iDEFENSE

VeriSign

// THE POWER OF INTELLIGENCE

Distributed Denial of Service (DDoS) and Botnet Attacks

An iDefense Security Report
April 28, 2006

INSIDE THIS REPORT

1 Executive Summary

The distributed denial of service (DDoS) attack is among the most potentially costly and intractable cyber threats facing technology-dependent companies today. DDoS attacks are also more frequent, larger and more costly than ever before, and the number of available "zombie" computers in the wild is greater than ever. The commanders of bot armies are more numerous, more sophisticated, harder to identify and have better tools than at any time in the past, and these trends will continue for the foreseeable future. This report discusses why and what DDoS mitigation and prevention strategies are used to keep technology-driven organizations in business today, and how early DoS attacks evolved into present-day techniques.

There are many different variants of DDoS attack, but these can generally be classified as one of two types: agent-handler attacks or resource depletion attacks. While each attack type has different strengths and weaknesses (depending largely upon the type of system targeted), one of the more relevant trends that attackers are currently showing is the propensity to use multiple attack tools in "waves" constituting a single attack. This technique increases the necessary skill required to launch such attacks, but it also increases their severity and problems for mitigation.

A wide variety of mitigation tools exist, either in the form of tools that can be bought and integrated into one's network or in the form of services procured from firms. Again, the size of the network, its available bandwidth and the expertise of its staff determines whether any of these mitigation tools would be redundant or unnecessary. However, given the increasing scope and seriousness of attacks, companies may wish to evaluate several of the most capable DDoS mitigation services, especially those provided by Prolexic and Arbor Networks. Moreover, the size and notoriety of most large corporate networks should ensure the focused assistance of its preferred ISPs.

2 Timeline Evolution of DoS/DDoS attacks

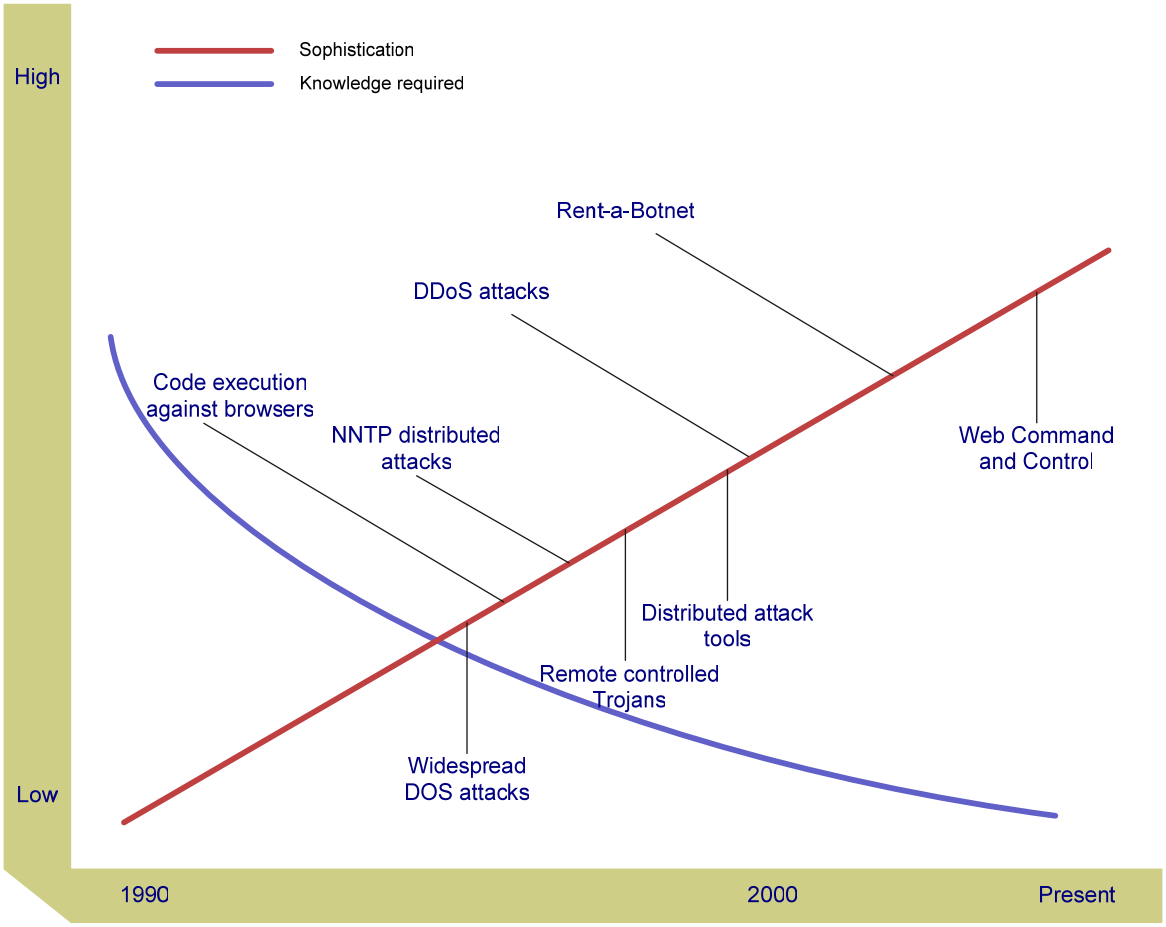

The evolution of DoS and botnet attacks

The sophistication of DoS attacks has increased, even though the average attacker need not be as technically proficient as in the past because so much of the pre-attack organizational process has been automated. Indeed, as soon as attackers master the basics of a new technique, automation tools tend to appear quickly thereafter. This, in turn, offers a wider range of options for lesser-skilled attackers (i.e., "script kiddies") to build or simply purchase bots. In addition, these tools enable attackers to take advantage of newfound vulnerabilities quickly, before they can be patched, and to "recruit" many more zombie computers. Indeed, since the mid-1990s, not only have bot-management tools grown in power and efficiency, but so also have the malicious code variants used to infect the zombies. Since the emergence of DDoS attacks, specifically flood attacks, botnet controllers have begun to rent their armies in full or in part to anyone willing to pay the price. It is generally suspected that many of these new bots-for-hire are used for spamming and mass visitations to “pay-per-click” sites, although the armies used for attacks have also increased. Ultimately, the result has been an increasingly diverse and growing bot population organized in more transient clusters that can be easily disaggregated and employed in a wider range of tasks.

3 The Current State of DDoS Attacks: The Worst Web Problem?

3.1 Preliminary Caveats: Why Experts Do Not Know More

Few topics in information security are as difficult to study or as hampered by self-interested obfuscation as the DDoS threat. The past several years has produced an extensive body of literature on DDoS tactics and mitigation techniques, though the problem continues to defy deep understanding. Information security officers, consultants, hackers, journalists, academics and government employees have all offered their experiences and analyses. There now exist dozens of books and hundreds, if not thousands, of book chapters, articles and web pages devoted to the investigation of DDoS attacks. Despite the hard work of so many researchers, expert consensus exists only on the most basic and general aspects of the issue. For instance, most experts agree on a generic typology of DDoS attacks, and virtually all agree that attacks are growing more frequent, more powerful, more nuanced and more difficult to detect. Beyond these obvious assessments, consensus collapses.

Of course, the greatest impediment to an accurate understanding of the threat is the marked lack of useful data. There are three important and presently insoluble reasons for this. First, any company that has experienced a DDoS attack has few incentives to report it. The traffic logs of each user could be valuable pieces of evidence were they ever aggregated with a multitude of other such records; individual logs alone can provide some insight, but nothing approaching an understanding of the true scope and severity of the problem. The primary disincentives to release such logs are, of course, negative publicity and potential legal liability for negligence.

Second, there is no widely agreed-upon method that researchers might achieve a reliable, indirect inference about global DDoS activity, although many have tried. Some innovative researches have crafted highly creative and mathematically complex models to approximate the activity of botnets. Honeynets and "backscatter" analysis are two of the more promising approaches, but even these studies only served to illustrate the utility of the methods, rather than uncovering conclusive patterns in the data (see for instance, D. Moore, G.M. Voelker, and S. Savage, "Inferring Internet Denial-of-Service Activity," *Proc. Usenix Security Symp.,* Usenix Assoc., 2001, and The Honeynet Project & Research Alliance, "Know Your Enemy: Tracking Botnets: Using Honeynets to Learn More about Bots", *White Paper*, March 13, 2005, http://www.honeynet.org/papers/bots/) In addition, since the data sets are drawn through random sampling, these inferential studies are likely to illuminate only the random, experimental types of attacks, rather than the more dangerous targeted types. Still, the achievements of such research are notable and should be built upon and drawn into the broader debate about the DDoS threat.

