First Insights from a Mobile Honeypot

Matthias Wählisch*, Sebastian Trapp*, Christian Keil†, Jochen Schönfelder†,
Thomas C. Schmidt‡, Jochen Schiller*,
*Institut für Informatik, Freie Universität Berlin, Germany
†DFN-CERT Services GmbH, Hamburg, Germany
‡Dept. Informatik, Hamburg University of Applied Sciences, Hamburg Germany
{first.last}@fu-berlin.de, {last}@dfn-cert.de, t.schmidt@ieee.org

ABSTRACT

Computer systems are commonly attacked by malicious transport contacts. We present a comparative study that analyzes to what extent those attacks depend on the network access, in particular if an adversary targets specifically on mobile or non-mobile devices. Based on a mobile honeypot that extracts first statistical results, our findings indicate that a few topological domains of the Internet have started to place particular focus on attacking mobile networks.

Categories and Subject Descriptors

C.2.0 [**Computer-Comm. Networks**]: General—*Security and protection (e.g., firewalls)*

General Terms

Security

Keywords

Mobile vs. non-mobile attacks, mobile honeypot

1. INTRODUCTION

Common attacks on Internet devices start with a connection to a random or particular TCP/UDP port. An adversary attempts to perform a denial of service (DoS) attack or to overcome system barriers with the intend to gain unauthorized access. The attack gets more effective when the attacker exploits a specific vulnerability of the target.

Mobile phones are particularly threatened by attacks. Their limited hardware resources allow for easy DoS disruptions, and the local system protection is less mature as compared to notebooks or PCs. However, it still remains an open question whether typical adversaries conduct context-specific port scans for intrusions.

In this paper, we report on a mobile honeypot, which collects external requests to an end device connected by a mobile operator. We compare the measurements of more than one month with the logs of four probes that are attached to different (non-mobile) ISPs. Our analysis reveals that a mobile device on average suffers from the same amount of attacks as a home network device. However, the distribution of attacks across autonomous systems, as well as the number of attackers from the same ASes is significantly more pronounced and could indicate Internet regions that specialize on attacking mobiles.

Current studies focused on the identification of attacker-friendly ASes [3] or the network-level behaviour of spammers [2] but did not differentiate between mobile and stationary access. Work in the field of mobile honeypots [1] deals with the secure *implementation* of honeypots on smartphones. We are not aware of any attack analysis based on a mobile honeypot.

Copyright is held by the author/owner(s).
SIGCOMM'12, August 13–17, 2012, Helsinki, Finland.
ACM 978-1-4503-1419-0/12/08.

2. MOBILE HONEYPOT

2.1 Background & Design

A honeypot is a trap for collecting data from unauthorized system access—in this analysis via IP—initiated by remote parties. Its intention is to learn about the nature and the characteristics of attacks. The novel term *mobile* honeypot is not yet well-defined. It is used for a probe that either (a) resides on a mobile device, (b) is running on a mobile operating system, or (c) is operated in the network of mobile devices. We argue that remote attacks are bound to the network layer and moreover do not focus on specifics of mobile hardware, but solely target at the system level.

For the mobile honeypot, we build our subsequent analysis of monitoring attacks on a Linux-based system that is connected to a mobile operator network for two reasons: First, a major part of currently deployed smartphones runs Android, which makes this platform appropriate for observations. However, current OS fingerprinting tools such as Nmap or Xprobe do not reveal specifics of the Android OS, but only report about a Linux system. Second, this approach bears the advantage of fully compatible results across platforms, as we use the identical honeypot tools for both, the mobile and the fixed Internet domain.

2.2 Measurement Setup

We deploy a low-interaction server honeypot based on the standard tools Honeytrap and Dionaea, because we are interested in the statistical analysis of attacks and not in dedicated threats. This mainly concentrates on attacks from the 'background noise', but even more intricated threats require the establishment of a connection to the target.

The honeypot runs at four Linux hosts, each connected via a different network access type: One UMTS network, an open university network, a DSL home network, and a darknet. At each site, the honeypot listens at a single *public*

