Dead Phish: An Examination of Deactivated Phishing Sites

Craig A. Shue and Erik M. Ferragut
Cyberspace Sciences and Information Intelligence Research Group
Oak Ridge National Laboratory
{shueca, ferragutem}@ornl.gov

ABSTRACT

Efforts to combat phishing and fraud online often center around filtering the phishing messages and disabling phishing Web sites to prevent users from being deceived. Two potential approaches to disabling a phishing site are (1) to eliminate the required DNS records to reach the site and (2) to remove the site from the machine itself. While previous work has focused on DNS take-down efforts, we focus on determining how long a phishing site remains on a machine after the DNS records have been removed. We find that on the day a site is reported, as many as 56% of phishing sites remain present on the hosting machines even after the DNS records have been removed. While many of these sites are removed within a few days, the DNS caching behavior at ISP resolvers may preserve the phishing site accessibility until the phishing site itself is completely removed.

1. INTRODUCTION

Phishing is a significant problem on the Web, with an estimated $3 billion lost in 2007 due to phishing attacks [1]. In a phishing attack, a phisher impersonates a legitimate institution, such as a bank or government, and contacts a victim to request sensitive information. Typically, the phisher requests that the user submit this sensitive information to a phishing Web site, which mimics a legitimate site. To conceal their identities and evade prosecution, phishers often use compromised or "bot" machines to send their phishing messages and host the phishing Web sites.

The impersonated institution suffers each time one of their customers is phished. These institutions leverage contractors and industry consortiums to aggressively fight phishing. They solicit users to report phishing and investigate any reported phishing sites. Once they find a phishing site, these institutions can take two approaches: 1) disable the DNS infrastructure used to access the site or 2) attempt to have the machine hosting the site cleaned. The first approach may be the fastest: domain registrars are familiar with phishing and removing DNS records can be done quickly. Once the records are removed, many users will be unable to resolve the host names associated with the phishing sites, thwarting the attack. However, the second approach is also important: if anti-phishing institutions only target DNS record removal, phishers can rapidly register new domains and reuse the same machines to host their phishing sites. Further, any organization or ISP that had a user visit the site before the records were pulled may have the DNS records cached, making the site accessible to all of their users until the records expire. Accordingly, a full solution requires both approaches: quick DNS take-downs to deal with the on-going attack, and end-host cleaning and patching to prevent the machine's involvement in subsequent phishing campaigns or from affecting users at sites with cached records.

While prior work has studied the success of DNS take-downs, no work has examined what happens to phishing Web servers after the take-downs have occurred. In this paper, we examine whether these sites are deleted quickly after the corresponding DNS records are removed. Such metrics can help institutions determine the effectiveness of their take-down efforts in different scenarios.

In this work, we collect URLs from phishing reporting feeds, quickly resolve the host names of these URLs to IP addresses, and, on a daily basis, connect to the original IP address and request the phishing page. Additionally, after each connection, we perform a DNS resolution on the host name and record the results. From this process, we can determine whether a phishing site is accessible on the hosting machine itself and the status of the DNS records, allowing us to see how DNS take-downs affect the site's presence on the hosting machine.

In this study, we find:

- The DNS records required to reach phishing sites are being removed quickly after the phishing site is reported.
- The phishing sites often remain on the hosting infrastructure even after the DNS records are removed, posing a risk to users at an organization with the DNS records cached.
- These sites quickly become unavailable after the DNS records are removed, indicating that anti-phishing forces are also removing the sites.
- Less than 20% of reported phishing URLs appear to be using compromised systems or Web hosting providers.

The rest of this paper is organized as follows. In Section 2, we provide background and related work. In Section 3, we describe our data collection efforts. In Section 4, we analyze the phishing infrastructure trends. We conclude in Section 5.

2. BACKGROUND AND RELATED WORK

Motivated by the financial impact and trends associated with phishing [2], researchers have studied phishing from a variety of angles, including the messages sent by phishers, the role of social context in phishing effectiveness, the effectiveness of filtering phishing messages, the effectiveness of blacklists, and the infrastructure used by phishers. With this broad background of work, we highlight some major works in the area.

A phishing campaign begins with a phisher searching for a victim. Prince *et al.* [3] found that phishers often use Web crawlers to obtain email addresses. Shue *et al.* [4] confirmed that Web crawlers are frequently used to harvest addresses but also found that some unscrupulous Web sites provide their users' email addresses to spammers. Upon finding victims, the phishers send unsolicited messages to the victims to lure them to phishing sites, often using botnets [5]. Jagatic *et al.* [6] found that phishers that provide context to the user, such as impersonating a victim's friend, can dramatically increases the user response rate for a phishing attack.

