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.

Network Intrusion Detection - Keeping Up With Increasing Information Volume

Threats to the security of a company's key business information come from many different sources. These range from natural disasters to accidental destruction or alteration malicious from activities of people inside and outside the company. The security of key business information stored on computer workstations and servers that are accessible through a local or wide area network can be enhanced through the use of various network security tools. The tools form a network security strategy called defense-in-depth. It tak...

Copyright SANS Institute
Author Retains Full Rights

AD

Network Intrusion Detection – Keeping Up With Increasing Information Volume

Timothy Weber

GSEC Practical Version 1.2f (amended August 13, 2001)

Introduction

Threats to the security of a company's key business information come from many different sources. These range from natural disasters and accidental destruction or alteration, through the malicious activities of people inside and outside the company. The security of key business information stored on computer workstations and servers that are accessible through a local or wide area network can be enhanced through the use of various network security tools. The tools form a network security strategy called defense-in-depth. No tool by itself will secure a network. It takes firewalls, access lists in routers, network scanners, security policy, host-based Intrusion Detection Systems (IDS), and other security devices all working together to secure a network. If an attack gets through one security device, then the idea is another level will catch it. Another critical part of a defense-in-depth strategy is a tool this paper focuses on called a network-based IDS. Innella and McMillan wrote a good article describing what an IDS is. They define a network-based IDS as a device (hardware or software) that detects possible intrusions onto a network by analyzing the data traversing the network and then notifying the proper individuals upon detection.[1] This paper will detail ways to help a network-based IDS cope with the ever increasing volume of information that threatens its ability to fulfill its role in a defense-in-depth strategy.

The Threat

The raw amount of data traversing networks is increasing as companies move from 10 Mbps to 100 Mbps to 1 Gbps speeds to keep up with the increase in information that needs to travel across their network. In addition to increases in local area network bandwidth, metropolitan area network (MAN) bandwidth is increasing also. MAN access speeds are on the increase feeding more and more data into WEB servers and other devices accessed from the Internet. MAN interconnect speeds of 2.5 Gbps are common and 10 Gbps are on the way. Some people suggest that there is more bandwidth available than there is traffic to fill it. The term used to describe this situation is called the bandwidth glut. Others don't believe a glut exists because traffic continues to increase at an unprecedented rate. One of the Internet's founding fathers, Dr. Lawrence Roberts, completed research that "suggests traffic on the Internet has been increasing as much as an unprecedented four times annually through the first quarter of 2001."[2] This increase has a direct, negative impact on a network based IDS tool, which in turn increases the vulnerability of critical business data.

This not-so-obvious vulnerability to business data comes from the increase in the number of bits traveling across a network. An increase in bits results in an increased load on the network protection devices in general, and the network-based IDS in

[1] Innella and McMillan
[2] Pastore

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

particular. In tests conducted by Mircom, Inc., an independent testing company, network-based IDS speeds for one vendor's product range from "690.86 Mbps detecting 98% of intrusions up to 986.94 Mbps detecting only 44% of intrusions."[3] An article referencing that study states "IDS equipment from other vendors hasn't fared much better in lab tests."[4] What happens to your IDS systems as the amount of data it has to process increases? Will your network-based IDS gracefully degrade in all areas of detection or will it experience an immediate total failure? If it degrades, how fast will it degrade and what kind of indications will it give? If your network-based IDS is installed in-line with your traffic flow so that all your data passes through it, will it fail open allowing all traffic through, or will it fail closed shutting down access to your network? These are questions a network security team must know the answers to because if a network-based IDS is not detecting intrusions, the network it is suppose to protect is vulnerable. The information war is still raging no matter what condition the network-based IDS is in.

A Potential but Non-Optimal Solution

The above picture is a dream come true for the product manufacturer's sales team. Their first thoughts might sound something like this, "I've got just the thing for you; it's a bigger, better, faster box." Their solution could very well be right. A problem with that solution is that at the rate traffic volume is increasing, there eventually comes a point where it is too expensive and too difficult to manage.

