ICANN
Transcription ICANN63 Barcelona
GNSO - CPH TechOps Meeting Session 4
Sunday, 21 October 2018 at 15:15 CEST

Note: Although the transcription is largely accurate, in some cases it is incomplete or inaccurate due to inaudible passages or transcription errors. It is posted as an aid to understanding the proceedings at the meeting, but should not be treated as an authoritative record.
The recordings and transcriptions of the calls are posted on the GNSO Master Calendar page
http://gnso.icann.org/en/group-activities/calendar

Marc Anderson: Hello, everyone. This is Marc Anderson. And welcome back to Tech Ops. For the next session, we have a presentation on ID4me. And I'll hand it over to Vittorio who will be providing the presentation. So, yes, take it away, please.

Vittorio Bertola: Thank you. Thank you for inviting us. I'm Vittorio Bertola from Upper Exchange, which is not a registry/registrar but it's the parent company for PowerDNS. And we are one of the initial promoters of the ID4me project together with DENIC and 1&1 and several others companies that have been joining throughout the life of the project.

So today I'm going to explain you what this is about. So what's the deal, what is this project, what we - and what are we trying to do and why this is relevant to the domain name industry in general. And then we will briefly see how it works and then of course we will encourage you to join the project.

So you - can you, next slide? Yes. Go to the next one then. Okay, this is fine.

Now, so this is a project about online digital identity. You might have noticed that in the recent few years there have been attempts to be something different than the traditional set of these domains and passwords that everyone uses, which are becoming increasingly insecure, impossible to manage and very hard for the user.

So I mean, smarter users maybe have password managers in their browsers or specific applications or even at best stuff like YubiKeys, I mean this kind of stuff. And so they can deal with different user names and passwords.

What most people usually do is just they have one or a few passwords and they recycle them everywhere, which is totally insecure. And in some cases, given the requirements of the different Web sites, which seems everywhere besides makes up different requirements, you will not be able to recycle them. And so you end up having to recycle, more or less, the same password with variance and (unintelligible) and not remember anything.

So this is why there are some single sign-on systems that are gaining ground, specifically the ones by the big OTTs. So if you noticed in the last three or four years, not more, most Web sites have implemented a sign-in with Google or sign-in with Facebook or a few others.

I mean, there are some that have a list of ten or 15 different social networks and other OTTs that you can log in with because this is much simpler for the user. I mean, the user only has to remember the center password, the one for the center single sign-on system. They just click and they are in, and they don't need to register for the Web site. So this is a much better experience.

But still, we don't like this. Well, first we don't like this because we think that the identity market is also - is important so we would like to - I mean, let more people can - more companies can be part of it. But also in general, I mean, just thinking of the - for the good of the Internet, these are all closed systems that are building walled gardens.

So for Web sites, you need to implement each and every login box separately so one login box for every single sign-on system you want to support. For users, you need an account in each of these systems. And in the end, all the data are centralized in like a single provider of each of these systems and you cannot move your data as a user from one system to another.

So that - so in the end, what we want to do is a single sign-on system that works exactly like the OTT ones like sign-in with Facebook. But it's open. It's federated. So it's open to any number of identity providers, and users can pick one and just even change it if they want and have more control of their data.

So that we - what we are planning to do is basically to build an open standard for this because we believe in open standard. And we think that there is an opportunity because there is a demand for more privacy around this service. So, I mean, people like single sign-on systems, but if you can provide a single sign-on system that is more privacy friendly and more flexible, people could be interested in it. And so we think so.

The next one.

So in the end, what we are also planning to do is not just to solve the problems of the user, which is important, and to keep the Internet open but we are also planning to increase the chances - I mean the opportunities for the future of the DNS.

So since we are all companies that work around the DNS space, we think that new users for the DNS have to be found. So if you want to drive the domain phase, you need more things for which the domain can be used.

So domains nowadays are registered for - mostly for Web sites and for e-mail addresses. But if they could also become your identifier, so your user name on all the Web sites and online services that you use, then this would be one more reason to register, especially personal domain names so this could be - well, both - I mean, a reason for people to register a domain name and a way for companies to increase their sales basically.

And in the end, we think, I mean, that - there are a lot of discussions around online identity. And if you are familiar with them, there is also lots of discussions around blockchain because people say we should do this with blockchains. Nowadays, blockchains are very popular.