Some organizations have sought to fight phishing at its source: they seek to disable the infrastructure used by phishers to stop attacks. Some recent ISP de-peerings, including Atrivo [7], McColo [8], and Pricewert [9], show that operators have had some success in shutting down ISPs that facilitate malicious activity. Another approach commonly used by institutions targeted by phishing is simply to disable the DNS records used to reach a phishing site. Moore *et al.* [10] examined the resilience of phishing sites to these take-down efforts by examining their provisioning approaches, including fast flux. In monitoring the phishing sites, they probed the site several times a day and noted any changes in content from the site, allowing them to detect when a phishing site is removed. In later work, Moore *et al.* [11] examined trends relating spam mails for phishing sites and the take-down efforts between the sites. They found that spam continues to be sent for phishing sites over a week after the phishing site first appeared, yet the bulk of spam for the site is sent while the site is still alive. Our work augments these studies by looking directly at the machines hosting phishing sites themselves, allowing us to continue monitoring the phishing sites after their DNS records were removed. The Moore studies focus on when the sites are first unreachable, either due to DNS errors or due to server errors. In these studies, the authors acknowledge some sites may persist for some users because of DNS record caching. In our analysis, we can determine when a site is fully disabled for all users.

3. DATA COLLECTION AND METHODOLOGY

We aim to investigate the availability of phishing sites and their DNS records to determine how DNS take-down efforts affect these Web sites[1]. To do so, we require real-time feeds of reported sites. The Anti-Phishing Working Group (APWG) [12] and PhishTank [13] have large feeds of phishing site URLs. The APWG granted us access to their data feed while PhishTank makes their data feed publicly accessible. We use these feeds to create our `APWG` and `PhishTank` data sets, respectively. Each hour, we extract host names from the URLs currently in the feed, and perform DNS resolutions on those host names to get the IP addresses associated with the host names. We collected this data from October 1, 2009 to November 30, 2009, yielding 61 days of input.

With our list of URLs and associated IP addresses, we probe each of the IP addresses by directly connecting to each IP address and manually issuing an HTTP request for the full URL associated with that address. In doing so, we bypass the DNS, allowing us to access a site even after the DNS records for the site are deactivated. This allows us to observe phishing sites after they are "dead" or inaccessible to most users. We record the results of each connection attempt, including the page content if it is available. We performed each probe from a single source machine and encountered issues with DNS resolutions and connection attempts hanging for several minutes before the resolution or connection timed out. The phishing sites were particularly unstable after being reported, forcing us to set a short time-out value (2 seconds) and preventing us from retrying failed probes. Further, upon encountering a time-out, we excluded any subsequent connections to the same IP address for other URLs in the same input file. While these constraints were necessary for our collection, some sites may be incorrectly reported as unavailable simply due to network congestion and packet loss or delay. When selecting out time-out value, we consulted the work by Jung *et al.* [14], which shows that between 70% to 90% of resolutions complete in 1 second or less. By doubling this value, we hoped to accommodate most network latency in resolution requests and connection attempts.

Immediately after each connection attempt, we perform a DNS lookup on the URL's host name using the `gethostbyname` system call and record the results of this lookup and any IP addresses returned. From this, we can detect if the DNS records for a site have changed or been removed. We perform the first connection attempt within two hours of the URL being listed in the feed and repeat the lookup process for each URL for each day in our data sets on a daily basis from October 1, 2009 to January 27, 2010. Accordingly, for each of the 61 input days, we have 58 days of probe results, allowing us to see the availability of each phishing site for almost two months after first being reported to a phishing feed.

3.1 Characterization of the Feeds

In our collection period, we had 543,549 entries with 29,600 hosts in 17,130 domains. However, these entries were originally associated with only 19,063 IP addresses. Upon manual inspection, many of the domains appeared to be randomly generated, indicating they could be registered by phishers. Many of these domains could be unregistered if found to be involved in phishing (though phishers could later reregister them). To determine if this is happening, we examined the zone files of eight popular generic top level domains (TLDs): .ASIA, .BIZ, .COM, .INFO, .MOBI, .NAME, .NET, and .ORG. We examined these zone files on February 1, 2010 to determine whether the domains were still registered two months after they were listed in phishing feeds. We found that 12,279 phishing domains were registered under these TLDs. However, only 10,957 (roughly 89.2%) were still registered at the point we re-examined them. This indicates that some domains were unregistered

[1]We specifically focus on phishing sites in this work. Other Web sites advertised via spam may have different properties and are beyond the scope of this work.

or allowed to expire after they were involved in phishing.

4. ANALYSIS AND RESULTS

We begin by analyzing our connection attempts to known phishing sites, categorizing any errors we encounter. We then examine the DNS records associated with a URL, including whether the records match. We then focus on sites after their DNS records have been removed and examine whether the phishing site remains available. Afterward, we examine the redirection behavior of the phishing sites. Finally, we examine the phishing site content to determine whether a phishing site is taken down using page content modifications.

