



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>

The Consumer Desktop – The Weak Link in Internet Security and Why ISP's Are Uniquely Positioned to Help

**John E.H. Clark
GSEC Security Essentials
Practical Assignment Version 1.4b Option 1
February 2003**

Abstract

The Internet community today is seeing a rapidly growing number of distributed denial-of-service (DDoS) attacks. At the same time the sophistication of these attacks is maturing, making defense more and more difficult. Common to all DDoS attacks is the requirement for “drones” or “zombies”, desktops or hosts that have been compromised in a way that lets an attacker utilize these systems as proxies to generate attack traffic while maintaining the anonymity of the attacker. The growing community of consumer desktops with “always-on” Internet connections provides attackers with a large source of potential drones. Securing the consumer desktop and choking off this source of drones is one of several ways to reduce the occurrence of DDoS attacks. This paper demonstrates why consumer desktops are particularly vulnerable to compromise, what options are available today to protect the consumer desktop and why Internet Service Providers (ISP's) are particularly well positioned to improve the security of consumer desktops.

© SANS Institute
full rights.

Table of Contents

1	Introduction.....	4
2	The Threat – “Denial-of-Service” (DoS) Attacks	6
2.1	“Vulnerability Exploitation DoS”.....	6
2.2	Flooding-Based Distributed Denial of Service (DDoS).....	7
2.2.1	Direct DDoS	8
2.2.2	Reflector DDoS.....	10
2.3	DDoS “Drones” and The Undefended Desktop.....	12
3	The “Consumer Internet User” - The Changing Face of the Internet User Community	12
4	Defending The Consumer Desktop Today	13
4.1	Resisting Compromise Using Consumer Anti-Virus Software	14
4.1.1	Host-Based Anti-Virus.....	14
4.1.2	Online Anti-Virus Scanning.....	15
4.2	Resisting Compromise Using Desktop Personal Firewalls	16
4.3	Resisting Compromise Using Intrusion Detection and Intrusion Prevention	16
4.3.1	Intrusion Detection.....	16
4.3.2	Intrusion Prevention	17
4.4	Resisting Compromise Using Home Network Firewall Appliances	18
4.5	Resisting Compromise Using Automated Software Updates	20
4.6	Resisting Compromise Using Desktop Vulnerability Checking	21
4.6.1	Host-based Vulnerability Scanning.....	21
4.6.2	Remote Vulnerability Scanning.....	22
5	The Alternative	24
5.1	A Case for ISP Action – ISP-Based Countermeasures.....	25
5.1.1	An ISP Community Response - Deploying Countermeasures at the ISP “Edge”	25
5.1.2	Viewing ISP Customers as an “Enterprise” Environment – a Potential New Revenue Source.....	27
5.1.3	Enterprise Solutions	29
5.1.4	An Example of ISP-Based Integrated Security Services.....	29
6	Summary.....	30
7	References	31

Table of Figures

Figure 1 - Internet Incident Report Counts 1988-2002 ¹	4
Figure 2 - Alternatives for Countering Flooding-Based DDoS Attacks.....	5
Figure 3 - The "Multiplier Effect" of DDoS.....	8
Figure 4 - Direct Distributed Denial-of Service Attacks.....	9
Figure 5 - Reflector Distributed Denial-of-Service Attacks	11
Figure 6 - Technophobia Among Internet Users ⁵⁻³	13
Figure 7 - Security Countermeasure Activity Domains	24

© SANS Institute 2003, Author retains full rights.

1 Introduction

The public Internet today is experiencing an increase in malicious activity. Carnegie Mellon University's CERT Coordination Center® is one of several agencies tracking this trend. Their report entitled "Internet Incident Report Counts 1988-2002"¹ shows an alarming increase in the number incidents reported. Figure 1 shows this trend.

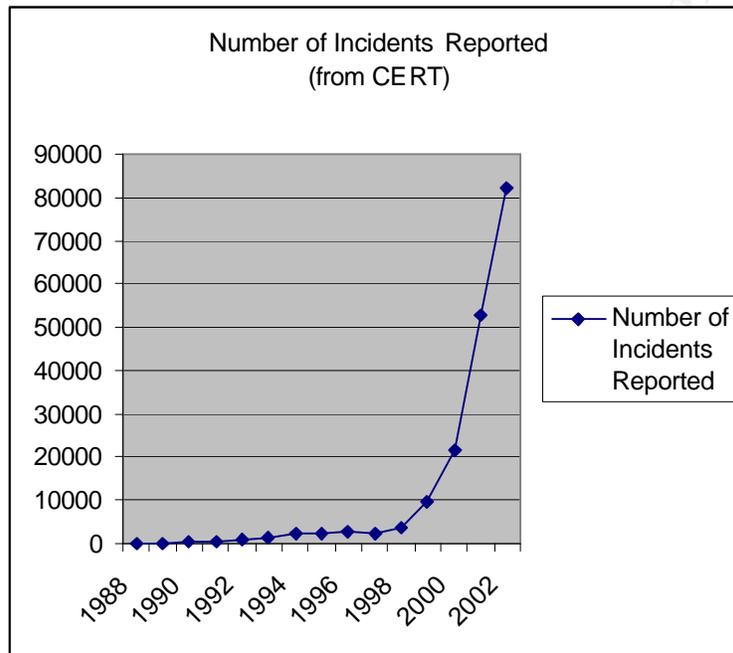


Figure 1 - Internet Incident Report Counts 1988-2002¹

In among this mix of assorted incident types are the class of attacks known as "denial-of-service" (DoS) attacks. Seeking ways to defend against DoS attacks has become a critical focus for many Internet security professionals. While "viral activity" may result in millions of dollars in lost time and data to individuals and enterprises, DoS attacks have the potential to undermine the very operation of the global public Internet. National governments, academic and financial institutions, enterprises of all sizes, consumers, all have a growing dependence on the public Internet. Numerous DoS attacks against individuals, institutions and government agencies have been recorded and many believe what has been experienced to date may just be the "tip of the iceberg". The potential for economic, social and even political disruption is enormous.

Disturbingly, many in the security community seem to have decided that DoS attacks will become a "fact of life" on the Internet and we must learn how to deal with that. This need not be the case. While today considerable effort is focused

on “detect and filter” rather than on prevention, a multi-pronged strategy to eliminate the source of DoS attacks while at the same time deploying tactical solutions to defend against, and mitigate the impact of, DoS attacks is feasible. Figure 2 depicts the range of countermeasures that can be deployed to reduce the occurrence of DoS attacks.

