Advisory on Utilization of Whois Data For Phishing Site Take Down

March 2008

Contributors

Rod Rasmussen, Internet Identity
Patrick Cain, Anti-Phishing Working Group
Laura Mather, Anti-Phishing Working Group
Ihab Shraim, MarkMonitor

Summary

Given fundamental policy changes regarding accessibility of both domain and IP Whois data currently under consideration by ICANN, RIPE and others, and the evolving environment surrounding the Whois system, the APWG Internet Policy Committee (IPC) has updated this industrial advisory, comprised of a set of real-world case studies in which Whois data has been instrumental in neutralizing phishing sites in order to help ICANN, RIPE and others comprehensively inform their policy deliberations.

The intent is to better inform the broader internet policy 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

example cases describes a specific event, but represents hundreds of analogous events that occur daily.

The report's sponsor, the 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. Among the institution's membership are a large number of 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.

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

Counter to popular 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 of the domain name(s) and Internet Protocol (IP) network addresses involved have been irreplaceable components of the take down process -- invaluable resources, in fact, necessary to the resolution of most of the cited cases. For cases in which legitimate machines or services have been hacked or defrauded, published domain name or IP network address Whois information is an important tool used to quickly locate and communicate with site owners and service providers. For cases in which 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 a phish site shutdown. The longer that a phish site is live, the larger the number of consumers who are defrauded. Most phish sites are shut down within 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 and IP Whois, in real-time.

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

In most cases, law enforcement is uninvolved in taking phishing sites. They are precluded by statute, capabilities, and/or manpower from taking on tasks required to remove a phishing site

It is the hope of the APWG's IPC that exposure to this information and the following case studies will allow the relevant committees of ICANN, RIPE and other governance bodies to make better informed decisions on Whois policy and promote policy modifications that will not result in reduced access to Whois data for 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 and IP address registration information (also known as "Whois") and its manifold utility 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 jurisdictional issues [??] as well as limited technological capabilities and/or manpower from taking on the tasks required to remove a phishing site from the World Wide Web.

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 greater the number of innocent users who 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.

Whois data can be tied to other bogus registrations and directly to victims of prior identity theft, allowing responsible registrars to take action on domains that are part of current or future phishing scams

In a majority of phishing cases, published Whois data on the domain name(s) or IP addresses involved has been a valuable part of the take down process. For cases in which 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 in which 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 Whois information.

IP network address 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 effect an actual take down of such a phishing site is to get the fraudulent domain suspended and removed from DNS.

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.

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

A great example of Whois information being an invaluable tool for rapid phishing site termination 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 call the United States-based hosting company where the site's server was located went unanswered. 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 turn-around time for getting the hosting company to respond and act. With the accurate Whois contact available, 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.

Some phishing groups use methods of attack that leave visible patterns in the Whois database. Email addresses are especially important in this regard, as they are often used for "drop accounts"

Tracking that information allows entities such as anti-phishing services or law enforcement agencies 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 shut down live phishing sites and suspend several domains that had yet to be erected into working phishing sites

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 Uniform Domain Name Dispute Resolution Policy (UDRP) process in order to allow the legitimate trademark holder to assert control over the domain. Since that process takes at least four months to complete, 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 – and 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.

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

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, 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 kind of problem occurred on July 1, 2006 with a phishing site targeting the customers of a major credit card company by presenting a

convincing counterfeit of the company's own website. 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, contacting 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 neither information about the owner or operator nor contact data for 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.

The owner answered within 10 minutes of being sent an email. 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" was largely illusory.

Case Study #4: Using IP Whois information to contact the Internet Service Provider hosting the phish site

There has been discussion about creating restrictions on the ability to look up the owners of IP addresses and IP address blocks. While it is imperative for the phish site take-down providers to have access to the domain Whois system, it is also important for these organizations to have access to contact information for the IP addresses hosting the phishing sites. Because Whois information on IP addresses is much more complete and accurate than Whois information for domains, losing this resource would have a huge impact on the anti-phishing community.

This is exemplified by a recent case where a lookup of the information for the owner of the domain being used for the phish site was not displayed because they had opted to use a privacy service as the information in the domain Whois record. In this case the phish site take down provider had to determine the organization that was hosting the website to have the website disabled.

Without the ability to determine the owners of the networks and subnetworks, it would not be possible to get the phish sites shut down in the case that the domain owner cannot be contacted

To get the site disabled, the phish site take -down provider had to perform a DNS lookup to determine the network IP address of the phish site. This required the provider to determine the geographic location associated with the IP address. Based on the geographic location of the IP address, the take-down provider went to the IP lookup service associated with the particular geography, in this case, RIPE, and used that service to lookup the owner of the IP address. Unfortunately, in this case it was not possible to look up the owner of the IP address, so the take down provider used a traceroute to determine the upstream provider hosting the internet traffic for that phish site. The take down provider then used the RIPE IP Whois

database to find out who the upstream provider was, based on the Whois for their IP addresses that appeared in that trace.

After the take down provider had identified the ISP, they used reverse DNS to determine other sites hosted by that ISP. This revealed that the phisher had registered several phishing domains targeting multiple institutions and that they were hosting these domains on the same network. The take-down provider collected this evidence and then contacted the ISP and used the fact that there were multiple phish sites, all with the same proxy domain Whois information, hosted on the same network. This evidence made it easier to convince the ISP that the domains were being used for malicious purposes and they were shut down within an hour of the ISP being contacted.

Case Study #5: Disabling a phish site that is based only on an IP address

A different scenario that phish site shutdown providers encounter is when the phish site does not use a domain, but is hosted directly on an IP address. An example URL that uses an IP address instead of a domain is http://123.123.123.123/phishlogin.html. Approximately 15% of all phishing URLs use IP addresses instead of domain names.

The process for shutting down a phish site using this kind of URL format is quite different from closing a domain-based phish site. An example of one of these cases is when the IP address resolved to a machine in China. In this case the phish site shutdown provider found the URL in question and determined that the IP address that was the main part of the phishing URL resolved to an IP address on a network hosted in China. The shut down provider used the Asia Pacific Network Information Center (APNIC) to determine the ISP that controlled that IP address.

Approximately 15% of all phishing URLs use IP addresses instead of domain names. The shut down process for a phish site of this format is quite different than for a domain-based phish site

The shut-down provider then contacted that ISP, through a translation service, and explained that the IP address was being used to host a phishing site. After much discussion about how the site was being used to defraud consumers, the ISP agreed to disable the IP address in question.

There are several cases where determining the owners of the networks and subnetworks is crucial. These include cases involving proxy domain Whois information, the inability to contact the owner of the domain, or a phish site hosted on an IP address. In these cases it would not be possible to get the phish sites shut down without open access to the relevant IP Whois database. This could result in a huge increase in the number of consumers whose identities are stolen due to phish sites being live for much longer timeframes.

Conclusion

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

The five example cases studies cited in this advisory are useful for categorizing the issues that can come up during a phishing site take-down operation. What's more, they exemplify the huge value the current Whois system can provide for facilitating phishing site shut downs.

Additional scenarios exist, but almost all of them rely upon having accurate Whois information available to investigators and first responders, in real time, 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.

Correspondent Authors:

Rod Rasmussen, Internet Identity: rod.rasmussen@internetidentity.com
Patrick Cain, Anti-Phishing Working Group: pcain@antiphishing.org
Laura Mather, Anti-Phishing Working Group: laura.mather@antiphishing.org
Ihab Shraim, MarkMonitor: ihab.shraim@markmonitor.com