4.1 Connecting to Known Phishing Machines

When attempting to connect to a Web server, regardless of whether it is a phishing site, a client can receive a number of different responses. An error can occur while attempting to access the machine, such as a connection refusal or no response, which eventually leads to a time-out. Even if a connection is successful, the Web server can return errors using HTTP status codes. Some errors refer to client-side errors (e.g. "Page not found" or "Forbidden") while others refer to server errors ("Internal server error" or "Bad gateway"). Even when accessing a URL is successful, requests can return either content, such as an HTML page or image, or a redirect message, indicating the client should consult another location to find the requested content. In analyzing the responses from a server, the type of error message is revealing: connection errors are independent of the site being accessed since they occur before the client specifies the host name or path of the URL it is requesting. Accordingly, these errors indicate that a machine is not operating a Web server on the given port for any phishing campaign. Client or server-side HTTP error codes may indicate that a particular phishing site is not available, yet does not preclude the existence of other phishing sites on that infrastructure.

In our analysis, we perform daily sweeps of the phishing sites listed on the 61 input days (Oct. 1 to Nov. 30, 2009). To be able to show aggregate results across input days, we group each result by the amount of time that has elapsed since the site was first reported rather than by calendar days. In the case that a URL appeared on multiple days, the first day listed was used and all subsequent duplicates were discarded. In Figure 1, we show the number of reported URLs that were inaccessible, either due to connection errors or due to client or server error codes. We see that the URLs in the `APWG` feed have about 44% unavailability even within hours of being reported. This may indicate that organizations reporting to the `APWG` feed aggressively pursue take-down efforts or report the site only after the take-down efforts have begun. In general, total errors in the `APWG` feed reach about 96% and then hold steady. In the `PhishTank` feed, the rate of unavailability starts at about 17%, climbs quickly to about 80%, then slowly increases to about 94% at the end of the collection period. In general, URLs reported to the `PhishTank` feed do not have their DNS records pulled quite as quickly or as often as the `APWG` feed, which may be a side-effect of the `APWG` feed belonging to a closed community.

Requests that do not result in errors can successfully return content or provide a redirect. In Figure 2, we show the percentage of attempts that successfully return content. As

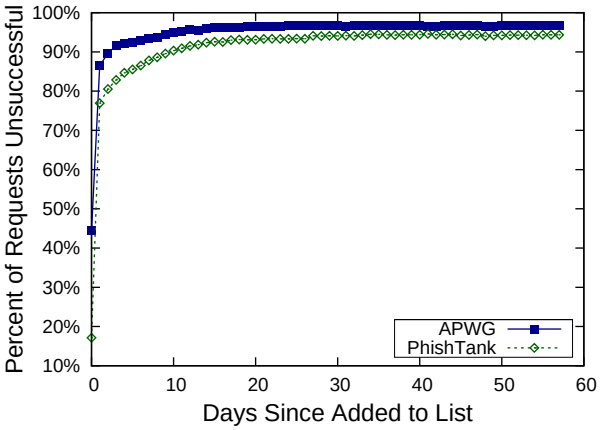

Figure 1: Rate of errors in retrieving URLs

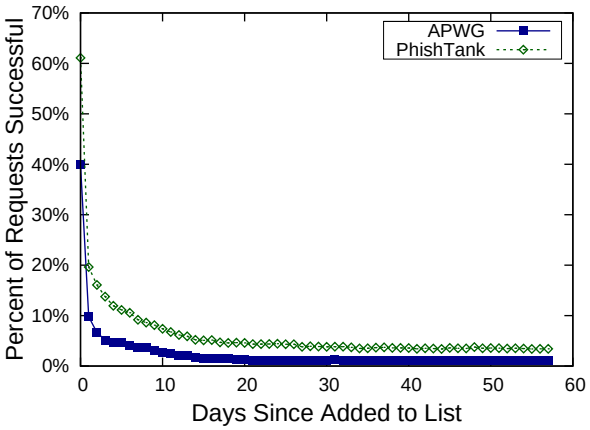

Figure 2: Rate of successful retrieval of URLs

suggested by the error rate, the greatest success for URLs in the `APWG` feed (at about 40%) comes on the day the page is first listed. Two weeks later, the success rate was down to 1.7% and remained in the 1.0% to 1.5% range for the remainder of the collection period. For the `PhishTank` results, the success rate starts around 61% and quickly drops to around 13% before slowly lowering to around 3.5%. As indicated in our data collection methodology, these results are a conservative estimate of phishing site availability as we were unable to retry connection attempts, making our analysis sensitive to network congestion. At the same time, this analysis only indicates whether a page was successfully obtained, but provides no guarantee that the page is the same as that used in the actual phishing campaign; however, in Section 4.6, we find that few pages have content changes.

