



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.

Denial of Service Attacks and the Emergence of "Intrusion Prevention Systems"

Firewalls and Intrusion Detection Systems (IDS) have been the mainstay of network security perimeters for many years and have evolved over time with increasing sophistication and technological advance to maintain protection of Enterprise Networks. These systems however are bearing the brunt of increased Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks from across the globe[1,2]. Where possible new techniques and technologies should always be considered to provide additional defences to prevent t...

Copyright SANS Institute
Author Retains Full Rights

AD



Denial of Service attacks and the emergence of “Intrusion Prevention Systems”

SANS GSEC Practical Assignment v1.4b
Option 1 (Re-Submission)
Adrian Brindley
November 1, 2002

Abstract

Firewalls and Intrusion Detection Systems (IDS) have been the mainstay of network security perimeters for many years and have evolved over time with increasing sophistication and technological advance to maintain protection of Enterprise Networks. These systems however are bearing the brunt of increased Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks from across the globe[1,2]. Where possible new techniques and technologies should always be considered to provide additional defences to prevent these devices and the Enterprise Network itself from becoming overwhelmed during such attacks.

The objective of this paper is to give a review of DoS / DDoS attacks, provide a list of basic network attack prevention techniques, provide a brief comparison of current and emerging Intrusion Prevention devices available and to give an example implementation scenario using one of these products.

Introduction

Since the first real world Denial of Service attacks that were publicised in early 2000 that caused significant financial losses to several major e-Business giants these attacks still persist and are constantly evolving in complexity and increasing frequency[9].

Whilst the Internet core providers and wider Internet community as a whole work towards filtering techniques as suggested in RFC 2827 to prevent IP spoofing attacks we all remain potentially at risk. This fundamental filtering change will obviously take time to become effective but in the meantime we need to maintain a vigilant approach and look at alternative methods to enhance the “defence in depth” posture for Enterprise Networks. Some ISP's are moving towards this Denial of Service filtering goal but at the same time many others do not want to provide this preventative measure because of the costs, technical requirements and legal ramifications that could follow if a client ever became the victim of such an attack. Ultimately your own organisation may have to face and control this menace using only the “assistance” of your ISP.

Overview of Denial of Service Attacks

There are many manifestations of Denial of Service attacks but they ultimately have the same objective – to deny or degrade a user's ability to legitimately access network or host based services. DoS attacks accomplish this by exhausting the limited resources of network bandwidth by packet flooding or exhausting host resources by consumption of CPU cycles, random memory, static memory or data structures[3].


DoS attacks can generally be classified as either a Flood Attack or a Malformed (or crafted) Packet Attack and that where attacks originate simultaneously from several compromised sources that these can be classified as Distributed DoS attacks[4].

Fundamental to the IP protocol every packet has a source and destination address field that is used to determine the originating and destination end points. The process of forwarding these packets by intermediate routers partly relies on the destination field; the source address will only be used when a response to the packet is required. This makes the implementation of DDoS flooding attacks easy to accomplish because fake or "spoofed" source addresses can be used, and packets will generally be forwarded unchallenged to the specified destination. This allows a DoS or DDoS attack to be carried out from any location and with total anonymity.

If an attack is underway from a single address then it is possible to arrange for a "block" of the offending source IP address at the ISP or the border router. However, when a DDoS attack occurs the problem is not as easy to resolve because packets appear to be coming from hundreds or even thousands of different hosts, there is absolutely no point trying to implement temporary Access Control Lists on routing devices or modify the border Firewall rulebase, it is too late – you are left at the mercy of the attack underway. The types of attack can also take the form of a single "one shot" crafted packet originating from a single host to thousands of packets per second originating simultaneously from multiple hosts.

Common DDoS attacks

The normal DDoS attack architecture works upon the basis that the required hosts to launch the attack from have already been identified and compromised via Trojans or "backdoors"[5]. In a DDoS scenario the Intruder (also called the Attacker or Client) issues control traffic to the Master (also called the Handler) which, in turn then issues commands to the Daemon (also called an Agent, Broadcast program or Zombie). The Daemons that are at the end of this command chain finally initiate the attack traffic against the Victim. This distributed architecture increases the attack capability many times over and allows the Intruder the means to remain undetected as shown in Figure 1 below.



