Interested in learning more about security?

SANS Institute
InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Intrusion Detection Systems: An Overview of RealSecure

Code Red. Denial of Service attacks. The never-ending flood of security bulletins from Microsoft. For a network security professional, drastic measures have been needed to stem the tide of attacks against your network. Wouldn't it be nice to have a tool that could warn you of potential threats to your network? A tool that would log intrusion attempts and notifies you in near real-time of attacks to your network? Intrusion Detection Systems (IDS) are tools to monitor networks for anomalies that could indicate an attack ...

Copyright SANS Institute

Author Retains Full Rights

AD

Intrusion Detection Systems: An Overview of RealSecure

Darrin Wassom
GSEC Practical
Version 1.2f

Introduction:

Code Red. Denial of Service attacks. The never-ending flood of security bulletins from Microsoft. If you are a network security professional, these "media events" of 2001 have caused many sleepless nights and the sudden realization that drastic measures were needed to stem the tide of attacks to your network. After many hours spent patching your servers and wading through countless CERT advisories, you might have come to the conclusion that it would be nice to have a tool that could warn you of potential threats to your network. A tool that would log intrusion attempts and notifies you in near real-time of attacks to your network. Does such a tool exist? Fortunately, yes!

Intrusion Detection Systems (IDS) are quickly gaining in popularity as a way to monitor networks for anomalies that could indicate an attack on your network. Typically, IDS is passive in nature and works by scanning packets for patterns that match a pre-defined signature base. The signature base contains information relating to a known vulnerability, threat or pre-attack probe. Most IDS platforms will also allow for the creation of a custom signature base that can scan for pattern matches (passwords, keywords, etc) or new threats where a known signature does not currently exist. [1]

Types of Intrusion Detection Systems:

Many of the existing IDS products on the market fall into one of two categories; host based intrusion detection (HIDS) and network based intrusion detection (NIDS). HIDS monitor files, access attempts, system logs or other definable portions of a particular host for suspicious activity that could indicate an intrusion attempt or successful entry into the host itself. NIDS, on the other hand, are usually placed on a particular network segment and scan the traffic for patterns that fall outside of configured boundaries.[2] Weighing the inherent advantages and disadvantages of an IDS will help you in determining which type of IDS, or both, will fit your individual needs and help assure network integrity.

HIDS have two distinct advantages in that they can detect attacks that are not normally detectable in a NIDS type configuration because they have view of events that are local to the host and they can operate effectively in a switched network environment. Maintaining a HIDS environment can be a very time intensive process because the sensors must be installed on every host that needs to be monitored. Since the host sensor resides on the local platform, the potential for OS performance degradation exists because of the processor time needed to perform the monitoring function. Another point to consider when deploying HIDS is that they are susceptible to being discovered by a hacker and they do not perform well (if at all) under certain denial-of-service attacks. A NIDS configuration can be very cost effective in terms of the number of sensors that would need to be deployed. A single network sensor, properly placed, can monitor very

© SANS Institute 2001, As part of the Information Security Reading Room. Author retains full rights.

large networks with little or no impact to the existing environment. However, they can be difficult to deploy in a switched network environment and at times of peak traffic, they could drop packets in order to attempt to keep up with the stream of traffic coming across the network.[3] IDS, in general, are not a tool that will eliminate intrusion attempts or make your system more secure but they will aid in vulnerability assessment, risk mitigation and provide another layer to a "defense in depth" strategy.

As you decide on which IDS product would best suit your needs, you will probably be deluged with an onslaught of marketing hype, client testimonials and a countless array of companies that promise to deliver the best IDS available today. There are a number of sites on the Internet to obtain a listing of commercially available IDS products but the most comprehensive listing seems to be at: http://www.networkintrusion.co.uk/ids.htm. Not only does this site reference commercial products, it also contains an exhaustive listing of open source products like SNORT (http://www.snort.org) and SHADOW (http://www.nswc.navy.mil/ISSEC/CID/) that have gained the respect of network administrators around the globe.