Rather than providing a definitive success or failure, a Web server may return a redirect status code and provide a location where the client can find the intended content. In Figure 3, we show these results as a percentage of URLs. Like the successful retrievals, redirections are highest on the day the site is reported and drop sharply, settling to between 2-3% of URLs after the first few days. We examine these redirects and their destinations in greater detail in Section 4.4.

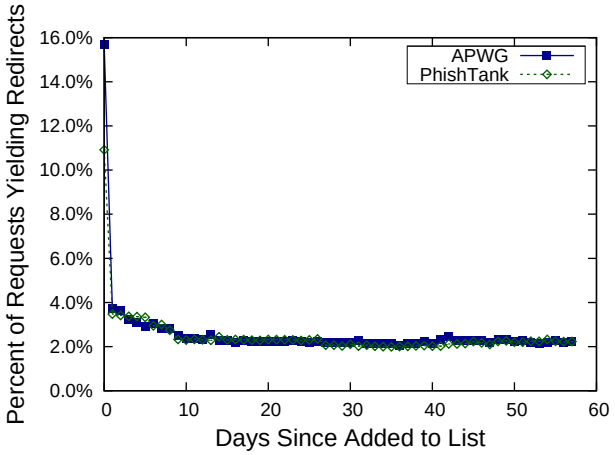

Figure 3: Rate of redirection when retrieving URLs

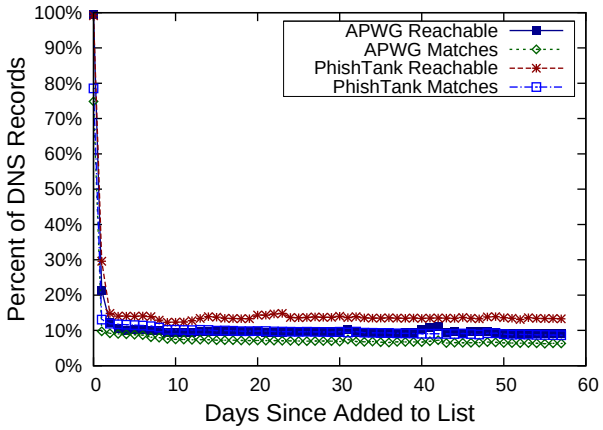

Figure 4: DNS reachability and matching

These results show that phishing sites are changing drastically after being reported. Most URLs result in a connection error state immediately after being reported with many of the remaining pages resulting in page errors. These results show that anti-phishing forces are being effective at combatting these sites.

4.2 DNS Availability of Phishing Sites

Each time we attempt to connect to a machine, we check the DNS records associated with the original host name from the machine's URL. This allows us to detect when DNS records have been removed. In Figure 4, we show the percentage of URLs with functional DNS records as "`APWG` Reachable" and "`PhishTank` Reachable." In both feeds, we see a rapid decrease in record availability during the first couple days after the site is reported, plunging from about 99% availability on the day the URLs were reported to under 15% availability on day 3. In the `PhishTank` feed, the records stay between 13% to 14% for the remainder of the collection period. The results of the `APWG` feed stabilize between 9% and 10% availability. The DNS record removal appears most likely to occur soon after the site is reported with little change after three days since the site was reported. Our results generally confirm the trends reported by Moore *et al.* [10].

While the DNS records for a host name may be available, they do not necessarily continue to point to the same host. ISPs and registrars could instead point DNS records to other sites in order to educate the user about phishing attacks [15] or phishers could be using DNS fast flux or other techniques, causing the records to not reflect the original machine. Accordingly, we examine these records more closely to determine whether these DNS results include the original machine. Naturally, if a record is unavailable, it cannot match the initial IP address; these results are strictly lower than the results we found in the reachability results. However, for clarity, we provide these results as a percentage of the total number of URLs in Figure 4. These results show a vertical shift downward from the simple DNS availability results. In particular, in the `PhishTank` and `APWG` feeds, we saw a 20-25% decrease in matching on the original day the URL was reported, showing that the DNS records changed within hours of being first reported. On the day after being reported, the matching decrease was about 12% for `APWG` and about 17% for the `PhishTank` feed. On subsequent days, the difference between total records available and those that matched was about 2.6% for `APWG` and about 3.9% for the `PhishTank` feed.