Examples of Flooding Attacks

Floodnet is a Java application that sends requests to non-existent pages on target hosts and queries to search engines, the target host will reply with various responses that inevitably increase network traffic and causes consumption of valuable CPU and memory resources by these inappropriate requests.

SMURF attacks involve flooding a target network with Internet Control Message Protocol (ICMP) "echo" response messages that are directed at the network broadcast address and with the source address being configured as the "victims" IP addresses (spoofed address)[6]. This attack system can "amplify" the original packet transmission a hundred or even a thousand-fold.

Trin00 is a DDoS SYN flood attack type and was one of the earliest attacks to be seen on the Internet, default ports used for communications between each component are 27665/tcp, 27444/udp and 31335/udp.

Tribe Flood Network (TFN) encompasses a range of DDoS attacks such as ICMP flood, SYN flood, UDP flood, and SMURF style attacks, TFN itself uses the ICMP protocol echo and echo reply to communicate between the masters and daemons.

Stacheldraht is a DDoS tool that uses encrypted communications between the intruder and master; it also allows automated agent updates via "rcp" type commands (514/tcp). Attacks are similar to the TFN type attacks – ICMP, SYN, UDP floods and SMURF, communication ports in use are 16660/tcp, 65000/tcp and ICMP Echo Reply.

Trinity uses a number of flooding attacks including SYN, RST, ACK, UDP, fragment and other flood types. Communications from the intruder to the daemon is accomplished via Internet Relay Chat (IRC) channels or ICQ. This attack uses port 6667/tcp and retains a backdoor capability on port 33270/tcp.

Targa3 sends malformed TCP/IP packets (fragmented, packet size, invalid option flags) to a target system causing it to crash or become unstable.

Mstream uses packet flooding attacks based upon TCP packets with the ACK flag set, the default ports used are 6723/tcp, 7983/udp and 9325/udp.

Shaft uses a TCP SYN packet flooding attack mechanism, the DDoS daemon and the client controls the size of the flooding packets and duration of the attack. Shaft uses ports 20432/tcp, 18753/udp and 20433/udp to communicate.

TFN2K incorporates advanced features to disguise communications between the master and daemon, it allows TFN2K traffic over multiple protocols such as TCP, UDP and ICMP. TFN2K attacks use flooding techniques and malformed packets (invalid headers) to cause systems to crash or become unstable (also similar to the Land and Teardrop/2 type attacks).

Naptha attacks are based upon the weaknesses of the TCP "state machine" handling that were identified by Bindview's RAZOR team [10]. Briefly, the TCP connection cycle consists of eleven various states - CLOSED, LISTEN, SYN RECD, SYN SENT,

ESTABLISHED, CLOSE WAIT, LAST ACK, FIN WAIT -1, FIN WAIT -2, CLOSING, TIME WAIT that occur during a session. Naphtha attacks are based upon IP stack or application exhaustion during state times other than the SYN RECD state (e.g. FIN WAIT -1 or ESTABLISHED). This attack is also referred to as "asymmetric" because resources on the attacking machine are consumed at a negligible rate compared to the victims; the attack machine uses raw packet access that allows direct creation and response to TCP stateful packets.

Reflective type attacks are used to forward large amounts of SYN ACK or RSTs from high bandwidth sites such as major Web sites or Universities back to a Victim's IP address (there is still some debate about whether this is a valid attack because although flooding occurs only small length TCP ACK or RST packets are forwarded from the reflecting site, and if the source address is already "spoofed" then the reflection process is not required).

Examples of Malformed (crafted) Packet Attacks -

Ping of Death uses ICMP ECHO request packets that are larger than the allowed maximum IP packet size to crash or reboot the end target.