Internet Security Systems - RealSecure

In the Internet age, time is measured not in years but rather days, weeks and months. Products come and go with nary a glance or mention but one IDS product that seems to stand the test of time is RealSecure from Internet Security Systems (ISS). ISS is an Atlanta, Georgia based security-consulting company that was founded in 1994 by Christopher W. Klaus. Their mission is to be their "customers' trusted security provider and premiere provider of security management solutions for the Internet."[4] RealSecure is considered by many in the industry to be the flagship product of ISS. The company claims that RealSecure is the most comprehensive and widely used IDS product on the market and, according to International Data Corporation (IDC), they are right in making this claim.[5]

The RealSecure IDS uses a distributed client-server architecture comprised of a workgroup manager and various sensors that can be used to monitor traffic on multiple platforms and network environments. The Workgroup Manager provides a secure channel of communications to administer, configure and generate reports from the installed sensors. The manager console also allows for alert monitoring and offers an online database of well over 2000 attack signatures.[6]

The RealSecure sensors can be installed at strategic points in your network and include network, OS and server sensors. The network sensor (obviously) is a NIDS and monitors network traffic for signs of malicious intent while the OS and server sensors fall into the HIDS definition because of their proximity to the host. Both the OS and server sensors can be used monitor inbound and outbound traffic to the host for any unauthorized activity.[7]

Testing RealSecure:

© SANS Institute 2001, As part of the Information Security Reading Room. Author retains full rights.

Building a test system is the tried and true method that most network administrators use when evaluating a product. We all know that companies would **never** exaggerate claims on the products ability but it always helps to test the product and put it through its paces before making a final determination of whether you want to purchase or move onto the next glossy brochure and lofty claims. Ideally, testing should be done on an isolated network so that it will not affect any processes running in production. I am very lucky in the fact that I have access to a lab network that allows me to spend hours tinkering with various products and testing methods. Using the following hardware, I configured the lab network for testing the RealSecure product.

Lab Components:

- Network connection: DS3
- Router: Cisco 7x series
- Switch: Cisco 5x series
- Firewall: Checkpoint Firewall-1 running on Solaris (E420R)
- Tap: Shomiti 10/100 Full Duplex Tap
- RealSecure Sensor: Sun Ultra 10 running Solaris
- Web Server: 486 workstation running Linux and Apache
- Web Server: Windows 2000 running IIS
- Workgroup Manager: Windows 2000

© SANS Institute 2001, Author retains full rights

© SANS Institute 2001, As part of the Information Security Reading Room. Author retains full rights.

