CLAUDIA RUIZ: Good morning, good afternoon, and good evening everyone. You are all welcome to LACRALO monthly teleconference on Monday, May 18th 2020 at 23:00 UTC. In Spanish, we have Sergio Salinas Porto, Harold Arcos, Adrian Carballo, Anahí Menendez, Antonio Medina Gómez. On audio, Augusto Ho, Humberto Carrasco, Lilian Ivette De Luque, Lito Ibarra, and Vanda Scartezini.

We have received apologies from Dev Anand Teelucksingh. From staff, we have Silvia Vivanco and myself, Claudia Ruiz, managing the call. Our interpreters tonight are Marina and Paula for Spanish, Bettina and Galina, and Claire and Jacques. A kind reminder: will you please say your name for the transcription and the interpretation?

SERGIO SALINAS PORTO: There is an unidentified person asking for connection. Who is there?

WLADIMIR DAVALOS: This is me.

SERGIO SALINAS PORTO: You are welcome. Claudia, will you please mute everyone so there is no background noise? I don't know if Harold is already with us. If he is with us, I will give him the floor. But before that, I should say happy Internet Day to everyone, and that’s all. If you are here, Harold, if you are so kind as to start with the agenda?

HAROLD ARCOS: Thank you, Sergio. While we work on the issues with the Zoom room, let's start with the agenda. We will start with the adoption of the agenda, then a webinar by Nicolas Antoniello on "the General Operational Resilience of the Domain Name System," then there will be a regional exchange on DNS abuse.

After that, there will be comments by our ALAC member, an update by Sylvia Herlein. After that, a few minutes for discussion on LACRALO's engagement tool by myself, then the update from the working group by the directors, and that is a regional update by Sergio Salinas. He will tell us about the elections, the rules of procedure, and our strategic remote plan.

Before the evaluation of the webinar, we have another item in the agenda that's available for any other business. I'd like to know if anyone wants to add any topic. I do not have the chat. I cannot see it. So, if anyone can help me if there is any topic to add to any other business? That's all.

SERGIO SALINAS PORTO: Thank you, Harold. There are no further issues. If there is any proposal made during the call we will certainly add it. Thank you for reading the agenda for us, which has been adopted as-is. Now, we will have our webinar. It is held by our region, together with the Capacity-Building Secretariat of our region and GSE.

With us, we have Nicolas Antoniello who is the Regional Technical Engagement Manager of ICANN for Latin America. He will have 40 minutes for the topic that I believe should be in everybody's agenda, which is the General Operational Resilience of the Domain Name System, DNS, which is a highly sensitive topic in ICANN's order of business, and I think we should be aware of what this is about in as much detail as possible.

Without further ado, Nicolas, it is a pleasure to have you with us. You are welcome and you have the floor for your 40 minutes to explore this topic and some time for questions on our side. Thank you, and you have the floor, Nicolas.

NICOLAS ANTONIELLO: Thank you so much. I hope you can hear me well. Probably that is the question everybody's is asking for the past two or three months. Can you hear me well? Well, I will try to share with you. How can I share screen? Is there a way to do that? There it is. I hope you can see the screen I'm sharing. Great.

So, in these 40 minutes I intend to make a general overview of the operation of the DNS. For those of you with years of experience in the management of name servers or involved on DNS, you will probably find this too familiar and, perhaps, somewhat boring. Nevertheless, it is aimed at the general public and we are assuming that there is no prior knowledge of the DNS.

After this basic overview of its operation, we will also talk in general terms, also. We will talk about some mechanisms used to make a DNS more robust, more resilient, and fault-tolerant: natural faults, hardware faults, or any other software problems, or intentional faults resulting from attacks, or other undesired activities. So, let's start.

Well, of course, before that, at any time if you have questions, or if you think I'm going too fast, or if you do not understand anything, please let me know and I'll slow down and explain. Otherwise, at the end of the presentation, certainly, I will be available for any questions and comments.

So, let's start with a little bit of history: how the Domain Name System started. As we all know, in Internet devices that identified over the Internet using IP addresses, both IPv4 and IPv6 version ... I apologize for this interruption.

So, as I was saying, IP addresses identify devices but people in general remember names much more easily than addresses. So, similar to a phonebook in the traditional telephony system, the same approach is used by any system translating names to addresses.

So, at the beginning, that was done through a text file that was stored in each device. You probably remember the file host and .txt, which basically contained the translation of a name into an IP address for the device.

That posed several problems. Though it fulfilled its original mission, it also had some problems—scalability issues, loading issues—because whenever a change was made of the file, it had to be overwritten onto all devices, and that resulted in problems of synchronization. There were also significant bandwidth issues because that file grew because it comprised all translations into a single file. There was no hierarchy.

As a result of all that, there was this idea brought up to create the Domain Name System, the well-known DNS, which is, indeed, a hierarchical system—it is not a flat system as that file originally was—and as addressing in directional telephony work.

This is a general system where the entire management is not centralized on a single location but it is distributed, as we all know, all over the world. So, that’s how the Domain Name System started. It’s a structure. The original structure when the system was created for standardization, it’s an inverted tree structure where the so-called “root” of a Domain Name System is represented by a dot on the apex of the inverted tree.