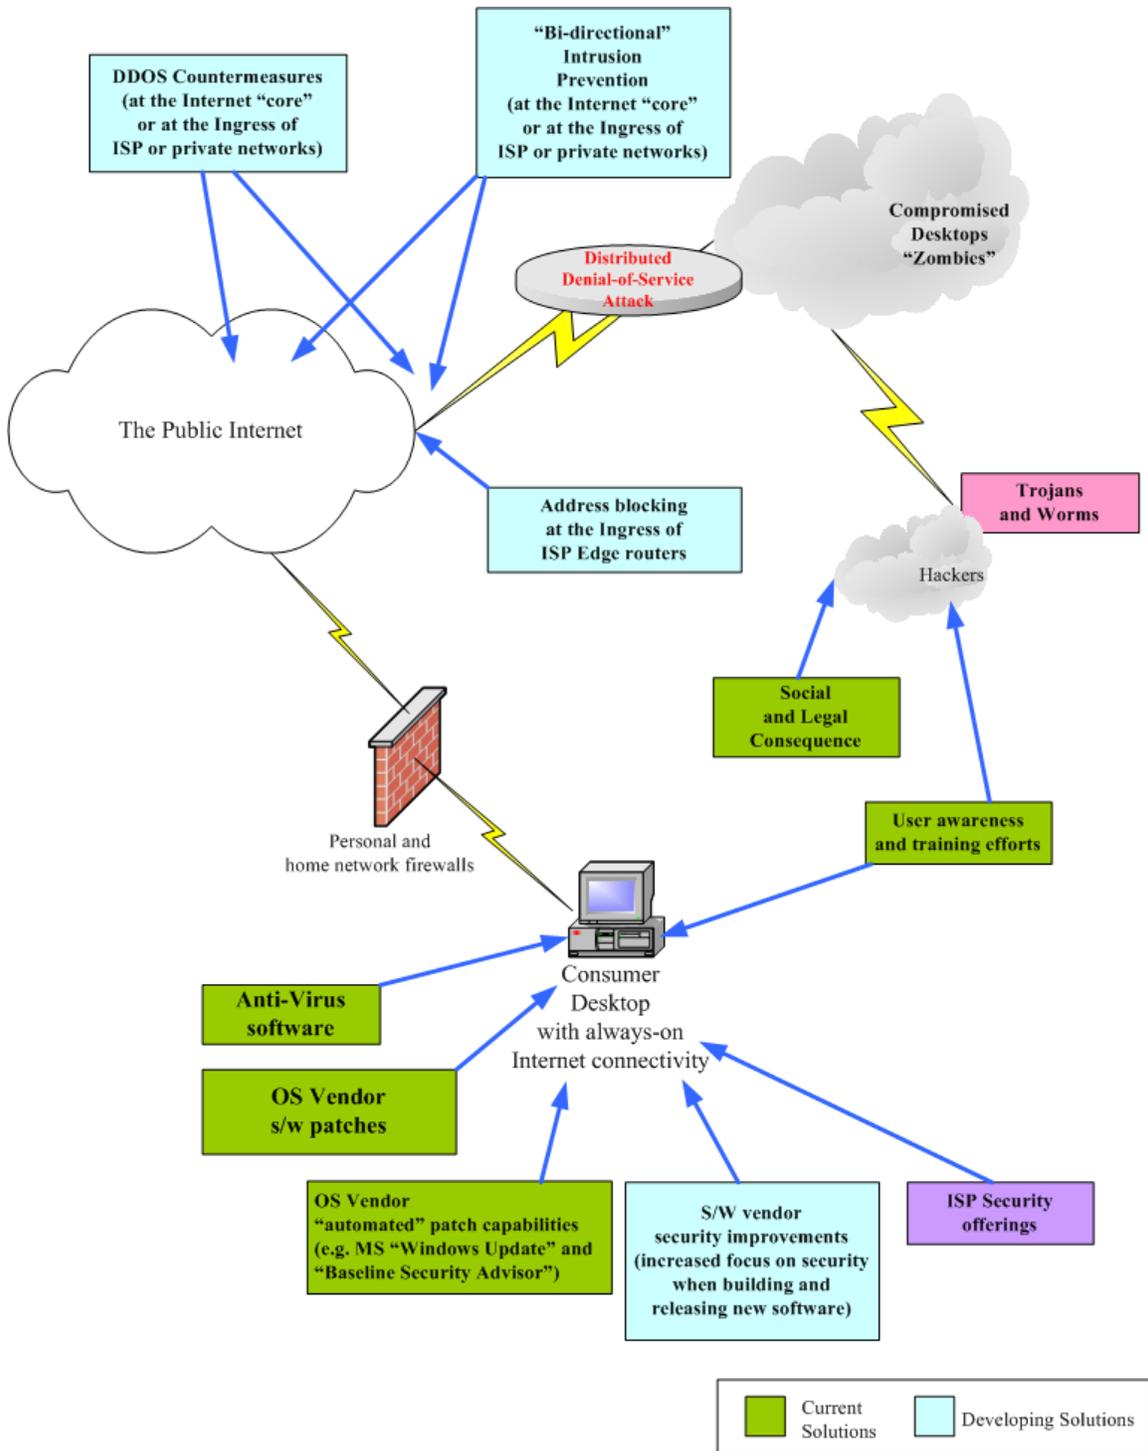


Figure 2 - Alternatives for Countering Flooding-Based DDoS Attacks

2 The Threat – “Denial-of-Service” (DoS) Attacks

A denial-of-service attack may be defined as any activity that seeks to interrupt a service either by causing the failure of the service (rob the bank), or by preventing access to the service (partially or completely close or lock the door to the bank without the bank’s consent). In the Internet world, DoS attacks have two pre-requisites: assurance of anonymity (who’d rob a bank without wearing a mask?) and an ample supply of drones (“thugs” to do the “dirty work”). Anonymity is achieved by using proxies (referred to as drones or the more colorful “Zombies”) to execute an attack. Drones are hosts or desktops that have been compromised in some way and are now pawns to a remote attacker. More often multiple levels of proxy are utilized, i.e. handlers managing drones, and communication between the attacker and the handlers/drones is further de-linked using anonymous chat rooms or billboards. Steven Gibson’s graphic tale of a DoS attack on his site (“The Strange Tale of the Denial of Service Attacks Against GRC.com”²) is an eye opening portrayal of how complex the attacker-handler-drone communication network can become.

Reiterating, there are two ways to interrupt a service: cause it to fail, or prevent access to it. You can cause it to fail either by leveraging a design vulnerability and causing it to fail, or you can deplete its resources to the point where it cannot service valid requests.

2.1 “Vulnerability Exploitation DoS”

Early DoS attacks were successful in attacking vulnerabilities in key software, that when invoked, simply caused the host to crash. Today these are relatively easy to defend against. But attacks have evolved and software vulnerabilities are now leveraged in a more devious manner, typically to “deposit” a piece of (difficult to find) executable code. When executed, this code may in turn copy more sophisticated code onto the compromised host or complete other unwanted activities (such as seek out other nearby hosts to compromise). In a particular form of compromise, code is loaded onto a compromised desktop in the form of normal application or utility software corrupted to perform unintended functions. These “Trojans” can form the basis for the drone network needed for a “distributed denial-of-service” attack (DDoS). Well known DDoS Trojans include Trinoo, TFN, Stacheldraht and Shaft (with, no doubt, many more to come). The CERT “Incident Note IN-99-07”³ provides descriptions of Trinoo and TFN.

A point to note is that, for this mechanism of “external trojan-loading” to succeed, there is a fundamental need for network connectivity and for a “vulnerability” to exist before a host can be compromised. If there is no network access to the host and there are no vulnerabilities to attack, a host cannot be compromised. As has

been noted by many, the best firewall is "about half an inch of air". Today, network connectivity of some form is a virtual "must-have" and, unfortunately, software vulnerabilities abound.

2.2 Flooding-Based Distributed Denial of Service (DDOS)

Causing a denial-of-service using a single drone will not likely be successful unless a fatal vulnerability can be leveraged, something that is less common now. DoS aficionados will more often turn to DoS attack strategies that attempt to overwhelm a server in some way, either by depleting its resources or clogging its network access. To execute these strategies requires a large volume of traffic. Unhappily, with a little work (and some help from the software industry) it's not difficult to build a network of drones that do just that, generate a large volume of traffic. By distributing the source of an attack, the attacker gains the potential for very large traffic volume (hundreds of compromised PC's churning out attack traffic via multiple network access paths instead of just a few). The "multiplier affect" (Figure 3) can be devastating on the target as it is flooded with traffic.

© SANS Institute 2003, Author retains full rights.

