Issues in Using DNS Whois Data for Phishing Site Take Down

The Anti-Phishing Working Group
DNS Policy Committee
(www.antiphishing.org)

May 2007

Summary

Given fundamental policy changes regarding accessibility of Whois data currently under consideration by ICANN, and the evolving environment surrounding the Whois system, the APWG DNS Subcommittee has produced this industrial advisory, comprised of a set of real-world case studies in which Whois data has been instrumental in neutralizing phish sites, to help ICANN comprehensively inform these policy deliberations. The intent is to better inform the broader ICANN community of the invaluable assistance the full range of Whois data provides in shutting down nearly 1,000 phishing sites per day (and climbing) at current rates. Each of these cases considered describes a specific event, but represents hundreds of analogous events that occur daily.

The report's sponsor, The Anti-Phishing Working Group (APWG) is the global pan-industrial and law enforcement association focused on eliminating fraud and identity theft that result from phishing, pharming and email spoofing of all types. The institution has a great number of members serving as security professionals who specialize in electronic crime detection and response for commercial enterprises. Over the past several years, these APWG members have shut down hundreds of thousands of phishing websites throughout the world. Counter to many people's perception, the vast majority of these phishing sites are *not* removed by the efforts of law enforcement. Site take down is usually accomplished by companies being targeted by the phishers and third parties, generally private security companies, working on their behalf who communicate with ISPs, hosting companies, server operators, registrars, and individual computer owners whose machines, services, and/or networks have been abused and/or compromised in the creation of the phishing sites.

In a majority of phishing cases, published Whois data on the domain name(s) involved, has been an irreplaceable part of the take down process, an invaluable resource, in fact, without which most of the cited cases would not have been successful. For cases in which legitimate machines or services have been hacked or defrauded, published domain name Whois information is an important tool used to quickly locate and communicate with site owners and service providers. For cases where domain names are fraudulently registered, the published domain name Whois information can often be tied to other bogus registrations or proven false to allow for quick shut down.

It is important to understand the timeframe of phish site shutdown. The longer that a phish site is live, the larger the number of consumers that are defrauded. Most phish sites are shut down within a few hours of being launched. Therefore, it is critical that the entities that are investigating and shutting down phish sites have access to the appropriate tools, like domain name Whois, in real-time.

Given potentially fundamental policy changes affecting domain name Whois under consideration by ICANN, the APWG's DNS Subcommittee has produced this advisory memorandum, comprised of real-world case studies that represent the most common applications of using domain name Whois data employed in phishing site shut downs. The intent is to better inform the broader ICANN community of the uses domain name Whois data provides in shutting down nearly 1,000 phishing sites per day.

It is the hope of the APWG's DNS Subcommittee that exposure to this information and the following case studies will allow the relevant ICANN committees to make better informed decisions on Whois policy and promote policy modifications that will not result in reduced access to Whois data to those who use it to respond to phishing events.

Background

The members of the APWG include brand owners who are being phished, commercial security companies that specialize in phish site takedown, developers of anti-phishing technologies, academic researchers, and law enforcement agencies. This wide range of experience puts the APWG - as a collective whole - at the very forefront of expertise on issues surrounding the relationship of domain name registration information (also known as "Whois") and its uses in combating the problem of "phishing".

Over the past several years, APWG members collectively have shut down hundreds of thousands of phishing websites throughout the world. Almost none of these phishing sites were removed by the efforts of conventional public-agency law enforcement. In most cases, law enforcement is uninvolved in the actual take down process of phishing sites. In fact, they are precluded by statute, capabilities, and/or manpower from taking on the tasks required to remove a phishing site from the Web. Therefore phishing sites are usually removed by employees of the impacted brands or by vendors that specialize in these services that are retained by the impacted brand owners. In addition, the longer a phish site is live, the more innocent users are compromised. This makes it imperative that targeted institutions and their representatives be able to obtain as much information as possible about the location, ownership, and hosting of phishing websites from publicly available resources as quickly as possible.

In a majority of phishing cases, published Whois data on the domain name(s) involved, has been a valuable part of the take down process. For cases where legitimate machines or services have been hacked or defrauded, published Whois information with open, accurate contact data is an important tool used to quickly locate and communicate with site owners and their service providers via email, phone, and fax.