In addition to just keeping up with the increased traffic flow going across a network, the increase in raw bandwidth highlights a vulnerability associated with "slow attacks." At ten times the data rate, an attack that took ten days to develop now can be accomplished in one day in the same amount of data. In another scenario, slow attacks can be set to complete in the same length of time but be an order of magnitude harder to detect because they are imbedded in ten times more data. So, not only does the network-based IDS have to process more data, but it has to be able to save it, and it has to find small pieces of an attack inside that increased amount of data. The task becomes significantly more daunting at a 100 times increase in bandwidth and even more daunting at a 1,000 times increase. Vendors have some ideas on how to deal with this problem by moving away from detection based on signature scans to such methods as anomaly detection, target monitoring, and stealth probes to detect suspicious activity.[5] (Research on other detection methods that vendors are looking at is outside the scope of this paper, but it might be a topic for research in a separate paper.)

There are other options in addition to more powerful network-based IDS tools to help mitigate the problem. These options will not completely solve the problem, but they could help a network-based IDS better keep up with the increased traffic rates on higher bandwidth networks.

---

[3] Mier Communications, Inc.
[4] Messmer
[5] Innella and McMillan

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

Weigh the Cost

Before implementing any of the options below, some preparation has to be done. As with any change, there is normally a cost associated with the time, materials, and procedural changes required to implement it. A risk assessment should be done to weigh the costs against the potential impact of not implementing the options. If there is no risk, then there is no need to allocate resources. If there is risk, then an effort must be made to determine if it is worth allocating resources to mitigate it.

If changes are implemented, they will impact the people who use the network. Some people will accept the changes quickly and others will resist or ignore them. The key is to constantly communicate the changes and the benefits that result from the changes to the people they impact. Get them to understand and agree to why the changes are needed and what the benefits are. This is one of those easier said than done tasks, but without buy-in, the changes will meet with limited success or totally fail to be of any benefit.

Turn Down the Traffic Volume

The following options fall into the category of reducing the volume of information that the network-based IDS is required to filter without reducing or cutting off the information that employees or customers need.

Awareness training. First, let employees know what potential damage could be done by not being able to detect an intrusion, and what they can do to reduce the traffic load. Show them different, less bandwidth intensive ways to do their job. A few ideas revolve around email. Make sure replies to emails do not include the original attachments unless they are absolutely necessary. Remember, the person being replied to sent the document out in the first place. Another idea is not to send email in HTML format. Besides the security and privacy issues of email in HTML format, an HTML version of an email is three or more times larger than the plain text version. If there are embedded graphics and pictures, it is even worse. Make users aware of what they can do because an informed user is a great asset on the information warfare landscape.

Enforce existing usage policies. Network access and usage policies state how users are to interact with the network. For the purposes of this paper, it is assumed that the policies in place are good policies and they strike a balance between security and business needs. (Whether that is true or not is a good topic for additional research, but it is outside the scope of this paper.) Some policies are written to specifically identify what is allowed, and to deny all else. Other policies are written just the opposite; allow everything unless it is specifically denied. In either case, some policies were put in place to ensure bandwidth is available for business purposes. To determine if a network-based IDS is using cycles looking at traffic that should not be there in the first place, do network auditing and review log files.

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

Divert what cannot be analyzed. There are certain packets that network-based IDS tools cannot analyze for intrusions such as packets that contain encrypted data. These packets cannot be analyzed because the network-based IDS does not have the encryption keys needed to decrypt the information. There could be malicious code encrypted in the payload portion of the packet and the network-based IDS would not be able to identify it.[6] So as much as possible, divert this traffic around the network-based IDS with a router or some other upstream device. If the encryption being used is IPSec or SSL, the upstream device could identify it by simply looking at the packet header.

Just because packets are diverted around the network-based IDS, it does not mean the network is unprotected. The previously mentioned defense-in-depth concept in network security comes into play here. There should be other devices in a defensive posture such as virus scanners and host-based IDS tools to identify malicious code. Those devices are placed in the network where they see the traffic after it has been decrypted. The other defense-in-depth tools also protect against packet header based attacks such as SYN attacks. They accomplish this by analyzing the packet headers of packets that are diverted around a network-based IDS tool because they contain encrypted data.