But in the end, we think that the DNS is really well suited to be the - basically the public database for identities because it's already there. It's already federated. It scales very well. It doesn't have all the problems that new blockchain technologies have. And it's already also controlled in a way that is not centralized. So there is some - still some guarantee that no one can take over the entire DNS and blank everything and.

So I mean, indeed, so it is a globally unique namespace. So it's actually perfect to use the DNS to give people identifiers.

And so basically in the ID4me project, we have chosen to use DNS source names as identifiers. They do not necessarily have to - I mean, to correspond to an existing host name. They are just names in the source??6:30. But they are the identifier that people can use.

Though you could also do the same with e-mails, converting them into host names, special host names. And we will see this later.

So next slide, yes. Okay. Go ahead, yes. Yes, no, go to the next one. This is supposed to an example. Okay.

So in the end, the difference between ID4me and the existing identity systems, it's actually not much on the technical level since to make this easier, we're trying to keep as many of the existing standards as we can. So we are using of course DNS and DNS Sykes for the database. And we are using OpenID Connect, which is the common standard for a single sign-on for the proper login process.

But we are editing the passwords as necessary to make this open, federated and public and to ensure people - to ensure that people have a choice of their providers. So they can choose and even change their provider over time.

And also, we are introducing a separation of roles which is mimicking the way the domain name industry works. So in the existing identity providers, I mean, you don't just have one single provider but the single provider that controls the world gathering of, for example Google or Facebook, also does everything, does the authorization and does the management of the user information.

And we think it - that things would be better if we could separate the two roles. So you have basically a registry so an identity authority, which is only dealing with the verification of the credentials.

So it's the place where the user creates the password and maybe there second factor authentication mechanisms and whatever is used to verify that they are there when they have to log in. And so the authority will - is the trust anchor in the system. And it is all tasked with verifying the authentication when the login happens.

And on the other hand, the identity agent is the equivalent of the registrars so it's the part that has the direct relationship with the customer, that manages all the information. So that's where the user enters their - all the pieces of information in their identity -- their name, address and so on.

And it's also the part that controls the distribution, so the authority will basically authorize the agent to share this information with third party Web sites.

So this, in the end, is well suited also for adoption by domain name. Usually, when you - in the standard, there is nothing really preventing a single entity

from landing both roles and so for being both the authority and the agent for the service.

So next one.

Sorry, yes, so this is in meaning slide. It's the summary of why this is different and how this is different from other existing identity projects. I mean, if you go online, there are a number of other identity projects, not just I mean the one by the OTTs.

There are several others starting in Europe but we don't like them. We are different because we are the only ones really willing to build this federation so build something which is an open standard and it can allow any number of providers. While there are other parties that just allow thinking of becoming just one more world gatherer and then controlling everything alone.

And it also means there are efforts to build public identities. In Europe, we have the eIDAS standard and COUNTER. They are doing the same. These are usually governmentally run projects, which are very secure, very heavy. They require a lot of verification of the identity of the user.

So I think these are well suited for public services for paying your taxes, maybe for logging into your bank. But they are not what the users are going to be everyday to log in to a forum to chat with people or fill out an online application or just a standard Web site or whatever.

So we are actually - I mean, we think we are complementary with these kind of services. So the idea is that you could have one identity which is eIDAS compliant in Europe for the strongly authenticated services and also compliant with ID4me for all other use of the Internet.

But I mean, we really think that we have a unique proposition in this space. And this is what makes it interesting also for companies.

So next one. Yes, okay. I think we already went through this. So let's go ahead.

So in the end, yes, the point is basically about convenience, about privacy because nowadays you don't have control. Actually, many of the companies, at least in the private spaces, that are doing these identity systems and single sign-on, they are doing it because they want to profit from the data. So they want to have the data from anyplace you log into because these are part of the profile they build and then they monetize it through advertising or through other mechanisms.

And so we really think that - I mean, we will see actually a demand from especially - I mean, from certain parts of the world from privacy friendly users that would like to use a service like the one by the big OTTs but don't like the idea of having companies that literally get most of the revenues over user data monetization, I mean to have the control of their logins.

So the next one. Yes, okay. Let's go ahead. So let's go to the next one. This is the title. Okay.

So this is how it works. This is - as I was saying, it's based on OpenID Connect which is basically the standard protocol which is used nowadays in most single sign-on projects. It's actually going to be certified by the OpenID Foundation.

