Board with Technical Experts Group
24 June 2015
ICANN 53 - Buenos Aires

DPRIVE Effort /Encrypted DNS

>>WARREN KUMARI: Basically what is this is trying to protect is over here between the user and her DNS server.

So, currently, if somebody is able to monitor a link, this is what they see. This is a simple DNS CAPTCHA with a tool called tcpdump. It is very clear. This user is looking up aa.com -- sorry, aa.org, Alcoholics Anonymous. Potentially something they don't necessarily want to share with anybody who can watch their queries go by on the wireless, on the wire, wherever.

What we hope to get to is an encrypted DNS.

What a pervasive monitor or an attacker would be able to see now is the user is looking up some sort of DNS but it's basically just a big opaque blob. This is actually a CAPTCHA from an example implementation that we have. So other than this being interesting and cool, why might ICANN care? So initially, DPRIVE is focused on encrypting from the stub resolver to the recursive operator. This has some implications for people who run recursives. Mainly ISP constituency sort of the main folk will be affected by this.

Once we get that actually deployed, we hope to have a second phase which will actually encrypt from the recursive server to the authoritative server. So, basically, most people who participate in ICANN are potentially going to be impacted by that. So, you know, root server operators, TLD operators, anybody who runs a name server might have some work to do here.

The impact is expected to be fairly small, a little bit more TCP, a small bit of overhead from encryption. But we believe that this will give a large privacy win and so is well worth the effort.

And questions? And as I say, I went through this really quickly so we had time for some. I hope it wasn't too fast for people to make sense what was going on.

>>ASHA HEMRAJANI: Thank you. Thank you, David. So you were a bit too fast -- Oh, this is Asha. My name is Asha Hemrajani. I'm on the ICANN board.

You were a little bit too fast for me on the part of why ICANN would care about DPRIVE. So could you explain that again slower, please. Thank you.

>>WARREN KUMARI: So the first phase doesn't really affect many ICANN participants. People in the ISP constituency who run recursive resolvers, they might have a little bit of work to do, you know, potentially deploying new name server code, potentially slightly more load for those folk.

Once we complete the encrypting from the sort of client machine to the recursive resolver, we hope to do sort of a second phase of this work. And that will encrypt from the recursive server at the ISP to all of the authoritative servers. So it will then be encrypted from the ISP to the sort of TLDs, for example, or to anybody who runs a name server.

This will give them the advantage of hiding what queries are hitting their servers which is potentially really good for them. I mean, that's business intelligence stuff. But they will potentially need to deploy new code. There might be some additional load from doing the encryption to support this. This is likely to be many years out. But, you know, this is supposed to be a forward-looking group, provide some heads-up of what things might be coming along.

Unfortunately, at the speed that the IETF currently moves, this is likely to be many, many years but, you know, potentially different code, a little bit more load, things like that. Nothing you need to worry about now.

>>DAN YORK: Dan York with The Internet Society.

One criticism or one question I've had asked is about -- and this is another reason why it may be of interest to ICANN is, is various law enforcement agencies and others who are used to using monitoring of DNS queries, et cetera, this would hide that from those agencies.

I don't know, Warren, if you have had much dialogue with any of those, that constituency or things. But this is one impact it would have.

>>WARREN KUMARI: Yep. The queries would be in the clear on the recursive server. And so if law enforcement comes along and follows due process and serves the correct paperwork to the recursive server, they will still be able to get access to this. That's definitely -- part of that is that's kind of what the point of this is.

>>STEVE CROCKER: So to pick up the thread from the last question, if you can observe essentially all the information there is to observe at the recursive resolver and you want to do it surreptitiously, not with a court order, how effectively is this entire thing, assuming it is all implemented, compared to having it all be open? How much -- how much of the attack surface have you reduced? And how effective is it? Because I just take the devil's advocate position here, that there's a clear roadmap

that tells the other side get ready to put your probes at the place that this is leaving alone. And they'll be ready for you long in advance.

And it's the queries that tell you everything. You don't need the answers, right?

>>WARREN KUMARI: Yeah. So, I mean, in order for people to see the things that are being queried, they would need to instrument the name servers themselves. They would need to go along to everybody who runs recursive name servers and say give us a stream of all of your queries.

Hopefully, the recursive name server operators will push back and say, you know, I will give you the queries for this user if you come along with a subpoena for this user, not we will give you everything.

If we end up in a world where we assume that all of the recursive server operators and all of the authoritative server operators, and all of the Web server operators just routinely provide all data to anyone who asks, then we've sort of lost. And we might as well give up and go home.

>>STEVE CROCKER: Let me ask two other questions about implementation. I'm usually in favor of crypto and the cost of it, the computational cost. When people say, oh, it is too expensive, I tend to take the other side of it.

Nonetheless, what's the -- it's not just the computational cost but how effective is it if I can watch a large number of queries all of which are being encrypted with the same key. I got to worry a little bit about the strength of this. That's one question.

The other is zone cuts disappear. And so -- if I understood this right, and maybe I've got this wrong. But where you had sort of mildly said a few extra lookups, that's because you're not getting the efficiency of multiple layers, multiple pieces.

Do we have any statistics broadly in the Domain Name System about how effective -- how much efficiency comes from that?

>>WARREN KUMARI: So the way -- at the moment, the way it looks like we're going to be doing the encryption is clients will first set up a TCP connection and will run TLS over this. They will then pipeline queries over the TLS over TCP for a set amount of time, you know, five seconds, ten seconds. And then they will tear this down.

That sounds fairly expensive except that we have a fair amount of experience from tune-in Web servers to deal with just this sort of thing. You know, random person arrives, sets up a TCP connection, does a few queries, you know, a few exchanges and goes away.

So, yep, there will be some overhead. We have a couple of drafts. One of them is actually DTLS over UDP, which means you don't have the TCP setup but you have to do a bit more work on each query.

We don't yet know exactly how expensive this is. There have been some tests done. Still slightly unclear.