Finally, the threat is inherently international and anonymous. It is more difficult to ascertain the identity of a "bot herder," especially one residing outside the US, than many other malicious cyber criminals. Indeed, clever organization, anti-forensic techniques and new command and control techniques are making it more difficult. Recent arguments put forth by Jim Crowcroft, professor of Communications Systems at Cambridge University, indicate that bot herders will soon be able to control their armies via Voice over Internet Protocol (VoIP) systems (Peter Judge, "Cambridge Prof Warns of Skype Botnet Threat", *PC Advisor*. January 26, 2006). However, for now, the botnet commanders keep their armies secret by maintaining different cohorts on different servers, and occasionally rotating different segments among several servers so that no individual bot remains on one server for very long. Add to this the increasingly popular practice of "renting" botnets to lesser-skilled users, and it becomes clear why chasing down a DDoS attacker is among the most difficult tasks facing law enforcement and IT security personnel.

In combination, these factors have made DDoS attacks a murky area of study that evolves quickly and grows even faster. As such, almost all of the following findings must be taken with some caution. This is not to say that the data is useless, but rather that they cannot be assumed to accurately represent the universe of cases; there are exceptions to almost every rule.

3.2 Attack Motivations

There are three basic types of motivations for DDoS attacks: experimental/challenge-driven, principle-driven and illicit economic gain (John E. Dunn, "DDoS Attacks Still Biggest Threat", *Techworld*, Oct. 13, 2005, http://www.techworld.com/security/news/index.cfm?NewsID=4570). While the first two likely provide the impetus behind much of the innovation in tools and tactics that have made DDoS attacks as powerful as they are today, it is only the economically driven attacks that should be of concern to large US financial institutions. Thus, the other motivations will be discussed only briefly while the economically motivated attacks will receive more attention.

3.2.1 DDoS as Experiment or Challenge

This is the most common motivation driving the script-kiddie attacker. Those hacking enthusiasts who are first learning "tricks of the trade" may find some value in learning how to execute DDoS attacks simply for the knowledge of knowing how to do so. This motivation is a main reason why the majority of attacks are of brief duration and why their targets are often randomly chosen. Of course, since many of these attackers are likely to be young and somewhat risk acceptable, they may incrementally advance to more sophisticated tactics simply because they enjoy the challenge or because launching large and successful attacks may glean some respect from peers.

Of course, youthful exuberance can easily transform into pernicious rivalry. According to researchers at Arbor Networks, one of the most common reasons for DDoS attacks are rivalries between online Internet gaming groups. The rationale is quite simple: adolescents who group to compete with each other do, from time to time, allow their rivalries to spill outside the boundaries of the game. This concept should be familiar to former athletes who may have played pranks on rival schools.

While this is a persistent nuisance to ISPs, these attacks are of almost no significance for large financial services firms. They should not be large enough to deplete problematic amounts of bandwidth, and will almost always cease within minutes on their own, or the target employs the most basic countermeasures (discussed in the following).

3.2.2 Principle-Driven Attacks

Many companies separate politically motivated attacks from protest flood attacks, but this distinction is generally only a reflection of researchers' own biases. iDefense makes no such distinction because it feels the core motivation is the same: a desire to silence someone whose values are seen as inimical to one's own. Thus, whether it is a Web "vigilante" flooding pornography sites, or a religious extremist using a DDoS to shut down the site of a rival sect's government, the motivation is the same.

There are several potential reasons that could inspire values-driven attackers to attack: anti-capitalism, anti-Americanism, anger with past service, sympathy with those in credit card debt, or ethical objection to some business affiliation or investment a given organization maintains. However, the first two of these could just as easily lead to attacks against dozens of companies, so they should not be regarded as likely. In general, attackers with this motivation will eventually find a more satisfying target or will remain on

the offensive long enough to be identified. In either case, they are ultimately less dangerous than attackers driven by the third type of motivation, sabotage and extortion attacks.

3.2.3 Sabotage and Extortion

Extortion is easily the most dangerous type of DDoS attack, and for many companies sabotage is the second. These attacks have in common an economic incentive and a reason to target the victim directly. They differ only in the profile of the likely attacker and in the fact that sabotage is intended only to harm, while the damage inflicted in an extortion attack is secondary to the desire for financial gain.

The saboteur would be either a former employee who felt cheated in some way or a competitor. Neither of these is very likely. Former insiders would have many other opportunities that could be achieved more easily and could generate more damage. Competitors would face extreme costs if identified as complicit in the attack. That said, sabotage or "competition-driven" attacks have occurred before, but all recorded cases have been between smaller companies and potential rivals (Denise Pappalardo and Ellen Messmer, "Extortion via DDoS on the rise", *Computerworld*, May 16, 2005, http://www.computerworld.com/securitytopics/security/story/0,10801,101761,00.html?source=NLT_SEC) .Interestingly, the competition is not only among companies, but also takes place among organized cyber crime gangs themselves, with botnets as their preferred tools of attack. This is a tactic long employed by credit card fraudsters against their rivals in the underground marketplace; with plenty of bots available for spamming purposes, they made an obvious tool to disable the competition and drive business to their own boards.

With these elements in mind, the most serious DDoS attackers are extortionists. Not only are they likely to be the most technically advanced users, but they are also probably the most experienced in extortion negotiations. There is no way of knowing how severe or widespread the problem of DDoS extortion is. iDefense analysts suspect that fewer than one quarter of all companies who suffer from DDoS extortion ever report the problem. This is perhaps the greatest strength of the extortionist, as security researchers remain without enough data to generalize about the true scope of the problem and the processes involved. Private means of tracking an attacker are usually more expensive that the payment demanded by the attacker, and alerting the police could create reputation costs that would be more expensive still.

Many commentators note that victims of DDoS attacks should never give in to attackers' demands. "If the companies pay" says one reporter, "the attacks will continue." This mainstream line of thinking is not so much wrong as merely incomplete. The truth is that DDoS attacks will continue whether companies pay or not. Moreover, many analysts often overlook the fact that once a DDoS attacker promises not to attack after the ransom is paid, it will be difficult to extract further payment if the attacker's word has been broken.

"Almost all of them have an international connection...There aren't many cases where people are doing this from the U.S., and many times it is a juvenile subject to the laws of another country", said FBI agent Paul Brasser (John E. Dunn, "DDoS Attacks Still Biggest Threat", *Techworld*, Oct. 13, 2005, http://www.techworld.com/security/news/index.cfm?NewsID=4570)

3.3 Frequency and Duration

With the aforementioned caveats in mind, it is not surprising that there is no authoritative figure pertaining to the frequency or duration of DDoS attacks over time. Some estimates reach as high as 8,000 attacks per day while others are as low as 1,000. It is not known how many companies have suffered DDoS attacks, although the most rigorous study to date puts the number at 17 out of every 100

organizations. More than 90 percent of ISPs surveyed cited simple "brute-force" TCP SYN and UDP datagram floods from zombie PC networks as their biggest daily problem (John E. Dunn, "DDoS Attacks Still Biggest Threat", *Techworld*, Oct. 13, 2005, http://www.techworld.com/security/news/index.cfm?NewsID=4570). By another measurement, a sophisticated and far-reaching backscatter analysis detected an average of 12,000 attacks against 5,000 unique targets over three 21-day periods (D. Moore, G.M. Voelker, and S. Savage, "Inferring Internet Denial-of-Service Activity," *Proc. Usenix Security Symp.,* Usenix Assoc., 2001)

While these numbers are significant, the problem is not yet so serious that affected companies' Web presences are made unavailable. One explanation for the current frequency of DDoS attacks is that the vast majority of attacks are of brief duration, most likely because they are "practice attempts" by novice bot herders.