And I mean, we have extended the standard by a few things but we will even try to make them standards. So we are approaching the OpenID Foundation and trying to make the extensions that we made a part of the official standard. And also we have submitted a couple drafts to the idea. So I mean, the idea is to - anything related to this project will be an open specification and an open standard.

Then we use the DNS both for the namespace so to give names to the authorities through the identities and the database so as the place where a Web site can learn which identity provider is managing which identity, which is necessary to perform the authorization process.

And then we will - basically we are creating a federation. I mean, the model we have in mind for - is the e-mail model. So I mean, if you look at e-mail, of course you have any number of providers. But if you have an e-mail address, you can send e-mail to everyone else, even people that have e-mail from other providers.

This is not the same for example with instant messaging, which you need an account on WhatsApp Tutor with people that have a WhatsApp account and then one on Skype and then one on... So I mean, this is really an attempt to create a traditional - I mean a service meeting the criteria that the architecture of the traditional Internet services from the (unintelligible) over the Internet.

So next one.

So these are the things that we are adding to OpenID Connect. And I don't know how many people here are more or less familiar with OpenID Connect. But if you know OpenID Connect, you know that there is actually a discovery process, which is the mechanism that allows you to discover the identity provider just by starting from the identifier.

But the one that it is in the standard is really cumbersome. No one is actually using it because it requires you to have a Web site, basically a Web server at - in each and every domain name that you use for an identity. So for a host that maybe has a - I mean, if you had one million domain names and identity if a million of domain names, you would need one million Web services up and running just to deal with the discovery.

While we do the same with a DNS cycle, which is much easier. I mean, you need to have the DNS zone anyway. So we are - we have replaced the original discovery processing (unintelligible) with a DNS-based process using - well, we'll see.

We are adding - so the - basically the federation so the capacity for a Web site to just implement a single login box and support all identity providers immediately. So this is even easier for the Web sites that increasingly have a problem of having ten different login boxes to support in different single sign-on systems.

We are basically creating a portability mechanism so as long as if you - as you the user on the domain, which is also an incentive for you to buy a personal domain name, you can port your identity from different providers. You can move from the current one to another one if you want.

We are adding the separation of roles that I already talked about.

We are adding the features that are necessary for management of user consent. So in our system actually the user gets (unintelligible) with the sharing of each news or piece of information.

So the first time you log into a new Web site, you are asked - I mean the Web site says, I want to know the name, the e-mail and all the list of the pieces of information, and you the user get to choose which ones you want to share with them, which is not something that is happening with the current single sign-on systems.

And also what we are doing is that we are adding more information feed. So the idea is that you would be able to put in your online identity any kind of information, not just the basic e-mail and name but maybe your frequent flier number, your weight, your whatever, your preferences for your seat on the airplane so that in the end, through this mechanism, you would be able to

basically share information with a click and not have to register or type information into a Web site anymore.

So basically this would make all the registration forms go away and replace them with just the agreement to share the data that you have loaded into your identity.

So the next one. No, below. Okay, yes.

So as I said, there are four roles in ID4me. So the - basically there's the user, the relying party which is any online service that wants to authenticate. Of course the initial test is with Web sites but we are also working on offline authentication, stuff like using ID4me to log into your e-mail server with IMAP or, which requires some of token exchange but we think we can be (unintelligible) securely.

And then the other two, as I said, are the identity agent and the identity authority, which are basically the real end of the registrar and of the registry in the domain name registration process.

So next one.

So in the end, when you want to create a new identifier, well first you have to pick the host name in a domain name that you control, which could be your own personal domain name if you are a smart user. Or it could be your - I mean some company, for example your registrar giving you an identity under their name. So this could also become a way for companies to spread their name, promote their name by having their customers sign in with - into any Web site with their - with a string that contains their domain name.

So but we are not actually forcing any business model. We also curious in a way to see what the market will want. So we are just creating the standard and basic conditions for this to work. And then different providers could try

different business models and different pricings or bundle this with other services. And we'll see what succeeds.

So in the end, the user will have to go to one identity agent, which will register the domain name if necessary or sell it even if necessary. It will set up the DNS records that are used for discovery, and we will see an example later. And at this point in time, we'll go to an authority which is their partner.

I mean, the authorities are - could be TLD registries but there is no strict relation between a TLD and an authority so an authority could just I mean work as an authority for any domain in any TLD if they wanted. So the - and also the agent could choose the authority they want.

And so the authority would basically receive the request for registration, verify that the DNS setup is correct and go from recreation of the identifier.

