
DUBLIN – Universal Acceptance of TLDs
Monday, October 19, 2015 – 15:30 to 16:45 IST
ICANN54 | Dublin, Ireland

EDMON CHUNG: Hello, everyone. This is the Universal Acceptance session. I understand there is a little bit of a delay in the technical; getting the slides up and all that. Please bear with us for another probably minute or so and we'll get under way.

Welcome, everyone. As the team is working on the technology part, I just want to make sure that if there is remote participation we don't miss it, so I was hoping to wait, but is that up? Is it just the... But the Adobe Connect and everything is up? So remote participation is good and audio, everything? Okay. In that case, then, I think I'll get the session underway once we have the slides it can show.

So, welcome again, everyone. I'm Edmon Chung from .ASIA. This is a Universal Acceptance public forum and we're very excited to have quite a number of experts on the matter and also people who are devoting a lot of time on the issue here with us on the panel but one of the more important things is really to get the community engaged in this journey of Universal Acceptance.

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

If there was a slide, then I was going to move to the next slide, which is going to talk – just a quick overview of the agenda for today. I'll be giving an update on what the Universal Acceptance Steering Group (the UASG) has been doing the last few months, and also we're happy to have Rinalia here to give us a little bit of perspective on why this is important. Why Universal Acceptance? Why now?

Then we will go to Mark who has been working with the UASG very closely and working on what we really mean by Universal Acceptance and how do we define being UA ready, being Universal Acceptance ready?

Then we'll go into a panel discussion. I won't introduce everyone. I think when we get to the panel, each one of them will talk a little bit about their initiatives from registries, from registrars, from other organizations, how different people are getting ready for Universal Acceptance.

So just in case you still don't know what Universal Acceptance is, then I apologize. I didn't explain this in the very, very beginning.

When we talk about Universal Acceptance it is about the ability for applications, programs, and the Internet really to accept internationalized domains, new top-level domains, internationalized e-mail addresses, and names that in previous times perhaps were unanticipated by some systems.

With the new gTLD program, with the IDN ccTLDs that are launched with e-mail that can have internationalized user names and domain names, this becomes very important. Really what we're talking about is an upgrade of the Internet. We're talking about getting the Internet updated for the next few billion of people to come online to access the Internet. So I'm on my third slide if you're trying to catch up.

So the UASG has been working really since between ICANN 52 in Singapore and ICANN 53 in Argentina. The UASG was formed and since then we have done a number of things. You might see a lot of things happening just yet but I think the best way – and Ram Mohan who is our chair, unfortunately he couldn't join us here, the way that he really explained it I think it's really good.

Right now where we are is that I think the ship is built, it's loaded, the sails are up, the wind is blowing, and we're about to embark on a journey. This journey's going to take a little bit more time, but there are a number of things that we are doing and we have been doing.

Now I'll quickly jump over a few – actually six things – that we have actually moved a bit forward on. First of all is EAI. What is called EAI is E-mail Address Internationalization and that really means the ability for having different languages as e-mail addresses. Why this is important, and we identified this actually

as like a top-line issue, is because this is kind of a superset of the issues themselves.

If the systems are able to use and recognize internationalized e-mail addresses, then more likely that will be able to handle IDNs, International Domain Names, very likely they will be able to handle new gTLDs and new TLDs. And that's quite important.

And the way that we are engaging, we are engaging with different organizations; Microsoft, Apple, Google, but also Postfix and other open-source software, projects that are [inaudible] utilizes, that helps us deliver e-mail. That's really what we're working on: to get them on board and also to identify good practices, sometimes called best practices, but there may be more than one way to do universal acceptance and we are trying to identify good practices. Next slide.

From the e-mail internationalization, what we found is a very interesting topic what is now called linkification. When you use your text editor, even with your messaging editor, sometimes you type in kind of like a domain name or kind of like an e-mail address – it magically turns itself into a link. But how do systems and how do applications recognize that, and how should those practices be and how should those experiences be? That, and how should the new gTLDs and new IDN and internationalized e-

mail addresses kind of inform that process is something that is very important for us. Next slide.

And of course one of the very important parts of universal acceptance is really outreach. And the outreach part is not just about us waving the flag, but also listening back and listening to the community on what needs to be done. And that hasn't been going on and that's one of the things we are targeting different organizations.

One of the core deliverables on the outreach is really a set of documentation, a set of materials, that can be provided out to the community at large and how they can fix it. Not only to tell them that these are new top-level domains, maybe you haven't anticipated in your software systems, but also how you can fix it. And in the coming year where we're going to have a number of different awareness raising events and we will be presenting at a number of different forums as well. Next. The connection is lost again.

The next slide is actually about registries and registrars. So one of the things that we've recognized here in the ICANN community is that I guess we need to be the leaders in terms of universal acceptance. So one of the things that we are looking at are blueprints or kind of the guidelines for registries and registrars to get ourselves universal acceptance ready as well.

And this is not only for registries and registrars, because registries and registrars ourselves are actually IT companies ourselves. We have CRM customer-facing systems. We have internal systems. We also have different types of systems that interact with the public. And there is a huge range of organizations from smaller registries or registrars to very large registries and registrars. So it actually gives us a very good opportunity to test how we kind of send the message out and make sure that they are relevant to other industries as well.

And then two more aspects of that, which I'm going to slide number eight.

UNIDENTIFIED FEMALE: I just got a message that the entire conference just lost the network so if it's important that this is being recorded, it's not currently. Now it's being recorded again. Sorry about that. Slide H?