Chargen uses the functionality of the UDP chargen (19) and echo (7) services to create an endless loop that can consume a large amount of network bandwidth between the victim and echoing system.

WinNuke also known as an OOB - Out of Bounds attack uses port 139 of Windows95 computers to forward packets to, the result is generally a system crash.

TearDrop (fragmentation) attackers sends multiple fragments that cannot be reassembled properly by manipulating the offset values of packets that eventually causes a reboot or halt of victim system.

Land involves an attacker sending forged packets with the same source and destination IP address. The victims system or IP stack will become unstable with the end result of the system crashing or rebooting.

Examples of Network Worm Attacks -

Network worms can be considered a form of Denial of Service because of the fact that network bandwidth can become saturated due to high system scanning rates once a host has become infected and it tries to replicate itself onto other systems.

Code Red I/II scans random IP addresses on port 80/tcp attempting to locate and replicate through vulnerable Microsoft IIS Web servers (signature example shown below).

```
/default.ida?NNNNN NNNNNNNNN NNNN NNNNNNNNN NNNN NNNNNNNNN NNNN NNNNNNNNN NNNNNNNNN NNNN NNNN  
NNNNNN NNNNNNNNN NNNN NNNNNNNNN NNNN NNNNNNNNN NNNN NNNNNNNNN NNNNNNNNN NNNN NNNN  
NNNNNN NNNNNNNNN NNNN NNNNNNNNN NNNN NNNNNNNNN NNNN NNNNNNNNN NNNN NNNNNNNNN NNNN NNNN  
%u9090%u6858%ucbd3%u7801%u9090%u6858%ucbd3%u7801%u9090%u6858%ucbd3%u7801%u9090%  
u9090%u8190%u00c3%u0003%u8b00%u531  
b%u53ff%u0078%u0000%u00=a
```


technology and even IDS type functionality they remain a “routing” element within the network.

Are Firewalls an Intrusion Prevention System? – Again, yes to a certain extent, Firewalls come in many flavours and levels of hardware and software sophistication, they allow packet (non-stateful), stateful, application and proxy based protection, content scanning and some have even incorporated methods for providing in-built Intrusion Detection and alerting capabilities. Their main purpose within the Enterprise though is to enforce Enterprise policy and maintain connection state information for legitimate users internally or externally and not to prevent high volume DoS/DDoS style attacks.

A “true” Intrusion Prevention System uses dedicated technology to provide increased levels of protection against DoS/DDoS and Network Worm type attacks. Once installed these devices sit there providing continuous 24 x 7 protection to instantly remove any unwanted traffic heading towards your Enterprise Network and hosts.

Requirements

As always there is a list of requirements that must be considered to ensure that this type of technology will be of benefit to your organisation, such as -

- Have you and your ISP implemented all possible methods to address Packet saturation of Wide Area Network links (the DoS/DDoS device within your security perimeter cannot prevent this)?
- Does this enhance your existing security solution and will it allow your existing network and security infrastructure to remain unchanged?
- Does it allow other systems to focus more on their own specific functions (e.g. traffic load balancers, Firewalls, Application Proxies)?
- Does it require advanced administration set-up or additional training costs?
- Does it allow your existing investment in security devices (IDS's and Firewalls) to be retained?
- Can it add beneficial Alert handling and logging functionality to existing Security Management Systems?
- Will it allow outbound traffic analysis and traffic blocking, which will prevent your own site from becoming a launch pad for DoS/DDoS attacks?
- Can initial cost and resource set-up requirements be justified?
- Will the possibility of limited resilience (such as fail-safe backup) and load-sharing functions be acceptable to your organisation?

Current and emerging Intrusion Detection and Prevention systems –

The new Intrusion Prevention System (IPS) products that are now emerging and that will increasingly start to work within security perimeters in the future will incorporate as standard high performance anti-DDoS capabilities and incorporate intelligent auto-update features to ensure that the latest protection is constantly available. The migration to dedicated hardware platforms that use high performance architectures such as Application Specific Integrated Circuit (ASIC) technology will also ensure that performance matches any increased bandwidth needs and that real-time threat analysis and prevention remains feasible at ever increasing wire speeds.

