pdf>

Cisco- Hot Standby Router Protocol Features and Functionality

<http://www.cisco.com/warp/customer/619/hsrpguidetoc.pdf>

Configuring SPAN & RSPAN

[http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw\\_5\\_5/cnfg\\_gd/span.htm](http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cnfg_gd/span.htm)

TCPDUMP documentation

<http://www.tcpdump.org>

SQL-Slammer-worm

[http://www.iss.net/security\\_center/static/11153.php](http://www.iss.net/security_center/static/11153.php)

© SANS Institute 2003, Author retains full rights.