ICANN75 | AGM – GAC Capacity Building and Outreach Workshop 6: DNS Roles and Responsibilities - Country Code Managers and Generic Name Relations
Sunday, September 18, 2022 – 13:15 to 14:30 KUL

[THIS SESSION IS BEING RECORDED]

GÜLTEN TEPE ÖKSÜZOGLU: Hello and welcome to the ICANN75 GAC Capacity Building and Outreach Session 6 on Sunday, 18th September, 2022 at 05:15 UTC. Please note that this session is being recorded and is being governed by ICANN's Expected Standards of Behavior. During this session, questions or comments submitted in the chat will be read out loud if put in the proper form. If you are remote, please wait until you are called upon and unmute your Zoom microphone. For those of you in the GAC room, please raise your hand in Zoom, and when called upon, unmute your table mic.

For the benefit of our other participants, please state your name for the record and speak at a reasonable pace. You may access all available features for this session in the Zoom toolbar.

. . . . With that, I will hand the floor over to Tracy Hackshaw. Tracy?

TRACY HACKSHAW: Thank you, Gülten, and welcome back to all from our lunch break and from wherever you are in the world, from our one-hour break. As we continue with our capacity building weekend, we have now a very interesting and exciting lineup of activities for you [chuckling]. in terms of what is happening in the community at ICANN.

During the session, GAC attendees will hear from community stakeholders responsible for managing the DNS, including operators of country code and generic name registries, including the geo registries and operators of gTLD registrars. These stakeholders will share their experiences, business practices and how they manage their relationships and their interactions with governments on a wide range of topics. So our first speaker will be Jia-Rong Low from ICANN.org to give us an overview of the Domain Name System, DNS, what it is and what it does. Jia-Rong, you have the floor.

JIA-RONG LOW: Thank you, Tracy. Good afternoon, everybody. My name is Jia-Rong, I am the Vice President for stakeholder and engagement and also the managing director for the ICANN Asia Pacific office. I'm based in Singapore and I look after the Asia Pacific region. So a big welcome to everybody, welcome to the region, welcome to Malaysia.

So I think it's in this capacity, being partly a host for the region, that I'm asked to speak with everybody. But the topic is not something that I am unfamiliar with, even though I think amongst technical experts they would probably call this an oversimplified version. But I think it suits the audience for the GAC members because we have many new GAC members rotating, coming in, so keeping it very simple for everyone so that you have a broad sense of the ecosystem, and later on when we hand it over to Bruce, he will dive much more in depth and that's where you can get into a lot more meat of the discussion. So without further ado, let's go to the next slide, please.

At this basic for fellow GAC members, from a policy maker or regulator's perspective, when you think about Internet, it seems to be a very broad thing. When we think about the internet, we tend to think about things like Spectrum, 3G, 4G, 5G or fiberoptic cables, or on the content side, we tend to think about things like Facebook, Google, or looking ahead, you hear buzzwords like Web3, BlockChain, AI, et cetera. A lot of these in fact are not really Internet, but they are more applications relying on Internet technology. Fundamentally, what is Internet? It is the ability for my device to connect to your device, and an easy way to describe how the Internet works is through DNS resolution. So again, the key word is DNS resolution.

Now, on this slide I will use DNS resolution. Here, this slide says DNS query; we will use that to illustrate what happens for my device to connect to your device, or my device to connect to a server which houses the information that I want to connect to. So here as a basic example, first of all, you open your browser and then you type a domain name. So in this case, we used as an example www.domain.org; so right at the top, you type www.domainname.org. What happens? The DNS, the Domain Name System, would convert that into an IP address to look for the IP address of this domain. Where is the IP address of www.domain.org? So that's what the system does immediately.

And assuming that your ISP server doesn't know it, it will go to step number 1 on the slide. So it doesn't know. So it does know how to go and find the nearest root server. So the root server will say, "Well, I don't know the address for www.domain.org, but I know where is .org. Go and ask .org." So then you go to step 2.

Then what happens is, the .org servers would then be able to say, "Well, I don't know the www.domain.org, but I know where is the domain.org server, go and ask them." So then it goes to step 3, and finally you get connected to www.domain.org is at IP address

192.0.2.0. So it is a series of referrals from root servers to the .org server to the domain.org server that allows you to eventually find the particular website you want to connect to. So fundamentally, this is what the Internet does.

Next slide, please. Thank you. So understanding how the domain name resolution process gives you an appreciation for the decentralized nature of the Internet because there are different ”players” or operators that are involved in making the Internet work. Now, I mentioned the first step that you do; you type www.domain.org in your browser. What happens at the back end is, your Internet service provider, so you see the top left I put an arrow, at that process your internet service provider recursive server is the one that will go to help you to ask the different root servers and .org servers, where is the IP address. So your ISP is already involved right at the start when you want to get connected.

Now, continuing this same process in step 1, you will go to the root server to ask: where is the IP address? The people or organizations that maintain and operate the root servers are called Root Server Operators, so actually here in ICANN you sometimes will meet a few of them, and all of them sit on the RSSAC. ICANN is also a Root Server Operator, so we operate the L-Root or the ICANN managed root server, IMRS. So there's a whole community of root server operators that are in charge of managing and operating the root servers.

Under Step 2 in this slide, for the .org servers, they are managed by the top-level domain registries. So I think within the ICANN space, everyone is very familiar with registry operators and largely, you can categorize them into two categories; one is gTLD operators and ccTLD operators.

Now, the third one is under step 3, domain.org server; so these are second level domains. So in our ecosystem a few players, as part of it, will be the registrars who sell domain names and help to register domain names, so registrants, and to put up the domain name for users, a website you need a hosting provider, and also the website operator could be a company itself, so it could be like for example Facebook or any other company that provides a service.

So within the whole ecosystem, there are different players in charge of different parts, and for policy makers and regulators, you have to understand that there are these different players as part of the ecosystem and all of them either play different roles or have limited roles, and understanding the relationship between the different roles is important to have a holistic view when you want to think about whether making policy or considering regulations.

Now go to the next slide, please. I just thought to add a slide here because I just mentioned the different players. But some of the things we are concerned about within our community would be on security. So when we talk about DNS security in relation to DNS attacks, it happens between the phases. So one example is you type a domain name, like www.domain.org, go from step one to go and ask the root servers, a malicious actor could intersect that query and then give a different response. So the intersection of that query is a DNS attack. Sometimes we call them 'the man in the middle' attacks.

Another kind of attack, and a pointed the arrow from the DNS attack to a server, is a DDoS attack, a denial of service attack. So basically you flood a server with millions of queries to the point where the server cannot manage the queries, and then the server basically can no longer provide service query returns. So then essentially the server shuts down. So you can block an entire top level domain by flooding it with queries using a denial of service attack. So when we again consider the ecosystem and think about security, within our space, DNS security is linked to these areas. So again, I think later Bruce will mention this, within the DNS, we're not talking about content at all.

Now, I will close with this final slide. So again, the ecosystem I had mentioned, the process, and within the resolution process, the different players that are involved, as a very brief overview. And now on the right going clockwise, you will see for each unique identifier there is a different platform that looks after the unique identifiers. Now, how I typically talk about this, n the Internet unique identifiers there are three key areas.

The first is that all the computers or devices need to speak the same language. In other words, we code them to speak the same language. In other words, they share the same protocol parameters. Every single device that's connected needs to have an address as illustrated in the DNS resolution process. So we need to have an IP address. But humans cannot remember the IP addresses so we remember names instead. So we type the names. So that's domain names. So the unique identifiers are these three key areas, and each area is managed by a different organization/platform, and in the case of domain names it is ICANN. But ICANN only manages gTLD. And the ccTLD space, the policies regarding ccTLDs is managed by your own country's ccTLDs.

And in the space of IP addresses, the distribution of IP addresses and how they're used is largely discussed within the regional Internet registries and for the APAC region, it's APNIC, and there are 5 regional internet registries. Now, for protocol parameters it is the IETF, Internet Engineering Task Force, and the other ones who came up with IP version 6 when the community realized that IP version 4 was insufficient for the 50 billion devices that were already on the Internet right now. So that really wraps up the ecosystem. I hope that was easy to understand, an oversimplified version that would at least give you a flavor of how the internet works and who are the players. Thank you very much.

TRACY HACKSHAW: Any questions? Questions, feedback?

ROBERT HOGGARTH: Maybe we could take questions at the end, Tracy. Do you think?

TRACY HACKSHAW: Alright, so I'll hand it over to Karel now for introductions to…

KAREL DOUGLAS: Okay. Thank you so much, Jia-Rong, that's a fantastic presentation on the ecosystem, domain names, root servers, the TLDs, registries and so forth. And it never gets old because every time somebody asks me how does the Internet work, I think this is a perfect example of that. So I do appreciate that so much. And this point in time, I'm going to move to Chris who has a panel with him who he shall introduce, and we will get into some of the ccTLDs and the gTLDs. So Chris, would you like to introduce your panel?

CHRIS DISSPAIN: Thank you, Karel. Hi everybody, I'm Chris Disspain, I am a former ICANN Board member, a former Chair of the ccNSO,a former ccTLD manager for the .EU ccTLD, former lots of everything and now I find myself back on the ccNSO council again. I'm going to introduce you to the panel who are here to talk about ccTLD and gTLD registries and their interface with governments.

And I'll try and do it in order, so I will start with Bruce at the far end. So Bruce is the Chief Operating Officer for .Au Domain Administration, the Au ccTLD. He was Chief Technology Officer and Chief Strategy Officer of Melbourne IT from '99-2017, and he too has served on the ICANN Board for a lifetime, which in ICANN Board terms is 9 years.

Then next to him is Edmon Chung, who serves as the CEO of .asia. which operates the .asia Sponsored gTLD, which is a not-for-profit organization. .Asia serves as the secretariat of the Asia Pacific Regional Internet Governance Forum, and supports the operations in the recent launch of the community gTLD .kids.

Next to Edmon is Ram Mohan, who is the Chief Strategy Officer at Identity Digital who is the new registry company from the merger of Donuts and Afillias. Ram has been deeply involved in gTLD and ccTLD issues for 22 years and has also spent a lifetime on the ICANN Board as a liaison for the Stability and Ssecurity Advisory Community.

And finally, Roelof Meijer is the CEO of SIDN, which is the registry Netherlands .NL, which has over 6.2 million domains registered, which makes it one of the largest and most successful ccTLDs. So, that's the panel. We're here today to talk about the distinction between the way that ccTLDs interface with governments and the way that gTLDs interface with governments. But before we do that, Bruce is going to take five minutes to just explain the role of registries in the context of the Domain Name System. And then we're going to come to the panel for some thoughts. But over to you first, Bruce.

BRUCE TONKIN: Thank you, Chris. Well, it was worth just identifying the different parts of the DNS system and also ICANN's role in each of those parts. And I will use a general analogy really of phone books. I think most of us are probably of the age that we would have had a large phone book sitting next to our telephone, so I will use that sort of White Pages phone book analogy.

At the top level, we have ICANN that manages the list of top-level names. And there are about 2,000 names at the top level. So clearly, ICANN's role is to manage that list of names. The next slide down then is the Root Server Operators. And effectively, they work through cooperation with ICANN and generally pre-exist ICANN probably by about 15 years or so, so they will have been operating much longer than ICANN itself has, and they operate a globally-distributed set of name servers, essentially computers, and really all they have on their computer is the list of names and where you should go to find out some more information, and in fact it is effectively a list of directories.

The next layer down we have gTLD registries. We have examples of those, they've been around a long time, so .com, .net, .org and in 2012 onwards, there's been another couple of thousand names added, some of the city names like London. Some are names of organizations; .afl is the Australian football league in Australia. And quite a large number of more generic names like .digital, which is operated by Identity Digital.

Now, this layer works by contract, so they have a contract with ICANN and ICANN sets a few sort of high-level policies relating to topics like WHOIS and how names get transferred from one registrar to another, so that's built into the contract and a lot of the discussion at ICANN meetings essentially is relating to roles set in the contract between ICANN and the gTLD registry operator.

The ccTLD registry is a bit different. Again, they predate ICANN by often more than 15 years, and so again, they're much older organizations than ICANN itself. And here the relationship is a cooperative relationship with ICANN, so there isn't generally any contract, and those registries share information with each other and are mostly operating through rules set through their local community in multi-stakeholder processes.

And so far, these are just directories of directories, so it's a list of names followed by the name of the directory where you should find further information.

The last row is the phone books themselves. So if I look at a large ccTLD registry like .au, we have millions of domain names listed in that directory, and within those millions of domain names, they in turn refer to a set of what are called authoritative name servers for each name, and there would be thousands of authoritative name servers that are basically the phone books for the domain names themselves.

So on that bottom row, phone books can be run by multiple organizations, so a few in some cases registrars would have the option generally included as part of your domain registration fee that they will run a name server, a default name server that that registrant operates. If you have purchased your domain name through a telecommunications company or Internet service provider, they also have their default effective phone directory. With hosting companies it's very similar.

DNS service providers is a little different. There are a few companies that specialize in just providing DNS, and when you purchase a domain name, you can choose to specify which DNS operator or which phone book you want your domain name to appear in, and you can choose these DNS service providers. And those phone books essentially contain the IP addresses of the website that might be associated with a domain name or the IP address of the email server that might be associated with the domain name. So that's really what your phone books are.

And what is interesting here is that ICANN doesn't have a role, so ICANN's role is very much at the level of the directory of directories, and the actual web server addresses and the actual [inaudible – 00:24:05] record addresses are managed in a very distributed way through many thousands of organizations around the world that operate those directories.

Next slide. I want to be clear here about where content sits. So the DNS itself doesn't contain any content other than what is in the domain name itself. You can have a short English phrase. I'll just give you an example here of have a nice day.au, but you certainly can't have a book in the DNS and you certainly can't have a video in the DNS.

So the very limited content and the only content is basically a name that's in the DNS, and as discussed, it's basically just a hierarchical set of directories, and you progressively go from one DNS name server and it says, go and ask this other name server, and eventually you get to the phone books, the authoritative domain name servers and they contain the web address and email address information. Generally, the website and mail servers are not operated by the domain name registry. So most domain name registries don't provide any content, they don't look after websites, they don't manage email servers for customers, generally.

And then the last point I'll make here is that often, registry operators get asked to remove a domain name, and I just want to be clear, that doesn't actually remove the content. All you have done is remove a line in the directory. So it's a bit like saying if I remove somebody's phone number and name from a white pages phone directory, I can still ring the person, it's just that they're not in the directory. So content that's on the Internet could be connected to many domain names, and of course, you can reach that content with the IP address directly, just like you can ring someone's phone number directly. They don't need to be in the directory.

Next slide. And the other role in the Domain Name System is resolvers; so these are basically computers that are usually operated by an Internet service provider, and that will send a query to these other DNS servers on behalf of the end user and they keep a copy of the result. So you ask for the very first time a brand new domain name, you ask your ISP, the ISP will find where that domain name is and will keep a copy of that result so that when the next person asks, they can generally give the information from a local store, and they will hold that information for a period of time and then they'll refresh it.

And again, these resolvers don't contain any content but they can implement security measures, so some internet service providers for example will put software into their DNS resolvers to protect their end users. They might block sites that might be hosting malware and things like that. So certainly there's some interception if you'd like that can happen at the level of these DNS resolvers.

And again, there's no ICANN contracts that relate to these DNS resolvers, but some government laws are implemented by Internet service providers using these devices. So I just want to be clear, it's nothing to do with ICANN, but certainly governments might well interact with their local operators of resolvers.

CHRIS DISSPAIN: Thanks, Bruce. So, as you could see from Bruce's first slide, there's a fundamental difference between a ccTLD and a gTLD, well there are several actually. One of them is that gTLDs have contracts with ICANN and exist in the ICANN arena thanks to that contract with ICANN. Whereas the ccTLD is effectively existing within its own sovereign territory or country and it's governed by the laws of that country. Now, the policy is maybe set by a multistakeholder organization in Australia for example; AuDA is a membership-based organization. The government is involved but government doesn't actually control the ccTLD. In some other countries, the ccTLDs are actually controlled, owned, if you will, by the government. But the involvement in ICANN is purely on the basis of voluntary involvement.

So because we believe in balance, we have a very balanced panel. We have two gTLD registries and two ccTLD registries. I'm going to start with you, Ram, if I may. And if you could just talk about the relationship of the gTLD registries or registry with governments when it comes to policy.

RAM MOHAN: Thank you, Chris. Ram here. It's a very different kind of relationship when you are running a gTLD registry. As Chris pointed out, the governing contract is with ICANN as well as with the partners, the registrars who sell the domain names on your behalf. So those are two legal contracts that really govern how the registry is run. But clearly, if you look at where the registry is incorporated, in which country it is incorporated, clearly in that country the registry is going to be conforming to whatever the rules and regulations are in that locale. Beyond that, as a gTLD registry, the government interaction is quite broad.

At Identity Digital, for example, we are engaged with many different governments, especially with folks from there on either the law enforcement side or with CERTs in the governments that work on security issues, so we work closely alongside them, but at the end of the day, with perhaps the exceptions of the jurisdiction where the registry is incorporated, all the other governments are treated on a somewhat level-playing field.

CHRIS DISSPAIN: And Ram, when it comes to policy, then, your government of incorporation is effectively setting the rules that you have as a company, how you account etc.; when it comes to policy under the terms of your contract, that happens here at ICANN.

RAM MOHAN: That's correct. Almost completely here at ICANN, and in particular for gTLD registries, those policies are made in the GNSO.

CHRIS DISSPAIN: Thanks, Ram. So let's go to a ccTLD. Roelof, perhaps you could talk a little bit about your ccTLD, or generally, ccTLDs relationship with governments in respect to policy.

ROELOF MEIJER: I think first of all, in this perspective, our role is a lot simpler than for instance in Ram's case because we have only one government that we report to and one jurisdiction that applies to us. So we have no policy that we apply that we have developed for instance within the ccNSO. An example of ccNSO-developed policy, for example IDNs, we don't have IDNs within .nl. So the way we develop our naming policy, wo who can be a registrant and who can be a registrar, and what conditions are being used, that is something that we develop with our local Internet community.

So with stakeholders, we used to do what we call domain name debates which were public debates where we invited stakeholders but everybody could join, and where we would discuss a certain topic, for instance WHOIS. One of the latest of those domains eventually led to what we now call a tiered WHOIS, so that's a good example, it's a different policy from for instance the policies that apply to Ram's gTLD with regard to the WHOIS.

We have no specific contract with any party that determines that we are the ccTLD, we're the operator for the Dutch ccTLD. So at the moment, as IDN became the registry, our government wasn't really interested in the Internet yet, so this was something that the business sector should kind of deal with, and of course now things are different so we have a very good collaboration with multiple ministries within the Netherlands. We have what we call a confinant with the Ministry of Economic Affairs; for us, that's the most important ministry that we deal with, it's also the ministry that supplies the GAC representatives for the Netherlands. I said we have a confinant, that confinant it's not about, as IDN we allow you to run . Nl, it is about neutral agreements about how we ensure that .nl always works; so it's about, this is continuity, it's not a license to operate .nl.

CHRIS DISSPAIN: Thanks, Roelof. Edmon, you are up next.

EDMON CHUNG: Thank you, Chris. So I want to emphasize what Ram said. So our organization, .Asia, operates .asia obviously, and as the introduction from Chris mentioned, we are also operating .kids. So the basic policies and processes are -- well, policies, especially in terms of .asia and .kids is actually from the GNSO and consensus policies that are created here.

However, we're on the complete other end of the spectrum, and we are quite special in the way that for example when .asia launched back in 2007 and 2008, and through the application process at ICANN, actually we had a very tight relationship and interaction with the GAC. In fact, through the application process, we were hit, if you will, by one of the first early warning from the GAC.

Some of the GAC members at that time raised certain issues about the use of .asia on the Internet, and since then we have through the start-up process actually worked very closely, especially with the Asia Pacific representatives from the GAC to set up additional policies at .asia; and they include reserve names, for example city names, province names, geographical indicators and so on. So as .asia was launching, actually .asia worked very closely with the Asia Pacific GAC representatives to solicit and work with the reserve names, so we have added a number of additional on top of what the GNSO and ICANN required.

But again, I want to emphasize, that is not part of the registry agreement with ICANN but it is something in terms of the responsibility of the registry .asia. And after about 12 years we are reaching out to GAC representatives again to see if there are any updates to those reserve names and geographical indicators and so on. And so I look forward to hearing your response.

At the same time, .kids, this is also one of the very few strings that the GAC was especially keen about; in fact back in the process between 2012 and around 2016, there was a lot of attention paid in terms from the GAC as including, especially from the European Commission and the UK about the nature of .kids. And our interest to operate it as a community TLD to include policies above and then beyond the ICANN agreement. Basically you can think of it as you would have additional requirements for things that happen in a park or in a school. So we want to put policies in place for .kids to also do those types of -- to mitigate against abuse.

So one of the things that, in fact, I would like to start – we are starting to reach out to GAC members about, is if there are for example obscene words that would be put into the additional reserve names list, we look forward to interface with GAC members to look into those types of policies as well. So these couple of registries are very different from the more open gTLDs, but the baseline remains the same, which is governed by the contract with ICANN and the contract with our registrars.

CHRIS DISSPAIN: Thanks, Edmon, that's great. And we will come back down the line in a little while and talk about the way that whether you're a gTLD, or a ccTLD, affects how you might interface with other branches of governments such as law enforcement, but before we do that, Bruce, your take on ccTLD and policy.

BRUCE TONKIN: Thanks, Chris. If we can just forward through a few slides back to where we were. One more. Thanks. So first thing just from a policy point of view is, ccTLDs such as .au attend ICANN but the policy discussions here are really about the rules for adding a country code name to the top level and the rules for removing a country code from the top level. So adding names in recent years has generally been internationalized versions of domain names that reflect the national languages in each country, so for example Singapore .sg, also has the top level name in Tamil and Chinese. And then the other topic is removal of country codes, which is generally when a country ceases to exist, so that’s usually unfortunately related to some form of warfare, and the results after the warfare is some redistribution of the borders of a country. But the operation of each ccTLD is managed through the local community.

Going to the next slide, I can talk a bit about kind of the trends there and certainly what we're seeing in Australia. One of the things that we have noticed in Australia in recent years is a split between the requirements around the operational infrastructure, which is really the DNS part that I spoke about earlier, and the requirements around naming roles.

With respect to the operational infrastructure, the DNS infrastructure, many countries now including Australia are starting to treat the database of domain names and the DNS name servers as critical -- it should say on the slide infrastructure, rather than infraction -- but similar to water, electricity, gas, and telecommunications.

The focus of sort of government discussions in this area is on security, and certainly the three elements of security; how do we protect the confidentiality of information that’s held in the registry, particularly information about private citizens might be contained in the registry. Focus on availability, making sure that the DNS infrastructure operates 100 percent of the time. And focus on integrity, which is making sure that there can be no unauthorized changes to the DNS settings.

And in Australia this conversation tends to happen with our Department of Home Affairs, and there was a recent legislation called The Security of Critical Infrastructure Act, and as a result of that, the .Au name system is now specifically covered under law with requirements to maintain the security of the infrastructure.

The other area is naming policy, and naming policy has a number of elements to it. One is eligibility, who gets a name. So some country code operators have fairly open eligibility. In the case of Australia, the eligibility is restricted to those who have an Australian presence. The next set of rules relate to allocation as in what name can an individual registrant have, and some ccTLDs don't have any specific roles in this area. In .AU we do have rules about what names registrants can have, and specifically, there's certain reserve names under Australian registration that limit what names a registrant can register.

Then there's the aspect of accountability, so governments are generally concerned with how can you hold a registrant to account if the use of the name is being used for some illegal purpose. And a lot of this then gets tied into what information does the registry hold about the registrant, is that information accurate, and how can the government request that information when they're needed.

The next element is transparency, which is really for the general public, how do they know who's responsible for each domain name. And most registries operate a WHOIS service, the information that's available in the service might vary, but that is a topic of sort of local rules.

So generally, the naming policy covering those elements of eligibility, allocation, accountability, and transparency, is developed, certainly in Australia, through multi-stakeholder mechanisms, so those policies where we engage with government, industry, the not for profit sector, the academic sector and the broader civil society. And , we have found that to be a very successful model in Australia and it's been successful in developing the naming policy and increasingly on Internet related topics, we're trying to encourage organizations and governments particularly to use that mechanism when they start looking at other aspects of Internet policy.

And then the final point here is that each ccTLD manager is subject to the local laws that are in their country. And in Australia the ones that most directly relate to us is the critical infrastructure laws and the privacy laws.

CHRIS DISSPAIN: Thanks. Bruce. So let me sum up where we're at and where we got to. It should be clear I hope that the place where gTLD policy is dealt with is in ICANN, and that has a system for doing that, and you guys in the GAC are a part of that, you liaise with the GNSO, et cetera. With respect to the ccTLDs, each country or territory sets its ownpolicy, and those policies vary widely, they vary from the point of view who is entitled to what name, they vary about what information is available and so on. They vary among gTLDs as well, but as Edmon has said, the baseline contract with ICANN and the policy is the minimum. There are some gTLDs who offer more.

When it comes to setting policy, the ccNSO's job is not to set policy for ccTLDs. Wwhat the ccNSO does in effect is sets policy for how ICANN deals with some of the aspects of ccTLDs. Such as how you retire a ccTLD or how you delegate a new one. So there are fundamental differences. And there were also some fairly important differences in the way you do things. If you take a simple example, we're going to talk about law enforcement and relating with other agencies.

If you live in Ruritania and you get a contact from the Ruritanian police or summons from the Ruritanian justice department, you kind of know that that is your police force and your justice department and you know what you need to do. But what do you do if it comes from another government in a different country or different territory and as a gTLD, what do you do about law enforcement generally in respect to policy as opposed to your corporate activities, which is a slightly different area.

So perhaps if we go -- maybe we do it in the same order, Ram, if you could sort of talk about from a gTLD point of view, how do you deal with law enforcement queries, how do you deal with intellectual property queries and so on and so forth, given that it's a vast array of countries and a vast array of possible registrants.

RAM MOHAN: Thank you, Chris. So one of the big challenges is that gTLD registries with their global scope have an expectation from many governments that they be responsive to that government's request, that government's needs. And it's an expectation that is very hard to fulfill. So as a result, the common ground is found in the policies that are set at the GNSO. And those policies are adhered to absolutely by the gTLD registries.

Now, when we get requests from law enforcement agencies from other countries, as I had mentioned earlier, there are really two avenues. One is to guide them through the registry's existing data disclosure requests policy, the data disclosure requests mechanism that exists. And the second is through relationships that are built with contacts on the regulatory channel. But in general, a request from a law enforcement agency from a given government, the fictional Ruritania government, we would receive it, we would view it quite carefully, but then we would end up walking that through, here is our standard policy and does it fit through that policy lens.

CHRIS DISSPAIN: So Roelof, you have the privilege in essence of having a clear line of inquiry from your sovereign territory, from your local police. So I'm guessing I know how you deal with those. If you could address that and also talk about how you might look at requests that come in from overseas law enforcement to you in the Netherlands.

ROELOF MEIJER: Maybe I should start with the latter. So if we get requests from non-Dutch law enforcement agencies, we refer them to Dutch law enforcement agencies. If we get orders from law enforcement agencies from the Netherlands, of course we respond accordingly, and if we get requests, then it becomes a bit difficult because that means that we have to kind of judge if it is something we should do or not.

We do a lot about abuse, so we have let's say a positive stand towards requests from law enforcement, but we are actually at the moment implementing what we call a kind of ethical Board where we create a team that helps us to decide on requests for data for instance where we are not obliged to provide the data but where it is allowed. So it is our decision.

And I think that's a good example of the difference between gTLDs and ccTLDs because I think for Ram, it would be absolutely impossible because he would have requests from so many parts he wouldn't know whether it was a valid request, and first it's very easy because we know the parties. And if we don't know them, they're probably not Dutch law enforcement so we don't respond.

Another good example is maybe the WHOIS; we had WHOIS discussions within our Internet community way before the implementation of the GDPR, and we had a tiered WHOIS as a result of those discussions way before the implementation of the GDPR. And then we became the registry back-end operator for .Amsterdam, which is a gTLD, a Generic Top Level Domain. I think

Amsterdam is one of the oldest brands in the world but it's still a generic gTLD.

So of course we, as a responsible registry, wanted to implement the same WHOIS as we're running for .nl because we run that type of WHOIS on the convictions and meanings of: law enforcement can get access to all WHOIS data but on the basis of a contract that they signed and where they promise that they will use that entrance proportionally and apply the subsidiarity principle.

And ICANN of course obliges us to run an open WHOIS for .amsterdam and it took a long time to configure -- well, in fact the GDPR had to come to convince ICANN that the gTLD, WHOIS obligations were violating European legislation. But I think it was a good example of where as a ccTLD you can very quickly react to your local government's requirements and to the needs for your local Internet community. And for a gTLD, you are in the process of ICANN and the GNSO, and that's a big crowd so it takes a very long time before you get any consensus on the outcome.

CHRIS DISSPAIN: Yes, that makes sense. And Ram, I'm guessing that, whereas for Roelof is quite easy in his territory to set up relationships with trusted notifiers, with people who you know who they are and trust the inquiries when they come, that they are bonafide, than it would be in the gTLD world to set up with trusted notifiers to deal with queries that come in.

It's a tougher challenge, I imagine.

RAM MOHAN: Yes, it’s a far more difficult challenge because the standards of determining who should be a trusted notifier are not even across multiple jurisdictions, it's multiple countries.

CHRIS DISSPAIN: So Edmon, let's go to you and talk about how you interface as a gTLD with law enforcement.

EDMON CHUNG: Yeah, I guess it's not so different from what Ram says. Speaking about .asia, then one interesting thing is that the one time that we worked very closely with law enforcements from Hong Kong where we are incorporated is actually in 2008 when the Olympics happened in Beijing, although one particular event was held in Hong Kong and there was actually a close collaboration to watch over the zone during that period of time.

But in general, it’s very similar to what Ram said but I would add one more dimension to it as gTLD registries. A lot of actions if they are taken, we would defer to the registrar. So therein also lies another additional consideration or perhaps sometimes complexity because the registrar could be in a different jurisdiction. So when something is referred to the registry on a specific domain, we tend to refer to the registrar of record, the sponsoring registrar, and the sponsoring registrar obviously could potentially be in a different jurisdiction and therefore they have to abide by the law enforcement in their jurisdiction.

So what we tend to do also is to kind of relay that request to the registrar. So in short, yes, being in Hong Kong obviously we respond to the Hong Kong law enforcement agencies. Outside of Hong Kong would be somewhat more difficult, but the added complexity is working with the registrars to actually take action on certain domains.

CHRIS DISSPAIN: Thanks, Edmon. And Bruce, if I could run to you lastly on this topic, then we might go to questions, but from the Australian point of view.

BRUCE TONKIN: I think generally for us, our interaction with law enforcement is a little bit like say a numbered plate agency; in that, if you commit a crime with a car and you've got a numbered plate and someone writes down the number plate when you have raided some 711 or something, so basically the law enforcement would come to us and we would provide them information that's not publicly published in the WHOIS.

So we have postal addresses, phone numbers, the data which the domain name was registered et cetera. But I think in Roelof's model, we don't provide law enforcement with direct unfitted access to the registry; they need to submit individual requests and that's generally managed under the Australian Privacy Act. So the law enforcement body would need to explain what the crime is and provide sufficient cause, if you like, and then we would provide the full information that we have in the registry.

The other thing that we get less commonly is requests from law enforcement to take down domain names. And usually, we would only do that if the information that' been provided at the time of registration is false, so effectively, that's given us false information and that's grounds for us to take down the domain name. But we also will take action on certain categories of DNS abuse, particularly phishing sites, so for example a site that's trying to look like a bank or sites that are set up primarily for malware. So a couple of areas where we would take down. But normally, our role is to provide the information on the registration. You mentioned overseas law enforcement; similar to Roelof, our approach would be that we would direct them to the Australian law enforcement agency and then we would respond to Australian law enforcement agency.

CHRIS DISSPAIN: Super, thank you. So I want to take some questions if there are any. But before I do that, Ram, if you want to just make a comment about regulation.

RAM MOHAN: And before I do that, well, I think what Bruce said and what I said before that really gives a nice insight into the differences that CCs have in their policy, they adopt to local legislation and local requirements because like Bruce said, we have contracts with every law enforcement agency about their direct access to WHOIS and those contracts have been approved by our Data Protection authority and we published them on our website so we're being very transparent about that and we do notice some takedowns.

But there is also you could say a more negative difference for a registry that runs a ccTLD as compared to a gTLD, and that is that of course we get local regulation. So we are the largest market party in the Netherlands within the domain sector, we have a market share of about 17 percent, .com has 20 something, but since NIS1, we are judged to be a provider of an essential service and I think we are, so it's very logical that we have a certain regulation that is aiming at ensuring that the registry always works and .nl is always available.

But .com has no such regulation for the domain names at Dutch companies and it is mainly the big companies that use dot com domains. So there is a difference and there is a risk that creates an unlevel playing field, where CCs are at a disadvantage, and some of the upcoming European legislation might have that effect with for instance NIS2.

CHRIS DISSPAIN: We might come back to that. Are there any questions that people want to ask or comments people want to make? Sir.

ABDALMONEM GALILA: This is Abdalmonem Galila for the record. Actually, mixing between gTLDs and ccTLDs in the same panel is a good idea. So the question comes from here. If I have a string of gTLDs, more than three letters have the same meaning like the ccTLDs and two letters, are there any regulations about that? That is the first question. Second question, sometimes I try to open the Internet and write a domain name, maybe this domain name is a [indiscernible] [indiscernible] for this domain name in large prices. So are there any regulations regarding this as well?

Last one. For slide number 6, you said that for the country they could be more strings for ccTLDs in IDNs; to which limit? Maybe some countries have more than 1,000 languages. So you will add the main string for a domain name and then add something, the Unicode for this domain name as well with different languages? Thank you.

CHRIS DISSPAIN: Let me try and deal with the third one, which is about – I think if I understood you correctly, you're asking about IDNs where a country has lots of different languages. The current rules in respect to ccTLDs are that if the language is an official language and the name applied for is a meaningful representation of the name of the territory or the country, then that is acceptable to go through the IDN process. And in some cases, Ram, I don't remember, India has 22 IDN effectively ccTLDs, whereas the UK has zero. So I think it depends on -- the answer is yes, there is no limit as long as it is an official language, official script, and a meaningful name. On the first point, first question, Ram.

RAM MOHAN: Thank you. That is a great question when you have a gTLD application whose meaning is similar or is the same as the name of a country. And in the previous round of new TLDs, that is actually where the GAC early warning process, all of those safeguards came in place to ensure that you wouldn't have that kind of a confusion to occur. I expect that those kinds of safeguards, they were good in the previous round, I don't think we had a great deal of names that had confusion with country names, so I expect those to stay in effect.

And similarly, if you look at dot Amsterdam, for example, or dot London or nyc or city names like that, even in those cases, the GAC helped create specific advice and shaped the guidelines on requiring governments to step in and either approve or not object before that could go forward on the gTLD side.

CHRIS DISSPAIN: And in respect to your second question, if you don't mind, that was a little indistinct. Can we take that one to the back of the room when we finish, and I will talk to you and answer that question then. Is that okay? Anyone else? Yes, sir.

TARIK MERGHANI: Hello, good morning. My name is Tarik Merghani from Sudan, I am here representing GAC and at the same time, I am the ccTLD manager for Sudan .sd, and even the IDN looked at Sudan in the Arabic language. And it's [inaudible – 01:03:44] by SIS, a non-governmental organization, it's the Sudan Internet Society, it's the chapter of Sudan there and it's managing .sd at the same time, but my real work is in [inaudible – 01:04:02] the regulator, and it is in national Sudan CERT.

So this makes a lot of confusion for me regarding Data Protection, regarding when we get a request in the CERT for some domain, if it's regarding .sd, if it's a Sudanese domain we don't have a problem but maybe it's registered to another country, an European country Data Protection. How can we deal with such things, that I can respond to myself that it's all okay. I don't know how to deal with it. Okay. Thank you very much.

CHRIS DISSPAIN: So if I understand you correctly -- Bruce, I think you may be able to answer this, because this is about how you would deal as .Au with GDPR when you are not actually in Europe. I think that is the crux of the question, which is how do you deal with regulation that's coming in from elsewhere and does that have an effect on you as a TLD? Do you understand what I'm saying?

BRUCE TONKIN: Yes, that is a tricky question actually. But with our roles in Australia, we have a requirement that you have an Australian presence, but that Australian presence can include overseas orgs that have a trademark in Australia and it can also include an Australian citizen that is living in another country, sucn as a European country. And we also have registrars that hold information about their customers that are both in Europe as well.

So when we look at a particular piece of rgulation, like a lot of the things on the Internet, you need to look at where is the registrant, if you like, the entity that holds the domain name, what country and jurisdiction are they in, and then you look at where is the registrar, what country and jurisdiction are they in, and then finally where the registry operator is. So in some cases, there are three or four country laws involved in that chain between the end user and the registry.

CHRIS DISSPAIN: I hope that goes somewhere to starting to answer what is a very complicated question. And again, we would be happy to talk about it when we finish the session if you want to carry on.

Anybody else have a question or comment that they'd like to make? Yes.

INDONESIA: Not a question perhaps, but a comment to Jia-Rong Low, if Jia-Rong is still here, hopefully; Jia-Rong's first presentation because I think from the government point of view, it is not only important to know how the data flows and so onwhen you connect to find domain.org for example, but who actually holds all the data in the CERT and root servers, for example, whether it is the companies, the organization and so on and also the functions of the organization that looks after the roots on file, for example, is it still Verisign under the contract with DOC or is it only transferred to somewhere else and so on?

Now, whether the function of IANA has been transferred to PTI and what kind of data is still being held by IANA organizations. So I think in addition to the process of the data flow and data information, it is also important to see, to know who actually holds the data both in the root server and Root Zone File, thank you.

CHRIS DISSPAIN: Okay. So there are different categories of data. But let's start with root servers and IANA. Ram, do you want to tackle that? And Bruce, Maybe you could comment as well.

RAM MOHAN: Sure, thank you. So if you take the root zone itself, the root is managed by IANA and the file, the actual root zone file is created from within IANA. Now, the distribution, the creation -- there are multiple root servers run by Verisign, so they create that first copy of that file and then they distribute it to all the other root zone servers, and each of the root zone servers have many, many Anycast instances of their servers and that provides a great deal of resilience across the entire infrastructure. But if you really look at it as Bruce mentioned at the very start, the addition of a TLD into the root zone or the removal of a TLD from the root zone, that is a function that is inside of IANA governed by the PTI Board rather than Verisign or any other private or other non-profit organization.

BRUCE TONKIN: Just to refer to a couple of slides earlier if the staff wouldn't mind winding back. The first thing to recognize, the query information that's going from the end user usually is not going directly to the top of the hierarchy. So it's not like the root server operator, if I type a domain name in here and I typed in Google.com or something, it's highly unlikely my query is going to go to the root server and that is because it's going through a hierarchy. That bottom layer I mentioned, there are probably hundreds of thousands of servers at that level so it’s very distributed, so there’s no one server that would have information on the queries. And if I could just jump to the DNS resolver slide...

So generally, when you are doing a query from your computer, your computer here, this laptop is actually not generally directly going to any of those servers. What this computer would do because I am sitting here at the ICANN meeting and my computer’s connected to the ICANN wi-fi, this would be connecting to a local DNS resolver, and that DNS resolver will store a copy of the results.

So if for example I was going to Google.com, it’s highly likely that the answer for Google.com is stored on this computer that’s maybe somewhere in this building, therefore never sends the query further up the chain, so most of your queries are being resolved in the Internet service provider that you use and at the top level, they would never see the query coming from these end users.

CHRIS DISSPAIN: Thanks, Bruce. We're coming to the very end of the session. I would like to say -- and I know you would like to say -- thank you to our panel for presenting today. And I also know that Bruce and Edmon and Ram and Roelof and I are going to be around for most of this week. And if you see us in the corridor and want to talk and ask us any questions, we would be delighted to help and answer them the best way that we can. And with that, thank you very much for listening.

[END OF TRANSCRIPTION]