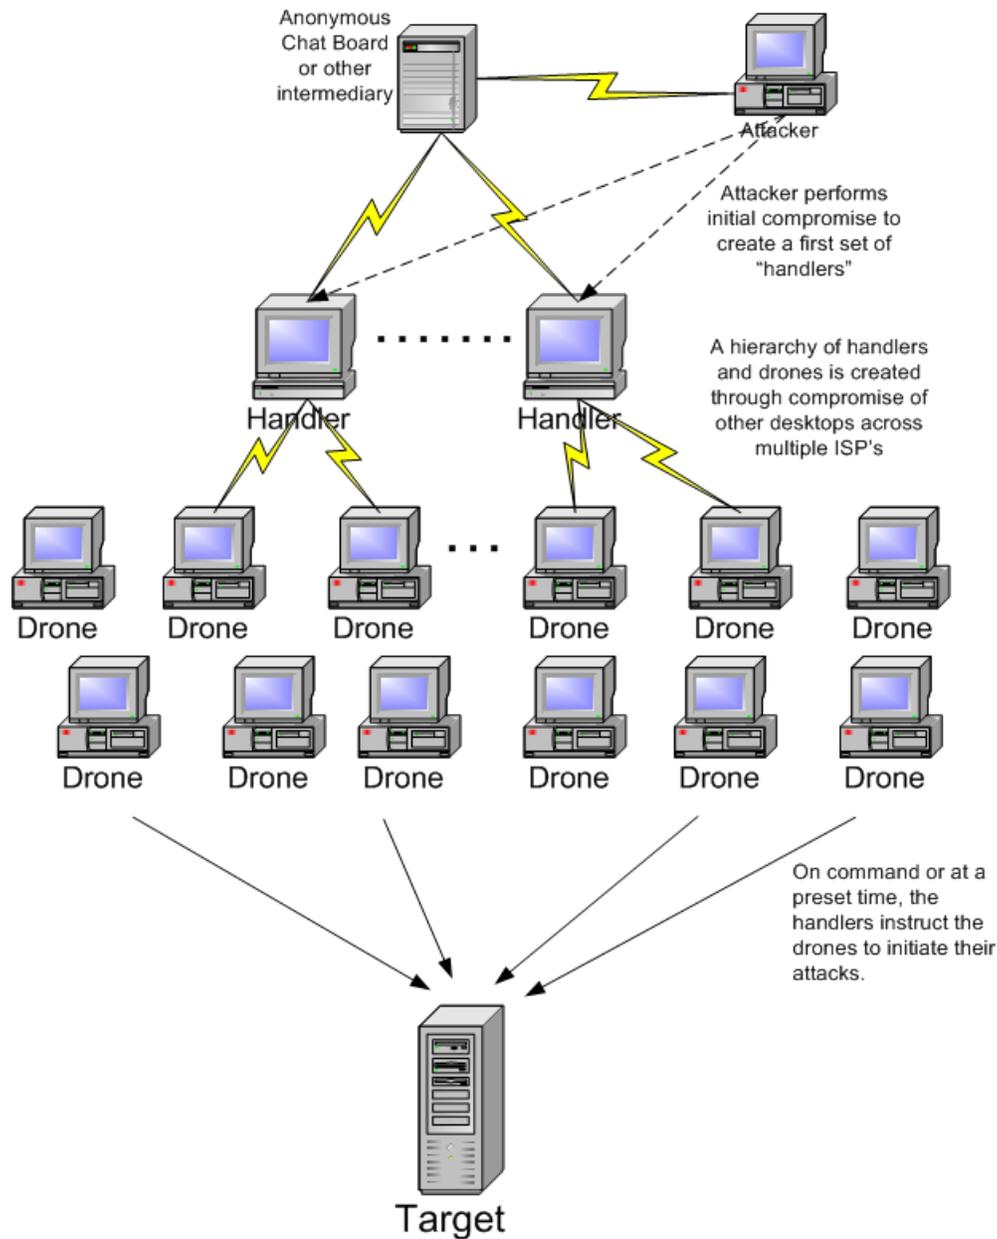


Figure 3 - The "Multiplier Effect" of DDoS

2.2.1 Direct DDoS

With direct DDoS attacks the drones flood the target with packets intending to consume target resources. A simple example is "SYN flooding" whereby the drones send a continuous stream of TCP SYN packets to the target. The target responds to each with a SYN-ACK packet and then sets aside resources for what it believes will be a successful completion of the TCP "3-way handshake" - a TCP session set up. However the completing ACK's are never sent (see Figure 4). Eventually the server exhausts its session resources holding open partial sessions and is unable to accommodate valid requests. With these attacks the

drones may or may not spoof their source address (artificially substitute a source IP address other than their own into the packets sent to the target). Since the drones are unwitting pawns of their handler/attacker, the attacker may or may not choose to hide the drone's identities.

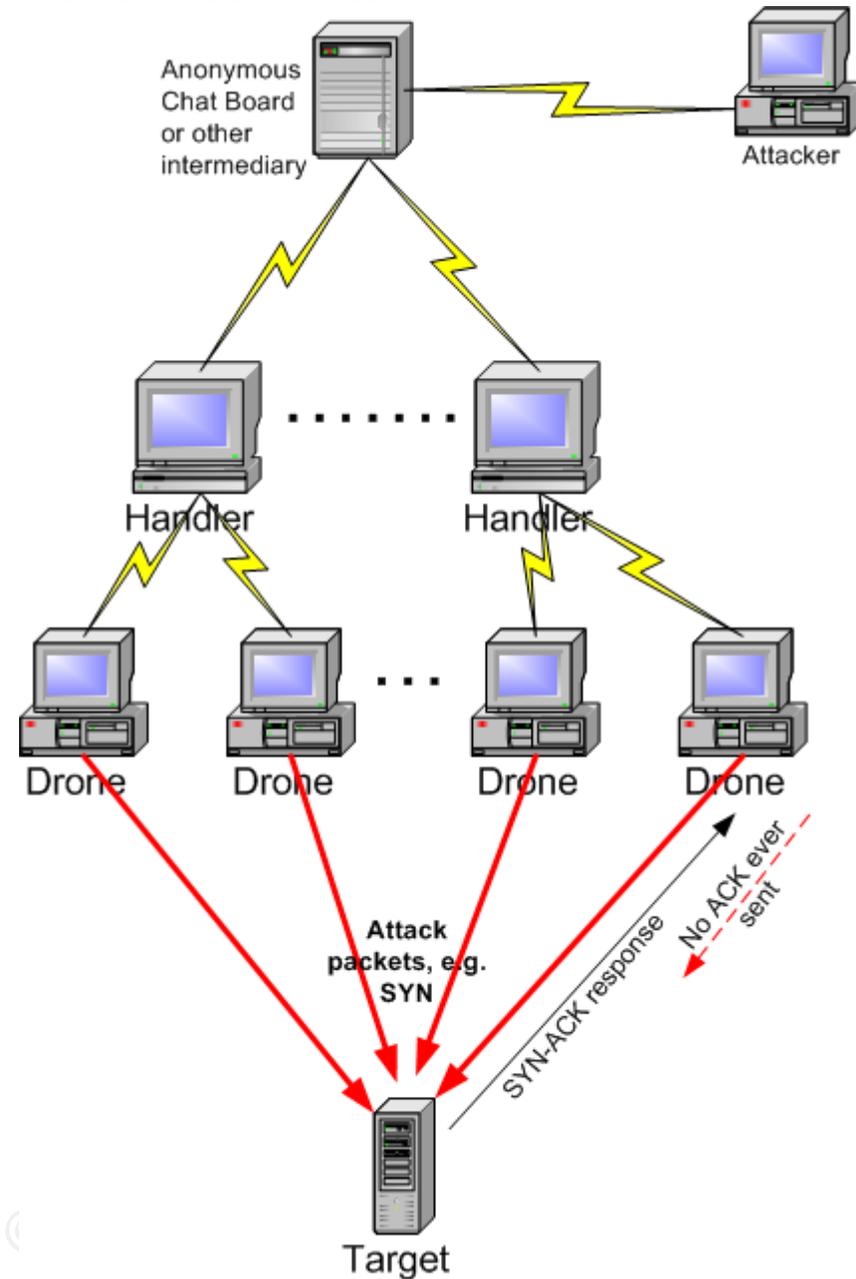


Figure 4 - Direct Distributed Denial-of Service Attacks

A point to note here is that to gain the traffic volume needed, the attack is most effective if many drones are dispersed across many networks. If it were harder to “enlist” drones, it would be more difficult to build (and retain) the attack network and keep it in place. Unfortunately today there is no shortage of potential drones.

2.2.2 Reflector DDoS

Direct DDoS can be countered by filtering packets based on the drone's source address once an attack has been identified. But what if the attacker could arrange for the attack traffic to come from legitimate servers? What if the drones could enlist an "innocent bystander" to direct traffic at the target?

With Reflector DDoS the intent of the attack changes. With this technique, the intent is to clog the network link to the target with unsolicited traffic and thereby impair legitimate access. The drones (on direction from their handler(s)) send legitimate requests to legitimate servers but with the address of the target artificially inserted into the source address field of the request packet (source address spoofing). The servers respond to the target (as they are designed to). The result is a flood of traffic directed at the target from legitimate sources (Figure 5). You can't filter it without impacting legitimate traffic from those servers to other hosts in the target network. There is any number of interesting variations on this theme depending on the type of request sent by the drones. A popular variety uses DNS requests as the DNS server reply is typically large adding to the impact on the target's network link.

© SANS Institute 2003, Author retains full rights.