With the difference between available and matching DNS records, we examined whether the IP addresses being returned were co-located with the original IP address, which may be common for large hosting providers, or whether the IP addresses were largely from remote parts of the network, which is a common indication of fast flux in botnets. To make such a distinction, we required some notion of co-located hosts. We leverage inter-domain routing information for this purpose. When participating in inter-domain routing with BGP, networks controlled by a single administrative operator are grouped together as an autonomous system (AS). Additionally, the individual networks that form an AS are each aggregated into IP prefixes to reduce memory requirements in inter-domain routers. We leverage this routing data, which we obtain from the Route Views Project [16]. We then loaded this routing table information into trie data structures and used longest prefix matching on the IP address to determine the AS and prefix associated with each IP address. Accordingly, we could determine whether the original IP address 1) matched the IP address in subsequent lookups, 2) belonged to the same IP prefix as subsequent lookups, 3) originated in the same AS as subsequent lookups, or 4) completely differed from subsequent lookups.

We tracked 249,608 unique records. As our earlier results show, we had a high DNS failure rate, with 83% of records failing at least 85% of the times they were queried, with 171,501 (69%) failing in every lookup after the original. For simplicity, we exclude any failures and track the results of successful DNS queries. In Figure 5, we show the percentage of records that match the original IP address. We examine the IP addresses that do not match the IP address, but match the prefix and AS, but these generally below 10% each day and we have excluded them for readability. We next plot the records that do not match the IP address, prefix, or originating AS in Figure 5. The results fluctuate greatly on some days, which may be a result of smaller record sample sizes on the given day. The results show that of host names that maintain a consistent match type throughout the sweep

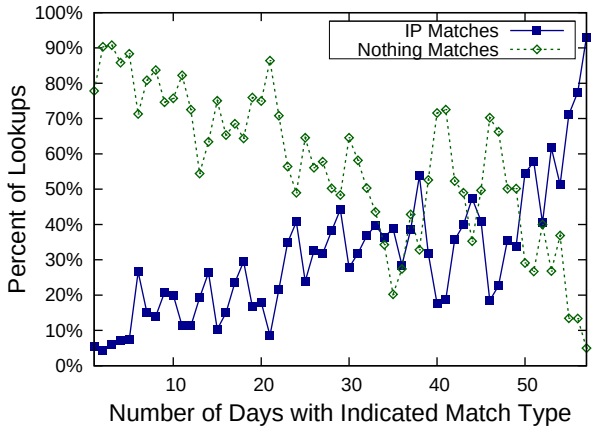

Figure 5: DNS records that match IP address and those that have different addresses, prefix, and ASN

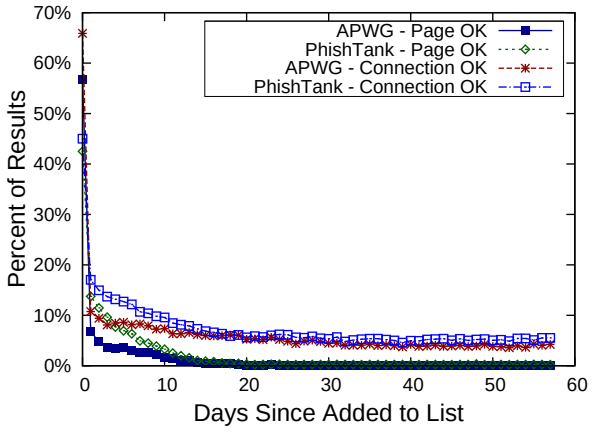

Figure 6: Machine status after DNS records removed

days, over 90% are exact IP address matches. In general, as the instability of a host name's matching status decreases, the rate at which the host resolves to a completely different IP address increases. The lack of stability in non-matching host resolutions suggests that the DNS records are not being used for fast-flux attacks. Further, hosting providers, registrars, and ISPs do not seem to be altering DNS records on a large scale to thwart attacks and merely correct the problem on the host or pull the DNS record in its entirety.

4.3 Availability After DNS Take-Down

DNS take-down efforts are a common practice in anti-phishing efforts. However, such efforts may not solve the underlying problem. In particular, if institutions only perform DNS take-downs and do not contact the ISPs providing connectivity to the actual phishing machines, these machines may not be cleaned and can be reused in later phishing campaigns. In this section, we characterize the extent to which this is occurring.

We begin by tying the DNS status of the host names in the reported phishing URLs. For each host name that returns a DNS error, we examine whether we were successful at accessing that same machine on the same day. We show these results in Figure 6. In particular, we provide two results for both the `APWG` and `PhishTank` data sets. The first, which we label "Page OK," indicates that a connection to the system was successful and we obtained a successful response from the server (HTTP status code 200). The second, labeled "Connection OK," indicates that a server accepted our connection, but it also includes redirects and page errors. The "Connection OK" results hint at two factors. It captures compromised machines whose campaign-specific phishing page was removed by the phisher when the DNS records were eliminated. At the same time, this result also captures Web hosting providers that have removed a phishing site.

