ICANN85 Mumbai | CF – DNSSEC and Security Workshop (1 of 3)
Wednesday, March 11, 2026 – 13:15 to 14:30 IST

KATHY SCHNITT

Hello and welcome to the DNSSEC and Security Workshop, part 1 of 3. My name is Kathy, and I'm joined with my colleague Dan, and we are the participation managers for this session. Please be advised that this session is being recorded and is governed by the ICANN Community Participant Code of Conduct, the ICANN Expected Standards of Behavior, and the ICANN Community Anti-Harassment Policy.

To help us keep the discussion moving smoothly, here is how you can get involved today. The speaking queue. To keep the queue fair for everyone, all participants, including those here in the room, should use the raise hand feature in Zoom. When speaking, please state your name and affiliation for the record and maintain a moderate pace.

For questions, you may also submit any questions for the speakers today through the Zoom Q&A pod. I will read those aloud as time permits and when directed by the moderator. While the chat is available for attendee interaction, please note that the comments or questions posted in the chat will not be read aloud. Official questions must be raised through the speaking queue or submitted through the Q&A pod.

Now I'm happy to hand the floor over to Mark Elkins.

MARK ELKINS

Good morning, everybody. My name is Mark Elkins. I'm part of the DNSSEC and Security Workshop Program Committee. We normally meet once a week to discuss the sessions, et cetera, that are going on, et cetera. So, this is the intro. I'm remote. I'm currently sitting in Pretoria in South Africa. So, it's 9:46 here where I am in the morning. So, not a bad time from my point of view to start a meeting.

Anyway, the normal process is we have the first presentation, which is going to be all about numbers. And Eric Osterwell is going to present that. That's the job he normally does. So, I will now hand over the floor to Eric. Thank you.

ERIC OSTERWEIL

Thanks. Hopefully everyone can hear me. And welcome to the DNSSEC and Security Workshop. So, as Mark just said, I'm Eric Osterweil. I'm a distinguished scientist at Verisign. And I want to start off by acknowledging and thanking the program committee who puts this organization together, puts this workshop together, Gautam Akiwate, Steve Crocker, Mark Elkins, Layal Jebran, Russ Mundy, Ram Mohan, Peter Thomason, [inaudible – 00:02:50], Dan York, and with the wonderful support from ICANN orgs, Daniel Gluck and Kathy Schnitt.

And also thank the support that's given for the organization from SSAC and the Internet Society. And without further ado, the

DNSSEC, DAIN and RPKI deployments around the world, or previously known as COUNTS, COUNTS, COUNTS.

So, first is our standard map of the DNSSEC global deployment in country code top level domains to represent areas around the world where DNSSEC is being deployed. What's noteworthy in this picture is the color coding where when there's been announced that there is an experimental deployment of DNSSEC, a country code top level domain, ccTLD, is marked orange. Light yellowish is used for when a TLD has announced that it is planning to do DNSSEC. Darker blue is when DNSSEC is beginning to be deployed. Lime green is when the DSSEC deployment has a corresponding DS record in the root. Darker green is when that ccTLD is accepting delegations for zones below it, secure delegations to DNSSEC zones below it.

And finally, the lighter blue is when DS automation is in place for a ccTLD. And this shows what is currently the status around the world. It's available at maps.dnssec.gmu.edu. These are views of that last global picture broken up into areas. So, in this area, we can see Asia, APNIC, and some areas of Europe. This is a view of more northern Europe.

This shows Central and South America, North America, and Africa. Just seen a lot of deployment recently. And here's a video showing time-lapse view of DNSSEC deployment from the beginning in 2005 through to roughly today.

And now just a slightly different view showing the size of the deployments around the world over a corresponding period as a heat map. Taking a different view of the DNSSEC deployment, looking at DNSSEC validation from the site.

I'd like to thank the hardworking students at George Mason University for the previous set of slides and as well as the database. But looking at results from APNIC's labs, what we can estimate is the global validation rate. So, as of the end of last month, the global validation rate is continuing to tick up less fast, but at currently estimated 35.14% of global resolution traffic.

