ABU DHABI – Joint Meeting ICANN Board & Technical Experts Group (TEG)
Wednesday, November 1, 2017 – 17:00 to 18:30 GST
ICANN60 | Abu Dhabi, United Arab Emirates

DAVID CONRAD: Okay. I guess we can probably get started. Welcome to this meeting of the Technical Experts Group that is a bunch of technical experts meeting with the ICANN board. And in particular this year, we have something a bit new and that is with the creation of the Board Technical Committee. The Board Technical Committee is now -- I believe this is a mandatory meeting of the Board Technical Committee, so they can't escape even if they wanted to.

So in the interest of time, I will keep my remarks short, although I did want to make special note of the fact that the reason that I am wearing a tie today is not because I am going on interviews. Rather, it was in honor of the last meeting by Steve for the TEG. At least as a board member, he is, of course, welcome to participate in the TEG moving forward, if he so choose.

So I would actually just like to say personally that I am deeply appreciative of all the work that Steve has done in actually creating this work and has driven it until the Office of the CTO was created and then magically it landed in my lap. But I do want to thank you, Steve, for the efforts that you've done in improving the technical stature of the organization and helping me just as CTO and improving the technology at ICANN and all the various other things that you've done.

So with that -- unless you wanted to say something.

STEVE CROCKER: Well, thank you for your kind words. You've actually done a huge amount of work. I think you've made a big difference over a long period of time. And the good news is you're still here. And it's -- and this group, this Technical Experts Group, I'm quite pleased about it. It came about in a sort of unexpected fashion; but I think it's an added source of expertise, insight, creativity, and a place for certs kinds of discussions that might not have had a place otherwise.

So I'm quite pleased that this group has been created and has a certain vitality. And I hope that I've shared that with -- both shared it with the people who are here and done the thing that you do when you're in bureaucracies, which is you build it up so nobody can take it apart afterwards. So take it away. This is good.

And David, I must say, has really taken on both the administrative aspects and the substantive, creative agenda building and gathering of people and so forth. So it's really -- it's really quite nice. And he has a nice tie.

DAVID CONRAD: Okay. With that, I guess we can start off with the content of the meeting. And the first presentation is by Fernando Lopez on persistent identifiers over the DNS, which will be a demo. And this is a prototype, proof of concept that uses persistent identifiers that are similar to what's described in something called DOA, or digital object architecture, that are implemented as an application on top of the DNS.

I believe this presentation will be done in Spanish. So the translators will provide you -- if you don't speak Spanish, will provide you with some help in that space.

Alain looks like he wants to say something. So, Alain.

ALAIN DURAND: Thank you, David. I'm going to introduce the first few slides that will lead into the demo. I will speak in English.

DAVID CONRAD: Or your version of English, as I understand.

ALAIN DURAND: I can speech French if you to. I'm not sure I will understand myself, so it's not really good.

[Laughter]

Are the slides up?

DAVID CONRAD: We have a briefing paper up right now, but if we could switch over to the slides.

ALAIN DURAND: Okay. I will start with some disclaimers. So this is work that the ICANN Office of the CTO has started. And it's aimed at demonstrating if DOA-like persistent identifier could be achieved simply with the DNS. So next slide, please.

So this talk is going to present the state of prototype that we have done in collaboration with the University of La Plata, and Fernando will talk about this later.

This is not an endorsement of DOA technologies by the ICANN organization. I really want to make this very clear.

So next slide, please.

So the context of persistency is claims have been made that URL can break for many reasons. There are organizational changes, company name changes, mergers and acquisitions. And after 12, 18 months, 24 months, a large amount of URLs end up failing.

Next slide, please.

There have been a number of industry solutions to this problem like URL redirects, tiny URL.

Next slide.

The solution from the DOA is to look at persistent identifier through the Handle System. And what they're doing is we have prefixes where we use numbers. So by not using names, they're not using something that has a mnemonic semantic attached to it. So the claim is that if organization change, the number can remain the same. If you have a mnemonic to your organization, you may want to carry it over or not carry it over to the change of an organization. If it is a number, it doesn't matter as much.

On the suffix side of the handle, instead of having a deep structure that somehow reflects the internal structure of the company, we recommend to use a flat space, no hierarchy.

Next slide, please.

The Handle System uses specific protocols that are not standardized in open bodies such as the IETF. And looking into them, it doesn't look like those protocols add anything to the persistent history. It's just a different way to resolve identifiers. So the persistency is really the result of a naming convention that is described above.

Next slide, please.

So can we do this with DNS? Well, our short answer after looking at resource capital seems to be yes. We need three things. We need a place in the DNS to anchor this. So we're going to call this persistency anchor or PANCHOR. There doesn't need to be only one. There can be multiple events to allow for competition. We need a naming convention that is somewhat similar to the one I described above.

So in DNS labels do not use mnemonic. Do not use anything that has a human semantic. You can use a number. You can use a hash. You can use random letters but nothing mnemonic. And do not map the organization's structure. Use as flat spaces as possible.

The third thing we need is a new record type. So we have introduced a new record type that we have called DOA, but it was like the first attempt. Probably we are going to rename this. And we are thinking about calling this DTA for data.

And in this record type, we are going to put the structure object. A structure object is something that contains some information that doesn't have to be a mapping from name to I.P. address. It can be all kinds of information we're going to see about it in the next slides.

Next, please.

So this is this record type. It contains data such as an enterprise number. This is a number that IANA allocates to enterprises. So if you want to have your own private data types, you can put it in there. The second field is media data type. It could be some pretty fine values or user defined. It will be a location. Either the data will be contained inside a record, or there will be a pointer outside to where you can actually find the data. When you find a media type, it explains what it is. For example, it is text encoded into a certain dataset; or it can be binary, whatever you want. And if there is data, then we can contain the data in there, knowing that it cannot be too big so that it fits into a DNS record.

Next slide.

So just go to next slide. This isn't the most important one.

Okay. So we have been thinking about a prototype for this and what is the usage of this. We thought about the domain of IoT. IoT, we have heard a lot as looking at identifiers that are persistent and that are tied to a type of device. So we have been thinking about a company, call it BigCo, that is creating IoT devices. So under this persistency anchor, we give a label to that company, label 12. And that company makes devices; and for this particular type of device we put a number, like 78902. We just made the first numbers.

And to give you an example of what we can put in both records, to describe the company, we can have a Web page that points to some information about the company, contact email address, or we can install the public key associated with the company.

Describing the object of a device model actually, not objects but device model, we can have the same thing but we can also add more interesting things like the firmware or pointer to a firmware or firmware signature or firmware version. So that's a device -- excuse me, a device that will look this up, and figure out am I running the correct version of a software. If not, then we will be appointed to go and download it.

Next slide.

So we're going to go over the demo now, and I will hand this over to Fernando.