Lab Switched Network Used for Testing

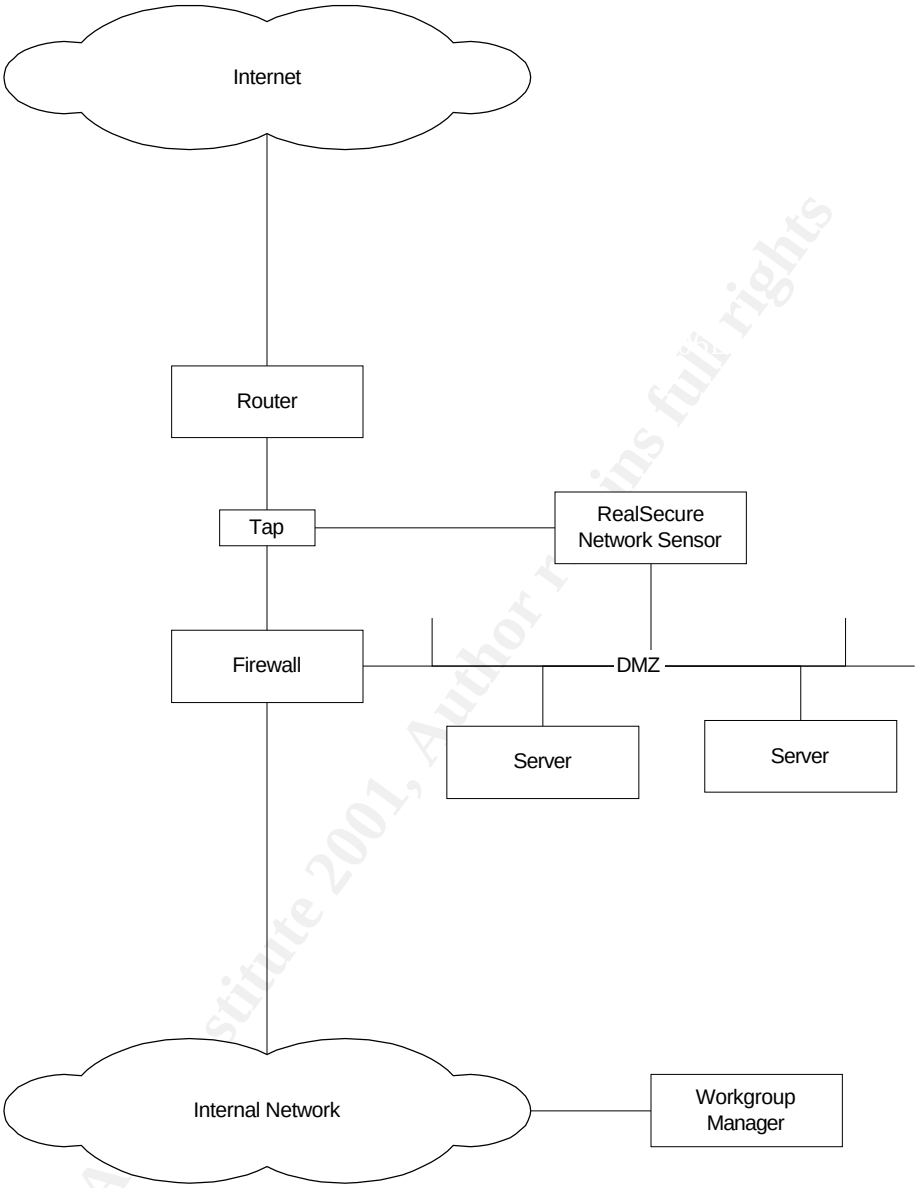

The Network Sensor:

Because of a limited budget and the fact I wanted to get the "most bang for my buck", I opted to use a network sensor for my testing. The RealSecure network sensor will monitor all traffic on the segment without any impact to the network itself. Because the network sensor must scan for thousands of known attack signatures, it must be installed on a dedicated machine. In this case, the network sensor is installed on a Sun workstation (Ultra 10) although it can also be installed on an Intel platform running a Windows operating system (NT or 2000).

© SANS Institute 2001, As part of the Information Security Reading Room. Author retains full rights.

The minimum requirements needed for the network sensor:

- Operating System: Solaris or Windows NT/2000
- Intel Pentium II 300 MHz for Windows
- 128MB RAM (minimum)
- 256MB RAM (recommended)
- 25MB disk space for installation
- 150MB disk space for log files and database entries
- Network Interface Card (NIC) capable of promiscuous mode and connected to the network segment to be monitored
- OPTIONAL: A second NIC connected to a secure segment for out of band management from the Workgroup Manager

The latest system requirements for any ISS product can always be found at: http://documents.iss.net/literature/RealSecure/rs_sysreqs.pdf

Sensor Placement:

Proper sensor placement is critical to the success of a defense in depth strategy with IDS. It is widely debated that most attacks to a network come from the inside by a disgruntled employee, errant code or simply non-malicious intent by a less than informed employee but this claim is often not supported by concrete data. Ideally, a network sensor should be placed in front of the firewall and a second sensor behind the firewall to monitor outbound traffic for signs of a compromised system.[8] However, given a limited budget and the opportunity to place one sensor, the best place is in front of your firewall. This placement will allow the sensor to see ALL traffic heading inbound to your network before it reaches the firewall and will be able to give you a better sense of the types of traffic being blocked by a properly configured firewall. Other options could include placing a sensor on any part of the network that you feel is critical but this can be costly and calculating the risk of not deploying a sensor should definitely be considered before breaking the bank on IDS.

I chose to deploy the network sensor with 2 network interface cards. This allowed me to use one NIC for managing the sensor from within the internal network while also allowing me to use the second NIC to see traffic on the external network in promiscuous or stealth mode. In stealth mode, the promiscuous NIC does not have a protocol stack bound to it, which means it does not have an IP address. In this mode, the NIC is relatively invisible to would be hackers. The management NIC is connected to a secure network and has a full TCP/IP stack with an IP address and IP services. This channel is used to communicate with the console. When deploying the network sensor on Solaris, the stealth interface is NOT enabled at system boot. To bring this NIC up every time the system boots then place the following command in the RealSecure startup script:

/usr/sbin/ifconfig nf0 plumb –arp

Use the command, **ifconfig –a** to see if the card is working properly after system boot. You should see an output similar to this:

© SANS Institute 2001, As part of the Information Security Reading Room. Author retains full rights.

