KOBE – Tech Day (2 of 3)
Monday, March 11, 2019 – 13:30 to 15:00 JST
ICANN64 | Kobe, Japan

EBERHARD LISSE: Okay, everybody please settle and sit down. We are now starting with the afternoon session and Shinta Sato from JPRS is going to talk a little bit about preparing for the disaster. They occasionally, which we will note later on, have got serious earthquakes and other issues here. So they have a plan, and it's always good to look at how other people prepare for disasters.

SHINTA SATO: Hello, everyone. I'm Shinta Sato from JPRS. Well, actually this was told that this is a host presentation, but it's just a Tech Day item. So my presentation is about .jp and .jprs, preparing for the disasters.

There's a little bit of a background here. There's natural disasters in Japan. Well, it happens sometimes. Not often. Maybe often, but it does happen. The upper left one shows the Great Hanshin-Awaji Earthquake which happened in 1995. It was very large, and actually the place of the earthquake it was just right here. Kobe was the very center of the earthquake.

In the lower right maybe many people know this one. In 2011 we had a Great East Japan Earthquake. That happened in March 11. That is today. This is the memorial day of this earthquake.

These two were very big, large [earth] disasters we had in the recent days. There are many others, some kinds of natural disasters, but this was a huge one.

So what happened? Mainly in the network part, in 1995 the earthquake made buildings and highways collapse. Roads, railways, lifelines, those were very damaged.

In the network situation, it was launching days of the everyday life of the Internet. The Kobe government and universities, they were the early adopters for the World Wide Web and they were making information to the public, and they were making information to the public through the Internet. But the main information sources for the users was still the TV and the radio and those kinds of things. That was the days of 24 years ago.

But eight years ago in 2011 the disaster was very huge [inaudible] similar to the last one. Actually, there was additional to the nuclear power plant disaster, but that's not in this case.

The network situation, this was very changed. Smart phones and SNS, the safety confirmation service, these were widely used on the Internet. Actually, some lives were saved via the information from the SNS. So the effectiveness, Internet was very effective to have the information through the disaster case. So it's being recognized in those events.

So .jprs, we are the ccTLD and we operate the DNS. So we are in the very core place of having the critical mission of the Internet.

So this is just some very brief introduction what we have [done] preparing for the disasters and also what we've done for after the disasters.

One is distributing the DNS server locations. It should not be in the same place. Actually, not only placing the serves, but we conducted some experiments for putting more local DNS nodes in the local regions as well.

We constructed a DR site. Registry systems and office systems both need a DR site. And conduct training for emergency situations. We need to establish headquarters and those kinds of things, and we need to make the test of the responses of the DNS shutdown as well.

What we've done, we've helped domain name registrants who lives in the place of the disasters. Dispense [with] the renewal fee of domains having registrant address of the affected areas. Details will be shown from the next slides.

First is about the DNS servers. Well, .jp DNS servers, it has eight names from A-H, DNS .jp. The location, the old days the location was all in Japan. But right now it is using the Anycast and with the cooperation with other organizations and actually JPRS it's ourselves, we conducted to make the many locations of the DNS servers, some in multiple places in Japan, some worldwide, and those kinds of things. For the .jprs, they are also Anycasted as well.

We conducted demonstration experiments about the continuous Internet services. A little bit on the background of the network in Japan. The Internet resources are very concentrated in two places, Tokyo and Osaka. Not only the network connection, data centers and hosting services, and those kinds of things are very concentrated in Tokyo and Osaka.

So what happens in disasters? Natural disasters may cut the link to Tokyo and Osaka. What happens to the local places is that their networks will be isolated within its regions. So no connection to Tokyo, no connection to the others worldwide as well. So we have conducted the test to put the local DNS server nodes to the domestic regional ISPs and joint research we are using the .jprs environment.

.jprs is an R&D platform. It is not a ccTLD. It's a gTLD operated by JPRS. So we are operating this place for the experimental environments that we can learn lessons from incidents, errors, and failures, and those kinds of things. The outcome of those results should be used to the .jp and the global community as well. However, the SLA of ICANN is very strict for the gTLDs, so we cannot do so much of those kinds of tests. But this is the place we use for R&D.

I'll tell a little bit about the joint research itself. We made a research with local ISPs as I said before. The goal is to keep the continuous access to the Internet resources in each area. What we provide is the continuous DNS resolutions. That's what we targeted, by local nodes of Anycast TLD DNS or .jprs.

We used one of the name servers which was operated by JPRS, but at the event of this research we placed local nodes to each ISP and conducted the test. The eight domestic ISPs participated in this research. These ISPs cover designated geographical areas without overlap in Japan. Because these ISPs were the subsidiary of regional electricity companies, their customers are similar to the customers of the electricity companies. And actually, they have very robust infrastructures for the power line and the data center and those kinds of things. With these eight domestic ISPs and JPRS itself collectively covered whole Japan. So this was a very good test [and] organizations collaborated.

This is a sample of the results. We simulated the loss of Tokyo and Osaka connectivity and saw the effectiveness of the local nodes. It's very simple. Number 1 is the normal state. The second one is partially disconnected. In that case, may queries go to the local nodes, but not all. But if everything is disconnected, all the queries go to the local node as just we thought of. And if recovered, queries go back again. It's very simple, but we need to test in the real environment and learn what we have. And there were some findings in these tests. More details are available in this place, but I think that's in Japanese so sorry about that.

That was the DNS. Next is the DR site. We constructed DR sites in two places in Japan. One is [inaudible] Tokyo and one is in Osaka. As I said, not only the registry systems but office systems as well. They are placed in both Tokyo and Osaka. The registry system, that is active in Tokyo data center. And Osaka data center, the office system is active. The JPRS office itself is mainly in Tokyo. That is to avoid the simultaneous failure of both systems at once. We don't want both registry systems and office systems together shutdown in the disasters, but one is active. So we only need to activate one other one. That's what we do.

A little bit about architect is that the architect should be very simple. We place the same hardware, same functions in both data centers. Very similar network we created and real time data synchronization in the database. But not only the database, but there's some file based synchronization as well. But the main is the [inaudible] the database.