Unfortunately, some encryption cannot be diverted around the network-based IDS. An example is packets carrying an encrypted file that is being FTPed from one machine to another. There is nothing in the header of those packets that indicate the payload is encrypted. In that case there really is not a way to keep the network-based IDS from spending cycles looking through the packet.

Block unwanted traffic. Anything that is not allowed on a network should be blocked as close to the first line of defense as possible. If connections are not allowed to originate from outside the network, then block them at the first device they can be blocked with. This device could be the Internet router or the first firewall. There is no need to check traffic for intrusions with a network-based IDS (or any other security device in the defense-in-depth set of tools) if the traffic is not supposed to be there in the first place. Another application of blocking unwanted traffic is if there is an area of the network with a network-based IDS protecting it, and that area should only be accessed from internal IP addresses or from a business partner's address space. If only certain IP addresses should be on a network segment, then only allow those IP addresses in and throw everything else away. The key point is to throw the unwanted traffic away before it gets to the network-based IDS.

Use information caches. Both proxy and client caches are useful in the effort to reduce traffic. "Web Caching is the act of storing copies of Web pages on a 'local' system. If the same pages are requested at a later time, and the cached copy is still valid, there is no need to contact the origin server again."[7] The cache could be on a separate server that is pointed to by client web browsers and used as a proxy, or it can be the built in cache that is part of most web browsers. By implementing caches, users will download information from a web page, the information will be checked by the

[6] Allen, et al., p. 54.
[7] Wessels

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

network-based IDS, and it will be displayed to the user. In addition, the downloaded information will be stored in the cache. The next time that information is requested, it comes from the cache and does not have to be checked by the network-based IDS. This only works if the cache server is positioned after the network-based IDS in the traffic flow, or if users are using the cache capability in their web browsers. Web browsers can be configured with how much disk space to use to store downloaded web pages. If this value is too low, it may need to be raised to be of any benefit.

Remove unused protocols. Preconfigured, out-of-the-box computer systems normally have multiple network protocols installed that are not needed. Uninstalling those protocols will reduce the load on a network and on a network-based IDS. If a protocol is not running on a network then obviously the network-based IDS will not have to spend cycles analyzing its packets. Removing protocols that the network-based IDS will not analyze will not provide as much benefit as removing those protocols that it will analyze. It is still a good security practice to remove them though. Most network-based IDS systems only analyze IP traffic and not AppleTalk, NetBEUI, or IPX. Network-based IDS tools normally consist of a packet collection engine and a packet analysis engine. The packet collection engine will pull all packets off of the network, including AppleTalk and IPX packets, but it will not send the packets containing protocols that the analysis engine cannot analyze to the analysis engine itself. So pulling the unused protocols off the network means the packet analysis portion of the network-based IDS tool will have fewer packets to process.

Other ideas

There are other potential methodologies to help a network-based IDS keep up with increased traffic flow. These might be topics for further research to see if they are feasible. They fall into the category of using existing network-based IDS tools more efficiently.

> **Divide and conquer the information streams**. Split the packet streams so that one network-based IDS looks only at HTTP traffic and another looks at FTP traffic. Network-based IDS tools may run more efficiently if they are configured to look at fewer protocols. This idea might mean more network-based IDS systems will be required because the packets that are not analyzed by the specialized tools still need to be analyzed.
>
> **Use load balancers**. Load balance the information streams across multiple network-based IDS systems. There are equipment manufacturers that have specialized devices for this purpose.
>
> **Use dynamic firewall rules**. Dynamic firewall rules can disrupt an intrusion by blocking all packets from a particular site for a finite period of time. There are problems that have to be overcome with this option such as the potential to accidentally block traffic from legitimate customers with legitimate business. This option also may not work well with distributed attacks.

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

Conclusion

Information load on networks and the threats to that information are both increasing. One security tool that is suffering under the increased traffic load, resulting in a reduction in its effectiveness, is the network-based IDS. Sometimes the only solution is to install a bigger, better, faster box; but at other times there are alternatives. This paper contains some alternatives that will reduce traffic and relieve some of the pressure on network-based IDS tools. Those alternatives are awareness training, enforcing existing usage policy, diverting what cannot be analyzed, blocking unwanted traffic, using information caches, and removing unused protocols. They can be used to bolster a defense-in-depth strategy which in-turn will help in the battle for control of the network.