And a tabular view generated the previous value from the same site showed that roughly half a million additional samples were taken in the period between the last presentation in October of last year and when these were taken at the end of February.

Looking back again at the DNSSEC deployment globally as estimated by the number of DS records and this data taken from dnssectools.org shows that there's continued to be an increase in the deployment, roughly 795,000 increase of DS records seen in the route since October last year. And now a slightly different view looking at the DANE deployments around the world.

This is one of the main use cases for DNSSEC that impacts users. What we usually show is the number of zones that deploy DANE for SMTP for mail, and that number has actually decreased. It's gone down about 85,000. But what's new in this presentation is taking a view at the number of zones that are actually using the records and

that number has continued to increase. In fact, has increased rather rapidly.

So, this is one of the wrinkles of DNS that a lot of us are very familiar with. And so, the number of records may not reflect the number of zones that are actually using those records. I can have a bunch of records from my zone. And they may not even actually be in my zone.

In fact, frequently they're not. So, if we look at the actual number of zones that are using DANE SMTP as shown on the dnssectools.org, the value has increased by over 11,000 between October of 2025, the last workshop and now. So, still success.

Switching gears to the RPKI, the Resource Public Key Infrastructure used for routed resource certification on the Internet. We can see that the route origin authorization counts are continuing to grow across all the RIRs. Some slight upwards and downwards ticks from one RIR, but for the most part, most of them monotonically increasing.

Looking at the coverage of prefixes by the ROAs, what we can see is that what we would hope to see in the deployment, that the number of valids, meaning RPKI is covering a routed prefix, is growing. And now about 61%. And that the corresponding number of not founds, those are not covered, is dropping as more gets covered, less is uncovered.

And looking at the invalids where there is a disagreement between what is certified and what is being routed, that rate is now very low. It's less than 1%, 0.9%, which is great. And then one additional thing to throw out there, as RPKI continues to gain traction, route origin validation is a process by which ROAs are considered in the routing system. And one view maintained by the site antd.nist shows the rate at which a ROV has changed for a given ROA, where it's gone from, for example, being not found to being valid, which is great.

So, looking at where proper functioning deployment can be seen measured by the ROV or cases where valid is not valid anymore, what does it turn into? So, if it's turning into not found, that could be just simply someone turning it off. So, as RPKI starts to become more mature and more relied upon, and starts to join other security contexts like zero trust architectures or something, this is a nice way to look at how well it's functioning and nominally whether it's moving towards the point where it's ready to embrace technologies like OCSP or whether that should be applied, etcetera.

But this is a new view that seemed like it was now that RPKI is reaching this level, it would be good to include this potentially in the COUNTS, COUNTS, COUNTS. And with that, I'd like to sort of give time to the rest of the presentations and pause and see if there's any questions.

MARK ELKINS

Okay, we do have one question in the question-and-answer section, Eric. It's from Om Prakash Sharma. Sorry if I've not said that properly. His question is, what does the white legend in the DNSSEC deployment map represent?

ERIC OSTERWEIL

The white legend in the DNSSEC deployment map. This one, I assume, is this is the color code. If this is what you meant. This is the color coding I was mentioning before. Let me get back to the -- this color coding shows, you look at this map and it's got a bunch of bright colors and each of the colors represent stages that a ccTLD is in the DNSSEC deployment path.

So, where someone has announced that they're planning to do something, they're going to just experiment with DNSSEC, that may or may not be measurable externally. And so, the DNSSEC deployment maps has cataloged since 2005 cases where someone said that they're for their ccTLD planning to, you know, experiment with DNSSEC.

When they've announced that it's going to go live, when it can be seen to be functioning with DNS keys and RSIGS, etc. And when there's a DS record in the route that points to that ccTLD and when it's operational, when it's actually got child zones that it's securely delegating to and accepting new registrations, new delegations.

And finally, when it's deploying the DS automation like CDS or CDNS key records in order to make things to keep the its own

delegation hierarchy automated and more-smooth. I hope that's what you're asking.

MARK ELKINS

Okay, unless there's a further comment on that, I'll assume that that's fine. There is one other short question from the chat, is that we keep hearing that DNSSEC protects the internet from fake websites and DNS attacks. But years later, many websites still haven't switched it on. Who is responsible for making sure this actually happens? And what happens to those who don't comply? Eric?

