

Understanding the Modern DDoS Threat

by Gunter Ollmann, VP of Research, Damballa

Introduction

The breadth of cyber threats that an organization must engage with and combat seemingly change on a daily basis. Each new technology, vulnerability or exploit vector results in a new threat that must be protected against. Meanwhile some forms of attack never appear to age - they remain a threat to business continuity despite years of advances in defensive strategy. One particularly insidious and never-ending threat is that of the Distributed Denial of Service (DDoS) attack.

DDoS attacks are the staple disruptive technique preferred by an increasingly broad spectrum of attackers.

Never far from the news headlines, DDoS attacks are the staple disruptive technique preferred by an increasingly broad spectrum of attackers. While they may be the oldest and most commonly encountered form of cyber attack, defenses against them are often non-trivial and even the best tried-and-tested protection can fail under a sufficiently well conceived attack.

This paper **examines the technology, coordination tactics and motivations behind the DDoS attacks** likely to pose a risk to Internet accessible businesses now and in the immediate future.

Denial of Service Concepts

As the name implies, DDoS encompasses the coordinated activities of multiple Denial of Service (DoS) agents and tools. The general concepts of DoS are simple – cause an action upon a computer or networked device which results in other processes, resources or activities floundering and failing to adequately respond.

DoS attacks can take on many forms depending upon the target system and objectives of the attacker. For example, an attacker may deny their victims the ability to log into their computer systems by intentionally supplying multiple incorrect passwords until the application locks the accounts out.

At the other end of the spectrum, an example would be the mid-1990's "ping of death" – in which an attacker sends a specially crafted network packet (in this case an overly large ping packet) – resulting in multiple victim machines crashing and eventually rebooting.

While many DoS attack techniques have their own nuances and specific naming conventions (e.g. ICMP flooding, teardrop attacks, reflected attacks, etc.), in general DoS attack techniques can be grouped into two major categories:

1. **Application** – Known application logic limitations, flaws and vulnerabilities are exploited resulting in a specific application failure or data corruption.
2. **Network** – Vulnerabilities in the way in which networking equipment, Internet protocols and routing are configured are exploited to deny Internet access to/from the victim machine or infrastructure components.

Leaving little to the imagination, DDoS attacks are where multiple participating devices engage in DoS attacks from a distributed assortment of locations. The specific attack traffic or vulnerability exploitation that causes the target(s) to become unresponsive is typically the same for DDoS as it was for singular DoS attacks. The distributed sources engaged in the DDoS attack often make the attack more difficult to defend against and are generally more successful against larger and faster responding targets.




Figure 1: Multiple attackers launch their DoS payloads to constitute a DDoS attack

The Distributed DoS

Compared to the many classes of complex threats routinely encountered by businesses, DDoS tends to be inelegant and one of the most easily recognized attacks. While the threat may be one of the most commonly encountered and discussed, many people are unaware of the fact that the label of “DDoS” encompasses multiple attack techniques – each with their own nuances and effect on the designated target.

In essence there are three main classes of DDoS attack:

1. Bandwidth Consumption

In this category of attack, the DDoS participants attempt to flood network chokepoints with more network traffic than they can reliably consume. The object of this flood is to prevent legitimate network traffic from reaching the targets network infrastructure and prevent that traffic from being successfully routed to the waiting applications and services. This is the simplest DDoS attack to perform as, in its rawest form, it is simply about pumping more network traffic at the target than the target has bandwidth to consume. If the target has a 10Mbps Internet connection, the attacker only needs to direct 10Mbps of DDoS traffic at it – which may only require the output from as few as a dozen bot infected machines.

2. Resource Exhaustion

In this category of attack, the DDoS participants attempt to exhaust system resources of the target; an attack class that typically requires less bandwidth to be successful. Internet services (whether they be web sites, email servers, etc.) are hosted on servers with physical and programmatic resource limitations.

For example, a single web server may be able to cope with 4,000 simultaneous HTTP-based user sessions or 500 HTTPS-based user sessions. Once 4,000 user sessions have been started, no further sessions can be made until some of the earlier sessions have expired. Therefore, an attacker can seek to exhaust the resources of the target in order to deny legitimate users or customers from being able to interact with the Internet services. A

botnet operator could instruct a few hundred botnet agents to each make several dozen simultaneous connections to a web server, seek to keep those connections open for as long as possible, and consequently prevent others from connecting to the web server.

