34 Burnfield Avenue
Toronto, Ontario M6G 1Y5
Canada

Tel: (416) 588-0269 Fax: (416) 588-5641
Web: www.LEAP.com

Leap of Faith Financial Services Inc.

October 9, 2022

ICANN Transfer Policy Review PDP Working Group

Subject: Initial Report on the Transfer Policy Review - Phase 1(a) - response to comments of Theo Geurts, and additional insights

Submitted by: George Kirikos
Company: Leap of Faith Financial Services Inc.
Website: http://www.leap.com/

Dear ICANN Transfer Policy Review PDP Working Group,

After the Transfer Policy Review PDP Working Group's public meeting at ICANN75:

https://75.schedule.icann.org/meetings/hNdkMxTP2FLu93z6h

during which I participated (as a member of the public, as I'm not a member of the working group), Mr. Theo Geurts sent an email to the working group's mailing list:

https://mm.icann.org/pipermail/gnso-tpr/2022-September/000574.html

with some "Input on the break through proposal", which has some of his "observations" on a proposal I submitted on behalf of my company during the recent public comment period. That proposal can be found in the public comments archive:

https://www.icann.org/en/public-comment/proceeding/initial-report-on-the-transfer-policy-review-21-06-2022/submissions/leap-of-faith-financial-services-inc-15-08-2022

https://itp.cdn.icann.org/public-comment/proceeding/Initial%20Report

%20on%20the%20Transfer%20Policy%20Review%20-%20Phase%201(a)-21-06-2022/submissions/Leap%20of%20Faith%20Financial%20Services%20Inc./LEAP-comments-Transfers-Phase1a-20220814-FINAL-15-08-202

Mr. Geurts' email did not generate any further discussion on the mailing list, as of the time of this letter. However, presumably his observations would be discussed on a future working group call (the first call after the ICANN75 meeting is scheduled for October 11, 2022).

As I have not yet been invited to participate directly in the working group (to correct the severe unbalanced and unrepresentative participation, as noted in Sections B and M of my prior submission), this letter publicly responds to Mr. Geurts' "observations" (and will also be posted on the FreeSpeech.com blog).

I also add additional observations and insights to the ICANN75 meeting discussions.

Sincerely,

George Kirikos

RESPONSE TO COMMENTS OF THEO GEURTS, AND ADDITIONAL INSIGHTS

by: George Kirikos

TABLE OF CONTENTS

A. THEO GEURTS EMAIL IN FULL

Below is the entire email from Theo Geurts, as posted to the working group's mailing list. [In subsequent sections of this letter, I'll respond to each issue raised by his email.]

https://mm.icann.org/pipermail/gnso-tpr/2022-September/000574.html

Subject: Input on the break through proposal

Hello,

Some high-level observations on the breakthrough proposal.

Step 1 goto gaining registrar.
This excludes resellers, making it a highly complex process as resellers use different registrars for different TLDs for various reasons. Coding this into systems will be difficult., if not impossible.

Also, registrants know who their reseller or hosting company is. They usually have no idea who the underlying registrar is. This issue is usually related to the wholesale registrar industry.

Generating the PTID.
The suggested method is to log in to the registrar's account and generate the PTID on the website.
While logical, Wholesale registrars have zero control/interaction with registrant accounts at a reseller level.

The proposal does not cover the complexity of sub sub resellers.

If the PTID is compromised, it is still possible for an attacker to set up an account at the registrar and continue the transfer to that registrar and move the domain name to another registrar when the lock period expires.

Regarding the proposal for making the losing FOA visible by using consent from the data subject, consent is a very shaky legal option.
Plus, I have a hard time imagining how this system would work, without creating all kinds of new risks and possible data breaches by people who did not understand what the consequences could be.

Again everything in an ICANN policy is public, and attackers will modify or create new TTPs to get around the barriers/security requirements mentioned in the policy.

And we have not considered all the operational and security effects of the proposals.