For cases where domain names are fraudulently registered as part of the phishing scheme, the published Whois information can often be tied to other bogus registrations – especially via email accounts - and even directly to the victims of prior identity theft through name, address and phone numbers. This allows responsible registrars to take action on domains that are part of current or future phishing scams.

In all, over 80% of phishing site take-downs involve using the domain name Whois system to find a contact for assistance via e-mail, phone and/or fax, or to prove the registration to be fraudulent through any or all portions of the available information.

IP Whois databases are also quite useful in performing shut downs. However recent trends in phishing sites that use fraudulent domains tied to "fast-flux" DNS to rotate the phishing site around large "bot-nets" (sometimes these bot-nets can have tens or hundreds of thousands of compromised and remotely controlled computers throughout the world) have created a difficult problem. Since a phishing site can be moved to hundreds of different servers around the world, the only way to affect an actual take down of such a phishing site is to get the fraudulent domain suspended and removed from DNS. For the

remainder of this document, the use of the term "Whois" will refer to domain name Whois rather than IP Whois data as that is the focus of this paper.

Recent trends in large-scale obfuscation or withholding of Whois data, either by legitimate domain holders or fraudsters taking advantage of obfuscation systems (both commercially available or easily duplicated) have made the phishing site deactivation process more difficult and thus slower. Of course, slower shut down of phishing sites leads to increased consumer exposure to such sites and higher monetary and personal information losses for both individual victims and the financial institutions being targeted.

Case Studies

Case Study #1: Use of correct, available Whois information in a domain name registration record to effect rapid shut down of an illegal phishing website.

APWG members have used accurate, public Whois data in thousands of cases to rapidly shut down phishing websites. While having the correct technical contact details for the entity providing the domain's hosting is the usual avenue for fraudulent content removal, in many cases, the domain owners themselves are the agents that perform the shut down of a phishing site attached to their domain. They are often more easily reached than their actual hosting provider (who can be deep within a reseller distribution channel) and Whois is often the only way to get the contact information for the owner of the domain.

A great example of Whois information being an invaluable tool for rapid phishing site shut down came on January 27, 2005 in shutting down a phishing site targeting a major credit card company. The site was asking for detailed information about the credit card, including card number, PIN, and the name of the card holder. The phishing site was embedded within a legitimate website that had been hacked by a phisher.

Attempts to raise the domestic hosting company where the site's server was located went unanswered. In this case this was the hosting company's first phishing incident and they had no established procedures or published contact information for such abuse reports (they did subsequently develop such procedures as a result of this attack and others). The real website that had been hacked had no contact information for the owner/operator available on it (i.e. no "contact us" section).

However, take down team personnel were able to quickly find the actual site owner due to his name and cell phone number being published in the administrative contact field of the Whois record for his domain name. This allowed for direct contact with the site owner who took immediate action to disable his website and clean up the hacked server. Without the phone number that was available in the Whois information, this site would likely have been active for well over 24 hours, as that was the expected the turn-around time for getting the hosting company to respond and act. With the Whois contact available and accurate, the site was taken down in just a few hours. This undoubtedly saved many individuals from divulging their credit card information and being defrauded.

In turn, this saved member banks the expenses associated with covering the fraud losses on those cards.

Case Study #2: Use of criminal pattern tracking in the Whois database to quickly shut down and even pre-empt launches of phishing attacks.

Some phishing groups use methods of attack that leave visible patterns in the Whois database. For instance, they often utilize a single or small set of unique names, addresses, phone numbers, or contact email addresses to control their portfolio of fraudulent domain names. Email addresses are especially important in this regard, as they are often used for "drop accounts" – email accounts that phishers use to collect and traffic in stolen credentials and personal information of their victims. Tracking that information allows entities such as anti-phishing services or law enforcement to quickly identify several different domain names as current or future phishing sites. Armed with that information, such groups can work with registrars to connect these illegal activities with specific domain registration accounts and act to shut them down.