EDMON CHUNG: No worries. It's good we're back online. And that's part of universal acceptance, I guess. Universal access, that would be. Anyway, down to the last couple of items. Measurement and monitoring. We need to understand what it is that we are working on. How we're measuring – how ready things are,

including longer ASCII top-level domains and also IDNs and e-mail addresses. So next slide.

And then for documentation. One of the key things here is that from some of the feedback that we've got, which we are actually on the third edition, third version of what we call a document that Mark is going to explain a little bit more is aimed at developers, but also creating documentation that could be aimed at CIOs, at an architectural level and different type of documentation that allows us to provide to the community at large how to address the issue and look at the problem. Next slide.

And quickly, one of the things that of course is very important for the initiative – and thank you ICANN for supporting it – is a budget that was allocated for the initiative and we are going through the process. Here are the few items that were listed, as I mentioned. Next slide. I couldn't see the slide myself.

So just as an administrative part of it, we're now formed – the UASG is well-formed now with the working groups. As I mentioned the different works that are going on, especially yesterday at the – next slide, please.

Especially yesterday at the Sunday workshop here we had more than 50 people participate from across the different regionals

and across the communities, and we have identified – next slide, please.

We have identified actually 16 task items, 16 action items that will be taken forward in the next few months and those include, as I mentioned, a CIO guide system, architecture guide, building on the existing introduction. We're going to break into documents that different people can actually use. We're going to establish better relationships with other organizations because the outreach effort is very important in this. We can't do it alone. We understand that. We're going to build a relationships with other organizations, like the MAAWG, which is the malware organization. ISOC, GSMA, the DNA, various other organizations.

There's also the issue of linkification that we are going to take more steps in scoping out what that really is and also creating a blueprint, as I mentioned, for this particular community and registries and registrars. With that, next slide.

Really, the next phase of the project is happening. We have been urging everyone to join and participate, to get it started. We're now started. We have identified the key aspects to work. Really, right now this is really when the rubber is hitting the road. We're scheduled for an intersessional in January also between now and Marrakech. Also we're hoping that, as mentioned in the last action items, there will be progress in documentation. Those are

the working groups that are being formed and they'll be delivering in the next few months.

So that's the update from where we are, moving from Buenos Aires to here. With that, I'll pass it on to Rinalia who will come back to talk a little bit more about why this is important as an initiative on its own. Rinalia?

RINALIA ABDUL RAHIM: Thank you, Edmon. As there is interpretation, I think I will speak at a slower pace for the benefit of those trying to hear interpretation.

My name is Rinalia Abdul Rahim. I am a member of the ICANN Board, but I am not here in that capacity. I believe I was invited to speak today based on the little bit of work that I have done on IDN matters. I was a member of the At-Large Advisory Committee between 2011 and 2013 and I was asked to look into the portfolio of IDN issues because no one really understood it. This was the end user community, and they needed to know what are the concerns related to IDNs that would pertain to end users. It was a deeply complicated matter, and that was when I encountered the problem of universal acceptance, which concerned me deeply.

In my previous life, I led a global multi-stakeholder organization where I built regional communities in all regions of the world, and that involved building multi-stakeholder partnerships to use information and communications technologies to help people in poor communities, basically to give them access to the Internet and services. That required providing services in local language and local scripts. The problem of universal acceptance stands in the way of that, which is why I delved deeper into this work.

This is a meeting on universal acceptance. I feel that I may be speaking to the converted, and this presentation may be a little basic. But if you bear with me, I will give you the perspective of Internet users on why this is important. Next slide.

So whenever you see a projection on Internet growth, inevitably it is always extremely rosy and positive. This is one projection. In 2020 we expect to have four billion connected people and more opportunities for revenue generation, more apps, more data, more embedded and intelligent system. But behind that there is a divide. Next slide.

There is the digital divide that we talk about. We have more than three billion Internet users right now against a population of about seven billion world population.

So this is a digital divide between the haves and the have-nots in terms of Internet access and ICTs. But within those who have

access to the Internet, the divide is with those who can have proper full access to the Internet and those who do not, and that pertains to the problem of universal acceptance. Next slide.

The other divide is when users are not able to register or use a domain name, users are not able to use domain names in their native scripts or language script of their preference, users ultimately not able to access and use the Internet, its services, and applications which are essentially bundled into the problem of non-universal acceptance of TLDs. Next.

So why do domain names matter to users? You all know this because you are experts, but when you speak to users in general, they may not understand this and you need to explain this to them. We know that Internet users are numerically addressed and domain names make it easier for people to access Internet resources without having to memorize numbers. This is particularly important in today's context because many of us don't even remember mobile numbers of our spouses or some family members. It's just too long. And IPv4 and IPv6 numbers are generally quite long. Next.

Now, why do domain names matter to users? First, it's accessing and using e-mail. Some people actually forget that in e-mail addresses there are domain names after the @ symbol. If you

speak to users, generally they're not aware of that and you need to bring that up to their attention.

Domain names are also used in accessing or using mobile apps. They're used for signing up for services. They're used for authentication of identity. When you're using multiple devices and they're not sure if you're the right person, they ask you for an e-mail address and you have to input it in.

But also domain names are used by application providers to serve users. So the question has popped up before in this community whether or not the value of a domain name is decreasing. For users themselves, there is a limitation. For those who provide services to users, the value of domain names is increasing. And if you speak to technical people within the Internet Engineering Task Force, they will confirm this.

Domain names are also using/accessing content via web browsers. Of course this was the traditional way of looking up information. Next slide, please.

So universal acceptance to me, this is an end state, where TLDs are usable in Internet applications regardless of script, length, or newness and it supports all Internet users around the world. And we are clearly not there yet. Next slide.