In these results, we see that both feeds have machines that return Web pages even after DNS record takedowns. In particular, on the day the URL is reported but after the DNS records are removed, 56.81% of `APWG` feed sites and 42.5% of `PhishTank` feed sites still have a fully functional Web site providing the phishing page contents. This plunges to 6.8% and 13.8% for `APWG` and `PhishTank` on the following day. These results show that DNS record removals tend to be an early phishing site take-down effort with actual site removal used as a follow-up. While a practical measure, with modern DNS caching behavior, the removal of these DNS records will have no impact if the DNS record is cached at a site or ISP's DNS resolver. As a result, a single user accessing the phishing site while the records are still active may cause the record to be cached and place all other users behind the resolver at risk until the site is deactivated. Savvy phishers may begin taking advantage of this caching behavior and set longer TTL values to keep their sites online as long as possible. Unfortunately, modern DNS does not provide a means for a site to send purge messages to caching resolvers in the case of a DNS take-down. While such messages could be added to DNS, the large installation base of DNS servers and resolvers would likely prolong deployment.

We also look at the Web server availability after the DNS records are removed. If a phisher uses a hosting provider or compromises a legitimate site, they may use their own DNS records to reach the machine's IP address. Accordingly, if the DNS records are pulled and the connection yields a page error, it could indicate the legitimate server remains available after the phishing site is removed. In these cases, we would expect the connection to remain stable even weeks after the site was reported. Like with page success rates, we see that connections are successful to a site over 40% of the time on the day it is reported with a sharp decrease the following day. However, in both feeds, the connection success rate drops to between 3.5% and 6% and remains stable. Accordingly, we can infer that phisher rarely use independent DNS and compromised or hosting provider servers to host their sites. Phishers could still use the legitimate site's own DNS records; however, we show that in Figure 4, only about 9% to 14% of host names return valid DNS records in the long-term. This would indicate the legitimate site's own DNS records would be unavailable as well. Combined, these results suggest that a maximum of 20% of URLs point to sites using hacked or hosting provider Web sites.

These results suggest that the phishing pages on these machines are eventually being removed. However, in many cases, the Web servers continue to operate on the machines. It is unclear whether the page and Web server removals are

the result of hosting provider or ISP involvement or the simple reprovisioning of the compromised machines by phishers. To investigate that question, we considered the null hypothesis that the machine connection results were independent of the DNS server results. We focused on the data where the site is reachable at first and eventually becomes unreachable, but does not again become reachable[2]. Among this restricted set, we estimated the multinomial distributions of machine responses from days up to and including the last day the site is reachable and from days starting with the first day the site is unreachable. We compared reachable sites' responses to that of unreachable data sites' machine responses by using the null hypothesis. We suspected that these sites came from a multinomial distribution. We used a Chi-squared statistic with six degrees of freedom (DOF). In the `APWG` set, the statistic was $977,992$ with 6 DOF, well beyond the threshold of 22.45 required for a 0.001 p-value. In the `PhishTank` set, the statistic was $234,832$ compared to the threshold of 20.51 for 5 DOF at a 0.001 significance level. (Two rare cases were combined, reducing DOF to 5.) We conclude that, in the aggregate, sites that are reachable via DNS have machine responses that differ from those that are not reachable.

In general, as sites are first reachable and later become unreachable, it is possible that the statistical significance follows from a date-dependent rather than reachability-dependent effect. To explore this, the same analysis was performed separately for each day, again referring to the number of days after the site was blacklisted. Computations of Chi-squared statistics resulted in statistically significant deviations between reachable and unreachable sites, even on a day-by-day basis and at an 0.001 significance level, excluding the first day which lacked data on unreachable sites. We conclude that the machine results for reachable sites are statistically significantly different from those for unreachable sites, even accounting for days since blacklisting.

In Table 1, we show the conditional probabilities for site availability given DNS reachability in the `APWG` data set. The key difference is that DNS available sites are far more likely to return OK or redirect than are the DNS unreachable sites. Also, connection errors, which are the most common result for DNS unreachable sites, are a third less likely for DNS available sites. The results for the `PhishTank` data set were similar and have been omitted for brevity.

Table 1: Machine response depends DNS availability (`APWG`)

Machine Response	Pr(Resp | Reach)	Pr(Resp | ¬Reach)	% Diff.
E-Connect	0.618	0.914	-32.4%
E-Page	0.118	0.064	+83.8%
E-Server	0.001	0.001	-9.06%
OK	0.152	0.014	+963%
Redirect	0.112	0.007	+1500%

From these results, we see that there is a strong correlation between sites being unavailable via DNS and decreased reachability, even after accounting for the amount of elapsed time. We discussed the matter with two leading phishing site take-down specialists. When performing a take-down, these groups notify both the ISP responsible for the machine and the DNS registrar at the same time. As a result, this relationship would not be causal in the case of administrative take-downs: the site take-down may simply take longer than the DNS record removal. However, this does not preclude phishers removing a site after it is no longer available. Unfortunately, we are unable to distinguish this case.

