Lending Clarity to Security Risk Definitions

by Dave Piscitello and Greg Aaron

In its Beijing Communiqué of 11 April 2013, the ICANN Government Advisory Committee (GAC) called on ICANN to have new gTLD registry operators find and act upon a variety of abusive activities occurring within their TLDs. This led to a requirement in all new gTLD contracts:

> *Registry Operator will periodically conduct a technical analysis to assess whether domains in the TLD are being used to perpetrate security threats, such as pharming, phishing, malware, and botnets. Registry Operator will maintain statistical reports on the number of security threats identified and the actions taken as a result of the periodic security checks.* [Specification 11, paragraph 3.b]

Due partly to time constraints, this language was purposefully left general and without detail. On 25 June 2013, the New gTLD Program Committee (NGPC) of ICANN's Board of Directors issued a resolution (Resolution 2013.06.25.NG02) that calls for a Framework to define details regarding how registry operators will conduct that threat analysis and what kinds of follow-up actions may be suitable.

For effective policy development, all parties must share a clear understanding of the definitions of these security risks. They should also understand the way that threat data are collected and reported for these risks.

In this paper, we provide definitions of four fundamental types of security or abuse risks for which monitoring and reporting can be implemented: phishing, malware, spam, and botnet command-and-control. We explain how threat data for these risks are collected and reported today. We examine relationships (similarities or characteristics) of certain security risks currently under consideration in order to lend clarity to the ease or difficulty of collecting or reporting threat data.

The ICANN SSAC has recommended that spam be added to the list of abuses in addition to phishing, malware, and botnets specifically mentioned by the GAC and in the registry contracts. The SSAC notes that domain names used for these four types of abuse "present greater security problems for Internet users, are obtained via registration systems, and affect user perceptions of Internet security and usability."

Data Sources

A variety of resources are available to identify problematic domain names and URLs.

Internet blocklists are lists of domain names, URLs or IP addresses that the blocklist administrator considers harmful, dangerous, or undesirable. Blocklists have been used to protect Internet users since the late 1990s. Each blocklist has a different purpose and therefore contains different types of content and has different criteria for adding and removing entries. Some blocklists are available for free use, while others may require a subscription. Blocklists also vary in their quality and the quantity of threats they identify.

Anti-spam, anti-malware, security suite software, firewalls, and gateway ("middle box") products often use blocklists to obtain threat data used to protect users of these products from falling victim to attacks.

Other services can identify problems with domain names or specific URLs. For example, there are services that scan web sites for malware.

Phishing and Pharming Attacks

Both phishing and pharming bring users to bogus web sites but these attacks are perpetrated in different ways.

Phishing is primarily a social engineering attack that lures a user to a fraudulent website, typically via a deceptive email message. This website often impersonates a trusted site, like that of a bank, and asks the visitor to enter sensitive personal information such as user credentials, account numbers, personal identifying information, etc. Phishing is a well-documented phenomenon, and the Anti-Phishing Working Group (APWG) publishes statistics about the number of observed attacks, the number of domains used, and so on.

Security professionals and victim users frequently report phishing URLs to blocklists such as the APWG eCrime eXchange, SURBL and PhishTank. Web browsers and other applications that allow users to click on hyperlinks can incorporate phishing blocklists into their phishing filters to alert users when they attempt to visit phishing sites.

Pharming is the re-direction of victims to a bogus site, typically via DNS manipulation. This can be accomplished via cache poisoning, or a malware infection on the victim's computer or device that redirect's the victim's computer to use DNS servers the attacker controls. As such, pharming has a distinctly different technical component. But like with phishing, the user ends up at a bogus site where he or she is enticed to give up sensitive information. A pharming redirection will work even if the victim has typed the correct web address.

Pharming is not as easily detected as phishing and it hasn't been as well documented. Since pharming changes the translation of domain names to servers controlled by the attacker, there are no URL blocklists dedicated to pharming.

Phishing and pharming use many different resource identifiers

Domain names clearly play prominent roles in phishing and pharming attacks.