FERNANDO LOPEZ: Hello, I will speak in Spanish.

I'm Fernando from La Plata University. I'm a professor, a researcher there. There is a team in my university working through Cabase. We started working to create a demo, an application for DOA records.

Next one. Next, please.

So Cabase registered a persistent.lat domain and (saying name), which is an agency working within La Plata University, set up a set of servers for names within that domain. Servers we're using provide DOA records, and they are a beta version of BIND but of no special changes. It's just a beta version. They have implemented DNSSEC.

Next one, please.

As regard devices, that is where we prepare the demo. We worked with devices called NodeMCU, using a low-cost microcontroller integrating WiFi. It's call ESP8266. The price of these devices including antenna and flash is $1.50 in large volume. And they are usually programmed on C ++ or different languages.

To implement the demo, we had to change LWIP library, which is a library providing network support to these devices. And within that library, we modified the DNS part so that it could send DOA requests with it to 59 records so they could process responses.

Next one, please.

So the demo, I'm going to summarize it. Originally, you have to set up a record. The record has already been configured so we'll skip second step and go to third step.

The device will start running. It will make a request to a DOA record, and from the server is going to get the response. And the response will include the latest firmware version available, a link to that firmware version, and a firmware signing. The firmware will be available then and will be updated automatically.

Can we go to next screen, please?

So this is a snapshot of the device. We're going to skip this part.

Can we go to the shared screen, please?

On the left of the screen, you are going to see the startup process. It's connected to my computers so that we can see the startup, and we can stop it at any moment. On the right, you will see the traffic of DNS queries and the responses.

In red, you will see DNS requests and in blue DOA responses with DOA records. So that's the first step, a query with the responses.

You can see the field with the description, the firmware field with a URL, version field, contact email, and the signature field.

The capture field is slower than usual so you can take a look at the fields. But on the left, you can see that the firmware also received this information. And if you go to next step, you can see that the firmware will be downloaded, it will be updated and it will start with the new firmware.

Depending on the connectivity, it may take a few more seconds. It's downloading the update. In the meantime, I can tell you that this change is relatively small. See, implementation means 300 modified lines. It's one week of research to understand DOA, to understand the library, and then only three implementation days.

On the left, you can see a device has already been updated, has rebooted. And you can see that we have a new version, 1.0.

ALAIN DURAND: I wanted to add something. When we tried to do this live demo, we had some hiccups on the network. When we switched to IPv6 in order to run this demo, it has been working.

I really wanted to thank the people from the La Plata University and the people from the Cabase organization that has helped us to put this demo in just about three weeks.

DAVID CONRAD: Thank you.

ALAIN DURAND: This is the device that we're talking about. It's a $1 device, dollar and a half, depending on how much you buy.

DAVID CONRAD: Dave, did you have a question?

UNKNOWN SPEAKER: (off microphone).

DAVID CONRAD: Yes, of course. We have about five minutes of questions. Actually, start -- Steve has a question.

STEVE CROCKER: I didn't mean to preempt you. The automatic update catches my attention. Obviously, that's a very good thing unless the update causes some problem and then you lose total control of the device.

How do you characterize the safety property so that the updates don't leave you in a worse-off position?

FERNANDO LOPEZ: This specific device has a (indiscernible) for updating. It only takes the update if the hash checking -- the (indiscernible) framework is valid.

STEVE CROCKER: That takes care of errors.

FERNANDO LOPEZ: Of network errors.

STEVE CROCKER: But what happens if the update itself properly transmitted still has a bug in it?

FERNANDO LOPEZ: Well, then you need to implement some other solution, like a physical reset or something like that.

ALAIN DURAND: So I want to say, this is not a product. This is just a proof of concept. So there are a number of issues that you go on to describe that would have to be fixed if it were a product.

STEVE CROCKER: I like proofs.

DAVID CONRAD: Dave and Jonne -- and Jonne and Rick. Okay.

DAVE PISCITELLO: First, this is very cool. Thank you very much for implementing this. I love the fact that you've pruned down something that is very hard to explain into something that can be implemented in a very, very short period of time and very -- with very impressive clarity.

Have you thought of allowing -- instead of focusing on the data level, maybe an object level so that you could have a dynamic registration of a device as part of the interaction because if you do that, we're basically catching up to botnets. So if you know what a dropper file is in a malware infection, what you could essentially have is the equivalent of dropper -- dropper firmware in, you know, anything that you would put on the IOT network. And the only thing that that device would be able to do would be to use the DNS to go and enroll itself and then receive instructions just as it would from a command and control. And I see enormous opportunities for taking DOA in that direction. So I might suggest that instead of you calling this DOA or DTA, just call it OBJ, object. And so, you know -- so we're just focused on that. But thank you. This is really, really impressive.

ALAIN DURAND: Thank you. If I may just quickly respond. Yes, we have been thinking exactly about what you're suggesting because as we showed the delegation to the company, and the company to the object model, we can have one next layer of delegation to the actual serial number and we can delegate this to be managed to the object itself and use DNSSEC to go and validate. Just want to remind you that in this demo all the zones have been signed with DNS section.

UNKNOWN SPEAKER: Yes. So what you are describing is a modification of the DNS system itself to allow for DOI, but currently the DOI system is using the DNS. They are using doi.org/ the flat space. So the issue of persistence is not an issue for DNS itself. It is a policy issue with organization and everything else. But the URL that work from 1990 and the doi.org is as stable as anything, any other registry that the DOI system is going to produce. So I'm wondering why you chose to modify the DNS system, given that they are using the DNS root. I mean, the ICANN DNS implementation.

ALAIN DURAND: Thank you for the question. First, we are not modifying the DNS system. What we have done is creating a new RR type. We haven't changed the servers. We haven't changed the resolvers. We haven't changed the set resolvers. We haven't changed any of the billions of DNS things that exist. The only thing we have done is create a new RR type.

In order to do that, we describe the type. We ask (indiscernible) implementation to do that. In a week time, we had already four implementation of that thing.

Now, to your second comment about the DOI using the DNS, actually what they do is use a proxy. They send over data over HTTP to a proxy that then translates this into the handle system. And we were thinking, can we avoid all of this? Can we avoid the proxies? Can we avoid the privacy concerns that come with those proxies and do something with plain DNS, plain exact technology that we have been using for 40 years. And the answer seems to be yes, we can.

DAVID CONRAD: Okay. So just in case any questions, the queue is closed. Running a little tight on time. Let's see. Jonne.

JONNE SOININEN: Yeah. Just wanted to quickly really emphasize that what Alain said. So the demo is not actually on updating the device. The demo is about having a persistent -- or using the DNS in a way that, for instance, DOI works and not having to modify anything and the implementation is trivial. And you get -- in addition to that, you get all of the good things, even in a very small device, that you'd have with DNS, for instance, DNSSEC. And these are not a burden or a problem for the capacity of those devices. But you can actually run the traditional protocols in a very tight package and with very limited implementation required for this.