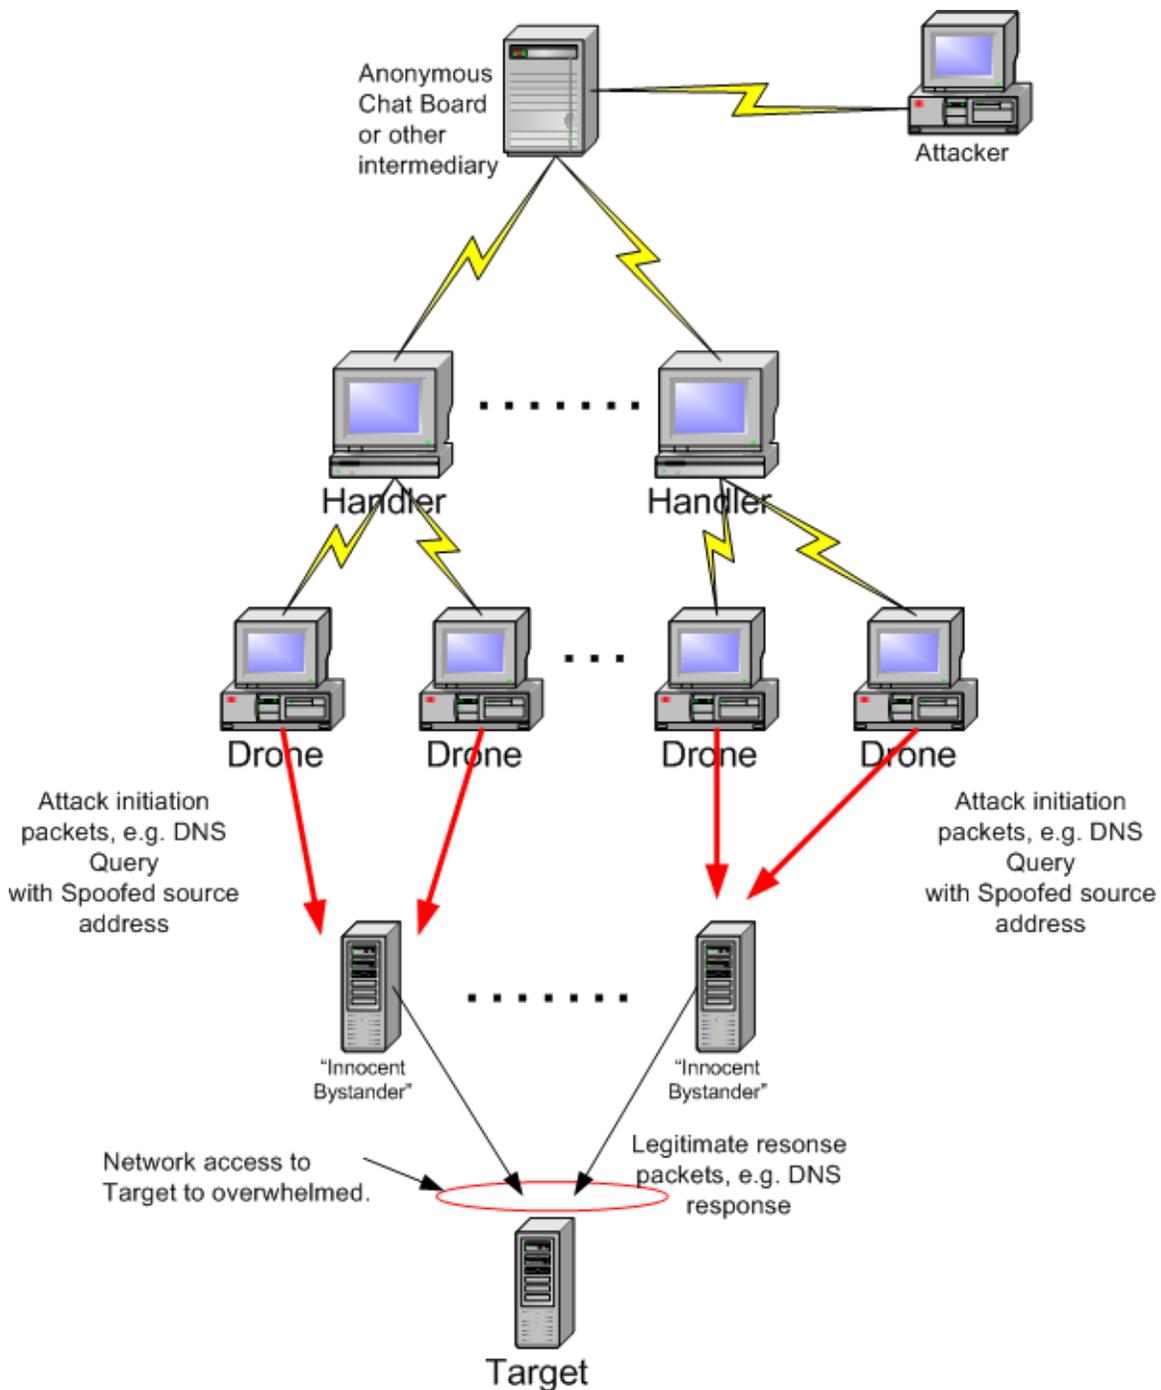


Figure 5 - Reflector Distributed Denial-of-Service Attacks

Direct and Reflector DDoS attacks are subject to intense scrutiny today as the security community works toward ways to mitigate or prevent these attacks. An excellent description can be found in "Internet Infrastructure Security: A Taxonomy"⁴ by Anirban Chakrabarti and G. Manimaran at Iowa State University.

A point to note here is that if source address ingress filtering were in place on edge routers (drop packets whose source address does not fall within the range

of addresses expected on an interface), reflector DDoS attacks could not happen unless:

1. The innocent bystander is in the same subnet as the drones (the request packet does not have to pass through any ingress address filtering which may exist), or
2. The target is in the same subnet as the drone (the spoofed source would successfully pass thru any ingress filter).

The Section 5.1, “ISP-Based Countermeasures”, briefly discusses ingress address filtering.

2.3 DDoS “Drones” and The Undefended Desktop

As mentioned, without anonymity there would not be any DoS attacks, and without drones the attacker cannot assure anonymity. Hence in any defense against DoS the incidence of compromised desktops and servers (i.e. drones) must be addressed. In particular the common consumer desktop may be the largest risk to Internet security today. Partly because there are so many of them, partly because more and more of them are moving to “always-on” Internet connections, and partly because the way we use the Internet is changing (refer to Section 3), the consumer desktop continues to be relatively vulnerable, providing a fertile place to grow DDoS drone networks.

3 The “Consumer Internet User” - The Changing Face of the Internet User Community

Since 2000 the UCLA Center for Communication Policy has been surveying a static population of about 2000 Internet users and non-users. The intent is to track the evolution of this new medium and how we use it (or not). The survey results (“Surveying the Digital Future”^{5-1/2/3}) document in 2000,2001 and again in 2002, an increase in:

- The number of hours spent connected to the Internet
- The amount of home time spent on the Internet
- Broadband access versus dialup
- The number of computers in the home
- The number of home computers that are connected to a home network
- The number of Internet non-users who plan to go online in the next year

While Internet usage is on the upswing, all levels of Internet users, new users as well as experienced users, report a degree of technophobia, “feeling discomfort about computers or any computerized technology”⁵⁻³. Consider the statistics from the “Year Three” report shown in Figure 6. Nearly a third of new Internet users expressed a degree of technophobia. More than 10% of experienced Internet users felt the same.

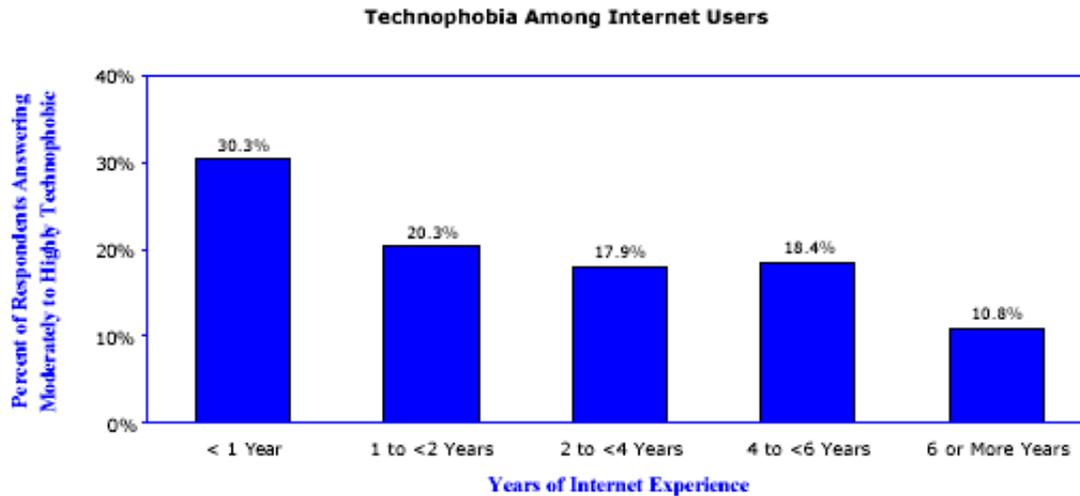


Figure 6 - Technophobia Among Internet Users⁵⁻³

While there is no lack of advice, information and instruction (see, for example, the step-by-step instructions provided by Carnegie Mellon CERT/CC¹¹), securing one's personal desktop can be a daunting task.

Expecting a new Internet user (or even an experienced Internet user) who has already expressed some discomfort with computers to take on this task would seem unreasonable.