I want to touch a little bit about the evolution of top-level domains, because with ICANN, we've certainly moved into this direction. Next slide.

We have a multi-lingual world. Do we have a multi-lingual Internet? It's a question mark. I don't think we're there yet, either. So we have more than 7,000 living languages in the world currently and about half of those are spoken in the Asia-Pacific region, which is the region that I come from.

There are numerous writing systems or scripts. There is no specific number to pin to it. There are certainly more than hundreds of them. And the root zone up to the year 2010 had only ASCII top-level domains, which is a subset of Latin. This, for me, is a problem because it doesn't reflect the diversity of our world. Next.

In terms of evolution of TLDs, where we are going is we are going towards longer TLD labels and beyond ASCII. In terms of longer TLD labels, we have more than two or three letters, labels and we have more script diversity.

From ccTLD, we are moving to IDN ccTLDs and here you can see the scripts of Egypt, in Arabic, China, in Han script, Chinese and Thailand in Thai script. In terms of gTLD, we're moving for longer labels, but also IDN gTLDs in multiple scripts. Next slide.

Since the launch of the IDN ccTLD fast track process in the IDN gTLD process related to the new gTLD program of ICANN, there are requests for about 18 scripts in total. And quite a lot of them come from the Han or Chinese script and also for Arabic. That's the bulk of it. And the rest are distributed across different scripts. Next.

This gives you an idea of the diversity we're talking about in terms of IDN country code top-level domains. I want to draw your attention in particular to India, .in. You would see that for this country alone there are seven official scripts, and that's just for the name of India. That's the diversity of the world that we live in. Next slide.

I want to raise your awareness, if you don't know this already. ICANN has this project called the root zone label generation rules project. So we have the demand for the new gTLD program for 18 scripts, and some of those scripts involve the use of variant top-level domain. To deal with the safe deployment of variance we need to go through this process of developing label generation rules. That involves different script community to come together and to define the script repertoire for their language script that would be safely deployed for TLD labels in the root zone and the rules that will govern the deployment of variance.

At the moment, there are more than 18 script communities that have started work to do this, and within each script generation panel, which is the group working on it, there's expertise in terms of technical, on DNS, IDN, and Unicode, linguistics expertise, as well as policy.

I want you to remember this slide, because at the end of my presentation, I'm going to touch on why this important to your work on universal acceptance. Next.

This is just a point to say that if you solve universal acceptance for one script, you actually can have cross-regional implication. For the Arabic script alone – and this is used for 50+ languages across regions, from Africa, to the Middle East, to Asia. And in terms of the shades of green, what they refer to is the degree of Arabic script as official national orthography.

In the deepest color, it's the sole official orthography of that country. In lighter shades, it tells you that it's an official orthography but they may recognize other languages or Arabic is one of several official orthographies.

In the lightest shades where you see India, China, Malaysia, is where it is official at a sub-national level. In my country, the Arabic script is used as an alternative script for my language, which is Malay. Next, please.

As a recap, Internet users are reliant on domain names to access Internet services and applications. Some of them may not realize this and we need to make them realize this because we need them to be advocates for universal acceptance.

Now, some Internet users and software applications have not kept up with evolution of TLDs, rendering them unusable, and therefore they block user access.

Software and service providers have been unaware of the problem. They've had little access to architectural guidelines on how to set up their systems for universal acceptance, and there is little [inaudible] or regulatory incentive to invest in solutions to solve the problem. Next slide.

I think you know this, but just to recap. To support Internet users worldwide, TLDs need to be made usable in applications regardless of their script, length, or newness, and that means acceptance of TLDs written all scripts, acceptance of TLD names longer than three characters, and support for IDNs or non-ASCII names in e-mail. Next.

So what does success look like? It's basically addressing the challenges I've just outlined. When any person can register and use a domain name in any language script in any length in widely distributed web browsers, in e-mail clients, in mobile

apps, in setting up online accounts for the Internet and other services. Next.

And this is my final slide in terms of how to achieve success. There is a role for the user community and there is a role for the technical community. The technical community knows that they have to address the issues through technology, but the user community may need a little bit of help to understand why they need to advocate for this what they can do.

One thing they can do is they can identify and prioritize the issues as clearly as possible, so that those who can solve the problem can start working on it, and to ensure that there is clearly demand for it.

I touch again on this slide on the label generation rules because there are more than 18 language communities currently working on it. The Arabic Script Panel, which I am a part of, has completed its work together with the Armenian one, but the Arabic group has decided to move on to tackle the problem of universal acceptance. And this could potentially happen for all of those script generation panels, and they can be your allies and partners in addressing the problem of universal acceptance. With that, thank you.

EDMON CHUNG: Thank you, Rinalia, and thank you for providing a little bit more context, especially on the IDN internationalized domain name issue as well.

I'll now turn to Mark who will give us an update on part of what I guess what Rinalia mentioned, but a more robust definition, really, of what we mean by universal acceptance and how we really tackle it. Mark?

MARK SVANCAREK: Thank you. I'm Mark from Microsoft. Here's what we'll discuss for the next few minutes. There are five criteria that we define universal acceptance by, and then let's talk about people, actual user scenarios, not just about technology but what people will really do and why universal acceptance is important to them. And some counter examples as well, things that you might see today that are examples of the failure to achieve universal acceptance. We'll talk about some more details of those criteria, those five criteria. We'll get into some implementation details, and somewhere along the way I will give you some technical definitions as well.