ERIC OSTERWEIL

Well, I think one of many voices is no more authoritative than any other, but I'd say, you know, nobody makes you comply with a choice to not buckle your seatbelt when you get in the car. I guess maybe that may not be true in the case of cars, but in the case of the internet, I don't think there's anybody that can compel you to enable security protections.

And it certainly leaves it to everyone to decide whether or not they want them. I'd say that one case where we can see people deploying DNSSEC, because it does benefit the end user, are the DANE deployments. In the case from these slides, when someone's deploying DANE for their email SMTP, that's because they want to be sure that their mail in flight is protected and secured and encrypted.

So, DANE enables that more smoothly. It gives security automation to that. So, that's a case where this is all people decide what they want for their own zones, but that's one motivation that prompts the use of DNSSEC.

MARK ELKINS

And then Warren had his hand up, but he seems to have dropped it, stating, I think I don't understand the previous answer. The map showed that Mexico is just DS in the route, but the website inso.gmu.edu shows 7,000, greater than 7,000 DNSSEC zones. So, is the data up to date, I guess?

ERIC OSTERWEIL

Okay, so you're talking about an interplay between inso.gmu and maps. There is an admin contact for each of those sites. I would encourage you, if you have additional data to drop a note, an email to the admin contact of those sites, so that they can be reconciled if there is or the discrepancy can be explained to you. And I'm happy to --

MARK ELKINS

No, that's fine. Yeah, I feel very, very privileged because I've been teaching DNS and advanced DNS for quite a few years with Johan Stenson. He used to come out to South Africa, and that certainly made me sit up and do DNSSEC long before the ZA domain name system in South Africa was enabled, doing look aside and stuff like that.

And anyway, never mind. I think we're at the end of this session. And as far as I understand, the next speaker, we're going to change the schedule slightly. And we're going to have Andrei next. I understand. Mark, can I add one sentence? Oh, all right. Sorry. Yeah, sorry. Peter?

PETER THOMASSEN I just noticed that on the DNSSEC stats tools website from Wes, there's a count of DS records, and it just crossed 25 million. So, that's a good number to remember.

MARK ELKINS Fair point.

ERIC OSTERWEIL Well, thanks, everyone. I just wanted to really quickly mention that the resources that we use to pull this data together are available on this page. So, if anybody's interested and having trouble finding them, they're all available here.

MARK ELKINS Okay, great. So, I understand that Andrei is the next person that's going to speak. Is that going to be fine? Kathy?

KATHY SCHNITT Yes, that's correct.

MARK ELKINS You're in the room?

KATHY SCHNITT Yes.

MARK ELKINS Okay.

ANDREI ROBACHEVSKY Okay, hi, everyone. I suppose I need to talk into microphone, yes? Okay. All right. Let me upload my slides very quickly. I hope you see the slides. Okay. I'm glad that Eric finished on RPKI stats, because I'm going to talk mostly about routing security in a workshop that starts with DNSSEC, so don't be confused.

I will start with sort of maybe a higher-level look at the Internet integrity in general. But I am Technical Director of Internet Integrity Program of Global Cyber Alliance. And who is Global Cyber Alliance? Well, we are one of those organizations that really care about trustworthy and secure Internet and make any effort to make it more secure and more trustworthy. We operate in mainly three domains.

One is about securing infrastructure, and we're talking here about basic infrastructure such as routing, DNS, and general connectivity. We also look at end users and small companies, Small and Medium Enterprises, where we provide toolkits for them to fight malice and

protect themselves. But we also care about the ecosystem, and you might have heard about some initiatives that Global Cyber Alliance spearheaded, such as Common Good Cyber, Cyber Civil Defense Initiative, and non-profit cyber.

And, yeah, things like Quad9, which Global Cyber Alliance spin off some time ago, probably also now. So, in a way, many bad things that we observe in the Internet are a continuation of all the good things that we enjoy in this wonderful infrastructure.

