

LOS ANGELES – Universal Acceptance
Wednesday, October 15, 2014 – 16:30 to 17:45
ICANN – Los Angeles, USA

FRANCISCO ARIAS:

Hello, everyone. This is Francisco Arias, Director of Technical Services within the Global Domains division at ICANN. We're going to give a brief update on Universal Acceptance followed by a discussion panel on the topic. This is the agenda for today.

So let's start talking what is Universal Acceptance. In the context of ICANN, we started this [inaudible] Universal Acceptance a few months ago, and one of the things that we first realized is that there were different opinions among the community about what the problem was that we needed to solve. So we started a discussion, and the community came back to us saying there are three issues that this initiative needs to take into account.

First is the lack of support for all TLDs within applications and Internet services. Second issue is lack of support for internationalized domain names in general – not just IDN TLDs, but IDNs at all levels.

And the third issue that the community wanted ICANN to work on is the lack of support for internationalized e-mail, seen as the potential clear application for IDNs, but a protocol in itself that has its own complications.

So with this, we published a roadmap describing the objective of the initiative, which can be described in these three [holes] that are not surprisingly directly mapping from the three issues that the community

identified for us. And so these are the goals of the Universal Acceptance initiative: to work on promoting the support in applications on services for all TLDs, IDNs, and internationalized e-mail.

We published this roadmap that was subject of public comment within the community, and the updated roadmap was published just a few weeks ago. Besides the goals, the interesting things to mention about this roadmap is that the intention is to foster discussion here [to] elevate the awareness of this issue to the same level as IPv6 and the DNSSEC have been for ICANN to promote, among other things.

In this roadmap we recognize that the issue of Universal Acceptance is far beyond what ICANN itself, the corporation, can do to solve the issue. This is a big issue that needs the cooperation of all the entities that may be interested in this topic, namely the very same TLD registries that are interested on having – perhaps the primary interested parties – in having this Universal Acceptance support.

We have a new website with some new content, and we are planning to have more content as we produce it later. So now let's talk about a couple of issues that we are dealing when working in this initiative.

We are working right now on developing a communications plan. In the roadmap, we identified seven groups of stakeholders that will need to be reached out in regards to this topic, going from perhaps obvious in the software and service vendors and providers to the end users and everyone in between. But one of the difficulties has been is how do we measure success? How do we measure Universal Acceptance?

We have been thinking on and have heard suggestions from the community, but I don't think we have yet a good idea on how do we go about this problem. We could have some kind of census in which we go on check all the applications to see what is the level of support they have for the three problems that we mentioned before. However, there are some limitations to this topic, is we probably cannot go and make a census of all the applications that exist over there. That's probably too much of a job. [Who knows?]

So one of the ideas that has come up – and we are by no means settled on it – it's just one idea to throw here and we would like to hear your feedback, is on this and any other ideas you may have is, we were thinking on some sort of assessment that we could define and that could be published and promoted with the different software and service providers so that they can assess themselves the level of support they have for Universal Acceptance. Like I said, we are more than open to hear comments and ideas on how to go about this problem.

Now, since we are planning to go and talk with a big number of people, these [70] stakeholders, and ask them to offer support for these technologies, we were thinking, “Well, perhaps before going and talking with them (or at the same time as we go and talk with them) we start leading by example and have our own house in order,” as we say in Mexico where I'm from.

So we need to lead by example, and one of the things that we need to have is ICANN itself, the contracted parties, the registries, the registrars, we should show the world what does it mean to offer Universal

Acceptance. Have our systems and services upgraded to support Universal Acceptance, support all TLDs, IDNs, internationalized e-mail.

So at least in the case of ICANN, we are already starting the work. One of the things we are going to do is have our standard contractor template to include language that requires that every new software or service that is delivered to ICANN will have support for these three technologies. And we are also starting work on a plan to upgrade our existing services/systems to support the goals of Universal Acceptance. So that's in the case of ICANN what we are starting to work.

And this week, and actually before and during what is called – there is an initiative in place by my colleagues in the registry services [area] called the Registry Roadshow. It's a service of three presentations, meetings. I'm sorry – one in LA, the other was in Tokyo, and there is one more coming in Istanbul. And that's directed to the registries.

In that roadshow we have been presenting about this topic and calling the registries for action in effecting these changes in their own systems so that they can show the work, how this can be supported.

And also another important thing is to create demand, show the providers that there is interest on having these technologies to be supported. During this week we also went to the registries and the registrars once again to ask them their support for Universal Acceptance.

Well, with this, let's now go to the panel. Let me pull this up. Let's start with Jothan Frakes, a Mozilla Foundation volunteer in the Public Suffix List.

Jothan?

JOTHAN FRAKES:

Thank you very much, Francisco. I was very delighted to see the growing efforts and the growing support and resourcing that's being put towards this important cause of Universal Acceptance for all the new TLDs and existing TLDs, and IDN TLDs, as well as international e-mail. It's very good to see that and it's very inspiring.

It is something that I've been doing as a volunteer since 2006. I got involved with IDN and was introduced to something called the Mozilla Public Suffix List. It's a list of top-level domains, or suffixes, that enhance the IANA Public Suffix List. The IANA list of TLDs, that include ccTLDs, and TLDs like com, gTLDs. And then the Public Suffix List was created to fix and remedy a security issue around cookies that would occur and descended further. Many of the ccTLDs have a more diverse and a variety of ways that they will work with their substrings.

So within the UK – and they've recently liberalized this – but they would have co.uk, or org.uk, or ac.uk. So whereas one list might say, “only UK,” the other would go in and descend into the variety of different strings. Because for security purposes, while you might if you're Amazon.com or another large online vendor want to be able to share cookies across Amazon.com, you may if you were Amazon.co.uk not want somebody other than you to have cookies shared across co.uk, and it created security issues. And so the browsers reached up, created this file that had this detail in order to solve the security problem.