These systems have advanced from the passive DoS / DDoS detection and alerting method to the attack prevention and mitigation stage where packets are actually blocked, and they use varying degrees of detection technology to accomplish this - signature based and traffic anomaly, combinations of both and In-Line (active) versus Tap (passive) operating modes. The purpose of these devices is to automate as far as possible the accurate identification and blocking of malicious traffic and to only allow legitimate user traffic to access your site, some of these systems allow very easy integration into the security perimeter whilst others require routing changes to be implemented.

An example list of "Intrusion Detection and Prevention" systems available today -

Vendor	Product	Detection Method	In-Line or TAP	DoS / DDoS Capability	Price (appx)
Captus	CaptIO	Anomaly detection	In-Line (active)	Detect and Prevent	\$15K
Lancope	StealthWatch	Traffic analysis	Tap (passive)	Detect and Prevent	\$20K
Mazu	Enforcer	Anomaly detection	In-Line and Tap	Detect and Prevent	\$32K
Radware	FireProof	Signature detection	In-Line (active)	Detect and Prevent	\$28K
Top Layer	Attack Mitigator 1.0	Anomaly detection	In-Line (active)	Detect and Prevent	\$10K
Webscreen	WS100	Heuristics	In-Line (active)	Detect and Prevent	\$23K

Table 1.

A brief summary of the products listed in Table 1 -

CaptIO The CaptIO uses anomaly-based network security policies to detect, identify, and stop Denial of Service (DoS) and Distributed Denial of Service (DDoS) using dedicated hardware devices[11]

Stealth-watch is a behaviour-based Intrusion Detection System (IDS) operating on flow-based architecture, this enables Intelligent alarming, advanced network monitoring, gigabit operation, recognizes unknown threats and creates forensic data of network activity[12].

Enforcer uses either traffic collection via passive taps (also port mirroring on switches) or active Enforcers, statistics are sent to a Profiling system. Active Enforcers can also filter traffic to mitigate security threats[13].

FireProof is a Security Application Switch, ensuring the integrity and operation of security devices across the enterprise. FireProof combines load balancing, optimization, and high

availability for firewalls, Virtual Private Networks (VPN) gateways, and Intrusion Detection Systems[14].

Web Screen WS100 uses heuristic algorithms to detect DDoS attacks. WS100 looks at the nature of the access rather than their exact signatures which remove the requirement for signature updates[16].

Example Enterprise implementation scenario using the Attack Mitigator

This Scenario assumes a single Internet connectivity point and is very similar to our own in-house environment that was used to test the capability of this product, when we evaluated the Attack Mitigator unit [15] it did not support propagation of stateful connection to secondary units or redundant power supplies but this now appears to have been resolved with later platforms. The ASIC architecture and CPU used a proprietary Operating System to run the device and we found few issues even with the initial v1.0 release firmware. In use the system constantly updated it's Web interface with traffic statistics that indicated whether traffic was trusted, suspicious or malicious and the packet analyzer also confirmed packets were being discarded correctly.

The location of the product is critical to the whole function of Intrusion Prevention / Mitigation and the unit should be located between the edge routing device and the Enterprise firewall, this ensures that all traffic flows are analyzed correctly and allows the unit to remove the unwanted traffic before it approaches the firewall as shown in Figure 2.

Enterprise to Internet connectivity example employing an Intrusion Prevention device -




Figure 2

Attack Mitigator Technical overview

The Mitigator is a layer 2 switching device and therefore allows transparent installation between the edge router and Corporate Firewall – it is literally plug-and-go, the unit is pre-configured with 15 (v1.0 firmware) of the most common DDoS attack methods and allows easy addition of user-defined signatures via a Web front-end and initial Wizard setup.

There are a number of hardware platforms that provide a range of 10/100/1000Mbps connectivity options but in effect the units offer a combination of 2 fixed Internal (Inside) ports and 2 External (Outside) ports. A further 6 ports can also be configured as either Internal or External LAN connectivity points.