3. Application Exploitation

In this category of attack, the DDoS participants seek to exploit weaknesses or vulnerabilities within the applications actually being served by the hosting infrastructure. To differentiate between the resource exhaustion of the hosting infrastructure and application exhaustion, the attacker seeks logic flaws and errors within the applications running within the targets environment.

For example, an application may accept three failed login attempts before it locks-out an account. The attacker may abuse this lockout logic by automatically submitting 3 bad login attempts to every user account on the system – thereby making the application inaccessible to all users. There may also be vulnerabilities within the application that causes follow-on processes to become stalled or operate inefficiently – for example, a search feature within the application where by requesting a search of lots of SQL join statements causes the applications backend database to stop answering all search requests for several minutes.




Figure: “Low Orbit Ion Cannon” DDoS tool. The tool is downloaded by the attacker and may be used for individual target attacks or it may receive centralized CnC attack instructions.

DDoS Botnet Business

While there are literally hundreds of public tools available for tinkerers and potential attackers to launch small-scale DDoS attacks against their targets, the preferred route for most professionally orchestrated and managed DDoS attacks is through the use of botnets. DDoS botnets can come in a variety of sizes – ranging from a small cluster of a half-dozen zombie machines, through to a globe-spanning army of a million-plus slaves – and access to them can be acquired for a nominal fee (should the attacker not already own and operate their own botnet).

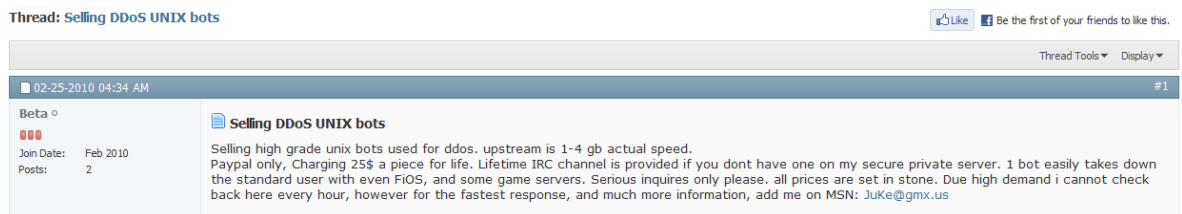


Figure: Typical DDoS botnet for sale

The prices to rent, sub-lease or purchase a DDoS botnet vary considerably. Professional botnet building and operations are a competitive business with rates fluctuating based upon demand and the type of DDoS attack needed. In general, DDoS rental fees typically start at around \$50 for day small attacks (a few hundred machines with a few Gbps attack output) – but daily fees can drop to only a few dollars if the timing is optimal.

Commercial DDoS operators – such as the IMDDOS botnet previously documented by Damballa – are plentiful and usually provision their services from former Eastern-block countries and Asia. Pricing can vary from \$5-\$10 per hour for non-distributed DoS attack services, and \$20-\$100 per hour for high volume DDoS attacks. A 24-hour DDoS attack can reach several thousand dollars for a large attack against a web portal (with the largest observed attack being in excess of 49 Gbps in 2009), while the cost to send a million spam emails is regularly in the order of \$40-\$150.

Many DDoS service providers also offer “try before you buy” services – where a free 3-5 minute trial DDoS is launched against a specified target as proof of the botnets capacity (and the owners credibility).

While the traffic generated by a DDoS botnet is the cause for bringing down the target (referred to as “Drop” by attackers – as in “to drop the website”), what really happens is that various equipment within the targets infrastructure fails or forces upstream providers to “null route” the targets IP address at the network level – with the end result being the IP address of the target getting dropped from Internet routing tables.