And the uses for this file have grown over the years. The file needed editing. Mozilla is a volunteer organization. It's a group of developers globally, every country, every language, and they didn't have an overlap though with this industry, within IANA and ICANN, and I did. And so I had the opportunity to step up and volunteer my time towards helping to mesh those worlds together.

Now, recently we've increased the velocity of introduction of domain names into the root, which is going at quite a nice pace, and I compliment ICANN on that pace. And I'm not complaining about the rate of adoption of TLDs up until 2014 in the root. However, the pace was quite notably different. Many more strings are being added.

And so what I'm talking about at this moment is about general TLD acceptance. The Public Suffix List and the IANA List each get updated and used by software, within browsers, within search engines, within a variety of different developers and integrators use this list for security purposes, for identification of what is a domain and what is a unique portion of a domain.

We were able, and grateful to ICANN for the help in this, to align the pace of the updates of the Public Suffix List through their help, with something that updates at the pace of the contracting rate that ICANN currently has. The team with Francisco –and thank you Francisco for the help from your team – we were able to have a csv file that gets updated twice daily I believe.

UNIDENTIFIED MALE: Four times.

JOTHAN FRAKES:

Four times daily – goodness. And then we're able to take that and update our file with more frequency to keep up with the pace of new TLDs.

As Francisco had demonstrated, there are browsers that have different behaviors. I can talk towards Mozilla, and I'm glad to talk with each of you. I have a very brief amount of time, but anyone who wants to talk afterwards, I am glad to.

But browsers have different activity that they perform. Some, like Safari by Apple or Chrome by Google, have an integrated search and domain lookup function within the location bar. Others, like Mozilla or Internet Explorer simply perform a DNS lookup or do less type of activity with the URI.

In the case of the Public Suffix List and how it's used within browsers, that's often something that's used to route where that traffic goes. For some of the new TLDs, as they're launching, when you go to type something.newTLD, depending upon what browser you're using and on what platform, you might see either a DNS error or you might see a search page that you're directed to as was evidenced on the slide. And that has to do with the presence sometimes in the Public Suffix List or in the IANA Suffix List and how the software will incorporate it.

Now, there were a finite number of domains for a very, very long time, and you can't fault developers for trying to be efficient, and many of them were. They incorporated a very efficient process to parse what was a domain or was not a domain. But in this new era of the Internet

where we have this pace of adoption and growth there is new ways that are necessary. And the Public Suffix List is one of the ways that people are doing that, to be able to validate what portion of the domain is a domain, and is it actually a valid domain. And a lot of times this protects you from adding spam or other things that you don't want to get, and also helps you to get to the places you do want to get.

To conclude, the work with the Public Suffix List is just the tip of the spear. There are many programming languages that incorporate it. There are libraries for Ruby and PHP and Perl. Because it's one, I say, the least awful resource out there that's publicly maintained and available. Many, many different open-source libraries include it and incorporate it.

And so with this small work that I'm able to volunteer my time with and with the help of ICANN, we are able to do quite a lot to affect programs form validation and other things that are derivative uses and incorporations of this file.

So I thank ICANN for their help in keeping this updated and the future evolutions of his file, because I think that this is going to be very helpful. When programmers are incorporating this library, it accelerates the process of Universal Acceptance. I'm grateful for your time. Thank you very much.

FRANCISCO ARIAS:

Thank you Jothan. Now I will pass the microphone to Brent London. Brent is a Program Manager at Google where he works on acceptance for new TLDs, internationalized domain names and Internationalized E-mail Addresses at Google.

The most recent change in this area is e-mail support for three part internationalized e-mail addresses. He'll be talking today about some of the unique challenges of internationalized e-mail addresses and why they present different problems than IDNs. Brent?

BRENT LONDON:

Thank you Francisco. So, like Jothan, I'm very excited about ICANN's proposal to coordinate some of the efforts and encourage adoption. I think that some of the problems that Google has run into while adding support for e-mail address internationalization are coordination problems. We need some authoritative body to help tell us what is right. In general, the IETF does a good job, but there are some gaps. So what I intend to show today is some examples of where we've run into those problems and where we need that coordination.

So, just for anybody in the room who's unfamiliar with the way internationalized domain names work, when you type in an internationalized domain name into your web browser what's actually going on behind the scenes is a DNS lookup is being performed on an ASCII encoded form. So the nice thing about this is the applications that you're using don't necessarily need to support internationalized characters.

So if you're dealing with systems that don't actually have the correct fonts or aren't designed to work with internationalized domain names, the user will end up seeing this Punycode form on the bottom. It's encoded and it's ugly, but it actually will generally work.

And in the context of e-mail, when we're talking about just internationalized domain name e-mail, much of this is exactly the same. So if you have an ASCII local part, and you have an internationalized domain name, the domain name itself can be encoded using Punycode.

Where we really run into problems is when we move to the left side of the @ symbol. So on the left side of the @ symbol is the local part. The local part has a few notable characteristics that make it different than just internationalized domain names. One is that, unlike general convention, it's case-sensitive. The interpretation of the local part is entirely up to the destination mail server.

So for example if you are a Yahoo mail user, and you send an e-mail to a Gmail user, Gmail disregards any periods that you put. So for example if you put, I don't know, `corgy.crazy@gmail.com`, it will go to the same location as `corgycrazy@gmail.com`. And that's because the Gmail product team made a decision to disregard that and it doesn't rely on any of the systems to support that. That's just all interpreted on the Google side.

In general, though, these local parts don't need to be case-sensitive, so capital letters could in theory be treated as a different mailbox. You could have spaces. If you enclose the local part in quotation marks, you can put a space in there. You can have special characters, so even though a lot of these are not common practice, they do exist, and it presents some sort of problem when it comes time to standardize activities, particularly when you're dealing with systems that don't support internationalized characters.

So in this case we have a fully-internationalized e-mail address with Cyrillic characters on the left side of the domain name. What you cannot do is encode this in Punycode and expect that it's just going to work. For the same reasons that I discussed before, the interpretation of whatever is on the left side of the @ symbol is totally up to the destination mail server, and there is no standard that says, "This is the way it's supposed to work for e-mail addresses."

