Metrics of collateral damage from blacklisting of domains exploited by “phishing” operations.

John Nagle
SiteTruth
February, 2008

Abstract

The abuse of exploitable security holes in major web sites by “phishing” attacks is a known problem. We have developed some simple metrics for measuring the problem, and have discovered the problem to be both smaller-scale than previously thought, and fixable.

Introduction

Phishing operations find it useful to place their sites, when possible, under prominent web domains. This conceals the identity of the phishing operator and allows them to exploit the reputation of a major domain to make their site appear legitimate. This attack can deceive both web users and automated spam filters, some of which detect phishing spams based on the URLs found in e-mail messages.

As part of our web site legitimacy rating system, we use existing lists of “phishing” URLs. We've chosen to take the hard line that any verified phishing URL anywhere in a domain results in the blacklisting of the entire second level domain. This is effective, but can result in collateral damage to vulnerable and innocent web sites. We thus needed to develop a means for measuring the impact of such blacklisting.

Approach

Our solution to the problem was to generate a list of “Major domains being exploited by active phishing scams”. This is not a list of sites being mimicked by phishing sites; rather, it is a list of major sites with vulnerabilities which make them useful in mounting attacks against third parties. This list is compiled by comparing the PhishTank database of “phishing URLs” to the English-language portion of the Open Directory, which contains most English-language web sites of any significance. The intersection of the second level domains[1] in these two sets is our list of exploited domains. PhishTank contains roughly 10,000 “active and online” phishing URLs at any one time, and the Open Directory contains about 1.7 million “major” web sites.

Results

As of early February 2008, only 45 domains are on this list.

When we started producing this list, in November 2007, there were 174 domains on the list, including some very well known names. Efforts to reduce the problem were made. We published the list, which is available at **http://www.sitetruth.com/reports** and is updated every three hours. We took other steps, ranging from privately contacting web site operators to publicity in the trade press. The Anti-Phishing Working Group provided some assistance in reaching key contacts at major sites. PhishTank's operators cooperated in more quickly removing sites from their “active and online” list once the problem had been cleared up. Some exploitable and exploited features on major sites were fixed as a result of this effort, and the list is now considerably shorter.

Characterization of exploited domains

These exploited domains fall into a few standard categories.

1. **Consumer Internet service providers**, inadvertently providing connectivity for corrupted computers. Some providers are more successful at dealing with these problems than others, and the list reflects this.
2. **Web hosting services** hosting phishing sites. Most major hosting services are now rapidly terminating accounts used to support phishing operations, and our lists reflects this. Hosting operators appear briefly on the list and are removed once the attack has been repulsed.
3. **Sites with "open redirectors",** used by phishing sites to craft URLs which will get through spam filters. This was a major problem when we started, with some well known sites providing open redirectors. Almost all those holes have now been plugged.
4. **Sites which have had a server break-in** by a phishing operator. These are usually fixed quickly; some site operators simply need to be told that they have a problem.
5. **Redirection services** (notably "tinyurl.com" and "notlong.com"). These services now pro actively terminate their redirection service for identified phishing sites, but are so easy to exploit that they usually have a few exploits outstanding.

The Honeynet Project, in 2005, divided phishing attacks into three categories, corresponding roughly to categories 1, 2, and 4 above.[2] Categories 3 and 5 reflect advancements in attacks since then.

Category	> 60 days	30-60 days	< 30 days	Total
ISP	5	4	3	12
Hosting service	2	2	8	12
Open redirect	0	2	1	3
Server break-in	3	0	12	15
Redirect service	0	1	2	3
Total	10	9	26	45

(Data as of 2008-02-10 0600 hrs PST)

These numbers are surprisingly small. The data indicates that most attacked sites are clearing their problems within 30 days. Only ten domains have been on the list for more than 60 days. Most server break-in problems are cleared within 30 days, often sooner. This is consistent with Anti-Phishing Working Group statistics indicating that the average life of a phishing site was 3.8 days as of June 2007.[3] Note, though, that the lifetime on this list is the time that the domain was observed to have any active exploit. It is not the lifetime of individual exploits. This is an indication that sites are fixing their underlying problems, rather than simply blocking previously identified attacks

When we began this effort, the conventional wisdom was that a much larger number of major sites were being exploited, and thus this was not a problem which could be solved by focusing attention on specific sites. This proved not to be the case. Only a few sites in each category are affected.

We note that, for each category, there are many sites competitive to the ones on the list which are not themselves on the list. Thus, there is no valid business case that a business must have a vulnerability which puts them on the list in order to perform their function.

Data quality

For each reported phishing URL, there are presumably a much larger number of unreported but similar URLs. We thus do not consider the number of reports per second level domain to be useful. One confirmed report is sufficient to indicate a vulnerability.

Additions and deletions to the PhishTank database are manual, and run behind actual events. We may add data sources based on “honeypots” to improve the timeliness of the data for real-time use.

Transparency

The means by which we develop this list is transparent. The data sources used are public, the reason behind each blacklisting event is fully disclosed, and an established process is in place for removing erroneous reports from PhishTank and thus from our list. The process outlined here does not require secrecy against the phishing attackers.

Further work

We have identified a choke point where a small amount of effort can be used to block an entire class of phishing attacks. It is worthwhile to search for other such choke points. At the opposite end of the scale from the major domains addressed here are the short-lived domains registered purely for phishing purposes. Those are being addressed in other efforts.[4]

Conclusion

Exploitation of major sites by phishing attacks is a problem which can be solved. The number of involved sites is small.

Our results indicate that the collateral damage from blacklisting entire second-level domains for a single phishing attack is acceptable. While some transient damage to legitimate sites does occur, very few sites incur a long-term impact. This is a price worth paying to stop this class of attack. Provided that the mechanism by which sites are chosen for blocking is transparent, as it is here, this approach is effective.

[1] We use the term “second level domain” to indicate a domain at the level at which registrars sell domains. This can be a true second-level domain, such as “**example.com**”, or a domain under a domain of a country code, such as “**example.co.uk**”.

[2] “Know your Enemy: Phishing -- Behind the Scenes of Phishing Attacks”, Watson et. al., The Honeynet Project & Research Alliance, (http://www.honeynet.org) May 2005.

[3] Phishing Activity Trends, Anti-Phishing Working Group, (http://www.antiphishing.org/reports/apwg_report_june_2007.pdf), June 2007.

[4] “ICANN Considers Plan to Stop ‘Domain Tasting’, London, Kirk, Jeremy, The New York Times, January 30, 2008.