As we are providing the DNSSEC, we need to put the KSK and ZSKs, private keys, to both data centers. That is delivered by the secure transportation service we are using. The site failover, that is very difficult. It's still partially automated but it's not fully automated. And it is triggered by hand operation as well.

That was the registry system. Also, we conducted training. Be prepared to take immediate action in events. Several times per year, each for different purposes. This is some of the menu we have done. One is building structure of emergency headquarters. Defining the decision makers and launch teams for many things. One of for walking to JPRS office. We tried to walk to the office and check the route, which road we can take, and share the security of the facilities, and potential dangerous places we looked for and public facilities as well.

And we tested for the failover of the DR site. A little bit different from the disaster, but the cyber attack response training, these are similar types of training we are using and doing.

There are some pictures for the training. Launching emergency headquarters. We assume that the power outage will have come in the disaster, so this is very dark. Decision-making board members. Board members gathered in the headquarters and we shared the information to have the decisions [inaudible] space. Walking training. There are very dangerous places. If you walk to the office, there might be some places like this one.

Other ones, pictures of training. The DR site failover. JPRS office is very near to the data center. So once something happens, we rush to the data center by walking. There is a small office space, and the failover operations are done in that place. Those are the training we have done.

The next one is what JPRS has done once the disaster happens. Dispense with the renewal fee of the domain names. In March 11, 2011, the disaster affected area was very wide and we worried about registrants in the area. They may not make the renewal operations for themselves. So what we did is that we made the renewal fees free to affected registrants for one year.

However, issues were find in identifying the target domain names and registrants. We wanted it to be done in automation [inaudible], but we couldn’t. The operations of checks and adjustments were done manually. We used WHOIS information. However, it was useful but it was not in some part. Something like changing the city names after the registration, that was not reflected [fully] so the old address was shown. And there are multiple ways of describing the same address in Kanji or [inaudible], and that was a very hard thing we faced.

This would be the last slide, the future works. We are going to continue training and DR site system improvement and those kinds of things. We prepare for the cyber attacks toward the large event is the 2020 Tokyo Olympics. And there may be cyber attacks and those kinds of things, so we are preparing for that and now we're working for that as well.

So this is the end of my presentation. Any questions if you have?

EBERHARD LISSE: Thank you very much. This was very cool. In particular, I liked to say test the walking way to the data center so you know where you're going to and you know where this works and so on. That's a point that is probably often overlooked, especially if you have a larger organization where staff members change. So it's always now they know this is what happens, here's the handbook, go there.

Patrik?

PATRIK FALTSTROM: Thank you very much. Patrik Faltstrom, Netnod, Sweden. We are doing similar work in Sweden and have been doing very similar investigation and just delivered a report a few months ago. I have a question regarding the involvement of the ISPs that you were talking about. You mentioned that they were involved and each one of them were covering different geographical areas. And because they were originally subsidies of energy companies, they had very robust networks each one of them. So do I understand you correctly that it was the network between the ISPs or between the ISPs and Osaka and Tokyo that you investigated? You did not investigate partitioning the ISP's network? So the ISPs network was still working as a whole. That was not partitioned, right?

SHINTA SATO: Well, the ISP networks, we didn't investigate about those things. But those ISPs, they connected to Tokyo and Osaka, just to those places and not connected with each other so much. So if the connection to the center is cut down, they will be isolated even if the other places are still connected.

PATRIK FALTSTROM: Thank you very much. Because when we did this investigation, we found that one of the weak points regarding partitioning networks in the country was partitioning of the control [plane] of the ISP. That's because how fiber is built in Sweden, that is a weak point. But it seems that you don't have the same problem. Thank you very much. I will probably contact you offline to compare the results of our studies. Thank you.

SHINTA SATO: Okay. Actually, there may be many connection to the ISP itself, but this time [they] focused into this kind of connectivity.

WARREN KUMARI: Warren Kumari, Google. First off, great presentation. And as Eberhard said, it's really common that people forget stuff like how do people actually get to the DR site. Would you mind going to Slide 9? That one. I might have missed it when you said it, but I didn't quite get how you simulated the loss of connectivity in this. Did you actually drop connections to those ISPs, or did you just pretend that you didn't get those routes or based upon IPs? How did you actually simulate it?

SHINTA SATO: It is simulated because we couldn't cut the real fiber line or those kinds of things. But actually, we blocked the traffic maybe using the network [gears].

WARREN KUMARI: Okay, so you just dropped their traffic and see where it goes?

SHINTA SATO: Okay, one more question.

UNIDENTIFIED MALE: My name is [inaudible] from [inaudible]. From the presentation, the registry in Osaka and Tokyo were both active. I want to find out if in your simulation you tried a situation where a registration on a particular domain name is tried at the same time in both locations.

SHINTA SATO: Well, actually we do not activate both locations. Are you asking like that?

UNIDENTIFIED MALE: No, from your slide, your registry one in Tokyo and one in Osaka were both active.

EBERHARD LISSE: He already answered this question. They're not both active at the same time.

UNIDENTIFIED MALE: Oh, okay.

EBERHARD LISSE: They have got the registry system active in one and the DNS on that location as a standby and the other way around.

UNIDENTIFIED MALE: Okay, thank you.

EBERHARD LISSE: All right, thank you very much. That was a very nice presentation. I very enjoyed it.

SHINTA SATO: Thank you.

EBERHARD LISSE: Tom Barrett, there we go.

TOM BARRETT: Great. Hi, everybody. My name is Tom Barrett. I'm founder of an ICANN-accredited registrar called EnCirca which I founded back in 2001. We are a specialty registrar focusing really on partnering with registries. So we focus on restricted and regulated TLDs. We provide validation services to those types of registries. We build white-labelled storefronts. And more recently, we've started to help integrate some of the ICANN TLDs into blockchain [inaudible].

Why do we care about blockchain and Internet of Things? In some way, these are the biggest innovations that are emerging in the

Internet today. And certainly, blockchain is probably the best invention someone quotes and says since the invention of the Internet itself. In terms of venture capital involvement, this is Internet of Things investments, about $1 billion last year. Blockchain was much bigger, so about $4.8 billion, about 375 deals in terms of venture capitalists.