So this presents some serious adoption problems, because if you have an e-mail system that supports internationalized e-mail addresses and you want to send a message to somebody who doesn't, there is no clear way that you can actually do that. Each mail system is required to come up with its own method of doing this.

One potential way to solve this problem is to issue separate e-mail aliases that are written using Latin characters to users who also have an internationalized e-mail address. But that doesn't help somebody who's trying to reach you, and say, only has your internationalized e-mail address but can't send using whatever their e-mail client may be.

Another example of where there's a lack of coordination, and it's going to potentially present some problems down the road, is with directionality.

So if we deal with a completely Latin e-mail address, it's pretty straightforward to how one should interpret this. In the context of, say, a fully Arabic e-mail address or any right-to-left e-mail address, you start on the right-hand side and read to the left, so in this case the domain name is on the left-hand side. This gets kind of confusing if you have

both Latin characters and Arabic characters. You could parse this and figure out, "Okay, there's a dot on the right-hand side. That one's definitely got to be the domain." Where it gets totally crazy is if you happen to have a local part that contains characters that also are a top-level domain.

In this particular example both shabaka and .cl are valid TLDs, so if you were to actually print this on a piece of paper and hand it to somebody, even if they completely speak Arabic, there's no way to actually figure out where this is supposed to go, whether the domain name is on the left or the domain name is on the right. Even the context in which it's printed doesn't necessarily give that information. So in a right-to-left context it's still not clear you should continue reading right-to-left. Sometimes you might want to jump from the left to the right. And many Arabic users will be used to doing this.

So this doesn't present a problem in terms of taking immediate next steps. It's possible to get an e-mail from point A to point B without having this solved. The problem is, if everybody starts interpreting these problems and solving them differently, we start to have lots of different systems behaving inconsistently.

Just as an example, over the course of the past three days I've had conversations with people who very strongly believe that e-mail addresses should never run from right to left, people who very strongly believe that e-mail addresses should always run in the language or in the context in which they're presented.

Smart people thinking about the same problem can come to completely different conclusions and we need an authoritative source about this, so I'm thrilled that ICANN is working on these resources. I'm also excited for Google to play a part in this as well.

FRANCISCO ARIAS:

Thank you, Brent. Now let's go to Wang Wei. Wang Wei is the CTO of CNNIC and co-secretary of the Chinese Domain Name Consortium. He's now serving as the chair of the CGP and member of the RSSAC caucus.

Wang Wei?

WANG WEI:

I'd like to introduce the application of the IEA (International E-mail Address) so far. Since the new gTLD we already have many internationalized domain names in TLD. And of course I expect there are going to be more applicants for the internationalized e-mail address because I think the EIA, I think people want IEA not only because for the e-mail application, but they probably want user localized e-mail address to apply for some ideal account for the social network. [inaudible] scenario, the localized ID means a lot. Its localized ID is welcomed by the end-user.

Anyway, we already have some IETF RFC listed on the slide. The RFC [inaudible] Framework and the SMTBN POP Extension, IMAP extension, and the e-mail header and [inaudible] or something like that. And I'd like to thank the support from Google, and the engineers from Microsoft, Opera, [SAN], and [AT&T].

And thanks for the [inaudible] who support EAI in July, and of course CNIC has sponsored this project. And if you are interested in the new [inaudible] version, you can download this version from this homepage. And thanks for Gmail, who support the EAI in August.

In China in 2012 we cooperated with the biggest e-mail software vendor in China, and its name [inaudible] who supports the EAI. And if every e-mail service provider can support e-mail address, I think we will have a better situation than what is now, but we still have some problems to solve.

The first problem is that there are some servers have support EAI, and the EAI mail can be sent and can be transported between these servers, but the major client software such as Outlook still doesn't support the EAI. I mean, the end-user will not choose to use that without client support. And just as I said, as I mentioned, maybe not only the e-mail client is supposed to support EAI, in the future some instant messages service or instant messaging client should be supporting the EAI as well.

The second problem is about the e-mail service registration. Some e-mail service provider has support for EAI, but there are still no free public registries now EAI. The end-user cannot register an internationalized e-mail address on the service.

And I just talked with our Google friend. I think maybe Google already has the plan when or how to open this registration service for the end user. So there's still lot to do.

CNNIC is holding APEC. CNNIC has initiated a project called Development of Multi-language E-mail Address Technology in APEC, and

we got the support from USA, from Russia, from Korea and from Malaysia. And this year end of October we will host the APEC EAI workshop in Beijing. If you are interested in this project, come to Beijing and join us.

And contact always Dr. Yao Jiankang who is author of the SMTP extension. This is his e-mail address. Thanks.

FRANCISCO ARIAS:

Thank you, Wang Wei. Now we have David Crocker. David Crocker is principal with Brandenburg InternetWorking. He has authored more than 50 IETF RFCs and served as an area director for the IETF. Over the past four years, he contributed to early Internet research, product development, and [inaudible].

Dave has created and operated two national e-mail services, designed two others, and was CEO of a community non-profit ISP. David is recipient of the 2004 IEEE Internet award for his work on e-mail.

In the late '90s David was appointed by Jon Postel to the pre-ICANN international Ad Hoc committee and [aided] its report. The committee coined the term gTLD, created the concept of registry-registrar split and the framework for the Dispute Resolution Policy.

David?

DAVID CROCKER:

Thank you. Good afternoon. I think they schedule a session like this which concerns a topic that has had some problems so that you can

have the last session of the day want you to get to the happy hour afterwards that much more.

My first full-time job was in the computer science department down the road at UCLA in the early ARPANET, the predecessor to the Internet. I was, at that point, a college dropout and I was doing user support and documentation. And the best document development environment that was available was up in the San Francisco Bay area at the Engelbart research center – that's the group that invented the mouse, and so I taught the secretaries to use that.

And we would print documents off on the high-speed line printer that was next door in the computer center. In order to connect them, we had to take the document from SRI and ship it down to Information Sciences Institute, the system they had there. That's in Marina del Ray, also down the road in a different direction.