The Internet is built on collaboration and interdependence. And that means that impact, sort of small efforts being honest, they create global impact, and innovation propagates globally. But it equally means that the global reach enables attacks and cyber security challenges that arise in one part of the globe be sort of impacting other parts of the globe and the Internet globally.

So, if we look at the Internet integrity landscape, we see reaching problems where bad actors can be anywhere. There is no specific jurisdiction where they are located. There is also no one organization that can solve the problem alone or by itself.

In almost all the cases, a collective action is required. And in many cases, the solutions are coursed by so-called collective action problem, where sort of ability to coordinate and move forward is sort of, you know, coursed by negative externalities. Regulation doesn't really work, and the universal mandates don't work. There is no single jurisdiction that covers affected countries or the globe.

And technical network, very unique, and technology is developing very fast, so any regulation will be outdated very soon.

So, changes have to happen, at least as far as we talk about Internet infrastructure, infrastructures at the level of network services. But they cannot be in coordinated, and especially when we talk about this collective action problem, we need action, we need coordination, and we need motivation.

So, what we use in sort of looking at those problems, we have what we call a blueprint for effecting change, where we appeal to the community, which is the industry actors that really need to agree to predefined behavior that we want to promote, and adopted to address the cybersecurity challenge. But that alone, sort of, you know, just speaking to this audience doesn't solve the problem and doesn't motivate them.

So, there are three things that we think are very important in this context. One is the agreed behavior itself. So, I think drawing line on the sand is very important when you can separate unacceptable behavior, unacceptable operational practices from acceptable operational practices. And that's different from any best current practices documents of sorts. It really talks about this kind of bottom line.

It needs resources, such as training materials and tools. So, as much as we can reduce the cost for implementing those behaviors or following those behaviors, yeah, we should invest in that. And tooling and training can certainly help. But also, a level of

transparency, what we call here dashboard, where behavior not only declared, but also verified.

I think that level of transparency mobilizes this collective action in the community. And of course, beneath all this, there's a lot of work with maybe not a very good term, social engineering. But that's how we all humans ultimately, while the Internet is bits and bytes, but we all humans and we react to certain human motivations and incentives.

Now, that might sound a bit abstract. So, I wanted to present one practical example of how that is being implemented in practice and bring up initiative that is, well, not very young by Internet standards. It's more than 10 years old, and it's called Mutually Agreed Norms for Routing Security or MANRS.

So, this initiative, to be just clear, it started not by the Global Cyber Alliance. It started in 2014 by the Internet Society. And then last year, it was moved to secretarial operational responsibilities were moved from Internet Society to Global Cyber Alliance. That's where this initiative now resides. And that's why I'm talking about it.

So, just to maybe step back and talk a little bit about what is the challenge of routing security? I think that pretty much relates to the challenges I outlined when talking about Internet integrity landscape. There are many, many different networks, independent networks that participate in what we call inter-domain routing or

Internet routing. More than 78,000 networks that are in routing terminology called autonomous systems.

And the thing is that the Internet roadmap is built, and there is no one single Internet roadmap. Each network builds the roadmap for themselves, and that's a great architectural sort of principle which enables resilience of the Internet. Unfortunately, they build this information by accepting information from other networks, so you know what networks you connect, and you spread this information to other networks, and they spread the information further, building their own map and enabling others to build roadmap of the Internet.

Unfortunately, the BGP, the underlying protocol that enables spread of this information, it doesn't have intrinsic authentication means, meaning that you can lie or you can make a mistake and declare that your network can be a point for reachability to certain destinations while you are not, for instance.

So, that makes routing system vulnerable, both to malicious attacks and human mistakes, and what is, well, I won't say remarkable, but what happens is pretty common to the Internet, that mistakes, they spread very fast, and they have global impact very often.

Now, on the left side of the slide, just to show you the complexity of the routing system and the Internet system, that's a wonderful chart produced by CAIDA, and each dot represents an autonomous system, and the closer it is to the core, the more sort of connections

the network has, a more sort of customer code, bigger customer code cone, and closer to the edge, well, that's just edge networks that connect to the upstream providers, but you see the complexity of the system, and of course the questions, how can we secure it?

