Network Bandwidth Denial of Service (DoS)

Angelos D. Keromytis
Department of Computer Science
Columbia University

Synonyms

Network flooding attack, packet flooding attack, network DoS

Related Concepts and Keywords

SYN flood attack, application-level DoS, algorithmic DoS, memory and state exhaustion DoS, distributed DoS (DDoS)

Definition

Network bandwidth denial of service (DoS) attacks seek to consume the available bandwidth or router resources at or near a target host or network, such that legitimate traffic cannot reach its destination. The primary means for achieving this goal by sending large traffic volumes (packet floods) that do not respect congestion control signals, such as that in the Transmission Control Protocol (TCP) or Explicit Congestion Notification (ECN). In wireless networks, such attacks can also be carried out through radio jamming.

Background

Network bandwidth DoS attacks have been seen on the Internet since at least 1996, with a TCP SYN flood attack against Panix, an Internet Service Provider. Large-scale DoS attacks against some high-visibility sites including Ebay, Yahoo, CNN.com and E*trade were widely reported in the news in February 2000. Since then, network DoS attacks are a common phenomenon, with as many as 700-800 observed in a single day. The effectiveness of such attacks has increased with the advent of large-scale botnets, compromised hosts that are under the control of a single entity. While early attacks appear to have been launched primarily for amateurish reasons, more recently DoS attacks have been used for financial crime (extortion, corporate warfare) and political purposes. Notable cases of the latter include the DoS attacks against government websites in Estonia (2007), Georgia (2008), Iran (2009), South Korea and USA (2009), against the Cloud 9 ISP in the UK (2002), and against a specific user on Twitter (2009). In addition to their targets, such attacks cause significant collateral damage, by affecting the available bandwidth and router resources in links close, but otherwise unrelated to the target site.

Application

Bandwidth DoS attacks target the availability of networked services by preventing legitimate users from successfully communicating with the target of the attack. The easiest way of achieving this goal involves the transmission of large volumes of traffic toward the destination. Packet-switched networks such as the Internet use statistical multiplexing to allocate limited resources to large numbers of users. In addition, the Internet is an "open" network: any host can attach and transmit traffic, as long as it obeys some very basic rules with respect to packet format. Because of these two design properties of the Internet, bandwidth DoS attacks can saturate network links and consume scarce router resources, especially at the edges of the network (*i.e.,* topologically near end-hosts and networks).

DoS attack traffic generally does not respect congestion control signaling, such as TCP's built-in exponential back-off algorithm or ECN packet markings. User Datagram Protocol (UDP) and Internet Control Message Protocol (ICMP) traffic, due to its inherent statelessness and lack of congestion control in the protocol, is extensively used in many such attacks. However, legitimate TCP traffic may also be used, especially when large numbers of hosts can be made to initiate connections. Several attacks also involve fabricated TCP traffic, *i.e.,* TCP packets that are not part of any actual connection but are generated by attack tools for the purposes of evading detection.

The volume of traffic that is necessary to successfully cause a DoS attack depends primarily on the link capacity of the target site. For well-provisioned sites, a single attacking host may not be able of generating enough attack traffic volume. With the prevalence of botnets in recent years, large-scale DoS attacks involving many thousands of attacking hosts have become common. In most cases, such DDoS attacks aggregate enough traffic to saturate the network links of their targets. Even when a DoS attack does not consume enough bandwidth to saturate its target or upstream links, quality of service may be significantly affected for legitimate clients. Depending on the type of traffic that such clients generate, the effect of an "incomplete" DoS attack may be equivalent to a fully saturating one.

The content of the attack traffic is not relevant to the attack itself. However, recent attack tools seek to create traffic that could have been generated by legitimate clients. The purpose is to avoid detection and blocking by various Intrusion Detection and Intrusion Prevention Systems (IDS/IPS). Given a large enough population of attack hosts, an attacker can simply instruct them to visit the target website (or other service) by faithfully following the relevant protocol interactions. In doing so, the attacker is simulating a "flash crowd", albeit one in which the majority of the crowd are automata under his or her control. One countermeasure against such attacks involves the use of Reverse Turing Tests, which seek to discriminate between human users and automata by posing (hopefully) problems that are currently hard to solve algorithmically. Examples of such problems include the recognition of words or characters in a distorted image, or the identification of an image with specific features from among a larger set of images (*e.g.,* "select the image that shows a dog running in a park").

DoS attacks sometimes use evasive tactics to hide their origin and to mask any distinguishing characteristics that could be used for block them. The earliest and most obvious strategy involved using a fake source IP address (spoofing). This was a necessity when only a few hosts were participating in the attack, and thus could be identified and easily blocked. With the advent of large botnets and the use of anti-spoofing techniques by ISPs, address spoofing is neither crucial for an attack's success nor as common. Other ways of hiding the attack's true origins involve the use of *reflectors*. These are hosts (and protocols/services) that will respond to traffic sent to them, without validating the source IP address. Typically, such services are connectionless, *e.g.,* ICMP packets or certain UDP-based protocols. Attackers can send traffic to these hosts, purporting to originate from the target site by using address spoofing, causing all response traffic to go to the target. Other than re-designing the protocol, the only countermeasure available to reflection involves traffic shaping, *i.e.,* limiting the amount of such unverified/unverifiable traffic that a host will generate per unit time. Reflectors become an even more serious problem when they also act as *amplifiers.* Amplifiers are hosts (and protocols/services) that respond with a large volume of traffic in response to a much smaller request. Examples of traffic amplification include Domain Name System (DNS) zone transfers over UDP, targeted subnet broadcast, and streaming media traffic (such as Voice over IP streams). Protocol re-design seems to be the best way of avoiding such vulnerabilities.

