



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>

A Thousand Heads Are Better Than One – The Present and Future of Distributed Intrusion Detection

Robert Zuver
GSEC v1.3
April 30, 2002

Summary

The rapid increase in worldwide Internet activity in the past half-decade has given rise to a host of new network security threats. Until recently these threats have been (more or less) successfully combated with a combination of antivirus software, firewalls, and intrusion detection systems. But the latest generation of distributed denial of service (DDoS) attacks and Internet worms has demonstrated the shortcomings of traditional host- and network-based intrusion detection systems: incomplete information and inadequate user knowledge.

One powerful remedy for these shortcomings is Distributed Intrusion Detection (DID), which facilitates the consolidation of intrusion detection information from many different individual sources. This aggregation of information allows the potential victims of malicious network activity to differentiate between harmless anomalies and actual attacks, and provides Internet service providers with the information and motivation they need to pursue and shut down hackers and worm-infested computers. Today's most popular DID systems, DShield and myNetWatchman, are an important first step towards realizing the full potential of distributed intrusion detection, but future implementations will need to address the issues of confidentiality, compatibility, and education if DID is to be widely adopted.

Introduction

One of the most interesting technological phenomena of the past few years has been the phenomenal increase in global connectivity. Since January 2000, worldwide Internet access has more than doubled.¹ In the United States alone, 21 million people have access to "always on" broadband Internet connections at home.² Corporations, government agencies, and educational institutions have all increased their Internet connectivity as well.

As a result of this prodigious increase in Internet access, attacks by hackers and

¹ Nua Internet Surveys.

² Miller.

network-aware viruses (also known as “worms”³) have skyrocketed. In 2001 the CERT Coordination Center reported 52,658 network security incidents, an increase of more than 140 percent over the previous year.⁴ A significant portion of this increase is due to the proliferation of Distributed Denial of Service (DDoS) attacks and worms such as “Code Red,” which use previously infected systems as additional vectors of attack against new victims.

Defense against these and other types of attacks falls into three categories:

- Antivirus software
- Hardware and software firewalls
- Intrusion detection systems

This paper will focus on intrusion detection systems in general, and specifically on two examples of the most promising new weapon in the battle against Internet hackers and worms: distributed intrusion detection.

Intrusion Detection Systems

The basic purpose of an intrusion detection system (IDS) is to detect and identify suspicious network traffic. This can be done for traffic to- and from a specific computer (host-based IDS) or all traffic on a network segment (network-based IDS).⁵ In either case, the IDS monitors systems and network segments for unusual or unauthorized traffic and reports this information to the network administrator. The administrator can then make changes to his or her firewall rules to protect against new types of attacks. Administrators also make use of IDS logs when reporting malicious network activity to the offender’s Internet service provider (ISP).

As broadband Internet access has proliferated, many home users have also begun using software packages such as ZoneAlarm and BlackICE Defender, which combine a personal firewall with intrusion detection capabilities.

While both traditional IDS and the new ‘IDS + firewall’ products perform their intended functions well, the information they provide is necessarily incomplete. Host-based IDS can only monitor and report suspicious activity on a single computer, and network-based IDS can report only on a specific network segment.

Sophisticated hackers and well-programmed worms can use this limitation to

³ Symantec AntiVirus Research Center.

⁴ CERT Coordination Center.

⁵ Lehmann.

hide their activity through various means, such as scanning the Internet non-sequentially and only scanning each IP address once. These methods provide the administrator of a particular host or network with little information as to whether a particular probe is a harmless anomaly or the product of a concerted attack.

In addition, most home Internet users lack the knowledge and experience necessary to properly analyze their IDS and firewall logs and report abuse to the proper authorities. And in many cases, when they do report them, their individual reports are discounted or ignored by harried ISP staffs.

As a result of these limitations, much malicious Internet activity goes undetected and/or unpunished.

The Next Step: Distributed Intrusion Detection

These limitations of traditional IDS's have led to the creation of what is sometimes called "distributed intrusion detection" (or, alternatively, "consensus intrusion detection"). Distributed Intrusion Detection (DID) is the aggregation of ID data from many different sources to provide a broader and more detailed picture of malicious network activity at any given time.

