Defending against Distributed Denial of Service

Yang Xiang
Central Queensland University, Australia

Wanlei Zhou
Deakin University, Australia

WHAT IS THE DDOS ATTACK?

Recently the notorious Distributed Denial of Service (DDoS) attacks made people aware of the importance of providing available data and services securely to users. A DDoS attack is characterized by an explicit attempt from an attacker to prevent legitimate users of a service from using the desired resource (CERT, 2006). For example, in February 2000, many Web sites such as Yahoo, Amazon.com, eBuy, CNN.com, Buy.com, ZDNet, E*Trade, and Excite.com were all subject to total or regional outages by DDoS attacks. In 2002, a massive DDoS attack briefly interrupted Web traffic on nine of the 13 DNS "root" servers that control the Internet (Naraine, 2002). In 2004, a number of DDoS attacks assaulted the credit card processor Authorize.net, the Web infrastructure provider Akamai Systems, the interactive advertising company DoubleClick (left that company's servers temporarily unable to deliver ads to thousands of popular Web sites), and many online gambling sites (Arnfield, 2004). Nowadays, Internet applications face serious security problems caused by DDoS attacks. For example, according to CERT/CC Statistics 1998-2005 (CERT, 2006), computer-based vulnerabilities reported have increased exponentially since 1998. Effective approaches to defeat DDoS attacks are desperately demanded (Cisco, 2001; Gibson, 2002).

Figure 1 shows a hierarchical model of a DDoS attack. The most common attacks involve sending a large number of packets to a destination, thus causing excessive amounts of endpoint, and possibly transit, network bandwidth to be consumed (Householder, Manion, Pesante, Weaver, & Thomas, 2001). The attack usually starts from multiple sources to aim at a single target. Multiple target attacks are less common; however, there is the possibility for attackers to launch such type of attack.

In order to launch a DDoS attack, the attacker first scans millions of machines for vulnerable service and other weaknesses that permit penetrations, then gains access and compromises these machines' so-called handlers, and zombies or slaves. Malicious scripts—such as scanning tools, attack tools, root kits, sniffers, handler and zombie programs, and lists of vulnerable and previously compromised hosts, and so forth—are then installed, and the infected machines can recruit more machines. This propagation phase is quite like computer viruses.

Next the communication channels between the attacker and the handlers, and between the handlers and zombies are established. These control channels are designed to be secret from the public, in order to conceal the activity of the attacker. TCP, UDP, ICMP, or a combination of these protocols is used to perform the communication. Recently, some attack tools exploited the existing infrastructure of Internet Relay Chat (IRC) networks, which are not as easily discovered as earlier versions, because they do not present a new open port that could be found by a scan or audit scheme (Houle & Weaver, 2001).

Staying behind the scenes of attack, the real attacker sends a command to the handlers to initiate a coordinated attack. When the handlers receive the command, they transfer it to the zombies under their control. Upon receiving attack commands, the zombies begin the attack on the victim (Lau, Stuart, & Michael, 2000). The real attacker is trying to hide himself from detection, for example by providing spoofed IP addresses. It makes it difficult to trace the real source of the attacker and filter malicious packets from the legitimate traffic.

Figure 1. A hierarchical model of a DDoS attack

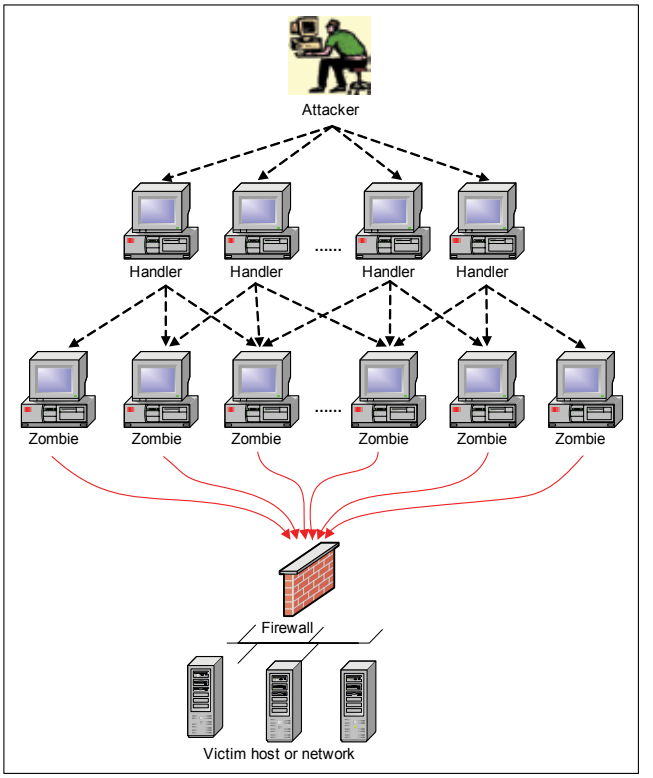

PASSIVE DEFENSE AGAINST DDOS ATTACKS

Passive Defense Cycle

We define passive defense as defense actions taken only after the DDoS attacks are launched. Hence, the target host or network is harmed to some certain extent before the attack source(s) can be located and handled. The traditional passive defense mechanism includes a protect-detect-react cycle (Householder et al., 2001). That is, after attack actions are detected, some reacting steps are taken, such as traffic limiting, blocking, and filtering. This method has advantages over the poor "lesson learned" experience, which responds to the attack only after the accident is over. However, it is far from enough. We need an active defense system with a surveillance-trace-control cycle, which will be presented in detail later in this article.

By deploying the passive defense system, an attack is usually detected by monitoring of inbound traffic volumes and other performance metrics. But ironically, the first signal of attack often comes from the external customer's report that shows the service is no longer reachable, instead of the alarm of detection system. Then apparently it is too late to protect the victim from the attack.

Current Passive Defense Mechanisms

Passive defense mechanisms can be classified into two categories: one is the detecting mechanism, and the other is the reacting mechanism. The common detection method includes monitoring traffic volumes and source IP addresses, and resource accounting. However, usually simply monitoring the traffic volume cannot tell accurately the real attack, because sometimes Internet flash crowds also cause network congestion (Jung, Krishnamurthy, & Rabinovich, 2002). So this method cannot differentiate legitimate requests or malicious requests. According to the characteristic of IP spoofing techniques of DDoS attack, monitoring source IP addresses is a feasible measure to mitigate the attack.

After detecting the malicious actions of DDoS attacks, the passive defense system turns into reacting