Available online at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/cose

Computers & Security

A framework for generating realistic traffic for Distributed Denial-of-Service attacks and Flash Events

*Sajal Bhatia**, *Desmond Schmidt*, *George Mohay*, *Alan Tickle*

Information Security Discipline, Science and Engineering Faculty, Queensland University of Technology, Brisbane, Queensland, Australia

ARTICLE INFO

Article history:
Received 21 July 2013
Received in revised form
30 September 2013
Accepted 16 November 2013

Keywords:
Synthetic traffic generation
DDoS attacks
Flash Events
IP-aliasing
Testbed framework

ABSTRACT

An intrinsic challenge associated with evaluating proposed techniques for detecting Distributed Denial-of-Service (DDoS) attacks and distinguishing them from Flash Events (FEs) is the extreme scarcity of publicly available real-word traffic traces. Those available are either heavily anonymised or too old to accurately reflect the current trends in DDoS attacks and FEs. This paper proposes a traffic generation and testbed framework for synthetically generating different types of realistic DDoS attacks, FEs and other benign traffic traces, and monitoring their effects on the target. Using only modest hardware resources, the proposed framework, consisting of a customised software traffic generator, 'Botloader', is capable of generating a configurable mix of two-way traffic, for emulating either large-scale DDoS attacks, FEs or benign traffic traces that are experimentally reproducible. Botloader uses IP-aliasing, a well-known technique available on most computing platforms, to create thousands of interactive UDP/TCP endpoints on a single computer, each bound to a unique IP-address, to emulate large numbers of simultaneous attackers or benign clients.

1. Introduction

In spite of considerable research conducted over the past decade, Distributed Denial-of-Service (DDoS) attacks in various guises continue to constitute a pernicious threat within the Internet community (Nazario, 2008). These attacks, using the same basic *modus operandi* as some of the earliest known attacks, still exist; the only difference being their magnitude, complexity and frequency, which have all increased dramatically. Hence, developing techniques to accurately and reliably detect DDoS attacks, and differentiate them from Flash Events (FEs) – in which a similar high-load is generated from a large number of benign clients accessing the server simultaneously – remains an active area of research. However, the requisite research in developing such techniques is hampered by the extreme scarcity of *recent* and *realistic* public domain datasets representing real traffic traces, whether attack or benign, for testing and evaluation purposes.

This paper addresses this challenge by proposing a traffic generation and testbed framework to synthetically generate a variety of realistic DDoS attacks, FEs, and other benign traffic traces. The proposed framework uses modest hardware and a software traffic generator, called Botloader, which exploits IP-

* *Corresponding author.*
E-mail addresses: s.bhatia@qut.edu.au (S. Bhatia), desmond.schmidt@qut.edu.au (D. Schmidt), g.mohay@qut.edu.au (G. Mohay), ab.tickle@qut.edu.au (A. Tickle).

0167-4048/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cose.2013.11.005

aliasing, a well-known technique available on most computing platforms for associating multiple IP addresses with a single Network Interface Card (NIC).

The remainder of the paper is structured as follows. Section 2 provides an overview of the various techniques so far proposed for obtaining or generating realistic DDoS and FE datasets, focusing on the limitations of the existing public domain datasets, and on problems intrinsic to the generation of synthetic traffic in the laboratory. Section 3 discusses the design of the traffic generation and testbed framework, consisting of a hardware testbed and software traffic generator – Botloader. Section 4 discusses the implementation details of the experimental testbed environment used for traffic generation. Section 5 presents the emulation results of a real-world DDoS attack and an FE, and evaluates how closely they mimic real traffic traces. Finally, Section 6 summarises the work and describes the future directions for this research.

2. Background and related work

Over the years DDoS attacks have evolved from a naïve low-scale to a more sophisticated large-scale problem. During this period there has been considerable research and development into attack detection strategies, but only limited research on testing these techniques against realistic data. One of the key reasons for this has been the legal and privacy issues associated with sharing captured data. As a result, the majority of published work on DDoS attack detection is evaluated by: (a) replaying publicly available datasets, (b) generating traffic using open-source traffic generators, and (c) using testbeds based on simulation, emulation or direct physical representation strategies. Each of these approaches has its own technical limitations, may not reflect current trends in DDoS attack scenarios, and may produce misleading results for the evaluation of the proposed detection techniques. Each of these approaches will now be discussed.