So the five criteria are to accept, and to store, and to process and validate and display all the top-level domains equally, and all the IDNs, and all the e-mail addresses equally.

Here are those definitions. This is probably well-known to most of you, but let's just cover it for completeness. So most of the Internet today is using ASCII which is a very old standard and it only covers 127 characters. They're all within the Latin script. Unicode, particularly UTF 8, which is the preferred format, supports millions of characters and lots and lots of languages all over the world.

IDNs are domain names that contain Unicode. If any of the labels contain Unicode, we can call it an IDN. EAI is an e-mail address which is solely comprised of Unicode.

Punycode is an algorithm. It's an algorithm for translating back and forth, to and from, between ASCII and Unicode. An A label – and I'm sorry to throw out terms like this. An A label is what happens when you use Punycode to translate some Unicode.

So the example showed there is [inaudible], everyone. And you can see the Punycode translation, results in a string that begins with XN dash-dash if you see XN dash-dash which is the ace prefix, you know that the remainder of the string is an A label. You see that [mina] started as three characters and it became L-A-J-T-G-9-B, which is substantially longer. This is a thing that happens.

Okay, with those definitions behind us, we'll move ahead. This first example, I think it's near and dear to you all, but we list it for

completeness. Somebody will want to register a new gTLD. It might be an IDN. It might not. In the example here, it's a .brand. I'm a retailer and I had a .com, but it wasn't serving my interest because people were squatting on it, phishing against it, and so now I've acquired a .brand and I completely control who has access to it. It reinforces my brand identity. I can use it to educate people.

Accessing that gTLD, a user might type it into their browser. They might click on a link somewhere. The expectation is that even though the gTLD is new, that it will work just like the old TLDs did. Next.

Using an e-mail address as an identity is something that we all do. Maybe your bank, maybe your airline loyalty program. It could be really anything – your Facebook account. You supply an e-mail address and that becomes your identity for accessing that service.

Universal acceptance means that even if that e-mail address contains a gTLD, a new gTLD, that it will still be accepted the same as if it were an older domain.

Accessing an IDN, this is similar to the gTLD example, except now it's a Unicode string. So it's not merely new. It's new and it's in an alternate script. Universal acceptance means that

applications, browsers, all treat these the same as the older domain names would've been treated. Next.

Likewise, what if your identity is based on an e-mail address where the domain is not just a gTLD, but an IDN? It's the same situation. I want to use my identity, which represents me and my native script, and when I sign up for a service I would like to use that. And when I enter it into the box not be told, "Please enter a valid e-mail address", but have it accepted like any other e-mail address.

We mentioned something called linkification. That's dynamically creating a link within a document. Imagine you're typing something that looks like a web address, something that looks like an e-mail address, into a document. We haven't really defined what the correct rules are for doing this. There's a certain amount of controversy as to when this should be done or even if it should be done. That was a minority view that we heard the other day.

But whatever those rules are, they should apply consistently regardless of whether the thing that looks like an e-mail address is an EAI address or not. Or the thing that looks like a domain name is an IDN or not, or something that contains a new gTLD or not.

Then the last user scenario is what if you're a developer? How do you create applications that are universally acceptant? Or even more difficult, what if you are trying to maintain an older piece of code and you're trying to update it and trying to make it universally accepted? Like your tools to have libraries and support APIs that do this automatically for you, so you don't have to understand how to do normalization, how to do Punycode translation and the like.

The examples, if you can't do any of the scenarios we just talked about, that would be a failure, of course. But there are other things that you might be doing now. You might see this somewhere that aren't so obvious I guess.

For instance, if you're doing the validation on your own, if you're using a heuristic like a top-level domain should only be a particular length – 2 or 3 or 4 or 6 characters – if you're using a heuristic like that, you could be failing universal acceptance.

If you're using an out-of-date source, if you've hard-coded a list of acceptable names into your application, you're probably out of date and you're probably going to fail. If you are consuming from an authoritative source, but not updating very frequently, you could be failing universal acceptance.

Remember I showed you those A labels? Even displaying those to users is, in my opinion, a failure of universal acceptance.

Users really do not need to understand what the Punycode algorithm is and they shouldn't have to see xn dash-dash anything, ever. Certainly we should never ask them to enter it into an input box for any reason, whether they're a user trying to type an international e-mail to send to their friends or a system administrator trying to configure a website for a hosting company.

Lastly, I mentioned down-leveling EAI to A label e-mails. During our testing we have encountered some e-mail solutions that can accept EAI, a Unicode e-mail address, and then convert it to something else, to a different format, because presumably the app developer believed that that would be more compatible with the ecosystem, and they may be correct but that's not the correct thing to do and it would be an example of not being universally acceptant.

So regarding the five criteria, here are some more details. I apologize for how dense these slides are.

Applications and services allow domain names and e-mail addresses to be entered into user interfaces and received from other applications via APIs. So if you're an application writer, you need to know what the new formats are. You need to know what the new standards are. Things are longer than they used to be. There are different scripts. You have to know which

normalizations to use. You have to know what a normalization is.

One of the criteria is validation. This is one of the places where validation usually occurs. When you're typing something into an input box, usually there's something that is validating whether it's a proper – a well-formed string or not. That's where a lot of the frustration comes from today. You enter in a string and it says, "I'm sorry, that's not valid." Next.

Secondly, storing. Computer systems have to store data whether on a transient basis or on a long-term basis, either within working memory or within a database, something like that. Regardless of the lifetime, you should be saving them in Unicode. Definitely some RFC defined format, not something proprietary. If you have to use intermediate forms, make sure you can easily convert them back or forth. And UTF-8 is the preferred format.