On the zone cut issue or the QName minimization, how much this is going to help or hurt, how many additional queries, we do have some stats. I can't remember them. However, one of the things that you potentially -- that offset some of the additional lookups is because you're doing a query per sort of label, you can potentially do a much better job of doing an X domain, understanding where a domain no longer exists. So as sort of a classic example, the root sees a lot of something something dot Belkin, something else dot Belkin, something something something else dot Belkin.

If you can cache the fact that dot Belkin doesn't exist, you now don't have to send those junk queries. If you can cache the fact that example.COM doesn't exist, then you can automatically stop throwing away food.dam (saying name). Hopefully that offsets it potentially. There's a lot of trickiness there with things like empty non-terminals, making sure that you do that correctly. The standard example for where this might cause an extra lookup is co.uk. Both of those have fairly long TTLs, so not many --

>>PATRICK JONES: This is Patrick Jones, but I'm asking this on behalf of Dave Piscitello who is following remote and trying to put it into chat.

Dave asks: Can you share how you see DPRIVE affecting passive DNS collection or investigating criminal domains?

>>WARREN KUMARI: I don't know if I heard the full question, but I think this goes back to -- go to the recursive operator where this passes them in the clear, serve them with paperwork, and they will probably be more than happy to help you. If you just want to watch the queries go by, sorry, that's not likely to work for you. Did that actually answer? Okay.

>>DAVID CONRAD: Any other questions? I will take the chair's prerogative and ask a question.

So you indicated that the pros of this were privacy and privacy. Is there also a potential benefit in reducing a risk associated with someone compromising one of the authoritative servers that is being queried to be able to respond with an arbitrary response for the full query? You know, someone, say, breaks into a com server or the root server, even more entertainingly, you could in theory imagine a

scenario where you would respond to the entire query with the redirection to your favorite man-in-the-middle attack. Would query name, QName minimization actually help in that situation?

>>WARREN KUMARI: I would need to think about that a bit more. I don't think so. Where it would help, though, is people inserting additional answers or doing cache poisoning on, like, a Starbucks network, for example.

If you go into the Starbucks network and make a query from your name server, from your machine, an attacker can see your query go out and so can forward your response. He has all of the information. So he doesn't need to do any sort of clever hacks. He knows the query I.D. He knows the query name. He can just respond. So this gives you sort of the bit that DNSSEC didn't.

DNSSEC gives you sort of the authentication of it, if everybody runs it. With this, you get the privacy potentially. You also get some of the good things that DNSSEC is supposed to provide as well.

>>DAVID CONRAD: Okay. And the status of this within the IETF right now?

>>WARREN KUMARI: QName minimization is in IETF last call. As far as I remember, it is past IETF last call. And the IESG is just faffing with it.

And then DPRIVE is a working group which got spun up, I think, two meetings ago. We've adopted -- we have one document which is with the IESG and two documents which we've adopted. We have a couple of toy implementations as well.

>>JONNE SOININEN: Just a question. How does this work actually with captive portals? Does it actually break them?

>>WARREN KUMARI: That's an interesting question. I'm running a buff and prog on captive portals. It breaks some of the DNS redirection, which some captive portals do. Many operating systems do some probing to discover if there's a captive portal. And so potentially they would try and do a couple of DNS lookups, encrypted. If these don't work, you know, fail back to a DNS lookup and see if a captive portal happens and then try to talk to it.

There is also a draft which Joe is responding, Joe Abley, which is an extension to DHCP. So when you first connect your machine to the network, DHCP tells you that you're behind a captive portal and how to contact it. The expectation is then applications would start up in a sandbox, and they would know not to do real DNS queries until you've satisfied it.

The captive portal buff is supposed to -- or a working group is supposed to do sort of a full interaction thing so you can talk to it. Good question, though.

I think Paul Wouters had a follow-on.

>>PAUL WOUTERS: Sure. Paul. So it doesn't break it much worse than if you're running DNSSEC on that same laptop at the same coffee shop. So it is the same problem basically. We have to solve it anyway, whether we do query minimization or not.

Security Challenges With Web Standards and DNS

>>WENDY SELTZER: Wendy Seltzer here. I'm the technology and society domain lead with the World Wide Web Consortium and W3C's representative here on the technical experts group.

And so I very much appreciate the opportunity to bring a question more at the brainstorming phase of discussion here, which is some of the questions around domain name persistence and their implications for Web layer technology.

W3C works at the -- the application layer. We do Web standards. And most of our interaction with ICANN and IANA functions is through IETF and -- as a user of the DNS.

There are a couple of places where we have more intertwined connection. One that I'll hope to address at a later phase of discussion is the security considerations where a lot of Web application security relies on the same origin policy, which is a domain-name-based segregation of resources on sandboxing and relies on names and the boundaries between names, and so I'm interested to delve further into the SAC70 report on public suffix lists.

And then we also rely on domain names for lots of the operations of the Web. For persistent identification, information location, linking, Web apps, all the things that make the Web a Web.

We get challenges when domain names lapse, when you've been linking to something, using something, and suddenly it goes away. And that can be a challenge for data on the Web. Articles disappear. Citations disappear. Harvard authors found 70% of links in three law reviews they studied were dead or no longer pointed to the right place a few years later. You get problems with Web mash-ups. You're trying to pull in multimedia or images and they drop out from under you.

Challenges get deeper with data mash-ups. Linked data on the Web depends on public data sourced from multiple sites. Ontologies that may be hosted elsewhere describing that data and Web applications that then mash those up to enable computation or data visualization.

End up with lots of Web templates, document-type definitions and namespaces retrieved from one place to fill in on a page, and script libraries and infrastructure services.

And, you know, because, you know, why copy when you can link? Better to refer to the canonical version of something, use the cache that a user may already have available, than to pull it down each time you're using the same resource, except when that goes away.