2.1. *Replaying traffic traces*

One obvious way to test and evaluate a workable solution for detecting DDoS attacks is to replay traces of real attacks (Ahmed et al., 2010; Buchanan et al., 2011). Although DDoS attacks have become common, real-world attack traces in the public domain are rare, because the sharing of captured live data (both attack and normal) is limited by significant legal and privacy issues.

The KDD Cup 1999 Data is one of the most referenced and possibly the only reliably-labelled dataset available in the public domain (Hettich and Bay, 1999). Use of the KDD dataset completely eliminates any legal and privacy-related issues that arise when using any *real-world* data. However, due to its age (1999) and other limitations e.g., its prime usage is for evaluating signature-based intrusion detection systems (McHugh, 2000; Tavallaee et al., 2009), it is not really suitable for testing DDoS detection and mitigation mechanisms. The set of network traces from Waikato Internet Trace Storage Project (Waikato Applied Network Dynamic Research Group) is another widely referenced dataset within the DDoS detection domain. However, these datasets have their own limitations: IP addresses are pseudonymised, packets are truncated 20 bytes after the transport header, and any retained payloads of UDP packets (except the DNS) are zeroed. A more recent addition has been the 'CAIDA DDoS Attack 2007 Dataset' containing an approximately one hour traffic trace of a DDoS attack which aimed at consuming the computing resource of the targeted server (Hick et al., 2007). However, the available traces have had their IP addresses pseudonymised using CryptoPAn and their payloads removed, thereby limiting their usability.

In the FE domain, a very limited number of datasets representing different types of FEs are publicly available. The 1998 FIFA World Cup Dataset, provided by the Internet Traffic Archive, is the most widely used and probably the only predictable FE[1] dataset available in the public domain.

In addition to these specific limitations, these datasets also share some characteristics which greatly limit their usability. First, excepting the CAIDA DDoS Attack 2007 Dataset, the other datasets are old and obsolete in the context of continuously evolving DDoS attacks with complex attack vectors. Secondly, most of the available datasets (except KDD and DARPA) are not labelled, which makes it difficult to cleanly filter out the attack from benign background traffic, and increases the training difficulty for detection strategies that rely on machine learning algorithms. Thirdly, even though the target addresses of the packets in these datasets can be altered and subsequently directed to a test machine while replaying, the real effects of the network trace still cannot be reproduced due to the non-availability of the services addressed in the original trace. Fourthly, the missing payloads completely eliminate the possibility of developing any *content-based* analysis technique. Lastly, the majority of FE datasets are web-server logs in Common Log Format, thereby making them difficult to use, e.g. it is not possible to replay them over a network to test FE detection algorithms.

2.2. *Traffic generators*

Given the limitations of direct replay of captured traffic traces, artificial generation would seem to be the only practical way to generate the required traffic, both attack and benign. Unfortunately both hardware and software traffic generators are far from being ideal tools for simulating such attacks.

The hardware traffic generators can do much the same as their software counterparts, but at higher speed and at greater cost. Hardware traffic generators like the SmartBits 600 (SmartBits, 2001) can artificially generate configurable flooding attacks, but they fail to interact with the application at the target machine. Our experiments with the SmartBits generator highlighted significant differences in how the computers and the other networking devices like switches and routers responded to hardware-generated and software-generated

[1] A predictable FE is one whose expected occurrence is known *a priori*, thus allowing network administrators to prepare for it by using various provisioning and load-balancing techniques. Some popular examples are product releases (e.g. by hi-tech companies like Apple), widely followed sporting events such as the Olympics, or online play-along websites for popular television programs. A detailed classification of FEs can be obtained in the paper by Bhatia et al. (2012).