The average duration of the average reported DDoS attack is as low as a few minutes, although simply stating the average does not provide a very good snapshot of activity. Rather, a synthesis of the best studies suggests a probability distribution heavily skewed toward brief attacks, the vast majority lasting between 10 and 25 minutes. In fact, fewer than .5 percent of all attacks last longer than an hour. While some sources mention prolonged attacks of weeks or even months, these constitute less than .01 percent of all known attacks. Nevertheless, iDefense retains case studies of organizations that have been under attack for no less than one year. Of course, these attacks are most likely to be the largest and most directly targeted against specific organizations.

Unlike other characteristics of DDoS attacks (e.g., bandwidth consumption and frequency), there is no evidence to indicate whether attacks are growing longer or shorter over time. The most that can be said conclusively is that the capacity exists to launch longer attacks, although there are few compelling reasons why bot herders might wish to do so, as this increases the chance they may be identified.

Another important consideration is the degree that the same targets are subject to repeated attacks. The majority of data suggests that organizations that have already suffered a DDoS attack are more likely to suffer additional incidents. One reason for this is that attacks tend to come in waves as bot herders attempt to find the minimum number of bots and the optimum configuration of different bot types to flood a specific network. Another factor is that targets proven to be vulnerable are more attractive to other malicious actors.

3.4 Bot Army Size and Bandwidth Consumption

It is clear that the number of total zombie computers and the number used in any given attack are increasing and have been since DDoS attacks were first noticed. The first publicized DDoS attacks in 1998 employed several hundred zombies, but within one year, the average attack cohort was several thousands of bots. Attacks of this size were initially capable of several megabits per second (Mbps), but this quickly grew to several hundred Mbps by 2001 as bot herders amassed armies containing tens of thousands of infected computers. Enabling the formation of such large armies was the rapid spread of continuous broadband subscriptions; with users' computers constantly online, bot herders found that they could infect and use them whenever they wanted. Between 2000 and 2002, the number of infected zombie computers increased by 10,000% (prolexic Technologies, Distributed Denial of Service Attacks, *White Paper,* Q4, 2004, p. 2, http://www.prolexic.com/downloads/whitepapers/Prolexic_WhitePaper-DDoS.pdf). The highest cited figure for a single botnet stands at 1.5 million, though the vast majority are surely not yet this large (Ryan Naraine, "Return of the Web Mob", eWeek, April 10, 2006 http://www.eweek.com/article2/0,1895,1947561,00.asp).

The results of an increasing zombie population are reflected in trends of bandwidth consumption. The latest observed record (from December 2005) for bandwidth consumption was 10 Gbps, a marked increase over previous estimates. Also, recent survey data from Arbor Networks indicate that several respondents had been witness to attacks consuming greater than 10 Gbps. To put these figures in perspective, 10 Gbps amounts to a full 500 percent increase over Cisco's best routing appliance, and far outstrips other hardware such as load balancers. Add to this the potential growth of DNS recursion-amplification attacks, and the 10 Gbps figure could easily increase by a factor of five, 10 or more.

These larger sizes required novel forms of organization to maintain them. First, bot herders learned that keeping massive bot armies together on one or two servers increases the likelihood of detection. As such, they began to split their armies into smaller units, usually a few thousand strong, and rotating these among a handful of servers. Second, the hackers began differentiating between the most capable and the slowest computers in their armies. Some began creating an "officer cadre" of fast, capacious computers that passed on commands to slower, older systems. Some iDefense hacking experts envisage that such functional differentiation of bots could be used in attacks themselves, with an "elite unit" of fast, newly infected and state of the art computers to complement a normal bot army. The elite unit could be tasked with overpowering the most sophisticated countermeasures, leaving the older, slower bots to do the more routine work seen in today's standard attacks.

While these sophisticated attacks are the most difficult to defeat, they are not nearly as common as simplistic, brute-force-type flood attacks. When advanced attacks do occur, they are conducted by experienced hackers, usually with some concrete goal in mind. The majority of DDoS attacks, by contrast, are conducted by "script kiddies" working with free tools supplied by more experienced users.

4 The Range of DDoS Attack Tools and Tactics

DDoS attacks were first made possible by tools clusters that enabled a would-be bot herder to infect large numbers of computers and retain command over some share of their processing power. Early tools such as Trinoo have since evolved into one of many components included with malicious code. There are two main models for DDoS attacks: Agent-Handler and IRC. The following image shows how an attack typically takes place.

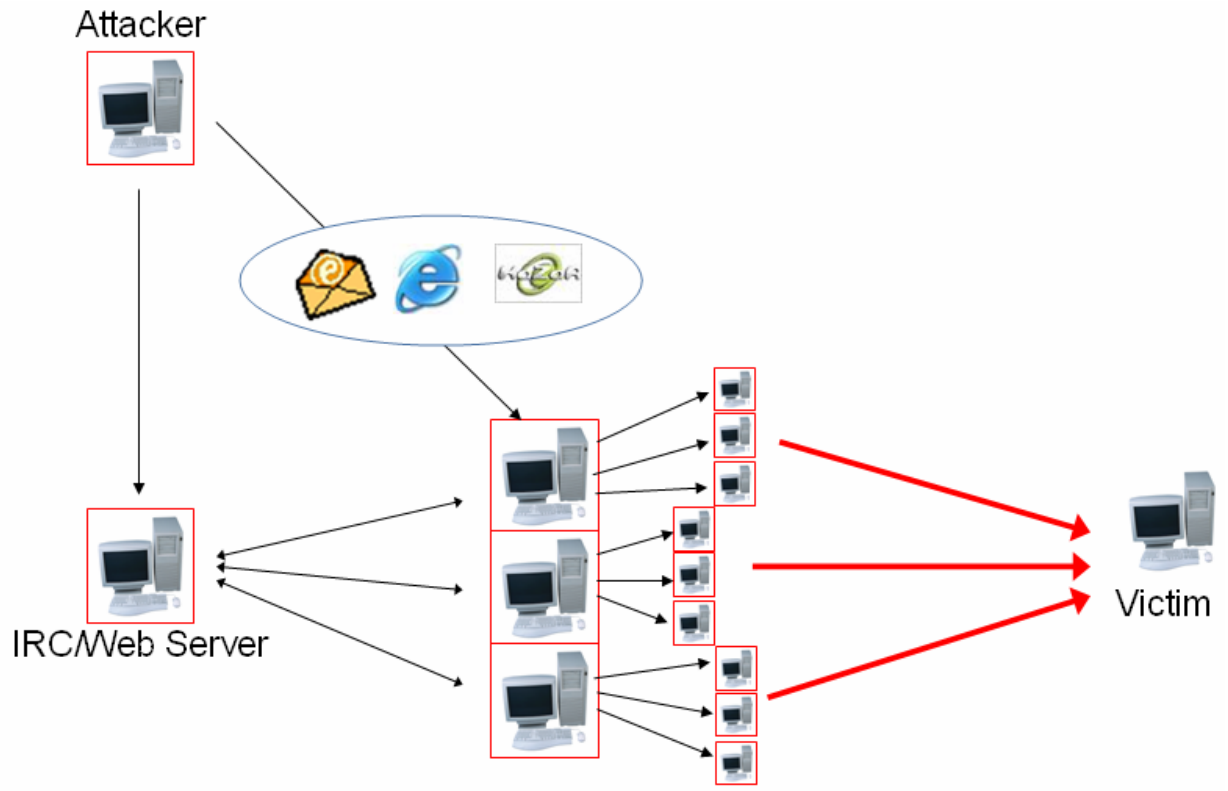

Typical DDoS Attack: Although DDoS attacks differ greatly the above diagram shows how an attacker forms a botnet by compromising many computers, which are then organized into a botnet. This botnet is then able to generate large volumes of traffic that can overwhelm the resources of a victim.

4.1 DDoS Attack Variants

There are two main types of DDoS attacks: bandwidth depletion and resource depletion. The former inundates a target with massive amounts of request data that eventually prevents the flow of legitimate traffic, while the later attempts to overwhelm the processing power of the target. Bandwidth depletion attacks have become the preferred method of most attackers because resource depletion attacks often depend on bugs or facile configurations in systems or networks that, over time, can be fixed.

4.1.1 Bandwidth Depletion

Bandwidth depletion involves flooding a victimized target with unwanted traffic. This prevents legitimate traffic from reaching the target. Examples of bandwidth depletion attacks include flood and reflection attacks, such as "ping of death," "smurf" and "fraggle."

