Journal of Computer and System Sciences 80 (2014) 502–519

Contents lists available at SciVerse ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Estimating the number of hosts corresponding to an intrusion alert while preserving privacy

Alif Wahid *, Christopher Leckie, Chenfeng Zhou

Department of Computing and Information Systems, The University of Melbourne, Parkville, VIC 3010, Australia

ARTICLE INFO

Article history:
Received 30 November 2012
Received in revised form 30 April 2013
Accepted 14 June 2013
Available online 2 July 2013

Keywords:
Intrusion detection
Address aliasing
Privacy protection
Statistical modelling

ABSTRACT

An inherent feature of IP addresses is the aliasing that arises due to dynamic address allocation. This creates a significant barrier to the estimation of the malicious host population from a set of intrusion alerts. In this paper, we propose a method for estimating the number of malicious hosts that may have bound to an alerted address, based on the correlation of different data sets that were collected independently and a probabilistic model of host-to-address bindings. We analysed a two week trace of real-world intrusion alerts along with a global survey of ping responses, and inferred that over 80% of malicious addresses were bound to multiple hosts. Such aliasing effects highlight the inaccuracy of assuming static bindings between hosts and addresses when exact host identification is not possible due to privacy protection. However, our method demonstrates that reliable inferences can still be made when a sufficient overlap exists between the correlated data sets.

1. Introduction

IP addresses are a limited resource for hosts connected to the Internet. This stems from the successful deployment of the IPv4 protocol [1] over the course of three decades, which also had the unforeseen consequence of exhausting the limited pool of globally unique addresses [2]. A large number of control plane protocols have also been deployed over the years to support the rapid growth of IPv4 infrastructure throughout the world. These include dynamic address assignments [3,4], translation of addresses at network gateways [5], domain name look-ups for decoupling addresses from servers [6], application layer proxies, and firewalls for security and privacy protection.

As a result, IP addresses are seldom tightly bound to individual hosts for long periods in the order of weeks or months, since they need to be recycled and/or shared by multiple hosts. This aliasing phenomenon raises a fundamental challenge from the perspective of network monitoring: *how can we reliably infer the number of hosts corresponding to an address*? The difficulty of this challenge is further highlighted by the legal requirement of protecting the anonymity of all users associated with any given host, such that observations of an address must never be *unambiguously* linked with human activity that could reveal user identities and violate privacy (refer to the "Information Privacy Principles" legislated in the authors' jurisdiction [7]).

Internet hosts generate and absorb all traffic that one can measure at the network layer. This means that any uncertainty underlying the population of hosts due to dynamic addressing has a pervasive impact on all analysis and inference tasks involved in network measurement and traffic engineering. Moreover, in the context of network intrusion detection and prevention, one needs reliable ways of identifying and tracking those Internet hosts that engage in deliberate malicious activity

* Corresponding author.
E-mail addresses: a.wahid@student.unimelb.edu.au (A. Wahid), caleckie@unimelb.edu.au (C. Leckie), cvzhou@gmail.com (C. Zhou).

0022-0000/$ – see front matter
http://dx.doi.org/10.1016/j.jcss.2013.06.007

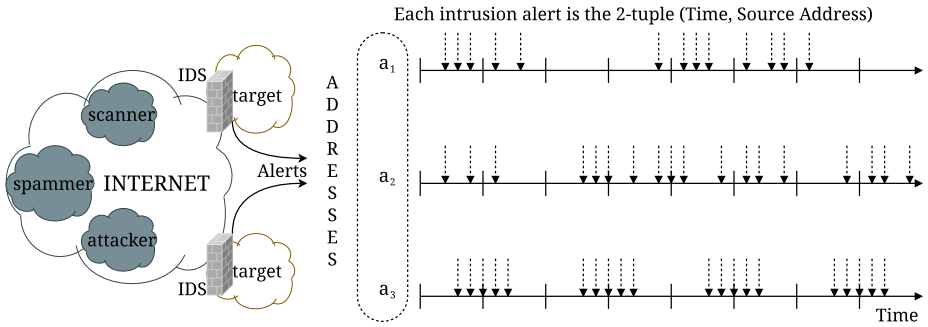

Fig. 1. An example scenario of intrusion alerts that need to be unambiguously mapped to malicious hosts for reliable blacklisting and/or other persistent countermeasures.

and cyber-crime over the course of many weeks and months, e.g., spamming mail servers, Distributed Denial of Service (DDoS) botnets, port-scanning worms, and so on. In these cases, accurately estimating the number of hosts corresponding to an address is of real practical significance since the threat scales with the population of malicious hosts and not the number of unique addresses [8].

The requirement of privacy protection severely restricts access to any data that might improve the accuracy of estimation (e.g., Ethernet MAC addresses that are tightly bound to their corresponding hosts). Therefore, any method of estimating the number of hosts corresponding to an address must find a trade-off that preserves user anonymity as well as providing accurate and reliable estimates of the host population. In this paper, we propose a solution to this problem and demonstrate its validity using real-world data sets. Our main contributions are as follows.

1. We develop a probabilistic model of host-to-address bindings that is dependent on one hidden variable and one observed variable. This model allows the calculation of conditional probabilities for the number of hosts corresponding to an address, given: a) the number of times the address appeared, and b) its latent probability of binding to a new host in each appearance.
2. We derive and test careful approximations of this model using a large data set of dynamic host-to-address bindings from a campus wireless network that comes with anonymised ground truth regarding the exact number of hosts corresponding to an observed address. However, this ground truth is purely for validation and our model is only dependent on the two stated parameters related to the observed addresses (as opposed to parameters related to the hidden hosts, which potentially violate user anonymity).
3. We correlate real-world intrusion alerts with ping responses in order to find a common set of active addresses that are attached to malicious hosts and subsequently apply the proposed model of host-to-address bindings. We find that more than 80% of these addresses bind to multiple hosts, which highlights the limitations of popular countermeasures such as static blacklisting.

The rest of this paper is organised as follows. We begin in Section 2 by describing some of the motivating use cases and operational context that highlight the practical importance of this problem. Then Section 3 derives our proposed model and presents various implications. We develop approximations of this model in order to efficiently estimate the corresponding number of hosts for large numbers of addresses in Section 4. These methods are subsequently tested and validated in Section 5 using a four month trace of sanitised DHCP logs collected from a campus wireless network. The correlation of intrusion alerts and ping responses is described in Section 6. The statistical characteristics of this correlated data set is presented in Section 7. We apply the proposed model of host-to-address bindings and present the results in Section 8. The implications of our findings with respect to the existing literature are discussed in Section 9. In Section 10, we compare our model with previous literature. Finally, we summarise and conclude the paper in Section 11.

2. Operational context

Consider the network security scenario of Fig. 1. A private network operator who happens to be among the targets of malware propagating on the Internet can install an Intrusion Detection System (IDS) to automatically monitor the transport layer packet-flows entering and leaving their network. Any packet-flow found to be malicious according to pre-specified criteria can be flagged and its corresponding source address blacklisted as a typical countermeasure. However, due to the temporal volatility of the binding between a malicious host and its corresponding address, such intrusion alerts are necessarily unreliable. Refer to [9] for an example of fast-flux service networks that deliberately exploit such volatility to remain undetected for considerable periods. In addition, a common form of attack scenario involves spoofing addresses due to the unaccountability of the IPv4 protocol suite [10], which only accentuates the aliasing and uncertainty of host-to-address bindings. Hence, the process of blacklisting can lead to adverse consequences whereby legitimate users are denied access (for example, the paying customers of a cloud service) [11]. A method of unambiguously inferring the binding between