The promise of distributed intrusion detection is well-illustrated by this account from SANS's Internet Storm Center, which describes a worm attack that began on March 22, 2001:

In the March 22 wave of attacks, thousands of organizations that had not updated their version of BIND were rewarded by being infected with a worm called Lion. Lion stole password files from infected machines and sent them to a site in China, and it installed a distributed denial of service tool so that the infected machines could be used in denial of service attacks. It also installed other malicious software and then forced each infected system to search the Internet for additional vulnerable machines to infect.

Had the worm attack happened even a few months earlier, the intrusion detection systems that recorded the probes would have been acting alone. There would have been little chance of raising a widespread alarm to offer an early warning to those who had vulnerable systems, and nothing would have been done for at least a couple of days, and probably much longer. During that delay, tens of thousands more systems would have been compromised and used by the worm.

But this time something different happened.

Hundreds of those intrusion detection sensors that were logging attacks had become part of regional and industry-specific security monitoring networks. They sent their logs to analysis sites in Cambridge, Atlanta, Reston, and Indianapolis. There the data was aggregated and charted automatically, and posted for analysis. Analysts immediately saw a spike in the number of attacks on DNS Port 53. One regional analysis showed an increase from two hundred probes per day throughout March to 50,000 probes on the 22nd. Some kind of man-made, electronic storm was sweeping through the Internet.

Within an hour, the analysts, all of whom were fully qualified as GIAC certified in intrusion detection, agreed that a global security incident was underway. They named a global incident handler who immediately sent a notice to a global community of technically savvy security practitioners asking them to check their systems to see whether they had experienced an attack. Within three hours a system administrator in the Netherlands responded that some of his machines had been infected, and he sent the first copy of the worm code to the analysts. Over the next 24 hours twenty-five more sites from Canada, Brazil, Venezuela, the US and the UK would report that they had been infected, but that first copy of the worm was enough for rapid action.⁶

Armed with this information, the analysts were able to quickly develop and rapidly distribute a patch to detect and fix the problem in infected computers and prevent it in uninfected ones. They also notified law enforcement.⁷

This capacity for rapid detection and cooperative action is one of the most attractive features of distributed intrusion detection. By providing administrators with a bigger picture of network activity than they would otherwise have, and by creating a communication channel for administrators across the globe, DID has the potential to provide for a quicker, more accurate, and more thorough response to incidents like the LION worm attack.

DID can also prove useful in cases of smaller-scale attacks. The people in charge of the distributed systems are better qualified to examine firewall and IDS logs and separate legitimate network traffic from suspicious traffic. ISPs are also more likely to take seriously the aggregated complaints of dozens (or hundreds) of users, sent to them from a qualified network administrator, than a single complaint from a regular user.

There are currently two popular DID programs, DShield (www.dshield.org) and myNetWatchman (www.myNetWatchman.com). We will now examine these programs, and their strengths and weaknesses.

⁶ SANS Institute, "Internet Storm Center."

⁷ Ibid.

DSshield

Founded in November 2000 in response to the indifference of many network administrators to individual reports of malicious network activity,⁸ DSshield (www.dshield.org) now receives nearly 30 million reports per month.⁹

DSshield provides free client software for use with most popular Windows firewall products, including ZoneAlarm and BlackICE Defender as well as the McAfee, Sygate and Norton personal firewall products. The DSshield client also supports logs from ipchains, Snort, and several other Linux/UNIX firewalls, as well as logs from Cisco routers. The client converts these logs into the standard DSshield format and sends them via e-mail to DSshield.

Users of ipchains, iptables, ZoneAlarm, SonicWall, and Raptor can also submit their logs by pasting them into a web form or simply e-mailing them to report@dshield.org, with the option to use PGP encryption for added security. For users of non-supported firewall and ID products, DSshield provides tools to assist in the creation of custom log-conversion scripts.

After receiving firewall data from a client, DSshield aggregates it with data from all of the other clients to produce a number of reports and graphs, including a geographical breakdown of attacks and “top ten” lists of offending IP addresses and most-attacked ports. These reports are then made available for viewing on the DSshield web site. Users can also run searches for activity from a particular IP address or subnet, or all activity on a given port.

Users who register their e-mail addresses with DSshield (which is optional) can also access a series of customized reports, including a list of the last fifty incidents reported by that user. Within this report, clicking on an IP address brings up a report with basic information about the offending system, and clicking on the port number displays a graph detailing the number of attacks on that port over the last 30 days.