In a series of phishing incidents targeting several of the largest US ISPs in late 2004 and early 2005, a take down service vendor was able to track ongoing phishing attacks utilizing domain names that had several common characteristics. The phishing sites were set-up to collect login information for major online services and then credit card details including number, PIN, name, address, phone etc. The domains typically involved the use of the online service brand in combination with a trusted word like "account", "login", or "password" (e.g. bigISP-login.net). These domains were registered in batches over several days in different months, utilizing a dozen or more registrars, but all with a very small set of unique registrant names and administrative Whois contact credential sets that included a rotated set of names, addresses and phone numbers, as well as specific email addresses created and used specifically for the phishing attacks.

Armed with this information, the vendor was able to work with registrars to not only shut down the live phishing sites, but also suspend several domains that had yet to be set-up as phishing sites. In many instances this prevented even a single victim from being lured in by a fake domain name. Without access to Whois information that showed this clear pattern, the domains in question would have had to go through a lengthy UDRP process in order to allow the legitimate trademark holder to assert control over the domain. Since that process takes weeks to months to follow, the phisher would have been able to easily start phishing scams on those domain names and steal thousands of user credentials.

Case Study #3: Obfuscated Whois information interfering with phishing site shut down – increasing the number of potential victims of a phishing crime.

Inaccurate, incomplete, or intentionally obfuscated Whois data is a hindrance to any investigation of an active phishing site. The Whois system was originally intended to aid in resolving technical issues regarding the Internet presence the domain name represents. A hacked server with a phishing site on it would certainly fall under that description, however with an unusable Whois entry, resolving these problems that impact the entire Internet community becomes a much harder problem – the exact opposite of the original intent for the database.

When this issue is discussed, the use of obviously fake data to set-up a phishing domain name comes to mind naturally, as anyone could fake their Whois data entry and criminals will often do just that. However, that tactic can often backfire against a phisher, as a registrar is more likely to terminate such a domain more quickly or not even register it in the first place, so "smart" phishers are sticking with realistic entries and email addresses that actually work. More problematic has been the recent widespread adoption and marketing of domain "privacy" services, which has created a method for scammers to hide illicit registrations. It's nearly impossible to track criminal registrations through such services, as they are created explicitly to make it difficult to contact a domain name's true owner. Beyond the obvious problem with hiding criminal registrations, the use of such "screens" makes it more difficult to track down a legitimate domain owner who does not know his site has been hacked. This can increase a phishing site's longevity, and ironically leaves the domain owner unaware of potentially serious issues regarding the very Internet presence they are trying to protect.

A good example of this problem occurred on July 1, 2006 with a phishing site targeting a major credit card company. The site was configured to steal a wide range of personal data as well as credit card information – a full identity theft kit. The site was located on a server that had apparently been hacked through a vulnerability in a commonly used blogging software package. Unfortunately, the hosting company did not have staff in-place to handle the incident at the time of the report, and did not respond to requests for action. This is an all too common issue, as many hosts – especially on weekends – can take 12-24 hours to read their abuse queues and may not answer their phones. Because of this, getting to a site owner is often the quickest way to resolve many phishing incidents.

In this case, the domain holder Whois information for the site being hacked was masked using a domain Whois "proxy" service. This made the domain owner unknown and unreachable, since the website itself contained no information about the owner or operator, nor how to contact them. Further investigation using alternative, time-consuming sleuthing over several hours by expert investigators eventually produced a reference to the site owner's email address. The owner answered requests for action within 10 minutes of being sent an email, and took down the phishing site right away. Had this information been accessible via Whois, it could have reduced the phish site live time by as much as 12 hours.

That translates into a large number of financial credentials and personal information sets that were likely obtained by the phisher in the interim. Ironically, if the owner was obscuring their contact information in order to avoid spammers finding his email address, investigators were able to eventually run it down on the Internet anyway. So he has probably had it "scraped" by spammers already, and the Whois "protection" is largely moot.

Conclusion

The APWG has thousands more sample cases of phishing activity every month in which the availability of accurate domain Whois information has played an important role in the determining how quickly a phishing site has been disabled.

The three cases cited above are useful for categorizing the issues that can come up during a phishing site take down operation and exemplify the huge value the current Whois system can provide for facilitating phishing site shut downs.

Additional scenarios exist, but almost all rely upon having information available to investigators and first responders to enable them to rapidly disseminate information about an online threat to the people who control the on-line asset being used to enable that threat.