Draft June 17, 2009:

DISCLAIMER:

This *ISSUES* document is still in DRAFT form and cannot be considered final. The APWG Internet Policy Committee (IPC) has gone through several major editing rounds on this document, with another to come. The issues discussed here are important to many APWG-IPC members, however this committee is still working towards better consensus and support for all the questions included here, or modification/deletion of some issues. A SEPARATE document will be generated that includes potential solutions or recommendations to deal with the issues outlined here. The opinions in this document do not represent the views of the Anti-Phishing Working Group as a whole, and participation in this APWG-IPC effort does not signify endorsement of all points in this document by the participating members or their organizations.

Potential issues in malicious use and abuse of the domain naming system created or exacerbated by the new gTLD expansion

Overview

The Anti-Phishing Working Group (APWG) is an industry association focused on eliminating the identity theft and fraud that result from the growing problem of phishing and email spoofing. The organization provides a forum to discuss phishing and e-crime issues, evaluate potential technology solutions, and actively engages a wide variety of organizations throughout the world to work on common solutions to these problems. The APWG's Internet Policy Committee (IPC) is a standing committee within the APWG that includes over 90 members representing the full spectrum of the APWG's membership. The mission of the IPC is to help developers of Internet policy understand evolving electronic-crime threats and assist in the development of domain name system (DNS) and other Internet-related policies that protect Internet users and organizations from e-crime. IPC members include people from security vendors, registrars, registries, academia, law enforcement, financial institutions, technology consortiums, and other APWG members. IPC members have been attending and briefing Internet policy makers at many forums since its inception, including extensive involvement with the ICANN community, and have managed to bring various ICANN constituency members into the APWG community

as well. Initiatives completed by the IPC include advising the ICANN WHOIS and Fast Flux working groups, providing use cases for how WHOIS is used in phishing site take-downs, publishing statistics on domain name use and phishing trends - including a study on the use of sub-domains by phishers, and publishing registrar best practices. Ongoing work includes creation of a registry-level domain suspension process, studies of website vulnerabilities that lead to phishing site creation, continued data studies, and launching initiatives to educate both users and web site operators on phishing. The IPC has a long history of working with the ICANN community and counts among its membership many from that community, including registrars, registries, members of the business constituency, network operators, government and many others representing all of the various ICANN constituencies.

From this perspective and experience, the APWG's IPC views the planned expansion of gTLDs to be an important event with potential impact on the e-crime space. This paper is intended to provide constructive input to the ICANN community on various issues the APWG's IPC feels merit attention and planning during the roll-out of the new gTLDs. These are not intended to be a list of objections to the entire process, but rather issues that may need to be addressed via policy, contracts, best practices, or education of new registry operators.

This paper was coordinated by the APWG's IPC and includes input from IPC members as well as other APWG members. The list of IPC members involved in the creation of this document is:

- Rod Rasmussen (Subcommittee leader), Internet Identity, Co-Chair, APWG-IPC
- Laura Mather, Silver Tail Systems, Co-Chair APWG-IPC
- Greg Aaron, Afilias
- Paul Diaz, Network Solutions
- Jeff Neuman, Neustar
- Mike Rodenbaugh, Rodenbaugh Law
- Joe St Sauver, University of Oregon
- Dan Schutzer, Financial Services Technology Consortium (FSTC)

Other Major Contributors (APWG members)

- Peter Cassidy, Secretary General, APWG
- Dave Piscitello, ICANN (representing his own individual views and experiences with malicious content and DNS infrastructure, and not representing ICANN's positions)

To quantify these issues, we have chosen to categorize them into three primary classes: new threats introduced with this roll-out, issues of scale – problems that arise because of the vast increase in the number of registries, and longstanding problems that can be addressed at the creation of a new domain registry rather than "patched" later.

This paper covers issues ONLY, and does not propose specific solutions. It is designed as a document to spur further discussion and prioritization of issues. A subsequent paper will offer prescriptive measures and thoughts on potential policy inputs or best practices.

