

Barbarians at the Gate: An Introduction to Distributed Denial of Service Attacks

Matthew Tanase 2002-12-03

Introduction

Recently, major news outlets reported that a [coordinated attack](#) designed to disable several of the Internet's root name servers had taken place. The attack, described as sophisticated and complex, is known as a distributed denial of service (DDoS). Although no serious outages occurred, it was a hot topic in the security world - again. Again? [Similar attacks](#) first made headlines in February 2000. Although discussed in security circles for some time before that, this was the first prolonged example of a DDoS, and prevented legitimate traffic from reaching major sites for several hours. Yahoo, eBay, Buy.com, and CNN were but a few major sites who were inaccessible to their customers for extended periods of time. Now, almost three years later, can it be that we're still vulnerable? Unfortunately the answer is yes. This article will explain the concept of DDoS attacks, how they work, how to react if you become a target, and how the security community can work together to prevent them.

What is a Denial of Service?

In order to understand the incidents described above, it would be helpful to take a step back and look at a more basic form of the same attack, the denial of service attack. A denial of service, or DoS, is a very basic category of attack in the world of security engineering, one which can be used in several scenarios. The term can be applied to any situation where an attacker attempts to prevent the use or delivery of a valued resource to its intended audience or customer. It can be implemented via multiple methods, physically and digitally. For instance, an attacker can deny access to telephone systems by cutting the major telecom cable feeding a building, repeatedly calling every available phone line, or cracking the switch that handles the PBX. In all three instances, the attacker succeeds by denying the users access to the resource, as all incoming and outgoing calls would fail.

The DoS concept is easily applied to the networked world. Routers and servers can handle a finite amount of traffic at any given time based on factors such as hardware performance, memory and bandwidth. If this limit or rate is surpassed, new requests will be rejected. As a result, legitimate traffic will be ignored and the object's users will be denied access. So, an attacker who wishes to disrupt a specific service or device can do so by simply overwhelming

the target with packets designed to consume all available resources.

A DoS is not a traditional "crack", in which the goal of the attacker is to gain unauthorized privileged access, but it can be just as malicious. The point of DoS is disruption and inconvenience. Success is measured by how long the chaos lasts. When turned against crucial targets, such as root DNS servers, the attacks can be very serious in nature. DoS threats are often among the first topics that come up when discussing the concept of information warfare. They are simple to set up, difficult to stop, and very efficient.

A Distributed Denial of Service?

This article is concerned with a specific type of DoS, which implements a coordinated attack from multiple sources. Known as a distributed denial of service, or DDoS, it's easily executed on a large network and can be frighteningly effective. A DDoS can be thought of as an advanced form of a traditional DoS attack. Instead of one attacker flooding a target with traffic, numerous machines are used in a "master-slave", multi-tiered configuration.

The process is relatively simple. A cracker breaks into a large number of Internet-connected computers (often using automated software known as an [autorooter](#)) and installs the DDoS software package (of which there are several variations). The DDoS software allows the attacker to remotely control the compromised computer, thereby making it a "slave". From a "master" device, the cracker can inform the slaves of a target and direct the attack. Thousands of machines can be controlled from a single point of contact. Start time, stop time, target address and attack type can all be communicated to slave computers from the master machine via the Internet. When used for one purpose, a single machine can generate several megabytes of traffic. Several hundred machines can generate gigabytes of traffic. With this in mind, it's easy to see how devastating this sudden flood of activity can be for virtually any target.

The network exploit techniques vary. With enough machines participating, any type of attack will be effective: ICMP requests can be directed toward a broadcast address (Smurf attacks), bogus HTTP requests, fragmented packets, or random traffic. The target will eventually become so overwhelmed that it crashes or the quality of service will be worthless. It can be directed at any networked device: routers (effectively targeting an entire network), servers (Web, mail, DNS) or specific machines (firewalls, IDS).

But what makes a DDoS difficult to deal with? Obviously the sudden, rapid flood of traffic will

catch the eye of any competent administrator (if the phone ringing and pager beeping doesn't!). Unfortunately though, all of this traffic will likely be spoofed, an attack technique in which the true source address is hidden. An inspection of these packets will yield little information other than the router that sent it (your upstream router). This means there isn't an obvious rule that will allow the firewall to protect against the attack, as the traffic often appears legitimate and can come from anywhere.

So what's left to do? Not much, other than to start an extremely frustrating process: the DDoS investigation. With each step up the chain of routers that handled the malicious traffic prior to your network is a new set of administrative contacts: more phone calls must be made, panic emails sent, and packet captures analyzed. It's very time consuming, which is amplified by the fact that the network or machine is currently down. Given the fact that the slaves can be located anywhere in the world, the sad truth is that the DDoS flood more often ends due to the attacker's whim than to any action taken by the targeted system's administrator.