Processing can mean pretty much anything. There are lots and lots of ways to process things. You might be using it to say everyone in New Zealand has a .nz. All the pharmacists have an @pharmacist e-mail address. A controversial one would be I'm going to block this domain because of its reputation. That could be done on a per-enterprise basis by corporate IT. That could be done at a national or regional basis because of national policy.

So that processing needs to be aware of what are the valid formats and be able to apply policies to it like that. So this is another place where validation typically occurs. Just a point that validation doesn't exist in a vacuum. It's happening during the other steps.

Then finally, displaying. So you have to be able to display all of the code points within Unicode, and usually you're depending on operating system resources to do this. So this is something that is usually provided by the platform, accessed through your tools, through common APIs. When it isn't, there will be application-specific transformations that you have to do. So some of this is happening at the application layer. Some of it you're depending on whatever platform provider you're working with.

We lost our connection again. I'll just wait a moment. Yeah, I don't know.

EDMON CHUNG: We apologize for the slight pause. It seems like we have lost connection. It's coming right back again.

MARK SVANCAREK: I think we're back. So taking all this information and thinking about it as requirements, like requirements you'd see in a spec,

if you look at this list of things as a developer, this is one way to determine if the app you're going to writing or the app that you're maintaining will have universal acceptance or not.

If you're approaching this from the perspective of an IT professional, this could be the list that you use to talk to your supplier or your contractor to determine if what they're giving you will support universal acceptance, whether a particular version is compatible or not, whether additional changes are required.

So an application or service, which supports universal acceptance, should support all domain names regardless of their length or character set, should allow the entry of international characters – that is to say, all the Unicode code points into all the UI inputs – can correctly render and display all of those code points from those strings, can do this correctly even if it's right to left or left to right. Next.

Communicates between applications and services through compliant APIs in compliant formats. And that would be your internal APIs as well as your external APIs.

Stores the data in these compliant formats, supports all the domain names from authoritative sources. The ICANN supported list, the public suffix lists.

Can send e-mail to recipients regardless of the domain or the character set. Can receive e-mail regardless of the domain and character set.

And finally, supports accounts that are associated with multiple aliases, whether those are a combination of ASCII or Unicode. I think this last one will be more common. In the interim, as the ecosystem comes online and brings in more and more universal acceptance, people who have ASCII e-mail addresses will want to add EAI addresses and find that some of the people they're communicating with don't have universal acceptance in their e-mail systems, and they'll wind up with multiple accounts. We need to support those configurations as well.

I hope that was helpful. That's it for me. Oh, just a last thing. As mentioned, we are developing some documentation that goes into more detail about the best practices for developers. We hope that you'll review that and contribute to it.

EDMON CHUNG:

Thank you, Mark. I guess this is a taste of some of the documentation that is being developed here. With that, I will move on to a panel that is looking at actually implementing universal acceptance and looking at the issue, what to actually do.

With that, I see that Ashwin is in the room, so allow me to invite Ashwin, CIO from ICANN, to also join us. We particularly want to thank the IT team at ICANN, who really for the last six months have been actively working on universal acceptance for the systems across ICANN and starting to understand the issue. It seems like a very thorny issue as we look into it.

Sometimes I recall a colleague [inaudible] that described the issue. I think it's very well-described. It's kind of like a Russian doll where you open one and then you see another one, and then you open another one and you see another one. You keep opening them and they keep having another Russian doll inside.

I guess in order to avoid that one of the things that's important is to really plan ahead and think about it. I'll start with a very small project. Sometimes you think about the entire – I'm sure Ashwin will talk about the entire ICANN in that scope.

But recently, my organization supported the ICG commenting forum. It's just a public comment platform that's on the web, but also taking e-mails as public comments. That small project itself included, I think what Mark said, how to accept internationalized e-mail addresses, how to store them, how to process them – because you have to send back the e-mail. If somebody used an internationalized e-mail address, how to

send it back, to validate it, whether it's really an e-mail address. Is that something? And then to display it.

That required us to look at the different libraries that we use. We found that a couple of libraries for validation of the e-mail address or validation of the domain name needs to be updated, because they didn't understand IDNs. They didn't understand internationalized e-mail addresses. We have to swap out the mail server. We have to also add additional Punycode conversion aspect to the platform to allow it to be universal acceptance.

Even in such a small project can actually entail a number of different areas that require a little bit of attention to universal acceptance.

I guess, with that, I'll turn to the panel. I'll turn to Ashwin first and then go across the panel. We talk about different organizations, what prompted some actions to be taken, how UA was originally looked at. "Maybe it's a very simple project. We can get this done in three months." And then what was done to explore it and then what insights or findings where the definition, the perception of UA, how that has changed. And most importantly, what's the plan and the roadmap going forward and the key challenges you see ahead.

I'll go with Ashwin first, and then go to the panel across.

ASHWIN RANGAN: If you can, for me, rephrase your question. Is it a specific question or just stream of consciousness that you would like me to go with here?

EDMON CHUNG: Right. Perhaps in a way, stream of consciousness. It was a general number of questions. How different organizations approaching the issue, what you found as you studied the issue, and then what's the roadmap forward and the challenges you see ahead.

ASHWIN RANGAN: Thank you. Thanks again for inviting me to participate on the panel. At ICANN, obviously we're close to the matter that is being discussed because this is something that we were intimately involved with as the whole process began.

Our first challenge was just to get a handle on our portfolio of services. It was interesting. It took much longer than I thought it would. It started out with a listing of services that we knew of, and as we started digging around, we started to stumble across things that were infrequently used or used by a few people frequently that didn't belong in the catalog.