And routing security is indeed very hard, because even if you secure your own network and you implement routing security practices, I think that brings some benefits, but it still doesn't make your network completely secure.

There is no perimeter of sorts that you can secure and solve this problem, at least for yourself. Really, security of your own network is in the hands of other network operators, and that makes the problem very challenging, because you really depend on other network operators, on their action, and you have no way to instigate this action.

The other way is also true, that if you do nothing for your network, and other network operators implement all security controls, your network magically becomes secure. You're no longer vulnerable, because the routing system is secure. So, that creates this kind of, you know, displacement of incentives and action.

That is what is called a collective action problem, when while we all realize, all networks realize, that they are better off by having a stable, resilient, and secure routing system, different priorities, lack of coordination, lack of incentives that move them towards that, doesn't allow them to achieve this goal effectively.

There are some solutions in the space of collective action problem. It's not just the Internet problem, I mean, climate change is another example of collective action problem, and we know how hard solutions are.

In the Internet, and I mentioned this, regulation doesn't really help. As I said, there is global span and dependencies. There is a risk of fragmented solutions, and there is a risk of limited solutions, because regulation can go only so far. A more promising approach that, well, GCA has embraced as well, and MANRS has embraced, is making good practices a norm. Normative behavior, which is widely accepted, not exactly a least common denominator, should make a difference, but also visible and measurable.

And you can see those components when they show the blueprint, well, you see them, the dashboard, the resources, the baseline. And in MANRS, basically, there are two things, two pillars that the initiative is based on. One is the norm, the actions, and the other is practices, not the best practices. That's really sort of distilled essential things that operators must do, what we call MANRS actions.

And transparency, which is a tool, what we call MANRS observatory, and it allows others to see how their other participants in MANRS initiative stand vis-a-vis their commitments. So, it's basically not just declaration of commitment, oh, implement those actions, but it's also verification whether those actions are implemented and to what extent.

The baseline is not the same for every type of network. Networks are different, and one size doesn't fit all, or at least they can be optimized for better impact. And therefore, after launching MANRS for network operators, we also realized there are other types like IXPs that can play a very important role in routing security. For instance, if ISPs or upstream providers can secure sort of relationship between them and their customers, the internet exchange points have an opportunity to secure a peering environment among their members.

We also engaged with content delivery networks and cloud providers because they have a lot of what we call peering power. Many other networks would like to extend their peering relationship with them, and therefore they can present certain requirements, I mean how they accept those peerings and how they accept peerings on what terms, including routing security requirements that MANRS states.

And finally, you can't implement good routing security controls, of course, if your equipment doesn't support necessary features. And we launched a network equipment vendor program in 2021, basically where vendors, participating vendors, they present how the equipment supports implementation of specific MANRS actions for other programs for other participants.

The community, well, I talked about the community, and that's probably one of the most important parts. We start with nine network operators in 2014, and now we are more than 1,300

network operators. There is growth across all four programs. Not the same growth, obviously, but also the sizes of these communities for these types of operators are different, and their appetite for coordination is also different.

It's really global initiative, so here I show some numbers on how initiative is being embraced globally. It's interesting to see Brazil, for instance, that stands out, and this is just to underpin the observation that there is no such thing as global community. Communities are essentially local, and it's very important that there is sort of ignited collaboration on local level, and that's what happened in Brazil, where Brazil and NIC, they created what they call Safe Internet Program, and they went out and promoted, among other things, very important things, I think DNSSEC was among those.

They promoted routing security practices and MANRS, and as a result, many of Brazilian network operators, they embraced this, and they joined MANRS.

Accountability and transparency are very important, because if you don't know whether others adhere to the norm, I think the effect of this, the gravity of initiative is not there. It's just a paperwork. So, what we did, we have this MANRS observatory that shows what we call MANRS readiness, and here I show you actions that sort of MANRS readiness with regards to specific actions in MANRS, such as filtering, anti-spoofing, coordination, routing information, and such, and if there are questions, I can elaborate on what it really

means, but you can see that depending on the region and depending on types of network operators, you can see different sort of readiness scores.

It traces history, it traces trends, and that's also very important to understand, whether we move in the right direction, whether there are improvements in the ecosystem, not necessarily whether MANRS is causing them, but I think ultimately our goal is to reduce number of incidents and make routing system more secure.