And actually it's not a big surprise in a way because DNS and much of the I.P. protocols were developed in a time where most probably that what you have in your hand would have needed much more space and would have been called a desktop. So it's -- this is a excellent example of that actually we should look at the kind of like what we have already today and in the new ways of how they can use -- be used.

ALAIN DURAND: Thank you.

DAVID CONRAD: Okay. Rick.

RICK LAMB: Yes. Wonderful. Very happy about this. I'm also an ESP8266 user, same chips. But I have a specific question. Did you modify the LWIP stack to support DNSSEC? In other words, were these lookups validated?

FERNANDO LOPEZ: No. And right now, it is not checking anything with DNSSEC.

RICK LAMB: That would really be cool. That same thing that you are using is used in all of the IOT devices you go out there and buy for DNS --

FERNANDO LOPEZ: A very important point for the solution. It has to be done.

ALAIN DURAND: So a quick comment. This whole thing started right after the LACNIC meeting about three weeks ago, and I had a discussion with my good friends from Cabase about this. They said, let's try to do this. So instead of going home from the Montevideo meeting, I took the boat and I went to Buenos Aires. And the next day they drove me to Universidad de La Plata and they said, Can we do this? So we had three weeks. So we had to cut this to the absolute bare minimum to make sure it worked. Phase 2 is we want to do what you described. We want to do what David is describing. There's nothing that will stop us from doing that. It's relatively simple.

DAVID CONRAD: Asha.

ASHA HEMRAJANI: Thank you. Very impressive. And in light of what we -- the discussion you just had, in addition to updating devices, another challenge with IOT devices is authenticating them so that you don't have -- so that you have proper credentialing done. So how simple would it be to extend this to authentication of the device?

ALAIN DURAND: I think actually it can be applied to quite a lot of things. We can do authentication of a device, but we can think about other type of application that needs a persistent identifier. Folks have been talking with me about, for example, medical records. We could do something like that. This is something that use the DNS technology, not as mapping names to IP addresses we classically think about but as a layer of indirection where we have an identifier that can be as persistent as you want it to be and that's going to point through this object to what you want.

And you decide what you want it to point to. So there's a fairly large domain of application.

DAVID CONRAD: Okay. Jay.

JAY DALEY: I am equally confused and horrified by this. DOA is to my mind, and I think we've discussed this in this group a number of times, deeply flawed technology, with a deeply flawed governance model, and a deeply flawed intellectual property model. Now, this is going some way towards attempting to fix the technology, but the governance model and the other bits around it are not likely to be fixed by this. And so I really have to ask, why is ICANN doing this?

DAVID CONRAD: I'll take that. So one of the activities that the office of the CTO is to look at new technology, new identifier technologies. DOA is a technology that has generated some interest in a bunch of different venues. Part of the project was to understand what exactly DOA was and how it worked and its governance model. One of the things that Alain identified as he was doing the research on DOA, just trying to understand it, was that it didn't appear to actually change much. The technology itself is a naming technology, but it is wrapped into a different governance model that, from our perspective, didn't seem to be necessary. Part of this work is to show that in actuality it is not necessary. The governance model is completely separate from the technology, of course, and the technology can be implemented on top of the DNS so you don't actually need that governance model. That's part of the demonstration of this technology and the point that we were trying to understand what the technology did, how it did it, just to make sure that we can relay that information to the community.

And with that, I think it's time to move on to the next presentation. Could you pop up the -- Okay. Is this --

LEONARD TAN: Hi, everyone. I'm Leonard Tan, volunteer for Ethereum Foundation. So today I'm presenting on Ethereum name service. And for those of you who don't know about blockchains, here's the two-minute overview. So basically blockchains are distributed letters -- ledgers. Thank you. And like all ledgers, it's basically just numbers plus and minus. And this has actually been implemented in the past by other protocols, algorithms like Paxos and PBFT. But blockchains scale relatively better. I mean, Bitcoin is all over the world. So they scale very well. And it compensates for the costs of verifying via mining and giving minus block rewards and they also disincentivizes attacks.

So a quick overview of how mining works is the transactions -- eTransactions input and output. And in order for the transaction to be valid, you have to reference a previous output. So how transactions look like would be something like this. Every single transaction refers to a previous transaction, all the way to a mine block where the first cryptocurrency Bitcoin is created, the first cryptocurrency mining.

And then the next one just refers to that transaction. And every single transaction's output is locked by a public key, and it can only be unlocked if you have the private key.

So this is where the blockchain comes in. The problem with this system is that you do not know if users have double spent the money, you do not know the transactions are properly formatted, so what happens is you have nodes, minors, taking a series of transactions, about 2,000 for Bitcoin, for example, and they verify all of them, follow these rules, and then they race to try to get a hash of this block below certain threshold. And when they finally do get it, they get rewarded with some Bitcoin.

So Ethereum is a little like Bitcoins, a blockchain. But on top of just storing data and transactions, we can also do computations on Ethereum blockchain. ENS is a system built on top of Ethereum.

So first off, I'll talk about why we needed to build ENS and then how it works and then I'll give you some updates on where we are right now.

So what is ENS? Primarily it's a way to map human-readable names to resources. One of the problems of blockchain is that identifiers are long hexadecimal strings. Their difficult to read and hard to remember and exposes users to phishing attacks. In general, they're just hackers' use. With ENS, we can now refer to records and addresses and contracts by using names. But that's not all. We can also use ENS to refer to other kinds of records, like Swarm and IPFS records or even public keys for identity attestation.

At a higher level, you can think of ENS as a distributed lookup service. It's resistant to DDOS attacks and because transactions in a blockchain are transparent and it was also designed to be upgradeable.

The internal architecture of ENS is split into two components, the ENS registry and its resolvers. The primary purpose of the ENS registry is to maintain a mapping between names and owners and its resolvers. So if you're an owner of a name, you can do one of three things. One, you can change the owner, you can reassign a name to someone else; two, you can change the resolver; or, three, you can create subdomains so you can see this hierarchical structure there.

As for resolvers, their responsibility is to answer questions about a name. For example, what address is associated with this name. What IPFS record is associated with this name. Name resolution in ENS is very straightforward. User queries the registry asking what is the resolver for, for example, foo.eth. The registry relies with Ox1234, which points to the resolver. The user then asks the resolver, what is the address of foo.eth, and then the resolver replies with Ox2345, which is the address the user is looking for.

So we did a soft launch for ENS back in May that lasted eight weeks, from 4th of May to July 12th. And over this eight weeks, we gradually released a bunch of popular names to users via auction process and released it gradually so as to prevent a spike in bidding activity that would flood the network and increase (indiscernible) costs. So over this eight weeks, at the end of July, 180,832 names were auctioned. From the graph, you can see a spike in activity and then a pretty level activity and then drops off towards the end after popular names have been taken.