Understanding the Attackers

Throughout the years, while DDoS tactics have remained largely unchanged beyond tweaks for specific target and delivery platforms, the motivations of the attackers have changed considerably. Originally, DDoS attacks could have best been characterized as opportunistic displays of displeasure by their instigators but, by the late 1990’s, as businesses found new ways of making money over the Internet, DDoS began to prove itself as a valuable component of an extortionists toolkit. A decade later, the categories of people involved in launching, managing and participating in DDoS attacks have become much broader – and their motivations more diverse than ever.

The people behind modern DDoS attacks can select from a huge pool of tools and methodologies. But, in order to efficiently respond and withstand an attack, security professionals and network administrators should seek to understand who the attackers are and their objectives. Despite the diversity of potential attacker profiles and objectives, things can be distilled down to three major categories of attacker:

1. **Professional:** Professional DDoS operators are the career criminals who tend to have adopted business models that generate a financial profit from conducting such attacks. The tool of choice for modern professional DDoS operators are botnets – and will regularly have at their disposal a handful of independent botnets, with each botnet consisting of a few thousand or tens-of-thousands of compromised computers. Some of the biggest botnets – consisting of hundreds-of-thousands of computers – are operated by professional DDoS criminals.

While the botnet(s) may consist of hundreds-of-thousands of compromised computers, the criminal operator does not necessarily employ the full power of the entire botnet in a single attack. There is typically no need to do so. The same botnet may be conducting multiple DDoS attack campaigns simultaneously.

The two most common vehicles for monetizing professionally organized DDoS attacks are subleasing and extortion.

- Gamerz:** “Gamerz” are a generic category of “misguided” or malicious individuals that often have a higher-than-average technical understanding and who employ DDoS tools and tactics for personal gain. This category of attacker often builds or operates dedicated DoS tools and uses them to gain an upper hand in online games – causing fellow gamers and Internet opponents to slowdown or drop their network connectivity – allowing the DoS instigator to win. In many cases Gamerz will utilize small botnets (of a few tens or hundreds of compromised computers) that they personally build and operate, and will task them with DDoSing specific application servers or individual gamer IP addresses.

It is not uncommon for multiple Gamerz to participate in small peer-based groups and collectively apply their DDoS capabilities against shared and pre-agreed targets. These kinds of collective attacks (beyond disruption of service availability to online games) are often characterized by the focused disruption of high visibility targets and the public release of messages loosely explaining why the particular target was selected. There is often a high degree of peer recognition and one-upmanship involved in the Gamerz participation decisions.

- Opt-In:** “Opt-in” attackers can often be thought of in terms of participants in mainstream protest movements. Individuals from diverse geographic regions are rallied together using various social networking platforms and media, guided in the acquisition and use of various DDoS tools, and provided with target lists from a coordinating source. In a growing number of cases the DDoS tools are in fact customized DDoS botnet agents – with a centralized CnC governing the list of targets that will be subject to the attack.

Opt-in botnets and their association with social networks, hacktivism and centrally-controlled protesting are covered extensively in the Damballa whitepaper “The Opt-in Botnet Generation”.



Figure: Online tutorial covering the deployment and use of DDoS tools for launching attacks against other online gamers.

DDoS Attack Objectives

For many organizations situated at the receiving-end of a coordinated DDoS attack, the objectives of the attackers may not necessarily be self-evident. In too many cases, organizations come under the cross-hairs of multiple DDoS campaigns – with the campaigns being launched and orchestrated by different groups.

From an instigators perspective, some of the most common DDoS objectives are:

1. **Extortion.** The attackers seek to cause key online business services to become unavailable at critical times and expect payment for an attack to cease. For example:
 - a. Preventing customers from placing bets within an online gambling portal in the run up to a major sporting event and requiring payment to prevent a DDoS occurring on game day.
 - b. VoIP calls are made continuously to an organization's phone and fax numbers thereby preventing any in-bound communications. Automatic messages are played to anyone answering the phone that payment is expected for the calling to cease.
2. **Espionage.** The attackers seek to cause key business services to become unavailable or unresponsive while reaping an award on another front. The DDoS attack itself is used as a method of disguising the real purpose of the attack or distracting the victim's attention. For example:
 - a. The operator of a botnet is paid to DDoS the email services of a local business operator to prevent that organization in responding to a competitive business bid.
 - b. A high-volume DDoS assault is conducted against multiple online business portals with the expectation that the targets incident response team will be too busy to notice a parallel stealthy attack against the true target as things are "lost in the noise".
3. **Protesting.** The attackers seek attention to the particular cause or public issue they are pursuing and work to force a particular change in policy or behavior. Attack participants are provided with target and coordination details from a central "authority". For example:
 - a. The global DDoS of a particular government's web sites in response to (perceived) unfair election practices.
 - b. A coordinated campaign of DDoS attacks against the web portals and email systems of any organization supplying materials to businesses and laboratories that conduct animal testing trials.
4. **Nuisance.** The attackers launch attacks against a broad spectrum of targets "because they can". The objectives vary greatly between targets but the DDoS attacks are typically short-lived, often reactionary to a perceived slight, or designed to gain some temporary advantage over named individuals. For example:
 - a. Opposition team member IP addresses are DDoSed during an online game so that the attacker's team can win or obtain the highest scores.
 - b. A student launches a DDoS against the schools homework submission system in an effort to cause other students to miss a specific homework due date.

Getting Inside the Attackers Head

One of the best ways to understand the factors and the dynamics behind the threat (and ultimately be able to respond to the threat) is to think like the attacker. The following are typical questions that an attacker will have worked through in order to deliver a successful DDoS attack:

- **Who and what is the target?**

This is perhaps the most important question and defines much of the attack methodology. The tactics used for taking down the network connection of large financial institutions are different to the tactics employed in causing opponents in an online poker tournament to lose their connections to a game.

In general, smaller targets can be taken down with less sophisticated tactics targeting specific servers – while large distributed corporate applications and portals may be attacked through application-level vulnerabilities and network architecture bottlenecks.

- **Where are the weakest components of the target?**

The larger and more sophisticated the target, the more reconnaissance is needed by the attacker. While a corporate mail server may be able to handle several thousand emails per second, it is likely hosted within the corporate offices and connected to the Internet via a leased line of only a few tens-of-megabytes of capacity.

An attacker bringing to bear a hundred or so botnet victims would likely saturate that corporate Internet connection – thereby preventing all emails (and other corporate communications) from being sent or received by the target. For attacks against organizations that have large Web applications hosted in cloud environments and able to withstand tens-of-Gigabytes of DDoS traffic, it is likely that DNS is hosted by a different service provider and may be more vulnerable to attack.

- **What are the best tools to generate the attack traffic?**

Depending upon the type of Internet service destined to be targeted by the DDoS attack, there are an extensive array of tools available for use – many of which are optimized for certain attack traffic profiles (most commonly HTTP and SMTP). There are hundreds of readily-obtainable DDoS tools that can be pre-configured for a specific type of traffic and array of targets, themed for the cause, and packaged for download by opt-in protesters.

Other tools may be optimized to make use of network and application configuration vulnerabilities that can amplify an attack – for example, the capability of sending a single data packet that will elicit a return response from hundreds of computers, and have all those responses automatically directed at the target. Meanwhile botnets offer great flexibility; the malware agents installed on botnet victims are sophisticated enough to be able to generate or replay any kind of network traffic that attacker wishes.

- **How many resources should be rallied and coordinated for the attack?**

Unless an attacker already has a stable of DDoS agents of sufficient size prepared for the attack, they will have to invest in infrastructure components to rally and command the hundreds or thousands of new members they acquire each day. Not all DDoS participants will be accessible or controllable at the same time. Just because a botnet has 100,000 members, it does not mean that the botnet controller can instantly launch a DDoS attack against the target. A large botnet can bring to bear tremendous volumes of attack traffic but, like comparing a Jet Ski to an Oil Tanker, increases in size comes with corresponding sacrifices in agility.

- **Does detection matter?**

Detection means a number of things to the attacker - Will the target be able to detect an attack? Will the target be able to enumerate (and consequently takedown or block) all of the

systems participating in the attack? Will the target (and law enforcement) be able to detect the instigator and controller of the attack? With sufficient preparation an attacker can evade many forms of detection – should they choose to do so. Some DDoS attacks, for example those undertaken by opt-in participants protesting a particular political issue, specifically want the attacks to be detected and do not care whether individual participants are identified. Meanwhile professional operators who rent or sub-lease parts of their botnets to subscribers for DDoS will go to great lengths to prevent detection of the command controllers and ultimately their involvement.