That's important to note because, as you know, venture capitalists like to have an exit. They tend to be impatient. So some of the things that since both blockchain and IoT require DNS to work, it also requires some new features that might not exist in the DNS today. So the question is, will they wait for those features to emerge, or will they create their own? And, of course, what role will ICANN play if, in fact, they start creating their own types of DNS protocols?

So Internet of Things. The world is exploding with devices. We're talking about physical devices. Not only cars, drones, but refrigerators, stoves. So Internet of Things has to do with connecting all those physical devices to the communication network like the Internet. There are constrained devices that, frankly, don't have the computing power to support the typical TCP/IP stack. Those are sensors, bar codes, etc. So they have special requirements in terms of how they will connect to the Internet.

Blockchain. How many folks here, by the way, own some cryptocurrency? Just a show of hands. Okay, so you're familiar that bitcoin and other cryptocurrencies are built on top of blockchain. Blockchain is an open, distributed ledger that will record transactions between two parties in a verifiable way. As you add transactions to the ledger, they add their blocks to the blockchain.

One of the key parts of this blockchain is that typically once your transaction has been added, it's very hard to change or remove that transaction. The idea is it's permanent once it has been added to the blockchain. So that creates some great applications.

My favorite actually is real estate. Every time in the U.S. when you buy a home, the conduct a title search to see who all the previous owners were for that home. Because they want to make sure that no one will come back later and say that land was stolen. So they conduct a title search, it's called. Now if that home sells every two years, they repeat that title search every two years and go back perhaps to the beginning of time as far as the ownership of that land.

So blockchain is a perfect application because once you've done that extensive search, you can put it into the blockchain. And the next time that home sells, you only need to add that last transaction because all those other transactions are permanent and accessible in the blockchain.

My perspective here, by the way, is an ICANN registrar, so I sell domain names. So what we see emerging from Internet of Things and blockchain is that they need a naming infrastructure that is very similar to the DNS. So the question is, how does DNS advance to support the Internet of Things? And in blockchain, will ICANN stay compatible with these advancements?

DNS is relatively old. Invented in 1983, that's 15 years before the birth of ICANN. It's advanced over the years. For example, security was not part of the original spec, but it's continually evolving. They've added DNSSEC. They've added IPv6. They've added IDNs, etc.

IPv6 is of particular interest because not only does it address the fact we're running out of IPv4, but it enables some new applications in Internet of Things. Because we're going from 4.2 billion addresses to 3.4 x 10^{38} so a trillion, trillion, trillion IP addresses are now possible with IPv6.

It remains to be seen, however, what happens to the DNS when you go from IPv4 which is billions of addresses to trillion addresses. Does it scale? A lot of these devices require zero or auto configuration. The availability and performance of a lot of these devices. Imagine you're trying to track drones in space to make sure they don't collide. They can't accept the same type of latency that you might have with a web browser. So it's slightly different requirements as it scales.

DNS, of course, isn't standing still. So there are evolving RFCs to support some of these needs, such as auto configuration, service discovery, and so on. But again, the main question is, this is all theoretical at this point, and what does happen when we start to add another few billion or 20 billion or 30 billion devices to support Internet of Things?

So there are venture capitalists out there already developing frameworks and technologies to address this need, and they're outside of the normal realm of DNS. One example we'll talk about which is of particular interest is geo-fencing. The idea is that we take the earth, the world, and we map it into a three-dimensional coordinate system. That allows us, for example, to track folks moving through a mall. It allows us to make sure we understand where drones are in our space and keep track of what's going on in the physical world.

Imagine a sugar cube. In fact, imagine that we took this room and filled it with sugar cubes. Then we assign each one of those sugar cubes an IPv6 address. Then let's extend that to the entire world, and we have no created a 3D coordinate system for the entire world based on basically a cubic centimeter or whatever the size of a sugar cube.

This is actually an application that is coming out later this year, and they're using the .PLACE TLD as a way to map those sugar cubes to IPv6 addresses. Something that would not have been possible with IPv4 because of the scarcity of IP addresses.

And the other thing to note. For those of you who have been around for awhile this is not necessarily a new idea. The first round of ICANN TLDs back in the year 2000, one of the applications was for .GEO. This was SRI International, and they were proposing something very similar. Although, it's not clear how they would have done it back then.

Let's talk about blockchain. Again, what we're seeing in blockchain are some new naming infrastructures that look like the traditional DNS but they're not moderated or regulated by ICANN.

There are a lot of applications [on] blockchain. I'm just going to talk about the blockchain wallet application. As you can see here, the use of blockchain wallets is exploding. There's about 35 million blockchain wallets as of the end of January 2018.

Here's your typical blockchain wallet. It looks like a 40-plus string of characters. With this wallet, this is where I will store my cybercurrency. If I want to, for example, sell a domain name or a refrigerator to you, if you want to pay me in cybercurrency, you need to know my wallet to send me currency.

If only there was a more friendly way to tell you what my wallet was, right? Ethereum, the leading blockchain out there, created something called the Ethereum Naming Service (ENS). Again, modeled after the DNS, they created this infrastructure to offer a secure way to assign human-readable names to those wallets. They even created their own TLD called .eth. They use terms like "managing the ENS root" and they have "root key holders." It looks like they're building their own version of ICANN.

After 18 months, how did they do? I've created a new acronym: BTLD (blockchain-only TLD). It's not a ccTLD, it's not a gTLD. They did 300,000 registrations, $28 million in deposits. They even sold one of these names for $3.5 million, although it probably was paid for in cybercurrency.

So compared to the new gTLDs, they would have been – whoops. My graphic didn't come through, but this was a list of the top 20 gTLDs. So they would have ranked in the top 20. Again, totally outside of ICANN regulation.

So what we have is this naming scheme on the blockchain where I can take, in this case, a .eth domain name – and I'll use my company name as an example – and map it to EnCirca's wallet.