Best.
Theo

B. RESELLER ISSUES

Mr. Geurts wrote:

> **Step 1 goto gaining registrar.**
>
> **This excludes resellers, making it a highly complex process as resellers use different registrars for different TLDs for various reasons. Coding this into systems will be difficult., if not impossible.**

This is a bizarre criticism. Our proposal doesn't "exclude" resellers. For simplicity and brevity, we used the language "Step 1: Go to gaining registrar and initiate a transfer." (page 11 of our prior submission) It should have been **obvious** to anyone with an iota of technical knowledge (which Mr. Geurts certainly possesses) that the registrant could go to a reseller, a sub-reseller, a sub-sub-reseller, etc. In other words, they just need to go to wherever they want the domain name to end up, just like they currently do when they place an order to transfer a domain name to a "destination" (whether that destination where they place orders is a registrar, reseller, sub-reseller, sub-sub-reseller, etc. is immaterial).

This criticism is even more bizarre because the term "reseller" only appears once (in a minor section on page 38) in the working group's report, despite policy changes that would also impact resellers, and sub-resellers and sub-sub resellers, etc. [My company was certainly more than familiar with the concept of resellers, given that was explicitly mentioned on page 27 of my submission, in a different context). We're a Tucows reseller.] Indeed, even the "Swim Lane Diagram" on the last page of the Initial Report fails to mention resellers.

Instead of "gaining" resellers (or registrars, or sub-resellers, as the case may be) taking an **input** of the TAC (as they currently do) to facilitate an inbound transfer, these "gaining destinations" would instead **provide** a PTID to the registrant at the time a transfer order is placed (which they would then take to the losing registrar/reseller/sub-reseller, i.e. whoever they are currently dealing with in relation to the domain). If it's a reseller or sub-reseller or sub-sub-reseller, that PTID could be obtained via API communications with the "higher up" entity that they're already dealing with (and ultimately, it should be the registry that generates it, to ensure uniqueness, prevent reuse, etc. as will be discussed below, and was already hinted at in the original submission).

While this is a change, it's not a "complex" change. When interfacing with the "next level up" (whether that be a higher level registrar, a higher level

reseller, etc.), an API call could be made that would return a standard response (a PTID), just like today the reseller or sub-reseller has to pass as an input the TAC (or AuthInfo Code) that they receive from a registrant.

The argument "Coding this into systems will be difficult., if not impossible." applies to **any change**, not just the changes proposed by my company.

For example, the working group proposed that the TAC be specified as per RFC 9154. This would obviously impact resellers, sub-resellers, sub-sub resellers, etc. too! Did the working group (or Mr. Geurts) complain about or consider all the new code that must be written to comply with RFC 9154?

https://datatracker.ietf.org/doc/rfc9154/

Some resellers (or sub-sub reseller) might be creating and syncing their AuthInfo codes (to be renamed "TAC") to private databases. All of those systems will no longer be functional! They'll need to write new code to handle a compliant TAC. **This is no harder than writing new code to receive a PTID.**

For instance, RFC 9154 states:

> Because of this, registries may validate the randomness of the authorization information based on the length and character set required by the registry -- for example, validating that an authorization value contains a combination of uppercase, lowercase, and non-alphanumeric characters in an attempt to assess the strength of the value and returning an EPP error result of 2202 ("Invalid authorization information") [RFC5730] if the check fails.
>
> Such checks are, by their nature, heuristic and imperfect, and may identify well-chosen authorization information values as being not sufficiently strong. Registrars, therefore, must be prepared for an error response of 2202 and respond by generating a new value and trying again, possibly more than once.

All of that requires new code for resellers, sub-resellers, sub-sub resellers, and indeed that might vary between registries. Note that the RFC didn't state that "resellers must be prepared for an error response", yet Mr. Geurts did not appear to criticize that omission (it would obviously be a consequence of the change).

So, if you're going to be critical of "**any change**" then you have to apply that same critiques/standards to the working group's proposals too, which are also proposing changes.

Frankly, it's mere **hyperbole** to describe the changes you oppose as being

"difficult, if not impossible", while not subjecting those changes to any objective standard, and then completely ignoring the fact that the working group is proposing changes of their own that require new code.