Also, during this process 168,595 ether, thereabouts, was deposited. That's about 50 million U.S. dollars.

How it works is users submit a bid, and if they're successful at getting a name, the ether is deposited, it's locked up for a minimum of one year, after which we return to them if they (indiscernible) their names.

So client adoption for ENS has been good. MetaMask, My Ether Wallet, Itoscan (phonetic), Leaf, iWallet, Mist, (saying name) and Status. And we expect even more clients to use Ethereum to continue using ENS as the technology matures.

We also held our first ENS workshop back in August 2017 in King College, London, with 27 participants. And at this workshop, we covered many of the issues that I actually covered here at ICANN, such as dispute resolution, permanent registrar design, how to secure subdomains, and also how to integrate with DNS existing system today. For the last one, we have made progress recently, and we've managed to do DNS integration via DNSSEC. So on the slide here, you can see a chain of trust starting with the hash of the root nodes DNSkey, that lets you verify the DNSkey at the root node, so on and so forth, all the way until you get text record, for example, for _ens.ethlab.xyz. And for those of you who are familiar with DNSSEC, it's essentially the same thing except that for last part, instead of securing an eRecord, you secure a text record with the value of the Ethereum address.

So how the workflow goes for this process is the user submits a proof using the steps that you saw in the previous slide to an Oracle and then he makes a claim to registrar that he owns the subdomain. The registrar then queries the Oracle, asking if the user did submit a valid proof earlier. Oracle then replies with yes or no, and depending on that, the registrar will then register the subnode to the user under the ENS system.

So we've been working on it, and we are working prototypes on the test net. And theoretically this can be done for any TLD that supports (indiscernible) hashing and (indiscernible). And that's three quarters of all TLDs today. So thank you, and stay tuned for more announcements for ENS at ethereum.com. Questions?

DAVID CONRAD: Yeah, thank you, Leonard. I guess we have some time for questions. Anyone have any questions for Leonard and Ethereum?

I'd guess one question I'd have myself is so, obviously, you use .ETH. And I'm curious what your plans are sort of moving forward with regards to the top-level domain or the identification that you're using for that.

LEONARD TAN: Right. So we understand that .ETH is a three-letter code for Ethiopia so it's probably out of the question. But we are still in discussions. Right now we are all looking towards integration with existing systems first and testing out whether ENS is functional. And then afterward, we'll see how it goes.

DAVID CONRAD: Yeah. Just to clarify, .ETH -- so three-letter codes are not reserved. So the fact that it's a three-letter code for Ethiopia, it doesn't actually mean that it's been reserved for Ethiopia.

LEONARD TAN: That's great.

DAVID CONRAD: So if the next round of gTLDs occurs, that might could be something that you could look into, or not. But, John.

JOHN LEVINE: Thank you. I guess I have a security concern because blockchains are as secure as the miners. It's pretty clear that Bitcoin is heavily influenced but not controlled by big mining pools in China that use it for whatever they use it for, maybe money laundering.

So, like, do you know who does the mining for Ethereum? Do you have any reason to believe that there are not pools of miners colluding?

LEONARD TAN: So, first off, for Bitcoin even though most of the miners are in China, historically all the miners have always moved with incentives. That's one of the reasons why blockchains work, is because miners respond to incentives.

JOHN LEVINE: If your pool controls more than 50% of the mining, they can lie.

LEONARD TAN: Right. But even then, the miners individually they want to do something they are incentivized to do.

JOHN LEVINE: Rather than arguing whether it's possible for -- like, I'd like to know, basically do you have any idea who the miners are? Or are you assuming that there will always be enough of them and there will be few enough pools that you don't have to worry about miner -- about a pool of miners taking over your blockchain?

LEONARD TAN: I mean, that is a goal, a goal to make it more distributed. So for Ethereum, for example, we don't have specialized mining rigs. So this is one way you can prevent people from coming together to create specialized mining equipment. So that's one way to make it more distributed.

But as to your question of whether you can make it such that people don't collude together, I don't think that is something you can stop others from doing.

So we can build it into the system, to make it so everyone can mine with equal opportunity, if that's the best way to put it. But as to stopping collusion, that would be quite difficult.

JOHN LEVINE: Okay.

DAVID CONRAD: Okay. Any other questions? Paul.

PAUL WOUTERS: Sure. Paul Wouters, IETF. So I have a question. Let's say IETF gets the domain IETF in this naming system and we pay our fees for a couple of years. Everybody uses the site. And then at some point, we forget to pay and the domain falls back into the pool and then somebody else registers it and we don't know where they are or who they are. Now I go to a court system. I get some legal opinion saying I own this trademark and now I want to get this domain back. Is there any way for me to get this domain back?

LEONARD TAN: So right now, the ENS industry, you can change it because it requires four out of seven people. Most of them are Ethereum developers. And it is a consensus for several of them to make any changes. So it is possible, but it is going to be a very difficult thing to do but it is possible.

DAVID CONRAD: Okay. One more question, and then we'll move on to the next presentation.

JORDI PAILLISSE: Hi. Jordi Paillisse from UPC BarcelonaTech. I would just like to remark that regarding mining the way that you were having before, there are a lot of approximations to mining that do not require special miners doing this process. It's now called proof of a stake that takes a different approach and uses the value inside the blockchain to generate the new blocks. So it is a different approximation that maybe should be also taken into account. Thank you.

DAVID CONRAD: Okay. Thank you.

And thank you, Leonard, for that talk.

Moving on, I guess Michael Palage and Pindar Wong. Who is going to be speaking?

PINDAR WONG: Thank you. Could I have the deck, please?

DAVID CONRAD: Can we switch the slides?

PINDAR WONG: Thank you very much for having us today. Michael and I are volunteers within The Internet Society's blockchain special interest group. The reason why we are here, about two months ago, we had been thinking about the evolution of this very young technology of so-called blockchain.

In fact, yesterday was not necessarily Halloween. It was the ninth anniversary of the publication of the Bitcoin white paper. So it's pretty young technology. You know, we're not quite sure if it works. But we're interested in this development of basically naming systems like the one you've just seen, right?

The ENS is one of the so-called blockchains. Yesterday there were 1,234 blockchains out there. And just now there are 1,244 blockchains out there. So it's kind of -- something is going on. And each one of them will face similar issues of you've got these 34-character addresses that are very essentially random, just random sort of characters that will need to eventually have an easier way of being managed, such as using names.

So I think the broad comparison that we would like to make today through the paper that we have just released, the six-page primer, is that there's something going on here that will require potentially names and governance. In this case, name-to-public key mapping instead of name-to-IP. address addressing.