4 Defending The Consumer Desktop Today

The tools for defending the consumer desktop are in many ways already available today. The issues are: who's responsible for doing it, who pays for it, and are there privacy or civil liberty issues that must be considered? Here are some of the solutions that can be deployed today:

- Consumer Anti-Virus Software
- Desktop Personal Firewalls
- Intrusion Detection and Intrusion Prevention
- Home Network Firewalls
- Automated Software Updates
- Desktop Vulnerability Checking

Within each of these there are a variety of products and approaches that apply.

As will be seen all of these require the consumer Internet user to "do something", often something they lack the skills, training or interest to do.

4.1 Resisting Compromise Using Consumer Anti-Virus Software

4.1.1 Host-Based Anti-Virus

These days, loading anti-virus software on your desktop is a “socially expected thing” (the Internet generation is more and more accepting of the concept of “safe computing”). Most consumers are at least peripherally aware of viruses and their threat and most would likely purchase anti-virus software. As well most new systems arrive with antivirus software bundled. But that extra step to keep the virus definitions up-to-date may be beyond the threshold of consumer acceptance. Even with the “live update” common now with antivirus software, a “subscription” often needs to be renewed, perhaps yearly. All too often the subscription lapses after the first included year, so although the consumer is still running the antivirus software, the protection afforded by the software begins to diminish daily as new vulnerabilities and associated viruses appear.

From Symantec’s most recent “Internet Security Threat Report”⁶ (February 2003):

Adding to risks associated with cyber attacks, the discovery rate for new IT product vulnerabilities accelerated substantially over the past year. The total number of new, documented vulnerabilities in 2002 was 81.5% higher than in 2001.

At this rate even with “current” virus definitions the “Day One” effect (a virus spreading prior to availability of signatures or other countermeasure) can lead to a compromise. A new virus can propagate significantly in the few hours it may take to detect the virus, develop a signature for the virus, and make the signature available to consumers.

A point to note about consumer desktop software is that, even with foolproof bulletproof installation and automated live updates where feasible, the consumer can’t be depended on to maintain that software. If they have to “do” something that does not pertain directly to what they want to do with their desktop, it may not get done. In the case of anti-virus software (and personal firewall software – see below) inadequately maintained software may be more dangerous than having no software at all (“I’m safe. My anti-virus software is running...”). The time between the release of a new virus signature and the installation of that signature on a significant population of consumer desktops might be measured in days, weeks or even months.

Host based antivirus software is available from several vendors. All provide a subscription-based virus definition update service. Example anti-virus products are listed in Table 1.

Symantec Norton Anti-Virus	http://www.symantec.com/nav/nav_9xnt/
McAfee VirusScan	http://www.mcafee.com/myapps/vs7/
TrendMicro PC-cillin	http://www.trendmicro.com/en/products/desktop/pc-cillin/evaluate/overview.htm

Table 1 - Example Host Based Anti-Virus Products

4.1.2 Online Anti-Virus Scanning

One way to get around the “out-of-date virus signatures” problem is to have the signatures maintained by someone else or let the signatures reside with a 3rd party that is requested to scan the desktop from afar. There are free and for-fee versions of this solution. The largest problem with remote scanning is that they are “static”; they scan a disk or memory for viruses but do not maintain a constant on-board vigil for virus-like activity. The “activity profile matching” aspect of anti-virus is becoming an essential component of the anti-virus war as it is the best solution to date for the “Day One” problem (countering a newly released virus for which a signature has not yet been developed or deployed). Nor can remote monitors scan incoming files as they arrive on the desktop (e.g. e-mail attachments or web page files). An alternative for incoming files is to delegate the antivirus checking to upstream resources (such as the ISP or a home network firewall). But for real-time onboard activity monitoring there is no alternative to host based anti-virus software.

Remote antivirus monitoring does not constitute a complete desktop antivirus solution. It must be coupled with on-board desktop antivirus software that handles the real-time monitoring. The ideal is to combine some form of upstream anti-virus detection with on-board detection and activity monitoring, a design in line with the “defense in depth” philosophy promoted by the security community. And the virus signatures, on-board and upstream, must be current.

4.1.2.1 “Free” Online Anti-Virus Scanning

The sources for “upstream” antivirus scanning are many. One source is the vendors of common anti-virus scanning software. All major vendors, perhaps recognizing the win-win scenario of additional revenues coupled with “faster-time-to-desktop” for security countermeasure, have released on-line security suites. As a loss leader, they often offer “free” remote antivirus scanning, typically with reduced functionality when compared with their for-fee service. Table 2 lists several free offerings.

McAfee "Freescan"	www.mcafee.com/myapps/mfs/
TrendMicro "HouseCall"	http://housecall.trendmicro.com/
Symantec (part of Symantec Security Scan)	http://security.symantec.com/ssc/
Panda "ActiveScan"	www.pandasoftware.com

Table 2- Example Free Remote Anti-Virus Scanning

4.2 Resisting Compromise Using Desktop Personal Firewalls

Just as antivirus software can monitor a desktop for what may be virus activity, a personal firewall can monitor the traffic in and out of a desktop for signs of what may be intrusion attempts. On detection the user is typically notified and asked to permit or deny the activity. Suitably configured (again, something a consumer may not necessarily be expected to do) a personal firewall may intercept malignant activity on the part of an internal or external process. Table 3 lists several personal firewall products.

Symantec Norton Personal Firewall	www.symantec.com/sabu/nis/npf/
	www.symantec.com/sabu/nis/npf_mac/index.html (Macintosh)
McAfee Personal Firewall Plus (Windows)	www.mcafee.com/myapps/mpfp/default.asp
Black Ice PC Protection	http://blackice.iss.net/product_pc_protection.php
ZoneAlarm Pro	www.zonelabs.com/store/content/catalog/products/zap/zap_details.jsp
Tiny Personal Firewall	http://www.tinysoftware.com/home/tiny2?la=EN

Table 3 - Example Personal Firewall Products

McAfee's offering is now delivered as a subscription-based service. For a fee this relieves the consumer of the responsibility for maintaining the firewall software and configuration. It is not available for Macintosh users.

4.3 Resisting Compromise Using Intrusion Detection and Intrusion Prevention

4.3.1 Intrusion Detection

Intrusion detection comes in two varieties: host based (HIDS) and network based (NIDS).

Host based intrusion detection systems monitor system activity for “suspect” activity and can generate alerts to the user. For example an attempt to modify a user password from a process not associated with the user could be flagged and alarmed. In the case of HIDS’ the configuration is OS-dependent and may be intimidating for a user not familiar with the internal workings of their desktop OS.

And, as with antivirus software, for both personal firewalls and host based intrusion detection, there is a reliance on the user to maintain the software. Other than configuration difficulties this may be the largest drawback to this solution.

Network based intrusion detection relies on an upstream entity to observe and detect potential intrusion activity, log the activity and optionally issue an alarm in some fashion. For the consumer user, NIDS may reside with a home network firewall or potentially with the ISP. NIDS’ by definition only passively monitor and record intrusion activity. A NIDS alarm means the intrusion has already been attempted and may or may not have succeeded. NIDS may feed to other tools to activate countermeasures, for example to automatically adjust the access control lists on a router to counter the observed intrusion activity.

As with host based intrusion detection, NIDS on a home network firewall may be too complex for a consumer environment. And there is the question of how to usefully notify the user of the intrusion activity. The only viable NIDS location may be with the upstream ISP.

4.3.2 Intrusion Prevention

While an IDS is a passive monitor, intrusion *prevention* actively intercepts and discards intrusion traffic. Intrusion prevention also comes in host based (HIPS) and network based (NIPS) varieties. Network based intrusion prevention is discussed further in “Bi-Directional Network Based Intrusion Prevention” in Section 5.1.1 of this document.

HIPS is similar to the activity profile monitoring of some anti-virus products but actively responds to any observed intrusion activity. Typically the HIPS “sits” between the kernel of the operating system and application or utility software issuing requests to the kernel. Some activity can be determined to be an intrusion to a very high degree of certainty and in such cases the HIPS software blocks the request, denying the access to the kernel.

HIPS products today are targeted at large server installations. Their complexity, OS dependence, and relative sensitivity to OS-application configuration suggest a widespread deployment of HIPS technologies on consumer desktops is unlikely in the near future. More likely HIPS technologies may appear as part of desktop OS functionality. Linda Dailey Paulson provides a good explanation of intrusion prevention in “Stopping Intruders Outside the Gates”⁷.