Rather than doubling down on all the poor design choices of the past 20+ years (i.e. "pull" system), one should instead bite the bullet now, and shift to a more sound foundation (which would be the "push" system we described, which is comparable to what many registrars already are doing for internal transfers). Changes have been made in the past to many other systems (e.g. due to GDPR), where resellers and sub-resellers, etc. also had to comply. This would just be one more occasion. As noted on page 15 of our initial comments, a new push system could work in parallel alongside the current system, to allow for a transition period (until the older and less secure system is deprecated completely).

C. MORE RESELLER ISSUES

Mr. Geurts wrote:

> **Also, registrants know who their reseller or hosting company is. They usually have no idea who the underlying registrar is. This issue is usually related to the wholesale registrar industry.**

As noted above, for simplicity and brevity we used the term "gaining registrar". It could apply equally to any reseller, sub-reseller or other "destination" that registrants go to place their transfer orders or to manage their domains.

Regarding whether registrants know the identity of their underlying registrar, note that the RAA says:

https://www.icann.org/resources/pages/approved-with-specs-2013-09-17-en

> 3.12.2 Any registration agreement used by reseller shall include all registration agreement provisions and notices required by the ICANN Registrar Accreditation Agreement and any ICANN Consensus Policies, and shall identify the sponsoring registrar or provide a means for identifying the sponsoring registrar, such as a link to the InterNIC Whois lookup service.
>
> 3.12.3 Its Resellers identify the sponsoring registrar upon inquiry from the customer.

Anyone can determine the underlying registrar for a domain name, via a WHOIS lookup. It's unclear whether Mr. Geurts seeks to decouple the identity of the sponsoring registrar from the resellers (and keep that from the registrants), but ultimately that linkage always exists, and for good reason. Ultimately, registrants can continue to deal with the reseller/sub-reseller, etc. they've already been dealing with.

D. EVEN MORE RESELLER ISSUES

Mr. Geurts wrote:

> **Generating the PTID.**
>
> **The suggested method is to log in to the registrar's account and generate the PTID on the website.**
>
> **While logical, Wholesale registrars have zero control/interaction with registrant accounts at a reseller level.**

As noted above, for simplicity and brevity we used the term "gaining registrar". It could apply equally to any reseller, sub-reseller or other "destination" that registrants go to place their transfer orders or to manage their domains.

Instead of inputting a TAC at the gaining destination (whether that's a reseller, sub-reseller, etc.), the registrant would be receiving a PTID (which they would then take to whoever they're dealing with currently, i.e. the "losing registrar", or the reseller at the "losing registrar", or the sub-sub reseller at the losing registrar, etc.).

E. COMPLEXITY OF SUB SUB RESELLERS

Mr. Geurts wrote:

> **The proposal does not cover the complexity of sub sub resellers.**

As noted above, for simplicity and brevity we used the term "gaining registrar". It could apply equally to any reseller, sub-reseller or other "destination" that registrants go to place their transfer orders or to manage their domains.

Any "complexity" of sub sub resellers would apply equally to existing policies/procedures, and also to any proposed changes (like the working group has already put forth, as noted above). It's applying a double-standard to suggest that the working group's proposals wouldn't create equal or even greater "complexity", given the term "sub sub resellers" doesn't appear anywhere in their own report (as noted above, the term "reseller" only appeared once in the working group's report).

F. PTID "COMPROMISE" CONCERNS

Mr. Geurts wrote:

> **If the PTID is compromised, it is still possible for an attacker to set up an account at the registrar and continue the transfer to that registrar and move the domain name to another registrar when the lock period expires.**

This analysis is not correct. There's no such thing as the PTID being "compromised", since it wasn't ever meant to be a secret. It represents **unique routing information for a specific domain name transaction (just like bank wire transfer instructions or IBAN info can be public, or be put into a contract or placed on a website)**. The PTID is linked to a **specific order (perhaps that wasn't clear enough in the original submission)**.