And so we have -- on the W3C level, our Web applications security group is working on mechanisms to help in checking the integrity of those resources pulled in from other locations, so that the sub-resource integrity spec is one to specify the hash of the resource you're expecting when you pull in a script from jQuery or an image from a Tumblr site, so that you as page author can give instructions to verify that what you've pulled down there is what you were expecting and block the loading of content if someone's managed to inject or change what you -- what you get there.

Security mechanisms can go away if domain name lapses or isn't maintained

But any of these things can -- can go away if the domain name is allowed to lapse or forgotten about, and for some of these services it can be really embarrassing to see your images replaced by somebody else's images.

Heinz was recently embarrassed by putting a URL on a Ketchup bottle that they failed to review and users who snapped the QR code got images of pornography instead.

Or they can range to the truly challenging for software applications and Web apps that depend on automatically linking resources from multiple sites and combining them. They can just fail entirely because those resources are no longer available.

So we can take various approaches to -- to dealing with these challenges.

For the purely text or reference resources, where you're just trying to find the information and not de-reference a link, a stewardship model like what Internet Archive does by crawling the Web and archiving lots of the stuff they find, or perma.cc which takes a more active response to the problem of academic reference link-rot by inviting people to archive content there and holding references, so that you can then either query the archive or just give the perma link in your reference.

But these are more challenging for some of the automated retrieval that software is doing expecting to find something at that link.

So we could consider -- so one question is, are there any classes of domains for which it would be reasonable to extend the 10-year maximum registration period.

ICANN, through its contract, says you can register domains for 1 to 10 years, and is it -- would it ever make sense to enable somebody to make a longer registration so that they didn't forget about that registration later on.

Of course there are some challenges to that.

First off, of course, there's plenty of non-domain-related link-rot, and we'd still need to solve the problems of redesign and technology changing and business changing and people simply deciding to decommission domains or particular resources.

There are legal requests for removal.

There are -- none of those would we be solving.

Even if we say the goal is to extend the life of a resource that somebody meant to register and forgot about, simply extending the life of an old resource that somebody forgot about could introduce security considerations of are we supporting somebody's defunct site and if they really forgot about it, are they also forgetting to patch something and is the script you're trying to include from there so horribly out of date that you'd be better off if it didn't resolve. Are we opening the door to spam or malware if these become attractive targets for takeover.

And are there -- are there other options that -- that we could consider like encouraging use of some reserved or special use domain at the top level for resources that had particular persistence concerns.

So with that, as I said, I'm inviting brainstorming and thought about the -- this persistence challenge.

Relative stability of "set and forget", vs. constant maintenance of DNSSEC expirations

>>STEVE CROCKER: One of the great values of DNS has been the fact that you could set it and forget it, basically, and that's what an awful lot of people do, and so you get stale links or you lose control of it because the system administrator is long gone and it's in a file somewhere and nobody knows how to maintain it.

So there's risks associated with the fact that it's entirely stable from that point of view, aside from the -- from losing the registration that you pointed to.

DNSSEC moved us into an entirely different domain where things expire and it requires active control to keep them in -- keep them going. It seems to me this falls -- this is addressing something that is roughly in that same area.

So what's going through my mind is, stepping back from this particular problem and the DNSSEC in particular and lumping them into a broad class, for things that require active maintenance on some schedule from a protocol design and operations point of view, perhaps we have -- we need to develop an attention factor to that, and when we design the protocol, if it's got a requirement for active control, we should also be designing the maintenance processes that go with that so that you can maintain that.

And this would apply to keeping a domain name registered every 10 years and it would apply to keeping DNSSEC records signed every 30 days, or whatever it is.

So that's just a thought that's going through my mind in terms of looking at a broad syndrome here.

>>DAVID CONRAD: So I actually have a couple of questions.

First, with the idea of persisting domain names, has there been any research on sort of what a good time frame would be?

You know, right now you can buy a domain name for 10 years. Steve, would you happen to know why a limit was placed on the length of time a domain name can be owned?

>>STEVE CROCKER: The answer is undoubtedly nontechnical because we only have technical people here and we don't have any answers, right?

>>WENDY SELTZER: Well, and I think for the -- for the technical resources that people are trying to link, a regular maintenance obligation makes a lot of sense, and something that is an even more frequent tickler about, you know, is this really up-to-date could be useful.

For the data-on-the-Web community where they want to assign a persistent identifier and have you be able to get to the same piece of data every time you link to that all the way into the future when we're using these things in the archives of a hundred years ago, they have shied away from URLs even though those would look like the right vehicle because of the sense that at some point, the lease will run out and our identifier will no longer work.

>>WARREN KUMARI: Yeah. So there are a lot of things that I dislike doing. Expense reports are one. Renewing my domain's another.

However, I always just do it for a year. And the reason for that is, then I'm vaguely kind of sort of likely to remember that this needs to happen.

If I do it for 10 years, there's no way I'm going to remember in 9 years and 11 months that something needs to happen.

If I did it for longer, you know, there's absolutely no way that there's still anybody working at whatever company who understands any of this process.

Do you think potentially that ends up causing issues or do you think that most people are just way more organized than me?

>>WENDY SELTZER: I think the -- sort of the process parts of those issues are ones that we solved in other domains, like there are leases that run for 99 years plus and people somehow manage to find the landlord at the end to renew the lease, so we -- we probably could find ways to make that happen.

>>JONNE SOININEN: So, Wendy, one question that I had in mind is that like you did mention in your presentation that you said that there are many other reasons for link-rot than just domain names.

Have you actually looked at how big a part is actually the domain name stability of this? And how much is that people just go and change their Web pages and change the URLs of stuff that used to be there?

>>WENDY SELTZER: I have not -- I think in the Harvard study, they found about 5% of the domains no longer existed. And some of the domains that did exist were now entirely different sites. And this is things that law review authors had chosen to cite to which may or may not be a useful sample of what's out there on the Web.

>>JAY DALEY: Thanks. So we do ten-year registrations. A small percentage of people take them up. They're fine. Accountants love them.