4.4 Resisting Compromise Using Home Network Firewall Appliances

In addition to the ability to share an Internet service, the shared firewall appliance offers several security benefits to the home network including:

- Traffic filtering based on IP port or IP address
- Perimeter anti-virus scanning
- Address hiding
- Rudimentary intrusion prevention
- Content filtering

The home network firewall can be the home's first line of defense against intrusion. Properly configured it can prevent unwanted traffic from entering (*or leaving*) the home network. Filters can be defined based on IP traffic type, IP port number or IP source or destination address. *This presumes that the consumer knows what an "IP port number" is (for example)*. Out of the box, most consumer firewalls provide good but not necessarily adequate protection. The user "management interface" is usually web-based with varying degrees of complexity and functionality. This of course presumes that the consumer *wants* to manage the firewall.

Several home network firewall vendors offer an option to perform perimeter antivirus scanning of inbound e-mail and web traffic on behalf of the entire home network. Its tempting to view this as an alternative to host based antivirus. For a family with four computers on a home network, one might question the need to purchase antivirus software and subscriptions for each home desktop if the perimeter firewall is scanning all inbound files. However perimeter antivirus is *not* a complete antivirus solution. For example the shared firewall can't:

- Monitor for on-board virus activity,
- Scan the floppy or CD that "Junior" borrowed from a friend,
- Scan traffic carried through an IPSEC tunnel (a logical connection between two sites carrying encrypted traffic that the firewall can't analyze),
- Scan encrypted e-mail attachments

Instead the firewall antivirus scanning should be viewed as a first line of defense with desktop antivirus as an additional line of defense.

(As an example Sonicwall offers an integrated firewall/antivirus solution that also *enforces* and maintains desktop antivirus implementations on the desktops protected by the firewall.

See <http://www.sonicwall.com/applications/antivirus.html> for details.)

The home network firewall can also provide IP address hiding by implementing network address translation (NAT) and/or port multiplexing (PAT). Most consumer Internet subscribers are provided with one or two *Internet-routable* IP addresses (the ISP "owns" these addresses and rents them to the subscriber).

Without NAT/PAT, this address is assigned to the single desktop attached to the ISP service. This address is visible (and traceable) on the Internet. In all cases the ISP can determine which subscriber is using which address at any particular time (even though these IP addresses may be dynamically allocated on a demand basis, there are always logs recording which IP address was allocated to which physical address over what period). A firewall implementing NAT/PAT acquires a single ISP address on behalf of all of the desktops on its internal (protected) network, and uses that address as the source address for all traffic emanating from the internal network. Internal network addresses never reach the Internet (typically internal networks can and should use “private” IP addressing that Internet routers should never route). The NAT/PAT firewall can perform one-to-one address translation or it can multiplex multiple internal addresses onto one external address by dynamically assigning IP ports to sessions initiated or accepted by inside devices. The port assignment is used to map “outside sessions” to “inside addressing”.

Most home network firewalls also provide rudimentary intrusion prevention and will drop packets associated with well known attacks such as “TCP FIN” scans and “sub seven” attacks.

A third capability of the home network firewall is shared content filtering, useful in the never-ending battle to keep the less desirable aspects of the Internet away from children. Firewall vendors may offer subscription-based content filtering.

For the most part, home network firewalls come “out of the box” well configured for basic operation. Adjusting the configurations to more stringent requirements may not be “intuitive” for some consumers. The care and feeding of the firewall appliance, like other security countermeasures, may not be a priority of many consumers. “As long as I can get to the Internet it must be working....”.

Acknowledging the additional complexity that a hardware firewall adds to a home network, it is still a significant protection in that it at least protects the network from the growing “background noise” of nuisance Internet attacks that is experienced today. And if maintained adequately can be a significant part of a home network security plan.

However if the consumer does not maintain the home network firewall its value as a countermeasure diminishes, as is the case with anti-virus and personal firewall solutions. The option then may be to push the home network firewall function upstream to the ISP.

Table 4 lists some available home network firewalls.

Sonicwall	www.sonicewall.com
D-Link	www.dlink.com/products/vpn_firewalls/index.asp
Linksys	www.linksys.com/products/group.asp?grid=34&scid=29
NetGear	www.netgear.com/products/routers/websafefirewall.asp?view=hm

Table 4 - Examples of Consumer Hardware Firewalls

As with an enterprise or small business, a home network can benefit from the use of a shared firewall. However there is an illusion that the firewall is all that's required to be secure. This is not the case. As with enterprise and small business environments, the concepts of "defense in depth" apply.

4.5 Resisting Compromise Using Automated Software Updates

As commented earlier, a key ingredient in the campaign to secure the desktop is a response to the proliferation of software vulnerabilities. For, if there were no vulnerabilities, the job of defending a host or desktop would be much easier. Today when installing new software on a desktop, one inevitably reaches the "software license disclaimer" step ("Click I Agree to Proceed"). Most disclaimers release the software vendor from any liability whatsoever from the consequences of their software failing or having a vulnerability that can be used to compromise the desktop. Few consumers would think to click "No, I don't agree". This is a difficult predicament. One can sympathize with software vendors who are under severe pressure to deliver new functionality quickly and competitively. It would seem inevitable that some software would get out the door prematurely. On the other hand, the rate at which software (and firmware) vulnerabilities are detected continues to climb (refer to "Symantec's Internet Security Threat Report Vol. 3"⁶).

Perhaps acknowledging that some vulnerabilities will inevitably get past the software test phase, some major software providers are trying to automate the update/patch cycle. Notably Microsoft ("Windows Update") and Apple ("Software Update") both deliver an automated or user-initiated "ET-call-home" type of software checkup. The idea is to make it as easy as possible to keep the core operating software and key applications current with the latest patches. While there is still an element of user involvement, this is an early but promising step toward reducing the vulnerability of the desktop. Table 5 provides links to the Microsoft and Apple descriptions of these services.

Microsoft Windows Update	http://v4.windowsupdate.microsoft.com/en/default.asp
Apple Computer	http://docs.info.apple.com/article.html?artnum=60504

Table 5 - Examples of Automated Software Updates

4.6 Resisting Compromise Using Desktop Vulnerability Checking

Even with up-to-date software patches, a desktop can still be insecure if it is improperly configured. Software, especially operating systems software it seems, often comes out of the box inadequately configured from a security perspective (ease-of-use still wins out over security and perhaps always will). Assessing a desktop from within to see how it compares with defined security models or assessing from outside to see how it responds to known attacks, can suggest essential configuration changes.

4.6.1 Host-based Vulnerability Scanning

Two examples of local desktop vulnerability assessment are Microsoft's "Baseline Security Analyzer" (MBSA) and the Center for Internet Security (CIS) Security Scoring Tool (for various platforms). Table 6 provides links to descriptions of these tools.

MBSA	www.microsoft.com/technet/security/tools/Tools/mbsahome.asp
CIS Security Benchmarks and Scoring Tools	www.cisecurity.org/bench.html

Table 6 - Examples of Desktop Security Assessment

MBSA is a somewhat user-friendly attempt to address the challenge of vulnerability assessment for newer Windows versions. MBSA compares the desktop on which it is running (or other accessible desktops) with defined Microsoft security profiles and with a Microsoft patch database and provides the user with recommended updates. *The user must then take the next step to apply the recommended changes.* While Microsoft makes the patch process fairly easy, this process is one a consumer may not be comfortable with. Microsoft targets this tool at "IT professionals".

The Security Scoring Tools from the Center for Internet Security (CIS) also compares the desktop with predefined profiles and is also targeted for use by security and IT systems professionals. Many consumer users would not understand the significance of the security profiles used by this tool and would not likely be comfortable using this tool.

Again, the question is, can a typical “new age” Internet user be expected to do this?

4.6.2 Remote Vulnerability Scanning

Enter the security service providers. Enterprises have long struggled with the same problems identified for consumer desktops. How can you effectively manage desktop security if your desktop community numbers in the 100’s or 1000’s of desktops, are geographically distributed, and may not be focused on maintaining their desktop? Several vendors have stepped forward with centralized security scanning and management tools targeted at the enterprise customer. Leveraging these same technologies vendors are also now offering remote security assessment for home users. Table 7 lists two such services.