And then, there are several levels below, several branches that stem off that dot. In each node after that in the DNS lingo is called “domain.” So, we have several levels. The so-called first-level nodes or domains, the second level, third level, etc.

So, among the top-level nodes we have .com, the .org domains, the domains assigned as country codes, the ccTLDs, and many more other ones that were created afterward.

The ccTLDs and the relatively few names with only three characters. Now, there are many more and they are not all only three characters. As an example, I'm showing here .coffee, the domain, and there are many more.

So, this inverted tree structure with several levels is the shape with which the distributed database was created. That database, we are called the Domain Name System.

So, what are the components of a name? There is a certain way to denote the domain name which is known as "Fully Qualified Domain Name," which basically works as follows.

To indicate a domain name, it starts with the name or the label of the node, and then all the following nodes backward to get into the root are indicated.

For example, this domain example, the right and complete full way of naming it is the following: example.com. As you can see, the labels comprising the domain name are separated by dots, and this root, which we do not name it, actually, is always there. This one, here, www, would be www.example.com, and that is the same with any other domain we would like to designate. So, let's move on. Are there any questions? I perhaps heard someone.

SERGIO SALINAS PORTO: No, Nico, it was just someone that was not on mute.

NICOLAS ANTONIELLO: What is a domain? A domain is a label of a given node at a given level of the tree, and by "domain" we typically mean that label and everything below it. In this example here, this domain, .com, the top-level domain .com consists of the entire tree that follows this domain. In other words, all these labels which are below it that make up the entire tree are below the domain .com.

For example, if we were talking only about these two, .mail and www, these two, these are within the domain example. So, this is the domain comprised, an example, and everything below it.

Now, if we are talking about the root or the label "dot," that domain is the entire DNS. The entire Domain Name System is contained below the root. How does this work? What is its management like in practice?

This database is not managed in a central way. So, parts of the tree are delegated. So, this activity, this action of delegating the management, the administration, results in the creation of various zones.

This typically leads to confusion. We might confuse what a zone is and what a domain is. A domain makes reference to the structure of the Domain Name System. A domain is a label, a position in the tree, and everything below it, while a zone, unlike a domain, is an administrative space.

It could be that the zone is the same as the domain but it may also be different. In this example, the domain .com … Remember, it used to be the domain was everything below .com. In this example, the manager of the .com domain delegated to three entities the management of the tree below it.

To one entity, it delegated just the management of the domain .example and everything below it. To another entity, the management of the domain .bar and everything below it. And to a third entity, the management of the domain .full and everything below it.

So, in this example we have a zone consisting exclusively of the domain .com, and in turn we have the delegation of three zones for the entire tree below the .com domain: .full, .bar., and .example. So, we have one domain, .com, with four zones in this case. The entity management .com and other three entities manage the three other domains. That is to highlight the importance of understanding the difference between a domain and a zone.

Now, let's talk about the internal composition of the database. In other words, how we store the information. We have already described its structure, we have already explained what a domain and what a zone is. Now, let's see how the information is stored there.

The mechanism to store the information is, basically, this thing we call "records." Records are containers of information within the DNS database and each record has a specific functionality. So, the type of information and the functionality or the use of each record will depend on the type of record I'm referring to.

The operational characteristics of those records that will describe the information within the DNS. Well, in the standards they have identified a type of file format, which has been called "master file format," which is a standardized format, to define a flat text file.

Although, at present, there are many softwares that implement a DNS' servers that use relational databases or information storage mechanisms which are easier than a flat file. Although, the majority of DNS servers implement the master file format as standard, which stores the entire data on a text file. If I manage a zone, I will have at least one text file that we have all the information for that zone. I will, afterward, give more details.

If I administer three zones, I will have at least three text files, one per zone, each with the information of each zone. What I will never have is a single or the same text file with information for more than one zone. There is at least one per zone. There could be more but at least one.

So, these records, these information containers, are called “resource records” in English. These are repositories or storage facilities that store certain information with certain functionality. There are different types of records. I'm not going to detail them all. That is not the plan at all. I'm just going to mention a few of them that may be used as examples.

The ones here on the slide are some of the most commonly used records in the DNS. The A record is the one that is used for the main functionality the DNS was conceived for. Initially, the DNS was designed to convert a domain name into an IP address in the IPv4 version or later on into the IPv6 version.

So, the record that is used to store the information is an A type of record. The record type that is used to store information for the IPv6 address for a certain domain name is part of the A record. And then, there are other records such as the NS record that is used to indicate the authoritative name servers for a given domain name. Later on, we are going to go into further details about this.

Then, you have the SOA record that is used for storing some information that is not of use for the end-user but it is used by the DNS system itself in order to self-manage the Domain Name System.

The SOA record is like an [inventory]. It doesn't have to do with the items that I have in that store, but without that information I cannot manage the whole store, the whole repository.

CNAME is a record used for generating ALSes in order to establish an equivalent between a domain name and an alias for that domain name. If I wanted to have two domain names, when converted into IP addresses they should convert into the same IP address. Or if I want to have a domain name that points to the same domain name, I could use the CNAME record for that purpose.

The MX record is another essential feature of the DNS and it is used mainly to store the name of a mail exchange server for a given domain name. We usually, in a web browser, just enter a URL or a domain name. We access a website and we can see the webpages for that domain name.