Second thing, though, I'm not sure that even ten-year registrations would solve the problem of persistence for things such as similar to URNs. You almost need a permanent identifier, don't you, that's going to be there? And I think that's an entirely new thing, and it needs a different economic model than the one we have if that's going to happen, I suspect.

URNs, URIs, DDDS, and reg.int

>>DAVID CONRAD: I remember vaguely -- and maybe Patrick might have a better recollection than I that there was this vast architecture, the DDDS, dynamic data directory system -- I forget what it stands for -- that had a vast array of URNs or TURIs to URLs. And the intent was to allow, as I understood it -- to allow a layer of indirection so that resources -- the actual URLs could be swapped out and varied depending -- with essentially a permanent identifier that would then reference a name or a name that references an identifier, which then would reference the locater. Did I hit my head once too many times ages ago?

>>PATRIK FALTSTROM: I think you had about the same number abstraction layers in your answer that it is in that design actually.

But I think -- let me say differently. Both the DDDS algorithm that is used in, for example, the ENUM and a few other places like generate URN resolution, that is one way of reaching that kind of abstraction layer.

We also in the IETF have the SRE record. We have a combination of SRE records to get together with well-known URIs. That is, that Web browsers at the moment are pushing forward.

And then we also have the URI resource record that also reference sort of a domain name to URI. So there are multiple ways so more collaboration -- collaboration is ongoing. We can do better between the IETF and, for example, W3C and their browser industry.

>>DAVID CONRAD: Right. I also wonder -- some of the things that Wendy mentioned, the permanent source libraries and that sort of stuff, it strikes me that it may make sense to explore with the IETF or actually rather with the IAB whether ARPA might be a good permanent location for the sorts of infrastructure related things.

I would note as a perhaps humorous historical reference that there is, in fact, still a registry that uses the DNS for permanent identifiers. The reg.int domain that's used for algorithms -- yes --

[Laughter]

Just had to raise that one as humorous. But it would seem that there may be infrastructure within the DNS and potentially relationships between W3C and ICANN, ARPA -- sorry, IAB and the IETF that could, you know, make a -- sort of a permanent home for the purely infrastructure, you know, identifier kind of objects. Something to explore.

>>PATRIK FALTSTROM: We in the IETF also have been using the urn.arpa for that generate kind of resolution using the DDDS algorithm that you're talking about.

But, on the other hand, these are things that the IETF used many, many years ago and before. And I absolutely do not recommend use of it. But there is -- I clearly hear the need for some kind of cleanup there.

>>STEVE CROCKER: This is a kind of humorous piece of irony. Even if you have permanent registration for the algorithms, the algorithms have finite lifetimes.

>>DAN YORK: Dan York. I would also say, I mean, we are talking about another abstraction layer really for this type of thing. And I think it's -- as we're talking about this, it occurs to me, too, there are also any number of ones that are currently in existence. If you think of the URL shorteners that are out there right now, not that I'm saying we should use those, but an example of the type of things that people are doing with some of those -- obviously many of those operated by commercial companies.

But there are those -- you know, the TinyURL, bit.ly, all of those types of spaces, that are again looking at how do we create an abstraction layer on top of URLs and URIs. So it is just interesting to add into this feature another mechanism.

>>WENDY SELTZER: I think there is -- the abstraction layer question is a little bit orthogonal to the ownership question of the -- of who -- can I assure that ultimate control of this name is either with me or that it's dead, for example, that it's not going to go to somebody else who will subvert what was once at that URL.

>>DAVID CONRAD: And perhaps in the next round of new gTLDs, the restriction against having numbers and TLDs could be removed so you could get .W3C. Ha, ha, ha.

>>WENDY SELTZER: Or we could just keep it reserved as it is.

>>DAN YORK: David, I think it is an interesting challenge that I think we do have because I've written any number of documents that reference DTDs that are in my XML docs that, you know, if those particular sites go out or change, then all my XML docs basically can't work.

So there certainly is a challenge there. I don't quite know how we do that in the existing thing, but I agree with Wendy. It is useful to consider what we do as we continue to go down this grand experiment we're on.

>>WENDY SELTZER: In the meantime as a humorous side, make sure the W3C stays solvent so at least the w3.org DTDs continue to resolve.

>>JAY DALEY: Sorry. Just a thought. I wonder if we're thinking about the URL or the URN to the DTD is the persistent thing. Maybe it should actually be a signature on it so that we no longer actually care where it is. But if the signature becomes the persistent thing that we know is there, it is the signature that we are actually putting onto our XML documents and then we are getting it from wherever we can find it at the time. Just a thought.

Older Business

>>DAVID CONRAD: Okay. Any further questions? Thoughts? If not, then we will move into actual old business. This was the result of a discussion that occurred back, I believe, in, was it, L.A. or Singapore? L.A. And this is addressing the potential lack of continuity in DNS service during domain name transfers.

Potential lack of continuity in DNS service during domain name transfers

>>JIM GALVIN: Thank you, David. So, just as some context to add to the context that you just offered in terms of this being old business, back two meetings ago, Francisco Arias from staff had volunteered or been volunteered at the time to consider this issue with me. I had brought it to the floor at that time and was making a sweeping statement about the lack of support for DNSSEC and, in fact, the effect of that in terms of continuity of service.

And so Francisco and I have spent some time together. I believe he's sitting in the back over here. I want to acknowledge him in the development of these slides and continuing forward with this.

So the first thing to notice and observe is what we have today in terms of documentation. What I want to also add here in way of preface is we had some discussion on the mailing list for this group that people didn't really want to talk about DNSSEC technical issues, observing that DNSSEC has a home where it gets a lot of attention, particularly DNSSEC workshop.

So the context for this presentation really is about highlighting some potential policy areas that could use some attention. So I will mention certain technical details but only in the context of trying to highlight or indicate where we could do a better job with policy.

So today there's a fair amount of detail documentation about how registrations can move between registrars. And there's an inter-registrar domain transfer policy. There is a new one which was actually just published in January of this year.