Symantec Security Check	http://security.symantec.com/ssc/
McAfee Security Center	http://www.mcafee.com/myapps/msc/default.asp www.mcafee.com/myapps/vso/default.asp

Table 7 - Example Security and Online Anti-Virus Services

4.6.2.1 Symantec “Security Check”

Symantec “Security Check” is a no-charge offering from Symantec Corporation. The user goes to the Symantec web site and requests a scan of their desktop. The service operates in two modes, “server side” and “client side”. Server side is driven entirely from the Symantec server but is limited to network-access scanning (e.g. is there a Trojan hiding on your desktop, or do you have exposed Netbios shares, etc). The client side service is more comprehensive but requires that code be loaded onto the desktop and that ActiveX and scripting be enabled on the desktop browser. “Security Check” can check for security risks, confirm the status of the anti-virus tool on the desktop, and optionally offers a trace back facility when trying to identify the source of a potential attack. (If your personal firewall notifies you that you have been attacked, the first question that comes to mind is “Who is this person who is attacking me?” The Symantec “Trace a Potential Attack” tool is an interesting place to start when trying to answer that question although, often as not, for serious attacks the source is not that easy to find.)

4.6.2.2 McAfee “SecurityCenter”

McAfee Security Centre is a free download of assorted security assessment tools and evaluation copies of McAfee security products which link to for-fee remote security services. Table 8 lists several of the security services accessible from “SecurityCenter”.

McAfee Privacy Service	http://www.mcafee.com/myapps/mps/default.asp
McAfee VirusScan Online	http://www.mcafee.com/myapps/vso/default.asp
McAfee Personal Firewall	http://www.mcafee.com/myapps/firewall/default.asp

Table 8 - McAfee Security Center Linked For-Fee Services

© SANS Institute 2003, Author retains full rights.

5 The Alternative

If it is presumed that a significant population of Internet users will not adequately secure their desktops, what are the alternatives? Figure 7 is a simple depiction of how consumers connect to the Internet and where security countermeasures could be invoked (“security activity domains”). We have discussed at length what countermeasures are available within the “consumer activity domain” and why we might not expect those countermeasures to be in place. It is therefore reasonable to consider alternatives. In “Defending against Flooding-Based Distributed Denial-of-Service Attacks: A Tutorial”⁸ the author describes, among other things, the possible locations for performing DDoS attack detection and filtering. The user’s Internet Service Provider is a key player.

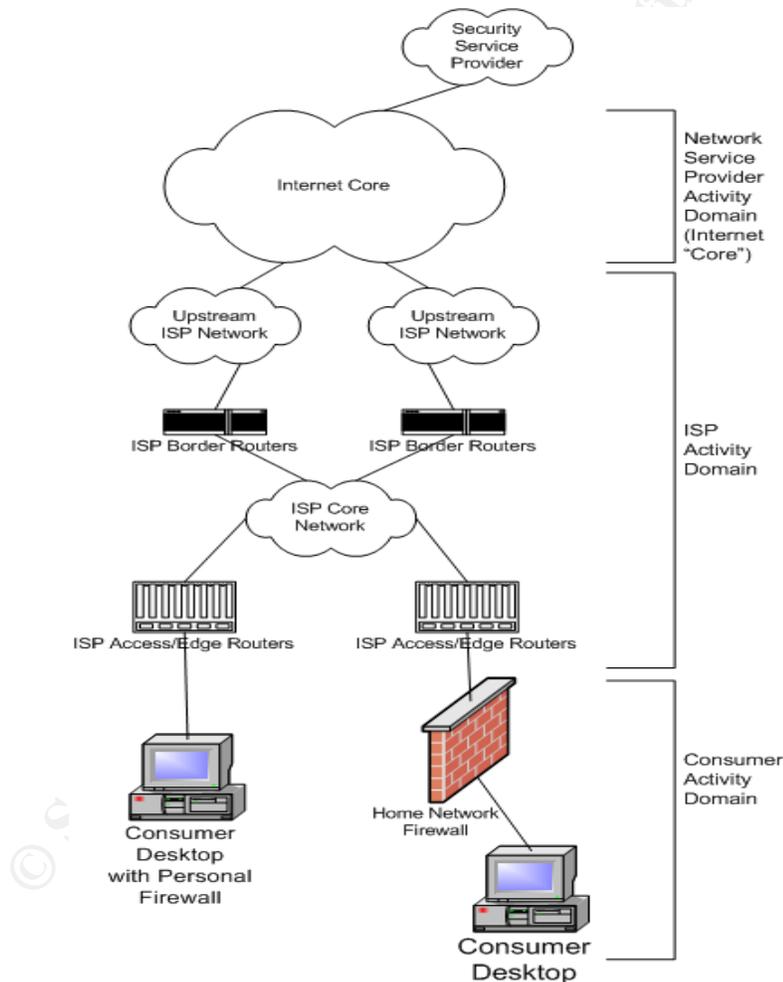


Figure 7 - Security Countermeasure Activity Domains

5.1 A Case for ISP Action – ISP-Based Countermeasures

The ISP is the user's conduit to the Internet. All traffic between the public Internet and the consumer's desktop will pass through the ISP's access and core routers. This uniquely positions ISP's to monitor traffic to and from a consumer desktop. Similarly, ISP's are uniquely positioned to identify and physically locate a consumer desktop. ISP's rent IP addresses to subscribers and depending on the type of ISP service (dial, DSL, cable modem), the address may "stay" with a subscriber for weeks or months. Regardless of the duration of the assignment, at any time the ISP maintains a unique association between the loaned IP address and the user's ISP account. This account/IP address relationship coupled with the physical cabling association between the subscriber and the ISP (e.g. modem port or DSL port) uniquely positions the ISP to offer and implement security services to the consumer and on behalf of the consumer.

Several ISP-based security countermeasures are apparent. Some might be described as just beneficial to Internet community in general. Others would appear to be win-win business opportunities for the ISP, representing a potential revenue source while at the same time enhancing the ISP's own security by protecting it from DDoS attacks against the ISP itself.

5.1.1 An ISP Community Response - Deploying Countermeasures at the ISP "Edge"

As has been stressed throughout this paper, there is significant challenge to securing the consumer desktop. While such initiatives must proceed, there is value in pursuing other "angles" to deliver the multi-pronged Internet-community-driven initiative needed to significantly address the DoS threat. Two Internet service provider-based countermeasures are:

- Bi-Directional Intrusion Prevention at the ISP access
- Ingress Address Filtering at the ISP access

5.1.1.1 "Bi-Directional" Network-Based Intrusion Prevention

Network based intrusion prevention (NIPS) technologies work on the premise that some Internet attacks can be recognized with better than 99.9% confidence and in such situations the attack packets can simply be dropped. Further, we assume that our ability to recognize attacks will improve over time, making intrusion prevention more and more successful at eliminating attack packets. At a minimum, intrusion prevention can reduce the volume of "nuisance" attack traffic that is evident today. For example, it might capture 50-60% of "junk" attack traffic, such as those using mal-formed TCP packets, or SYN-ACK packets that have not been preceded by a SYN packet.

Network based intrusion prevention is deployed in-line with the traffic being monitored (as compared with network based intrusion detection which passively monitors a mirrored copy of the traffic). Typically deployed at the interface between an enterprise and the ISP, the device is usually targeted at inbound traffic from the Internet to the enterprise. However there is no reason why such a device could not also examine outbound traffic. In this way the enterprise protects itself from Internet-based external attacks and at the same time, assists the community in general by limiting the initiation of attacks originating from inside the enterprise network.

If we now apply this concept of a “bi-directional intrusion prevention” facility to a consumer ISP service, treating the consumer access in a fashion similar to that of an enterprise network, we can filter the “low hanging fruit” of attacks at the edge of the ISP (before it is passed on to the general Internet or on to the rest of the customer access network). This could reduce the amount of attack traffic reaching consumer desktops (where we are most vulnerable) and at the same time significantly reduce attack traffic entering the Internet in general. This would trap obvious hacker traffic and any attack traffic generated by compromised consumer desktops.

The limitation today on this solution may be the processing capability and availability of intrusion prevention products. Placing NIPS intelligence on or near core ISP routers will be the least expensive but will require very high performance NIPS processors. As network based intrusion prevention is still in its formative stages, products capable of processing such traffic may not currently be cost effective. While the core network traffic of large ISP’s may be measured in multiple gigabits, NIPS intelligence could alternatively be located on the upstream interface of access routers where traffic is measured on Mbps. This would require less NIPS horsepower but would likely be more expensive (from a capital and operational perspective) due to the greater number of NIPS instances.

Netscreen’s Intrusion Detection and Prevention (IDP) product (<http://www.netscreen.com/products/idp.html>) is an example of a network based intrusion prevention product.

5.1.1.2 Address Filtering at the Edge Router Ingress of ISP Networks (The Spoof Killer)

As described above under the description of Reflector DDoS, the attack is carried out using request packets whose source address field contains the address of the target of the attack. This spoofing of the source address is fundamental to this type of attack. Source address spoofing is a way to make it more difficult to trace the source of an attack, and in the instance of reflector DDoS attacks, is also the mechanism for directing the attack traffic at the target. If source address

spoofing could be minimized or eliminated, a significant tool in the attackers tool kit could be impaired.

