Can DNS-Based Blacklists Keep Up with Bots?

Anirudh Ramachandran, David Dagon, and Nick Feamster
College of Computing, Georgia Institute of Technology
{avr, dagon, feamster}@cc.gatech.edu

1. Introduction

Many Internet Service Providers (ISPs), anti-virus companies, and enterprise email vendors use Domain Name System-based Blackhole Lists (DNSBLs) to keep track of IP addresses that originate spam, so that future emails sent from these IP addresses can be rejected out-of-hand. DNSBL operators populate blocking lists based on complaints from recipients of spam, who report the IP address of the relay from which the unwanted email was sent. To be effective in blocking spam, information in the blacklist must have the following properties:

1. *Completeness.* The blacklist must contain a reasonable fraction of all spamming IP addresses.
2. *Responsiveness (*i.e.*, low response time).* We term the period of time between when a host first starts sending spam, and when it ultimately becomes listed the *response time*. The blacklist must have a low response time so that other recipients can subsequently block spam originating from the respective IP addresses.

Despite the widespread use of DNSBLs, to our knowledge there has not been a thorough evaluation of the effectiveness of blackhole lists in blocking spam. Although our previous work has briefly surveyed the *completeness* of DNSBLs for various types of spamming techniques (specifically, botnets, short-lived BGP routes) [5] at the time each piece of spam was received, neither this study nor any other that we are aware of have studied the *response time* of DNSBLs.

DNSBLs have proved to be an effective mechanism for blocking spam when spammers were less agile (*i.e.*, when they sent spam from a smaller number of open relays). Previous studies, however, have suggested that spammers are becoming increasingly agile, distributing the spam "workload" more widely across mail relays [4]. The recent rise of *botnets*—large collections of compromised machines under the control of a single controlling user—suggest that spam is being sent from an increasingly larger set of IP addresses, that the distribution of workload would have an even longer tail, and that each spamming host is relatively transient (recent work notes that most spamming bots are transient, at least from the perspective of a single domain [5]). This transience implies that, for DNSBLs to be effective at all, they *must* be responsive. This paper presents a preliminary evaluation of the responsiveness of blacklists for a specific set of spamming IP addresses that are known to come from a spamming botnet that spreads via the "Bobax" vulnerability [1].

CEAS 2006 - Third Conference on Email and Anti-Spam, July 27-28, 2006, Mountain View, California USA
.

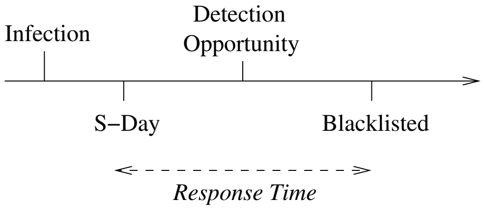

Figure 1: A conceptual view of a spamming host's life cycle.

2. A Model of Responsiveness

Figure 1 presents a model that shows four distinct phases of a host's life-cycle as part of a spamming botnet. Although we acknowledge that spam can certainly originate from uninfected machines (*e.g.*, rented machines for email marketing campaigns), we focus on studying the responsiveness of DNSBLs for blocking spam from infected machines (*i.e.*, likely botnet "zombies"), which send the vast majority of spam on the Internet today [6]. Initially, the host becomes a member of a spamming botnet; subsequently, the host begins to send spam (listed as "S-day" in Figure 1). After some time, the host's activities are detected, investigated, and recorded, which ultimately results in the host being blacklisted. Our goal is to determine not only completeness, but also *response time*, as shown in Figure 1, which is challenging given the lack of any ground truth: validating the time at which a host becomes infected or first sends spam is difficult, but we can still estimate lower bounds on response time.

3. Data Collection

Two datasets—a trace of DNSBL lookups and a trace of spamming botnet activity—allow us to establish a lower bound on response time: the difference between the time the host first becomes listed in the Spamhaus blacklist and the first time a host appeared after November 17, 2005 (*i.e.*, the time that we know the host has been infected). We have packet captures of DNSBL queries to a mirror of the Spamhaus blacklist [2] for November 29 and 30, 2005. This mirror sees approximately $1/17$ of all Spamhaus queries, most of which originated from hosts in the south-eastern United States (where the mirror is located). The domains being queried, of course, represent the entire population of spamming hosts. To derive some ground truth for hosts that are known to be spamming bots, we “hijacked” the authoritative DNS server for the domain hosting the command and control of the botnet and redirected queries for this domain to a machine at a large campus network, as in previous work [3]. We monitored Bobax drones [1], whose sole purpose is to send spam, and observed 2,042,991 distinct IP addresses[1] over the course of around 46 days. We observed 81,950 DNSBL queries for 4,295 of these hosts in our Spamhaus trace.