And as the last set, the user gets redirected to the authority's Web site because one of the key features in terms of security is that only the authority gets to see the password and the other credentials of the user. So there is - I mean, even the agent doesn't see it and also the online services that use this authentication never get to see the password.

The password and the other credentials are always entered only at the identity authority's Web site or up. And so this is - basically removes the risk of you, I mean, entering the same password in a hundred different places and then just well one of them is not so they get cracked or leaked or to compromise your password everywhere.

So the next one.

The next one is about how the login mechanism works. So it's pretty easy because I mean on the Web site you only have to enter your identifier. You don't enter your credentials. And then the reliant party, the Web site, will use

the DNS to discover who is managing your identity. And once they do so, they will redirect the user's browser to the identity authority, which is where the actual login happens.

At that place, I mean, the authority can ask for the password but there could be a secure section already open. So it could even be that the user doesn't even need to enter the password. And also, at that stage the user will be asked for consent to share the information, so they will see which information has been requested by the reliant party and will decide whether they want to share it or not.

And after this phase, then the user will get back to the - basically to the reliant party. The reliant party will also, if necessary, gets a token that they can use at the identity agent to retrieve the information about the user.

So at this point in time, the reliant party gets the pieces of information like the name and e-mail from the identity agent and can use them locally to create an account or if it's the first time that the user logs in or even to update the information if the information has changed in the meantime. And then the login is completed and the user is in.

Next one.

Yes, so we already have - I mean, the project has been going on for I'd say almost a couple of years now from the - when we started from the very beginning in terms of experiments. So we have actually developed working implementation and this - some URS that you can try if you'd like to get a copy of the presentation.

We have some test identity agents and the authority provided by DENIC and 1&1 and some best reliant parties so you can actually create your identities and try to use them for login and see how this works.

And of course, this is still experimental so we have not launched the service yet. We are in the phase where we are actually explaining what we are doing to other people and companies and getting more people to participate. So we have a good number of companies that have joined and we're looking for more before starting, possibly next year, to launch this on the market.

So next one.

So this is actually how it looks like. This is our own Webmail application, OX App Suite. And as we see, there's our inbox which doesn't have the password. It just has a specialty user name. We are actually into ways to - I mean, to get the reliant parties to recognize your user name automatically, in some cases at least. But as of now, you have to enter your user name and push the sign-in button.

And so next one.

This is what you get next. So you basically get redirected to the identity authority, so DENIC in this case, where you find of course your user name that you already entered is already pre-entered. And you find your - a space where you can enter your password. At this point in time, you can actually have two-factor authentication or other security mechanisms.

And one of the nice things about this mechanism is that as soon as the identity authority implements a new security feature or another factor of authentication or whatever, it immediately secures all the Web sites you log in. So you log into (unintelligible), since the login now always happens at the authority, there's just one place where the login has to made - to be made secure.

Then, I mean, after you enter your password again, you could even not have to enter it if you already have an open session. But let's say you have entered your password, again there's a second check. If this is the first time you are

logging into a Web site, you need to give consent to the sharing of your information. If it's not the first time, possibly you already, I mean, filled this form and you decided that you want to keep it in memory, so you just skip this phase as well.

But as you see there, you get a list of all the pieces of information that the Web site has requested and you see which ones are mandatory and which are not. And the Web site can even provide reasons or explanations to tell you why they need a specific field.

And then you can pick - I mean, for each individual field, you can choose whether you want to share that information with the Web site or not. Of course, if you don't share something which is mandatory, the Web site can refuse service - this - I mean, like the European laws and the GDPR now works. But at least you have complete control of what is shared with the Web site.

So the next.

And then at this point in time, we have no - we don't have the image but you - basically you're in. And once you're in, the Web site already has your information. So I mean, the screen that you cannot see there is my Webmail with a new blank account just created because it was the first time I was logging in with my name in it because the Web site has retrieved my information from the agent. And so it already just populated my account.

So these are the DNS records that we are creating. As you see, they are nothing special. So there is one which is really the basic one, which is the discovery record. And it's a - I mean, it's a TXT record, a format which is now standard and being used by several protocols like DKIM and other protocols.

It's like a set of name value couples for - I mean, for each identifier. In this case, my identifier, my user name is VB.Bertola.eu so there's this special

name and this card and this OpenID dot identifier. You - and I mean, and you as the provider just need to create a TXT record. That the first is our version number. And the other two are to point us to the identity authority and the identity agent.