One dedicated port provides the secure management functionality and 3 maintenance ports can be configured to provide traffic mirroring and packet analysis capability for discarded (mitigated) packets. RADIUS or local user authentication methods allows secured and controlled access to the Mitigator Management Interface.

The Attack Mitigator provides a fast start-up Web Interface to enable basic prevention measures immediately but also provides menu configuration parameters that can fine tune specific timeouts, connection types, bandwidth allocation and connection rate limits.

Planning Considerations for deployment

Understand the resources that must be protected

Critical to the success of deployment is to identify all Operating Systems on Servers, Workstations and Desktop devices and Network Services that are running. Complete the Excel worksheets (supplied with the in-box documentation set) to record and document your system prior to installation.

Understand the application protocols in use

Identify all inbound and outbound application requirements (either by Network Sniffer analysis or Network Management trend analysis) and consider whether applications need to be bandwidth limited, complete worksheets to record and document your system.

Determine User Groups

Confirm whether users are classified as Internal or External and determine normal traffic flows between each particular group, consider whether additional limits or controls may be required. Complete worksheets to record and document your system.

Deployment Steps

- 1) Document your Network (using the worksheets and any other pertinent information that is available). Ensure that you have determined methods of additional logging that may be required e.g. Syslog, SNMP via the dedicated management port.
- 2) Power-up the device and connect the dedicated management port to your Network Management LAN.
- 3) Connect to the Web Management Interface and run the "Getting Started Wizard". This guides you easily through setting up port roles, subnet allocation and edge device allocation.
- 4) Complete the Network cabling to the Internal and External ports as required for your network.

5) Verify transparent operation of the AM unit (e.g. the Internal LAN can see External LAN, all Services are accessible and that Network monitoring statistics indicate correct functionality).

6) Run the unit in Monitor mode for a number of days (4 -5) to obtain normal event and traffic flow data, this mode will not attempt to block traffic – it purely serves to create a “baseline”.

7) Following the calibration period enable Basic Attack prevention by moving from Monitor mode to Mitigate mode for: Fraggle, Fragment Restrictions, FTP Bounce, HTTP URI, ICMP Restrictions, IP Address Filters, IP Source Restrictions, Land, Smurf and UDP Bomb.

8) Review the Event Logs and re-configure as necessary to remove any “false positive” Alerts.

9) Once Alerts are understood and stable then complete the Advanced settings options – Application Group Limiting (this allows limits to be set on flows from Outside to Inside for a particular application group).

Host ranges (this allows limits to be set on flow connection numbers “to and from” a host range).

SYN Flood parameters (this allows configuration of thresholds for incoming connections – i.e. Trusted, Suspicious or Malicious).

10) Enable the Advanced Attack Mitigator mode for Connection Limiting, IP Application Rate Limiting and SYN Flood Mitigation.

11) Review the Event Logs and re-configure as necessary, continue with on-going monitoring.

Conclusions

The Intrusion Prevention device is an additional hardware defence that can in some cases be simply and seamlessly integrated with existing IDS and Firewall architectures but it is in no way a single “means to an end” against all types of DoS/DDoS attack. It does however allow the IDS and Firewall systems to focus their full processing capabilities on more esoteric type attacks. Hopefully these dedicated high performance platforms will become as commonplace as Firewalls and Routers to provide much needed counter-DoS techniques and will be of major benefit overall within the security perimeter. In essence this high-speed filtering allows high volume “irrelevant or malicious” traffic to be automatically discarded before it even hits your Internet facing Firewall and IDS’s.

It is not a matter of “if” the next DDoS attack will occur its “when” and we must ensure that we are in the best possible position to react effectively and robustly with any tools available to stop this ever present menace – IPS’s are now becoming viable systems within their own right.