So if you want to pay me in cybercurrency, all you need to do is know how to send money to EnCirca.eth. And if I want, I can create my own TLD. So instead of EnCirca.eth, send it to Tom.EnCirca. Now I've got a dot brand on the blockchain. So blockchain is solving a problem here for digital wallets that DNS solves for websites, all outside of the ICANN world.

So what's different about blockchain? Clearly, there's no ICANN. There's no Internet governance feature. Users have more control over their names. It's tougher for intellectual property or governments to come after them and shut down their particular domain name. And of course, it enables payment of cybercurrency.

So that's not the only one. There's a bunch of alternative TLDs out there, all live today. I talked about .eth which came out two years ago. This year you'll see .zil, you'll see .crypto. And there's a bunch of others out there as well, all alternative TLDs, all outside the regulation of ICANN.

What are their "benefits"? I put benefits in quotes because a) a lot of these are theoretical and b) obviously we have a lot of the attributes of DNS today is not static. They're continually evolving and improving security, etc. But they're claiming the blockchain is more secure. It's more private. It's censorship resistant. It's more scalable, etc.

This is the difference that we're seeing right now. As a registrar we're asking, should we be selling these domain names? But they're obviously very different between the blockchain TLD and the ICANN TLD. So there's no Internet governance, as I mentioned. Every node on the blockchain side is equal. There's no hierarchy like the 13 or 20 root zones that we have, the root servers. Owners have more complete control over the records so they can't be shut down by their government. And again, the theory is they're less prone to hacking than the current DNS. Certainly, without the hierarchy, they don't have any man-in-the-middle type of problems that you might see.

But here's the hitch. They only work with a plug-in. So again, we've seen this before, New.Net back in the early 2000s. There are some workarounds that are happening, .XYZ and .LUXE which are ICANN TLDs have integrated into Ethereum. There's another ICANN TLD coming out later this year, .PID.

But here's a question I want to leave you with, two questions. What would happen if a dominant market force like fill-in-the-blank were to enter the cryptocurrency market on a massive scale and begin accepting cryptocurrencies as a payment method? It might be WeChat. I know they had some problems last year in China. It might be J.P. Morgan who has just announced they're planning to introduce their own cryptocurrency. Facebook rumor has it will also be introducing their own cryptocurrency. Alibaba. Amazon which controls online shopping.

What if they decided to support cryptocurrency? And what if they decided it wouldn't be .eth. It would be their own dot brand on the blockchain. You would instantly enable those TLDs in your browser via a plug-in and the consumer wouldn't know the difference. They wouldn't know whether or not ICANN was governing this or not. There wouldn't be, for example, GDPR. There might not be WHOIS. There might not be any intellectual property protections, etc. But the consumer is pretty unable to discern if this is an ICANN TLD or a blockchain TLD.

So it gets back to ICANN's mission, a secure, stable, and unified Internet. And the question is, how are they going to adapt to blockchain? Do they bring it under its umbrella? Do we create a blockchain NSO just like the GNSO? Do they get to appoint board members and have someone on the NomCom? Or does ICANN try to attack it?

Let's continue the conversation. I'd love to hear from you folks. This is my e-mail and LinkedIn profile. And I thank you for your time.

EBERHARD LISSE: Thank you very much. That was a very quick and fact-filled introduction in a short period of time. We have got one question if there is one. Are you asking a question or just walking around? Okay. Good. There is no question. The address is on the agenda. The address is also in the presentation. All presentations will be on our website. Thank you very much.

TOM BARRETT: Thank you.

EBERHARD LISSE: Okay, now we will hear from Chuan Guo about Alibaba's cloud DNS practices.

CHUAN GUO: Okay, good afternoon. My name is Guo Chuan from China. I'm a [DNS] engineer from Alibaba, and this is my first time to attend the ICANN public meeting. I'm the ICANN [inaudible] for the Fellowship, so I would like to use this chance to share the Alibaba [inaudible] with you.

First, who we are. I would like to talk more [inaudible] this question. Who we are. Who am I? This is a very old question. We call ourselves the cloud DNS [inaudible] like Google cloud DNS. [inaudible] DNS protocol we have the root server, we have TLD, we have the SLD, and we have the [registrar], we have the [registry].

But we didn't have a name to describe an entity like us, the cloud DNS, because I think the inventor [inaudible] DNS protocol in the last century 1970s I think, and they assumed that everyone who would have the ability to access the Internet will run their own DNS servers. But in the 1990s there was rapid development of the Internet there was a lot of change in the technology. And one [inaudible] many of the applicants or the domains wouldn't run their own DNS server.

So at first it was the registry provide a [free] DNS resolution for the domain owners. But afterwards the domain owners asked for high quality for the DNS resolution service so [inaudible]. So we have [inaudible] more specialized DNS resolution service provider. So we are DNS service provider and the DNS [inaudible] designed for [inaudible] naming problem of the Internet. But we still have our own name.

I think how do you describe who we are, it's easy to understand what we are doing. And [inaudible] DNS resolution service, I call this the SLD authoritative server. And the local resolver server, we use the local resolver server to provide DNS resolution service for the Alibaba cloud computing environment. And the public resolver server, we have our own public resolver server that's like Google for [inaudible].

This is some [statistics] for our authoritative server. How many SLDs in our cloud DNS? There are 14 million domains. In fact, there are more about 40% websites in China have their DNS configuration in our cloud DNS. How many queries per day? There were 160 billion per day. Compare this to [data] we can find that many domains that are registered and never used in fact. And security attacks happen every day.

Our goals. I think the goals are similar with yesterday in this room with the RSSAC's goal to the root server system. We are stable, fast, safe, and because we are [faced to] end user, so we have some customize.

For the fast goal, this is our brief [architect], our authoritative DNS servers for the [inaudible] the client configure [inaudible] in the portal of the [inaudible]. We wanted to make it fast between the user change the configuration its effect here in the end [inaudible] point we can figure out a way we try our effort to make it in one second which [inaudible] key point in the data distribution [inaudible] and there will be many network problems. And we cannot guarantee that 100%, but in most [inaudible] we can hit enter in one second the data will be [inaudible].