Mitigation Techniques

Depending upon the nature and likely duration of the DDoS attack, there are various mitigation techniques available to the targeted organization. A core component in withstanding a DDoS attack lies with understanding the type of attack being launched and the objectives of the attacker. Armed with that information, an optimal DDoS mitigation strategy can be devised.

Based upon the class of DDoS attack, the following mitigation strategies may be appropriate:

1. Bandwidth Consumption

- a. Bandwidth over-provision. For low bandwidth DDoS attacks it may be economical to simply over-provision bandwidth to the targeted systems. This can be achieved by working directly with the ISP or hosting provider and to moving to a higher tier of service offering.
- b. Black hole routing. When an attack is detected, the malicious traffic can be re-routed with the aid of the ISP(s). With “source-based” black hole routing, a null route is promoted causing all traffic from a specific list of IP sources to be dropped. This can be effective if the number of DDoS participants is small. With “destination-based” black hole routing, the targeted system(s) IP address(s) are null routed – effectively taking down the victim while other safeguards are actioned.
- c. Distributed hosting. If the attackers are distributed, then it may pay dividends to similarly distribute the targeted system(s) via multiple hosting facilities or cloud provisioning services. By distributing the hosting (and effectively creating multiple copies of the service in different geographic locations), it becomes considerably more difficult for an attacker to overwhelm all facilities simultaneously. There may also be cost reduction implications as well – for example, the hourly cost of mitigating a 10 Gbps DDoS attack against a single site may be considerably more than dealing with ten 1 Gbps DDoS attacks against separate hosting facilities.

2. Resource Exhaustion

- a. Server patch management. It is important that all Internet accessible services are correctly patched with the latest vendor software. All systems should be hardened (with unneeded or superfluous services removed rather than simply disabled) and configured to close idle connections.
- b. Rate and connection limiting. Temporarily apply rate limiting techniques to inbound traffic. Round-robin allocation of system resources based upon source IP address can help ensure that legitimate users of the system can still gain access and perform functions – albeit at reduced speed.
- c. Connection aging. When idle connections fill up connection tables within firewalls and application servers, aggressive aging of the connections can help flush the systems quicker and make resources available for new connections.
- d. Load balancing and prioritization. Load balancers and many security appliances have the capacity to balance and prioritize inbound traffic at the edge. These technologies can reduce the volume of attack traffic being propagated between devices within the targeted environment.

3. **Application Exploitation**
 - a. Secure application development. The inclusion of secure development practices into the development lifecycle can go a long way in mitigating potential vulnerabilities that could be exploited by remote attackers. Development teams should also run regular “ethical hacking” exercises and security assessments against the online applications to identify new flaws and verify that the application is in fact secure.
 - b. Apply application controls. When dealing with DDoS attacks designed to mimic regular Web traffic, a combination of application-level controls and anomaly detection techniques are necessary from within the custom Web application or portal.
 - Filtering and blocking of known “bad actors” (e.g. IP addresses, user ID’s, open proxies, etc.).
 - Threshold alerting and subsequent dropping of single source/sessions with excessive request volumes or non-human timings between data submissions and page navigation.
 - Inclusion of “speed bump” tactics such as CAPTCHA’s designed to slow down automated attacks.
 - c. Traffic filtering. High capacity edge filtering technologies (such as Intrusion Prevention Systems (IPS) and Web Application Firewalls (WAF)) can be used to prevent known malicious traffic from propagating to the sensitive web application. These technologies should be the first guard in filtering out unwanted attack traffic and reducing the volume of inbound requests to the application.

Participation Mitigation

In some cases your organization may not be the targeted victim of a DDoS attack, but may in fact be a participant in the attack. Large organizations are frequently compromised with botnet malware and the corporation’s assets may find themselves under the control of a botnet operator – and subsequently be used to launch an attack against another organization or target of the criminals choosing.