Internet routing works because all routers have information about the networks to which they are connected. In particular, edge routers “know” what networks or sub-networks are on each of their interfaces. In this way they are able to advertise to their peer routers which networks are reachable via their interfaces. ISP access routers and ISP core routers know what subnets reside on their interfaces. In many cases the ISP “owns” the subnets and is renting these addresses to their customers. If an ISP access router sees a packet inbound from a customer-facing interface with a source address that does not fall within the range of addresses defined for that interface, that packet can (must) be dropped. It is either a packet from an improperly configured host or it is a spoofed packet. In either case it should not (must not) be forwarded onto the general Internet. Simply by dropping such packets ISP’s will be doing the Internet community an enormous favor. This discussion does not consider the possibility that an edge router (or any router) can be compromised and its routing tables corrupted.

Filtering of this nature will not catch all spoofing. For example, if the spoofed address is a valid address on that interface, it will pass the filter successfully (for example, a reflector DDoS in which the drone/attacker is attacking a target on the same network). However it would catch enough to warrant deployment of such filters. In the author’s opinion, this is a virtually no-cost adjustment to most ISP edge routers.

5.1.2 Viewing ISP Customers as an “Enterprise” Environment – a Potential New Revenue Source

Recognizing that ISP’s are generally not non-profit organizations, it is reasonable to question how ISP’s could finance any of the “community-driven” security countermeasures outlined. It is also reasonable to assume that most consumers want a secure desktop and would value a service that would help to “secure” their desktops for a reasonable fee.

A brief comparison between Enterprise network environments and ISP customer environments reveals some basic similarities:

- The desktops may be geographically dispersed making physical visits costly; remote access for management is preferable;
- The desktops may implement a wide variety of OS versions and applications;
- The desktop users may possess differing levels of computer literacy;
- The desktop users are totally dependent on the network for access to Internet-based services;

And some fundamental differences:

- Enterprise security policy is enforceable (in theory) due to the business relationship between the enterprise and its users; ISP's are not in a position to enforce security policies except in very limited circumstances (as might be required by privacy guarantees or emerging national and international law);
- Enterprise network users are typically not able to change "providers"; ISP customers often have a selection of providers from which to choose; ISP's may not enforce security if there is a risk of losing customers.

This loose comparison suggests that enterprise-scale centralized desktop security management solutions might be viably applied to an ISP's customer base. It may only be a matter of scale (i.e. cost). The requirements are similar:

- Enforce desktop antivirus by pushing updates to the desktop and denying Internet access to desktops with dysfunctional antivirus onboard;
- Perform network based anti-virus scanning of network traffic on behalf of the user;
- Perform remote vulnerability scanning with notification and logging;
- Perform firewall services on behalf of the user desktop; the enterprise would deliver this as a single firewall service, while an ISP would have to deliver this on a per-user basis as each user desktop could require a different firewall configuration;
- Perform network based intrusion detection and prevention on behalf of the user with notification and logging;
- Perform network based content filtering on behalf of the user.

The challenges to deploying centralized security management in an ISP environment are considerable. Among the obvious would be:

- Variations in desktop configurations (hardware and software) will make pushing security measures onto the user desktop risky; ("My %#\$#% ISP just blew up my desktop again! I'm switching to the other ISP.");
- Some user desktops may be incompatible with the ISP's management solution;
- ISP legal liability (this could be the subject of a completely separate paper);
- Costs for core/edge/access router upgrades to implement security measures could be significant; (but the payback period may be short?);
- Architecting the insertion of security countermeasures in a manner that does not disrupt or impair service would be "interesting".

The challenges notwithstanding, the rewards could be substantial, both for the ISP financially (ISP's may live or die based on their success in building services-based revenues) and for the Internet community as a whole.

5.1.3 Enterprise Solutions

Is it feasible to transport an Enterprise security management solution into an ISP environment? A number of Enterprise solutions are worth considering. Table 9 lists several current offerings.

Symantec Managed Security Services	http://enterprisesecurity.symantec.com/SecurityServices/content.cfm?ArticleID=682&EID=0
Trend Micro Enterprise Protection Strategy	http://www.trendmicro.com/en/products/eps/eps/evaluate/overview.htm
Freedom Security & Privacy Suite	http://www.freedom.net/products/suite/index.html
Enterprise Privacy Manager	http://www.zeroknowledge.com/business/epmproduct.asp

Table 9 - Enterprise-Scale Centralized Security Management Products

5.1.4 An Example of ISP-Based Integrated Security Services

TELUS Communications is a Canadian telecom carrier, Network Service Provider and ISP (www.telus.com). Through an affiliation with Zer0Knowledge™ (www.zer0knowledge.com) TELUS is offering a security service to its DSL and dialup customers (“TELUS to Provide Online Internet Security Services For Consumers”⁹). Referred to as “Freedom® Internet Security Services”¹⁰, the service is a co-branded package of Zer0knowledge™ products offered to TELUS customers on a subscription basis. The offering includes desktop based personal firewall, antivirus and content filtering, pushed to user desktop and managed centrally by TELUS. The offering is currently only available for Windows 98/ME/2000/XP. It is not available for Macintosh or Linux. This is an early offering. It is, for the most part, desktop based and does not introduce the core and edge router countermeasures discussed in this paper. It does however provide a potential proof-of-concept for deployment of ISP-provisioned security services.

6 Summary

This paper has attempted to demonstrate the threat to the public Internet represented by the continued vulnerability of consumer desktops. It has also attempted to demonstrate why the ISP community is uniquely positioned to address the problem. Whether ISP's will step up to this challenge is debatable. One might envision a public Internet split apart into three or more interconnected "security domains":

- The "wild" Internet to which anyone can connect without restriction and within which all desktops are responsible for protecting themselves (or not);
- A "restricted" Internet where connectivity is only *permitted* if the connecting desktop (or desktops in the case of an autonomous network) comply with a defined security "profile";
- And a "secure" Internet where end-to-end security is guaranteed and maintained by a provider.

Only time will tell.

© SANS Institute 2003, Author retains full rights.

7 References

- [1] CERT/CC Statistics 1988-2002
Carnegie Mellon University
Software Engineering Institute
CERT Coordination Center®
http://www.cert.org/stats/cert_stats.html
- [2] “The Strange Tale of the Denial of Service Attacks Against GRC.com”,
Steve Gibson, Gibson Research Corporation
<http://grc.com/dos/grcdos.htm>
- [3] CERT/CC Incident Note IN-99-07
Carnegie Mellon University
Software Engineering Institute
CERT Coordination Center®
http://www.cert.org/incident_notes/IN-99-07.html
- [4] Chakrabarti, Anirban and Manimaran, G., “Internet Infrastructure Security: A Taxonomy”, IEEE Network, November/December 2002, Vol. 16 No. 6: pages 13-21
- [5-1] The UCLA Internet Report: “Surveying the Digital Future,” UCLA Center for Communication Policy, November 2000, <http://WWW.CCP.UCLA.EDU>
- [5-2] The UCLA Internet Report 2001, November 2001
- [5-3] The UCLA Internet Report – Year Three, February 2003, page 24
- [6] “Symantec Internet Security Threat Report Attack Trends for Q3 and Q4 2002”
Volume 3, February 2003
<http://enterprisecurity.symantec.com/Content.cfm?articleID=1964&EID=0>
- [7] Paulson, Linda Dailey, “Stopping Intruders Outside the Gates”
IEEE Computer Magazine, November 2002, pages 20-22.
- [8] Chang, Rocky K.C., “Defending against Flooding-Based Distributed Denial-of-Service Attacks: A Tutorial”, IEEE Communications Magazine, October 2002, Vol. 40 No. 10: pages 42-51
- [9] “TELUS TO PROVIDE ONLINE INTERNET SECURITY SERVICES FOR CONSUMERS”, zer0knowledge™ Media Press Release, September

18,2002;

<http://www.zeroknowledge.com/media/pressrel.asp?rel=20020918>

- [10] Freedom® Internet Security services, by TELUS
<http://www.telus.com/cgi-ebs/jsp/viewitem.do?content=FreedomInternetSecurity&Region=A>
- [11] Rogers, Lawrence R., "Home Computer Security"
Carnegie Mellon University
Software Engineering Institute
CERT Coordination Center®
<http://www.cert.org/homeusers/HomeComputerSecurity/>

© SANS Institute 2003, Author retains full rights.