Okay another fast measure. Our authoritative server are developed based on the DPDK architect for the performance, and the [pop] [inaudible] use the Anycast around the world. This is friendly for the Internet user.

Okay, this is stable. For stable reason we [must] focus on the [inaudible] [strong] and the data recovery and the database [inaudible] backup. The two point is the [public] cluster management and the data recovery [inaudible] one point. And the data consistency, we try our effort to make it all data in all the [pop points] are consistent. In fact, this is some kind of difficult problem. This does stable.

And safe. The security is always the problem. [inaudible] security for [inaudible] user logins [inaudible] more and more DNS security [inaudible] in the portal [inaudible] user information content passwords are stolen. We'll have a very [inaudible] configuration [inaudible]. Our own [system] security, this is what [inaudible] solve in our team. The interface security problem [mainly] for the DNS-related network attack.

This is a customized example which is not a [based on] DNS protocol, but this is the [inaudible] requirement. So we can add the weight field in the [inaudible] [card]. If there have three [inaudible], we have two or three will return four x and the one [inaudible] will return the 4 y. So the DNS will have the function to load traffic management.

Okay, the next is the local resolver. We use the local resolver to provide DNS resolution in all the Alibaba cloud data centers [and] the traditional architect cache and forwarder. But the cache, we developed a kernel module for the performance. And in the local resolver we also have some customized configuration data like the [inaudible].

The public resolver, our public resolver address. But we try to have the public resolver and we guarantee their speed fast and stability. And we didn't do much work in data analysis in the public resolver. So we know that the data center is very sensitive in there because more and more people didn't run their own resolver so they would use the ISPs resolver or public resolver. So it's more and more important, especially the dataset is very valuable in the [inaudible] DNS. Okay, the second sentence should be [inaudible] from the latter page.

One reason we are not a specialized security team, and we think that the dataset in the public resolver is so sensitive. So we [inaudible] to the data analysis.

Okay this is DNS in the private cloud. Besides the public cloud, [sometimes] we provide the private cloud service. In the private [inaudible] strong need for the stability and a high availability. So we [are trying the] new architect [inaudible] ETCD structure and the [backend] database and provide the stateless API. And in the private cloud environment [inaudible] servers run the Anycast network [root].

So I think this is a brief introduction of the Alibaba DNS technology. And in the future, this is our first time to attend an ICANN public meeting. I will [probably] to attend other community more and more in the future. We compare the other entities in the community, I think we have the advantage to directly connect with end users. They may have new requirements of the DNS, and we can give that to feedback to the community.

And [inaudible] have two problems in my personal thinking about DNS. One is the security [inaudible] too much I [inaudible] implementation and have influence with [inaudible]. So I in fact in earlier time I have thinking this problem. DNSSEC I don’t think there [inaudible] push to the end user. At most, it will be pushed to the public resolver. So the [inaudible] HTTPS will resolve the client user and the security channel between the client user and the public resolver and the DNSSEC will provide the security channel with the DNSSEC with the resolver and the authoritative server. So who if [inaudible] guarantee the security in the public resolver or the resolver provided by the [inaudible]? There will be still a problem, I think.

The second question is from the invention of the Internet, we have the e-mail system, we have the World Wide Web. They are all based on DNS and they are always the most popular. [inaudible] sometimes I even think that if we have reached the limit of the application based on the Internet technology. So I have a strong interest to find the next killer app. And of course, I would like it to use the domain name system also.

Okay, I think that's all. If you have good idea or a question about our cloud DNS service, feel free to contact with us. Okay, thank you.

EBERHARD LISSE: I have one question from the chair. How many users to you have? How many people or individuals or companies or whatever are making use of this DNS?

CHUAN GUO: Our cloud [inaudible] users. You mean how many users? With that I can tell there have [14 million]. The user [inaudible]. And about 40% [inaudible] in China. And there [inaudible] another [inaudible] provided by the [inaudible] which also [inaudible] DNS service provider.

EBERHARD LISSE: 140 million scales much differently than the 2 million inhabitants of my country. Anyway, anybody else from the floor? Thank you very much.

CHUAN GUO: Okay, thank you.

EBERHARD LISSE: Okay, next one is the SSAC standing presentation. Tim April will hold it, ably assisted by Rod Rasmussen.

ROD RASMUSSEN: All right, good afternoon, everybody. We are here to talk about – not we. Mainly Tim. I'm here to just provide a quick intro and frame things, and Tim's going to dive into it. Many of you are aware of the recent high profile domain name hijackings that have gone on that have been in the press and affected government servers and things like that. So there's been a whole bunch of interest in that.

One of the things we're interested in, in SSAC is taking a look at that issue. We've had a lot of publications in the past discussing various ways to protect against these kinds of things. But we thought, given the notoriety and the sophistication of some of these things, we'd have a session here talking about that and perhaps even taking a look at some new work.

So anybody not familiar with the Security and Stability Advisory Committee in the room? Good. We've done our PR work well. I'll skip this slide. Here's the agenda. I'm just going to give a little background. And then Tim's going to go into the rest of this and actually give you some sage advice as well as talk about some of the technical aspects.

And then if Jay's here, he may – ah, there we go – we will sit up here and have a discussion because that's one of the things we want to do is get some feedback on this because we're contemplating work and we have some questions for the community as well.

So on the recent domain registration hijacking incident[s] just at a higher level, there was a series of attacks where the attackers were able to modify domain registration records at the registry typically via some sort of compromised login credentials. That could have been done in many ways because, as you know, credential management is a tricky thing and there are many ways of getting ahold of credentials. Parts of the attack were attributed to some specific malware. They also used a DNS tunneling thing that was kind of cool. But anyway, there are multiple ways of doing this.

The net of that though is the attackers changed the DNS records, name server records and A records for those in particular to use the attacker's DNS servers, which then gives you control over those. And then one you have that control, they were able to impersonate various services.

So it's not just the typical thing you might think of, of putting up a fake website and make a phishing scenario. They were actually using it on other services like e-mail. And then obviously in that position, you can do a man-in-the-middle attack to intercept that traffic. And that allows you to slurp up certificates, access credentials, all kinds of goodies that could then be used for subsequent attacks.