But the thing which is interesting in all of that documentation is there are, in fact, no DNS or DNSSEC transition requirements for registrars. The reason why this is important to the ICANN community in general is you can make this split and observe that a large proportion of the domain community, registrants -- I think it's safe to say more than 50%, but I hesitate to put an exact number on it. I don't have any actual facts to back that up with -- depend on their registrar to provide them DNS services. And so it's a bundled service from registrars on the gTLD side to provide both registration services and DNS services.

There is a much smaller percentage of the community which don't depend on their registrar for DNS services. But, nonetheless, they do need certain DNS features or DNS support from a registrar in order to effect any transition that they want to make between providers. And that's what I want to highlight here. This lack of requirements actually does affect the community at-large. And the opportunity here is, given our environment of registries and registrars, I think there's an opportunity here for the community to consider some transfer requirements.

The potential issue, of course, that this -- the lack of those requirements creates is that there's a potential risk of lack of continuity of your DNS and DNSSEC services when you're trying to do a transfer.

These in particular are the potential harmful behaviors that are seen at both registries and registrars that contribute to a lack of continuity in your DNS and DNSSEC services.

The name deactivation at the top is an interesting one. You find that when a registration that's bundled with DNS services is moving, it is actually -- it does actually occur with registrars that when they see that transfer request, they have five days for that transfer to actually take place. And some registrars are known to remove that name from the zone on that day. And what that causes is after TTLs expire, that registrant no longer has DNS services. And so for a day or two or perhaps three, depending on things, they have no DNS service. So they have no Web site, no other services, whatever they're offering until the registrar they're moving to actually then has the authority to reset name services.

The last item here is equally interesting. I mean, what you're really looking for in all of this is for the old provider to continue to offer services until the transition has completed. And that's probably the most significant requirement that you would want to have happen if you want there to be no loss of continuity.

Intent here is to highlight the fact that these are the consequences. These are the things that can happen given that you have no requirements on the old service provider, the DNS service provider. You do have the opportunity to have no DNS services. And, thus, none of the services that you would offer as a registrant,

whether it's just your own content provisioning, any other -- whatever services you offer at your domain name would be unavailable because your DNS is unavailable.

The reason for the five days there is because the transfer period is that five-day period. Typically, the actual downtime as a practical matter would be more like a day or two because of the way the TTLs will time out over time.

The one item to highlight here is the last column where it says DNSSEC goes insecure. So when DNSSEC is present, the situation is actually worse than when you don't have DNSSEC. That's because what you've done when your DNS doesn't become available is created a security incident, which is very different than the robust and stable DNS that we're used to under ordinary circumstances. People would ordinarily figure, well, the DNS is down. I can't get to the server. It comes back eventually and things just move on. And for a large part of probably many registrants, all of that's okay, or even Internet users. It's not highlighted to them what really happened if they suddenly can't get at a service they couldn't have. They just sort of move on once it comes back.

Unfortunately, when you have DNSSEC and your old service provider goes away, you create an incident which under ordinary circumstances -- you know, browsers is probably the easiest application to talk about -- they suddenly, you know, would put a pop-up up that says, Hey, wow. This is insecure. Don't know what this is. Can't find the server. Can't find the keys. Can't validate the signature. You are doing something really bad here. You don't want to do this. That has the effect of creating a bad impression about the place that you're trying to get to. And, of course, whatever business that's using that particular domain name.

To touch a little bit on a solution set to be considered here, is there are two things that would make a dramatic difference in ensuring that you had continuous DNS service and DNSSEC service.

The ability to be able to pre-introduce your new key into the old service provider is important, owing to the behavior of certain resolvers that are available out in the wild. So whatever your new provider is going to be, they're going to resign the zone. You're going to do a key roll, and you want the ability to put the new key out there and have that visible from your old service provider.

But probably more importantly in the large, what you need is to create the details for a cooperative transition between the old service provider and the new service provider. And although this is more or less -- the details of this are understood in the technical community. So if it was just between two DNS providers themselves, it would be fine. The introduction of DNSSEC creates this need for the new key to be moved over. And so you need access to the registrar to make that happen.

But for the large part of the market which depends on their registrar for DNS services, a consideration here is whether or not we want to ask those registrars in the same way that they have to manage the transition of a registration in a certain way, you know, should there be some requirements on them to manage a DNS transition in a certain way to ensure that we don't have this discontinuity in service.

So here -- this is my last slide, the next slide. There are three questions that I bring really to this group to consider. The first is whether or not this is a real problem. I think from a technical point of view this is actually a real problem. There are real issues here, and there is a real gap that happens.

The more important question is probably whether or not this is something we want to solve. Going back to what I said in the very beginning with the first slide, if you look at the community from two very broad categories, even today it's fair to say the large part of the domain registrants depend on their registrar for their DNS services. And they're used to the idea that some set of things come and go and there's probably no dramatic effect. There may not be one, so maybe we don't need to really concern ourselves with solving the problem for that big part of the community.

On the other hand, there are that smaller community, high-value domain names, if you will, which are very concerned about being able to maintain continuous service. And we all have our favorite for what is a high-value domain name, your favorite large service provider that does things.

And the fact is they cannot effect a continuous transition, or maybe they can be because they can lean on people. But there is nothing in the system that ensures that they can get continuous service as they move their name from one side to the other.

And so the questions to consider here are whether or not we want to solve this problem and if so, what would be the path to do it and how would we want to approach that. Thank you.

>>JIM BASKIN: So Olafur has -- the technical solution is one side of it but whether or not -- his solution has its own set of issues. So I'm not sure if you want to talk about some of those things here. One of the problems that he runs into is part of what he's pressing for is access to registries in order to effect changes directly into the registry. That, of course, would be incongruous with the model that we have which is that you can't have access to the registry because all of the legal instruments in place would not allow that. You would need to be able to do this through a registrar and not a registry.