So we would like to go through this by, first of all, thank you for having us. I want to thank Michael personally for raising this to my attention. I originally thought this was a long-term horizon risk. And in this -- since the time that we basically started the discussion, things have begun to come to a head.

Also liked to thank -- many of these slides are from -- the Decentralized I.D. slides are from Drummond Reed of Evernym and Manu Sporny of Digital Bazaar.

So today is the third engagement with the ICANN community as part of a process where we try to outline horizon opportunities and the horizon risks of the new group of the Board Technical Committee on Friday.

Yesterday we expanded on these blockchain naming systems and Decentralized I.D.s. And today we would like to go through and use this time to discuss why that is relevant at all -- if at all, to ICANN. Talking about the evolution and evolution for discussion.

So as you know, the Internet, we would like to keep it loose and in small pieces. So the trouble with permissionless innovation is that surprise things happen, things like Bitcoin. And so in this case, are the disrupters being disrupted, right? We're no longer finding a telephone system. We're grown up now and this blockchain technology is coming from the edge, right? Bitcoin is not -- it didn't come out of the standard's process that anyone knew about. It was first issued as a code and a white paper. And in some sense, it's very similar to the Internet in that it's bad-boy tech.

The point I would like to make is that I think it's of interest to the board and this committee because it might change some assumptions such as, you know, a global, resolvable single root, for example. But more importantly, it might lead to changes in market structure which obviously we have the whole DNS ecosystem.

But the main thing I'm concerned about right now and just using today's example of ten new blockchains since yesterday is the rate of innovation in this space and the rate of adaptation and adoption.

So in the last three -- two contact points, we've used four examples of the ENS, which you've just heard. That hopes to work with the existing DNS. We've chosen another example of what's called BNS by a company called Blockstack, which is completely outside of the DNS system as sort of to book end that discussion.

We already have name collisions within the blockchain naming systems. We have a different one which is one that I'm involved with regarding the Belt and Road.

What I would like to focus on also yesterday is the decentralized identifiers which are actually within the standards processes that we're familiar with, right? In this case, the W3C groups.

So there's innovation occurring both inside traditional fora, which we would probably be aware of in terms of horizon risks and horizon opportunities, but there could be systems such as Bitcoin which are off radar. And this opportunity here in the paper that we've circulated is try to bring everyone up to speed and try to provide some thought leadership in that regard.

So the challenge here, as you know, with the telephone era was that there was an assumption before the Internet came along. And if you understood this one assumption, you could make a lot of money through ISPs, right, which was that distance equals cost. Long-distance phone calls, if you remember those days. The distance you called was basically how much you paid. And the Internet, we completely destroyed the economics of that. Right now we can have video conferences on all the time and we don't care.

Pre-ICANN, similar one. Governance is bilateral treaties. Governance post-ICANN we have global governance is multistakeholder. Prebitcoin, we have another one, time is money. We were always told that. Bitcoin is a very specific example where data specifically in this case can be matched to money, right? It's now 6,500 U.S. dollars per Bitcoin.

So what I argue is preblockchain, the assumption we should be aware of and might be concerning us is that we assume that centralized should equal secure, right? Just build the firewall high enough, wide enough. And blockchain may be that, in fact, being decentralized may be more robust, may be secure, and in this case potentially persistent with these decentralized identifiers.

So the point here is that the development process, for example, Ethereum, the development meeting is happening today actually in Mexico. And so here are some pointers there. And they have a development process called the EIPs, Ethereum improvement proposals. And Bitcoin is very similar. It has the Bitcoin improvement proposals. And they have their technical conference actually also beginning tomorrow in Stanford.

And incidentally, David, this is actually modeled off APRICOT when it was created.

The verifiable claims to work, which we went through yesterday, theretofore rebooting the Web of trust and root materials there. There is the group which is the credentials community group inside the W3C process, and they are meeting at I think the TPAC meeting beginning on Monday next week.

Blockstack is again somewhat out of that process, and they have their white paper. This is purely there to provide some resources to demonstrate the variety of fora that these discussions are occurring. And I would argue that most of them other than the credentials group is probably not within the existing standards process. So we just wanted to make sure that you're aware that there are these groups out there outside of the traditional standards-making process. And they're innovating quite rapidly.

The two that may be of interest are the -- I would argue the public blockchains. You have heard about Ethereum. You have heard about Bitcoin. There's another one, IOTA, which deals with the Internet of Things. But you can look at the different governance models which might actually provide an opportunity for ICANN to consider its role given some of the issues that we've identified in the paper.

So with that, I'm going to pass on to Michael who will go through the paper of which you have a very abridged summary in front of you.

MICHAEL PALAGE: Thank you. As Pindar said, the ISOC BSIG, one of the initiatives we're trying to do here is to raise awareness of this emerging technology and its potential impact on ICANN. And ICANN should be applauded for reaching out in its first emerging identifier session in Copenhagen earlier this year and having Namecoin participate. We've now had Ethereum ENS.

The third other major technology which is covered in the paper that will be circulated or formally published later this month is Blockstack. Each of these three technologies have a different potential impact on ICANN. It could -- as the title alludes to, it could either be evolutionary or it could be revolutionary. And that's one of the reasons that we tried to raise this.

One of the other things we've done in the paper is we've tried to look at what some of the community members are doing, specifically with regard to patent filings. One of the things that caught my attention in my research is that VeriSign has filed for three patents -- three patent applications in the U.S. PTO earlier this year that potentially has impact. The specific title of those patents are -- deals with DNS trust anchors to objects outside of the DNS. And in the specific embodiment, they specifically reference utilizing public ledgers and blockchains.

In addition, another ICANN community member, Bill Manning, has also filed for a patent application. So this is significant. This is happening. And this was only our initial analysis in connection with U.S. PTO filings.

Over the course of the next couple of weeks as the article is peer-reviewed, we're going to be looking to see whether there are any other increases with regard to this technology on an international basis.

What our research also uncovered is that there are a number of other international fora that are actively involved this. As Pindar alluded to, the W3C in connection with their verifiable claims, the ISO TC3O7 with regard to potential standardization of the blockchain, and our colleagues from the ITU have been very active in SG17, SG20.

So, again, the purpose of this paper is not to direct ICANN to do anything. It is purely -- and, Cherine, you'll appreciate that. Tricia Drakes who we both know and I had the honor of serving with on the ICANN board, Tricia always talked about thought leadership. So the initiative here of ISOC BSIG is to take some thought leadership and ask ICANN to be aware of this technology. We understand that you have a lot on your plate. But this is something that we really wanted to have a call to action to prevent it from falling through the cracks. So, again, we hope to have this -- hopefully have it published by the end of the month.