So this was a multilayered attack where various underlying infrastructure providers were not the end goal but they were used and attacked and compromised in various ways in order to get to this setup where these attackers could very cleverly set up these man-in-the-middle attacks.

And by the way, those man-in-the-middle attacks were done over short periods of time to make it harder to detect them as well. So that makes it really difficult for standard monitoring to be able to find that because they didn't just deface a website or something very obvious and siphon off all the traffic or what have you. They very cleverly put these things up for a small amount of time, slurped up things, and then put it back to where it was.

So I'm going to turn that over to Tim to walk through some of the things you might do about that.

TIM APRIL: Most of the stuff that's going to be covered that we're going to talk about [in a couple] minutes has been talked about to no end, and many companies talk about this every year for their security awareness training and things like that. We've also got a few publications we'll list at the end where we've gone into depth on even more of these. But most of the advice we have here is for securing the registrant-to-registrar channel. It can also apply from the registrar to the registry and within the operators of the registry themselves.

First of all, credential management. Registrant credentials are one of the most critical for maintaining the security of any zone. While the attack that Rod was just talking about was more sophisticated and it was targeting different levels in the registration process, every day there attacks where domain registrant credentials are phished, compromised in some other way, taken from some data dump that's on the Internet. And it's used to modify a registrant zone.

So the key things to that are strong passwords everywhere you can. If your service providers offer multifactor authentication, you should always use that. If they don't, consider switching the provider that you're using or ask them for it and maybe they'll even give it to you.

Another key part of this is e-mail security where whenever – the password reset function is a more often abused thing. Where if you can get the e-mail account but you can't get the password for the registrar account, they can send a password reset and then get control of whatever they want from yours. So don't forget the e-mail password is a key factor.

As I was just saying, multifactor authentication or two-factor authentication is becoming more and more prevalent with all sorts of different services. Many e-mail providers will let you do that now, and I'm starting to see more registrars offer it as an add-on, often a free add-on, where you can sign up for that. You can use your phone as the authenticator for it, and then you can prevent some sorts of attacks from being fruitful for the adversaries.

When you're considering using a multifactor authentication system, there's been some guidance from the national standards institute in the United States where they've ranked multifactor methods from most – they provide some advice of how to use it, which multifactor methods are better than others.

At the top right now is the universal second factor (U2F). This is commonly implemented by [U company called Yubi] where you get a physical token that provides you an additional layer of authentication. And then going down the list you can go to time-based or hash HMAC-based one-time passwords where you can get an app for your phone or an actual physical token like the old RSA keys where you can use that to authenticate.

And then there's SMS and phone-based verification, but there have been recent attacks where there's a thing called SIM swapping where an adversary can call up your phone company and pretend to be you and give them a new SIM card that they're saying is your phone and get your cell phone number swapped to theirs and take over whatever entities you're using your phone as the second factor. So at this point, SMS and phone-based verification where they call you have been suggested to never be used whenever possible. They're still better than nothing, but if there's any other method for second factor, you should prefer that over those two.

When we were building these slides, we went back and forth about what advice to give people for password management. And then we resulted in just going with the trusty old XKCD method of it's really hard to come up with a random complex string of characters that you're never going to remember how to type. But you can use an arbitrary set of words that have some meaning to you but not to other people, and that often provides a sufficient level of security. This sort of thing won't help you if someone phishes you or you share it with someone else. So keep password hygiene in mind in that case.

Basic dos and don'ts for passwords. You've heard this many times before, I'm sure. Use a strong, unique password where strong is kind of a weasel word where I'm not going to go and define what actually it means for you.

ROD RASMUSSEN: Not "password123."

TIM APRIL: Yeah, not "password123." Not your cat's name or anything like that. Password managers are up and coming. And in my case, I only know one password to get into my password vault, and then everything else I have pretty much never seen before. That way they're strong, independent, and unique and I never share them across websites. So if one gets compromised, I'm not going to think, "Do I have to go change my passwords everywhere else?" Using MFA like we said.

And then, never share your passwords with anyone. There are some minor exceptions to that. If you have to have a roll account for your registrar in your company where you have three people that manage your domains but you can't get individual based accounts, you may want to use a shared password manager for that sort of thing.

And then never reuse your passwords across multiple accounts. Just a couple weeks ago, there was a large dump of about 2.2 billion user name and password pairs that was released on the Internet. I searched for 10 or 15 people that I knew, and they all showed up in this database with passwords that they were still using. Mostly from large recent breaches that have come up.

Now we get on to e-mail security. As I was talking about a little while ago, e-mail accounts tend to be the front door into a lot of things. That's where you get your password reset notifications. That's where you get your e-mail from your bank. It's also where all the phishing comes in. so e-mail accounts have to be treated very carefully. And as we were saying before with passwords, don't reuse them everywhere, anywhere.

And then if you are sending e-mails, so if you're a registrar or a registry in the room, please strongly consider using DMARC with SPF and/or DKIM for any e-mails you're sending to consumers. And if you are a receiver of e-mails, so if you run a mail server or you're the person making the decision for what e-mail server your company should use, you should see and try and enforce the use of verifying DMARC and SPF whenever e-mail comes into your organization. This helps verify that the e-mail is coming from who it says it is rather than just accepting anything that comes across the Internet.

And then there's one point. I've seen many companies that have had this happen in the past where if you looked at the WHOIS information before May of last year, you may have seen that people would use Gmail addresses or Yahoo mail addresses for their critical domains for their company. This is usually a bad idea because if someone gets into your personal Gmail account, they can go and wreak havoc with your corporate domains. Or if that person leaves the company, they may not be required to give you that domain back. So you may lose control of corporate assets.

One of the links you'll see in the bottom here in a little bit suggests using a roll based account within your organizations where multiple people on a mailing list get e-mails related to the domain changes that are going on so that if one person leaves the company or is on a plane or something like that, you can still get in touch with the registry or the rather, depending on how it works.