Flags=8c3<UP, BROADCAST, RUNNING, NOARP, MULTICAST>
mtu 4352
inet 0.0.0.0 netmask 0
ether 8:0:20:f0:0:ba

RealSecure in a Switched Network Environment:

Utilizing IDS in a switched network environment poses an interesting set of challenges. The most common method of deployment is utilizing the Switch Port Analyzer (SPAN) port to monitor traffic. The SPAN port works by mirroring traffic from a switched segment to a pre-defined port.[9] The advantage of using a SPAN port is that it can be easily configured and doesn't require additional hardware. There are some limitations to this type of deployment that could cause it to be more trouble than it's worth. A typical switch will only allow for one SPAN port to be used. This means that to monitor more than one port, a range of ports must be spanned to the single SPAN port. In a network with heavy traffic, monitoring more than one port would quickly overwhelm the SPAN port and cause it to start dropping packets.[10] There is also the inability of a SPAN port to mirror errors such as Runts and Giants, which could indicate a network attack.

Until recently, ISS recommended the use of a network tap when deploying RealSecure but they now recommend installing a server sensor on each server connected to a given switch. Could this be an attempt to generate more revenue by recommending that more sensors be purchased? A carefully constructed search of the ISS Knowledgebase will reveal a few articles that are old enough to contain references to the Shomiti Century Tap, the recommended network tap. (http://www.iss.net/customer_care/knowledgebase/ - Reference # 010215-000018 Modified on 2/15/2001 03:24PM – This site requires a login ID and password.) The advantage of using a network tap is that it is extremely fault-tolerant and will not impact the flow of traffic to the network.[11] The use of a network tap also prohibits a direct connection to the network sensor, which can protect it from various forms of attack. The use of network taps will add additional cost to an IDS deployment and it is also important to note that some functionality is lost. RealSecure has a tool that will terminate connections that meet pre-defined criteria. The use of this "kill" is not supported if used in conjunction with a network tap. This is because a tap will only allow the sensor to monitor traffic in one direction.

Because of the limitations associated with using a SPAN port, I chose to install the Shomiti Century Network Tap. Finisar Systems recently acquired Shomiti Corporation but the taps are still available for purchase.[12]

The Workgroup Manager:

The Workgroup Manager (console) provides a central point of control for management and scalability. A single console allows an operator to monitor up to 50 separate sensors at one time. The console communicates with the sensor via an encrypted channel, which ensures network integrity. The console performs a multitude of tasks:[13]

© SANS Institute 2001, As part of the Information Security Reading Room. Author retains full rights.

- Display events
- Inspect events
- Edit policies
- Push policies to the sensors
- View logs
- Generate reports
- Event response
- Propagate events
- Session playback and review
- Update sensors
- Database of known signatures

I encountered numerous "false positives" with the default installation and it required a few hours of detailed configuration to filter out traffic that wasn't a threat to the network. Extreme caution should be used when filtering out these alarms because they could, in fact, be trying to tell you something. False positives below are the Stream_DoS and SYNFlood.