[2]Other cases exist where sites enter an error state and return to a reachable state. We cross-referenced these cases with the relevant zone files and found the domain NS records were intact. We were unable to determine why the name server returned an error in these cases.

4.4 Redirects

When we visited some phishing sites, they provided a redirection instead of an actual Web page. We were interested in determining whether these redirects were provided by the phisher or by network operators, as recommended by the APWG [15]. We further analyze whether these redirects changed or were consistent across the data collection period.

We recorded the redirects associated with a URL across each of the sweep days. We found many entries where there was no record of a redirect, due to a transient outage or machine unavailability. Excluding such unavailable events, we examined the number of unique redirect destinations associated with a given source URL. In general, redirect destinations were constant across sweep days. For 86% of our URLs resulting in redirects, the redirect destination was the same for each day we accessed the URL. For another 10%, the redirect went to two unique destinations. However, at the opposite end of the extreme, some sites (0.35%) redirected to a unique destination on each access attempt. When manually examining these redirects, it was clear they were encoding a dynamic session identifier in the URL.

Across all our data sets, we had $79,515$ unique redirection instances which led to $37,276$ unique destinations. Excluding the query string from these destinations, which often contain unique session identifiers, we find only $10,443$ unique redirect destinations[3].

Of the original $79,515$ unique redirection instances, 74 redirect back to themselves. If the DNS entry for such host names were cached, such as in the DNS pinning technique, these redirects would yield infinite redirection loops. In other cases, this may be used with fast-flux to dynamically redirect visitors to different machines hosting copies of the phishing site. We found that $44,319$ redirects are to a path on the same host, with $24,825$ unique paths across each reported entry ($5,306$ when omitting the query string). An additional $3,291$ redirects are to a different machine in the same DNS domain. Another $1,796$ redirects are completely dissimilar in their host name, yet have the same path. Finally, $30,035$ are completely dissimilar. However, across all instances, these completely dissimilar URLs yield only $8,820$ unique destinations ($3,553$ when omitting the query string).

By examining the URL in the redirects, we were sometimes able to infer the semantics behind the destination address. Rather than providing simple "page not found" errors, hosting providers often redirected the visitor to a customized Web page explaining the error. Some simply indicated the page was not found (e.g., "404.html") while others explicitly indicated the page had been removed for administrative reasons (e.g., "suspended-page.html"). In our analysis, we found $11,799$ unique instances of URLs indicating such errors. Of these, $2,748$ occurred on a redirect

[3]Of these, $18,886$ were unique relative-path redirects ($3,537$ when query strings were eliminated).

to the same machine (118 unique paths, excluding query strings). Another 8,098 instances occurred to completely dissimilar destinations (1,474 unique destinations, excluding query strings). Accordingly, about 41% of the unique redirects to a completely different destination were to a page indicating the original page was not available.

During our data collection, we were contacted by an administrator of a URL shortening service because we repeatedly probed a link reported as phishing. When the link was first reported, he temporarily changed the redirect to a financial page at a popular Web portal to prevent users from being victimized. While this approach thwarted the attack, it did not remove the site from the vulnerable sites list in anti-phishing reports, leading to difficulties with the hosting provider. He has since modified the service to return an HTTP 404 error to ensure automated processes can confirm phishing sites have been disabled. While this was the only such instance we encountered of self-reported atypical behavior by an administrator, we recognize that anti-phishing efforts can complicate systematic phishing studies.

4.5 Reuse of Phishing Infrastructure

When a machine is compromised, it may be reused for multiple phishing campaigns. Accordingly, it would appear in our feed in multiple entries. We excluded duplicates where the URL for an IP address is exactly the same for multiple days. In our input, we found 16,368 unique IP addresses in the `APWG` feed and 10,646 unique IP addresses in the `PhishTank` feed. Of these, about 73% of `APWG` IPs appeared only on a single day and 78% of `PhishTank` IPs appeared on only a single day. In general, IP addresses appeared on only a small number of days, though some IP addresses were present for over half of the days in our analysis period. From this, we can conclude that most phishers are not rapidly reusing compromised infrastructure, though some appear to be doing so. We show the full results in Table 2.