Some more mail security tips. If you're operating a mail infrastructure that is available outside your corporate firewall, strongly consider adding some protections for man-in-the-middle. So using MFA, requiring TLS, pinning TLS certificates if your mail clients allow that so that you can know when the certificate changes. And then also protecting your users from phishing and other sorts of attacks that may come over e-mail. So spam filtering. Some companies offer feeds that show servers that are sending a lot of phishing e-mail and things like that.

Moving on to your domains themselves, while not all registries offer this service, registry locks are very helpful in protecting your domain. For the people that aren't aware, there are two different types of locks that are available in the domain registration systems.

There's client or registrar locks which are things like if you log into your registrar, you may be able to click a button saying please lock my domain. That's a registrar lock. That's usually a free service that's added by most registrars where they'll try and add some level of prevention from changing your records. It helps prevent updates to the record, domain transfers, deletes, or renewals in some cases.

There's also the other side where the registry can lock where the registrar and the registry have an agreement where whenever the lock needs to be applied the registrar will initiate a communication with the registry to turn on the locks.

This process can take some time. Some of the SLAs are in the three-day range where you're unable to make any changes to your domain during that time when the registry locks are on. But this also prevents the attacker going and turning them off or doing something to your domain. Not all registries offer this. Check with your registrar. See if it's offered by the registry. But it does come most of the time with an extra charge for you.

One of the things that the SSAC has been talking about in the last couple weeks is possible understanding of how to help standardize this process across registries because it may vary dramatically based off the TLD that you're dealing with.

In the recent attack, there was a noticeable difference in the effectiveness of the attacker for zones that were signed using DNSSEC. This doesn't mean it's the silver bullet. It's something that was just some [inaudible] data about this attack. But this is a good point to say you should consider signing your zone. It requires also that your users be using a validated resolver. That's becoming more prevalent. Some of the larger open public resolvers are doing validation for you. But if you're running an end user resolver, consider adding or turning on validation.

This helps with integrity protection. It does not help you with availability protection. So you could still have a man-in-the-middle could be dossing your queries with DNSSEC, but it's verifying that the integrity of the zone is ensured.

DNSSEC signed zones also acted as a canary for this recent attack where there were some indications of the attack going on by DNSSEC failures. So if you start to see DNSSEC failures on your zone, this may be that you forgot to resign your zone and the timer elapsed on it or it could be a sign of an adversary interacting with your zone. You can use tools like DNSViz which is now hosted by DNS-OARC to try and identify what's wrong with your zone if you're seeing something fishy going on.

This is also a place to pitch the DNSSEC workshop that's happening on Wednesday if you'd like to come and learn more about DNSSEC.

Then getting into the more tricky pieces of this, there are some intricacies with how your zone relies on the other TLDs on the Internet. There's a website that's hosted currently by Verisign of trans-trust.verisignlabs.com where you can go and enter your domain name and it will show you the transitive trust that you instill in the DNS for your domain. So it will show you that maybe your TLD name servers are in a handful of different other TLDs, so you're relying on the security of those TLDs to ensure your domain security. We don't have any specific advice on this at this point, but it's something to consider looking at in the near future.

As I was saying before, there's nothing here that's a silver bullet, but monitoring your domains is key. So if you're a domain operator, monitor your DNS infrastructure. Check the logs. Make sure no one is logging into your machines without you knowing about it. Monitor for the liveness of your domain servers, and monitor your DNS zone entirely. So check to make sure that the records that are in your zone are what you expect there to be, and make sure that the delegations from the root zone and the TLD are directing your end users to where you expect them to go.

EBERHARD LISSE: In about half a minute, I would appreciate if you could all rise, remove your headgear if you wear one. And I think now is the time we can do that and read the statement that should take a minute so we [inaudible].

Thank you very much.

TIM APRIL: So beyond monitoring your zone and all the delegations and all that, monitoring the WHOIS for information about your registry and registrar locks can also provide some sort of information about if someone is trying to attack or target you.

And then the more tricky things to do are certificate transparency logs. In the most recent attack – or in the recent attack. I can't say most recent. There were some certificates that were issued that were not requested by the actual domain owners. This is because the certificates were validated using what's called domain validation where if you control the server that's pointed to by the DNS, you can get a certificate for the name.

Some of the triaging that was done during the attack found that there were certificates that were issued that were valid for some of the [mail] host names that were used to then man-in-the-middle traffic going to the end users were sending. So monitoring certificate transparency is a very important thing to try and figure out who is trying to get certificates for your zone.

And then monitor for DNSSEC validation errors, like I was talking about before. And then monitor the name server records in all of the levels of your domains. So from the root all the way down.

Some relevant SSAC publications going back as far as SAC040. There were more, we just ran out of space on the slides. A lot of these are still very relevant advice for anyone that's either operating a zone or runs a registration system to try and protect both you and your end users.

And then the conclusion. This is not the first time that this sort of attack has happened. This definitely will not be the last. Security in depth can be helpful for this sort of situation. There is no silver bullet, holy grail, whatever phrase you want to use in that place. Securing the credentials for your registrar, your e-mail, your registry, your end users are key to protecting your assets. If you have the option to use MFA, do so. If you don't, try and request it or possibly consider switching providers.

E-mail address security can be a very useful tool in preventing these sorts of attacks. Deploy DNSSEC domain signing and validation and then registry locks in addition to registrar locks which should be a no-brainer for registrar locks. And then monitoring of your infrastructure. Very open, broad subject and not very easy to define but that's the best advice we have at this point with the current environment.

ROD RASMUSSEN: Thanks, Tim. I'm going to put this one out here which is more the heat one which I'm a little bit easier to take. But we really want to have a conversation around this. We're up here saying things for the most part that a lot of the people in this room already know and we've said things about in the past. This is a lot of review.

I mean, there were some interesting things here around the use of certificates and things like that and the DNSSEC signaling that went on in these attacks and actually in some cases prevention of damage. So those were all some really interesting bits to this. But we're in a situation where this is a lot about the hygiene and the things that we all know we should be doing anyway.