PINDAR WONG: We had the monthly call yesterday evening. We had a discussion within the BSIG. You all have the six-page report. Again, it's a draft. We will welcome comment on that. It's again not -- it's just there to provide some framing. It may be completely wrong. We don't know. At least it's an attempt to provide thought leadership and some examples of, again, the range of activities that are out there for those that would work with the DNS from those that are completely off grid.

The interesting thing with, again, these decentralized identifiers is that they are persistent. They use URNs. So persistency can actually be there.

So the question then becomes: Will we have -- at least at the consortium I have started, we have this notion of chain marks and how chain marks relate to trademarks, specifically for corporates. So the question here is when you have one of these addresses for Bitcoin or what have you, how do you actually know that it maps to -- if ICANN had a Bitcoin wallet, for example, how could I be sure that address actually maps to ICANN? Or Coca-Cola, or any of that.

So in trying to be ahead of the curve, we are trying to provide thinking in terms of what could be playing to ICANN's strengths, which in many ways is the community you have already established, both in terms of the ADR processes, the familiarity of dealing with I.P.

And that's going to be important if we're already dealing with potentially thousands of different blockchains where there will be effectively trademark infringement or trademark confusion by people registering front-running and all of that. Then it's not going to be good for that industry. But it would also, again, potentially distract many of the existing ICANN community members to participate in other fora just at the time where, again, there is the opportunity to go forward with ICANN.

Questions?

MICHAEL PALAGE: If I can just take -- I think we still have two minutes left out of our 15-minute block, if I was counting properly.

One thing just to show you the way this technology is emerging, Leonard had pointed to the fact that, I believe, there was 180,000 names registered in ENS. If you were to compare that to the current number of new gTLDs who were approved in 2012, that would rank in the top 50 out of the over 1,000. So, that to me is an important data point that should not be ignored.

One other data point is if you look at the total market cap of the domain names or the dot -- domain names that were registered, I believe under current valuations is around 55 million.

If you look back historically, that is equivalent to around 1998 when ICANN -- the green paper, the white paper. At that point in time, there was a total of 1.3 million domain names registered in the world at about $35 a year. So I think it's important to look to the history as we go forward, to look for some parallels. So thank you. And we look forward to your questions.

DAVID CONRAD: Okay. Are there any questions for Pindar or Michael? Wendy.

WENDY SELTZER: Thanks to both of you. And thanks for your participation. You mentioned some W3C groups, and W3C uses community groups as a similar sort of laboratory environment for things that we are watching over and exploring for.

Is there something standards track coming out of this work? So the active W3C blockchain community group is looking at some of these technologies, and the verifiable claims working group is looking at a particular piece of the standard vocabulary for claims about attributes.

And so, yeah, we've also been looking at this at the stage of -- there's lots of excitement and interest. And at the data layer, we're looking at when there's an interest in standardizing some components on top of that, we would look to pick it up.

DAVID CONRAD: Thank you, Wendy.

PINDAR WONG: Just to make an observation. I mean, a lot of the discussions immediately focus on the actually blockchain and the proof of stake and the proof of work and the different consensus algorithms. Most of the broad-based generalization here, but the security models of these systems are incomplete. They are incomplete because there's a security economic layer which is actually not that well understood. And so we started the BSafe Network with 22 universities now trying to figure this stuff out. Why that's relevant is because there are these tokens which have -- potentially have economic value. The Ethereum is an example. And the economics, i.e., the economic value of the token, actually affect people's decision-making, whether or not it's going to be speculative, whether or not you're going to hold or what have you. So it's not clearly about the technology.

There is an economic incentive which leads to some very strange behavior, but when you start having, again, when this tech becomes, what I would say, more mainstream, where you want to have a mainstream company take these digital assets as value in a wallet, then they're going to want to enforce or be very sure that the general public knows that that wallet belongs to them. And so with the Belt and Road blockchain, which I'm the chief architect of, the two and a half years or where we basically singled in in that one area to how we ensure for legal entity identifiers for corporations, the mapping between that and their digital identity online with these systems for any of these blockchains because we don't know which blockchain will with the blockchain.

So the opportunity here for ICANN is to consider whether to what degree it will engage in other fora and/or what relationship, when it does touch the DNS, what modality of engagement will it consider? Will it will proactive, reactive, or some combination thereof. I think we're early enough in the stage that you can get ahead of the curve, if you do see it as an opportunity more than a threat.

DAVID CONRAD: Okay. Thank you. And with that, I guess we will move on to our final presentation for the evening, and that is a talk by our esteemed chairman for another, what, 24 hours or so on tamper-proof root zone management. Go ahead, Cherine.

CHERINE CHALABY: Just before, if you don't mind. So what I hear from the corner there, clearly here, is that the DNS technology, as it stands today, will not survive. That's what they're saying. It's either going to be enhanced or it's going to be replaced by this new technology. Do you -- do you agree with that or -- as our CTO?

DAVID CONRAD: So I have always been of the opinion that nothing remains the same. That the technology is going to evolve or it will die. The direct vectors of change are unclear to me. There's a lot of fascinating stuff in the blockchain world. Whether or not it directly applies, there are plenty of arguments that suggest that it can. But I'm maintaining a certain level of doubt, simply because I don't fully understand the technology well enough to feel confident in asserting an opinion. So, you know, that's one of the reasons that I've been encouraging my team to investigate blockchain technology, to try to understand it, to understand its implications and how that will impact the larger ICANN ecosystem and beyond that the identifier system upon which the Internet depends. Yeah. So Pindar, what do you think? Will blockchain replace DNS?

PINDAR WONG: The issue of scale is serious. These systems don't actually scale well, okay? So let me just -- so we started scaling Bitcoin, or a group of us started scaling Bitcoin modeled on APRICOT, to try and scale the Bitcoin protocol. And we succeeded after two and a half years a few months ago. We're now layering the protocol. Instead of three to four transactions per second on the chain, we can now do -- the Visa network provides, what, 40,000 transactions per second. On layer two prototype networks on the Bitcoin layer, two networks can process 100,000 plus. And so the question here is that when assumptions change, market structure changes. And what I'm more interested in is the technical communities rather than the specific technology today. The Bitcoin community -- I'm a Bitcoin maximalist and because that Bitcoin community is amazing. It -- I'm the stupidest person in the room, which means I'm in the right room.

Now, whatever will come out of it, I don't know. But something's going on, and the something now, in order for that technology to be successful, has to be more usable. Typically we talk about naming. And that itself might be a mistake when we move to screenless computing. So who knows, but right now, to be more mainstream, it has to be more usable. We're dealing with QR codes. DNS is the first example that's easiest to understand, but machine-to-machine transactions don't need the DNS. So what exactly are we talking about? If the assumption is a single unique root that's definitive, what system -- if you imagine a world where we don't -- where it's statistical certainty and that's good enough, that kind of breaks a lot of assumptions. And so I think the main thing right now is to not get too involved with the tech other than to surface what are the assumptions and how does it affect ICANN's assumptions. It's 20 years in, you've got the car keys now, but if the assumptions which are sort of in the ICANN DNA get changed, that's a big deal. And I would strongly encourage ICANN develop its own blockchain strategy, just like everyone else. Everyone else is trying to figure that out.