The secretaries complained to me that this was only working about half the time. And I was at a loss to understand this, so I asked my office mate who was finishing his doctorate – a fellow named Jon Postel – and John said, "Well, this is interesting. Let's try to figure out what might be failing."

We went through a sequence of accumulative probabilities. We considered how reliable our machine was, how reliable the computer center machine was, how reliable SRI was, and ISI, and the network itself. And we made some fairly rough guesses in each case, and astonishingly the accumulative probability came out at .5, which is exactly what the secretaries were reporting.

This was my introduction to critical path analysis. That is a factor that we often ignore when we design changes for the Internet infrastructure.

The discussions you've had so far have been essentially in the deep details of IDN and EAI, and so in typical Internet fashion, I'm going to go from the bottom up. They've been doing the bottom. I'm going to go up. And I'm going to go high enough up that my comments are not directly about IDN or EAI, but about getting changes to any infrastructure.

The title says DNS. The examples refer to the DNS. In fact, I thought it was ironic that there was the reference earlier in one of the slides to making this topic – the current topic – elevated in awareness to the same level as IPv6 or DNSSEC.

For those who don't know, work on DNSSEC began in 1990. We are 24 years later and have what I think should be called reasonably poor adoption. IPv6 also began in 1990 and has also some problems in its adoption. And so raising things to that level I think mostly raises the challenge of gaining widespread changes to the existing global infrastructure. And speaking of changing infrastructure... there we go.

This slide seeks to make clear that when we want to change an infrastructure we have quite a lot indeed to change. We have the client and the server applications, of course. The user might well need changing. The examples that we're given in this slide of what users might have to type is an example of users needing educating.

The folks who administer those systems need to change their procedures, but also their administrative tools need to be changed. We

had the example of needing to be able to register EAI addresses and having that be a barrier.

And the administrative tools for changes to the DNS infrastructure is often a major barrier, we've run into that. Any time people wanted to put in something unusual for DKIM, even using txt record, just being able to have an interface that allows users to specify that, never mind an interface that actually makes that easy. And this is true for every single change.

And so for years we've had the core of the DNS experts community saying that there is no trouble supporting new RRs, new Resource Records. But in fact the support surrounding the mechanisms that we need for RRs have often been the barrier. Underneath these applications, then we have the entire DNS itself. And for many of these changes, such as IDN, we need to change every single resolver and caching resolver, and every single server that's in the sequence. And that's just the beginning of what makes this fun.

More serious is the fact is that all of these different systems are under independent control of different people. And we have to recruit each one of those people, each one of those groups, to make these changes. And if any one of them doesn't, the system does not work.

You'll notice that I have down users and operations and developers and servers. I don't happen to have designers, because I took it as a given that the challenge here is gaining adoption for the existing specifications.

But the larger challenge in the world that I tend to work in in places like the IETF is worrying about how to design things that minimize the number of places that need to do the adoption.

There are times you have no choice. You must make it change to the deep infrastructure. And for something like IDN we are pretty much in that place, but there are times when we do have choices, and this slide tries to make clear if you can avoid people in the critical path, you need to do that.

Given that we have all of these people in the critical path, then the question is: who are they? Who are each of the adopters? We've had examples given throughout the other talks of browser modules and servers and administration who are each of the actors that need to make changes.

But then comes the hard part. What's the motivation for making the change? Each of these places will not make a change just out of the goodness of their heart. We have open-source well-intentioned activities in the Internet. We rely on that. But that's not enough for the kind of complexity that I had in the previous slide.

And so while the idealistic efforts have served a critical role and will continue to, that's not enough to get a global change to a critical infrastructure service. We need business cases that motivate people to make the changes.

Sometimes the business case is, "My users are beating me up." And in fact that's one of the best business cases. Businesses are very responsive to users who might go away. Even in free services where the

users don't supply money, the users represent money – the fact that the provider makes the money somewhere else. If users are sufficiently irritated that they leave, then that's loss of revenue.

What we have tended not to have tended not to have the case of IDN or EAI or many of the other infrastructure changes is developing that business case sufficiently that it motivates the business to make the changes. That's the piece that we need to focus on. And it's a marketing kind of work; it's not a technical kind of work. And that's all.

FRANCISCO ARIAS:

Thank you Dave. Well, after this presentation, the work ahead seems daunting, I guess because it is. The analogy to the work on IPv6 and DNSSEC and the difficulties there I think it's a good one. I think it gives a perspective of the timeline that we had ahead of us in solving these kinds of issues so we can scale.

So, with this I would like to go to the discussion part of this panel. I have a few questions here, but anyone in the public, please feel free to jump to the microphone and fire.

DON HOLLANDER:

I think I sat in the right place. My name is Don Hollander and I will be from the APTLD in this discussion. I have many questions and comments from your presentation. I have many questions and comments from your presentation.

First of all, I agree with Dave Crocker that if you're looking to elevate this to the achievement pace of the IPv6 or DNSSEC that's very worrying

indeed, because it's taken many, many, many years to get to a very, very low acceptance space in the DNSSEC and the IPv6, so I'd hope that Universal Acceptance goes a bit faster there.

You asked for measures, how to measure, how to identify progress. I do point you to the EURid-UNESCO report that they produce each year on IDN achievements. They've been producing it for a number of years. They have clear metrics on how they measure acceptance. It remains poor, but there is an existing methodology, and metrics. And there's history there. As the EURid-UNESCO report.

I quite like that you're keen for a sense of urgency from ICANN. I would suggest that ICANN try to take the topic to Silicon Valley, given that ICANN has an office here. It would have been nice to have a meeting in LA during the ICANN meeting – didn't happen. But perhaps your group can identify an opportunity in the Silicon Valley working with Google. I know that Google has a small community of other people focusing on IDNs already, so take it to that community and then share the results with the rest of the community and perhaps your remote participation – that is very good.