>>STEVE CROCKER: So for the benefit of everyone else here, Olafur used to work for me. Jim and Olafur and I and a few others have all been through this loop in great detail when we were shaking down the process of moving from one registrar to another and from one DNS provider to another in order to make DNSSEC work.

There's an interplay here of, as you say, access to the registry. It raises sort of policy issues or what the model is, but as you said right at the beginning, this is fundamentally a DNS operator problem which is just immediately outside of the fence that we have and so there's complexity because there are no -- there is no venue for what the rules are for DNS operators except I would say in the IETF basically. We could step into it a bit, but I don't think we could own it. Would have to be kind of in concert, would be my reaction.

>>JIM GALVIN: So that's why I wanted to make the careful distinction between whether or not you have your DNS services at a registrar or even if you have your services at a third party what's important is that the system that we have here today, okay, that DNS providers have a dependency on us. And if we have an obligation to a secure and stable Internet and a secure and stable DNS, to ensure that that's true, we do have a role here, I think. We have the opportunity to address a particular problem that you're right, DNS operators who are not an integral part of this community, they have an issue. And ensuring that they can solve their problem in some way and in particular being able to get your public key information in through a registrar into a registry is really a detailed problem. It's a detailed part of the solution that would be important here.

>>JAY DALEY: We did all of this about two years ago in .NZ. So we introduced -- I didn't do this. It's our policy body. We introduced the contact with DNS operator and specified what they were and we then gave special specific requirements for the losing DNS operator and the gaining DNS operator and made registrars who take DS records responsible for some of that part of the process.

Now, it's easy for us because we have a smaller number of registrars, 80-something, and it's a cleaner, you know, system to do. But it's basically just exactly what you've been saying. We just had to implement that is the only way to ensure that they transfer correctly. And for all 25, I think, signed domains that we have, we now have that security that if they want to move registrars it's going to go well.

>>JIM GALVIN: So I think the only thing I want to say is bring it back to what I said in the beginning, okay? There are solutions to this problem and there are multiple things that are, you know, moving forward. There are other ccTLDs that have created solutions to this problem and have solved it. This -- what I didn't emphasize before is this is an issue primarily in the gTLDs because there are certain ways in which they do things. So it does have an effect on the solution set that's available. And then the other side of it is, there is no -- the question that I'm asking is, do we

want to solve this problem and ensure that continuous service is available to the community? I mean, right now it's not there. And so there are some things that the community could do, if it wants to take on that priority. And so that's really the primary question here. Not about whether or not there are solutions available. There are details to talk about as we identify what we might consider best practices for exactly how to do the things that need to happen in the gTLD community, but the question is, do we want to.

>>MARGIE MILAM: This is Margie Milam with ICANN staff. I was involved in the negotiations of the 2013 RAA and Jim, I don't know if you've taken a look at the additional registrar operations specification and would that language somehow help if a solution was to be -- you know, attempted to be followed? The language that I'm thinking of is something that says -- I won't read the whole thing. You can all look at it, but there's a section on DNSSEC and it says that the requests shall be accepted and processed in a secure manner and according to industry best practices. So there's a placeholder for -- without actually defining what -- you know, what is involved there, and perhaps that's one avenue that could be followed.

>>MARTIN LEVY: This is Martin on the phone. Can I get some time?

>>JIM GALVIN: So Margie, thanks for looking out and putting that placeholder in there. I think to actually make use of it in any realistic way two things have to happen. One is the part that you've, I think, anticipated which is the creation of those best practices. And then the other is an operational test that those best practices have, in fact, been implemented. So it's not an argument. I'm just sort of -- to help everybody understand what it would really mean to breathe life into that, that if they say they've implemented best practices, there ought to be a test where you register something and then you move it purposefully to a point of making sure that those things work. Trust but verify kind of thing.

>>JAY DALEY: So Jim, yes, you do have to do it. Anything else would just be, I think, leaving people in an unacceptable state. It's not easy. There is no simple solution. You have to do something that just puts difficult requirements on people. And the RAA text I don't think is sufficient because you also have to get into the operations of DNS operators separately from registrars. You've got to extend beyond just registrars.

>>MARTIN BOYLE: Yeah, the placeholder -- so Martin Levy at CloudFlare. The placeholder aspect is interesting, although best practices at the moment would be very poor practices. However, the reality is that if you are just a DNS operator and not a registry or registrar, you don't really have a voice to effect this issue and to move it forward with inside ICANN. However, if you are -- well, either a registry or registrar but you are the DNS operator of a zone outside of that, again you don't have permission or essentially mapping that to a voice to move this issue forward.

So this is, in fact, actually not only a technical issue but also an issue of who has the ability to effect this. And in Jim's presentation, it covered a lot of that except for the subtlety that you can easily be a DNS operator who is not party, except for the nameserver pointing to you, which is a vital thing, gives you nearly absolute control. You're not a party to make any changes. That's the point I wanted to bring up. Thank you.

>>JIM GALVIN: So there is an important distinction to make, Margie, going back to your comment about the specification that you had highlighted. That particular clause in the contracts is referring to normal operation. So you're at a registrar and what your obligations are as a registrar in terms of how you support DNSSEC. It's possible that that's a placeholder that, you know, maybe there is an expansion opportunity there. When you do the best practices, you can add there, you know, requirements for when a transition has to occur. You know, there's some opportunity there. But what's important is again, when I started on the first slide, there is a separate policy and documentation with respect to transition and transfer of domain names, and so one could make the case, too, that more logically any policy change for DNS services would go there, too. So I just want to draw that distinction.

I think just to say one more thing then to add on to what Martin said, you know, that really is the point here. We tend to focus a lot on registrations and domain names and, you know, obviously that's important. We don't really have any DNS requirements. The DNS requirements -- DNSSEC requirements that are currently in the contracts are all about how you work with a given registrar. But if you want to move, you know, there are some technical issues. And the question is, do we want to address that. You know, do we want to agree that security and stability, you know, really is a commitment. And we want to ensure that the Internet community at large has the opportunity for continuous services. And in that respect, I think that there are a couple of -- there are a few particular requirements that we need to pass down the chain because registrars have a role in supporting the community at large. Even if they're not going to offer the services directly themselves, there is a requirement there, if you want to support the community at large. Thank you.