DAVID CONRAD: Okay, Jay.

JAY DALEY: Thanks. I think when bearing in mind the answer to Cherine's question, we should separate out DNS from the registry business. Blockchain has a potential to make significant changes to the registry business, whether or not it has a potential to make a significant change to DNS.

DAVID CONRAD: Agreed. Okay. Now, Dr. Crocker, if you would.

STEVE CROCKER: Thank you. In comparison, this is a very simple -- this tackles -- what I'm about to tell you about tackles a very simple problem with the existing technology without trying to put in new paradigms that are affected across the entire ecosystem. And so it's really retrograde, by comparison with all of this advanced stuff. Next slide.

Oh, that's me. Sorry.

DAVID CONRAD: I did the same thing.

STEVE CROCKER: There's no voice recognition in this thing.

[Laughter]

All right. So this work -- this work is borne out of some conversations that I've had to endure over a number of years where -- this is all true, where I've sat face-to-face with senior officials in different governments and they speak as if it is a serious and real threat that either the U.S. government or ICANN or some combination or some equivalent version would make an abrupt change, the simplest of which is to simply remove their entry from the root. So country code, you know, XQ or whatever, pick one, all of a sudden disappears and so references into the root for that return nothing or return nonexistent. And their imagination is that this might happen in a period when there is serious political tension and so it would be an offensive move and they worry about that.

So what I've often done in those situations is I explain why that isn't going to happen, all the checks and balances, what our processes are and so forth, and further, that if it were to happen, the impact would be relatively slow, incremental. We have 48-hour time to live on the entries in the root so the effect would be a 2% per hour degradation in the caches. The news that that's happened, however, would propagate about 10 or 15 times the speed of light around the world. Not quite that fast, but you get my point. And that Boy Scouts and system administrators and others would seek to repair the problem with work-arounds and the effect would be quite different from whatever stupid senior official in the U.S. government thought that this was the right thing to do. I've even gone so far as to suggest, not directly but indirectly, that if some government wanted to cause the U.S. to suffer great embarrassment they would snooker us into doing that so that we would be embarrassed because the effect would be disastrous on ICANN's credibility, on -- and the U.S. government's credibility.

Nonetheless, that -- and I've come to understand that the people who are pushing on this actually do understand everything I just said. That they're not -- they're not stupid. They're not uneducated, but because you can conceive of the problem, it becomes a coin of the realm. You can trade on that. You can argue that that's an issue. You can make -- you can make a big deal about it. And so I said well, is it possible to counter this, not with political stuff, not with organizational stuff, but with really strong technical protection so that it becomes absolutely impossible, absolutely impossible for the -- that kind of scenario. Now, let me set that in context.

As I said, the nightmare scenario is that a -- an entry gets yanked out of the root abruptly. We do make changes to the root all the time. We have a process for doing that. So you can't -- it's not as simple as saying, well, we'll make no changes ever. But you want the process of changing to involve the affected party such that they have to agree to it. And if they don't agree to it, then things don't happen.

Well, that's not quite sufficient either, because there will be circumstances under which they eager cannot agree or will not agree but nonetheless the change has to be made.

So that's the overview. And now I'll take you one level through the detail and not through lots of stuff. So the motivation is, as I said on there, and the question, is it possible to design and field a system that precludes this nightmare scenario. And the answer is yes. Next.

The basic concept is built around a sealed system that cannot be tampered with. We have such systems in place, even in the -- the current system of DNSSEC, and what happens is if you try to tamper with the system, break into it, you get the private key, it zeros itself out. It says, not going to happen. Now, that makes the system inoperable, but it doesn't cough up the -- the private key. Well, what happens when you make the system inoperable? Then you can't make any more changes at all. So it's a trade-off between making inappropriate changes versus not being able to make any changes or in the vernacular a trade-off between false positives and false negatives or to use other terminology, type one versus type two errors. And in this case, as in many, many cases in real life, there is a big difference in the impact of one kind of error versus another, and in this setting, making an error by making an inappropriate change is much worse than not being able to make a change at all. We have delays built into our system so you can't tell how long it's going to take. Well, now you can with SLAs and so forth, but we -- we've lived for a long time with quite a bit of flexibility in how long it takes to make a change. So that's -- that trade-off works fine.

So the basic concept here is a sealed system that can't be tampered with. Our current update system has a split control, we have a database maintained by PTI, by IANA, and then one maintained by VeriSign and communication between them. And all of that has to be in sync. It is possible to build a sealed system that encompasses all of that. It's a little easier if you put it all in one, but it's not impossible to do it across both.

So the next statement is, you can think of the root zone as being divided into thin portions with one portion for each top-level domain. So a little bit of information for each top-level domain. And the main information associated with every domain is what is the set of name servers associated with that. And then with DNSSEC you also have what are the keys, what's the DS record associated with that. There's a little bit more detail and slightly more complicated discussion to have about the address records that are relevant, so-called Glue records. A little more discussion about handling the records for the root servers and the start of authority record that's in there. These are easily discussed. They don't pose any super problems. They just take a little more detail. Again, not appropriate for this level of discussion. Next.

So the next concept which I've emphasized a few times already is that no change should take place to the TLDs portion of the root without that operator's concurrence. This does not address other potential complaints from ccTLD operators and their governments such as sometimes the TLD operator says, I want to make a change, and they're unhappy that it doesn't happen immediately or that it doesn't happen at all. And that's a different class of problem. And one can also imagine that one sets up a system that is as robust and as strong as we're talking about that there may be some other things that emerge that we'd have to talk about. Next.

So if you think about it, we're now talking about what amounts to having a hardware token ideally, can be done with software but let's imagine that you have a hardware token, a device that you hand out to a TLD operator and they have to use that device to authenticate and authorize a prospective change. And if they don't agree -- do that, then the change can't take place. Well, how do they get that in the first place and how do you make the association between the appropriate operator and that device? That's a different process. That does require sort of breaking this -- this idea that no change takes place without their concurrence because they're not there in the first place. So the initial assignment has to be done through a more laborious process, and we have something -- we have key ceremonies and other processes that are similar where you get a whole bunch of people to agree that this is the right thing and those people are enough independent so that you can avoid collusion and you can avoid other forms of pressure on them and you make the whole thing slow, deliberate, visible, documented and so forth. So that would be one of the class of things that you would want to use this slow process for. Another is if you had to take it away from somebody because of a hostile reassignment. Another situation is, they lose it or it burns up in a fire or something and you have to make another, so it's a class of things. For all of those, the general solution is when you go through a laborious, slow highly-visible and documented process. So that's the multi-party political control.