In New Zealand we have a saying – you call it, "Getting your house in order." We call it, "Eating your own dog food." And you've indicated that ICANN is going to require all new software solutions to be IDN compliant. When do you expect to have IDN compliance in your e-mail systems, so that somebody with an IDN address can subscribe to ICANN's newsletters and so forth?

I think the idea of creating a clearinghouse for bug reports is marvelous, and I would like to send balloons and rockets up with that idea, because there needs to be some place where people can report that – that clearinghouse – and then you can pass that on to users. That's marvelous, so thank you very much.

Just two other questions. Jothan, we heard that IANA updates its CSV file four times a day. How often does Mozilla load that up? So that's another question.

And Brent, I was a little confused between you and [inaudible] presentations, because you're saying there is no standards on how to interpret an e-mail address, and we heard from CNNIC that there are these six IETF standards. I would like just a little clarity there very much. So those are my questions. Thank you very much.

FRANCISCO ARIAS:

Thank you, Don. I got three questions for me, but would you like to quickly answer your question?

JOTHAN FRAKES:

I have an easy answer, so let me get that out of the way. Don, thank you for asking. The Mozilla Public Suffix List is updated as frequently as every week, or two weeks, or monthly, as the volunteers have an opportunity to go and do it. We certainly have expedited it in the time since. However, that's about the pace it is.

DON HOLLANDER: I thought that you said it was an automated process?

JOTHAN FRAKES: No. We have automation to generate the section of the new TLD section, and then we basically add that in. We have a rather complicated process because we use a change management system so that we can roll them back, because it has the opportunity to affect quite a large number of things, so we are very careful. Yeah.

[DON HOLLANDER]: And the other question was, how do I reconcile the list of RFCs with the lack of standards, so I'll just clarify what I meant. There is a lack of standards regarding what's allowed in the local part of the e-mail address. In a left-to-right e-mail address, it's to the left of the @ symbol. That's because that string or that substring is locally interpreted by whatever mail server receives the message.

Gmail does not care what you put on the left-hand side. Only the destination mail server cares. Gmail meaning if you're sending from Gmail, the client I'm using doesn't care.

And then there were two particular examples of where we could use some standards. One is the bidirectional e-mail addresses, and then the other is how do deal with Legacy systems that don't support internationalized characters. So yeah, there are plenty of standards defining how e-mail works. It's just this one component could use a little bit more rigor.

UNIDENTIFIED MALE: [inaudible].

UNIDENTIFIED MALE: The question was, “Is that an opportunity for the IETF?”

JOTHAN FRAKES: I think it’s potentially an opportunity for the IETF, although I don't think that it needs to be constrained to one organization if there are other opportunities to collaborate. Really just being coordinated and aligned is more important than the branding from my perspective.

FRANCISCO ARIAS: I think Dave wants to chime in.

DAVID CROCKER: So this is a comment about the Suffix List, and it doesn't offer any immediate relief, but there is an effort that is trying to get started in the IETF to find an alternative method of having the DNS itself contain the information that is equivalent to the Suffix List. And there is some active discussions going on, the effort itself is not moving very quickly. So if there is progress made, it won't be any time soon, but it's worth you at least knowing about it.

UNIDENTIFIED MALE: I think it’s just...

FRANCISCO ARIAS: Yeah, [inaudible] [a mailing list]. I got three questions to ICANN, and let me see if I don't miss anything here. Regarding the EURid report, we have looked at it. My recollection of it is that it's about the penetration of IDNs in part, and where those IDNs are supporting applications, which is what this is about.

Regarding the comment about the conference, point taken. What we have thought so far – and please let me know if you think differently – is that instead of trying to bring the software vendors and service providers to one place where they may not have the business case or the interest to come is rather we go there where the events they have and give the message there. We think that's probably a better use of resources.

Second, regarding the support for internationalized e-mail address, yes, we have it in the plan. Do you already have that in your systems? Just curiosity.

UNIDENTIFIED MALE: I don't know. [inaudible] we have Gmail for assistance, so I suspect we do.

FRANCISCO ARIAS: Very good. I usually get the answer "no" when I ask this question, but you beat me this time.

[Winnie]

[WINNIE]: We have a few remote comments and questions. I'll read three and then queue up again to give others the turn. The first one's from Edmon: "It is very encouraging to see the progress on this important project. I look forward to this topic making progress at the ICANN Board level, as well to make sure that the [GAC] report they have received early this year could be formally adopted. I believe it should add even more concrete measures and substance for ICANN on this issue which speaks directly to the issue of consumer trust. In terms of the priorities as listed by this roadmap, I have mentioned this previously as well. I urge ICANN to focus on the Universal Acceptance of IDN, including IDN in e-mails, [and] that I am confident will automatically address the issue of ASCII TLD acceptance already.

"By focusing on IDN, however, we have the opportunity to take it to a level that it would enjoy community-wide support, bringing together the ccNSO, ALAC, GAC, even the RIRs to join the cause. Ultimately we are clear that this is not an issue that ICANN nor the immediate ICANN community alone can solve. But we, if acting in a united front across the community, can proactively drive the acceptance of IDN."

FRANCISCO ARIAS: Thank you. I am not sure there are...

UNIDENTIFIED MALE: It sounded like a statement.

FRANCISCO ARIAS: I didn't here exactly a question there, but thank you.



UNIDENTIFIED MALE: May I comment?

FRANCISCO ARIAS: Sure.

UNIDENTIFIED MALE: So what Edmon said – Edmon Chung has done a great deal to advance the cause of IDN, and he should be recognized for that. One thing that I do think needs to happen, though, is the parallel of the ASCII as well as the IDN, because the frequency and volume of each are different. So we can't pause the efforts on the ASCII while the IDN things are being addressed.

In the case of the Public Suffix List, because of the pace, ICANN has been very efficient about contracting and making sure that the TLDs are coming in. If we were to focus purely on IDN and not be resolving those along the way that those may not get caught up. And so I agree with him, and yet I also say there's a nuance to that, that we do need to keep the pace going.

FRANCISCO ARIAS: Please go to the microphone.