So over a period of probably three to five months, we were able to consolidate a list that was a representative and definitive list. So it was “the” list.

Once we had that, we started to break it down with different slices in different slices. One was to look at which ones of those had been coded in-house versus which ones had been procured off the shelf or had been leased for use if it were a SAS service.

The two are very different problems. If you think of ICANN, we are probably representative of a small to mid-sized company for an average software company. We’re not a large company in terms of revenue or footprint or any of the usual measures that they would use.

Our ability to influence the roadmaps of off-the-shelf software providers is rather limited if we were to approach them singly. It’s also, as I shared with the UA group yesterday, a bit of a double-edged sword when we go back to someone who has a contract with us to provide us with software.

We found that many of them were unaware of this problem. As we look at our contracts, not all of our contracts – none of our contracts, as a matter of fact – have standard clauses. So we’re in the process of creating standard clauses to use.

As we approach the suppliers, their level of awareness varies, and depending on where they are in their roadmaps, they may choose to work with us to understand the problem, or perhaps even to charge money for it in order to solve this problem.

The reactions that we're getting are across the board. Depending on who we talk to, it's a different reaction, so that prompts us to be careful as we take our next steps ahead.

And when it comes to the in-house developed software, there are three or four different variables that are impacting the speed with which we're able to move. The first and foremost is, like most other companies perhaps, we have a variety of platforms that have been used to develop software and place them in service.

And these platforms have gone through multiple revisions over years. And the farther back we go, the less documentation we are finding. The jeopardy there is even if we have skills to modify code, we don't want to do something for the benefit of UA and break a service that's actually working but doesn't have documentation. It's sort of like we want to do something right, but in the process we don't want to break something that's actually working and find that we are unable to move forward and just go back to what we previously had.

The fact that we have multiple platforms is in itself a challenge. Because some of these platforms are inter-connected, we need to be sure that we are understanding those linkages and making sure that the UA retrofits are not breaking those linkages.

The third challenge is when we look at the RFC, they are so far taking us to a particular point, but not necessarily getting us all the way through to the done/finish line. So there is more work that needs to be done in terms of defining the full comprehensive set of what does it take to get from “we have a UA need” to “we are done and we are now UA compliant.”

So there’s a variety of things that we are finding as we go through the processes that we are, with that part of the portfolio which is internally crafted.

EDMON CHUNG:

Thank you, Ashwin. I guess, as you mentioned, the scope is somehow often bigger than what originally think. So with that, I’d like to perhaps go to the end of the table and then come back. I’m conscious of time. I apologize, but please keep it to about two minutes or so, so we can take a couple of questions as well. Also, the experience from your organization and tackling or looking at the issue of universal acceptance. Jeff?

JEFF HUSTON:

Thanks, Edmon. Jeff Huston here from APNIC. We operate in the Pacific. We're the regional addressing registry for that area. Obviously we have clients and customers drawn from a wide variety of scripts and languages.

This is not just about trying to figure out how to put scripts and languages into domain name labels. That's the easy bit. The hard bit is actually thinking about Unicode and taking an ASCII computing system and converting it to a Unicode environment. And everything you ever assumed does not hold.

The glyphs collide. There are many ways to represent the same thing. They look distinct from the ones and zeroes. They display the same on your screen. They don't go from left to right or right to left. They go in arbitrary directions. You can't put a Unicode string together of any value and expect applications to make sense of it. This is horrendously difficult, and it's not just domain names.

So when we did an audit, the first thing we looked at was e-mail. Can we actually find any old user name, which is not the domain name part – it's the user name part. Can we accept it? The answer is right now no. Will our vendor do anything about it? No.

Why not? Well, from the vendor's perspective, it's of low importance. Little commercial value. Nothing will happen. Fine.

Let's go and look at the other thing we do a lot of, and that's our registry, and in particular WHOIS. Venerable protocol. Really old. Can it do Unicode? No. Will it? No. Why not?

Sorting and searching, really big problem. Syntax. Protocol elements that say "I'm going to talk to you in Unicode, not ASCII." We have to change all that. But it's not just us. All the clients out there have to change. They're not going to change. Will that happen? Probably not.

So the bleak part is that most of our front end isn't going to change. And it's not IDNs. It's Unicode that's the problem. So the only good news that we found, our database does UTF8. Ripper. Thank you.

EDMON CHUNG:

That's a very realistic and bleak situation, but this is why we're here and this is why we're trying to make it better. Another bleak situation perhaps.

UNIDENTIFIED MALE:

Hi, I'm [inaudible] from GoDaddy. When looking at the UA environment, what we found is basically our contact with the registrant starts on the purchase path, and you need to allow – or registrars, first of all, need to allow – the customer to enter in

their data, their contact information, in the language and script that's native to them that they understand.

Most of the world will not enter their data in ASCII, so they'll enter in a UTF-8 language and script, and then the registrars need to allow that to happen, and then to store that in their database and make sure that their database is supporting UTF-8.

The next issue that comes up is that not all registries support UTF-8 characters to be used to register the contact, or to create the contact. If you can't create the contact, you're not going to be able to register the domain name for the customer. That means that the information would have to be translated or transliterated.

The next issue that we came up with was the e-mail address, as Mark was saying and Jeff has alluded to, too. E-mail addresses, if they're entered in UTF-8 and not ASCII, there's a problem with making sure that the customer's going to get an e-mail to be able to verify the domain name according to the RAA specs. And if they're not, the domain is going to have to be deleted, according to ICANN policy. So that creates a huge issue.