So in the end, you can have at least more details than this. I can choose the port and the path and so on. But in the end, it's a very simple pointer. And this is all you need to create.

I mean, there is another DNS record that gets used in the creation process. This is - I mean, currently we do this with ACME but there could be other ways of doing it using even standard domain name provisioning protocols because this is what is used because the authority and the agent to check that the domain is really under the control.

So I mean, the authority, before creating the identifier, works to check that actually you're not pretending that you can use that domain name but you can use the domain name. And so it asks the agent, which is usually managing the zone, at least for consumers, or could be the user if it's a smart user, to create a challenge record, in this case for ACME using a standard ACME DNS as a zero one challenge. So you need to create this for the creation to work.

Next one.

Yes. No, I'm not going to go through this but you have it in the presentation, if you get a copy of it so if you want really to know how this works in OpenID Connect terms. I don't know, again, if anyone is familiar with. That's the entire login flow and described in detail.

So go ahead. Let's go to the final part. Okay. So this is where we are with the project. Next slide.

So these are companies that have already joined or declared support and so on. So we've - I mean, we've got now nominate. So we - I mean, of course, the project started from three German companies so of course it started in Germany and possibly Germany will be the first country where we launch.

But we really want to make something which is global, which is also a way to allow - I mean, because there are also some national ID projects like the Czech one, there's MyID, which is also possibly becoming interoperable with us. There's a Swiss one. There are several national OpenID projects that are trying to start.

The problem is that maybe - even if you start with the support of the biggest telco, their TLD registry in the country and all that, then maybe you will get supported by all the big Web sites in your country. But as soon as, I mean, your Swiss account - a Swiss citizen wants to use your account to buy a flight on Lufthansa, it will not work.

So this is why we think that this is also interesting for all these smaller national identity projects or local identity projects because in the end, we will need a way to federate all these projects. Otherwise, they will never fly or they will stay very local, just used for very few services.

So this is why, I mean, we're really looking forward to making this a global idea and to making this an open standard, possibly even a DIPF, once it gains adoption and so on.

So the next one.

This is basically where we are. So as I said, we have a working experimental platform. We're still adding features to it and checking, making it work better.

Now we are discussing the launch. So the idea is to launch in Germany, as I was saying, because we think it's good. I mean, for a project like this to have

success, you need some critical mass when you start. So we think we can form that mass more easily in Germany and start there and start building a user base and with that experience, go in other countries, possibly France and the U.K. as the next ones.

But again, this also depends on people to join. So we're in talks with several big (unintelligible) and we aren't saying must become (unintelligible) in the world. So there could be more.

So the idea is that in the first half of next year I think we want to try to launch in Germany and start building an actual user base. And then we'll see how it goes. But we - hopefully we can launch in other countries soon as well.

So next slide. We are almost at the end so.

So yes, so basically this is the message. If this is interesting to you… I think should be interesting to anyone in the domain name industry because I think we need more reasons for people to use domain names, and it's also good for the Internet in general.

So if you want to do so, I mean, we now have founded an association in Belgium, a nonprofit association, so that there is a nonprofit, independent home for the standard that can manage that - I mean, the further development plans so that this is definitely not owned by a few companies. So anyone can become a member.

Even if - I mean, if you don't become a member or before becoming a member you just want to have a look, you - I mean, you have a WhatsApp ID for me to talk, you can join the working groups.

So if you go to the next one.

We actually have three working groups, one technical, one policy and one about the adoption and outreach and market launch. Yes, actually we're changing the chair of the adoption working group now. But yes, anyway, you can - we can - so you can join and participate in the working groups.

So in the end, the last message I want to give you is that we think this is interesting. Come to use if - please join. We have a workshop on Tuesday. I think it's already fully booked but if you want to get more information, just come talk to us and, yes, we will be happy to support you to give you more explanation and in the end to convince you to join and join the effort. Thank you.

Marc Anderson: Great. Thank you, Vittorio. We have about ten minutes before our next presentation so I'd like to open it up. If you have any questions or anything you'd like to ask or comment on, please do so.

I'm not seeing any questions. Thank you very much.

Vittorio Bertola: Thank you for inviting us. I'll be around anyway.

Marc Anderson: All right.

This is Marc Anderson again. Lisa came in. Hi, Lisa. Why don't we - do you want to take like five minutes to get settled in and we can start with the next presentation? So we'll do a quick five minutes while we swap speakers, I guess, and start back up.