When we send out an e-mail, you may recall that there are two portions in that e-mail address: what comes before the @ sign and what comes to the right of the @ sign. So, you have the user that you want to reach out with that e-mail. What you have to the right is the domain for the site where that user has his or her own e-mail address.

So, in order to be able to route the e-mail I need to know the IP address for that mail exchange server. Therefore, in the DNS, the record that stores the name of the mail exchange server associated with a given domain name is the MX record.

And then, you may have an A or an AAAA record that will translate that mail server name into an IPv4 or an IPv6 address in order to route that e-mail.

And then, you have a PTR record that is used for reverse mapping. That means that when you have a query in inverted order—that is, you have the IP address but you want to find out the domain name associated with that IP address—you use the PTR record for that.

And of course, that also has a special implementation within the DNS tree but I'm not going to go into all those details because it will take too much time. Right, then.

I just added this screenshot in case you are more curious about the number of records that exist, the different types, what they are used for, and the standard that determines the functionality for these records.

On this webpage, this is an IANA webpage where you can find a list with all the DNS records, what they are used for, and the RFC or the standard for the use of those records. Some of them indicate that they are obsolete and the MD and MF records are almost not used anymore. They have been replaced by the MX record.

So, in that site, you can find a list with all the DNS records and the defining standards. So, there you can find a very detailed explanation of all those records, and how they work, and what kind of information they store.

This is just an example of an A record. This is an example of the A record for the domain name example.com. This IP address, 192.0.2.7, is the IP address that I will have to use in order to enquire into that A record for the domain name example.com. If I were to run that query against the AAAA record, the system would give me back this IPv6 address that you see here on the screen, 2001:db8::7.

So, these are two independent records. One stores the translation into the IPv4 address, and the other one the translation into the IPv6 address. Here, you will find some examples of some traditional record types.

Toward the end of the presentation, we are going to talk about DNSSEC and you will see the extension of the Domain Name System and the addition of some record types in order to use different types of information—the signatures—in all that is necessary in order to implement properly the DNSSEC protocol. So now, let's go over … Are there any questions so far? I'm not looking at the chat.

SERGIO SALINAS PORTO: No, we don't have any questions so far, Nico. I will certainly shoot at you some questions afterward.

NICOLAS ANTONIELLO: Perhaps you are falling asleep or you are entertained. I don't know. Okay. Let's move onto the resolution process. Let's talk about the resolution process. What happens when a web browser, for instance, received a URL name with a domain name and you want to access a certain website? What happens then? The kind of magic that occurs there.

So, before talking about the resolution process, let me go over a few concepts that are important. The DNS is a distributed database, as we said. So, in order to resolve a URL, in order to find the IP address for www.nicolasantoniello.com.uy, I'm not going to run the query to a centralized site and the website will be able to download the web page.

No. There are some other steps because the system is not centralized. So, I will have to run a number of queries throughout the DNS tree until I get the answer that I'm looking for. In my case, it will be the IP address associated with a domain name of the web page that I'm trying to access.

So, there, you will find two types of servers involved: the recursive or resolver server and the authoritative server, the "name server" in English. The name server or the authoritative server usually answers queries. It doesn't pose any questions. It just answers queries.

Name servers are authoritative servers for a given zone – one or more zones. The name server for one zone contains all the information pertaining to that zone. So, that is all the information, all the records, for all types of records that I have mentioned and many more, related to information for a given zone will be contained in the same name server.

So, if I wanted to find that address, the IP address for nicolasantoniello, I will have to run a number of queries against the name server for nicolasantoniello.com.uy so that I can get to that address in order to access that website.

And the recursive resolver servers usually do not contain information. The resolvers are responsible for running the query through the DNS tree in order to provide an answer to the user. So, the resolver is going to send a query to find out what is the IP address for nicolasantoniello.com.uy and it's going to provide an answer.

So, this is a practical example. This is the usual scheme but let's look at the practical details. But before doing that, there is something that I need to say. In this specific example, I'm pointing at … Can you see the pointer? Can you see the mouse pointer on the screen?

SERGIO SALINAS PORTO: Yes, we can.

NICOLAS ANTONIELLO: Okay. So, this is a client. This is a mobile device that has a browser. Somebody has entered a URL containing a domain name. So, the browser needs the IP address for that server in order to bring that back to me.

So, based on that name, it needs to run a query against the DNS. All devices connected to the Internet as part of their operating system contain a function that is called “stub resolver.”

The stub resolver will receive the queries against the DNS system and will route that query to the resolver that has been configured for that device. Who determines the resolver that will be configured for the network of my mobile phone or my laptop?

My ISP will allocate the public address so that I can connect to the Internet and also allocate the IP address for the recursive resolver to which all the queries will be directed. So, the recursive resolver 4222 is the resolver configured for this mobile device and it will be the resolver that will receive all the queries against the DNS.

So, let me continue. Let me move on with this example. So, there is a query against the stub resolver and the stub resolver directs the query to the recursive resolver. So, I want to find out the IP address for www.example.com. So, this is sent to the recursive resolver. The recursive resolver will not have that information so I will have to find it out somewhere else, asking other recursive resolvers.

So, another name service. If this resolver doesn't have the information and has to resort to some other servers in order to get that information it will need the IP addresses of those servers, and we said that the resolvers have no information.