[JAN]: Hi, I'm [Jan]. Actually I'm other co-chair of [JAG]. So I just want to add a little bit on what Edmon is trying to say. His point, because he's always so talkative, so probably sometimes he lost focus.

So his first point is he would like the ICANN Board to put a resolution on the [JAG] report, because that has been submitted to the ICANN board since early this year.

And also probably his second point he appreciates the effort and also the progress we already made here, it's a much more active project right now, Universal Acceptance in ICANN.

And also like it's been said, the IDN probably should be put more into priority at that point. Thanks.

FRANCISCO ARIAS: Thank you. [Winnie]?

[WINNIE]: The second comment's also from Edmon, and it was raised during Wang Wei's presentation: "Self-assessment should be a good idea, especially with some tools available. I remember there are web-based MTA, DNS, etc. tools that provide a pretty good information for self-assessing how well those services and configurations are functioning. Something for IDN Universal Acceptance would be great. Maybe we can build on those older tools too."

FRANCISCO ARIAS: Okay, thank you.

[WINNIE]: Third one. A question from Mark [Sfinkerak]: “China has been a major driver for EAI. [Corpmail] delivered the support very early, and CNNIC played a leadership role in the creation of the RFCs. Wang Wei, you have mentioned that EAI needs a push. Can you please tell us your feeling whether government mandates or regulations requiring Chinese script EAI support would be considered in order to create such a push?”

WANG WEI: Well, I think just as Dave Crocker mentioned that the whole ecosystem needs support from the majority. Maybe not all, but the majority support. I don't like the word "leader." The leader means you're [alone]. Even if we developed – we worked together with [Corpmail] and published a software version in 2012, but the EAI application cannot be [done] based on few software vendor and few software [service], I think. That's why I especially want to give thanks to Gmail.

I think with Gmail support it makes it now we are not alone. The team have two members. And I think the governments, the regulation – yeah, there needs some regulation or suggestions from the government, but I think without the community support the government's regulation is only a regulation on a website or on a document.

When we talked about community support, actually I have ideas or a suggestion for ICANN. Because I know the ICANN has many different types of members – the registry and the registrars and stuff like that –

and I know many registrars, they don't make money on DNS. But based on DNS they provide many other services just like the website service and e-mail services.

It is possible that supporting EAI could be a specification, a certificate, a requirement for these registrars.

FRANCISCO ARIAS:

So on that regard, we are starting with asking the registries and registrars to support these technologies. So that's the first step regarding going to the next level. That's a more complicated question to answer, which I don't have an answer right now. Dave, would you like to say something about this?

DAVID CROCKER:

A brief comment about government requirements to adopt technologies like this. It's an appealing thought. It's one that I've heard cited for this topic a number of times in the last two days. But I'll remind people that there were government regulations to adopt the OSI protocols, and government regulations to adopt the X.400 e-mail standards, and didn't those work very well.

So the Internet succeeded in spite of not having those supports, and in fact those supports being martialled against the Internet protocols. So the formal regulations by governments and other organizations might be helpful, but only as part of the more organized and industry-wide. And the specific idea that you don't want anyone standing alone – you want a community working on it – I think is exactly the right model.

UNIDENTIFIED MALE: And sometimes we are proud when we say Internet is a market-driven system or end-user driven system. But we can't forget it when we have problems, and then we say, "Okay, we need the support from the government. We need the regulation from the government."

FRANCISCO ARIAS: Thank you. So I just wanted to point out regarding Mark who was making the question in the chat. Unfortunately, he could not come to the meeting, but he is one of our contacts at Microsoft. I was exchanging e-mails with him yesterday when he told me that he is starting a new role in Microsoft as the leader in regards to IDNs and internationalized e-mail. So there are hopefully interesting times coming from.

Tony?

TONY HARRIS: Good afternoon. My name is Tony Harris. I am a member of the ISPCP, the Internet Service Providers' Constituency. One of my functions in the constituency is coordinating outreach on ICANN issues that affect the members of our industry. Two weeks before this meeting we conducted two outreach events – one in Guadalajara, Mexico, at the [WSIS] event; and the other at the NANOG event in Maryland, which was last week.

At both events we made presentations on the issue of Universal Acceptance. We consider this something very critical to our activity as

ISPs, basically because we're the people that are going to get the complaints when they don't work.

I'm not too sure I didn't hear about the clearing house which was mentioned, I was a few minutes late to this meeting, but just in case my idea on something you could do, [inaudible] something that happened yesterday.

We had our ISP meeting. I invited specifically a new registry – I won't say who it was – who I knew was having problems with Universal Acceptance, to come and present what was happening to them to our constituency.

Amazingly, this person put up on the screen an array of instances where users had attempted to interact with home banking, with two very large banks, or with platforms that had to do with electronic commerce, or let's say the purchase of services and so forth.

The answers usually came back, "You are not using a valid e-mail address when filling out your details. Use a .com, .net, or .org." I mean, this is brutal. When I see that, we should all be running bringing the [firemen] in. This is something that has to be fixed.

And basically I think it's great that we're getting this group. You obviously know very well what has to be done, and going to the right events where you have all the industry involved with software development.

Don't forget the developers [inaudible] for cell phones. Those are thousands of people. I think those need to be addressed too.

All that's going to be very useful, but we have complaints coming in every day. Are those going to some help desk which will resolve them? Are we producing trouble tickets on them and keeping those trouble tickets open until they're resolved? Because if we're not, this is a wake-up call.

And if ICANN is not chartered to have, let's say, such an animal within their organization as a help desk, well, let's subcontract it. I mean, I think people in the industry – I can't speak – well, I'm going to be a registry tomorrow, I'm signing my agreement. But I think new registries particularly would be interested even in contributing to this if it will solve the problem.

But if we're going to hear about events of rejection of new TLDs, and say “this is interesting, and this may have happened because of this or that,” but we do not go to these banks or these e-commerce platform people and ask them to solve this, I mean, what's the point?

Getting to your question, Dave, about the financial benefit, I don't think it's to the benefit of any home-banking service or e-commerce provider of any sort of service or product to have people come in and not be able to complete their transaction. There is a business case there, and if it's presented to them I'm sure they would pay attention, which brings me to the last point.