The next issue was, as Jeff talked about, was with WHOIS. WHOIS clients are not standardized. Some WHOIS clients will display UTF-8 characters, but others won't.

EDMON CHUNG: Thanks. Please?

ELAINE PRUIS: Hi, I'm Elaine Pruis from Donuts. Before I start, I want to know if anybody here knows this person who lost their passport from Finland. Anybody? Any Finnish people in the room? No? Okay, I'll keep looking.

So I'm the vice president of operations at Donuts and my main job is to make sure things work. Part of that is to make sure that our domain names that we sell to our customers work for them. We have a customer support system where we receive tickets from registrants very rarely, and usually in the case where they can't get their issues solved through their registrar. We're sort of the backstop whenever there's a problem.

When we first started TLDs a year-and-a-half ago, we got a customer service inquiry from somebody who tried to create an online bank account login with a dot new TLD. I don't know if it was .email or .domains or something like this. It was Chase bank and it didn't work.

So this issue is a problem for our customers, which makes it a problem for us. In response to that, we tested Chase bank's online forums and then took it upon ourselves to look at the top

100 most frequently visited websites back in 2014 and see – could we just create an account using a new TLD? In that exercise, we found that 30% did not recognize new TLDs, and that wasn't even an IDN. It was just not .com.

That was our first brush with this issue and we became involved with the DNA's Universal Acceptance group as well as the efforts that ICANN has made in this arena to try to help solve this issue.

But our in-house current approach to the problem is quite informal, but I find interesting. We have 38 employees who, whenever they find something like this, they pass it along. We make an effort to contact anybody we know at that organization or just their own customer support and say, "Hey, look, you've got a people here. People can't use their domain name. They can't use their e-mail address." Whatever it is. So just basic "please help your customer."

In the last couple of months, we've started a project where we look at the ALEXA 500, so the 500 most trafficked websites on the Internet, and we tested 320 of them. Some of them were duplicates – google.jp or google.de. We didn't double test those, because Google actually handles UA pretty well.

Of the 320 we tested, we found 252 could accept a new TLD in their forms or their account creation forms. And we found 70

that had a problem. Either they didn't recognize it or the e-mail wouldn't come through their mail servers.

The outreach we did there was just to contact their customer support, raise the issue with them, get it on their radar. We had three of those 70 were fixed within a couple of days, and the rest we just need to keep pressing them on the issue.

I've heard that it's not really of material interest to vendors to fix this problem, but considering that there are four billion Internet users predicted for 2020 and 50% of them will be in Asia, I think there might be some impetus.

Thank you for inviting me.

EDMON CHUNG:

Thank you, again. That's why we're here, trying to change that situation. Jennifer?

JENNIFER STANDIFORD:

Thanks, Edomon. This is Jennifer Standiford. I'm with web.com. Taking a little bit from [Jody's] perspective talking about the registrar, the front line to the registrant, the customers, on a daily basis.

We are addressing this problem just from an ASCII perspective, not even tackled the issues around IDNs since most of our customers are English-speaking.

They range anywhere from Fortune 500s down to small businesses that also have a substantial group of customers themselves. So I'm going to take the approach of just talking about what we've... Our activities. Mostly grassroots activities and focusing around the awareness, outreach, and education.

From the standpoint when someone comes through one of our purchase flows of a registrar, if a new TLD is purchased – and especially if a new TLD plus an e-mail account is purchased – they get a custom e-mail that talks specifically around some of the issues focused around universal acceptance.

What we're asking of our customers is to allow us to partner with them as they experience similar situations to what Elaine at Donuts was speaking of and reaching back out to Citi or Chase or whoever it might be so we can leverage our relationships that we might have through them through resellers or affiliates.

We partner with our customers to go out to the vendors, harbor companies, and IT companies to bring universal acceptance awareness and help them drive it on the roadmap.

We've also faced our own issues with vendors that we use in-house – CYSCO, Juniper, the like, as well as software vendors – and we're trying to leverage our relationships with them to ensure that they get it on their product roadmaps. But we have faced challenges with that as well.

So I would say, for the most part, from our customers coming back to us, we have created our own internal directory of various different web properties in which we've encountered errors and issues around using the use of top-level domains, new top-level domains. So we continue to track those and do any sort of outreach efforts possible to ensure that they can get it on their roadmaps. But like everyone else at the table, we face those issues on a daily basis.

From our own internal perspective, similar to some of the steps that GoDaddy has taken, is we had to go through our own storefronts – our CRM tools as well as our systems – to ensure that we were universally accepted. I can tell you, as of three months ago, we weren't. But as of today, we are.

So taking the same approach, going to my CTO or my CIO, similar to what I would take to go to an external CTO or CIO, is we're selling a product. We want to make sure our customers will be able to use it and support it through our website, and we try to deliver the same message externally.

EDMON CHUNG: Thank you, Jennifer. At least we're seeing some progress now. So going back to what Jeff was saying, perhaps as we go along, more and more of those companies would start to be more interested, hopefully. Then of course, let me go to Dennis from VeriSign.

DENNIS TAN: Hello. Dennis Tan from VeriSign. So I really don't have anything new to add, other than to echo all of you what you just said. As VeriSign, as a company, we have to – I see universal acceptance on two different levels. One, we are a company as anybody else is a company, but we also are a registry. So of course we have the specific registry/registrar/domain name industry specific which is how we store, process contact information from our registrants, and we rely on standard specs to do it and we don't want to break stuff, so we pretty much have to follow those.