This is Marc Anderson again. We're going to go ahead and start here. We're a little bit ahead of schedule but after a quick conference here, we decided that starting early and ending early wouldn't be a horrible thing. If we could wrap up a little bit early today, we'll go ahead and take that.

That said, I don't want that to pressure - I mean, we're pressuring anybody into not asking questions at the end. We have Lisa Box here from Domain Connect, who's going to take us through our next session and is happy to field questions at the end.

But with that, let me turn it over to Lisa.

Lisa Box: Thank you so much for having us here. Just in case I have not met you, my name's Lisa Box. I'm with a company called WPEngine. But the reason why I'm here today is to support an open standard called Domain Connect.

And if you want to.

So the thing that WPEngine does is we are a WordPress platform so we are on the hosting side of the business. And we run into challenges all the time with clients not being able to configure DNS properly.

So if you want to.

And so we have gotten together with a group of other companies -- Google, Microsoft, GoDaddy, 1&1 -- to come to ICANN to advocate for Domain Connect. And that is to help service providers and DNS providers provide a better Web experience for customers using open standards.

So Domain Connect was created -- you can go to the next slide -- Domain Connect was created several years ago by a group of engineers, and GoDaddy was one of the original authors. And it was to help customers who are not technically inclined be able to configure and automatically activate their Web sites quickly. And as you can see from some of the statistics, Microsoft found that 75% of Office365 users were unable to configure domain names properly.

And we have several people from all of these companies here today. If you want to raise your hand, so that everyone knows who you are, if you're with Domain Connect. So we have all of the companies represented.

And the next slide.

So what we did was we came together to make the standard open so that anyone can join and participate, making it better, and we've adopted it across all of our companies. And we've seen significant results. And so we encourage everyone here to consider this a standard of operability for Web DNS providers as well as the hosting or the service providers, including anything that needs DNS configuration to operate.

And this is a little bit about the Domain Connect project. And Arnold, did you want to speak to the project itself?

Arnold Blinn: Yes, this is Arnold Blinn. I can give a quick overview of how it works maybe. That might be useful.

So the way Domain Connect works is a user who comes to a service like a Web hosting site or an e-mail site and sets up and purchases, for example, their application, usually what happens historically is the application would say, hey customer, you know, in order to use this product with your domainexample.com, go into DNS and add this A record, add this MX record, add this TXT record.

And often, if the service provider were good and had a lot of resources, they might even go so far as to try to figure out by querying whois data and registry data, who the DNS provider actually was, who's running DNS from my domain. Based on that information, again if the service provider were really good and on their game, they would provide a bunch of instructions and help articles to the customer.

The irony of this to me was that if the DNS provider changed or updated their interface, often the instructions are out of date. So the instructions in trying to be helpful in some ways are often actually detrimental to the customer being successful.

So Domain Connect was created as an alternative way, an easier way, to solve this problem. So with Domain Connect, the service provider or the customer at the service provider types in their domain name. And the first step of the protocol is a discovery process to reliably figure out who's running DNS. Now, that's very simple because all you need to do is do a DNS query for a known record in DNS, and that tells the service provider who's running DNS.

After that's done, then the service provider can message the customer, hey it looks like your domain is running at, in one example, GoDaddy. And in order to configure your domain, just click here. And when the customer clicks that link, what happens is is a Web browser tab or window typically opens up and it says, hey, welcome customer, please sign in to your GoDaddy account. The customer would sign in to GoDaddy.

We would verify then that the customer own the domain name because obviously if they didn't, we can't make the change. Once they - we verify they own the domain name, we simply ask the customer -- and all the DNS providers implement it the same way -- we just say, hey customer, do you authorize us to make changes to the domainexample.com to enable whatever service is being configured. When the customer says yes, we simply apply the DNS changes right away.

So the customer - it's secure because the customer logs in and authorizes the change. And it's easy because we just apply the change. We don't make the customer mess around with DNS records.

Now, there's a little bit more to the protocol than I'm describing here. There are some advanced use cases that use the OAuth protocol and an API that some of the service providers choose to use. And there's a couple other smaller nuances around security and how templates are created and how they're secured. But the overall flow from a customer experience is largely what I described. So that's an overview.

Lisa Box: Thank you so much. So let's talk a little bit -- if you want to go to the next slide -- about what this means and what the impact has been. So we at WPEngine have actually just implemented this in the last few weeks and we've already seen around a 30 to 40% reduction in calls that apply to DNS, which is our number one call driver. So it's hundreds of hours a week that we're saving.