Presenting this to a home bank, such as the one we saw, Bank of America, if a new registry calls them and talks to their CTO about this, he's going to blow them away. They have to hear from an entity which sounds like an entity they should respect – something that ICANN sets

up or which really can rapidly convince them that this is serious and that they're talking to knowledgeable people, and they should pay attention to this. These are just some suggestions.

And finally I would say it's great to suggest things and walk away from the microphone. Well, we're here to help. When I say, "We're here to help," I speak at least for all the ISPs in Latin America who I'm here to represent on one hand, and whatever resources we can deploy or contribute, as we have done with these two outreach events, we're here to help. This should be solved and – well, I'm just putting this on the table. It's not a question by the way.

FRANCISCO ARIAS: Thank you, Tony. Thomas?

THOMAS RICKERT: My name is Thomas Rickert and I'm representing ECO, which is an Internet industry association based in Germany with roughly 800 members from more than 60 countries, including registries, registrars and ISPs. And I would like to speak to ICANN's role in Universal Acceptance.

When ICANN is approached regarding this topic, some people in ICANN would respond that ICANN can't do too much about it because they're not the software developers. They are not the ones that are developing the OS.

But looking at the new gTLD program and looking at ICANN's bylaws, I think requires ICANN to think about this a little bit differently because

this new gTLD program has been rolled out to provide for competition and consumer choice. This program is going to be evaluated.

We know of numerous instances where not only IDN new gTLDs but also ASCII new gTLDs are failing. They can't be used by users. And I'm waiting for the first incident to be delivered through every user through the mass media which states, "The new gTLDs are great. You can register them, but whether or not they're going to work, there's a lot of uncertainty with that."

So I think that ICANN needs to not only talk about the sense of urgency, but operationalize that. Looking at what ICANN has done we see the [JAG] report which still hasn't been adopted by the board. The [JAG] report had some concrete measures, in my view the roadmap that was published for public comment didn't have any of these concrete measures, but just a way forward. And talking about a methodology on how to measure success I think would be the second step after concrete steps have been taken.

So I know Francisco that you're working very actively on it, and this is certainly not to criticize you or your team, but it's a call-to-action to the ICANN Board to allocate appropriate resources to this.

I think ICANN, particularly in the light of the assessment or the evaluation of the new gTLD program, ICANN needs to build up a track record of concrete actions that have been taken. ICANN needs to be able to list speaking opportunities where they spoke to software developers or asked developers, browser manufacturers to ask them to do something.

I can report about one case where a new gTLD operator has brought to my attention that they've approached ICANN and asked for help. We know that, let's say, five e-commerce sites aren't accepting our e-mail addresses as valid e-mail addresses. And the response was, "Well, you're better placed to make contact with them and get that resolved." I think that's not enough for ICANN to do.

As an organization, we are trying to act as a clearinghouse. So whenever we are notified by member or non-members we would send letters to the respective companies to ask them to change their system so that they can provide for Universal Acceptance.

Discussions like this are good, but everybody in this room knows about the issue. We need to reach out to those that don't know. And the more anchors we can find from which to bring this to their attention, the better. And again, I think ICANN at the next meeting I think we should be taking about what plans we have, but I'd like to hear a report as to what exact measures ICANN has already taken.

FRANCISCO ARIAS:

Thank you, Thomas. Just a brief comment on that, perhaps you missed the first part when I presented about what we have been doing. We started defining the objective of the initiative in cooperation with the community. Now that we have that now defined, the next step in which we are working now is [inaudible] a communications plan to reach out to the [seven] stakeholders we've identified in the roadmap. So that's the next step, and hopefully we can report more on that in the coming meetings. Thank you.

WERNER STAUB:

Werner Staub from CORE. First of all, I must say that I was really impressed about the quality of the presentations here and the content. It really deserves to be spread. But can we rightly say that, in [inaudible] – and it stands in great contrast to what we've been able to do in terms of the deservedly empty roadmap that we have presented, because we haven't figured out anything yet, so basically the roadmap that's been published about the Universal Acceptance is, for all practical purposes, empty. So we should update that with ideas as they've been expressed in here.

And I might add some of them. One of them is specifically related to what Jothan has explained about the Public Suffix Listing and the Dbound working group of IETF. I think this is a development that ICANN should really – I don't know the word. It should make the most of it possible by putting enormous energy into it, and they can because it has IANA as one of the resources, and it has the ability to get the registries to contribute.

I give you an example. Rather than just making it a Public Suffix Listing, which is important, we can make it a resource that people will use just because it's good and is useful. Such as useful against phishing, many people have a good reason to do that.

The example I can give, for instance, just recently they .bank TLD went forward, and .bank has to have a very high security policy. If a resource like the future version of the Public Suffix Listing could have a computer readable policy statement – I say computer-readable. When I say

computer readable I mean this is very old technology. I was going to pull out a coin and show it to you.

The idea of making something intended for humans at the same time interpretable by machines, is about as old as civilization. The machines will be used to weigh coins. That was a machine. It was very powerful. At the same time you could actually see, "Oh, that's a coin." The emperor's image is on it. So both could actually recognize it.

Domain names have this capability, and this is extremely powerful. So people could say, "Oh, it's a .bank," and kind of learn how this should be used. But none of us know, if we use a dollar note we don't exactly know what are the properties that the dollar note is supposed to have so that the counting machine at the bank would recognize it and accept it. It is not necessary, just so long as it's consistent. So I don't have to know all the things as a user if there is a computer readable method to pick it up.

So the DNSSEC only, for, instance could be such a policy. So if a .bank domain is seen somewhere by any machine, that machine would have a resource where it could, "What's the policy for .bank? Oh, all the domain names must have the DNSSEC," that's relatively easy to implement. It would help about getting actual deployment of the DNSSEC because people could say, "This is really required." And if it's not there, it should be rejected. It should be an alert message on the screen.

You could just do the same thing for TLD-wide e-mail security specifications. You could say there are certain things that in this TLD

people may do with e-mail, because all the registrars adhere to that policy. If they don't, it's a rogue - reject it. Because in a .bank they could easily commit to something, such as no hidden URLs behind, don't click here, that kind of stuff. No hidden URLs behind links. This should be apparent. That's easy to police. Machines can do it.

Or you could ask for e-mail validation. Again, this would be, people like to validate e-mail. If they could basically just tap the resource depending on that list, where registries themselves put in the recommended regular expression by which in this TLD e-mails might be evaluated. That is easy to pull. It would help.

Now about IDN e-mail. Actually, IDN e-mail could be added to that thing as a mechanism. IDN e-mail. I've recently seen that Mozilla has tried to make a step forward, and it was a very bad misstep. So they tried to introduce – it deserves to exist – an intermediary solution. You can just type something with the IDN extension. You actually send it along rather than saying, “It’s a mistake.” But what the implementation of that that I've seen with Mozilla did, it put the part before the @ sign in Unicode, assuming that the server had been patched. Hardly any server has been patched.

So basically the first reaction that a person actually trying to do was to get an error message back from the server, and the next time that person is going to try is probably ten years later. If he got an error message, they're not going to try again.

So if you could just get a best practice statements for an intermediary level of introduction, there would be client software only to say that,

"Yes, do accept what comes after the @ sign and translate it into Unicode." You might, in doing so, highlight the TLD, which will resolve the problem [and] we don't know which case to read. It could be highlighted. A client could easily do that.

And at the same time, leave the part in front of the @ sign alone, because another solution, people always have these labels for e-e-mails and it's actually not really needed. Most people have a label explaining all those things – doctor something, and they have all kinds of explanations that are in there, including the name in the respective scripts, so it's not really so critical.

We could just go forward with that just working on clients, and then we would not be dependent for going forward on updating all the e-mail servers in the entire world.

FRANCISCO ARIAS: Thank you, Werner. We're running out of time and I'm going to close the queue. Do you have quick responses here?

UNIDENTIFIED MALE: I had a comment about the Mozilla feedback that you had, and I was wondering, was that in Thunderbird?

UNIDENTIFIED MALE: [inaudible].



UNIDENTIFIED MALE: So that was in Thunderbird. So I think what you are describing where users were placing Unicode to the left of the @ symbol and hitting send without having any kind of feedback from the application is less of a reflection of Mozilla thunder at mail client, and more of just the problem that I was pointing out in the presentation, which is there is no clear way to handle the situation where you put in a Unicode local part and the server you're sending it to doesn't respect that or doesn't accept that. I think that's something that can be solved, but it needs coordination.

FRANCISCO ARIAS: Okay. [Winnie], one more question.

[WINNIE]: It's a comment from Edmon: "In response to Dave Crocker on raising the level of Universal Acceptance to that of IPv6 and DNSSEC, that is directed much more at ICANN's Board and ICANN's executives' willingness to speak of, advocate, and push for Universal Acceptance. For example, to mention it in public speeches, in meetings with officials, in its own supplier requirements, etc. Hopefully we would have more accelerated adoption."

UNIDENTIFIED MALE: So I want to stress that I think doing those things is fine – everything helps – but that we have very extensive experience with efforts like that for the examples that I gave, and they are not anywhere near enough to make a big difference. As an adjunct, as an ancillary activity, that's fine.

But none of the those activities – and frankly in my opinion nothing by ICANN directly is going to make the change.

The change is going to be the implementers, and the operators, and the administrators changing. And the most groups like ICANN or ECO or any other group can do is create an environment that encourages the information, which is what the presentation's comments are intended for, but also create a discussion about the incentives.

FRANCISCO ARIAS: Thank you. Marc?

MARC BLANCHET: I just wanted to make sure that people understand. There were many times that people were talking about Dbound. So Dbound is not a working group at the moment in the IETF. I co-chaired the [inaudible]. There was a single session discussing it. There were a lot of interest. There have been a few additional meetings of some group of people about it, but not yet a working group, so we should not say...

And I guess they were probably looking for a few people that would drive the work. I'm not saying it's died, but I for one promised to write a problem statement and I didn't. So I encourage people to actually go to the next IETF which is on Honolulu, which is a pretty nice place to go, and [inaudible] or contribute for the Dbound work, which is I think an important part of this puzzle. But right now it's not a working group.

FRANCISCO ARIAS: Thank you, Marc. One of the persons that is involved in that discussion is [Louis] my team, the Technical Service Team at ICANN.

With this, are there any quick closing comments from any of the panelists?

WANG WEI: I think if we want to promote the Universal Acceptance, ICANN itself should do something. We have a consensus, but we need to change the consensus into a strong requirement. Maybe not a compulsory requirement, but a strong requirement.

In 2012, I was often asked, "When will CNNIC deploy the DNSSEC?" Or, "Does your DNS server support IPv6?" But I was never asked, "Does your e-mail server support EAI?" So if we are the member of ICANN as registries and registrars, why don't we support our own standards? That should be a requirement.

UNIDENTIFIED MALE: So I have just one statement. I'm delighted to see the resourcing put behind this. I saw [Cyrus] briefly was in the room, and I started to see occasionally Board members taking a look at this.

I parrot probably what Don said about trying to elevate this perhaps to a better pace than IPv6 or DNSSEC [inaudible] because of its important. And unfortunately it's an unpopular – we see we're at the scheduled time. The attendance unfortunately is lost to other things, but this is a very important thing, and it's nice to see ICANN putting effort and

energy with [Louis], yourself, Francisco, and [Cyrus], toward Universal Acceptance.

It will be delightful to see ICANN participation in developer conferences and getting the word out so that people can start to work on the things that they don't know that they don't know are broken, similar to the collision lists.

FRANCISCO ARIAS: Okay, anyone else?

UNIDENTIFIED MALE: No.

FRANCISCO ARIAS: Well, with this, thank you everyone. With this, we close the session.

[END OF TRANSCRIPTION]