From a company standpoint, we're looking at what are the systems that really touch on storing e-mail addresses in general and whether they can't be internationalized e-mail addresses. Same as domain names, whether not it's ASCII or IDNs.

But we find ourselves, same problem as Ashwin is encountering, is that there are inter-dependencies across this system. So these

are systems that we have not built ourselves, so we have to contact and pretty much what you are doing, reaching out to these companies and explaining what the problem is. But depending on the level of engagement, contracts, all of that to some extent limits our ability to solve the issue.

So I think we are all on the same boat, but as a group, our voices are getting stronger to some extent, and hopefully we don't shy away just because there's rejection.

I used to work in sales, and there was a saying for 20 nos, you're going to get one yes. I think that right now we are in that stage where we're going to get nos several times, but just one yes that we need.

EDMON CHUNG:

Thank you. That seems to be the case, right? Hearing from different organizations, probably those from the registries and registrars, we're taking first step to make sure this is out there, and hopefully the world will continue to follow.

With that, I'm conscious of time and I wanted to open up to the floor and see if there's any questions, rather than just continuing on the panel. Any questions from the floor for the panel? If not, actually, I see a question. Oh, there is a question.

ALEX LEE: For the record, Alex Lee, .trademark in Chinese. It's great to hear this wealth of experience and expertise as far as it's come so far in terms of universal acceptance. What I find is interesting and also a little bit disconcerting is the lack of – I wouldn't say representation, just experiences in terms of IDNs, as it relates to universal acceptance. Maybe that's just because of the current makeup of the panel. I'm sure that within all of the relevant working groups that there are IDN representatives from either registries, companies, or software providers and whatnot that do have experiences that would be valuable to share in an open discussion.

But I'm wondering, of the panelists that are here today, although most of you have shared experiences mostly with regards to ASCII universal acceptance issues, have any of you experienced, come across, or are experiencing anything related to IDNs? And if so, could you share a little bit about those specific experiences?

DENNIS TAN: At VeriSign, we operate IDNs at the second level as well as we apply for 11 IDN TLDs. I'm not going to talk about those. Something that's dear to my heart, because I'm the IDN project manager at VeriSign, and oftentimes I have to – don't tell anyone. Oftentimes I find myself working with my marketing

team because they put out a campaign where you have a search [inaudible] to find your domain name, and just as one domain name – and it’s always an IDN – I type the IDN. Sorry, you are not putting a valid domain name. Shocking!

So it’s a problem that I think internally it’s becoming more and more of something that they now have to pay attention to. That’s internal. Outside, what we’ve seen, I think what Mark in his presentation, is how browsers tend to treat IDNs, and the major browsers right now process IDNs pretty much well. Depending on where you are and depending on your settings, that might change a little bit. And implementation varies across the board.

I think in the USG what we are trying to do is to get a summary of those good practices from major software providers and try to put it out there so that small developers right now trying to launch to market mobile browsers for a niche market, they follow the same process and they offer the same user experience.

I’ve seen out there mobile apps processing domain names, but they process IDNs halfway, meaning they can receive an IDN, but it turns back the Punycode label and that’s what you see. So big browsers, okay. Small browsers, not there yet.

I think what we want to do is to bring everyone on the same boat and implement all the practices in a consistent way.

EDMON CHUNG: Thank you, Dennis. I note that we have actually run out of time. But one comment from Jeff, and I note that there's one comment from online as well. Jeff? Keep it short, thank you.

JEFF HUSTON: For APNIC we conducted a widespread test of 300 new gTLDs. They included 10 IDNs. We actually found that the IDNs had a significant error rate, that users from all over the Internet were, to a significant extent, unable to retrieve them.

This is actually a result of a particular widely used scripting language used inside browsers. The problem is that the scripting language couldn't tell the difference between the Punycode and the Unicode. Got terribly confused and decided that it wasn't going to fetch anything anyway.

This is part of the problem with IDNs, that internally there are two representations that are equivalent, and sometimes you get internal code [parts] that lose that equivalence and get horrendously confused. Evidently, bug reports have been put in. The scripting vendor concerned will fix its problem, but it's

pretty obvious that we're early days with IDNs, and more bugs are yet to be uncovered. And yes there are problems. Thanks.

EDMON CHUNG: And we are here to solve that again. Last comment from online.

UNIDENTIFIED MALE: This is a comment from Ram Mohan, chair of the UASG and Afilias. It says, "When Afilias launched .info in 2001, we inadvertently became the poster child for universal acceptance issues. Over the first 10 years of its existence, .info struggled to become accepted universally in applications, software, [inaudible] equipment, and enterprises. We are at a much better place now for .info, but we are now the registry provider for around 250 TLDs that are all greater than three characters, and we find that many systems are still overly rigid in their acceptance criteria for both domain names and e-mail addresses. Our largest corporate clients who are launching TLDs in their own corporate name are finding that their TLD is sometimes having problems resolving inside their own corporate networks. Not only is this embarrassing, it creates an acute business justification issues for those who advocated for the TLDs in the first place. This is a global effort and we should think of this as a marathon, not a sprint."

EDMON CHUNG:

Thank you, Ram, and thank you for joining us here. I apologize for overrunning a little bit. I blame the little bit of technical glitch that we had earlier on. I guess as Ram concluded, in his remarks, this is going to be a long journey. We understand that. But this is the community that needs to take the first step, and I think the UASG is one of the first steps, and now we have identified a number of work items. Please come and work together with us on those items, and hopefully the bleak situation that Jeff mentioned will, over time, be ironed out much better.

So with that, thank you for joining us and please give a round of applause to the panel and to the speakers.

[END OF TRANSCRIPTION]