We note a few limitations of our techniques and describe some mitigating factors. First, our Spamhaus mirror packet captures does not reflect all DNSBL lookups, but because the mirror serves an area that includes many large ISPs, (*e.g.*, Cox, Earthlink, BellSouth), we believe our sample is representative. Second, a host may be blacklisted for a reason that is unrelated to being the member of the Bobax botnet that we observed (*e.g.*, some hosts may have multiple infections). Because Bobax drones *only* send spam, we believe that it is reasonable to assume that IP addresses that appear in both the Spamhaus blacklist and our botnet trace were listed because of their activities as part of a spamming botnet. Third, we cannot determine with certainty the time a host first becomes infected, since a host may have been infected before the beginning of our botnet trace; still, our data still allows us to determine a *lower bound* on the response time. Finally, our measurements are specific to Spamhaus; while other DNSBLs may have different response times, we believe that studying the response times for Spamhaus are still useful, given its widespread use.

4. Preliminary Results

In this section, we summarize our preliminary results on DNSBL response time. We perform a joint analysis on the Bobax and Spamhaus traces described in Section 3 to study the following four questions: (1) What is the completeness of the Spamhaus DNSBL? (2) What is the responsiveness of the Spamhaus DNSBL? (3) How many distinct domains are targeted by a spamming bot before it is blacklisted? (4) Does the frequency of spam from a particular IP address change after it is blacklisted (*i.e.*, what happens to spamming bot behavior after response)?

To answer these questions, we analyze the Bobax traces to determine the first time we see each Bobax host in our traces. We then use the queries at our Spamhaus mirror both as “backscatter” traffic to indicate which clients are receiving spam from a particular spamming botnet and as an indication of whether a spamming host is listed or not. We can then calculate a lower bound the response time of the DNSBL by subtracting the time that we first saw activity from the spamming bot to the time when a Spamhaus response first indicates that the IP in question has been blacklisted.

Of 4,295 Bobax IP addresses that were queried at our mirror, only 255 were blacklisted, even after 46 days of continual activity (earlier studies have also noted that IP-based blacklists are often incomplete [4, 5]). Of the 88 IP addresses that became listed during the two days of our DNSBL trace, however, 34 were listed after just a single detection opportunity, suggesting that the relatively small number of IP addresses that are detected often have a low response time. The spamming behavior of these hosts did not change significantly after they were listed at Spamhaus.

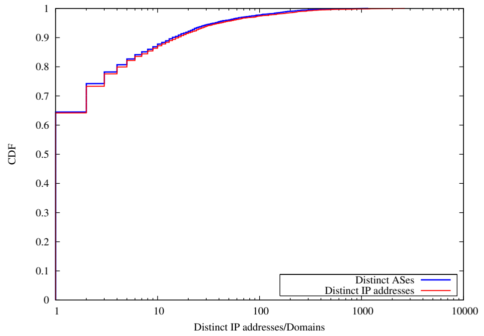

Figure 2: The number of distinct domains looking up each Bobax host

Figure 2 shows the number of distinct IP addresses and distinct Autonomous Systems (ASes) that queried our mirror for a specific IP address present in the Bobax trace. Over 60% of IP addresses were looked up by just one domain (AS), which suggests that most Bobax bots are low-volume and spam very few domains (effectively decreasing their chances of being reported to the blacklist as a potential spam originator). Around 10% of bots appear to generate lookups from a large number of distinct domains, which is cause for concern, since only about 5% of all bot IP addresses are ever blacklisted at Spamhaus.

5. Ongoing and Future Work

In our ongoing work, we are studying the above questions over longer durations. We intend to perform a more extensive study of how the spamming patterns of bots change (or do not change) after the IP addresses have been blacklisted (Question #4 above). A drop in spamming activity after blacklisting would suggest that botmasters are performing counter-intelligence on blacklists to determine whether various bots have been blacklisted. We intend to further analyze the Spamhaus traces for evidence of counter-intelligence.

REFERENCES

[1] Bobax trojan analysis. `http://www.lurhq.com/bobax.html`, March 2005.

[2] Spamhaus, 2006. `http://www.spamhaus.org/`.

[3] D. Dagon, C. Zou, and W. Lee. Modeling botnet propagation using time zones. In *Proceedings of the 13th Annual Network and Distributed System Security Symposium (NDSS ’06)*, 2006.

[4] J. Jung and E. Sit. An Empirical Study of Spam Traffic and the Use of DNS Black Lists. In *Proc. ACM SIGCOMM Internet Measurement Conference*, pages 370–375, Taormina, Sicily, Italy, Oct. 2004.

[5] A. Ramachandran and N. Feamster. Understanding the Network-Level Behavior of Spammers. Georgia Tech TR GT-CSS-2006-001.

[6] ZDNet Security News. Most spam genrated by botnets, expert says. `http://news.zdnet.co.uk/internet/security/0,39020375,39167561,00.htm`, Sept. 2004.

[1] This includes machines which are counted more than once due to factors such as DHCP churn. The actual strength of the botnet is estimated at around 100,000 hosts.