Available online at www.sciencedirect.com

ScienceDirect

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

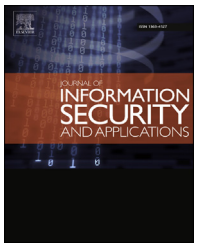

PeerRush: Mining for unwanted P2P traffic

Babak Rahbarinia [a,*], *Roberto Perdisci* [a], *Andrea Lanzi* [b], *Kang Li* [a]

[a] *Dept. of Computer Science, University of Georgia, Athens, GA 30602, USA*
[b] *EURECOM Institute, Sophia Antipolis, France*

ARTICLE INFO

Article history:
Available online 24 April 2014

Keywords:
P2P
Traffic classification
Botnets

ABSTRACT

In this paper we present PeerRush, a novel system for the identification of *unwanted* P2P traffic. Unlike most previous work, PeerRush goes beyond P2P traffic detection, and can accurately *categorize* the detected P2P traffic and attribute it to specific P2P applications, including malicious applications such as P2P *botnets*. PeerRush achieves these results without the need of deep packet inspection, and can accurately identify applications that use encrypted P2P traffic.

We implemented a prototype version of PeerRush and performed an extensive evaluation of the system over a variety of P2P traffic datasets. Our results show that we can detect all the considered types of P2P traffic with up to 99.5% true positives and 0.1% false positives. Furthermore, PeerRush can attribute the P2P traffic to a specific P2P application with a misclassification rate of 0.68% or less.

Published by Elsevier Ltd.

1. Introduction

Peer-to-peer (P2P) traffic represents a significant portion of today's global Internet traffic (Madhukar and Williamson, 2006). Therefore, it is important for network administrators to be able to identify and categorize P2P traffic crossing their network boundaries, so that appropriate fine-grained network management and security policies can be implemented. In addition, the ability to categorize P2P traffic can help to increase the accuracy of network-based intrusion detection systems (Haq et al., 2010).

While there exist a vast body of work dedicated to P2P traffic detection (Gomes et al.), a large portion of previous work focuses on signature-based approaches that require deep packet inspection (DPI), or on port-number-based identification (Sen et al., 2004; Hayes). Because modern P2P applications avoid using fixed port numbers and implement encryption to prevent DPI-based detection (Madhukar and Williamson, 2006), more recent work has addressed the problem of identifying P2P traffic based on statistical traffic analysis (Karagiannis et al., 2004; Karagiannis et al., 2005). However, very few of these studies address the problem of P2P traffic categorization (Hu et al., 2009), and they are limited to studying only few types of non-encrypted P2P communications. Also, a number of previous studies have focused on detecting P2P botnets (Gu et al., 2008; Yen and Reiter, 2010; Nagaraja et al., 2010; Coskun et al., 2010; Zhang et al., 2011), but with little or no attention to accurately distinguishing between different types of P2P botnet families based on their P2P traffic patterns.

In this paper, we propose a novel P2P traffic categorization system called PeerRush. Our system is based on a generic classification approach that leverages high-level statistical traffic features, and is able to accurately detect and categorize the traffic generated by a variety of P2P applications,

* *Corresponding author.*
E-mail addresses: babak@cs.uga.edu, babak.rahbarinia@gmail.com (B. Rahbarinia), perdisci@cs.uga.edu (R. Perdisci), lanzi@eurecom.fr (A. Lanzi), kangli@cs.uga.edu (K. Li).
http://dx.doi.org/10.1016/j.jisa.2014.03.002
2214-2126/Published by Elsevier Ltd.

including common file-sharing applications such as μTorrent, eMule, etc., P2P-based communication applications such as Skype, and P2P-botnets such as Storm (Holz et al., 2008), Waledac (Nunnery et al., 2010), and a new variant of Zeus (Lelli) that uses encrypted P2P traffic. We would like to emphasize that, unlike previous work on P2P-botnet detection, PeerRush focuses on accurately detecting and **categorizing different types of legitimate and malicious P2P traffic**, with the goal of identifying *unwanted* P2P applications within the monitored network. Depending on the network's traffic management and security policies, the unwanted applications may include P2P-botnets as well as certain specific legitimate P2P applications (e.g., some file-sharing applications). Moreover, unlike most previous work on P2P-botnet detection, **PeerRush can reveal if a host is compromised with a specific P2P botnet type** among a set of previously observed and modeled botnet families. To the best of our knowledge, no previous study has proposed a generic classification approach to accurately detect and categorize network traffic related to both legitimate and malicious P2P applications, including popular applications that use encrypted P2P traffic, and different types of P2P-botnet traffic (encrypted and non-encrypted).