There are two primary types of bandwidth depletion attacks: flood and reflection. IRC bots are commonly used in flood attacks, sending multiple packets to the same target at the same time to congest traffic to a specific website. A reflection attack is very different, where many packets are sent out to many computers, with a spoofed source address. When the computers all respond to that spoofed source address, the victimized computer in this case, traffic becomes congested and the computer may crash.

4.1.1.1 UDP Flood Attacks

UPD packets are often referred to as "send-and-forget" since it is a connectionless protocol. UPD flood attacks are simple, and send a large volume of UDP packets to a target to saturate the targeted network. Attacks commonly take place against random victim ports so that UDP packets cannot be easily filtered by a network under attack. If no services are running on a port targeted in a UDP flood attack, an ICMP packet (destination port unreachable) is normally sent to the source IP address, further tying up resources on the targeted network. The source IP may also be spoofed in a UDP attack to better conceal the location of the attacker and to improve performance of a DDoS attack. This frees up zombies and victimized computers to send out UDP packets in an attack, not having to reply to any ICMP packets or return traffic from the target of the attack.

4.1.1.2 ICMP Flood Attacks (Ping of Death)

This is similar to a UDP flood attack but involves sending out multiple ICMP_ECHO_REPLY (ping) packets to a target. This naturally overloads the targeted computer when it attempts to reply to each ICMP packet. In the case of a "Ping of Death" attack, a single large ICMP ECHO REQUEST packet is sent to the target, which may cause the target to become unstable or crash. This is a very old technique that simply involves the attempt to send a single packet that exceeds 65,535 bytes.

4.1.1.3 Reflection Attacks (Smurf and Fraggle)

Reflection, or amplification, attacks involve the sending out of many packets with a spoofed IP source address. This results in a large volume of return packets being sent to the spoofed IP address, the target of attack. These are older types of attacks that are fairly easy to mitigate with updated routers, firewalls, and other solutions today.

A smurf attack involves sending out of many ICMP ECHO REQUEST packets to many computers (handlers) with a spoofed IP source address. Each handler then returns an ICMP ECHO REPLY packet, congesting traffic of the targeted network. A "fraggle" is similar to a smurf attack but involves UDP ECHO packets instead. More information is available from CERT at http://www.cert.org/advisories/CA-1998-01.html.

4.1.2 Resource Depletion

Resource depletion, or protocol exploitation, involves a target that specifically attempts to deplete resources on the targeted computer or cause it to become unstable and crash. Examples of resource depletion attacks include synflood and teardrop.

4.1.2.1 TCP SYN Attacks (Synflood)

A TCP SYN attack involves sending multiple SYN packets to a target in an attempt to overload it. This is the type of attack that took place in a recent court case involving IRC zombies, which are capable of performing many types of DDoS attacks in most cases.

The TCP handshake involves SYN (synchronization) and ACK (acknowledgement) packets to establish communications between two resources. A SYN packet is first sent to a resource, awaiting an ACK response. Once this process is completed between the two computers the TCP handshake is completed and communications may commence between the two resources. During this process each computer awaits handshake packets, and if they are not received in a timely manner, may generate new SYN and ACK packets respectively. By sending a target many SYN packets, it may successfully overload the target as it attempts to process each handshake request.

4.1.2.2 PUSH and ACK Attacks

This is similar to a SYN attack but involves TCP packets with PUSH and ACK bits set to a value of one. This instructs the target to load all data into a TCP buffer to then send an ACK when finished processing. When many packets of this nature are sent to a target, it may overload the buffer and cause the target to crash.

4.1.2.3 Recursive HTTP Flood (Spidering)

This attack involves "spidering" a website via the HTTP protocol, in a recursive manner, to deplete resources on the targeted Web server.

4.1.2.4 Teardrop (Bong and Boink)

This attack exploits TCP/IP IP stacks that do not properly handle overlapping IP fragments. More information is available from http://www.cert.org/advisories/CA-1997-28.html. This may result in a host crash. This is an older attack that is now easily mitigated by most updated firewalls and systems.

4.1.2.5 Land

A land attack involves a specially crafted IP packet with the source address and port set to be the same as the destination address and port. See http://www.cert.org/advisories/CA-1997-28.html for more information.

4.1.2.6 Naptha

This attack attempts to exploit vulnerable TCP/IP stacks using crafted TCP packets. The attacker must create large numbers of TCP connections and leave them in certain states in an attempt to deprive the host of resources to the point of failure. Naptha does not use a traditional network API to set up a TCP

connection. It does not keep any record of a connection state. It responds to a packet sent to it based upon the flags in that packet alone. More information can be found on Naptha attacks at http://www.securiteam.com/securitynews/6B0031F0KA.html.

4.2 Major DDoS Tools

Attack tools of yesteryear, such as Trinoo, and have steadily been replaced with more agile and easily controlled IRC bots and similar codes of today. For details on historical attack tools, view extensive documentation online at locations including:

Trinoo, http://staff.washington.edu/dittrich/misc/trinoo.analysis.txt
Tribe Flood Network (TFN), http://staff.washington.edu/dittrich/misc/tfn.analysis.txt
Stacheldraht (German for 'Barbed Wire'),
http://staff.washington.edu/dittrich/misc/stacheldraht.analysis.txt
Trinity, http://www.securiteam.com/securitynews/5GP011F2KO.html
Shaft, http://www.securiteam.com/securitynews/5AP0F000IM.html

By far, IRC based bots are the most common DDoS tool and technique today. For example, the HangUP Team is well-known for major attacks involving Trojans and bots including Korgo in 2004. Following is a screenshot of a bot room called "#Waffen-ss" taken in 2004 when Korgo variants were spread in the wild.

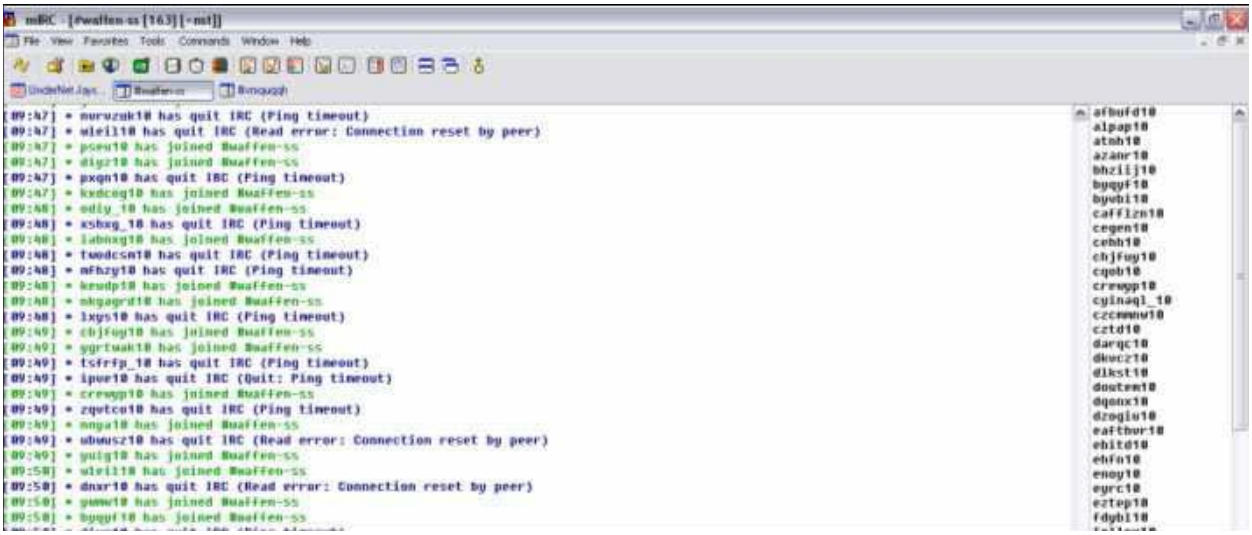

Today the sources codes of some of the most successful malicious codes to date are readily available on the underground, including PhatBot, MyDoom, Bagle, Cabir, CodeRed, LoveLetter, Kournikova, BubbleBoy, Kak, FunLove, Melissa, Happy99/SKA and more.

These types of source codes can be combined with newer exploits and DDoS scripts to increase the functionality of an attack. Some even have DDoS functionality already built into them, such as PhatBot and MyDoom.