There are two primary mitigation techniques:

1. **Attack traffic detection.** Using IPS, Anomaly Detection Systems (ADS) or firewall logging, it may be possible to detect the attack traffic being generated by the compromised hosts. Hosts that have been identified as participating in a DDoS attack could then be individually blocked or shutdown pending full manual remediation.
2. **Command and control detection.** Using network based sensors, malware infected hosts that attempt to communicate with known or suspected CnC services can be enumerated as belonging to a particular botnet.

If CnC detection techniques are being used to identify and track botnet membership within the corporation (or ISP), there are several additional actions that can be performed to mitigate the threat.

1. As botnet members begin to participate in the DDoS attack, traffic to and from the CnC server could be blocked. By doing so, no new commands are sent to the botnet victims and they typically cease their attacks. In addition, any other botnet members within the network who have not yet been tasked to participate in the attack will similarly not be able to receive instructions.
2. Walled Gardens can be selectively initiated around the infected botnet population – blocking just the ports and protocols being used (or likely to be used) in the attack against remote targets – without applying the same blocking to all hosts or subscribers within the network. For example, a botnet may be tasked with DDoSing a popular financial services web portal using a HTTP-based payload. It would therefore be important to only block the attack traffic and allow legitimate traffic through. A walled garden approach could be used in this

- scenario without having to utilize Deep Packet Inspection (DPI) to differentiate between the attack and legitimate traffic.
3. The ability to differentiate CnC server activity at the domain name level is important for botnets that utilize fast flux infrastructure to distribute command over large numbers of IP addresses. If recursive DNS services are provided by the organization to their enterprise hosts or subscribers, an alternative DNS response could be sent to the botnet victims – e.g. making botnet.badness.com.cc resolve to localhost (127.0.0.1).
 4. If DPI or PCAP capabilities exist within the organization, they could be selectively deployed to catalog the criminal communications between the botnet members and the CnC server. This detailed evidence of the attack (including the commands being sent by the CnC) can be used for takedown or prosecution purposes.
 5. If the botnet malware agent is relatively unsophisticated or if the CnC server itself is vulnerable to third-party takeover (e.g. a hacked server that the legitimate owner regains control and can now issue commands to the botnet, or if the Botnet CnC portal code contains remotely exploitable vulnerabilities), it may be possible to issue commands “on behalf” of the criminal operator instructing all the botnet members to stop their attack and to automatically uninstall the malware agent.

Conclusions

DDoS tactics and mitigation strategies will continue to develop as the threat evolves and the motivations that lead to an attack similarly change. By understanding an attacker’s motivations and the tools they have at their disposal, it becomes easier to formulate mitigation strategies that are both efficient and effective.

As new platforms get adopted and become prevalent, they will become both tools and victims to DDoS attacks. Already mobile communication platforms are being subjected to new forms of DDoS and the attackers are uncovering new ways to utilize mobile devices for new classes of attack. The compromise of Smartphone platforms opens the door to new personalized DDoS strategies (for example, the targeting of an organization’s executive team and making their mobile systems inoperable) as well as new victims (e.g. several thousand infected handsets all trying to call 911 emergency services at the same time for sustained periods).

Further Reading

“Ping of death” - <http://insecure.org/splotts/ping-o-death.html>

“The Opt-In Botnet Generation - Social Networks, Cyber Attacks, Hacktivism and Centrally-Controlled Protesting”, Damballa, http://www.damballa.com/downloads/r_pubs/Opt-In_Botnets.pdf

“Continuing business with malware infected customers”, Gunter Ollmann, 2008, <http://www.technicalinfo.net/papers/MalwareInfectedCustomers.html>

About Damballa, Inc.

Damballa is a pioneer in the fight against cybercrime. Damballa provides the only network security solution that detects the remote control communication that criminals use to breach networks to steal personal and intellectual information, and conduct espionage or other fraudulent transactions. Patent-pending solutions from Damballa are platform and system-agnostic, protecting networks with any type of device including PCs, Macs, smart phones, as well as mobile and embedded systems. Damballa customers include Fortune 1000 companies, Internet and telecommunications service providers, government agencies and educational organizations. <http://www.damballa.com>

Copyright © 2011, Damballa, Inc. All rights reserved worldwide