And for network operators that participate in initiative, they have a much more detailed view, where they can see how they stand sort of in absolute numbers, what kind of incidents contribute to their score, but also can see how they stand against other participants, whether they're sort of preferably in the range of ready. But as you can see on this slide, there are some networks that need to improve, because they're not in the ready state, they're either lagging or aspiring.

I wanted to show this slide, and full disclaimer, there is no causality, that's just correlation, and we don't know with which it correlates, but it's a nice slide that sort of warms up our heart at MANRS.

The upper blue line shows the growth of MANRS community. The green line shows the deployment of ROAS in RPKI, so it's similar, I think, to one of the lines that Eric showed in one of his slides. And the blue, dark blue and orange charts, they represent number of incidents, monthly number of incidents, and monthly number of culprits. And you can see, I think the positive message here is that

the trending is right, so the two bottom lines go down and the upper lines go up. Again, we like to think that there is some causation here as well, but ultimately that's just correlation.

And to finish up my presentation, I just wanted to say that, as I mentioned in the beginning, at Global Cyber Alliance, we're not only looking at browsing, we're also looking at domain space, and we are looking at numbers, which is traffic and IPs.

So, we have projects dedicated to each of the areas. One of the projects is called Domain Trust, and what we are doing now, we are trying to use a similar approach in the area of domain name abuse to bring communities together, in particular registries and registrars, and try to create a baseline.

Again, not shooting for the sky, but a baseline that at least the leaders can agree and implement probably already, and even more than that, and add some transparency in this space as well, so we can stimulate some collective action here. I think that ends my presentation, and thank you very much. I'm open to questions.

MARK ELKINS Thank you very much. Are there any questions? I haven't seen any questions pop up.

MACIEK PIASECKI Mark, we have one in the room.

MARK ELKINS

If not, that question in the room.

MACIEK PIASECKI

Maciek Piasecki, EURALO. Do you think this kind of intentional misuse of BGP should have been included in the definition of DNS Abuse, or is it a wholly different problem, and it would create even more? And also, as a second question for that, I think I see several people in the room who are members of the European Commission's special group for the best practices in the implementation of the NIS2 directive.

Do you see a chance in this for resolving this kind of problem, especially that many jurisdictions intend to be compliant with the EU law, the so-called Brussels effect?

ANDREI ROBACHEVSKY

To answer your first question, I think DNS security and routing security are, of course, interrelated. The impact of routing security incidents is that it's so fundamental infrastructure in the Internet that in many of the attacks you can see that routing is used to attack upper layers, DNS including application layers, including crypto wallets. I'm just giving you some of the examples. But I think as far as definitions are concerned, I think those problems, they have different nature and different communities participating.

Sometimes they overlap, of course, obviously. But I think maybe addressing them requires slightly different approaches, slightly

different communities. So, I wouldn't mix them, if I understood your question correctly.

On the second, on this multi-stakeholder platform that the European Commission launched recently, we're also the member of this platform, so we are curiously looking what happens there. I think the intention is good, but as the first meeting showed, where you also were present, there are a lot of discussions what kind of shape and form those recommendations could make.

If it's just yet another, or let's produce yet another sort of recommendations, there are plenty of those, I don't think that adds any value. If it's, oh, let's just regulate things, I think we need to be very careful, as I mentioned, what that really means. Whether we just say, oh, you need to implement RPKI, whatever that means on the high level, or whether you're looking for specific outcomes.

So, this is yet to be seen, so quite curious and we also there, as a community, we can separate good approach from not very good approach and we can also discuss what impact that might have, good or bad.

MARK ELKINS

Great. We're actually running out of time a little bit there. The next speaker is Steve Crocker, who is also on the Board, and if we can try and keep it to less than 20 minutes, I'd appreciate that. Thank you. Steve.

STEVE CROCKER

Any questions? That was a joke. Let's see. Thank you. Okay, so we're ready here. I want to talk about what we call Project Jake. The focus of attention is on the DNS registration data, although equally what we're doing applies to a number of, the registration data related to IP addresses and ASN numbers.