A variety of malicious code tools have emerged over the past few years, whether designed as standalone DoS tools, toolkits or components to be used with other codes. Data from Prolexic Technoliges and CERT indicate a steady increase in codes since the major media event of Mafiaboy attacks in 2000 (Prolexic Technologies, "Distributed Denial of Service Attacks, *White Paper*, Q4, 2004, http://www.prolexic.com/downloads/whitepapers/Prolexic_WhitePaper-DDoS.pdf). The massive peak in 2004 is likely due to the emergence of source codes and heavy trading of tools and code during that time period.

4.3 DDoS via Recursive DNS Queries

Early in 2006 many nodes of the global DNS system were used by a malicious actor to conduct large-scale DDoS attacks. These attacks began in January and peaked in early February used widely available name servers that are configured to allow openly recursive DNS queries. This technique provides significant amplification of attack traffic and affect the DNS system as well as the intended victim. Specific information on these attacks can be found in ICANN's SSAC Advisory SAC008 located at http://www.icann.org/committees/security/dns-ddos-advisory-31mar06.pdf .

By design the DNS architecture is recursive in that name servers refer each other through the DNS cache hierarchy when retrieving a particular DNS record. In order to answer a DNS query several name servers communicate before locating the desired record. However, just as mail servers were once set to relay mail by default, many name servers' default configurations are said to be openly recursive, meaning that they will respond to queries from non-trusted sources. This open recursion setting allows an attacker to use the name server to reflect a response to a target computer. Reflection is achieved by altering the packet containing the query to appear as if it had been sent from the target, thereby causing the DNS server responds to the target instead of the attacker.

In addition, queries for a DNS record are relatively small in comparison to the responses they generate. Although the query-response ratio varies by query and by response record it is possible for a response to be 73 times larger than the originating query (SSAC Advisory SAC008 DNS Distributed Denial of Service (DDoS) Attacks, http://www.icann.org/committees/security/dns-ddos-advisory-31mar06.pdf, March 31, 2006). Known as amplification attacks, this technique takes advantage of the increased response size, enabling an attacker with limited bandwidth to overwhelm a network of with many times more capacity. For example, a T1 connection could be consumed by attack traffic originating from a dial-up connection. To ensure efficient amplification, an attacker will often one of the hundreds of thousands of name servers and replace a legitimate DNS record with an inflated one, usually of 4k in size, the maximum.

By submitting small queries known to return large records to openly recursive name servers, and by spoofing the originating IP address, an attacker can send an overwhelming amount of traffic to the target. Although researchers and malicious actors have been aware of this method of attack for many years, there has been in recent months an unprecedented renewal of interest and utilization for attacks. Although DDoS attacks of this type remain rare, this method is incredibly effective and substantially reduces the resources needed to attack well provisioned targets. Whereas an attacker may have needed 100,000 bots in a direct DDoS attack, fewer than 2,000 would be needed if using openly recursive name servers. This not only makes DDoS accessible to many more potential attackers but also makes even the most well-provisioned networks vulnerable to this threat. For this reason, the potential severity of this threat cannot be overstated.

Just as the use of openly recursive name servers to reflect and amplify malicious traffic has been well understood for many years, sound recommendations for minimizing the threat have existed for just as long. However, there is neither strong incentive for ISPs or name server administrators to take the necessary preventative measures. Name server configuration and egress filtering by ISPs would

undoubtedly help to reduce the possibility of these attacks, although it remains highly unlikely that a sufficient number of administrators will enact these simple measures.

Name servers need not act as relays to be used at will by attackers. Moreover, no technical barrier exists to minimizing or eliminating this threat. The real impediment to solving this problem is the failure of collective action to obtain. Despite strident recommendations from security professionals, including US-CERT (The Continuing Denial of Service Threat Posed by DNS Recursion, http://www.us-cert.gov/reading_room/DNS-recursion033006.pdf, March 30, 2006) the vast number of name servers that are configured for open recursion--estimated at well over 500,000--is far more than attackers need to cripple the most capacious and robust networks. Even if half of the currently open name servers were reconfigured to only answer requests from other authoritative name servers the problem would remain. As this issue gains attention among security professionals, especially if the attacks using openly recursive nameservers increases, mounting pressure on administrators of these servers may provide the necessary incentives to reconfigure their servers.

ISPs are also able to take steps to severely reduce the threat of these amplification attacks. Many ISPs continue to transmit packets that contain originating IPs that are outside of the network range. These packets almost certainly contain spoofed IPs, since all packets originating in the network should contain IPs belonging to that network. ISPs currently have little incentive to enable egress filtering as the amount of bandwidth these packets consume is trivial. However, as with name server administrators, if these attacks become commonplace and impact customers ISPs will feel increasing pressure to prevent their networks from hosting these attacks.

Neither name server administrators nor ISPs can be expected to invest any level of effort in this issue in the near term. All indicators suggest that in the near term openly recursive name servers will continue to be prevalent and that spoofed packets will successfully be delivered. In short, there is no near-term solution exists and there is very little other than lack of motivation to restrain attackers from perpetrating the largest denial of service attacks ever seen. As such, this issue should be regarded as one of the most serious threats to any nodes in the global telecommunications infrastructure.

4.4 Botnet Command and Control

Bandwidth growth has helped protect networks against small scale DDoS attacks and has made botnets an important component of modern day DDoS attacks. As bot herders have continued to recruit zombies and their botnets have grown several different methods of command and control have emerged. Bot herders strive to evade detection while maintaining flexible armies of bots with ease of command. Some challenges that have driven botnet command and control to evolve include bandwidth usage, communication interception or interruption, and stealth. Many bot herders rent their bots to others, whereby ease of control allows less sophisticated users to become customers. Also, botnets are constantly under attack from law enforcement, vigilantes and competing bot herders making command continuity a primary concern.

4.4.1 Agent-Handler

Agent-Handler DDoS attacks refer to the interaction among clients, handlers and agents. The client is what the attacker uses to communicate with handlers. Handlers are software packages on remote computers that are abused by the attacker for a DDoS attack. It is common for an attacker to launch DDoS attacks from a victimized computer (handler) to make it more difficult to trace the attack back to the attacker (client). Agents are the software on a computer that actually performs the DDoS attack. This may be considered a sub-set of the handler software residing on the same system. Multiple agents may be involved in DDoS attacks. Protocols used in such DDoS attacks may include TCP, UDP or ICMP.

Owners of infected or victimized (hacked) computers are often unaware of the situation and use of their computer in the DDoS attack. This further hinders investigations into DDoS attacks, giving the attacker additional anonymity. By using many handlers in an attack, each computer must only have a few resources available for a DDoS attack. This helps to conceal the malicious abuse of a handler computer.

4.4.2 Internet Relay Chat

Internet Relay Chat (IRC) DDoS attacks have quickly become the tool of choice for DDoS actors over the past few years. They are typically installed with a bot; bots commonly spread automatically against multiple vulnerabilities, weakly protected network shares, and through other methods. Once installed, a full backdoor typically resides on the system, including an IRC component that connects the computer to a remote IRC server controlled by the attacker. These bots or "Zombies" are agile, easily controlled, easily created and easily leveraged for cash in multiple ways (proxy sales to spammers, DDoS attacks, extortion, installation of ad/spyware, warez and theft of CD keys, credit card theft and more).

To launch a DDoS attack via IRC an attacker simply logs into a malicious IRC server, authenticates, and issues commands to many zombies at once or to individual bots within private windows. It is trivial to start and stop DDoS attacks via hundreds or thousands of zombies using this method. IRC botnet operators tend to keep their botherds smaller in size, by rolling out updates and many minor variants of code to have dozens if not hundreds of smaller botnets on various servers. This helps to avoid any single point of failure, such as when authorities shut down a hostile IRC server. It also helps to avoid detection on IRC servers, where high traffic is common and smaller botnets are not easily identified.

IRC based attacks may involve many different codes and protocols. As seen with Agent-Handler DDoS attacks, IRC attacks may also involve TCP, UDP or ICMP protocols.