New threats

These threats aren't necessarily "brand new", but are being viewed as threats that have not had to be addressed in previous ICANN TLD roll-outs. The threats conveyed in this section weren't dealt with due to the relatively small size of those previous TLD expansion efforts or were not a major concern at the time.

Registry control/ownership

There is widespread belief that at least one former domain registrar was involved in supporting organized crime via its registration practices. There is an even greater level of damage that could be done if a domain registry were to be owned/influenced by criminal elements. Of further concern is that with the expansion of the domain space, there will likely be an increase in domain registrars as well, allowing further opportunities for organized crime to gain a foothold into control of a direct feed into the domain name space.

Under current rules, a registrar or registry is only examined for involvement of a felon in the ownership of that entity at the time of application or renewal of their domain registrar agreement. The current process also does not appear to involve a standardized, thorough background and reference check for such companies and individuals. This allows many loopholes for members of organized crime or other known criminals to gain access to or control of registries and registrars. This could be through subsequent change in the ownership, non-scrutinized employees, or deception.

It is our understanding that a new Registrar Accreditation Agreement (RRA) currently addresses this in part by providing more frequent, in-depth scrutiny of registrars. So this issue may already be addressed in-part through current policy development. However the basic issue is very important given the potentially large numbers of new domain registry operators and the strains on the ICANN staff charged with scrutinizing these applicants. How do we ensure that criminal organizations do not gain control of a domain registry?

Introduction of TLDs with intrinsic potential for abuse

With the plethora of ideas around new ideas for TLDs that have already been demonstrated, there are many that by the very nature of their structure may require deeper scrutiny given their potential for abuse. Primary amongst these are efforts to create TLDs centered on industries that are already attacked heavily – financial services, ISPs, e-commerce companies, and various business related activities. Such TLDs can imply a higher level of trust than others, as they are tied to industries or infrastructure that the average consumer will naturally tend to trust more, and indeed, early "marketing" of ideas for these TLDs seems to be based on this idea. Some members of our community assert that anyone running such a TLD should come under particularly heavy scrutiny and perhaps even regulation or audit to ensure that the TLD is run meticulously. Other than polling members of the industry in some as-yet-to-be determined method, there does not appear to be any consideration given in the currently published process for TLDs of this nature.

Another issue that will undoubtedly be covered by people with intellectual property concerns is the capability to create a TLD that could, by its very string of characters, be used as a direct substitute for a well-known financial institution or other infrastructure provider. The APWG's IPC has no comments in this paper on the intellectual property considerations here, but there is clearly some security risk in allowing formation of TLDs that also have such names. For example, .citi, .poste, or .chase could be seen as attractive for many reasons, but are also names or derivatives of multi-national financial institutions. It is our understanding that such names should be discovered during the review of applications, by the required independent examination of proposed strings. However, we wish to emphasize that this aspect be thoroughly covered in this process.

Business model expansion

With an entire set of new rules and registration procedures possible with many of the models that have been proposed for new gTLDs, there are serious questions to address around controls to keep abusers from exploiting these new processes. One of these is who actually controls the registration process and interacts with the registrant?

Ownership and access to point-of-presence registration data

With current domain registration models, there are already operational challenges for first responders and law enforcement, who need to determine what parties actually create and maintain a domain name registration on behalf of the registrant. New business models may create even more challenges. For example, collection of evidence necessary for investigating criminals' registration of domain names requires information only available at the on-line point-of-presence of the domain registration process. However, with alternative distribution models, such data is often not held by the responsible registrar since these registrars use reseller and even multi-level reseller arrangements to distribute domain name registration

services. This has been somewhat obviated by the new RAA, but transaction information is still held solely by resellers rather than registrars in most cases. This also means de-facto responsibility for handling domain registration activities is diffused, making it harder to investigate and mitigate malicious activities. While these challenges are being dealt with by law enforcement and first responders today, we foresee the possibility of even more complex models and further difficulties with incident response and investigation.