>>WENDY SETZER: Yeah. Thanks, Jim. A question and comment. A question, is this something that a knowledgeable registrant can fix on his own from any registrar? Is there always the option to choose an external DNS provider and sidestep the transition problems which doesn't necessarily address the -- the questions of sort of our role toward end users of those services. But that would be sort of -- is the problem unfixable without additional thought to the transition process?

>>JIM GALVIN: Yes, the problem is unfixable without the participation of a registrar in the gTLD community.

>>WENDY SELTZER: Okay. And then the comment, then it sounds like a place for where we should spend more time thinking about how to intervene because it's -- you know, at the point of exit from a registrar is the place where sort of the registrar's incentives are not necessarily to help fix the problem, and so we might want to give some stronger nudges to improve the transfer process.

>>JIM GALVIN: And I should add one detail, just to be fair. I've been throwing registrars under the bus, if you will, quite a lot. But I do want to be fair here. There are some things that registries might be obligated to do also that should be considered if you're going to really look at this problem. In fact, the TTL issue is the number one thing that comes up. Registries have, you know, rules about the TTL on DS records and NS records in particular and -- and Glue records and that, of course, affects your ability to transition. So there's -- you know, there's some opportunity there for requirements to be put down or, you know, guidelines for what should happen best practices. And then that, of course, causes a relationship to happen between a registrar and a registry because you need to know in order to make that change for a registrant to do it. So anyway. Didn't want to leave -- exclude registries from this mix here.

>>DAVID CONRAD: As a -- one question is whether or not you see this as a security stability-related issue and if so, whether or not this -- do you believe this would be appropriate for SSAC to provide an advisory to the ICANN board on, advisory recommendation.

>>JIM GALVIN: I don't think so. I mean, I think that SSAC could highlight the issue but in a sense I'm sort of doing that here. Isn't that part of the reason for this group's interaction with the Board and saying that an issue exists. It really is a community problem. You know, would you -- I mean, SSAC could highlight the fact that there's a problem, I suppose, is a potential response and I really should defer to Patrik as chair as to how he wants to, you know, address that particular question. But I don't think that SSAC would want to take on the responsibility. I'll wear my vice chair hat and say that I wouldn't expect that we would take on the responsibility to propose the solution because the solution involves registries and involves registrars, the actual discussion of the details of what we would want and would need really should happen in the community. So it's not -- that's why I was hesitating and saying, it's not clear to me that SSAC saying anything to the Board is useful or important. We could just go forward and try to get the community to pick this up, if we want to.

>>PATRIK FALTSTROM: Patrik Faltstrom, chair of SSAC. In this kind of situation where there had been some design needed together with SSAC advice on whether a particular design have whatever quality that we think is needed, like you have to be this tall, which I never fulfill, or something like that, what we have done more or less successfully, different successful in different cases is that we as SSAC have

participated in, for example, work that GNSO has taken up so that we have been able to sort of help along the road on the other end, the SSAC advice is something that we're given either to respond to specific questions or we might evaluate, for example, the preliminary result from that kind of work that is done within one of the PDPs. But I definitely see that, just like James said, these kind of things might have impact on the contract and otherwise which means that should -- from my perspective also it should be done as one of the normal PDPs or otherwise.

>>DAVID CONRAD: And to be clear, I wasn't suggesting that SSAC should propose a solution. Rather that SSAC recommend to the Board that there is a potential issue associated with this particular situation and then allow, you know -- make a recommendation that, say, the Board direct staff to figure something -- or the community figure something else out or something like that. Go ahead. Ram.

>>RAM MOHAN: Thank you. I actually think that's a little too much formality that is needed here. We've just raised -- it's just been raised. There are Board members here, there are staff members here. So, I mean, at least from my point of view we should just take this as raised. We should figure out, you know, where it falls in priorities. And then either we work on it or we don't work on it and we should come back to this group and say what we're going to do with it. But I'm not sure we need, you know, to get all of that mechanism going.

>>DAVID CONRAD: And to be clear, this group's primary role is to provide input to the Board. You don't actually have to respond to us. We don't -- we're just a bunch of friends getting together and talking about fun stuff. Yes, Jay.

>>JAY DALEY: One of the things that happens in other types of technical environments is where the business people get together and agree a requirement for something and then ask the technical people to go and implement it. This, what we've really come up with is actually a business requirement, not a technical requirement. But it could -- and we are talking about solving it through business changes, through rules there, but it is potentially solvable through innovative technology, some new technology. Is it -- can you foresee a time when, through the ICANN process, we have people coming up with a specific business requirement and then some form of request going to the IETF people, you know, to put it loosely, saying can you have some technology that delivers this for us, please?

>>DAVID CONRAD: So speaking personally, my experience in the past has been when ICANN has gone to the IETF with a particular requirement, it was -- it didn't go particularly well. The IETF deals with persons, individuals, and communities of interest on particular technologies. If ICANN found a bunch of other folks who were interested -- or not ICANN, but the ICANN community developed a community of interest around a particular technology requirement, then that might be an appropriate -- then the IETF might be an appropriate venue in which to pursue that.

But ICANN itself, you know, as an organization, we can't, you know, go to the IETF and say, "Please solve this problem for us," right?

>>JAY DALEY: Perhaps, though, just a Web page of problems that we think could be solved by technology, you know, that we have a required business need for within ICANN, just so that at least we can publish that. Because this is something, I'm sure, that someone could invent some new technology to solve.

>>DAVID CONRAD: Then I think we are on to any other business, if I remember correctly. And I'm going to start. It's the chair's prerogative and all that.