Fig. 1 provides an overview of PeerRush, which we discuss in detail in Section 2. The first step involves the identifications of P2P hosts within the monitored network. Then, the P2P traffic categorization module analyzes the network traffic generated by these hosts, and attempts to attribute it to a given P2P application by matching an *application profile* previously learned from samples of traffic generated by known P2P applications. If the P2P traffic does not match any of the available profiles, the traffic is classified as belonging to an "unknown" P2P application (e.g., this may represent a new P2P application release or a previously unknown P2P botnet), and should be further analyzed by the network administrator. On the other hand, if the P2P traffic matches more than one profile, an auxiliary *disambiguation* module is used to "break the tie", and the traffic is labeled as belonging to the closest P2P application profile.

The application profiles can model the traffic characteristics of legitimate P2P applications as well as different P2P-botnets. It is common for security researchers to run botnet samples in a controlled environment to study their system and network activities (Egele et al., 2008). The traffic collected during this process can then be used as a sample for training a specific P2P-botnet application profile, which can be plugged into our P2P traffic categorization module. In summary this paper makes the following contributions:

- We present PeerRush, a system for **P2P traffic categorization** that enables the accurate identification of *unwanted* P2P traffic, **including encrypted P2P traffic and different types of P2P botnet traffic.** To achieve these goals, we engineer a set of novel statistical features and classification approaches that provide both accuracy and robustness to noise.
- We collected a variety of P2P traffic datasets comprising of P2P traffic generated by five different legitimate P2P applications used in different configurations, and three different P2P botnets including a P2P botnet that employs encrypted P2P traffic. We are making these datasets *publicly available*.
- We performed an extensive evaluation of PeerRush's classification accuracy and noise resistance. Our results show that we can detect all the considered types of P2P traffic with up to 99.5% true positives and 0.1% false positives. Furthermore, PeerRush can correctly categorize the P2P traffic of a specific P2P application with a misclassification rate of 0.68% or less.

2. System overview

PeerRush's main goal is to enable the discovery of *unwanted* P2P traffic in a monitored computer network. Because the exact definition of what traffic is unwanted depends on the management and security policies of each network, we take a generic P2P traffic categorization approach, and leave the final decision on what traffic is in violation of the policies to the network administrator.

To achieve accurate P2P traffic categorization, PeerRush implements a two-stage classification system that consists of a *P2P host detection* module, and a *P2P traffic categorization* module, as shown in Fig. 1. PeerRush partitions the stream of *live* network traffic into time windows of constant size W (e.g., W=10 minutes). At the end of each time window, PeerRush extracts a number of statistical features from the observed

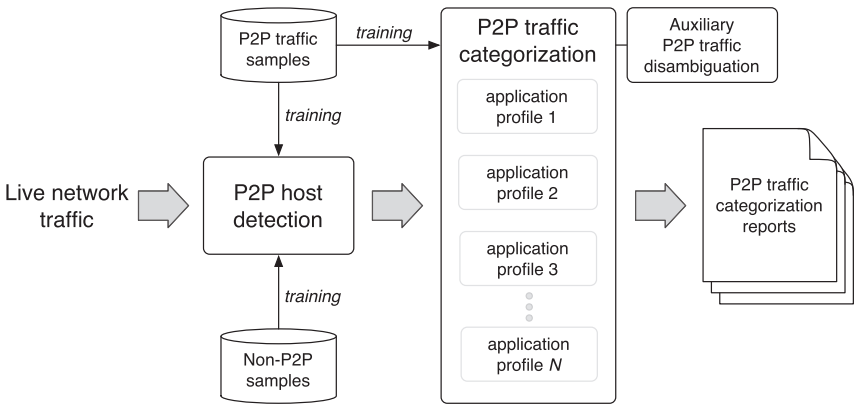

Fig. 1 – PeerRush system overview.