Anti-abuse policies and procedures

A large number of ccTLDs have significantly different business models than most gTLDs. Past behavior in targeting various ccTLD operators to exploit registries who lack strong policy and/or technical prowess indicates that similar issues will arise with new registries not pre-hardened to these abuse tactics. Similar precedent exists with companies and groups offering subdomain registration services. Such providers have a wide variety of business models, but often have little or no real infrastructure behind them. Such services have seen a rapid increase in abuse over the past two years,

The implementation of new registries on a large scale with a wide variety of new reseller/distributor arrangements may necessitate new, well-defined controls and defined roles in the domain registration process.

Changes to registrant qualification and "rights"

Given some of the early proposals for new TLDs, there is strong potential for creation of TLDs where an intrinsic "right" to purchase and/or operate a domain is conferred upon a specific group of people or organizations. The question is whether these "rights" would be conferred without regard to potential abusive or criminal behavior.

New, diverse business models are being proposed that would call for such things as a right for any citizen of a city, member of an activist organization, or graduate of an institution to obtain a domain name within a specific gTLD. Currently, there are many TLDs with eligibility criteria in their contracts, including gTLDs, sTLDs, and ccTLDs, and all spell out the eligibility requirements and recourses for non-compliance. The concern here is that there may be unintentional creation of a situation where a domain cannot be suspended or an abusive registrant blocked from access if rules for new TLDs aren't created with abuse issues in mind. Without caveats for limiting access or revoking domains for various abuse issues, there is the chance that domains registered by criminals for illicit purposes could not be prevented or revoked easily. Further, this type of model could easily lead to automated abuse techniques, where individual identities are fraudulently used in order to create new domain registrations. For example, one could use easily obtainable lists of names and addresses of residents of a city, alumni of a university, or other groups in order to obtain domain names with fully verifiable, yet fraudulently presented registration data. Revocation of such names could be

extremely difficult, and consideration for such issues should be considered in planning for the business models of such registries. Does ICANN have a role in mandating such considerations?

Vulnerabilities and software problems created by potential TLD strings

As has been seen in the past with the introduction of various new TLDs, the very nature of a new TLD can cause systems and software to "break". In some cases, even though precedent existed, four letter TLDs presented unique challenges as many applications wrongly assumed incorrect rules for TLDs. Even worse, some applications automatically try to append .com or other TLDs to what are interpreted as non-complete hostnames, and could thus send visitors to completely unintended websites. Many of these issues have been overcome with software updates, but older software remains. A clever criminal could take advantage of such configurations, and some software may just break badly, leading to buffer overflows and other exploitable conditions. Domains ending in common software extensions could have the same affect, especially on older software and systems. If longer labels than currently exist or "reserved" words in various OS's are employed (e.g. .exe, .pdf, .mp3) there are likely to be a slew of unintended and unanticipated consequences with some (especially older) software and systems interacting with the Internet. That often leads to security vulnerabilities that attackers can exploit for system break-in, obfuscation of malicious content, or other attacks. There are apparently plans to account for this issue as part of the new TLD string examination process, so this point is being made to emphasize the importance of this issue to members of the APWG's IPC.

Are there processes in-place to ensure that a new TLD isn't introduced that causes significant problems with a large number of computer systems? Are some TLDs being excluded already based on some such criteria?

Attacks based on the new TLD name

It has been pointed out that by using TLDs that can be identified with various activities and groups allows criminals to prioritize targets and provides tools to use for social engineering tricks. It would be far more preferable to rob a .bank or .store than a generic .com. A spammer knows to use French and current events in Paris to target someone using a .paris domain. These are just the facts of life whenever new TLDs are introduced, it's just important to understand that these kinds of attacks may occur.

In addition, a TLD like .exe could confuse consumers when a criminal sends a URL that ends in .exe (and, therefore, links to an executable file).

What can be done to ensure new gTLDs that relate to a given concept are adequately protected against attacks that leverage the TLD?