Registered users also have the option to join the “FightBack” program, which gives DSshield permission to forward the user’s log entries (including IP address and e-mail address) to an attacker’s Internet service provider. Without this permission, DSshield keeps all individual log entries confidential, using them only for data aggregation purposes.

⁸ Lemos.

⁹ DSshield.

myNetWatchman

Like DShield, myNetWatchman (www.mynetwatchman.com) aggregates intrusion data from a number of different sources. MyNetWatchman was created with three goals in mind:

1) Minimize effort required to report events

Backtracing attack events to their source and emailing the responsible party is extremely labor intensive. Many ISPs have their own reporting preferences that aren't immediately obvious. Tracking down foreign sources (e.g. Korea, Taiwan, China, etc...) are even more challenging due to the frequent lack of backtrace data (e.g. DNS, Whois, etc.).

2) Avoid false reports

Personal firewalls are essential, however, they are notorious for generating "attacks" that are completely bogus. ISP abuse departments are already overwhelmed with SPAM complaints. The last thing we need to do is flood them with port scan complaints that aren't real or aren't serious enough to warrant escalation. Ideally, we need to focus on attacks that clearly indicate a compromised host, or a hacker who is doing very broad port scans over hundreds, thousands or even millions of addresses.

Problem is, a single firewall log doesn't provide enough perspective to differentiate between false and real attacks and they can never indicate the breadth of the attack.

3) Provide aggregated attack report to responsible party

Depending on the type of attack, an ISP or system owner often needs supporting evidence from multiple sources before action is warranted. However, even if many people report the same source IP address, the responsible party may lack the tools to correlate this information and recognize a pattern. Ideally we need to combine multiple evidence sources into a single incident e-mail escalation. This minimizes the number of individual email reports and enables immediate action by the recipient.¹⁰

MyNetWatchman's free client software supports a number of firewall and intrusion detection products, though not as many as DShield (specifically the Tiny, Sygate and Norton personal firewall products are not supported). There is no provision for submitting logs via e-mail or the web, although there is a web

¹⁰ myNetWatchman, "myNetWatchman Vision."

form for entering single incidents.

Unlike the DShield client, which requires the user to launch the program whenever he wishes to submit log entries, myNetWatchman runs in memory and submits log entries as they happen, in near-real-time. The entries are then aggregated to produce reports similar to those produced by DShield. When a certain threshold of suspicious activity has been recorded from a single IP address, the myNetWatchman system automatically generates an escalation e-mail to the appropriate ISP.

Included in this e-mail are the dates and times of each suspicious probe, as well as a list of the IP addresses that were probed (with the last two octets obscured for privacy—e.g., 129.219.x.x). These escalation e-mails are formatted in a standard way, allowing ISPs to automate their processing of them and (perhaps) reduce response times.¹¹ The e-mail addresses of the reporting clients are not disclosed.

Users can log into the myNetWatchman web site to access a number of reports similar to those offered by DShield, but myNetWatchman's reports provide more, and more useful, information than DShield's. In addition to basic information (also provided by DShield) such as DNS lookup and total event count (i.e., the number of probes recorded from a particular IP), myNetWatchman's Incident Detail reports also include the date and time of every probe recorded from the suspect IP address by *any* myNetWatchman client and a copy of all correspondence between myNetWatchman and the ISP.

Each suspect probe in the list also includes a link to information about the nature and relative risk of the probe, based on the protocol and port. While DShield also makes this information available on its web site, it is not linked directly from the incident reports and takes some digging to locate.

As an aside, myNetWatchman's custom reports load much more quickly than DShield's, which can slow to a crawl (or simply refuse to load) at times of peak usage.

DShield vs. myNetWatchman

So which DID system should one use?

Network administrators and some "power users" may prefer DShield's greater compatibility, as well as its anonymity (since individually identifiable reports will not be forwarded if the administrator does not opt into the "FightBack" program). The fact that the client is not memory-resident also gives the administrator

¹¹ Arquette.

additional control over which log entries to submit, and when. If a Linux or UNIX administrator prefers automatic submission, she can configure a cron job to process and submit the logs at a chosen interval.