Any Other Business

IANA Transition

I just wanted to raise with this community, the technical experts group, that you might have heard of something called the IANA transition. I know it's -- I don't want to surprise anyone, but there is this IANA transition thing going on. And I thought it would be worthwhile to point out -- or to suggest that, you know, the organizations that you represent or the -- your own personal interests in this particular topic, that you actually make your positions known or make your organizations' positions known. You know, we have, on the TEG, the TLG members, which are ITU, ETSI, W3C, and the IETF. Clearly, the IETF has a position on this, but, you know, I would encourage the other TLG members and the TEG members to, you know, discuss amongst -- you know, internally amongst themselves, come up with some position, and just make that position known.

One of my -- speaking personally, completely taking off all hats, one of my concerns is that there appears to be a case where a relatively small number of rather large voices have -- tend to drive a particular set of agendas that may or may not be in line with the larger community.

This may or may not be an accurate perception. It just is my perception. And I think, you know, this community, in particular, has some ability to provide input into the transition that could be very helpful in helping the community reach a good resolution.

So I just wanted to throw that out as any other business, and will open it up to others.

>>DAN YORK: Dan York from the Internet Society but actually Jim sort of addressed the point. I was just going to raise to the board members that are here that if they

are not aware of the ongoing discussion that's been happening around this space with the role of operators and others. Jacques Latour, who is in the room here, just gave a presentation earlier today at the DNSSEC workshop and also at tech day on Monday, but there is this -- I'll give you a context in which -- and concrete example of the challenge that's currently faced.

And, yes, this is a DNSSEC thing but it has, you know, ramifications here.

CloudFlare is currently in the process of looking at how do they sign 2 million domains with DNSSEC. The challenge they have is if they, as the DNS operator, go to do that, the way to plug that into the chain of trust is to communicate through the registrars to be able to get the DS records up to the top-level domains in which they are.

For them to do that, they have to interact with X thousand registrars. I think it was 4,000-something when they actually looked at something. And, you know, some hundreds of registries. And there's not an easy path to do that.

So there is a -- there is this ongoing discussion, debate that's happening. Some of it is within IETF. There are some technical ways that this can be done which turn out to be fairly simple, in some ways, but there's more policy, procedure ways that -- are -- provide much greater challenges.

So more my point here was to say if ICANN board members and others are not aware of this, this is an ongoing discussion that impacts the kind of way the current registrar/registry model works right now.

>>DAVID CONRAD: And I will note that I spoke with Olafur about actually participating in this session, TEG session in Buenos Aires, and he was unable to participate but did indicate an interest in participating in the Dublin meeting, so this will be a topic on the agenda at Dublin, assuming that Olafur doesn't solve all of our problems by then.

>>PAUL WOUTERS: Paul Wouters. IETF. From the IETF community, we've seen a few issues with protocol compliance with respect to EDNS options and recent or new resource record types where TLDs are either not answering correctly or are answering with the wrong data.

And this is going to be a problem once we -- especially in the DNSSEC and DANE working groups when we come out with these new resource record types, that they just start dying as soon as they hit a TLD server. And so we're seeing more and more of these problems, and I just wanted to raise that issue and see if there's anything that can be done from ICANN to give a little pushback on people to make sure that they implement the DNS protocol correctly.

>>DAVID CONRAD: Yes. And I know staff has been and may be in contact with Mark, but is aware of, you know, the particular issues that Mark Anders raises in the -- in relation to the EDNS.

>>WARREN KUMARI: And a follow-on to that, which I think was mentioned at a previous TLG meeting, is it's very hard to introduce new record types that do useful things if you can't enter it through the registrar interface. And there doesn't seem to be a requirement that registrars allow users to be able to put in, you know, records of new types.

It seems like perhaps that would be useful.

>>WARREN KUMARI: Not registrars. DNS hosting people who are often, you know, hosting --

>>FRANCISCO DA SILVA: Francisco da Silva, ETSI. You have made this call for the TLG meetings for -- they were welcome to contribute to this. Perhaps this would be reflected in the minutes because not all TLG members are here today. I don't think ITU-T is here. So if we have this, it will be easy, okay?

>>DAVID CONRAD: Right. And to be clear, it was a personal suggestion that it might be of interest to everyone if you would make your -- your interest known with regards to the transition.

YETI

>>DAVEY SONG: Okay. Thank you. Davey. And some of you may notice that there's flyers in your bags in this committee, and it's just an introduction of the announcement of this Yeti DNS project.

This project is aiming to build a large-scale live test bed to discover some limits of this -- the DNS root system and to provide some scientific output and some proof.

ICANN -- the background of it is very simple. It is that there's Y camping this March, and several DNS people are joined together and have a brainstorming and want to try some new ideas, and so went out to build a test bed, so that's -- and we've issued three nodes. Now, there are seven of them. So the basic idea, research engine on this test bed, is simple. It is that -- how it survive in the v6 owning environment, how we can change the DSK and KSK more frequently, say to test up say to 5, 7, 11, and also some issues related to the renumbering issue, so can we add a delay through root system operate as frequently, such as every month add more and delete -- delete or add.

And how -- and how many root services -- servers is enough. How many too many.

So that -- that questions not addressed so far.

And also there is an issue about how -- can IANA name space can be served by more than one set of root name servers.

So you can find there's some technical issues and particularly for the current system and some are a little controversial about the policy related issue.

So that test bed we want to give some proof, and where we have three-year time line for the expert -- for the test -- for some experiment and want to find out some technical findings here.

And why I present here is that I think one of ICANN's missions is to coordinate the operation and evolution of the DNS root server name system, so you do not -- ICANN do not exclude the possibility of change of current system, so I just ask for ICANN's attention and people's attention here and also ask for and invite some interesting parties and individuals to join us. Thank you.

>>DAVID CONRAD: Right. Okay. And with that, if there's no other business, then I will call this session of the TEG closed and everyone can escape to wherever it is -- wherever their next meeting is going to be held.

Thank you very much and see you in Dublin.