Actually, that is not entirely true. All recursive resolvers have at least some minimum amount of information. That is part of the software for the recursive resolver. The minimum amount of information contains the IP addresses of all the DNS root servers. Later on, we are going to talk about that, as well.

There were initially 13 different addresses for 13 different root servers. There is some technical explanation associated with that. We are not going to go into all the technical details but the recursive resolver needs to have some minimum amount of information to start that search: at least the IP addresses of the root server.

So, as all root servers contain the same information, it is going to take one of those, the IP addresses, and the recursive resolver is going to send a query to one of the root servers. In this case, it would be the L-Root server which, by the way, is one of the root servers that is operated by ICANN.

So, the recursive resolver will send a query to the root server. In this case, it is the L-Root server, but it could be the A, B, C-Root server. Any root server.

So, usually, you have letters to identify the root servers, all the way up to the F. So, here, this resolver has chosen the L-Root server. What is the query that is sent out? Exactly the same query that it received.

This is very important. The query that is sent out will always be the same query that has been received. They won't be different based on the servers they direct the query to.

Does the root server have the answer for that query? No. So, when a name server of a given zone—in this case, the root server—doesn't have the reply, it will try to give the best possible answer. All servers know what delegations have been made. That is, which are the children of the branches of that tree that we saw before?

So, the root servers are the [CDO] level servers, the root which we represented as a dot. So, the root servers know the IP addresses of the servers of all the top-level name servers: .com, ccTLDs servers.

So, the root servers are going to provide an answer back to the recursive resolver, and that will be the IP address of the name server for that .com domain. The recursive resolver is going to ask exactly the same question to the .com name server.

So, the question will be, what is the IP address for www.example.com? The name server for .com is not authoritative. So, for example, .com is only authoritative for .com, so it's going to provide the resolver with the best possible answer. Since it is a name server for .com, it knows the IP address of the authoritative server; for example, .com. So, the .com name server will provide back an answer telling the resolver which other servers it can ask the question to.

So, the resolver will ask the authoritative example.com name server about that address. And in this case, this is the authoritative name server for www.example.com and it will provide back the IP address as an answer.

So, the resolver will get that IP address that will be the accurate answer and will give it back. It will forward it to the stub resolver. The stub resolver will give it to the browser and the browser with the IP address will be able to display the website. That is the end of the resolution process.

So, let's move on. Another characteristic of the recursive resolver is that each time a query is received they have to follow this resolution process in order to get an answer. At the same time, it is very important to consider that name servers reveal their load and you also have the local cache.

The local cache is a local memory storage that the recursive resolvers have. Each time they receive an answer from the different servers, they push the query against who … They will store that information in their memory.

So, if somebody goes back with the same query then the resolver will not have to go out to get that answer because it already has the answer stored in its memory. So, it will not have to go back through all these steps. It will save time and it will also reduce the load on these servers.

So, the information is stored by the recursive resolver but it is not stored there forever because if there is a change in those names it will never be able to reflect those changes. So, it will just store that information for some time.

After some time, the information will be eliminated from the cache memory and it will have to go back through all these steps in order to get an answer in case it receives the same query.

As we can see here, there is an example. The same client is asking for the IP address associated with ftp.example.com. So, remember that the resolver already has some information in the cache memory so it's going to go straight to the example.com server.

So, the query is going to be against the authoritative example.com server. It will not have to go through all the other servers. So, it will get the answer. The server will give an answer back and it can forward the answer to the stub resolver.

SERGIO SALINAS PORTO: Nico, we only have five minutes left and I want you to be able to finish your presentation and also to be able to get some questions.

NICOLAS ANTONIELLO: So, now I'll conclude. So, now the idea is to talk about the robustness mechanisms. This shows the complexity of the system. We can see that this is distributed all across the Internet and the queries [count across] various autonomous systems. Responses go back and forth. There may be a server hosted elsewhere, and so on.

So, this system, which is distributed throughout the Internet, is quite extensive. So, the potential fault points or the target points in case of intentional attack, they are not individual or there isn't a single fault point.

In the DNS, there are several potential fault points, and here in this diagram we can see the entire ecosystem from the registrant, which is the individual who has registered a domain name; the registry, which is the seller selling the domain name to the registrant; the domain name operator that operates or manages the authoritative resolver; the communications between the recursive and the stub resolvers, etc. So, several mechanisms.

All these various mechanisms provide greater resilience, and one of these mechanisms is the so-called Anycast. Anycast is, basically, assigning the same IP address to more than one server.

I could have several servers, up to 25-100 authoritative servers for a given zone. I distribute them all across the Internet and they all have the same IP address. When they all have the same IP address and there is a query, the query will be sent to the nearest server in network terms.

So, I do not have all the information in a single server but distributed across several ones, all with the same information, and that splits the load between many/increases resilience. Because if any of those servers are attacked, there will be others that will not be attacked.

If I have an operational problem with any of them, I could shut it down and the queries will be routed to the other servers, and that is transparent for the user and for the network as well.

Another mechanism that is widely used in addition to Anycast. There are other mechanisms and this is the last slide. Here, you can see several mechanisms that are applicable to various fault points that increase robustness and resiliency.