4.4.3 Web-Based

Despite IRC's continued dominance as the preferred method for botnet command and control; web-based reporting and command has been emerging over the past two years. Some bots simply report statistics to a website, while others are fully configured and controlled through sophisticated PHP scripts and encrypted communications with bots over port 80 and the HTTP protocol. Among several advantages web-based controls offer over IRC are:

- Easier to set up and configure website
- Better reporting and command functionality
- Distributed load uses less bandwidth and allows larger botnets
- Use of port 80 helps to conceal traffic and makes filtering more difficult
- Not susceptible to botnet hijacking via chat room hijacking
- Ease of use means it's easier to rent out

Multiple malicious code authors are now starting to make use of Web reporting and control features, including HaxDoor, Torpig and Briz. This trend is likely to continue now that this type of code is ever present and support for such functionality is trivial compared to several years ago. Most importantly, Web-based command-and-control is much more scaleable for control of larger and more conspicuous bot-herds. This new development will probably force security professionals to consider how to best identify and filter out such traffic over TCP port 80 and how to rapidly respond to the changing threat environment.

4.5 Major Bot Families

IRC Bots are the most common form of malicious code used for DDoS attacks. These are normally semi-automated codes that can be carefully and remotely controlled by an attacker as a "robot." A significant increase in IRC bots began in 2004, as shown in the following graph and continues into 2006, exacerbated by the increased availability of bot source code on the underground. Highly prevalent bot families include AgoBot, Phatbot, Rbot, SDBot, SpyBot and others. MyNetWatchman, at http://www.mynetwatchman.com/tp.asp, confirms this bot activity with current statistics from its global aggregate firewall project, which revealed in early 2006 that Sasser/Agobot/GenericBot is the most prevalent port traffic seen over TCP port 445:

PORT	Percentage	Legitimate Services	Potential Maliciousness
TCP/445	29.2	Microsoft SMB/CIFS	Sasser/Agobot/GenericBot
TCP/139	26.0	NETBIOS Session Service	*multiple
TCP/1434	9.1	ms-sql-m	SQL Slammer Worm
TCP/135	8.6	DCE endpoint resolution	MSBlast/Nachi
TCP/1433	5.2	Microsoft SQL	Spida Worm
TCP/80	3.0	HTTP	Nachi/CodeRed/Nimda

The Honeynet Project estimates that as much as 80 percent of all traffic in their research was over TCP ports 445, 139, 135 and UDP port 137. This is largely due to both worm-driven "contagious" bots that spread over these ports and exploits that have emerged against such services in the past three years. Not coincidentally, these are the ports used for resource sharing on various versions of Microsoft's Windows 32-bit operating systems, making it a prime target application for persistent attacks.

Normally, bots attempt to compromise well-known vulnerabilities. Performing regular audits against network resources for these highly targeted services is critical to any well-defended network. The Honeynet Project identifies the following top bot-targeted vulnerabilities:

- 42 - `WINS` (Host Name Server)
- 80 - `www` (vulnerabilities in Internet Information Server 4 / 5 or Apache)
- 903 - NetDevil Backdoor
- 1025 - Microsoft Remote Procedure Call (RPC) service and Windows Messenger port
- 1433 - `ms-sql-s` (Microsoft-SQL-Server)
- 2745 - backdoor of Bagle worm (mass-mailing worm)
- 3127 - backdoor of MyDoom worm (mass-mailing worm)
- 3306 - MySQL UDF Weakness
- 3410 - vulnerability in Optix Pro remote access trojan (Optix Backdoor)
- 5000 - `upnp` (Universal Plug and Play: MS01-059 - Unchecked Buffer in Universal Plug and Play can Lead to System Compromise)
- 6129 - `dameware` (Dameware Remote Admin - DameWare Mini Remote Control Client Agent Service Pre-Authentication Buffer Overflow Vulnerability)

5 Defense Against DDoS Attacks

The utility of different defenses against DDoS attacks vary greatly depending on the size of the target organization, the resources available and the strength of the attack. Ultimately, many defenses, once employed, can later be circumvented by the attacker. It has always been and remains easier for an attacker to adapt their attack vectors or simply increase the number of attacking bots than it is for the defenders to mitigate the attack, to increase resources, or to recover.

In this section, iDefense will discuss only the more modern DDoS attack components, including DoS tools such as Trin00, Tribe Flood Network, Stacheldraht, TFN2k, Shaft, Tinity and other more automated and highly destructive attacks carried out by large groups of zombie PCs or botnets. These botnets can include as few as several hundred zombies or as many as 400,000 or more hosts that carry out attacks through some sort of an automated malicious code or an IRC-based command-and-control structure or even the more advanced web-based control utilities.

Approaches to defending against DDoS attacks mostly fall into the categories of "prevention" (i.e., measures intended to make DDoS attacks impossible) and "mitigation" (or appropriate measures taken to successfully detect and react to an actual attack).

The use of hardware, consultants, outsourcing and other resources all offer avenues to aid in preventing and mitigating DDoS attacks. The length of a sustained attack can be a major variable and factor to consider when examining possible mitigation scenarios. The longer the attack, the more costly it can become to muster the resources to sustain an active defense. This can lead to costs in excess of the value of the service availability.

5.1 Case Study: DDoS Attack against US Financial Services Firm

In early 2004, a notable US financial company suffered a sustained DDoS attack caused by a worm that subsequently remained a significant issue for almost a year after the initial incident. iDefense conducted an interview with the company about the attack and how it was successfully countered. Details regarding the name of the code, the company and the attacker have been omitted from this report.

The attacks began in March 2004, when the company recorded levels of 2,300-2,600 server requests from the worm, which peaked over the summer to levels of 3,000-4,000. The attack died down after that, although occasional spikes in the hundreds were recorded. At the attack's peak, the attack consumed about 5-6 MB of the company's bandwidth, which was worrisome but nowhere near the company's limit.

The company used a Syslog appliance to manage firewall Syslog messages. A trigger was set to notify engineers of a possible DDoS attack at a level of 1,000. A watermark was set at about 50,000 requests within an 8-12 hour period to identify a potential DDoS attack. A "black list" was then implemented during rates of excessive traffic to successfully mitigate all attacks to date. About six legitimate users had been identified on this list and were removed following verification.

The company also initiated redirect implementation. A number of other, more sophisticated solutions were considered, but discarded as unnecessary. The greatest impact the worm had on the company was reportedly the labor hours spent to study and combat it. Although the attack continues (since some of the worms in question are still in the wild), the bandwidth consumed by it is now minimal.

The company also contacted a variety of law enforcement agencies, working primarily with the New York City Cyber Crime Squad, who eventually confiscated computers in Canada and the US. Interestingly, the company's investigators determined that the worm initially spread from computers in the Middle East and South America, although they were never able to trace it to a specific individual.

A deep analysis of multiple variants of code used in the attack traced it back to a known malicious code author with a long hacktivist history. He is a Muslim who claims to be part of the Al Qaeda network. His motive may have been to target the economics of the US. Extensive research into the individual and his code, means and motives proves that he is not a terrorist but a sympathizer motivated by his religious and political ideals. His coding skills are not considered advanced but he is heavily networked within various underground communities.

While this actor's worms were not very successful in spreading, the DDoS attack did gain ground over a period of time and proved troublesome for the targeted organization. Eventually anti-virus software updates helped to curb prevalence of the worms and the company was able to identify specific packet filtering and other strategies useful in mitigating surges in DDoS activity. In the end, the company suffered the greatest expense in terms of labor hours required to assess and respond to the persistent attacks.

5.2 Internal Approaches to DDoS Mitigation

Internal approaches to prevent and mitigate DDoS attacks rely heavily upon well-trained, experienced staff members and potentially expensive appliances specially crafted to deal with DDoS situations.

This approach can lead to expensive purchases and requirements for additional personnel that may not be economical due to the variability of attack schedules and changing demands of attackers. The internal approach to defense relies heavily upon internal teams, resources and the regular reexamination of purchasing products, technologies, additional bandwidth purchasing adjustments and re-engineering network architecture to better handle attacks.

5.2.1 Adjusting Network Architecture and Rules to Mitigate DDoS Attacks

The simplest and most commonly implemented protection plan is to overprovision resources. Ensuring reserve bandwidth is an important component of resource planning due to the substantial increase in traffic caused by a DDoS attack, but planning must also account for the additional processing by the servers and routers. Ensuring that servers have the necessary hardware (network interfaces, processors, memory, etc.) to handle excess capacity if needed, will not only permit the network to tolerate a larger number of bad packets but also benefit end-users with higher daily performance. Although this approach will increase the likelihood of withstanding a moderate DDoS attack, it is not a sufficient defense strategy on its own.