Issues of scale

Capabilities of new registries

Most current registries either employ large, well-supported infrastructures or outsource core systems and support to well-provisioned third parties. This stability has been helped by the slow, controlled roll-out of TLDs. The proposed large-scale roll-out of large numbers of new TLDs at once could easily lead to unprepared entities being given direct license to create and maintain entire TLDs. While this is being addressed via technical assessments, there are still concerns with untested operators entering the market en-masse. From untested legal counsel, to inadequate/inexperienced support staff, to the lack of ability to detect large-scale registrations of abusive domains, there are many potential issues creating attractive venues for criminals to engage in mischief. Past behavior in targeting various ccTLD operators to exploit relatively less security conscious registries who lack strong policy and/or technical prowess indicates similar issues will arise with new registries not pre-hardened to these abuse tactics. Similar precedent exists with companies and groups offering subdomain registration services. Such providers have a wide variety of business models, but often have little or no real infrastructure behind them. These services have seen a rapid increase in abuse over the past two years, and offer a cautionary tale for allowing registry operators that lack anything beyond basic capabilities to proceed without some sort of standards and training.

Adding orders of magnitude to the system's complexity

Currently there are 200+ domain registries worldwide and a few hundred active registrars for law enforcement and first responders to deal with. Further, while there are many registrars, they themselves only have to deal with a few gTLD operators, and a selection of ccTLD operators for whom they choose to provide registration services. Only large, technically proficient registrars typically deal with a hundred registries or more, allowing most registrars to act fairly efficiently with respect to dealing with abuse issues, and any rules/regulations surrounding affected TLDs. With the potential of hundreds of new gTLD registries being formed, the rule set, contact list, and technical back-end that both registrars and those investigating crimes involving domains is going to increase by an order of magnitude or more, as there will be competitive pressures for registrars to offer as many gTLDs as possible.

Adding more actors, restrictions, or new processes to any system will complicate that system – usually in a multiplicative way rather than just linearly -- due to the complexity of the interactions required. A more complicated system is typically slower to operate, and can often lead to more breakdowns in processes. The impact is magnified in this case since the new gTLDs promise to offer a wide variety of new characteristics as well, and not just a large number of new terms to the right of the "dot" that are run the same way. Given that, there is a concern that just understanding and managing all the different requirements, characteristics and processes involved with a large number of new gTLDs could significantly increase the costs to e-crime responders and potentially make reporting e-crime events to the proper providers much more difficult.

Is there a way to introduce new gTLDs without creating a huge new burden on first-responders and registrars to understand and interact with the new registries?

New registrars arising

With a slew of new TLDs coming online, there will likely be several new domain registrars as well. The same security concerns regarding registry ownership and operations apply to the new registrars. Can the issues already outlined for expansion in registries be addressed by the roll-out with respect to new registrars that come about as a byproduct of this process?

More data sources to consult

Counter e-crime operators regularly consult with data resources published by registries and registrars in order to determine where criminal activity is taking place, who may help mitigate it, who may be a victim, and who may be a perpetrator. These resources include TLD zone files, authoritative DNS servers, and WHOIS services. Currently that represents a relatively small number of services to consult and monitor for availability and accuracy. Expanding to scores or hundreds of TLDs will increase the system requirements, contract reviews, system-to-system communications channels, monitoring of data flows, and other operations necessary to ensure suspicious domains and web presences are discovered and quickly diagnosed. This increases costs while also increasing the likelihood of breakdown in these sorts of systems that could lead to missed early detections and thus more e-crime in general. If this expansion is widely successful, these affects are further magnified.

While this is largely just a "fact of life" in the expansion of any successful enterprise, we want to highlight these issues, as there are real impositions being placed on 3rd parties by the decision to greatly expand the TLD space. For example, currently if an enterprise wants to have access to gTLD zone files, they must sign an access agreement, which of course must be reviewed by an attorney. For the current level of registries, expenses for these reviews have run between hundreds and thousands of dollars per gTLD. With hundreds of new gTLDs, that could balloon to over $100,000 – just for contract reviews. A concrete example came in the period when domain "tasting" was first employed and grew rapidly. Many companies and individuals who were tracking the zone files for fraudulent registrations simply could not handle the massive volume increases and gave up, or had to invest heavily in new, more powerful systems.

Can the new TLDs be introduced in such a manner as to lessen the impact on third parties who rely on receiving data like zone files from registries to conduct operations?