Table 2: Number of days an IP address appears in feeds

Days	`APWG`	`PhishTank`
1	72.87%	78.26%
2	13.48%	11.37%
3	5.19%	4.01%
4	2.58%	2.21%
5	1.60%	1.15%
6	0.92%	0.70%
7	0.66%	0.62%
8	0.50%	0.39%
9	0.41%	0.23%
10	0.37%	0.13%
11	0.29%	0.23%
12	0.18%	0.14%
13	0.09%	0.08%
14	0.08%	0.06%
15	0.12%	0.08%
16	0.07%	0.08%
17	0.07%	0.06%
18	0.05%	0.06%
19	0.05%	0.05%
20	0.02%	0.03%
>20	0.39%	0.09%

4.6 Content-Based Changes

We now examine whether take-down efforts alter the content of phishing pages or simply remove them. We analyze the content of the Web pages to determine how often the document changes significantly from the first retrieval. Ideally, we would like to compare the full document at each phishing URL for each day we accessed it. Unfortunately, these documents were frequently truncated due to the low socket time-out required to complete the study. Accordingly, our analysis focused on the headers for each file and compared the site values across days. We only examined days that had multiple successful page retrievals, resulting in 43,335 records.

From our analysis, we found that 54-56% of sites exactly matched on the 1,000 character header we used in the two data sets. We then looked at the snapshots that did not match and found that in over 99% of cases, the non- matching entries were solely due to one of the snapshots being truncated, not actual changes in the pages. These results suggest that administrative page substitutions are not a common practice to remove phishing pages. While these substitutions could occur before our first access attempt, it is unlikely that in practice all administrators would react quickly enough to alter the page before our first retrieval, which happens within hours of the page being listed.

5. CONCLUSION

We examined phishing sites from a host perspective, allowing us new insight into how phishers and anti-phishing groups affect machines used to host phishing infrastructure. We find that anti-phishing efforts quickly result in the removal of DNS records for most sites but that machine cleaning efforts generally take longer. This can leave users exposed to the attack if the phishing site's DNS records were cached at the user's organization's DNS resolver. However, after the DNS records are eliminated, either the phishers or anti-phishing forces are eventually removing the actual phishing sites.

In our analysis, we found that long cache times in phishing-related DNS records may keep a phishing site active at some networks while shorter cache times may provide phishers with the ability to load-balance a campaign. As fast flux detection improves, phishers may move to longer cache times to resolvers to evade detection and to lengthen the site availability.

6. REFERENCES

[1] A. Litan, "Phishing attacks escalate, morph and cause considerable damag," 2007. [Online]. Available: http://www.gartner.com/DisplayDocument?ref=g_search&id=562912&subref=simplesearch

[2] Z. Ramzan and C. Wüest, "Phishing attacks: Analyzing trends in 2006," in *Conference on Email and Anti-Spam (CEAS)*, 2007.

[3] M. B. Prince, L. Holloway, E. Langheinrich, B. M. Dahl, and A. M. Keller, "Understanding how spammers steal your e-mail address: An analysis of the first six months of data from Project Honey Pot," in *Conference on Email and Anti-Spam (CEAS)*, 2005.

[4] C. Shue, M. Gupta, J. Lubia, C. Kong, and A. Yuksel, "Spamology: A study of spam origins," 2009.

[5] Y. Xie, F. Yu, K. Achan, R. Panigrahy, G. Hulten, and I. Osipkov, "Spamming botnets: Signatures and characteristics," in *ACM SIGCOMM*, 2008.

[6] T. Jagatic, N. Johnson, M. Jakobsson, and F. Menczer, "Social phishing," *Communications of the ACM*, vol. 50, no. 10, October 2007.

[7] J. Hruska, "Bad seed ISP Atrivo cut off from rest of the Internet," 2008. [Online]. Available: http://arstechnica.com/security/news/2008/09/bad-seed-isp-atrivo-cut-off-from-rest-of-the-internet.ars

[8] B. Krebs, "Major source of online scams and spams knocked offline," http://voices.washingtonpost.com/securityfix/2008/11/major_source_of_online_scams_a.html, 2008.

[9] J. Cheng, "FTC forces hive of scum and villainy ISP offline," 2009. [Online]. Available: http://arstechnica.com/tech-policy/news/2009/06/ftc-forces-hive-of-scum-and-villainy-isp-offline.ars

[10] T. Moore and R. Clayton, "Examining the impact of Website take-down on phishing," in *APWG eCrime Researchers Summit*, 2007.

[11] T. Moore, R. Clayton, and H. Stern, "Temporal correlations between spam and phishing Websites," in *USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET)*, 2009.

[12] APWG, "Anti-phishing working group," http://www.antiphishing.org/.

[13] OpenDNS, "PhishTank," http://www.phishtank.com/.

[14] J. Jung, E. Sit, H. Balakrishnan, and R. Morris, "DNS performance and the effectiveness of caching," in *ACM SIGCOMM Internet Measurement Workshop*, 2001.

[15] L. Mather, "APWG CMU phishing education landing page program." [Online]. Available: http://education.apwg.org/r/about.html

[16] University of Oregon Advanced Network Technology Center, "Route Views project," http://www.routeviews.org/.