Individual home users are more likely to prefer the convenience of myNetWatchman's automated log submission. In addition, myNetWatchman's approach to privacy protection (automatic submission of all logs, but without disclosure of the client's e-mail information and partial suppression of its IP address) seems to strike a better balance for most users than DShield's all-or-nothing approach. And the greater detail in myNetWatchman's incident reports should be especially useful to users who are less security-savvy.

Obstacles to Overcome

If Distributed Intrusion Detection is such a good thing, why isn't everyone participating?

There are three main obstacles to be overcome if DID is to achieve wider acceptance:

1) Confidentiality. Many corporations, network administrators and even home users are reluctant to allow the release of potentially sensitive information like e-mail and IP addresses. DShield and developers of future DID systems would do well to adopt myNetWatchman's approach to privacy protection, which would likely result in greater participation by individuals and organizations alike while still retaining DID's notification and enforcement advantages. Conversely, future versions of the myNetWatchman client should provide an encryption option like DShield's PGP.

2) Compatibility. DShield sets a good example in the area of client compatibility, with its array of submission options and tools for users to develop their own conversion scripts. But compatibility between DID systems is also important—ideally the data from various DIDs would be accumulated into one “metaDID” for the greatest possible detail.

In fact, steps are already being made in this direction; both DShield and myNetWatchman submit their data to the Griffin Project, which stores the information in a master database.¹²

3) Education. Many people simply are not aware of the need for (or the existence of) distributed intrusion detection. In fact, half of all home broadband connections are not even protected with a firewall.¹³ But as high-speed Internet connections become more common, the level of sophistication of the average

¹² SANS Institute, “Welcome to the Griffin Project Site.”

user is likely increase as well.

Conclusion

The history of network communications has been one of increasingly sophisticated attacks, countered by ever more advanced defensive tools and techniques. Distributed denial of service attacks and Internet worms represent the latest offensive weapons in this constantly escalating battle.

Distributed Intrusion Detection shows great promise as a defensive tool against these new weapons. While there are obstacles to its wider adoption, they are not insurmountable, and the successes already achieved by the first generation of DID systems demonstrate that overcoming them will be worth the effort.

¹³ Thorsberg.

References

Arquette, Brett. "A 'Watchman' to Stand Guard Over the Net." eWeek. 17 Sep. 2001.

URL: <http://www.eweek.com/article/0,3658,s=1868&a=14859,00.asp> (31 Mar. 2002).

CERT® Coordination Center. "Cert/CC Statistics 1988-2001." 10 Jan. 2002.

URL: http://www.cert.org/stats/cert_stats.html (29 Mar. 2002).

DShield.org. 31 Mar. 2002.

URL: <http://www.dshield.org> (31 Mar. 2002)

Lehmann, Dirk. "What is ID?" Intrusion Detection FAQ.URL:

http://www.sans.org/newlook/resources/IDFAQ/what_is_ID.htm (29 Mar. 2002).

Lemos, Robert. "The Net's new Neighborhood Watch." ZDNet News. 5 Dec. 2000.

URL: <http://zdnet.com.com/2100-11-503197.html> (31 Mar. 2002).

Miller, Leslie. "Web growth slows, but time online rises." USA Today. 28 Mar. 2002.

URL: <http://www.usatoday.com/life/cyber/tech/2002/03/28/net-statistics.htm> (29 Mar. 2002).

MyNetWatchman.com. "myNetWatchman Vision."

URL: <http://www.mynetwatchman.com/vision.htm> (31 Mar. 2002).

Nua Internet Surveys. "Worldwide." How Many Online?

URL: http://www.nua.ie/surveys/how_many_online/world.html (29 Mar. 2002).

SANS Institute. "Internet Storm Center."

URL: <http://www.incidents.org/isw/iswp.php> (31 Mar. 2002)

SANS Institute. "Welcome to the Griffin Project Site."

URL: <http://www.incidents.org/cid/partners.php> (31 Mar. 2002)

Thorsberg, Frank. "Half of U.S. Broadband Users Unprotected." PCWorld.com.

16 Jul. 2001. URL: <http://www.pcworld.com/news/article/0,aid,55154,00.asp> (31 Mar. 2002).

WanWall.com. "Defeating ddos attacks."

URL: http://www.wanwall.com/about/ddos_explained.html (29 Mar. 2002).