Accurate packet filtering is the most desirable solution to a DDoS attack, but is difficult to attain. Ideally, the network would correctly identify and block all attack packets without impairing normal users. In reality, this is very difficult to achieve and success depends largely upon the location of the packet analysis algorithm within the network and the resources available to perform the analysis without impairing network performance. The closer the analysis is performed to the target node, the higher the accuracy, due to a larger proportion of attack packets to legitimate packets. In addition, the closer the packet identification is done to the target, the fewer the users that will be affected by false positive results. One of the disadvantages to this design is that attack packets continue to pass through the primary pipe, which does not reduce the strain on the rest of the network. If the attack is large enough to

occupy a substantial portion of the available bandwidth, having packet filters after the choke point will provide insufficient protection.

Packet filtering at the point of network entry would have the benefit of reducing unnecessary network traffic and protecting the entire network. Unfortunately, the catch is increased levels of false positive results. As such, the costs and benefits of the two architectures will depend not only upon the specific network configuration but also on the specific circumstances of the attack.

Research is currently being conducted to increase packet analysis efficiency and accuracy with the aim of deploying such systems at central points of the Internet architecture. Although this research goal is admirable, it is likely that mitigation of widely distributed attacks would be best countered with distributed responses. Centralized packet filtering would require an extremely low tolerance for false positives to be practical, suggesting that DDoS defenses will need to remain close to the target node, at least in the near term.

5.2.2 DDoS-Ready ISPs and Over-Provisioning Resources

One commonly taken approach toward mitigating DDoS attacks is to simply maximize bandwidth to be larger than the attack itself. This simple approach requires that the upstream ISP has the bandwidth available and is willing to work with the target to mitigate any attack.

Many ISPs attempt to turn a server off or null-route the victim in the presence of a DDoS attack, usually to avoid collateral damage to other clients.

The following is a list of ISPs that market themselves as knowledgeable and willing to work with clients in "fighting through" DDoS attacks:

Cybercon: *http://www.cybercon.com/*
Based out of St. Louis Missouri, USA, Cybercon boasts great skills when dealing with DDoS attack, but is reportedly expensive.

DDoSProtection: *http://www.ddosprotection.com/*
Human monitoring service through its DDoS Shield product, this service boasts availability of a 100 Mbps mitigating channel.

EV1Servers*: http://www.ev1servers.net/*
EV1 servers utilize a technology called FireSlayer, which is a combination of EV-1 developed and commercially available anti-DDoS technologies to help protect its clients.

GigeServers: *http://www.gigeservers.com/*
Based out of Chicago, IL, GigeServers markets their ProxyShield technology in combination with a staff that claims more than nine years of experience dealing with DDoS attacks.

RackSpace: *http://www.rackspace.com/*
RackSpace claims to be capable in dealing with DDoS attacks and uses a Cisco-powered, Zero-Downtime Network™ that has unique self-healing attributes that allow them to deliver on a 100 percent infrastructure availability guarantee.

Staminus: *http://www.staminus.net/*

Staminus boasts about its dedicated servers with a 99.9 percent network uptime guarantee so you can have peace of mind.

The Planet: *http://www.theplanet.com/*
Based out of Dallas, Texas, it offers 19 Gbps available bandwidth using multiple commercial anti-DDoS technologies from Arbor networks, Cisco and Tipping Point.

AT&T-Internet Protect: *http://www.business.att.com/service_fam_overview.jsp?repoid=ProductSub-Category&repoitem=eb_internet_protect&serv_port=eb_security&serv_fam=eb_internet_protect*
Claims a security alerting and notification service that offers advanced information regarding potential real-time attacks, this service also has a DDoS defense option that allows identification and mitigation within the AT&T backbone.

Broadwing- DDoS Mitigation Service*: http://www.broadwing.com/*
Large global provider with the ability to sell large amounts of bandwidth.

COLT- IP Guardian: *http://www.colt.net/*
Colt is a Pan-European provider of business communications services and solutions. The IP Guardian service utilizes commercial Cisco XT 5650 Guards and Arbor Peakflow monitors to mitigate DDoS attacks.

TELUS- Managed DoS Services:
http://businesscontent.telus.com/webcontent/content/Products/internetData/secureNetworking/managedDDoS.jsp;:Gbie
Telus markets itself as a managed DDoS provider keeping networks up through monitoring and response.

5.3 External Approaches to DDoS Mitigation

Outside experience can provide timely implementation of a robust anti-DDoS strategy by using the experience of DDoS specialists. Combining the knowledge of outside DDoS expertise with internal network expertise, the solution will likely be more robust and will avoid the expensive learning curve of a home-grown anti-DDoS system. In addition to being able to provide guidance in developing an anti-DDoS strategy, many of the following companies offer emergency response services to assist in mitigating an ongoing DDoS attack. Looking to outside consultants may be a viable option if internal expertise is not already available.

5.4 Anti-DDoS Companies and Consultants

5.4.1 Prolexic Technologies

http://www.prolexic.com/
Prolexic is headquartered in Miami, and was formerly called DigiDefense International. Prolexic appears to be the leader in private consulting regarding DDoS-related threats. Members of the Prolexic team have been named in many operations involving successful mitigation of DDoS attacks. Its current CTO Barrett Lyon is fairly well known in the security scene, and once mitigated a particularly high-profile extortion attack, which was covered in an article found at: http://www.prolexic.com/news/20050501-csomagazine.php .

Prolexic is currently waiting for patent approval on what it calls an Intrusion Prevention Network (IPN) that offers "clean pipe" DDoS mitigation services. Additionally, Prolexic offers a four-part DDoS security assessment to prove its worth against other commercial approaches to DDoS protection and mitigation.

Prolexic markets its services as an end-all solution to DDoS attacks, beyond what cookie-cutter manufactured conventional appliances can offer.

5.4.2 Black Lotus

http://www.blacklotus.net/ddos/

While attempting to contact Black Lotus for details into its anti-DDoS technology and services, iDefense received no answer. There is no confirmation on the effectiveness of this solution or indeed its legitimacy.

6 Conclusion

As can be seen, no array of purchasable countermeasures, hardware, software or "solutions" can ensure the prevention of a DDOS attack. Indeed, the vast majority of the mitigation tools listed in this document are of no substance or power; they exist simply because so many undereducated DDoS victims are desperate for help. A determined attacker can launch such an attack against the largest networks anywhere. That said, Prolexic and Arbor Networks are the recognized leaders in the field of DDoS mitigation, although procurement of their services is a considerable expense, ranging in excess of $100,000 dollars for a significant attack on a large company. Indeed, with the advent of DNS recursion-amplification, attackers could conceivably disrupt a victim's Web presence for months, potentially causing some businesses to fail.

Just because a bot herder could attack the largest networks, does not necessarily mean that such an attack will occur. Determined attackers will almost always have some reason driving them. Having such a motive narrows the range of suspected attackers, and therefore increases the chances of identifying them. Bot herders can make thousands of dollars per week extorting relatively modest sums from small companies while incurring little risk of identification. Bot herders know that attacking a big company means conjuring the focused enmity of very wealthy and skilled organizations. The risk is simply not worth it.

Thus, if a bot herder wished to attack a major corporation, it would likely be for some reason other than extortion. This leaves political motivations, youthful mischief and insider revenge. While the revenge motive can be a powerful one, it is also the most easily detected, a condition deterring many such attackers. Youthful mischief is also a possibility, but such attacks are more randomly selected than others, making these unlikely to occur, but quite serious if they do occur. Again, attacker identification is very expensive but is generally easier to do so against adolescent thrill seekers who, despite their considerable skill, often do not have enough experience to know how to cover their tracks. This is illustrated nicely in the case of MafiaBoy, the Canadian hacker who successfully used a DDoS to attack Yahoo!, Amazon and other e-commerce giants in 2000.

Regarding effective mitigation strategies, most large networks are likely well-resourced enough to fend off most DDoS attacks and are not likely to suffer the largest, although DNS recursion-amplification attacks could prove an exception. The utility of any given product or solution will be based in part on its cost. Even using conservative estimates, iDefense sees no technologies that could usefully increase large corporate defenses by more than a marginal degree. However, one cannot rule out the possibility of an attack of hitherto unknown size and sophistication and a company's internal inability mitigate the attack. Thus, iDefense recommends Prolexic and Arbor unless government agencies and insurance companies are available to lend support.