Example: Jane initiates transfer of Example.com to GoDaddy on October 9, 2022 at 12:59:08 UTC, which is currently at Tucows/OpenSRS. At GoDaddy, she received a unique PTID of GODADDY:EXAMPLE.COM:123456

(the PTID should best be generated by the registry, and is unique to represent "routing" for that specific order)

If an "attacker" Melanie has knowledge of that PTID, it's worthless to her. If she creates an account at GoDaddy and initiates an order on October 10, 2022 at 1:25:09 UTC, that's going to be a **separate** order, **with its own PTID**! (e.g. GODADDY:EXAMPLE.COM:785369). Melanie doesn't get to select the PTID that is generated at the gaining registrar! When the true registrant inputs GODADDY:EXAMPLE.COM:123456 at the losing registrar (or reseller, sub-reseller, etc.), the transfer will be successful [although the losing FOA should be retained, as an extra check], and it would show up in **Jane's** account at GoDaddy, not Melanie's. An attacker will be successful only if they can get **their** PTID (routing info) to be input at the losing registrar.

(Aside: a registry can guarantee uniqueness by having a counter incrementing across all domains, so the PTIDs would always be unique, without reuse; so GODADDY:EXAMPLE.COM:100 might be followed by MARKMONITOR:SHOES.COM:101, followed by OPENSRS:EXAMPLE.COM:102, followed by DOTSTER:GAMES.COM:103, etc.; or one can use time within the digits, since that's always increasing and would be unique, and unable to be reused)

The suggestion that there's a "lingering order" that one can copy the PTID

from to hijack the transaction is incorrect ---- the PTID is not perpetual; it should be deleted if the order fails (on page 14 we mention that as an option; perhaps better to make the TTL mandatory); an attacker can't manufacture a specific PTID either (it's best created by the registry, and linked to a specific transaction).

Contrast all this with what happens if/when the TAC is compromised. When that happens, it's **game over** under the working group's proposal (which eliminates the losing FOA), as the attacker would succeed immediately!

If Mr. Geurts is concerned about the PTID being "compromised" (a non-existent threat, by design), where is his concern about the TAC being compromised, where (under their current proposal) the transfer would immediately complete to **ANY REGISTRAR**! (i.e. the attacker can use a compromised TAC anywhere, and there'd be no losing FOA to prevent that misuse)

This is why the "push" system is so much better than the "pull" system, as the attack surface is much smaller (that's why "push" is used for high value monetary wire transfers, not "pull"; crypto also uses push, for good reason, as they're very security conscious).

G. IMPROVING THE LOSING FOA BY MAKING VISIBLE THE "BEFORE" AND "AFTER" WHOIS INFORMATION

Mr. Geurts wrote:

> **Regarding the proposal for making the losing FOA visible by using consent from the data subject, consent is a very shaky legal option.**
>
> **Plus, I have a hard time imagining how this system would work, without creating all kinds of new risks and possible data breaches by people who did not understand what the consequences could be.**
>
> **Again everything in an ICANN policy is public, and attackers will modify or create new TTPs to get around the barriers/security requirements mentioned in the policy.**

The proposal (section G of my initial comment submission, starting on page 23) was to make the **future WHOIS info** potentially visible to some people (e.g. the current registrant), before the transfer is complete, on an opt-in basis, during the Losing FOA step. [the Losing FOA itself wouldn't be visible; this is about WHOIS]

"Very shaky legal option" lacks specificity to make any comment. I'm able to consent to make my WHOIS public (unredacted) at Tucows/OpenSRS, for example, and do so on behalf of my company on an opt-in basis. Tucows obtained valid consent for that. I fully understand the consequences.

There'd be numerous ways to implement the same for a pending transfer, to show what the WHOIS would become (before the transfer has even completed).