A domain name may be used to identify a location where an attacker has published a fraud or impersonation page. Such a page is accessible via a hyperlink, such as:

```
http://www.payyypal.com/login.html
```

This example of a domain name that a criminal actor has registered for the purpose of perpetrating a crime is what we call a *maliciously registered domain name* or *malicious registration*. When identified, such domains are candidates for immediate suspension by the registry or registrar since they have no legitimate purpose and the takedown will prevent further victimizations.

A criminal actor might also use a *legitimately registered domain name* hosted on a compromised web server over which he has gained unauthorized administrative control. With this unauthorized but undetected access, the criminal can publish a fraud or impersonation page at a hyperlink such as

```
http://www.personalsite.com/paypalcom/login.html
```

Such phishing attacks are best mitigated by the hosting provider, which can remove the phishing pages and patch the server vulnerability, thereby helping the innocent domain registrant and preventing additional potential phishing victims at that website.

Phishers may also publish many fraud or impersonation pages at a single, compromised web server. Here are some simplified examples of such hyperlinks:

```
http://www.marysembroidery.com/chasebank/login.html
http://www.marysembroidery.com/barclays/login.html
  http://www.marysembroidery.com/hsbc/login.html
```

Note that in the case of pharming, the hyperlinks to which victims are redirected can also be constructed with maliciously or legitimately registered domain names in the same fashion.

Spam

Spam is the commonly used term to describe unsolicited bulk email: mail that is sent to thousands or millions of recipients who have not consented to receive such mail. While national laws around the world vary in their definitions of spam, spam is considered a serious security and abuse problem for several reasons:

1. Spam is a primary means used to advertise and perpetrate a variety of harmful activities. These include phishing attacks, frauds (e.g., advance fee), illegal or counterfeit products, the delivery of malware via attachments, or hyperlinks to malicious web sites that have been embedded in the email message..
2. Industry measurements generally find that between 70% and 85% of all email sent worldwide is spam. Spamming is an illegitimate activity that places enormous burdens on consumers, network operators, and ISPs.
3. Spammers typically employ a range of fraudulent and deceptive practices in order to carry out their work.
4. Blocklistings indicate that spammers register and use more domains for spamming than are registered for other types of abuse such as phishing and malware distribution.
5. Much of the world's spam is sent from botnets.

Spam and Domain Names

Spammers register enormous quantities of domain names, which they advertise in the bodies of their emails. These destinations may host websites or may redirect victims to other destination sites. Spammers know that the domains they advertise will likely be blocklisted, so they continually register and use more domain names in an attempt to stay ahead of defenders.

Spammers may also send email from domains they register, commonly faking or spoofing the "FROM" address in their spam messages. It is very easy to forge the sender address in email. As such, instead of relying on the FROM address, security and blocklist providers tend to focus on two more reliable indicators:

1. The domain that is being advertised within the email – the destination that the spammer wants recipients to visit.
2. The IP address from which the spam was sent. This address is often a resource that the spammer controls. The IP address can often be derived from the full email header.

Spammers often control blocks of IP addresses, from which they send out spam. Spammers also transmit spam email messages from compromised computers and from compromised email accounts.

Spam blocklists are used by virtually all entities that operate email servers, including providers such as Google and Microsoft, ISPs, companies, and universities. These entities use blocklists to filter out spam and malware-bearing attachments so that they do not reach end-users. Prominent spam blocklists include SURBL,

Spamhaus, and URIBL. Registry operators and registrars can also use these lists to identify problematic domains and therefore problematic registrants.

Distributing or Hosting Malware

Malware (malicious software) is executable code that is installed without a user's knowledge or consent and that allows an attacker to perform malicious activities. Malware may be delivered as an attachment to a spam email message or it can be hosted at web sites, cloud or file sharing hosting sites. Spam again plays a key role: hyperlinks where malware may be accidently downloaded are often distributed in spam email messages. Thus, when considering malware as a security risk, spam again becomes an important, complementary security risk that merits monitoring or reporting at registry and registrar levels.