One of them is DNSSEC that authenticates the source of a query. When I add the IP address, if it has DNSSEC implemented, I make sure that the reply is the correct one and there is not a replacement.

There are other protocols such as DoT and DoH, DNS over TLS or DNS over HTTPS, that complemented the DNSSEC at a different fault point, which is the communication between the stub resolver and the recursive resolver.

When the client sends the query to the recursive resolver in its operating system., DoT or DoH encrypt the information. If anyone captures the information packet as it is encrypted he will not be able to see the data or the source of the originator.

So, DoT and DoH provide encryption, DNSSEC validates the source of the reply, and there are other mechanisms and recommendations both for registries and registrars to increase resilience and security in the Domain Name System.

Well, basically, that was my presentation, which was actually summarized to a great extent. But nevertheless, I will send you my PPT slide and you have my e-mail address.

SERGIO SALINAS PORTO: Thank you, Nico. As we read in our WhatsApp group and other messaging systems we have, we are all very pleased having heard this presentation. I don't know if there are questions. We don't have much time but I'm sure there will be many. I have a question of my own. There's another one. "What part of the Internet traffic goes through the Domain Name System and what part does not?"

NICOLAS ANTONIELLO: What part of the Internet traffic goes through the Internet to the Domain Name System and what not? Well, all Internet traffic that has to do with the transfer or with the query of an IP address involving a domain name, all that goes through the Domain Name System.

So, nobody remembers, none of us, in any application, "to access a service write the IP address." What we do is we write a name. So, virtually all traffic or all points of access to all services provided in the Internet starts with a query to the Domain Name System.

Of course, the reply, once the access to the server takes place, the Domain Name System is not involved. It's not in that resolution. Let me give you an example. If I want to watch a Netflix movie, when I select the film and I want to start watching it, all videos have been identified by a domain name and I need the IP address of the video server. That's a query to the Domain Name System that returns the IP address, and once that is confirmed the traffic with the film itself is not on the DNS.

But the DNS is critical for the Internet not in terms of operating issues, but if the DNS shuts down the entire Internet will not work because nobody remembers the IP addresses. If I told you, now, there are no more telephone books, I delete all your contacts, what happens? Nobody can call anyone because nobody remembers anybody's phone numbers.

SERGIO SALINAS PORTO: So, after the mobile telephone has started this issue is more relevant?

NICOLAS ANTONIELLO: Yes, this is the critical resource because this is one way of finding resources in the Internet.

SERGIO SALINAS PORTO: Alejandro Pisanty has a comment to make. Alejandro, you have the floor.

ALEJANDRO PISANTY: What I want to say is that, although all traffic requiring names … The traffic of the Netflix stream or in any Zoom conference does not go through the Domain Name System itself. So, on DNS, you have IP address to IP address. Nevertheless, this is a confusion that many people have because most people believe that everything goes through the Internet and the DNS acts as a central Internet server.

SERGIO SALINAS PORTO: Thank you, Alejandro, for your comment.

NICOLAS ANTONIELLO: Exactly as Alejandro has explained. What I wanted to say is exactly that. What part of the traffic goes through the Domain Name System was …

It's making an analogy with a telephone book. Once I know the IP address, once I have found the number in the directory, that's the end of the participation.

SERGIO SALINAS PORTO: I have, actually, two questions now. The first is related to DNSSEC: how DNSSEC improves the Internet user's security? And the second question, if you can briefly talk about cache poisoning?

NICOLAS ANTONIELLO: Well, these are two things that are not necessarily related. DNSSEC provides assurance or security to name resolution as follows. Let me try to explain it by way of an example.

Imagine that every time—and that happens increasingly more—whenever we use the Internet to access various services, such as financial services, whenever I access the website of my bank, the bank where I have my money, and I make a transfer or I check whether my pay has been deposited, or any transaction on my bank website, what I do is I type the bank name on the URL browser box. So, that's a query on the DNS.

The server returns the IP address of my bank's website server and I have access to the IP address, and therefore to my bank account. When I make a query with my bank name and I get the IP address, what's the guarantee, what's the assurance I have, that the IP address I get is my actual bank's IP address and not someone that is acting as my bank?

If that’s even attempted to bank website … Someone has designed a similar website. I will probably give all my credentials and give access to my bank account to that person who has managed to commit that act.

This is not an easy thing to do but it can be done. I will see no difference, and by no difference I will not be able to see. This is not phishing. It’s not that the URL has changed the name or anything. If an attacker does that, the user cannot notice.

So, what DNSSEC does is, through a digital signature, somehow I can check the signature through a public and private key system and [trust and core] system. I'm going to explain this very quickly, but with this system the person receiving the IP address associated to a name is sure that it comes, actually, from the truthful person.

It’s not someone that has set up a parallel website. It is digitally signed by the owner of the domain name. I can verify the signature and make sure that I'm actually accessing my bank’s website IP address and not another IP address of an attacker that is not my bank. That’s very important.

And for that to be successful, all authoritative servers that are queried in the resolution process must have their domain names signed with “DNSSEC,” and the recursive server provider has to provide a DNSSEC resolution, as well. So, it would be great if we demanded this from our ISPs and from recursive resolution providers – those who manage the zone. And the other question was?

SERGIO SALINAS PORTO: Yes. The other question was cache poisoning.