So as this series of attacks has shown, the adversaries out there have gotten fairly sophisticate and understand the DNS sometimes better than the people running parts of the DNS it seems from the way they were able to manipulate it and time it and all that kind of things. There was some very sophisticated timing on these things where they were able to understand when different things would maybe kick in where somebody would see something and they quickly were able to anticipate that and not be detected. So a lot of these things weren't detected until well after the fact.

So our adversaries are getting smarter. So isn't it time for us as an industry in general to step up our game and think about some operational standards, either best practices or some sort of requirements around adoption? I put that out there as a point to spur discussion, not as something we're necessarily recommending. But it's a situation where we want to make sure there's trust in the overall system. And then you have attacks that are going to the extreme of going after TLDs and the TLD operations in order to get to a target farther down the DNS stream, it's time for us to pay attention or more attention than we have been.

So with that, we'd love to take questions. Go ahead.

TIM APRIL: More actual details about the attack if you're interested will be talked about on Wednesday at 11:00 in Portopia Hall. The "Coming Up With Best Practices to Improve Security in the DNS Ecosystem" if you're interested.

ROD RASMUSSEN: Yeah, and that's an important session where the same question and conversation will probably happen. But we'd like to have it here with all you folks at Tech Day because I think you may have some really valuable input.

EBERHARD LISSE: Thank you very much. Let me kick this off from the floor. We use [inaudible] tools like many other small ccTLDs. We mandate two-factor authorization, and we find that our registrants regularly lose their cell phones. So we institute a one-time replacement, second time [charge] replacement policy that has resulted in much more attention being paid to that part of the thing.

What we actually see is that some registrants are systematically [polling] our EPP system. It looks like random names, but when you sort them alphabetically they're systematically [polling] it. So that's not really a drama. We then turn them off and sort them out.

We haven't forced our registrars to use [PGP]. I think that's one thing that you can easily use because if they lose their credentials, that's not so much of a problem as long as you have a known e-mail address.

We find however that being a small registry that we get e-mails, "Can we change the credentials?" "Yeah, sure. It must be coming from your registered e-mail address." That usually helps a lot. If they're nice, we tell them what their registered e-mail address is. If they're not, they can figure it out for themselves.

But generally speaking, we haven't seen any real phishing on our registry attempted, but we're a very small operation, family-run business. So all three of us doing this, we communicate very much with each other and we read the [roll] account e-mail. So all of these things that you're saying which were targeted against end users instead of registries, also apply to registries.

We have seven minutes, so I can take these three questions easily.

YOSHIRO YONEYA: This is Yoshiro Yoneya from JPRS. Thank you for raising this issue to the public. I cannot attend Wednesday [inaudible] so I give you some comments regarding this. I think we have already a lot of security mechanisms to protect our domains, but I believe that we don't have good document or textbook for end users to learn about the security mechanisms they should use. So I think accumulating some successful prior work within the industry and share such kind of text and translate it into the local languages and distribute to the end users would be a very good practice to do among us.

EBERHARD LISSE: My view on this is that you will not be able to sort out 3 billion end users. When have even any one of us changed configuration of their mail program of their browser recently once. Very rarely that you actually go to. It must go from the operators. In the ccNSO or ccTLDs we don't have mandated regulations, but most of us try to go by what the contracted parties have because there is at least some standard that we can abide by even though we are not mandating it. But to go to the end users is a good idea, but they don't understand this. My Internet is not working, and that means Facebook doesn't work and so on. My Internet is not working is something different to my wife than to my son or to me.

ARTEM GAVRICHENKOV: Artem Gavrichenkov, Qrator Labs, also here as a Fellow. This is not a question, more a comment. Thank you for the quick outline. If anyone were to look for a definite guide on passwords, I would happily recommend them to read NIST Special Publication 800-63B, Digital Identity Guidelines, especially the Section 5.1.1.2 Memorized Secret Verifiers, also known as passwords. I would personally be happy seeing such or similar requirements enforced at registrars. That's it. Thanks.

ROD RASMUSSEN: The last two comments kind of were related. I think one of the things that we're seeing in various spaces is an adoption of higher level controls. If you think about a lot of the social network sites and things like that, they're now offering multifactor authentication of the types that we're talking about. So from a consumer education perspective, they're getting more and more used to seeing these things on any web service they're seeing.

So I think there may be some work there from the perspective of, let's say, something that's facing an end user which will be a registrar being able to adopt some of those strategies that we're seeing in those places that are more consumer facing. Because they make it a lot easier to approach how to do this, how to manage it, things like that. So I think there may be some lessons that can be learned in other spaces that are having to do this to protect their user credentials and adopting it in the domain registrar community.

ARTEM GAVRICHENKOV: Completely agree. Thanks.

EBERHARD LISSE: Okay, one last, and be brief. You're preventing us from going on a break.

RICHARD ROBERTO: Okay, I apologize.

EBERHARD LISSE: No, no, you don't have to. I'm just joking.

RICHARD ROBERTO: Richard Roberto from Google. In the presentation you mentioned I think that the SSAC is interested in pursuing standardization around registry lock. Is that correct?

ROD RASMUSSEN: Yeah, that's a potential thing for us to look at.

RICHARD ROBERTO: How does one, say, stay informed about that or register interest in that process?

ROD RASMUSSEN: Well, we're going to have our public meeting.

TIM APRIL: Talk to your coworker.

RICHARD ROBERTO: Okay.

ROD RASMUSSEN: And then you have somebody that happens – do you want to wave?

TIM APRIL: Wave for him.

RICHARD ROBERTO: Okay, thank you.

ROD RASMUSSEN: Yeah, that. But just in general for the rest of the audience, we do have our public meeting Wednesday afternoon I think it is.

TIM APRIL: Yes.

ROD RASMUSSEN: Yes. So it's on the schedule. I don't remember the exact time. And we regularly update the community on what the things we're actually working on are. So we aren't taking that work on right now. It's just one of those things under consideration. But if we do, we'll be announcing it.

RICHARD ROBERTO: Okay, thank you.

EBERHARD LISSE: Okay, after we give them a big hand, we'll meet here in 15 minutes at quarter past.

[END OF TRANSCRIPTION]