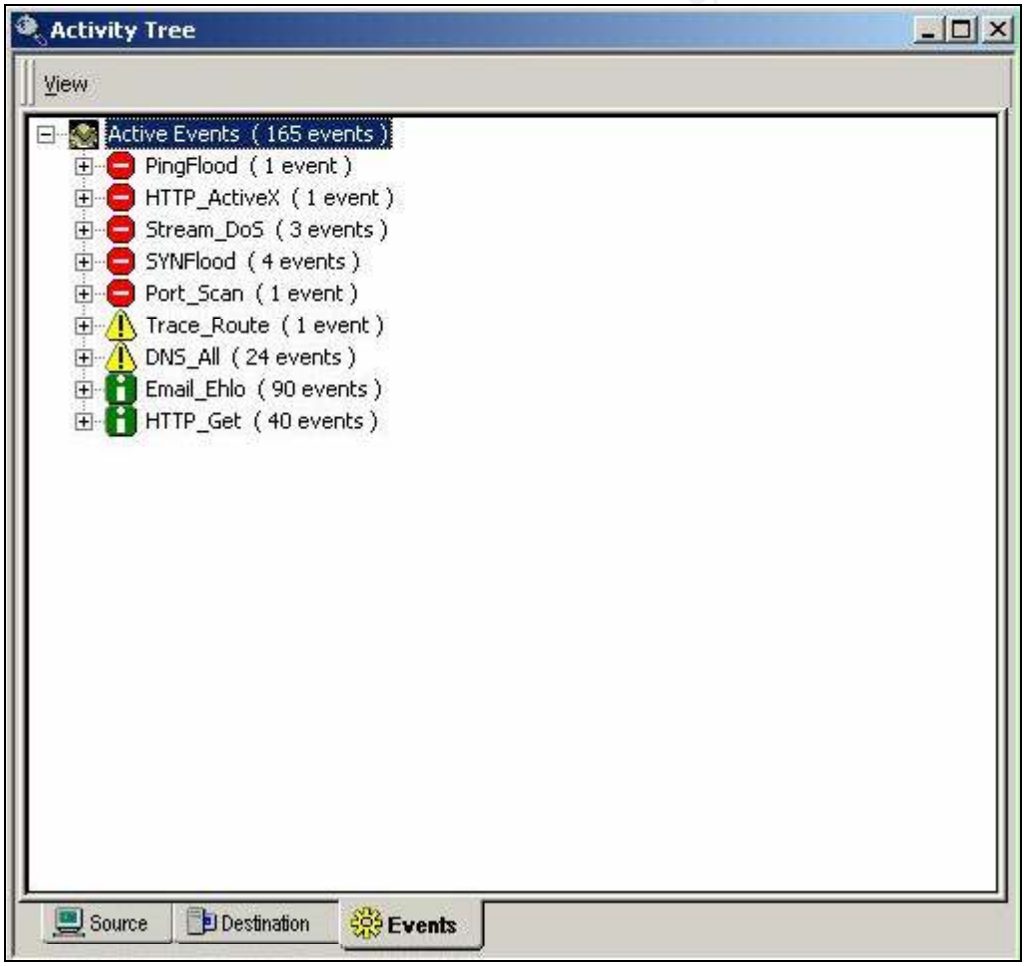

False Positives on Lab Network

Conclusion:

© SANS Institute 2001, As part of the Information Security Reading Room. Author retains full rights.

Intrusion Detection Systems can be a valuable tool when employing a "defense in depth" strategy to your network but diligence is required to ensure success. Spend time filtering out the false positives and always refer to product documentation for complete instructions and explanation of concepts. The RealSecure family of sensors can provide a great deal of security but there are a number of other options available in the commercial marketplace or open source movement. All IDS products have their plusses and minuses so it is important to spend quality research/testing time in order to find a product that will best suit the needs of your network.

References:

[1] Bruneau, Guy. *The Importance of Intrusion Protection.* 2000. August 15, 2001. <http://www.sans.org/newlook/resources/IDFAQ/IPE.htm>

[2] Elson, David. *Intrusion Detection, Theory and Practice.* 2000. August 15, 2001. <http://www.securityfocus.com/focus/ids/articles/davidelson.html>

[3] *ITL Bulletin – Acquiring and Deploying Intrusion Detection Systems.* 1999. August 18, 2001. <http://www.itl.nist.gov/lab/bulletns/nov99.htm>

[4] *Internet Security Systems – Corporate Fact Sheet.* August 18, 2001. <http://www.iss.net/company/profile/fact_sheet.php>

[5] *Internet Security Systems – Company Profile Sheet.* August 18, 2001. <http://www.iss.net/company/profile/>

[6] *Internet Security Systems – The RealSecure Advantage.* August 18, 2001. <http://documents.iss.net/literature/RealSecure/rs_ps.pdf>

[7] *Internet Security Systems – RealSecure FAQ.* 2001. August 18, 2001. <http://documents.iss.net/literature/RealSecure/rs5_0faq.pdf>