NICOLAS ANTONIELLO: Well, you remember that when the recursive resolver made a query, it received a query, stored it in its cache just in case someone has made the same question, the same query. So, in very general terms, what attackers do, and succeed, in cache poisoning: what they do is make a recursive resolver store in the cache something that is not true.

For example, through any means, I get my recursive resolver that I have configured to do DNS queries stored in its cache. I am actually not accessing any authoritative server but the local recursive resolver that initially has no information, and I make it store in its cache an association of my bank's website with an IP address that is not the bank's IP address but my fake website.

So, when a query is made to the recursive resolver on what is the IP address, the recursive resolver will not go through the authoritative route but already has it in its cache. It was injected with a poisoned reply. So, I will give that reply to the user, which is not right, which is wrong, and the user will access a wrong site.

Again, if DNSSEC is well implemented with all configurations proper, and there are additional precautions implemented for recursive resolvers, cache poisoning, although in some cases it might happen anyway, if I verify the signature, well, that signature should provide a wrong result, no verification, and the reply would be rejected, and there will be an error displayed, and the user will know that the site is not accessible because the reply is invalid.

SERGIO SALINAS PORTO: Thank you so much, Nico. We have to move on with our agenda. But before that, thank you so much, Nicolas, for being with us today. There are many requests for a second opportunity to invite you to have you with us. So, probably with Adrian Carballo and the Capacity Building Group, we will again go back to you for a second round.

NICOLAS ANTONIELLO: Yes, I will be pleased to do that. On the topic of this presentation we could continue, it would be great if we had more time for questions and also for comments from those of you, such as Alejandro, who have years of experience so that you can enrich what we are saying. But yes, we might need at least a couple of hours for the detailed version of this, to go into further detail.

SERGIO SALINAS PORTO: Well, there will be a second chance and Adrian will be contacting you, as well as Anahí and Monica. So, probably they will contact you to arrange a second round.

NICOLAS ANTONIELLO: It's a real pleasure to share with the LACRALO community and you know where I am. You have my e-mail address. So, anything you need for specific training or any other topic, if it's not me I can find the expert in ICANN for a presentation. So, thank you very much, again, and thank you for your time. Here I am for anything you might need.

SERGIO SALINAS PORTO: Thank you, Nicolas. Now, let's continue with our agenda. Let me go back to it very quickly. Just a second. If you can please, Claudia, scroll down? Now, we have a regional exchange on DNS abuse. We have implemented this with Harold as an additional way to go into detail on the knowledge in the region. Harold will lead this phase. Harold, you have the floor. You have six minutes for this.

HAROLD ARCOS: Thank you, Sergio. Can you hear me?

SERGIO SALINAS PORTO: Yes, we can hear you. Go ahead.

HAROLD ARCOS: I would like to share some information with the region. Rather than being the leader for this phase, I would like to remind you of some information. This is the phase that has resulted as a process that we followed over time to use our monthly meeting to discuss topics of interest and find agreements, and disagreements, consensus, and try to have some statement on the topics that are being discussed within ALAC.

So, the plan is just to spark the discussion. We are not going to find solutions to everything here. We are just going to share our views, our positions, and then we can continue working on this document.

We will post it on the proper space and we will continue the discussion over the mailing list. We will try to come up with a position statement for ALAC. Claudia, can you please share my screen? I'm getting a message indicating that I cannot share my screen. Can you enable that for me, please? I would like to have that.

SERGIO SALINAS PORTO: Harold, you can share your screen because you are a co-host for this meeting.

HAROLD ARCOS: Yes, I'm trying to do that but, for some reason, I cannot activate the option to share screen. Well, Claudia, perhaps we can have a blank page there. We will write down the information. Sergio, I'll give you the floor so that you can run this discussion and give the floor to the members of the region who would like to make contributions about the topic of DNS abuse. This is one of the key topics included in ICANN's strategic agenda so it is very important for us to share our views on this.

SERGIO SALINAS PORTO: Thank you, Harold, but we don't have enough time. I don't know if there are any immediate contributions. Perhaps somebody can make a brief comment now? Any of you who'd like to take the floor? I don't see any raised hands.

VANDA SCARTEZINI: Perhaps we can continue this discussion through the mailing list.

ALEJANDRO PISANTY: I think it’s very important to talk to experts on security about the use of local names. I hope that in the next few weeks we can gather some information after contacting some of these experts.

SERGIO SALINAS PORTO: Thank you for your contribution, Alejandro. If there are no other suggestions, I move that Harold take note of this and, given the fact that we only have 21 minutes left in this call, and since we still have to go through a lot of topics, I would suggest that we continue this discussion over the mailing list so that, now, we can move onto the next agenda item, LACRALO participation tools. Harold, is this what you were talking about?

HAROLD ARCOS: Yes, that’s right. Can you hear me, Sergio?

SERGIO SALINAS PORTO: I can’t hear Harold. Okay. So, let’s move on with our agenda. We have the directors of the different working groups in the call.

HAROLD ARCOS: I was using the mic.

SERGIO SALINAS PORTO: I'm sorry, Harold. Go ahead.

HAROLD ARCOS: Today I wanted to show you how to access the Wiki space and how the ALSes could work with their members, in order to send a request, in order to get a user created so that they can participate in the Wiki space.