The motivation is that we've watched, over a very long period of time, attempts to bring some order and precision to the accuracy of the data that's collected during registration, and then more particularly access control and management of who gets access to that data and under what circumstances.

The process grew without any coherent management over a very long period of time, essentially the entire lifetime of networking, and it's been kind of a sideshow that didn't get a careful design process. And then the effects which we've all seen is that a lot of the data that was available through WHOIS was abused, spam, harassment, fraud, and so forth, and then a counterreaction of hiding the data, putting in false data, proxies, and so forth.

And then legislation, GDPR among them, probably the most prominent, but not alone, of privacy regimes imposed by governments on a fairly crude, broad-brush basis that improved privacy in a sense, but didn't actually improve the precision and utility of the whole system.

So, having watched that over a long period of time, I began to look at that and gathered together a set of people who shared the frustration and insights that I had looking at not only the

technology, but the politics and the business aspects, trying to build a fresh framework for not making policies, but providing the tools for making policies and the tools for expressing those policies. So, that's the basic picture.

Let me take you through where we are. If I can get this to work. All right. We've named the project after Jake Feinler, who was the original manager of the registration data going all the way back to ARPANET days, a really expert person who's still around in her 90s and has generously agreed to let us name our project after her.

So, if one looks at sort of practice today, particularly in the ICANN environment of contracted parties and the use of the Registration Data Request System, the basic picture that one's presented with is that each and every request has to be examined by a qualified attorney to make a decision about what the risk would be in disclosing the data and whether or not the requester is authorized to use that data.

This is an expensive, slow process that is sort of awkward in Internet times when you want to build systems that respond in seconds or milliseconds, and now we're talking about days and weeks sometimes to respond to requests. So, from just a practical, business-like perspective, there's a tremendous amount of efficiency that can be gained by aggregating things along a couple of dimensions.

One is that various kinds of requests, say from an intellectual property attorney or from a law enforcement agency, will have a

similarity repeated. So, an intellectual property attorney, for example, will be interested in finding out who's behind a particular domain name registration because they're looking at infringement as a particular use case, or a law enforcement officer who wants to track down who's behind a domain name that may be implicated in some aspect of their investigation.

And the repetitive aspect is that those same people are going to be making the same kinds of requests over and over again because that's the nature of their business.

So, it's natural to suggest that the decision process about whether to disclose that data has a commonality that once you've done that once, you can embed that as a relationship and say, look, you don't have to re-vet who the requester is or what the purpose is. It's the same decision over and over again. So, that's one element of efficiency.

Another one is that there's a collection of people in that same business. So, when you talk about a particular intellectual property attorney, to take the same example, you talk about the rest of the intellectual property attorneys in the same business, their purpose and the rules that they operate under, again, are likely to be similar.

And one can aggregate those together and sort of squeeze out the extra work that's done to re-vet things. So, those are two elements. One can make rather substantial improvements in the efficiency of the whole system. Yet another aspect is how do you know who

you're talking to? How do you know that you can trust who the requester is?

So, there's identity providers and all kinds of tools for authenticating who the user is that we've not yet begun to make a strong use of. So, with all of those in mind, we've designed a framework that identifies and provides the machinery for operating across these groups. So, it's groups of requesters, groups of data holders, a term we use to refer to both registries and registrars because depending upon the environment you're in, sometimes the registrars are the ones that respond to requests and sometimes it's the registries.

And so, I'll take you through where we are. We're at the point where we now are running a pilot, a live operation with Taiwan, with TWNIC, the operators of the .TW. So, here's just another general picture emphasizing the idea of groups of requesters, groups of data holders. And a very, very key point is that the formation of those groups is organic in those communities, not driven by any top-down or forced organization.

So if we talk, for example, about law enforcement agencies, it's not a one-size-fits-all. It's those law enforcement agencies that cooperate and work under the same umbrella. And so there will be multiple law enforcement groups, hopefully not a very, very large number, but definitely not restricted to be just one.