BGP take-over attack of a TLD zone

A very interesting if esoteric scenario for registry zone take-over "BGP Spoofing in the Episode: Stealing Your (cc)TLD" was recently proposed by Berislav Todorovic of

KPN at a recent NANOG meeting. In brief, via Border Gateway Protocol (BGP) spoofing, an attacker can divert the DNS traffic for an entire TLD's authoritative nameservers to a malicious authoritative server, send out poisoned answers, and use settings to prolong the attack nearly indefinitely. There are various ways to make such attacks more difficult, but no complete cure is currently deployed. One of the primary defenses against a wholesale take-over of a TLD in this manner involves announcing more specific routes for Autonomous System Numbers of TLDs. That works at the current level of TLDs, but may cause problems with routing table growth with a large number of TLDs to protect in this fashion.

What best practices and operational procedures can be put into place to lessen the chances of this kind of take-over scenario and quickly mitigate one if it were to occur?

Longstanding issues addressable within new gTLD at inception

WHOIS policies

The APWG's IPC is well aware of the long history of debate around issues pertaining to the WHOIS database and does not look to re-open that discussion. However, WHOIS access and accuracy remain key issues for a majority of APWG'S IPC members. Fighting cybercrime requires access to timely, accurate information in order to mitigate attacks and track down perpetrators. WHOIS plays an important role in many of those processes. This is an issue that is very important for legitimate domain registrants as well, since a majority of phishing sites are located on compromised servers with a real domain name. Thus being able to contact the domain registrant is critical so they can be notified that they have been compromised and are at risk of having their personal information and credentials stolen. This notification is as important as protecting the potential victims of an attack. This section outlines some of the potential issues that the new gTLD process brings to the fore in this area.

Published character set

Currently, gTLD WHOIS information is presented in US ASCII as an artifact of earlier limitations to the protocol. There are proposals to switch to Unicode and not necessarily publish in ASCII any further for some registries. This could break the universal access model to WHOIS, making it much more difficult to investigate and mitigate crime and abuse issues. People and processes around the world rely on a universally understood character set to work with WHOIS information in a wide variety of ways, not least of which is contacting victims of e-crime. Allowing gTLD registries to publish WHOIS without an ASCII version could have a serious negative impact on not only anti-crime efforts, but also domain registrants themselves.

Is there a way to ensure access to universally readable WHOIS information for all gTLDs while accommodating the beneficial goal of providing WHOIS information in a script more useful to a particular user base?

Lack of identification details – contact to registrant for mitigation

The APWG's IPC has been assured that WHOIS is "not on the table" for this new gTLD process. We include this point here to express many of our members' desire to keep it out of the discussions around new gTLDs.

Various proposals have been made in the past to significantly change or even eliminate WHOIS altogether. It is imagined that some new registries will attempt to include such provisions in their plans as well. It is of vital importance for law enforcement and first-responders to e-crime and abuse to have access to accurate contacts for the people and organizations responsible for the use of and online presence for domain names. This is a debate that has been in progress for many years within the ICANN community, however there is great concern amongst the anti-e-crime community that the new gTLD launch will provide a "back door" opening into dramatically changing WHOIS publication process, circumventing the policy making process dedicated to WHOIS issues.

Proxy Registrations

There is another debate in the realm of WHOIS over the use and appropriateness of proxy registrations. The current impression amongst many of our members is that there are no set standards being used to regulate how proxies are handled by registrars, and with the notable and new exception of data escrow requirements, no rights are conferred to domain registrants using such services. Law enforcement and first-responder concerns can be addressed while potentially providing a robust marketplace for such services, and helping clearly establish and protect registrant rights and privacy.

It is our understanding that this issue is likely to be dealt with as an overall look into this practice for all TLDs in ICANN's purview. However, this remains an issue of keen interest to many APWG IPC members. Is there policy development in the offing or that could be proposed to address this issue overall?

Known issues addressable at the inception of a new registry

DNS authentication

It has been widely published, and recently explored in great detail at the DNS Symposium jointly sponsored by Georgia Tech and ICANN, that the DNS system has several vulnerabilities and issues surrounding authentication that lead to security and stability issues. DNSSec is a standard that is thought to address many of these concerns and is being implemented by some TLDs already. The creation of new TLDs affords a unique opportunity to build in this protocol from the ground-up.