SERGIO SALINAS PORTO: There is a lot of echo, Harold. Perhaps you are connected to both devices at the same time.

HAROLD ARCOS: Okay. Now, it’s better. I just wanted to share my screen so that you could see specifically what are the different sections within the Wiki page that we can use for participation purposes. Well, once again, I cannot share my screen. So, we will continue this over the mailing list.

SERGIO SALINAS PORTO: Okay. Now, we will start with the working group reports. I’ll give the floor first to Adrian Carballo so that he can give his report on behalf of the Capacity Building Working Group, and Sylvia Herlein, and Jose, and all the directors. I don't know if anybody else would like to talk. Adrian, you have the floor.

ADRIAN CARBALLO: Thank you, Sergio. I just wanted to let you know that I have already sent the preliminary draft on the ICANN capacity-building program. This is just an approach in order to get to a final version, in order to reach consensus with all the ALSes and with the other directors.

The goal is to have a first version of the course running by March next year once we have reached some agreement with all the participants. I am very pleased with the webinar we had today with Nicolas' presentation. It was highly informational. This is a very interesting topic to all of us so, as Sergio said, it would be nice to have Nicolas coming back to this call.

With ICANN Learn, let me go over the topics for the courses that we are considering: Internet, Internet security, combating cybercrime, among others. So, we are working on almost 20 different topics for ICANN Learn courses. I hope that I can soon have a briefing for you on the outcome of this process.

And just a point about the Southern School of Internet Governance. We were planning to launch that call in Buenos Aires but, of course, with all the pandemic outbreak, we have postponed that for October. So, I'm glad to take any questions you may have.

SERGIO SALINAS PORTO: Thank you, Adrian. Let me check whether there are any questions. There are no raised hands, so let me give the floor to Lilian. I think Lilian was not feeling well. Lilian is there, okay. Go ahead, Lilian. You may leave after your report.

LILIAN IVETTE DE LUQUE: Good afternoon/good evening to you all. Let me be brief because I have a terrible headache. As I told you in the last meeting, we are working to meet a short-term goal that we established for our communications working group. We wanted to have a newsletter and we have been working toward that goal.

We have been trying to compile information. We still need to collect some more information. I have requested that information from some other colleagues. So, we are hoping that we can have a working meeting with all the members of the group.

It would be the second meeting of the year. I think that the group has to adapt to the new circumstances. We need to find a way to work in a coordinated manner with the Social Media Working Group for the coming meeting of Kuala Lumpur that will be held online.

So, these are the three goals. We want to gather all the remaining information so that we can launch our newsletter. We want to plan another meeting for the working group, and then we want to have a meeting with the board in order to present all the information that we plan to publish.

And then, we also need to determine how we are working to coordinate our efforts with the Social Media Working Group for ICANN68. Thank you. That will be all. Any comments?

SERGIO SALINAS PORTO: Thank you, Lilian. Can you please send me a private message indicating who owes you some information so that we can conclude with that?

LILIAN IVETTE DE LUQUE: Okay. Thank you, Sergio.

SERGIO SALINAS PORTO: Are there any questions? I don't see any hands. "Multilingualism on IDNs;" Sylvia Herlein has the floor.

SYLVIA HERLEIN LEITE: Good evening to you all. I just wanted to tell you that we are ready to send our LAC TLD/LACRALO project. We wanted to take the opportunity of working together with LAC TLD and the fact that, now, we have access to all the LAC TLD members. With one of the co-chairs of the group, we decided to take this opportunity to enhance the questions in the survey to include some questions about Universal Acceptance.

[Graeme] has opened this new section in the survey. I had a conversation with Harold yesterday and really need to thank him because I didn't know that he was such an expert in surveys.

So, we worked together on the survey, finalizing the survey in order to make sure that we have all the necessary questions there and, also, to determine how we are going to use the results of those surveys so that we can make the most of the surveys.

So, next week, we will launch this project. You have heard me talking about this project several times so I'm going to update this information on the Wiki page this week. That's all I have to share.

SERGIO SALINAS PORTO: Thank you, Sylvia. I think that now I am going to give the floor to the Governance Working Group. As you know, we have completed the rules of procedure and we are waiting for your comments. We have a Google Doc.

The “Rules of Procedures” document has been sent to the mailing list and you have up to May the 22nd to post your comments. We will gather all those comments, and we will assess them within our working group, and we will come up with a final proposal that will be submitted to a vote in the region. I see that Sylvia is asking for the floor, so Sylvia and then Vanda.

SYLVIA HERLEIN LEITE: Just a quick note to let you know that I would like to appreciate Sergio and Harold’s efforts and also the efforts of all the rest of the Governance Working Group in all the sub-regions because you have worked very hard on this document.

It has a lot of detail; it talks about LACRALO governance mechanisms, how votes are going to be counted, how quorum is reached. So, I want to compliment you on this wonderful piece of work that you have done. It is available as a Google Doc and waiting for your comment. So, thank you for this tremendous effort.

SERGIO SALINAS PORTO: Thank you, Sylvia, for these kind words. It is true, it has been a lot of work. All the members in the sub-regions have worked hard but we always try to see how, from a regional standpoint, we could strike a balance. So, of course, there is always room for improvement, but we were just trying to reflect, here, how the region actually works. Vanda, you have the floor.