But up here, you can see that there's actually more benefits to all parties, not just the service providers. And so - and I know that Microsoft has seen all various different types of benefits, including more domains being associated with the services, higher renewal rates, reduced tickets in support. And so everyone actually benefits in this chain. And so therefore we don't view this as a competitive differentiator but something that uplifts the entire industry.

And does any of the companies want to speak to some of the benefits that they've seen personally? (Florentine), I don't know if you want to speak to that.

(Florentine): Yes, I think the benefits are pivoted based on customers, providers and DS providers and service providers. As you can see, in each of the areas, I mean, there are benefits for all.

Some of the numbers that are not present there that are interesting is when a customer is trying to perform the operation that Arnold described, what we noticed is it takes roughly on average four attempts per domain if you don't have Domain Connect.

If you have Domain Connect, it's usually one. And this speaks about the customer satisfaction that we see so our end sat is increased by 20% for our customers.

And it goes without saying that those happy customers are going to renew the service with the service provider as well as they're going to go to the DS provider and either attach domains that they already have or go purchase new ones. So all these benefits just go across the segment for everyone.

Lisa Box: So one of the things - I've been in the industry - if you don't know me, I've been in the industry for a very long time both on the domain owner side, the registrar side and the service provider side. And I haven't seen any technology protocol that actually makes an impact quite as large as Domain Connect.

So one of the reasons why I'm here advocating for that is to make this the de facto industry standard across all DNS providers as well as service providers. And so we're - our companies are all investing time, money and resources to make this the standard of operability. So that's one of the reasons why we're here, because this does benefit the entire ecosystem and significantly helps the clients that we're trying to delight and earn every single day.

So if you want to move to the next.

These are a sampling of all the companies that are actually involved in Domain Connect today. And you can see, you know, we're representing many of them who couldn't attend today. We will have a case study at the booth. We're at Booth 15 down the hall at ICANN. And we have a case study that was co-authored by HubSpot. We are working with all of these companies to make this the standard.

So if you work with any of these customers or your customers work with any of these providers, which they probably do, this is something that will benefit both commercially and from a support perspective.

We can go to the next slide.

So I know this is a big hard but this is the time where we can open it up for questions if anyone has questions about the protocol itself, the standard or joining. This is an open standard so anyone can participate and actively improve -- you know, whether it be going to DomainConnect.org and looking at The Spec, joining our slack room -- there's different ways to support the initiative.

So I'll open it up for questions at this point.

Man 1: I don't want to ask - get too technical. But I was curious, on that first step where they're determining the DNS provider, what is the - like you mentioned there was a record that helps accomplish that, is that on the individual domains or? Like how is that accomplished?

Arnold Blinn: Yes, we don't mind getting technical in this Technical Operations group. This is Arnold again from GoDaddy. So I'll just describe it in detail and I'm sure it'll answer your question.

So the discovery protocol, the first step that a service provider does, is he queries a record from DNS. In the domain that you're querying, it is the underscore Domain Connect record. Now, the DNS provider can choose to implement that however they want. They can put that record in every zone if that's what they to do. If they implemented their own custom DNS server, they can just inject it in by responding to the query with the proper value.

You can put a C name in every zone and then point that at one master zone so you don't have to - if you have to update it, you know, you can do that as well. That's actually what GoDaddy does.

So we don't dictate how you implement it. What we do - what the spec does ask is that the domain responds to a DNS query for underscore Domain Connect. The contents of that value, what it returns is actually a fully qualified host name. So in the case of GoDaddy, it's something like API.DomainConnect.GoDaddy.com. That's actually not what it is but it's something like that.

What the service provider does with that then is the second step of the discovery the call a rest API. That rest API returns a JSON structure that contains all the information necessary for the service provider to implement the protocol.

We did it as a two-step process because putting a bunch of JSON in DNS versus a simple, you know, sort host name seemed a little bit more scalable, right?

You can look at all this. I don't believe Lisa mentioned, there's a completely functioning application on the Internet at ExampleService.DomainConnect.org. That is a fictitious Web hosting company that will host a Web site in your domain. The sophistication of the Web site it renders is literally a simple string so you can type the word hello and have that hosted on your Web site.

And the implementation of ExampleService, if you run through it, it actually shows you each step along the way and the details that the protocol is passing back and forth. So it's really not intended for consumers as much as it is for developers who want to understand how the protocol works.