The use of spoofing, reflection and amplification make it difficult to impersonate a legitimate host and to blend attack traffic with legitimate traffic for the purposes of evading detection. One exception to this involves the redirection of large numbers of clients toward the target. The means of such redirection depend on the particular type of traffic and protocol that is being redirected. One example of such redirection is causing web browsers to contact the target site by including either static content pointers to the target or active content (*e.g.,* Javascript or Flash scripts) that repeatedly attempts to connect to the target without the legitimate user's action or consent, by including such pointers or active content in a popular web page, often in a compromised server. Another example involves redirecting requests for searched-for content in a peer-to-peer network toward the target site, causing a large number of download requests from clients that seek to retrieve said content.

Bandwidth DoS attacks are often combined with other types of DoS attacks. For example, they may be combined with memory/state exhaustion attacks, as is the case with TCP SYN flood attacks that seek to fill up operating system tables, or with algorithmic DoS attacks that try to induce excessive computation at the target host, *e.g.,* by repeatedly invoking operations that result in expensive database lookups. In this way, the impact of the attack is multiplied. Furthermore, even if the attack volume proves insufficient to saturate the target's available bandwidth, the secondary effects may be sufficient to deny service to legitimate clients.

Another way to amplify an attack composed of legitimate-looking traffic, is to seek access to large resources, *e.g.,* big files or media content. Because modern network links are full-duplex (*i.e.,* traffic can flow in both directions without interference), most network DoS attacks affect primarily the target's downstream bandwidth (*i.e.,* from the Internet to the target). By seeking to access large-volume content, attacks can also saturate the upstream bandwidth. In some cases, attacks primarily focus on the upstream bandwidth. Such attacks can be particularly effective, since they require relatively little attack traffic to cause disproportionately large responses. However, such DoS attacks are easy to identify and remedy at the target site.

Tools for launching bandwidth DoS attacks can be commonly found on the Internet. Their sophistication and effectiveness varies. Practically all botnet software allows its controllers to launch bandwidth DoS attacks, sometimes offering sophisticated features such as "pulsing attacks" (where an attack is turned on and off, in an effort to confound detection and defense), rotation (where different subsets of the botnet may participate in the attack at different times, again towards confounding defense), scheduled attacks, *etc.* In some recent cases, especially where DoS attacks were launched for political purposes, tools were made available for sympathizing users to download and purposely use to contribute to the attacks.

Defending against bandwidth DoS attacks is often difficult for the target site, because the congestion usually occurs upstream (farther in the network) from any equipment that the site controls (*e.g.,* a router or firewall). For an effective response, a target site typically needs to coordinate a response with its parent ISP. If the attack traffic is easy to characterize or otherwise "stands out", such as a UDP packet flood against a web site, blocking at an appropriate upstream location by the ISP is relatively straightforward. When the attack traffic is not easy to characterize, or the necessary router resources or features are not available for filtering, ISPs resort to *blackholing*. In that case, a routing entry for the target's network prefix (which may include other, otherwise un-targeted sites) is injected into the ISP's routing protocol. That entry points effectively causes all traffic to the target, both legitimate and attack, to be dropped by the ISP routers. In this way, the attack traffic is dropped as soon as it enters the ISP's network, thereby avoiding link congestion. Since legitimate traffic to the target site is also dropped, blackholing does not help restore access to said site.

Another practical defense against network bandwidth attacks involves the use of Content Delivery Networks (CDNs). For relatively static content, this allows a website to effectively create many different instances of itself distributed around the Internet, making it difficult for an attacker to completely deny access to all of them. However, this approach does not work well for sites with dynamic or interactive content.

Modern routers and firewalls offer some defenses against bandwidth DoS attacks, such as terminating TCP handshakes to mitigate against SYN flood attacks and applying statistical means to detecting network traffic anomalies. The effectiveness of such schemes varies against the different types of attack traffic, as does their impact on router performance. Many of these features are available on customer-premises equipment (as opposed to ISP-grade equipment), reducing their relevance in defending against attacks that cause congestion in upstream links.

Other types of DoS attacks, such as those that exhaust available memory on a system or cause excessive CPU utilization by exploiting worst-case algorithmic behavior via carefully selected inputs, can be used to the same effect. However, these typically require an good understanding of the particular system and application that is being attacked. In contrast, bandwidth DoS attacks can be launched against any host or network with little or no knowledge of the services that are available in the target.

Open Problems

Considerable research has been conducted in the areas of traffic characterization and anomaly detection, proactive network architecture design, reactive mechanism retrofitting, attack tolerance and avoidance, and attribution [MDDR2005]. The problem of detecting and defending against bandwidth denial of service attacks in open networks, and specifically in the Internet, remains unsolved.

Recommended Reading List

- [MDDR2005] Mirkovic J., Dietrich S., Dittrich D., Reiher P., Internet Denial of Service: Attack and Defense Mechanisms, Prentice Hall PTR, 2005.