VANDA SCARTEZINI: I just wanted to remind the Governance Working Group that I sent some petitions and some comments that came to my mind. I hadn't realized that we needed some more items included with regard to the processes. So, perhaps Sergio can give me an answer or we can have another meeting this week just to discuss whether we are going to include those other items or not, or how we are going to go about them.

SERGIO SALINAS PORTO: Thank you, Vanda. Let me say this publicly. You are part of the authors of these rules of procedures. After May the 22nd, we will hold a call with all the working group members. We are going to gather all the comments received, all the contributions, and probably some comments will be included, others won't.

It's not because we want to be whimsical but just because, sometimes, those comments cannot be included now because it is too late and we have already gone through that process. But we will try to reflect all viewpoints. We might leave some of them aside.

So, it is important for you to go over the rules of procedure, share your views, make your contributions, so that we can finalize the document. Thank you. So, please put down your hand.

Let me check. It’s my turn, now, according to the agenda. We need to talk about the elections and a regional update. As you know, we are starting an elections period. We have elections for the chair-elect position. It will be a position for one year, and then the chair-elect will become the chair.

There are three candidates for this position and I want to congratulate two members, Vanda and Tracy, as ALAC members who will join us after November in their new positions. We know that all of you can make very significant contributions, so I just want to acknowledge your incorporation and I'm looking forward to your contributions. Then there is another topic and that has to do with the coming elections. Perhaps Silvia can refresh my mind. Can you talk about the timings of this?

SILVIA VIVANCO: I hope you can hear me. Basically, next Thursday we have a call for the candidates to the chair-elect position. We have three candidates. They will make a presentation and take questions from the participants. This is Thursday. We kindly invite Tracy and Vanda to be there in this candidate call.

SERGIO SALINAS PORTO: Sorry, Silvia. I forgot to say that I have made Vanda know that this was happening but not Tracy, so I'm saying it now. It would be very good if you are in this call because you can tell us what you can do and how we will work this year, both Vanda, who has already made a comment on our mailing list, and also Tracy. We want him to be there to make a brief introduction in addition to the candidates for the chair-elect position. Go on, Silvia.

SILVIA VIVANCO: So, that is for Thursday. On Friday, there will be an e-vote, an electronic voting, from Friday 22nd up to Friday 29th, ending at 23:00 UTC. This is for the position that is being in competition, for the chair-elect.

These positions, after the count, the winner will be declared who will hold office in ICANN's AGM in 2020. In the case of the chair-elect, he or she will take office immediately because the current chair-elect has resigned. I apologize because I was reading the wrong calendar. So, in the case of LACRALO, the current chair-elect has resigned. That is why.

SERGIO SALINAS PORTO: Let me make a comment here because the work for our chair-elect and the workload they have is huge, both in LACRALO, in ALAC, in ICANN in general. They require a significant amount of time, and although the work is performed by a team this team has to be led by chair-elects and chairs. So, we have to be prepared for that. Thank you, Silvia, for your input. The last point, I don't know if Harold is with us. Are you there, Harold?

HAROLD ARCOS: I'm here, Sergio.

SERGIO SALINAS PORTO: So, the last item of our agenda is the Remote Strategic Plan. So, Harold, if you can give us an update on this?

HAROLD ARCOS: Thank you, Sergio. This is a plan that was presented by Rodrigo Saucedo from the GSE Team for Latin America and the Caribbean. It's a cycle of webinars. We have been asked that each sub-region should establish contact with the academia. We have several representatives from the academia in our region. Why? It's because we want to see our four sub-regions coordinate a direct engagement for the next 12 months of the fiscal year.

In order to hold webinars and digital meetings, that falls under the remit of LACRALO's Capacity Building Group and will be part of the outreach work in these times of contingency where everything is being held remote, online. So, we are kindly inviting each sub-region to fine-tune our calendars with the Capacity Building Working Group for coordination. Thank you, Sergio.

SERGIO SALINAS PORTO: Thank you, Harold. The last point, the last issue, is the evaluation poll. If you are so kind as to answer this survey, Claudia or Silvia? I don't know who is going to read.

CLAUDIA RUIZ: Well, you can read it on screen. The first question is, what was the timing like of the webinar? You can start answering this. The second question is, how was technology used in this webinar? Are you voting? Are you answering? Because I do not see any answers yet.

The next question is, has the speaker shown mastery of the topic? The last question is, are you pleased with the webinar? Are you satisfied with the webinar? The fifth question is, how many years of experience do you have in the ICANN community? And the last question is, what topics would you like us to cover for future webinars? Here, it's not a simple answer. You cannot choose and answer but you can write it in the chat.

SERGIO SALINAS PORTO: Thank you, Claudia. Well, we have come to the end. We have gone beyond by only two minutes but I'm going to write, now, on the chat, that I'd like to see a second round on DNS. Let me say it has been a pleasure to share this call with you.

There is a lot of work ahead of us to do. There are elections in LACRALO. It's very important for all of us to participate. We can make our input to the rules of procedure. The working groups are eagerly waiting for you to join them.

So, let's continue working to improve our participation and our discussions on ICANN policy. Happy Internet Day, again. It was yesterday, but for me, it's an Internet Week. See you soon. Thank you so much.

[END OF TRANSCRIPTION]