[8] Roamer. *Network IDS Sensor Placement.* 2001. August 19, 2001. < http://www.securityhorizon.com/whitepapers/technical/IDSplace.html>

[9] *Glossary of Terms.* September 27, 2001. < http://www.cisco.com/univercd/cc/td/doc/product/iaabu/csids/csids7/unix_cfg/gloss.htm#xtocid721415>

[10] *Intrusion Detection FAQ.* September 27, 2001. < http://packetstormsecurity.org/papers/IDS/switched.htm>

[11] *Century Tap Family.* September 27, 2001. < http://www.areanetworktools.com/ethernet_10_100_1000/century_taps.htm>

[12] *Finisar Systems – Taps and Splitters.* 2001. September 27, 2001. < http://www.finisar-systems.com/products/taps_and_splitters.html#10/100>

[13] *RealSecure Console Users Guide.* September 27, 2001. < http://documents.iss.net/literature/RealSecure/RS_Console_UG_5.5.pdf>

© SANS Institute 2001, As part of the Information Security Reading Room. Author retains full rights.

Last Updated: September 4th, 2015

Upcoming SANS Training

Click Here for a full list of all Upcoming SANS Events by Location

SANS Crystal City 2015	Crystal City, VAUS	Sep 08, 2015 - Sep 13, 2015	Live Event
SANS Network Security 2015	Las Vegas, NVUS	Sep 12, 2015 - Sep 21, 2015	Live Event
SANS Seoul 2015	Seoul, KR	Sep 14, 2015 - Sep 19, 2015	Live Event
SANS Tallinn 2015	Tallinn, EE	Sep 21, 2015 - Sep 26, 2015	Live Event
SANS Baltimore 2015	Baltimore, MDUS	Sep 21, 2015 - Sep 26, 2015	Live Event
Data Breach Investigation Summit & Training	Dallas, TXUS	Sep 21, 2015 - Sep 26, 2015	Live Event
SANS Perth 2015	Perth, AU	Sep 21, 2015 - Sep 26, 2015	Live Event
SANS ICS Amsterdam 2015	Amsterdam, NL	Sep 22, 2015 - Sep 28, 2015	Live Event
SANS Bangalore 2015	Bangalore, IN	Sep 28, 2015 - Oct 17, 2015	Live Event
SANS Seattle 2015	Seattle, WAUS	Oct 05, 2015 - Oct 10, 2015	Live Event
SANS DFIR Prague 2015	Prague, CZ	Oct 05, 2015 - Oct 17, 2015	Live Event
SOS: SANS October Singapore 2015	Singapore, SG	Oct 12, 2015 - Oct 24, 2015	Live Event
SANS Tysons Corner 2015	Tysons Corner, VAUS	Oct 12, 2015 - Oct 17, 2015	Live Event
SANS Gulf Region 2015	Dubai, AE	Oct 17, 2015 - Oct 29, 2015	Live Event
SANS Tokyo Autumn 2015	Tokyo, JP	Oct 19, 2015 - Oct 31, 2015	Live Event
SANS Cyber Defense San Diego 2015	San Diego, CAUS	Oct 19, 2015 - Oct 24, 2015	Live Event
SANS Sydney 2015	Sydney, AU	Nov 09, 2015 - Nov 21, 2015	Live Event
SANS South Florida 2015	Fort Lauderdale, FLUS	Nov 09, 2015 - Nov 14, 2015	Live Event
SANS London 2015	London, GB	Nov 14, 2015 - Nov 23, 2015	Live Event
Pen Test Hackfest Summit & Training	Alexandria, VAUS	Nov 16, 2015 - Nov 23, 2015	Live Event
SANS Hyderabad 2015	Hyderabad, IN	Nov 24, 2015 - Dec 04, 2015	Live Event
SANS Cape Town 2015	Cape Town, ZA	Nov 30, 2015 - Dec 05, 2015	Live Event
SANS San Francisco 2015	San Francisco, CAUS	Nov 30, 2015 - Dec 05, 2015	Live Event
Security Leadership Summit & Training	Dallas, TXUS	Dec 03, 2015 - Dec 10, 2015	Live Event
SANS Milan 2015	OnlineIT	Sep 07, 2015 - Sep 12, 2015	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced