ICANN82 | CF – How it Works: Root Server Operations
Sunday, March 09, 2025 – 10:30 to 12:00 PST

NICOLAS ANTONIELLO: Hello. Good morning, everybody, and welcome to this How It Works session on root server operations. My name is Nicolas Antoniello and I'm the participation manager for this session. Please note that this session is being recorded and is governed by the ICANN expected standards of behavior and ICANN community anti-harassment policy. During this session, please type your questions or comments into the chat, and I will read them aloud during the questions and answer part of this session. If you would like to speak during the questions and answer period, please raise your hand in Zoom and, when called upon, unmute your mic to speak. Please state your name for the record, the language you will be speaking if speaking a language other than English, and speak a reasonable piece. Our presenters today will be Ken Renard, Wes Hardaker and Lars-Johan Liman, who will later introduce themselves. With that, I will now hand the floor over to Ken.

KEN RENARD: Welcome. My name is Ken Renard. I work with the US Army Research Lab. We are one of the root server operators, and I'm also the vice chair of the RSSAC. Today we're going to talk to you about how the root server system works, some background. We update

our slides every once in a while. We've just done that, so this is new to us. We're going to go through this. We love questions and ask that you keep them until the end of the presentation. Write them down. Don't forget. We love questions. So if we can go to the next slide. So we'll talk about the domain name, some background, how it works, ones and zeros on the wire, things like that. We'll get into Anycast, which is very important specifically for the root server system, talk about how the root server system looks today, get into a little bit of how the RSSAC works, how we participate in the ICANN community and then get on to principles and governance and then some … Over the next slide and the next one after that. So the technical background here to set the context. So the IP address is really the fundamental identifier on the Internet. That's how computers talk to each other. It's a strings of ones and zeros, but usually—it's in the lower right there—it's represented numerically like 192.0.2.1. That's your IPv4 address at the 2001DB81. That's an IPv6 address. So that's how computers talk to each other. If we go to the next slide. So when we look at the original problem –why develop DNS in the first place?—it's because those numbers are hard to remember and they actually change quite often. With DNS, we don't see that. We've solved those problems of remembering an IP address. As the Internet has evolved, there are new problems here , especially for websites and services. IP addresses could be shared or there can be multiple IP addresses for a single service. DNS hides that from the user such that we don't have to remember them and we don't have to be involved with all this changing and distinguishing between all the different services on a specific IP

address. So DNS is a protocol and a service for translating names into other identifiers. So the classic example is translating a name into an IP address. It does a lot more things too. It does IPv6 addresses. It helps us find email servers, encryption keys. There's lots of other things in the DNS that are critical to how we use the Internet. A primary example is looking up a name into an address. In this diagram, we see the hierarchical nature of the namespace. So the names are organized into different authorities here, where we have the root and then top-level domains. These are .uk, .org, Some good examples of internationalized domain names. There's about 1,400 names [that] mark the transition of authority into who administers and controls those domains. So, next slide. We're going to walk through the resolution process of how a device looks up a name and the players and pieces involved in that process. So we start with this: either a web browser or your phone. This could be your laptop as well, your big server in the computer room, or your coffee maker or your refrigerator. Almost every device on the Internet [ends up] on this roll at some point to look up a name in the DNS. And they do this by connecting to or by contacting a resolver. We call it a recursive resolver. Now, this recursive resolver is the workhorse of the DNS operation. So this is something that's configured into your laptop, your phone, your device, either manually, often automatically when it joins a network. So this is assigned to you. This recursive resolver is run by your isp, by your organization, by the coffee shop you're visiting. Lots of different ways to do this. You can even set this to some of the resolvers that are completely public that you can use pretty much from anywhere

on the Internet—for example, Google's DNS. So when you need to resolve or ask a question of the DNS, you [inaudible], and there's a really good chance it's just going to return that answer to you. Done. That's it. It already remembered the answer from a previous query. So that cache down there, that memory, is used very heavily to optimize the DNS. If we go to the next slide. Now we'll transition back to almost the worst-case scenario where that cache is completely empty. And we'll walk through the process of how that cache can get populated and how the system can work here. So your device there on the left has asked for the address. The IP address of www.example.com. The solver has no information at all except for the basic information that it's configured with, and that is the location of the root servers. So that recursive resolver is going to do step one and ask server, what's the address of www.example.com? And that root server is going to respond with, well, I don't know the answer, but I know where dot com is. Next slide. Now your recursive resolver does the same thing, goes over to the dot-com servers and says, what's the address of www.example.com? It says the same answer: I don't know, but here's the address of the example.com server. Next slide. And the last step is that you, your recursive resolver asks the example.com name server what's that address? And now it does get a response of address. So this process happens very quickly. On a very slow day, this could take 100 milliseconds. That recursive resolver has now learned three pieces of information. It knows where the dot-com server is, where the example.com server is, and what the address is of www.example.com. All three of these pieces of

information are now stored in the cache. If somebody asks any of these same questions again, they can be answered immediately without all these any of these transactions. Now, that cache information has a Time To Live or a TTL. So that the information doesn't get stale, each piece of information has a time allocated to it or time specified that allows that recursive resolver to cache that information for a certain amount of time or it's considered old enough that it should look it up again. So that cache is very important. Now, one really important piece of this is that your device is not the only one resolver. There could be tens, hundreds, thousands, maybe even millions of other users using this recursive resolver. So think about all the questions that are asked of that recursive resolver and all the information that it caches. There is a very high probability that most of your answers are already in that cache. Now, on the example.com server, somebody might look up www2.example [and go out] again. But if somebody else looks up anything else in the.com domain, it already knows where.com is, so it doesn't have to go to the root. So as you go from the bottom right from example.com all the way up to the root, the frequency at which you end up querying those nameservers becomes less and less. So we can try and think of how many times DNS queries happen on the Internet: billions, trillions, quadrillion. We don't even know. Very few of those queries actually result in a query to the root server system. Because of that caching, it is extremely rare. Time To Live for entries in the root zone is typically 48 hours. So this recursive resolver only has to look up the dot-com, the location of dot-com, once every two hours. And how many queries

go to Google, DNS or your ISPs recursive resolver in that time? All of that is answered out of cache on that recursive resolver. So recursive resolvers are or that workhorse of the DNS resolution process and they have this cache which makes things very, very efficient for [inaudible]. [They] have this TTL or Time To Live that determines how long those entries can stay in that resolver's cache before they should be refreshed. And the root servers now only serve data about [mains.] They only tell you where.com is or.uk, or.org. They have no information that's deeper in the namespace. So really all that the root servers do is just serve those top-level domains. Again there's about 1400 of them. Go to the next slide. So these are some of the refinements, some relatively new things that have happened in DNS since its inception 40 years ago. Security is a very important. DNSSEC, the DNS security extensions, provide cryptographic signatures over the information in DNS databases. So it eliminates the risk of spoofing when used and validated by that recursive resolver. So it's up to that recursive resolver to get these signatures and validate and make sure that they're correct. Security is a very, very important piece of the DNS world and brings with it a lot of trust in the information that you get. It doesn't matter where you get the invalidate that signature. You know that it's correct. So privacy is another piece of the DNS world where the queries you make can potentially go … In that last example, the question—what's the address of www.example.com?—went to the root server. That went to the.com server and went to, you know, all three pieces. [There's a technology] called QNAME minimization, which can allow that

recursive resolver to minimize the query so that when it asks the root, it knows that the only information it's going to get is about that top-level domain. So it only queries the root with, what's the address of the dot com server? So now the root has no idea that you're going to www.example.com. So it's a way to reduce information that's really unnecessary. And [we] don't have a good way of measuring the uptake of how many recursive resolvers are using QNAME minimization, but it's definitely there and we can see some of the effects of it. Resilience is inherent in the DNS architecture. So we have the ability to really optimize the system to mitigate things like failures of systems, failures of networks. And we'll talk about that here in just a second.

LARS-JOHAN LIMAN

My name is Lars-Johan Liman. I work for Netnod, another one of the root server operators. We're based in Stockholm, Sweden, and I'm going to talk a little about Anycast and also how the root service system is. So next slide, please. Traditionally, on the Internet, every computer connected to the network has its own unique IP address. And [inaudible] by using that address in your packets, when you communicate with that computer. This is good for very many things. But for some services that are based, that are part of the infrastructure on the Internet, we need to provide service in a different way. We need to be able to provide service that has really a lot of redundancy and resiliency and also compute power. We're using a trick. All the root of operators use this trick called Anycast. And Anycast is not related to the DNS or the domain name system in particular. It's actually routing technology on the network where

the routers, the components on the network that direct the traffic across the various links on the network, talk to each other. And what we do play a little trick there. So with Unicast you have a single location for a specific address where all the … That means that if you try to overload that service, you have a pretty good chance of succeeding with that because all the traffic homes in on that specific server somewhere. It's relatively easy to overload that system. With Anycast, we have multiple instances that operate using the same IP address. And that kind of sounds like a contradiction [in] terms, but it actually does work under certain conditions. So a certain client, a certain resolver, that [inaudible] only be able to reach one of the servers using the same IP address, because the routers on the network will direct the packets to the nearest one. And that means also by having multiple servers across the entire Internet, there's one that is the closest one that's chosen in this communication. So the resolver will send the packet towards the IP address of the root name server, and the routers on the network will send it to the closest [inaudible] and also it will distribute attack traffic. Instead of all the attackers homing in on one single server, they cannot control where the traffic goes because it's up to the routers on the network to direct the traffic, and it will be spread over the entire set of servers. So each and every server will only receive a fraction of the attack traffic and hopefully it [inaudible]. So in normal Unicast situations, traffic from a source travels across the network to the single destination. We only have one server or a very local cluster of servers that receives and handles the traffic. Next slide please. Anycast. We can

have multiple servers in blue here that use the same IP address. And the routers, the circles in the middle, will forward the traffic to the nearest one. So that means you can get an answer very quickly. The path to the server that you're talking to is going to be fairly short, hopefully. And the routing, the path selection, is performed by the routers, the circles in the middle. And the end nodes cannot control that. There's no way that they can direct certain paths on the network. Next slide please. And again, if you have attack traffic, it will only hit the nearest server. So if you have a network of attacking computers, it might be possible for them to overload one node, but all the others across the globe will still be operating normally, hopefully. And this goes only for a single root nameserver. Now you have to multiply this with 13. So for each service, for each service identity, each IP addresses that you can reach, this scenario kicks into place. So there are 13 layers of this routing system. So it's extremely redundant and extremely resilient. So next slide please. And again, next slide. So the root server system today. We're going to start with a few definitions. The Internet Assigned Numbers Authority, which is an affiliate of ICANN (it belongs to the ICANN sphere of organizations) is the manager of the root zone. And what's the root zone then? Well, the root zone is a part of the total domain name system database, which is absolutely staggeringly huge. There's no way that you can fit that on one computer. It's hierarchical as described. We have different levels in the system, and the very top level is called the root. And the root zone is the piece of the big entire data [inaudible] at the top level. And again, to repeat what Ken said, it contains only

information about the top-level domains: com, net, org, but also [inaudible] the total. Did you see the number? It's 1400. 1400 names and a few records, DNS records, database records for each of them. So it's actually a fairly small piece of information. That database or that part of the database is the data that is served by each and every root server. And the data is served identically. So it doesn't matter which of the Anycast instances you talk to and it doesn't matter which operators service you talk to. You get identical answers. It's all extreme service system is the set of all machines, all servers, that provide this service. In total it's over 1,900 [inaudible] 13 IPv4 addresses. Each is operated by a specific root server operator, and again 13 IPv6 addresses. Every identity has both. And the total combined set of servers operated by all the root server operators is referred to as the root server system. Next slide, please. So root server operators. These are 12 organizations that operate 13 identities. Okay, one organization operates two for old, historical and hysterical reasons. So the root server operators each operate a set of servers in Anycast mode. And all the servers in that root server instance or a root server Anycast instance is one of these servers operated by specifically one root server operator. And again, a root server operator handles many instances. And together all the root server operators and all their respective instances build up. The root server operators come together in a couple of different ways, but one way here in the ICANN context is the Root Service System Advisory Committee, RSSAC, and it's comprised by each and every one of the 12 root server operators and then a number of liaisons to other parts of ICANN and, also

outside of ICANN, for instance, the Internet Architecture Board, which is the organization in the IETF standardization part of the Internet sphere. Next slide please. So here is, just for your information, a list of all the host names or identities, as we sometimes refer to it, that you can reach for this Root service. They're just named A,B,C,D up till M. And here are also their respective IP addresses, both IPv4 and IPv6. And also here you see the list of organizations that operate the various instances, various identities. So you will see that Netnod, my company, operates I-root. One thing to observe here is that the type of organizations is very diverse, and this is one thing that we try to stress. The root service system needs to be extremely diverse because we don't want any single points of failures. We don't want any single points of vulnerability where external [inaudible]. So we operate different operating systems, different types of hardware, different types of network equipment in front of our servers. The only thing that is really common between all the root server operators, is the root zone. [They operate] on exactly the same data, and the data is public. You can download it on your web browser here and now. I can show you afterwards if you want to know how to do it. But also we have diverse types of organizations. So you will see here government related organizations, you will see universities, you will see private companies, you will see Internet service providers. It's a big variety, and that creates more resilience in the system. So that's actually a good thing. Next slide please. This is just a quick overview of the history of the root service system. [It started operation] in 1984. So that's actually more than 40 years ago now,

and it has increased in numbers of servers gradually. So in 1993—no, that should be '91—the first server outside the United States was put into operation. And since 2001 we have deployed Anycast. We started in 2001 with a few more servers out of the 13. Back in the year 2000 we had only 13 machines, but it was a different Internet 25 years ago. Since then, more and more Anycast nodes have been deployed, and now we are up to over 1900. So I'm not going to say that it's impossible to take out the root server system, but I'm going to say [inaudible]. The root server operators have a common website, rootservers.org, where you can find information about the respective root server operators and their instances. And there's also a clickable map where you can see where these instances are located. So you can zoom in by clicking on this map and find cities where root servers are deployed and where your nearest root server is located. Next slide, please. And also a little history over the query load, how it changes over time. [We] find that the query load increases in ways that we don't really understand why. So we start to investigate. Do remember that the roots operators work very closely together. I know Ken quite well, I know Wes quite well, and I know [all the others] well. And we have very good ways to contact each other in times of crisis, which happen extremely seldom. But there are such. And also we work together when we see these anomalies in traffic patterns or query patterns and whatnot. And sometimes we manage to [inaudible] certain implementation of a certain piece of software that is misbehaving, and we can talk to the implementer and say, "Hey, you actually cause a lot of traffic going to the roots just with your software, so

please fix it. There are better ways to change the software." Next slide, please. There are a few myths about the root server system which we like to try to dispel here. One is that the root server system control where Internet traffic goes. No, they don't. They only tell you where your destination is. [That's] guided by the routers. Again, the circles in the middle in my previous picture. That goes for all Internet traffic. The root servers are involved in finding out where the address of your destination [is] as a step in this process of finding out DNS data, but they do not control where the packets go. Most DNS queries are handled by the root servers. That's not true. Ken pinpointed very carefully that the caching mechanism in the resolvers is very, very [inaudible]. That offloads a lot of traffic from the root servers. So most traffic is actually handled by the resolvers in most cases by finding stuff in the cache, which is already there, in some cases using shortcuts in the system. If they know where the com server is, they go directly there. If they know where the example.com server is, they go directly there. So they shortcut the system as much as they can. Very few queries actually go to the root. The root zone is not handled by the root server operators. It's handled by the IANA and it's implemented by the root zone maintainer. And all the root server operators get exactly the same copy and serve it [inaudible] able to change the data. Well, we could, but immediately that would be recognized by the resolvers that do DNSSEC validations. The root server operators do not [inaudible]. DNSSEC is there just precisely to give the clients a good tool to validate that the data is correct. So if the root server operators were to [inaudible] who would flag that and

say oops, signatures don't match. So all the root server operators serve exactly the same data, exactly as delivered from the IANA via the root cell maintainer [inaudible]. The root server identifiers have a special meaning. No, that's not true. You have the letters A, B, C through M. They're all alike, none of them [inaudible]. It's not the case that any one of these is the primary source for the root zone. They all get the root zone from servers that are hidden behind … They're up in the system from which each root server operator downloads the zoned file, the root zone file, and copies it to their respective Anycast nodes. So it really doesn't matter which one you talk to. The quality of the data is exactly the same, and it's updated at the same time. It contains the same DNS signature. Everything is identical between them. There are only 13 root servers. Well, that was through in the year 2000. We're now a quarter of a century further down the line and we now have more than 1,900 servers providing this service. Root server operators conduct operations independently. That's not true because we coordinate very carefully [inaudible] how to provide this service. If there is a major systemic change to the system (for instance, we want to add service over IPv6), that's carefully coordinated between all operators. Or if we [inaudible] the introduction of DNSSEC. That was a systemic change, carefully coordinated again between all the root server operators. And we meet regularly at least three times per year and have technical meetings where we coordinate and make sure that everything is running smoothly. So it's actually a very, very well coordinated effort. [inaudible], when you do that. That was the traditional way of operating DNS, how it was first

designed back 40 years ago. And with the modern ways of handling this, the QNAME minimization, we see more and more traffic which does not contain the entire query. We only saw where is the com server, where's the org server, where is the SE server for Sweden or NO for Norway and so on. So we don't see the entire query because it's really not necessary. And as you mentioned, it's a [inaudible] root service, and the clients really don't need to send that information to the root servers. So that's a trend going down. It's not gone entirely yet, but it's definitely [on the way]. I think that's the end of my part—no. The root server is the root server system. The root zone again is the data, the data that is served by the root service. The root zone is the data. The server system is the set of machines that provide access to the zone file. And the [inaudible] point [is] the list of top-level domains and their nameservers. If your top-level domain is not listed in the root zone, it simply doesn't exist on the network. No one can find information. So it's the authoritative list of operational top-level domains. It's managed by ICANN per community policy. So the policy for the root zone is set in the ICANN community by the IANA. So that's where change requests come in. They are sent to the IANA and then IANA informs the root zone maintainer who produces the zone, signs the zone. And the root server system is the system that you can access to obtain the data in the root zone. It's currently provided by 13 IPv4 and 13 IPv6 addresses from over 1,900 servers across the entire globe. We have servers in places I hadn't even heard of before we deployed servers. So you will find them in a lot of strange corners of the world, really. The root server system is a

purely technical role. This is just infrastructure. This is just operating machines to make sure that there is access to the root zone content. There is no real policy involved in operating a root server. We get the data from the IANA through the root zone maintainer. We don't make any changes. We don't have any opinions about what's in the root zone. We just receive it and put it on our servers. And our job is to make sure that all users on the Internet have access to the root zone data as delivered from the IANA. The responsibility to operate the root service system comes down to the root server operators, these 12 organizations that operate, the 12 diverse organizations, I will add, which have [inaudible]. We are located in different parts of the world and we collaborate very closely to make this happen. And this is my last slide. The next slide, please. Last slide. And this is a picture of how the root zone is distributed. Again, a change request may come from a top-level domain operator. They may want to change their set of nameservers or they may want to use a different crypto key or something. That change request is sent to the IANA function which evaluates it and makes sure that it's a correct [inaudible]. And then the change request is sent to the root sign maintainer. That's a role that's currently filled by the corporation VeriSign. And they in turn distribute [inaudible] own updated and signed to the root server operators. And each root server operator then distributes, to each, its set of Anycast nodes or instances. So the part that takes time here is for IANA to validate the request and make sure it comes from the right source, that it's a reasonable change to the system. And as soon as it's here and made available

through the root server operators, it's a matter of minutes before that's deployed on all the Anycast instances. It's really fast. For Netnod, the time is less than [inaudible] that there's a new version. It's 15 seconds until it's updated in all the corners of the world. So it's a smooth operation that works really well. Now I'm done. Now you get to listen to Wes.

WES HARDAKER

Good morning, everybody. So my name is Wes Hardaker. I am from the University of Southern California's Information Sciences Institute, which was actually where the DNS was created long before I got there. But we do house the original, which is why in one of the slides it said that we had two of the root servers back way back in the beginning, because we started it. And I'm going to talk to you today about what RSSAC is; so how the root server operators are represented within ICANN and the RSSAC Caucus, which is how anybody can help us out. Next slide, please. So what is RSSAC? The role of the Root Server Advisory Committee is to advise the ICANN community and Board members on matters relating to the operation, administration, security and integrity of the Internet's root server system. So that's strictly straight out of the bylaws, I believe, and it's a very narrow scope. We don't talk about operations within ICANN. Go on to the next slide, please. We are a committee that produces advice. It's primarily to the Board, but we produce documents for the community to read as well about how the root server system works overall. And within that, the root server operators are represented within RSSAC, but in RSSAC itself, actually, we don't talk about operations and how to do things. We

talk about how the community needs to interface with us. Next slide. So RSSAC the organization is sort of divided into two parts. There's the RSSAC, the advisory committee, which has a primary and a secondary and alternate representative from each of the root server operators. So that's 24 members, two from each of the 12 root server operator organizations. And then we have incoming liaisons from other bodies. So I'll get to that more in a minute. The other half of RSSAC is really the RSSAC Caucus, where we do the vast majority of our work. And this is where anybody with expertise can come help us. And it's a body of volunteer subject matter experts. You can apply to be a member, and the members are confirmed by RSSAC based on the statement of interest that they submit and why they want to help. Next slide. The current chair and vice-chair of RSSAC are Jeff Osborn. So normally we actually have more root server operators in the room. It turns out that we're split today because there is another discussion about root server operators for the Fellows I believe up there. So Jeff actually would normally be in this room, but he's giving a presentation about the root server system in a completely different room right now. But Ken Renard is our vice-chair, and he's from the US Army DEVCOM Research Lab. Next slide. So a little bit more about the RSSAC liaisons. We do communicate with a lot of other entities that have investments in the root server system to make sure that it's functioning properly. So we have incoming liaisons, liaisons from other organizations that sit in our meetings and partake in our discussions. One is from the IANA functions operator, which is PTI. One is from the root zone maintainer, which Liman talked a little

bit about, that actually works with IANA to produce the root zone that all of the root server operators make use of. We have another incoming liaison from the Internet Engineering Task Force, the IETF, and that is that liaison is appointed by the Internet Architecture Board, which is why it's on there. And Andrew, I was realizing this earlier today: we probably need to change that to IETF as appointed by the IAB, just for sake of clarity. So that person's actually responsible for representing all of the DNS standardization that happens within the IETF. The IETF is where the protocol actually gets modified over time. And then we have another one from the Security and Stability Advisory Committee, SSAC. SSAC is the security community within ICANN that we talk heavily with. We meet with them on a regular basis because we're very closely aligned in the types of things. And then RSSAC also produces a number of outgoing liaisons from RSSAC to other entities. The first is the ICANN Board. So we have a direct director that sits on the ICANN Board along with the other directors as a non-voting director. That's me. I currently serve in that role. We have a liaison to the ICANN Nominating Committee and the Customer Standing Committee and the Rootzone Evolutionary Review Committee which we haven't talked about much yet today. But they're responsible for if changes do need to be made to the contents of the root zone, not from TLDs but from … Let me go through an example. The most recent change that they made was to include a new cryptographic checksum of the contents of the root zone so that you could verify that you have the complete zone. That process got developed in the IETF and then RZERC recommended

that that change implemented by the root zone maintainer. And that happened actually only in the last couple of years, I want to say. Next slide. So what is the RSSAC Caucus? That is where we do most of our work. In the advice that RSSAC produces, we want the widest set of opinions to help produce that advice, the widest set of experts. So, far beyond just the 12 root server organizations, we want people from across the DNS industry to help us out. So the membership exists of over 100 DNS experts at the moment. As I mentioned earlier, we take a public statement of interest. Those are all published so that you can go read them; all the members and what their background is. And one of the primary goals of doing this is we give public credit for the individual work. Go read any of our RSSAC documents. At the bottom of the document is a list of the people that help produce it. And it's not just members of RSSAC. It's members across the DNS community they've helped out. So what's the purpose of the RSSAC? As I indicated, it's experts that bring a diverse set of expertise and opinions to our publications. We provide transparency for who does the work, and it's really the framework for how work gets done. The RSSAC produces documents. They often come up with ideas that they want to document. We work on it, usually over the course of a year or two, and then it gets published. So RSSAC eventually decides to approve it as a document to go forward. Next slide. I'm going to dive a little bit into the principles. What does RSSAC stand for? What are our core values? And how does the governance system work within RSSAC? Next. So one of the things that we did starting in RSSAC037, which I'll come back to in a minute, is we realized that

nobody actually had documented, what are our core principles, what do we stand for, what are our values? And we produced 11 principles in RSSAC037, and I'm not going to read them all here to you, but I will dive into a couple of them. And then even after we produced RSSAC[03]7, which was really a document about how the governance structure should work, we realized that we could even flesh them out even further and describe them better. So we produced RSSAC055, which adds additional bits of information to each of those 11 principles. Next slide, I'm going to dive into a couple of them. The first is principles two and four. So principle two states that IANA is the source of the DNS root data. IANA is the only source of data. The root server operators are committed to always publishing IANA data. And as Liman noted earlier when discussing DNSSEC, you'd notice if we didn't. It's actually detectable through cryptographic verification that we're not publishing the IANA root zone. So we are firmly committed that that is where the root zone will always come from. This principle number four is: the diversity of the root server operations is a strength of the overall system. And what that really means (and this was highlighted again before) is, by having a diverse set of organizations, we bring a diverse set of skills, we bring a diverse set of funding sources. Nobody is actually paid to do this. It's entirely a volunteer type of position. But if one segment of the economy in the world collapsed, it wouldn't take out all the root server operators because we're all from different venues. I'm from an academic institution, we have some from government agencies, we have some from organizational agencies, some from [martial] agencies. And that actually that diversity really

helps ensure that the root server system will always be functioning. Next slide, please. So Principle 9 (and these are the last two I'll talk about) is that RSOs must collaborate and engage with their [stakeholders]. That's why we're here, right? That's why we come to ICANN. There's many root server operator representatives in ICANN. It's why we go to the IETF. We participate in the development of the DNS protocol and how it works. We do a lot of outreach and participation in the standardization. That often means really long trips for us. This trip is back-to-back with the IETF, which is next week in Thailand. And it's going to be the longest I've been away from my wife in 30 years. So, it's a hard call sometimes. We don't like it when they're back-to-back. Often we get a week in between. But last November and this trip are both unfortunately back-to-back. And RSOs must be anonymous and independent. That doesn't mean we don't communicate. As was discussed previously, we communicate all the time. Through email, through a chat system, through regular meetings, we communicate all the time. But I never tell Liman how he should deploy his next node or here's the software he should run. We make our own decisions autonomously. And again, that's a strength. That independence means that we're all running slightly different software, we're all running slightly different hardware. If there's a major vulnerability in one piece of hardware or something like that, you just turn it off. The rest of the system will take over. So there's no single point of failure in the system. And that's by design. Next, please, please. So the one hiccup is that, when the root server system was last updated, it was when Jon Postel was alive. He was

the one that was responsible for adding and removing root server operators. And he passed away, unfortunately, 20 years ago, earlier than planned, and earlier than he planned. And the end result is there's no follow-on instructions for how to make changes to the system in a way that everybody agrees with. So RSSAC in June of 2018 produced RSSAC037, trying to set the seeds for, okay, how do we come up with a governance system for making changes to server operators? What happens if somebody quits? How do we deal with that from a structural perspective that the ICANN community can get behind? The end result of publishing RSSAC037 and 038, which actually had the recommendations, is it led to the creation of the Governance Working Group, which is meeting here this week. It was a couple of years before that started getting stood up. And they're taking the next steps. RSSAC037 was sort of a basic model, just circles and arrows on a diagram with descriptions of the functionality, but not down to, okay, this is how we actually vote, this is how we actually make changes. The Governance Working Group is now responsible for setting the next step forward. That group is making good progress and it's composed of people not just from the root server operator community like RSSAC37 had, but also we have representatives from the other stakeholders, including things like the TLD operators and, and there's people there from the IAB-appointed liaisons. Again, the ICANN Board has liaisons in it, as well as other members of the community of ICANN. These sessions are open to observers. They're all publicly documented. The minutes are published. You can actually attend the meetings this week. They're long and dry sometimes, but

you're welcome to join us if you are interested. Next slide, please. And actually this slide describes better the makeup of the Governance Working Group. So there are members of the ccNSO, there are two members appointed by the Internet Architecture Board, there's one member from each of the RSOs. I represent the University of Southern California there. There's two members from the Registry Stakeholders Group, two liaisons from the Board, one liaison from the root zone maintainer, and one liaison from IANA. So we try and get again a broad perspective in everything we do. I will say that, on the discussion today about what the root server system is (RSSAC is), this is not actually the root server system in RSSAC. It's related to it, and we can answer some questions, but we are not representatives of the GWG. Brad Verd is the chair of that, and he would be the right person to ask more detailed questions of. So how does this affect you? How is the root server system helpful to you? What should you do to be better integrated with it? Next slide, please. So first off, as we mentioned, the best way to ensure that you're always getting authentic data from any root server operator you talk to is use a validating resolver. If you are running a network, you can make sure that you're using one of these. If you don't know what this means, don't worry about it. If you're an Internet service provider though, you want to be using a validating DNSSEC resolver. Many of the [inaudible] that people are using these days, like Google's and Cloudflare's and Cloud9, are actually all validating resolvers. They will not give you an answer that they did not check is actually authentic, that came from IANA. You can participate in DNS technical communities. The root server

system and the DNS in general are widely talked about beyond just the root zone in all operating groups: NANOG, APRICOT, and RIPE. For the RSSAC Caucus, anybody, as I mentioned, is allowed to apply to become a member of the RSSAC Caucus and join our mailing list. We actually have recently established an observer role where you can read the RSSAC Caucus mailing list. It's all publicly archived. You can open it up on their webpage and look at it at any time. And then the Internet Engineering Task Force is where the DNS is actually modified at the protocol level. And then if you're interested in hosting an Anycast instance, if you want something near you, our recommendation is generally that you don't need all 13 of the identifiers right next to your you, but it's always good to have three or four. And as Liman said, they're actually all over the world now. You already have three or four most likely near you. If you think you need another one and you have a data center you want to put it in, reach out to us. You can send mail to ask-RSSAC@ICANN.org, and somebody will respond with details about how you could. And then finally, additional resources. So if this wasn't enough information overload for you today, we have more for you. We have a webpage, RSSAC.ICANN.org, of course. On it, there is an FAQ that [have questions that were] actually answered in many of the slides today as well. But it's another good place to go to look for additional information about what RSSAC is and what the root server operator community is like. And then as I mentioned, there's an email address you can write general questions to. And we do write back. And then the RSSAC Caucus has a webpage as well that talks about it and how to join as well.

So that [inaudible] and what it takes to become a member of the RSSAC Caucus if you have a skill set that we would benefit from. And that is it. So with that, I think we have made good time. We have some time left for questions. With the newer slide deck, we weren't 100% positive how much time that it would take to run through everything, but it turns out we still have quite a bit of time left. If anybody has questions on anything, it was a long past hour. Whether you have a question about the technical half earlier on or the organizational structures that I just talked about, please come to the table and we can start taking questions.

NICOLAS ANTONIELLO: Thanks very much for the presentation. We are now in the Q&A section of the session. So again, for any of you having questions, comments, if you are online, you can post it in the Q & A pod or raise your hand if you want to use the mic. If you are here in person, the same: you can just raise your hand, pick up the mic and make the questions.

UNIDENTIFIED MALE: Who decides what would be the Time To Live, TTL, which is 24 to 48 hours? I mean, assuming the growing traffic, supposing it's to be made [inaudible] to, say, a week.

WES HARDAKER: Yeah, good question. So there's actually a couple of timing considerations. The TTLs on all of the root zone data is two days, and that has been that way forever. That's not quite true. But if that

change was going to be made, if somebody wanted a different time, that would actually probably be done through the RZERC, I assume, which would be making that decision. So again, that's the body that I talked about. If there was technical changes to how the root zone was constructed, that would be done through RZERC, which is a different constituency within [ICANN], and it's open for public observation as well. I will note that a lot of software actually doesn't do two days. They actually do a one day. So different implementations do different lengths of time. The other thing to note is that DNSSEC signatures are not two days. They are 14, I think ,for the root zone. I'd have to go double check. Every zone has a different TTL. But the root zone actually has a longer TTL than most other zones on the Internet.

UNIDENTIFIED MALE Is there merit in increasing the TTL?

WES HARDAKER It's unlikely we would want to increase it. But one nice thing about very long TTLs … And one of my colleagues at ISI actually did a research paper on this [inaudible] that showed … Let's say you were going to some website under some TLD, and that TLD was under attack. So they weren't responding. Let's say it was.com—so huge it would be nearly impossible to take out. If your local resolver had already memorized dot-com, you're going to be able to get to anything in that for two days. So the TTLs actually provide you a buffer of protection against things like DDoS attacks [inaudible] really long. And there's been studies that show that longer TTLs

actually protect you from events on the Internet; outages, whether it's a regular outage caused by a catastrophe, or any kind. So longer TTLs are somewhat better. A lot of [zones] on the Internet don't have long TTLs, but the root zone does. So that's actually a reason for leaving it long.

UNIDENTIFIED MALE What's the process for adding a new RSO?

WES HARDAKER That is what I discussed; there isn't one. And that is why, in 2018, RSSAC produced a document to sort of start that ball rolling. It is currently being continued to be worked on within the Governance Working Group within ICANN. They are meeting later this week, all day on Tuesday. It's going to be a very long day. But that is not done yet and that is under discussion right now.

UNIDENTIFIED MALE So that's essentially something the Board would have to act on.

WES HARDAKER Eventually [inaudible] is going to produce advice to the Board that is in their charter, and then the Board will act on it as well. There's sort of these stair-stepping actions for how to get that process underway. And the first one is how do you create a governance system that can even make decisions on their own? So, unfortunately, TBD. We're not there yet.

UNIDENTIFIED MALE Thank you.

UNIDENTIFIED MALE So I would like to add in another thing. You just explained, like in the outrages, if you increase the TTL so it is useful … How long can we make it? Like, 7 days? Or 10 days?

WES HARDAKER You always have to balance longer TTLs against the need to know when things do change. And the reality is that changes do happen to the data as TLDs decide to add and remove IP addresses or keys or things like that. So two days seems like a good balance, and nobody has really had the desire to increase it or decrease it. Whether it's perfect or not is subject to debate, but I would be hesitant to make it longer.

UNIDENTIFIED MALE For example, if I increase it to, like, 14 days, will the DNSSEC validation work in that scenario?

WES HARDAKER It will, yes. There are two time limits. And so both the TTL and the DNSSEC time signature length both matter, and you will always take the shorter of those two that are coming.

UNIDENTIFIED MALE So I was reading that, like, seven days is the permitted limit because after seven days, your signature will be invalid. The DNS resolver decide[s] to increase the time for seven days, and beyond the seven days, then the signature will be invalidated. Is it the correct? I'm just asking you.

WES HARDAKER It won't be invalidated. The resolver will recheck. So even if the TTL has not been reached, but the DNSSEC signature time length has, it will go recheck. So there's really sort of two Time To Live values, almost.

UNIDENTIFIED MALE If it's unable to reach to the root server? In that case?

WES HARDAKER If it's unable to reach the root server, then it will likely return a failure. But as Liman has indicated, there's 1,900 instances around the planet, and it would be really hard to fall into that case.

UNIDENTIFIED MALE Thank you very much.

NICOLAS ANTONIELLO We have a question posted online. I'm going to first read it as it is and then do a small slight disclaimer before handing it to you. The question says, “What's the Board's thought about the RSO given only to a few selected ones?” So the disclaimer is that the Board

probably [is turning] to something that we can answer in this room, but I can change it a little bit. The question is, what's your thoughts about the RSO given to only a few selected ones?

WES HARDAKER

Yeah, fair point. And as the RSSAC liaison to the Board, you know, I cannot speak for the Board. Tripti Sinha is actually also part of a root server operator, but she is the chair of the ICANN Board, and she would have to speak. Right now, the ICANN Board is waiting for the GWG to get done so that those types of decisions can be made. As I indicated previously, when Jon Postel passed away, no changes have been made [inaudible] then because there is no processes, as one of the previous questions asked. So I think everybody would like the GWG to finish so that we actually have a process that everybody can trust and stand behind. But that's where we are today based on history of how the DNS was created a long time ago.

KEN RENARD

I'll follow up. The thing to keep in mind here is that the root server system is a technical [inaudible] end servers out there. There's 1900 of them, and we as RSOs really concentrate on making it available, correct and the best performing with respect to the environment, which includes bad actors. And we've done a pretty good job over the past 40 years, and we try to keep things in perspective of that technical availability and service. So the number of operators isn't

as important to getting answers to your DNS queries as the total number of servers out there.

WES HARDAKER: One other point to that is with the advent of DNSSEC, which wasn't around 20 years ago, nobody can lie about the data or alter the data. I mean, all the servers are equivalent, and you would know if they were not. And so the introduction of DNSSEC means that it doesn't really matter who is operating a server. If it's producing the right data, you would know whether it was incorrect or correct.

UNIDENTIFIED MALE: I just wanted to ask a clarifying question. There are RSOs, but the slides mentioned 1900 kind of servers or whatever you want to call [them, for] the distribution of the actual technical infrastructure. Are those all operated by those or are they subcontracted out? Because in one of the other slides it said contact us if you would like to offer an Anycast kind of system. So it seems like there are multiple layers here. There's 13 and there's 1900 and there's something else there.

LARS-JOHAN LIMAN: Good question. There are actually 12 organizations that operate 13 identities because one organization operates two identities. But to answer your question, yes, these 1900 are operated by the 12 organizations. Yes. So for instance, at Netnod, we operate something like 85 nodes across the entire globe, and it's operated

by our staff. [There are] varying numbers of Anycast instances per organization, but in total it's 1900. Thanks.

UNIDENTIFIED MALE [inaudible] during some of the engagement with the different community stakeholders [inaudible] something that never happens, but you know, the what-if questions about the architecture of the root zone. What happens if the root zone maintainer, say, temporarily loses contact with the root server operators? And how long can the root server operators run without being able to communicate with the root server maintainer?

LARS-JOHAN LIMAN What will happen in that case is that the updates to the root zone will possibly be stalled, but the root server operators will then continue to operate the data they have. It might be slightly outdated, but it will still work. And it will work as long as the signatures for DNSSEC are valid. And that's a timestamp. So I think the signing time is, what, two weeks. I need to check up on the exact number. Wes will check for us, but the order of magnitude is a couple of [weeks]. During that time, things will continue to work, but they will not change. So if a top-level domain operator wants to change some data in the root zone, it will not be possible to build during that time. That said, the system that we use [on root zones] is not the only way to send updates. Worst case, someone can call me across the ocean and read the changed data into my ear and I will type it into my computer and make sure that it's all validated and [has] signatures. So the existing systems that we use is only

one method. There is a number of engineers and we have been working with these systems for, in my case, over 30 years. There are a number of tricks, alternate methods, because again, it's actually not a huge amount of data. And the only thing a root operator actually needs to know are the changes to the data, and they are often very small per day. So it's a system that changes very slowly—only very small changes every day. And we can find alternate methods to distribute that and also distribute things among ourselves. So if only one root server operator has access to the root zone maintainer, they can then distribute it to the rest of the root server operators using other methods or different parts of the network or something. We will invent ways to get this going because that's what we do. Thanks.

WES HARDAKER Yeah, the root zone maintainer has a very extensive distribution network just to get to us. And then as Liman just said, we can give data to each other. I did check. It is two weeks. My statement earlier was right. And the reason we don't remember is we get two or three updates a day. So the signatures are always valid for a very long time beyond today. But it is a two-week signature at this point.

NICOLAS ANTONIELLO I have no more questions in the chat. If there are no more questions in the room, I would like to thank very much Ken, Lars-Johan and Wes for the presentation. I think you achieved something that is not easy, which is to put all of these technical and non-technical stuff

into something that it's very easy to understand for someone that is new or that not necessarily knows about this. So, really, thanks very much for your presentation. We can stop the recording and the meeting. This is the end of the session. Thanks very much.

WES HARDAKER And thank you all for coming.

LARS-JOHAN LIMAN One last thing. If you have more questions during the week, we are around for the entire week. Please just stop us in the hallways. We are happy to talk to you. We like questions. We like to try to inform you about what we do. So just reach out to us.

[END OF TRANSCRIPTION]