That was probably more information than you asked for before.

Man 1: No, that was perfect. Thank you.

Arnold Blinn: Yes, thanks, yes.

Lisa Box: Other questions. Well, some of the questions we get is there's no commercial or economic value to the actual - it's free. It's an open standard. I just want to make sure that's clear. The list on the provider side is relatively small in terms of engineering effort to get sort of a crawl up and running. And so this definitely is a standard that is achievable to have everyone create as the standard across the industry.

And I would love it if we came to ICANN in one year and 90% of the DNS providers and providers were using this as the de facto. I don't know if it's possible but I'd love as a group for us to, you know, try to achieve that.

And then the next slide is how you can participate. So again, we have a booth. It's Booth 15. Please come and talk to our engineers, product managers and all of the champions at each of the companies. We'll be there throughout the week.

We're offering a networking reception on Tuesday at a place called Bananas, which, you know, is just to talk more about Domain Connect and interact with people who have implemented it. And then you can always go to DomainConnect.org and everything is there and available. So it's that simple.

Marc Anderson: Any other questions for Lisa or anything else about Domain Connect?

Lisa Box: Yes.

(Mark): Yes, this is (Mark) from WPEngine. To add to your kind of simple how to, so we're a service provider and have partnered to the other side of what Arnold has described for implementation.

And for our - the previous slide where it talked about cost, we're just using that example service and along with that is the open source pipeline code for that. So with a few lines of code, we are able to take that and essentially mimic what we would stand up as our own service as a service provider interfacing with the DNS side of it.

So it was easy for us to proof of concept it before we built our longstanding implementation of that. So it was easy to spend a little bit of time just to put that together.

(Florentine): Well, speaking on proof of concept -- and this is (Florentine) from Microsoft -- last - this year at Crawfest we had a hack-a-thon. And Plesk was there and they implemented a proof of concept for the product that they have today, over two days.

Man 2: They did it in half that so.

(Florentine): Yes. If you want to try it out quickly, it doesn't cost much.

Arnold Blinn: I'll just throw one more thing out about the hack-a-thon that's kind of -- (Pablo) and I think it's pretty interesting. (Pablo's) from 1&1 and I'm from GoDaddy. Last year, at the hack-a-thon the other project we did is kind of cool too.

So we built - you guys remember back in the day they had dynamic DNS. It was common in people's homes and a lot of routers, built into their router, a protocol from DimeDNS that was a custom proprietary protocol to update an A record in a random domain with your home's IP address as it changed.

So we decided to build one based on Domain Connect. So we now have a - this isn't built into a router. It's actually a service that can run on either a Windows machine or on a Linux machine that will monitor your home's IP

address and if it ever changes, will update an A record in a domain that you specify using the Domain Connect protocol.

So the cool about this little app that we built… And it's also, the ExampleService, this app we built, they're all open sourced and up on the Domain Connect GitHub. The cool thing about this is is that it'll work with any DNS provider that stands up Domain Connect. You automatically get a DimeDNS client as part of this, which is kind of a…

I don't know how popular those things are these days but it's still kind of intellectually interesting so.

Marc Anderson: Thank you and a last call. Any other questions?

All right. Thank you, Lisa.

Lisa Box: Thank you. Thank you so much.

Marc Anderson: All right. With that, we're at the end of our scheduled agenda. So to keep me honest here, am I missing anything?

Woman 1: Do you want to talk about (unintelligible)?

Marc Anderson: I'll say - I'll just say real quick, you know, GDDD Summit seems a lifetime away, especially since we're in the middle of ICANN 63. But planning has started for the GDDD Summit. Just like last year, Tobias and I volunteered to represent Tech Ops on the planning committee for GDDD Summit.

So now is the time, if you have ideas for how you'd like to see our Tech Ops session go at the next GDDD Summit, you know, like before, please funnel them to Tobias and I. You know, and we'll represent - you know, we'll represent Tech Ops on the GDDD Summit planning.

I know we had a successful summit in Vancouver. We'd like to continue that and hopefully have another one. We don't want to try and tackle that today though but keep that in the back of your heads, if you have ideas. And like before, we'll be looking for session leaders, session topics and volunteers to help us with GDDD Summit.

Anything else? All right. With that, I know we're a little bit ahead of schedule, but like I said, I don't think that's a bad thing if we give you guys back some of your time. So with that, we're adjourned.

END