References

- [1] Kevin J. Houle, CERT/CC, George M. Weaver, CERT/CC “ Trends in Denial of Service Attack Technology” Oct 2001
URL: http://www.cert.org/archive/pdf/DoS_trends.pdf (Aug 30, 2002)
- [2] Stefan Savage, “ Inferring Internet Denial -of-Service Activity ”
URL: <http://www.cs.ucsd.edu/~savage/papers/UsenixSec01.pdf> (Sep 7, 2002)
- [3] Asta Networks Inc, “ Flood Attacks ”
URL: http://astanetworks.com/resources/types/flood_attacks.html (Aug 31, 2002)
- [4] “Denial Of Service Attack Swords ” May, 2000
URL: <http://www.searchlores.org/dod1.htm> Sep 3, 2002)
- [5] Internet Security Systems, “ Distributed Denial of Service Attack Tools ”
URL: <http://documents.iss.net/whitepapers/ddos.pdf> (Sep 19, 2002)
- [6] SAINT Corporation, “ Packet Flooding Problems ”
URL: http://www.wwdsi.com/demo/saint_tutorials/packet_flooding_problems.html (Sep 8, 2002)
- [7] A Householder, A Manion, L. Pesante, George M. Weaver, CERT/CC “ Managing the Threat of Denial -of-Service Attacks ” Oct 2001
URL: http://www.cert.org/archive/pdf/Managing_DoS.pdf (Aug 30, 2002)
- [8] Eric Cole, M. Newfield and John M Millican, GSEC Security Essentials Toolkit : SANS Press, Mar 2002
- [9] ISS X-Force, “Internet Risk Impact Summary” Sep 27, 2002
URL: <https://qtc.iss.net/documents/summaryreport.pdf>
- [10] BindVIEW Razor Team, “The Naptha DoS vulnerabilities paper” Nov 30, 2000
URL: http://razor.bindview.com/publish/advisories/adv_NAPTHA.html
- [11] Captus Networks, “CaptIO data sheet”, 2002
URL: <http://www.captusnetworks.com/images/pdf/captio.pdf>
- [12] Lancope, “StealthWatch data sheet”, 2002
URL: <http://www.lancope.com/XFRM.asp?RTN=Data/M100 &XML=products.xml&XSL=products.xsl>
- [13] Mazu Networks, “Enforcer data sheet”, 2002
URL: <http://www.mazunetworks.com/products/enforcer.html>
- [14] Radware, “FireProof data sheet”, 2002
URL: <http://www.radware.com/library/pdfs/products/FireProofSAS.pdf>
- [15] Top Layer Networks, “Attack Mitigator data sheet”, 2002
URL: http://www.toplayer.com/pdf/TLN_Attack_Mit.pdf
- [16] Web Screen Technology, “WS100 data sheet”, 2002
URL: http://www.webscreen-technology.com/the_solution/solution_ws100.html



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS London 2009	London, United Kingdom	Nov 28, 2009 - Dec 06, 2009	Live Event
SANS WhatWorks in Incident Detection Summit 2009	Washington, DC	Dec 09, 2009 - Dec 10, 2009	Live Event
SANS CDI East 2009	Washington, DC	Dec 11, 2009 - Dec 18, 2009	Live Event
SANS WhatWorks in Data Leakage Prevention and Encryption Summit 2010	New Orleans, LA	Jan 07, 2010 - Jan 12, 2010	Live Event
SANS Security East 2010	New Orleans, LA	Jan 10, 2010 - Jan 18, 2010	Live Event
SANS AppSec 2010 and WhatWorks in AppSec Summit	San Francisco, CA	Jan 29, 2010 - Feb 05, 2010	Live Event
SANS Phoenix 2010	Phoenix, AZ	Feb 14, 2010 - Feb 20, 2010	Live Event
SANS Tokyo 2010 Spring	Tokyo, Japan	Feb 15, 2010 - Feb 20, 2010	Live Event
SANS Geneva CISSP at HEG 2009 Autumn	OnlineSwitzerland	Nov 23, 2009 - Nov 28, 2009	Live Event
SANS OnDemand	Books & MP3s Only	Anytime	Self Paced