For example, a simple implementation might be for the gaining registrar (or reseller, etc.) to generate a WHOIS-passcode for the prospective registrant, at the time they place their order (or the prospective registrant can select it). Suppose I want to transfer Example.com to GoDaddy from Tucows. When I place the order at GoDaddy, they can provide a (strong) WHOIS-passcode of h17Bmm-732!b@Bm-7 and a transaction ID (generated by the registry) of 87864148 that can be used at:

whois.godaddy.com

When someone does a WHOIS for "example.com", in addition to displaying the current WHOIS (via Tucows), their system would recognize that there's a pending transfer into GoDaddy, and say something like "There's a pending

transfer of this domain name -- do you want to see what the WHOIS will look like after the transfer is complete?" Clicking that link would take one to a form, with 2 inputs (one for a transaction ID generated by the registry and visible at the losing registrar when there's a pending transfer, and one for the WHOIS-passcode). [you would need both inputs, in case there are multiple competing transfer requests with different transaction IDs at a given registrar!] If you get the transaction ID and password correct, then GoDaddy would display what the WHOIS would be after the transfer is completed successfully.

Extending this to pass the links (without the transaction ID and password filled out) to the losing registrar/registrant is straightforward, via the registry.

If a buyer (with a change of registrant) of EXAMPLE.COM won't share the WHOIS-passcode with me, so I can ensure that the "After" WHOIS is correct (matching a sales contract, for example), then I might make my own choice whether or not to allow the transfer to go through. [I'd be very suspicious, though] If it's a transfer of a domain name to myself (i.e. I'm the future registrant at GoDaddy too), then I'd be "sharing" the WHOIS-password with myself, a non-issue.

The above is just one possible implementation. An alternative implementation would be merge it into the planned SSAD. Note that the above doesn't collect any new personal identifiable information. It's what the WHOIS would be after the transfer completes, but shown (on a selective basis, password-protected) before the transfer is actually completed, while it still can be NACKed).

As for the last part "***Again everything in an ICANN policy is public, and attackers will modify or create new TTPs to get around the barriers/security requirements mentioned in the policy.***" TTP appears to mean "trusted third party" -- there's no "trusted third party" being created. This is the gaining registrant providing specific consent to the losing registrant to see the WHOIS "before and after". The losing registrar and losing registrant already know their own before WHOIS (as does the public). So, the only info being shared is the new WHOIS to the losing registrant (or anyone else provided with the appropriate password and transaction ID by the gaining registrant), at the location of the gaining registrar. It's whatever the WHOIS would become after a transaction succeeds, but is being displayed on a selective (password-protected basis) before the transfer actually completes.

H. BROAD OPERATIONS AND SECURITY EFFECTS STATEMENT

Mr. Geurts wrote:

> **And we have not considered all the operational and security effects of the proposals.**

That's a broad (and perhaps a "throwaway") comment that lacks specificity to respond to, but it certainly should also apply equally (and even more so!) to the working group's own deeply flawed proposals. I've already given explicit examples of the negative security impacts of the removal of the Losing FOA (including ones that don't involve compromise of the control panel of the losing registrar!).

Wire transfers are all push! I've pointed out how the working group's proposals are inconsistent with SSAC reports, too. Re-read the entire past submission. Hold your own proposals to the same level of scrutiny as counter-proposals by others.

I. ADDITIONAL INSIGHTS

Reviewing the ICANN75 call again, I have the following observations and insights:

1. The working group justifies the removal of the Losing FOA by suggesting that control of the registrar control panel is sufficient to provide all consent, since an attacker can reroute Losing FOA messages if they have access to the control panel. That is simply wrong.

A properly designed registrar system would do out-of-band verification for all critical changes (including changes to where a losing FOA would be sent, so an attacker couldn't simply reroute the Losing FOA without any confirmation). The SSAC reports said as much, but the working group has likely not even read them, let alone considered them. A properly designed registrar system should actually **contemplate** that it **will** be compromised (and thus have appropriate counter-measures, rather than simply saying "Congrats, you're now a Superuser and can do everything!")

Do all registrars do that verification? No, because many of them are horribly designed, and not following best practices.

2. Removal of the Losing FOA means that the losing registrar (and/or reseller, sub-reseller, etc.) and current registrant don't have access to one of the most critical anti-fraud signals, namely the destination of the domain transfer (i.e. the identity of the gaining registrar).