Malware and Domain Names

Some blocklists specifically identify domains or specific URLs being used for malware distribution. A domain name may be used to identify a web server or file sharing service where an attacker has published a malware executable. Such a page would be accessible via a hyperlink, for example, "trojan" malware hyperlinks might look like this:

```
http://deleondeos.tld/img/script.php?tup.jpg
     http://gov.f3322.tld:1717/syss.exe
```

These are examples of likely maliciously registered domains (registered for the purpose of perpetrating a crime); however, malware may be hosted at compromised sites as well.

Operating Botnets

Botnets are networks of infected computers or devices. These computers are commonly infected via spam email containing malicious attachments or hyperlinks that cause recipients to visit the attacker's website. These malicious attachments or hyperlinks install a malicious executable that allow an attacker to gain control over the victim's machine.

The creation and operation of a botnet is considered illegal in virtually all countries. Botnets are used to perpetrate a variety of crimes, including distributed denial-of-service attacks (DDoS), sending spam, online extortion schemes, and so on.

Botnets and Domain Names

The infected machines reach out and contact *command-and-control domains* – domains referencing attacker-controlled servers. These servers, controlled by the

attacker or "botmaster", are where instructions destined for the infected machines of the botnet are placed. This allows the botmaster to control the botnet. The instructions sitting at the command-and-control servers tell the infected machines what to do, e.g., emit spam, probe or attack other computers, exfiltrate data, or contribute to denial of service attacks.

Command-and-control domains are therefore weak points in a botnet. If these domains are suspended, the botnet may not be able to function, at least for a time.

Botnets often use *domain generation algorithms* ("DGAs"). In these cases, the botnet malware contains a system for generating many (sometimes hundreds) of domain names on a daily basis. The domain names created via these DGAs may be viable command domains that the infected machines will check in with on a given day. Only one or two of these may actually get registered by the botmaster and become functional command-and-control domains. The other possible domains are decoys. This strategy is designed to make it difficult to disrupt a botnet by suspending the viable domains or pre-empting their registration. Botnet operators' DGAs force interveners, registry operators, and law enforcement into complicated and protracted mitigation operations (see Conficker).

Some blocklists list active botnet and malware command-and-control domains.

Security researchers maintain DGA lists. These are compiled by reverse engineering botnet malware, a process where experts are able to identify what algorithmically generated domain names will be viable in the future. These lists are updated daily (or whatever periodicity the DGA exhibits) for multiple algorithms and can be voluminous.

Some registrars and registry operators have blocked command-and-control domains proactively. ICANN's Expedited Registry Security Request (ERSR) process was created specifically to allow registries to block botnet command-and-control domains.

Choosing Measureable, Meaningful Security Risks

Important observations may be drawn from this discussion:

1) Reliable data about domain names being used for phishing, malware, spamming, and botnets are available.

2) Each type of abuse has its own indicators and characteristics. Tactics for monitoring and mitigating each type of abuse should therefore be tailored to address the appropriate indicators and characteristics.

3) *Detection* is obviously a pre-condition to *mitigation*. To be effective, anti-abuse procedures must identify problems and then address them. Critically, problems must be mitigated in a timely fashion since the longer problems are allowed to exist, the more damage they cause.

4) It is important for responders to understand the difference between maliciously registered domains and those that have been compromised. Maliciously registered domains are candidates for immediate suspension. The hosting provider is best able to mitigate compromised domains or web sites. It may be challenging at times to conclusively establish a domain has been registered maliciously.

5) Communication is critical. Importantly, a registry operator that identifies a compromised domain can push information about it to the sponsoring registrar and/or the relevant hosting provider, who can then mitigate the problem for the registrant.

6) Data can point out problem registrants for special attention. Malicious domain name registrations are often made in batches. The "80/20 rule" often applies, with a small number of miscreants creating the majority of problems. Proactively addressing these "hot spots" can be a good business practice for registries and registrars, and can do much to improve security and user trust.

7) Threat data focuses on malicious or criminal activity, not on the technical means by which the activity is perpetrated. Certain threat data – for example phishing and pharming – does not discriminate based on technical means.