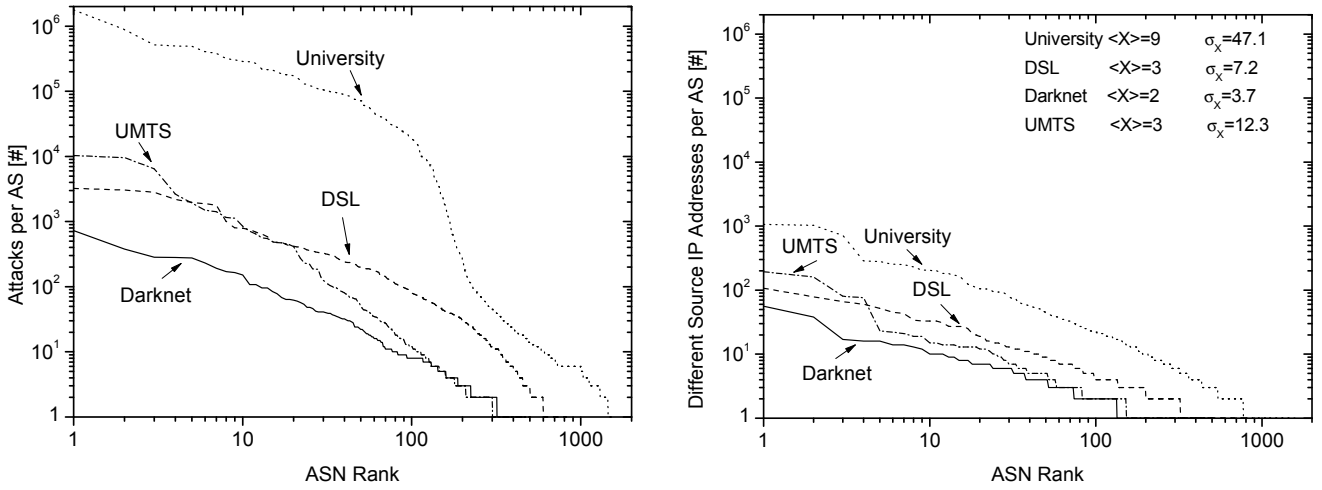

(a) Attacks per AS

(b) Attackers per AS

Figure 1: Comparing the amount of requests per autonomous system separately ranked per network access

IP address. Every external IP connect is considered as an attack, its source IP address is called the attacker.

3. PRELIMINARY RESULTS

The subsequent results represent measurements from mid December 2011 to the end of January 2012.

General observations Most of the attacks target at the (open and stable) university host and the minority at the darknet node. The DSL and UMTS probes measure on average 46 bogus requests per hour. Between 93% and 99% of the attacks use TCP with ports 22, 80, and 222 among the top-ten. Only 35 IP source addresses have been seen on all four monitors.

Attacks per AS We map the source IP addresses of the attackers to their origin autonomous systems (AS) and rank the ASes by the number of attacks. It is worth noting that the ranking is conducted separately for each type of network access. For example, AS 23650 and AS 8402 are ranked first in case of the UMTS and university network, respectively.

In general, most of the attacks have been initiated from the same small set of ASes (cf., Fig. 1(a)). The top-5 ASes mainly originate from China and Russia over all providers. It is clearly visible that these few ASes have a more pronounced impact on the mobile regime. In general, the distributions of attacks among ASes is of similar shape for the darknet and the home network, while the university encounters an enhanced and widened distribution of attacks. In contrast, attack statistics from the mobile network are significantly narrowed.

Attackers per AS We further quantify the number of different source IP addresses per AS to evaluate the number of distinct adversaries (cf., Fig. 1(b)). The AS ranking is conducted, again, separately. This measure balances individual intensities of attackers and consequently reduces the maximum values by some orders of magnitude. Nevertheless, the characteristic shape of the curves of Fig. 1(a) becomes even more significant. Attacks on the UMTS network remain significantly more concentrated to specific ASes than those of the fixed networks.

4. DISCUSSION

This paper presented first ideas towards a better understanding of the nature of mobile-specific attacks. Our preliminary results for an UMTS-connected device show that on the overall offenders show an intensity similar to home networks, whereas regions and originators of attacks are better pronounced and operate at higher intensity. This could be an indication of specific topological regimes that start to focus on mobile attacks.

We admit that any IP-level analysis is biased due to the problem of spoofed source addresses. However, identifying the spoofing of active networks at the end system is fuzzy, as well. Estimating the error is part of our future work. On the other hand, most of the addresses belong to the same IP range/origin AS (cf., also [2]) and thus do not affect our observations. In the future, we plan to perform more subtle correlation analysis of how specific groups of attackers behave with the aim to identify individual patterns of mobility-related aggressions. We will also analyse attacks per port in more detail. Due to limited statistics, though, these considerations will require a much longer range of observation.

Acknowledgements

This work is supported by the German BMBF within the project SKIMS (http://skims.realmv6.org).

5. REFERENCES

[1] Mulliner, C., Liebergeld, S., and Lange, M. Poster: HoneyDroid - Creating a Smartphone Honeypot, 2011. Poster at IEEE Security & Privacy.

[2] Ramachandran, A., and Feamster, N. Understanding the Network-level Behavior of Spammers. In *Proc. of ACM SIGCOMM'06* (New York, NY, USA, 2006), ACM, pp. 291–302.

[3] Shue, C. A., Kalafut, A. J., and Gupta, M. Abnormally Malicious Autonomous Systems and their Internet Connectivity. *IEEE/ACM Trans. Netw. 20*, 1 (2012), 220–230.