What can be done to encourage or require implementation of DNSSec for new gTLDs as they go into production?

E-mail authentication provisions

While e-mail standards and other non-abusive domain usage issues traditionally fall outside ICANN's policy-making authority, many potential new gTLD operators have expressed their interest in incorporating and requiring such standards as part of their operational requirements to registrars and registrants. Multiple e-mail authentication standards that would vastly reduce the volume of spam and e-mail abuse on the Internet have been developed over the past several years. One, DKIM, is an IETF accepted standard. Another, SenderID, has been adopted by a number of major e-mail senders. Overall adoption has been slow though, allowing these problems to continue. One of the key reasons adoption has been lagging is that these standards are reliant upon the DNS system for implementation, and often require domain registrars to provide facilities to allow for implementation. However, registrars have been largely unresponsive to this issue, with very few providing the tools necessary for their customers to easily implement these standards for their domains.

Can the new registry paradigm be leveraged to facilitate adoption by at least some registries of one or more domain-name-dependent e-mail authentication standards?

Prevention of fraudulent registrations

Many registrars have implemented extremely good systems and procedures to prevent clearly fraudulent domain registrations from occurring under their registrar account. Several registry operators have noted that abuse issues they have to deal with usually come from a handful of domain registrars. Typically such registrars have weak policies, poor staffing levels, or a lack of understanding of how criminals are utilizing their systems and exploiting their lack of personnel to respond in order to register large numbers of domains.

Which practices can be codified as policy and/or best practices across registrars seeking to offer new gTLDs?

Malicious Fast Flux and other DNS based attacks with domains

There has been a growing trend in criminal use of the domain system using techniques like Fast Flux hosting. There are several recommendations in the industry and coming out of the GNSO's Fast Flux Working Group that could be implemented to prevent, detect, and quickly mitigate malicious domains. The introduction of new gTLDs provides an excellent opportunity to address these issues, and even experiment with different methods to address them.

Which of these methods can be encouraged, tested, and/or required by new registry operators to their registrars?

Standards for domain suspension for abusive activities

Over the past few years, several gTLD and ccTLD registries have adopted new anti-abuse policies. They have universally resulted in direct impact on the amount and severity of abuse taking place within their TLD namespace. This has been

documented empirically in several studies. Many of the provisions adopted by these registries are basically the same – the right to suspend domain names based on specified criteria surrounding criminal and abusive behavior.

Can this learning be put into practice with the new gTLDs to make them more secure at inception? This question may involve voluntary practices, allowing registries the freedom to set relevant terms of service, and/or may involve policy issues that will require policy work in appropriate ICANN fora.

Registries or registrars obtaining names under their own account

Currently registration service providers of various sorts can be their own customer: registrars, their resellers, and even some registry models (e.g. .name) can purchase their own domains. This type of arrangement inherently breaks the practical possibility of the registration service provider acting as a control point for those domains, since this would imply saying something like "I've been bad, I must punish myself by suspending my own domains." In other industries, providers with similar conflicts of interest may be required to use a third party provider so as to maintain an arms-length relationship between a service provider and customer. There have certainly been complaints about perceived abuse on the part of some registration service providers that is and has been the subject of other policy development within the ICANN community (e.g. AGP, transfers). In the case of e-crime, the concern is that a bad actor may "trade on their own account" to create names used strictly for criminal activities that never are suspended.

Are there policy or contract considerations around this concept under consideration or that could be proposed to address this potentiality?

Orphan glue record removal

A recent APWG study found that approximately 3% of domains used for phishing were using "orphan nameservers". These nameservers are simply left-over "glue" records from a domain that was previously removed from a registry – often for abuse itself. This can create a potential "safe haven" nameserver entry in that TLD's zone file that abusers can continue to use to support criminal domain registrations – based in any TLD. Typically, registries and/or registrars remove such glue records at the time of domain removal as a best practice.

Can this practice be codified and disseminated to new TLD operators (and old) so that these safe havens are eliminated?