Appendix A - Anti-DDoS Technologies

The following are anti-DDoS technologies that compose the key players and the smaller, not-so-well-known individual companies in the appliance market. DDoS appliances are but one component in the fight against DDoS technology; attackers are constantly performing their own testing in attempts to find ways to break or work around the defensive measures provided by these devices. In some instances, the only way to truly combat large DDoS attacks is through combinations of appliances with professional services and a solid partnership with the ISP.

(Note: Every appliance has a physical limitation on traffic load and could be susceptible to large ongoing DDoS attacks.)

Arbor Networks

Arbor Networks PeakFlow SP: (Perimeter Defense)

http://www.arbornetworks.com/products_sp.php
Peakflow SP is a scalable platform that claims to be able to secure networks from DDoS attacks and worm outbreaks and provide operational reports for things like traffic utilization and routing events from a single device.

The PeakFlow SP appliance:
Detects and mitigates DDoS attacks
Detects and mitigates worm outbreaks
Delivers network-wide traffic topology information and traffic data
Aggregates network traffic
Practical route analytics
Allows reports to be obtained via XML, CSV, XLS or HTML

Arbor Networks PeakFlow X: (Internal Defense)

http://www.arbornetworks.com/products_x.php
The PeakFlow X system utilizes both a collector and a controller appliance. The concept is to establish a baseline of acceptable use within networked devices and resources; if anything falls outside of this, it is flagged. This is a sound approach to monitoring internal network activity.

The PeakFlow X appliance:
Can potentially stop emerging threats
Recognizes appropriate traffic levels in the event of attack
Segments and hardens network resources
Contains automatic update option

Analysis of Arbor Network Technologies

Arbor Network's PeakFlow SP and PeakFlow X are well-known anti-DDoS technologies that are heavily implemented by some of the largest DDoS-ready ISPs. Reviews of these appliances have shown them to perform well in certain circumstances.

Captus Networks

http://www.captusnetworks.com/products/ips_4000.html

Captus IPS 4000

Captus's IPS 4000 claims to recognize DDoS traffic before it reaches the protected network. The appliance attempts to remove packets from the data stream before the routers and firewalls can be disrupted.

Features include:
Threat validation, including screening against multiple policies
Traffic screening that assesses the context of each traffic flow
Controlled responses based on event types and severity
Successive iteration of responses to achieve the desired outcome
Rich policy options for advanced administration
Real-time response capabilities

Analysis of Captus Networks

Captus Networks is a major player within the mitigation of DDoS attack market. Multiple reviews have given the Captus Networks IPS 4000 high ratings, showing that it can perform well in various scenarios.

Cisco Systems

Cisco Guard XT 5650

http://www.cisco.com/en/US/products/ps5894/index.html
The XT 5650 offers multi-gigabit performance to protect against DDoS through per-flow-level attack analysis. The appliance also identifies and mitigates specific attack traffic. For the best defense against DDoS attacks, Cisco recommends that the XT 5650 be combined with the Cisco Traffic Anomaly detector XT.

Features include:
Two versions supporting 10/100/1000BASE-T Ethernet and 1000BASE-SX multimode fiber optic
Processes attack traffic at speeds up to 1 full gigabit per second.
Anomaly recognition
Source verification
Anti-spoofing technology
Cisco "Zombie Killer" technology

Analysis of Cisco Systems

The Cisco Guard anti-DDoS technologies are rated highly by ISPs as being capable in aiding to mitigate DDoS attacks.

Mazu Networks

Mazu Profiler

http://www.mazunetworks.com/solutions/internal/
Behavior-based network security appliance designed to protect internal networks. This system analyzes the behavior of hosts in the network instead of threat signatures to detect threats.

Features include;
Real-time network data
Detailed impact analysis creation
Rogue service termination
Policy enforcement
Policy optimization

Mazu Enforcer
http://www.mazunetworks.com/solutions/perimeter/
Mazu Enforcer is a heuristic-based perimeter appliance that watches for network congestion caused by threats. The appliance then singles out individual packets to filter them. Mazu claims the Enforcer can dynamically adept its filtering behavior throughout the lifecycle of a sustained DDoS attack.

Features include:
Administrator notification capabilities
Packet analysis and tracking
Flood mitigation using datagram protocols
UDP, ICMP and SYN flood mitigation
Fragmentation mitigation

Analysis of Mazu Networks
The Profiler and Enforcer appliances by Mazu Networks appear to be textbook approaches to tackling the DDoS threat.

Minor Players

McAfee
IntruShield: http://www.mcafee.com/us/products/mcafee/network_ips/intrushield_appliances.htm
IntruShield is a network IPS appliance integrating McAfee patented detection techniques with multi-gigabit capabilities.

Cs3 Inc.
MANAnet Shield
http://www.cs3-inc.com/ps_shield.html
MANAnet Router
http://www.cs3-inc.com/ps_router.html
MANAnet Firewall
http://www.cs3-inc.com/ps_fw.html
MANAnet Reverse Firewall
http://www.cs3-inc.com/ps_rfw.html
MANAnet FloodWatcher
http://www.cs3-inc.com/floodwatcher.html
MANAnet Infrastructure level DDos Defense products work in combination to thwart DDoS attacks. The Shield, Router, Firewall, Reverse Firewall and FloodWatcher work together to create an in-line and out-of-line DDoS monitoring protection system both inside and outside the protected network.

F5 Networks Inc.
BIG-IP
http://www.f5.com/products/bigip/
BIG-IP is a fast level 7 switch with built-in DoS protection.

Fortinet
Fortigate series
http://www.fortinet.com/products/
Fortigate is a series of products described as ASIC or accelerated multi-threat security systems. This is Fortinet's approach to the real-time network protection system.

Foundry Networks
Switch BigIron: http://www.foundrynet.com/products/l3backbone/bigiron/index.html
Router NetIron: http://www.foundrynet.com/products/routers/index.html
ServerIron
http://www.foundrynet.com/products/webswitches/serveriron/index.html
Foundry Networks markets the *Iron products as high-bandwidth, fast-switching routing technology that has bandwidth capability to handle DoS attacks.

Juniper Networks
Netscreen firewalls and Routers
http://www.juniper.net/products/integrated/
The Juniper Networks series of Netscreen products offer built-in DoS mitigation functions.

Melicor Inc.
Cyber Warfare Defense Layer
http://www.ddos.com/index.php?content=products/content.html
Melicor's Cyber Warfare Defense Solutions are built upon Barbican technology, which is an in-house technology specifically for DoS mitigation and real-time network protection.

Citrix
NetScaler Application Delivery Systems
http://www.citrix.com/English/ps2/products/product.asp?contentID=21679
NetScalar application delivery solutions offer many features, including attack defense in a single-network appliance.

Radware
Radware Application Security Solutions
http://www.radware.com/content/solutions/application-security/Default.asp
Radware provides integrated intrusion prevention and DoS protection with its APSolute OS and Defense Pro systems, based on behavior- and signature-based technologies.

Tipping Point
TippingPoint Intrusion Prevention Systems
http://www.tippingpoint.com/technology_dos.html
Tipping Point offers DoS protection through its appliance against SYN and established connection floods.

Green Gate Labs
DDoS-Guard Product
http://www.greengatelabs.at/
The DDoS-Guard product claims to be effective against 1,500,000 packets per second.

TopLayer
Attack Mitigator IPS
http://www.toplayer.com/
Attack Mitigater ISP is primarily an Intrusion Protection System with built-in DDoS protection mechanisms.

Webscreen
Webscreen family of Network security products
http://www.webscreen-technology.com/products.html

Webscreen claims its security products are specifically designed and optimized to detect and prevent DDoS attacks through heuristic algorithms that separate malicious from legitimate traffic.

CyberShield Networks Inc.
http://cybershieldnetworks.com/
CyberShield uses patented Intrusion Prevention and Diversion Management (IPDM) technology to fight attacks.

SysMaster
SysMaster Firewall
http://www.sysmaster.com/s_net_dos.htm
The SysMaster Firewall claims to be able to prevent the main types of DDoS attacks and map well-known ports to any port in a NAT.