References

Allen, Julia; Christie, Alan; Fithen, William; McHugh, John; Pickel, Jed; and Stoner, Ed. "State of the Practice of Intrusion Detection Technologies." CMU/SEI-99-TR-028. January 2000. URL: http://www.cert.org/archive/pdf/99tr028.pdf (December 18, 2001).

Innella, Paul and McMillan, Oba. "An Introduction to Intrusion Detection Systems." December 6, 2001. URL: http://www.securityfocus.com/infocus/1520 (December 16, 2001).

Messmer, Ellen. "Intrusion Alert – Gigabit-speed Intrusion-detection Systems Miss Attacks on Faster Nets." December 3, 2001. URL: http://www.nwfusion.com/news/2001/1203ids.html (December 16, 2001).

Mier Communications Inc. "Lab Testing Summary Report." Product Category: Intrusion Detection Systems, Vendor Tested: Intrusion.com, Product Tested: Intrusion.com SecureNet Gig. Report number 300401. April 2001. URL: http://www.mier.com/reports/intrusion/Intrusion-comPerfVal-5-04-01.pdf (December 16, 2001).

Pastore, Michael. "Internet Traffic Continues to Grow, Despite the Glut." August 15, 2001. URL: http://cyberatlas.internet.com/big_picture/traffic_patterns/article/0,1323,5931_866931,00.html (December 16, 2001).

Wessels, Duane. "IRCache FAQ and Users Guide." 2001. URL: http://www.ircache.net/FAQ (December 18, 2001).

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

Last Updated: June 20th, 2013

Upcoming SANS Training

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

Digital Forensics & Incident Response Summit 2013	**Austin, TXUS**	**Jul 09, 2013 - Jul 16, 2013**	**Live Event**
SANS London Summer 2013	**London, GB**	**Jul 09, 2013 - Jul 16, 2013**	**Live Event**
SANS Rocky Mountain 2013	**Denver, COUS**	**Jul 14, 2013 - Jul 20, 2013**	**Live Event**
SEC528 SANS Training Program for the CompTIA® New Advanced Security Practitioner (CASP) Certification	**Washington, DCUS**	**Jul 22, 2013 - Jul 26, 2013**	**Live Event**
SANS Mumbai 2013	**Mumbai, IN**	**Jul 22, 2013 - Jul 27, 2013**	**Live Event**
SANS San Francisco 2013	**San Francisco, CAUS**	**Jul 29, 2013 - Aug 03, 2013**	**Live Event**
SANS Boston 2013	**Boston, MAUS**	**Aug 05, 2013 - Aug 10, 2013**	**Live Event**
SANS SEC 560: Network Penetration Testing @ Bangalore 2013	**Bangalore, IN**	**Aug 05, 2013 - Aug 10, 2013**	**Live Event**
Critical Security Controls Summit	**Washington, DCUS**	**Aug 12, 2013 - Aug 18, 2013**	**Live Event**
Industrial Control Systems Security Training - DC	**Washington, DCUS**	**Aug 12, 2013 - Aug 16, 2013**	**Live Event**
SANS Virginia Beach 2013	**Virginia Beach, VAUS**	**Aug 19, 2013 - Aug 30, 2013**	**Live Event**
SANS Thailand 2013	**Bangkok, TH**	**Aug 19, 2013 - Aug 31, 2013**	**Live Event**
SANS Melbourne 2013	**Melbourne, AU**	**Sep 02, 2013 - Sep 07, 2013**	**Live Event**
SANS Capital City 2013	**Washington, DCUS**	**Sep 03, 2013 - Sep 08, 2013**	**Live Event**
Network Security 2013	**Las Vegas, NVUS**	**Sep 14, 2013 - Sep 23, 2013**	**Live Event**
SANS Canberra 2013	**OnlineAU**	**Jul 01, 2013 - Jul 13, 2013**	**Live Event**
SANS OnDemand	**Books & MP3s OnlyUS**	**Anytime**	**Self Paced**