Defending against Flooding-Based Distributed Denial-of-Service Attacks: A Tutorial

Paper by Rocky K C Chang, The Hong Kong Polytechnic University

Published in the October 2002 issue of *IEEE Communications Magazine*

Report & Critique by

James Gaskell
10 October 2011

CS 577
Prof Robert E Kinicki
Worcester Polytechnic Institute

Abstract

The focus of the paper is the “Flooding-Based Distributed Denial-of-Service Attack.”

> “Flooding-Based” is implied by the “Distributed Attack” and is soon dropped and is replaced everywhere with the simpler “DDoS” term.

DDoS is described as a “very serious threat.” In Feb, 2000, eBay, Amazon and Yahoo! were taken down.

So the author intends to:

1. Describe, via a tutorial, how a DDoS “works”

2. Describe current common “solutions,”

3. Look at newer, more comprehensive defenses.

Effect of the Attack

The basis of the attack is to "harness" the vast resources of the Internet to overwhelm a selected "victim."

The core idea is to have various "masters" place malicious code in many, many unprotected computers sitting in homes all around the country (world) that will, upon command, unleash a torrent of requests (or other messages) to the victim's server.

Thoughtfully, the author does not give any information about how to accomplish this step!

Once the attack is unleashed, there are two basic effects:

1. The sheer volume of incoming requests overloads the ability of the victim's server to handle the requests.

2. In a "SYN" attack (to be described later), the resources of the victim's server are exhausted by (1) having a client start a connection (to which the server allocates resources), but then (2) not allow the connection to complete and (3) having the server sit – often for many seconds – for a time-out to occur at the server's end and (4) to finally allow the server to release those resources so other connections can be made.

Attack Hierarchy

One or more “Attackers” begin the process.

The Attacker installs small programs in unprotected computers that are open to the Internet. These computers are called “Masters” or “Handlers.”

The purpose of these Masters is to infect other (and more numerous) computers called, variously, “Agents,” “Daemons” or “Zombies.”

Then, at some point, the Attackers instruct the Masters to instruct the Agents to begin flooding the Internet with their pre-ordained packets.

While the general idea of how the process works is described in the paper, no details are provided (Thank goodness!).

Direct v Reflected Attacks

A “Reflector” in an attack is simply a computer with access to the Internet whose address is known to the Attacker (via ping?); it does not need to be “infected.” The Reflector IP address is then spoofed into an attack packet as the source address. When it later gets an unexpected message from the Victim, it makes its normal reply (usually error message) back to the Victim, further using up bandwidth at the Victim’s node.

Types of Packets

UDP (User Datagram Protocol)

ICMP (Internet Control Message Protocol)

TCP (Transmission Control Protocol)

TCP SYN-ACK (TCP w/forged SYN addresses)

TCP RST (TCP w/forged Resets)

How Many Packets?

Microsoft Win 2000 Advanced Server

9 second release

BSD

Unspecified, but from graph, approx 100 secs.

Linux (2.2.29-19)

309 sec

For a Win2K server that allows 10K half-open connections, 1100 requests/second will "stall" the server.

Special note: In a SYN-ACK attack, the attacking node AUTOMATICALLY sends out additional requests when its SYNs are not ACKed.

General Flooding

1.544 Mb/s will jam a T1. That's 5K pings/sec. Or, for an echo request w/long echoes, many fewer.

Defenses

1. Prevent Masters & Agents from ever being initialized.

    Microsoft, Symantec, McAfee, etc.
    Firewalls.

Possibly being able to snoop traffic on the Internet for Master/Agent probes and back-trace any found.

2. After the fact tracing.

IP Traceback allows finding the source even when the source address in the IP header is false.

    xxxx

3. Have ISPs drop out-going packets whose "source" address is not in the ISP's realm.

    Good luck.

4. Have a victim "know" it's being attacked.

    Not always that clear!

5. Once an attack is underway, try to develop a "signature" of the attack packets.

    NPSR ("Normal" packet survival ratio)
    No documentation of any reasonable successes.

6. The Victim manually notifies up-stream ISPs to block packets with the offending signature.

    Needs ISP cooperation
    Needs a "good signature"
    Needs procedures in place before-the-fact.

A New Approach: An Internet "Firewall"

Two approaches, both employing distributed nodes within the Internet to (1) detect attacks and (2) do "local" packet filtering.

1. Route-based Packet Filtering (RPF)

Compares source IP address, destination IP address & BGP (Border Gateway Protocol) routing information to detect suspicious packets and drop them.

> Recent changes in routing might cause "legal" packets to be dropped.
>
> This would necessitate a change to the BGP messaging.
>
> Even if implemented, the problem is not eliminated but only reduced. The resulting reduced traffic is still easily crippling.

Also, packets from to and from Reflectors will never be detected.

2. Distributed Attack Detection (DAD)

Uses a wide area set of Detection Systems (DSs)

> "Significant" deviation from normal.
> Placed in "strategic" locations
> "Local" and "Global" detection
> An inter-DS communication system

Questions?

Ideas?

Comments?