Surviving DDoS Attacks

That said, there are steps that can be taken to mitigate the effects of a DDoS attack. As mentioned in the previous section, the first thing to start is the investigative process. Determine which core router (a router that handles Internet backbone traffic) is passing the packets to your border router (a router that connects your network to the Internet). Contact the owners of the core router, likely a telecom company or the ISP, and inform them of your problem. Ideally, there will be a process in place which can expedite your requests for help. They, in turn, need to determine where the malicious traffic reaches their network and contact the source. By that point, it's out of your hands. So what can be done in the meantime?

Since it's not likely that you'll be able to quickly stop the DDoS flood, there are a few steps which might help mitigate the attack temporarily. If the target is a single machine - a simple IP address change can end the flood. The new address can be updated on internal DNS servers and given to a few crucial external users. It's not an elegant solution, but a quick one which works. This is especially useful for key servers (i.e. mail or database) under attack on your network.

There is a chance that some filtering techniques can help. If the attack is unsophisticated, there might be a specific signature to the traffic. A careful examination of captured packets sometimes reveals a trait on which you can base either router ACLs (access control lists) or

firewall rules. Additionally, a large amount of traffic may originate from a specific provider or core router. If that's the case, you might consider temporarily blocking all traffic from that source, which should allow a portion of legitimate activity through. Keep in mind, however, that you'll also be blocking "real" packets, or legitimate traffic, but this may be an unavoidable sacrifice.

A final option, one which might be available to larger companies and networks, is to throw more hardware or bandwidth at the flood and wait it out. Again, it's not the best solution, nor the least expensive one, it may provide a temporary fix nevertheless.

It's important to stress that the investigative process should begin immediately. Without a doubt, there will be multiple phone calls, call backs, emails, pages and faxes between your organization, your provider and others involved. It's a time consuming process, so get the ball rolling. It's taken some very large networks with plenty of resources several hours to halt a DDoS, so plan accordingly.

Preventing DDoS Attacks

The DDoS problem can only be remedied by a community effort and stricter security standards. First, administrators and home users alike need to make sure their machines are secure. The slaves used in DDoS attacks are often the product of autorooters, programs which scan thousands of machines, crack vulnerable ones and install software. Keeping patches up to date, closing open services, and implementing basic firewall filtering can help keep your machines from falling prey and participating in such an attack.

The major difficulty in defeating a DDoS lies in the spoofed IP addresses of the attackers. This problem can be solved using a technique called ingress filtering on routers. Ingress filtering inspects packets destined for the Internet at the border router, one hop prior to the core router. These routers should know the address of every device behind them; therefore, anything outside of this range is spoofed. Spoofed packets should be dropped **before** they reach the Internet backbone (or core router). If they reach that point, it's too late. If network administrators implemented such filtering by default, spoofing a packet would become nearly impossible, eliminating the timely identification process in the DDoS investigation. Unfortunately, most networks do not have these crucial filters in place, and spoofed packets abound. IPv6, which will be deployed in the future, also has security features in place that address this fundamental networking problem.

The community also recognizes the difficulty of reaching the proper technical contacts on neighboring networks and is actively working on a solution (see Bugtraq). You should have in place a list of administrative and technical contacts at your ISP. Additionally, determine if they have a procedure in place for identifying and dealing with DDoS attacks on their own backbone network. Some of the major providers have sensors in place that can identify sudden increases in traffic at certain points, which serves as a useful alarm for discovering and isolating major DDoS incidents. If you're currently shopping for an access provider, ask them about dealing with DoS attacks. If you already have a provider, ask the same question. The response should determine if you need to be shopping for a new one.

Conclusion

The distributed denial of service is a very effective attack, one that is difficult to stop. The ultimate solution will require a vigilant networking community that enforces strict standards. Currently, the best defense techniques lie in anticipating such an attack. Having a DDoS incident response plan in place is crucial. And the use of ingress filtering and tight security standards should ensure that a machine under your control does not contribute to the problem. An active, aware, informed community can make the DDoS headlines of today a relic of the past.

Matthew Tanase CISSP, is President of [Qaddisin](#), a network security company based in St. Louis. His company provides nationwide consulting services for several organizations. Additionally, he produces [The Security Blog](#), a daily weblog dedicated to security.

Relevant Links

[Distributed Denial of Service Attack Tools \(PDF\)](#)

A good overview of several common DDoS attack tools.

[An Introduction to Autorooters](#)

An article on autorooters (often used to crack DDoS slaves) and how to defend against them.

[Always On, Always Vulnerable: Securing Broadband Connections](#)

An article how SOHO users can protect their networks.

[Bugtraq](#)

A security mailing list which is working on contact lists for dealing with DDoS attacks.

[GRC Denial of Service Attacks](#)

Steve Gibson's informative account of the DDoS attacks on his site.

[David Dittrich DDoS Page](#)

DDoS and security researcher David Dittrich's page of links and notes.

[Privacy Statement](#)

Copyright 2006, SecurityFocus