The next thing that might come to mind is well, not everybody's going to want to be able to do this. Not everybody is going to be able to do this all at once. We only have, what, 1,500 root -- TLD operators at the moment. Ten of them will be ready today and the other 1,490 will take a little more time and the last thousand of them will take two years or ten years or whatever. So it will be whatever some -- some transition process. So in any kind of design, you have to be prepared to operate with the current system and with the new system, and that would be fine. That's really no issue there. You could have notional. You could think of retooling the system completely and then having one of these devices for every TLD operator sitting on the shelf, and the ones that are ready to take it on, you send it to them. And the ones that aren't, then you do it for them and they don't see the difference. Next.

So this a picture of what the root zone process looks like. Changes come in from the TLD operator box on the left, they go to the IANA function. This is PTI, which validates the request as it looks like it says it's okay. Sends it on to VeriSign, the box on the bottom which does -- at least in this diagram -- two things. They edit their database and then they generate a root zone. I purposely left out key generation and signing which adds a complication to the diagram but for simplicity I'll just leave it out. And then it moves over to the distribution process twice a day. And a new version of the zone is made available to the 13 lettered root server operators. So that's the update process today. And the change that we're talking about. Next slide. Would be to encase that middle column in a sealed system, or in the case of two different groups that are operating, it would be physically two hardware systems, both of which are tamper-proof and have a -- a robust protocol that connects them together. Next slide.

So just repeat, there would be two kinds of transactions: So-called ordinary transactions for the regular changes in NS records and DNS key or DS records and associated glue records. And those would go through in the -- in the fast path where TLD operator would say, I want this change and here's my authorization for it. And that change can be approved or not approved, but it can't be tampered with. And that change gets made or doesn't get made.

For the bigger changes that changes the control and so forth, require this more elaborate process.

Next slide.

And so this is just another attempt at diagramming what an ordinary change would look like.

And next slide.

And my ability to craft what ought to be up there was a little limited. But think of that oval at the top as like a big conference table with a bunch of people around and all kinds of processes associated it.

And so it's similar -- very similar to how we do our key ceremonies with our trusted computer -- trusted community representatives and so forth.

Next slide.

And here I've basically just said what I've already said, that the oversight body would be a trusted set of people.

Next slide.

So if we wanted to explore this, the next steps would be to flesh out a conceptual design and document it and circulate it and socialize it and so forth. As it turns out, you can divide the work up into three parallel paths of laying out what the process would be. You could actually build prototypes of the interaction between the TLD operator and the system and you could build a prototype of how the system itself would work. And those three steps could all be done in any order in order to get more information, if you wanted.

Next slide.

There we go. Next slide.

I think that may be it. Last slide. That's the concept. I have been sitting on this for quite a bit of time, for a few years. I can't remember when I first started working on it. But I stopped talking about it completely and just set it aside. When the transition process started up, I thought it would be a complete distraction and cause too much confusion. The transition has passed so here we are again. Thank you.

DAVID CONRAD: And before we get to questions, I did want to point out that my team is actually planning on issuing an RFP for two different types of changes to the root zone management mechanisms. One is an evolutionary change to take the existing system and tweak it to make it, you know, perhaps better. Hopefully better.

And the other is a revolutionary approach which sort of restructures the way we do things. And one of the thoughts was to incorporate this -- the tamper-proof idea into the revolutionary approach.

And with that, any questions? And we only have a couple of minutes for questions unless people want to skip the cocktail which will make Kathy very unhappy.

[Laughter]

Rod, I guess.

ROD RASMUSSEN: So this is, I think, an interesting approach. I'm wondering, this is a root zone management -- I'm over here, Steve. I would also say that dot brand TLDs would be very interested in something like this potentially. Obviously a different governance model. But why stop at the root?

STEVE CROCKER: Yes, precisely. The same technology is obviously applicable at every level. And full stop, yes.

DAVE PISCITELLO: So -- Dave Piscitello. The only threat model that it seems -- that's instigated this is people are worried that something will break. I would really like to see a threat model and a cost-benefit analysis to get a good sense of what -- what is the end result of changing what we have? And what threats do we actually mitigate that we can't do today with the existing model? Because I don't see -- I don't see the threats as obvious as perhaps somebody who says, you know, the U.S. is going to take away our delegation.

STEVE CROCKER: I tried to address it at the outset. If you are sitting close to the center of the world like you and I are, we think this is absolutely ridiculous. It ain't going to happen. And nobody should be worried about these kinds of changes.

Move out into -- you know, far away and all of a sudden, it looks like, "Oh, my God, we're at risk and our whole country's economy is going to go down the toilet overnight." That's the threat model.

DAVID CONRAD: Okay. One last question?

STEVE CROCKER: That's the threat.

Paul.

LARS-JOHAN LIMAN: Lars Liman, Netnod. Does this prevent against denial of service attacks in the form of "I want to make a change, no, it won't happen?"

STEVE CROCKER: I'm sorry. Say it again.

LARS-JOHAN LIMAN: Does this prevent against denial of service --

STEVE CROCKER: No, no, it does not. And I tried to allude to that earlier. I had an interesting conversation several years ago with a major TLD operator. I said, how long does it take to make a change -- how long should it to make a change to your operation to a new name server? He said maximum 48 hours. As opposed to what happens with.COM when you can do it in a few seconds or a minute.

I said, so -- and how long do you plan for it? He says, six to eight weeks. Right? This is -- this is quite real. Because back in those days, what he was expressing was -- he goes -- makes the request and he isn't sure it's going to happen or it will go hold up for some time or whatever.

So we are in a much better state these days, much, much better. Nonetheless, my sense is that TLD operators that are going to make changes in their name server configuration are not in a position where they need to be done instantaneously and that if they put in a request and it doesn't happen, then they get to escalate and there's normal processes for escalating and for dealing with that.

So your characterization of it doesn't happen is actually the beginning of a process where it doesn't happen this minute and it then goes through a more extensive process. So it turns into delay as opposed to an absolute negative unless there's some reason to turn it down, in which case they would be the right answer.

LARS-JOHAN LIMAN: My reason for turning it down in this example would be political, that there's pressure that prevents the escalation.

STEVE CROCKER: But if it's political, then you get to sort it out on the political level.

DAVID CONRAD: Okay. With that, we get into any other business. And the only any other business that we have is everyone should -- who's interested in the cocktail should make their way to the grand stand entrance. And there will be buses to take us to a name I'm not going to try to pronounce because I will get it wrong. But it's a beautiful observation deck at a lovely building that we've probably seen pictures of in fine magazines and travel guides. So thank you, everyone. And I guess we'll see you again in Puerto Rico.

[END OF TRANSCRIPTION]