In any "push" system, that destination is obviously known **by design**, and risk systems can be triggered if it's suspicious.

In a "pull" system, it would be **reckless** to remove that important signal from the losing registrar (and losing registrant), but that's what the working group has proposed!

3. To expand on the "bearer bond" metaphor that I mentioned in my initial comments in the August 2022 submission, what the working group has produced via their proposals is actually a **secure withdrawal system**, not a secure transfer system!!

All the emphasis/verification/confirmation by the working group's proposal takes place **before** the TAC is generated. That's like going to a bank,

wanting to withdraw $1 million. The working group, by eliminating the Losing FOA (in the incorrect belief that authentication before the TAC is generated is sufficient), is really only protecting the **withdrawal** of funds (or withdrawal of the domain). At that point, the registrant (or holder of the cash/bearer bond) is on their own! The working group completely ignores the fact that the TAC can be compromised after it is legitimately generated, just as a holder of cash or a bearer bond can be robbed after they've left the first bank (where they made the withdrawal) on their way to a destination bank (where they intended to make a deposit).

Registrants deserve far better than a secure withdrawal system! **We want a secure transfer system, that is secure all the way to the end.**

The working group points to their TTL change and RFC 9154 TAC complexity as major accomplishments. They're simply not (I can **already** set a complex TAC, and use the Lock/Unlock too), and they don't overcome the fact that a compromised TAC can be used by an attacker to move the domain name anywhere they want to go. That's an inherently poor design (which is only defended at present by the Losing FOA ability to NACK, since the Losing FOA displays the destination of the transfer, i.e. identity of gaining registrar).

In contrast, a push system is far more secure, by design, and has a smaller attack surface. With a wire transfer, I specify the destination bank (like a destination registrar), intermediary banks, bank branch and account (just like resellers and registrant). That's why I can make an irrevocable transfer of $1 million or $10 million or $100 million, as the bank will verify with me all the details (including the destination) before they push the funds.

What the bank **won't do** is say "Yes, we've verified you want to withdraw $1 million. Here's a secret code. Take that anywhere you want, and we won't stop where the money ends up, and won't ask any further questions!" The bank doesn't do that because the bank recognizes that that code is inherently insecure, that's just not how to transfer valuable assets. Yet, this is what the working group has recklessly proposed, via elimination of the Losing FOA and insistence on a pull system.

4. Even with a push system, a losing FOA is desirable, in the event that the control panel is compromised (i.e. if an attacker's PTID is input into the control panel at the losing registrar/reseller, etc.). Why is this so? Even with a bank wire transfer, if I made a large transfer request, that would trigger additional verification procedures at the bank before the funds were moved, even if I was able to successfully input everything online. As the SSAC report rightly stated, these additional checks are essential for critical

changes, and they explicitly state:

"Treat transfer attempts as a security event (check and re-check)."

(as I noted on page 39 of my August comments) Given how poorly most registrars are designed (not necessarily confirming all critical changes), a losing FOA, at least on an opt-in basis, would be an appropriate counter-measure to protect against certain attacks.

5. The working group at times seems to think they have a communication problem, that the public doesn't understand their proposals. That's incorrect. We actually **do understand** the enormous negative implications of their proposals. They don't need to be explained better, as we already understand them.

Instead, their proposals need to be radically altered, in line with the feedback that's been provided, to actually listen to the concerns that have been expressed.

It would be better to do absolutely nothing, throwing the entire report in the garbage bin, rather than to adopt the recommendations as they stand. There's no way to sugarcoat it.

6. I encourage the working group, if they're not going to listen to registrants (which appears to be the case, given I've not been invited to directly participate) to at least seek out greater involvement by the entire SSAC. Hopefully with enough eyes on the problems (not just those from the registrar constituency), that would encourage real debate and real solutions.

J. CONCLUSIONS

In conclusion, there is a lot wrong with this working group's report and ongoing deliberations, too long to summarize briefly. The public deserves more than mere lip service during an ICANN75 meeting. We deserve active engagement throughout the remainder of the working group's efforts, especially given the unbalanced participation at present.