Same true for each of the other kinds of communities, whether it's security researchers or intellectual property attorneys or any other

kinds of groups. And the second thing to say is that it's not intended that this provides a home for every possible use case that you can imagine. It's a matter of taking care of the common ones and then slicing off those, and then you see what the remainder is. And at worst, you're no worse off than you were before if you don't have a place for a particular use case.

Critical in all of this are agreements between groups of requesters and groups of data holders. Those agreements cover the legal aspects and verbiage and then mechanical aspects of what parameters are included in a request, what level of access is provided for that kind of request, what data elements are expected to be returned, and all of those things so that they can be processed very quickly.

So, we're talking about a hierarchy of protocols. The base protocol is a request and response protocol that is based on RDAP as the basic transport and then additions to that, profiling and extensions, that carry the additional weight here. And then the provisioning mechanisms above that that put all of the agreements in place.

So, here's the pilot that we're running with TWNIC. Here's a picture of the software that is involved in all of this. The critical aspects are the OpenID server that is run by or on behalf of the requester group and an agreement database that is duplicated on both sides so that once the agreement is in place, you don't have to retransmit all of

the information. Both sides have access to what that agreement has been.

So from sort of the beginning to reaching all the way to making a request, there's a process for creating a group, and then there's a process for populating the group, shown only the requester side, but basically the same thing on the data holder side, adding members or removing members as needed. And then the interesting part of forming agreements between groups, what has emerged as the likely sequence of activity is that a data holder group would offer a set of templates of sort of agreements that need additional information filled in.

So, a group of data holders would say, look, we're willing to work with law enforcement agencies of the following type, and here's what the agreement looks like. Please fill in your pieces, submit it to us, and if we're comfortable with that, we'll agree to that, and if not, we'll say so. And then requester groups have to then go look for those agreement templates, fill them out, and make agreements as it were.

And then with all of that in place, the actual process of making a request is that a requester module, represented on the bottom on the left, initiates a request by getting a bearer token from the OpenID server, and then including it in the horizontal red line that is the RDAP request going over to the data holder, and the data holder then receiving that request, which includes the bearer token, can verify that that's the right, that that's authenticated, and

then make decisions about what data to return, if any, and other aspects of all of that.

One of the important elements of the design is that there's a persistent process running on behalf of the requester. So, it's not a case of a requester simply logging in or reaching a website of the registrar or the registry, typing in a request and then going away.

If the data holder wants to look at the request and review it manually, that takes some time that is well outside of the time frame that an RDAP process can persist. And so there needs to be a way to have a deferred response at a later time. For that and for some additional reasons, the idea of having a persistent process on behalf of the requester makes a big difference in making all of this work smoothly.

That's another diagram of the same thing. Two other concepts that are critical to the way we think about these things is that every, let me skip past this a little bit. Every data element that is collected, we take the position that either explicitly or implicitly receives two attributes.

One is how sensitive is it, which is a sort of broadening of the concept of public versus private. We listen to a lot of conversations about how to treat the data, and including such details as, well, if somebody is asking for some data but they're not allowed to see it because they don't have access at that level, then return “redacted”. So, that suggests sort of a simple two-level model of

sensitivity of the data. But we've listened even further, and in some cases, we hear something that sounds a little peculiar.

Well, that data, even if it exists, if you're not allowed to see it, it won't tell you whether it's there or not. So, there's an inference there that that data must be, in a way, more sensitive.

And then looking at sort of a broader set of use cases, we settled on a four-level model, which we think is sufficient but needs to be tested, obviously, with you can think of the lowest level as the equivalent of public, the next level up as the equivalent of private, and then sort of a very private and the top level would be things that, in practice, likely would only be accessible for somebody who has complete access, a court order or equivalent sort of thing.

So, that's the sensitivity attribute assigned to each data element. And as I say, either explicitly or implicitly, none of the systems, essentially none of the systems that are in use were on the order of 400 million registrations across the entire Internet, have such markings on them at present. However, one could ask, well, what is the implied sensitivity by practice?

And we think that it's possible to label, to infer what the label would have been if it had been assigned at registration time, and that comes into play later at the time of disclosure.

The other attribute that is treated similarly is how accurate is that information? More precisely, what level of validation was that data subjected to at the time that it was collected? And again, not in a

forced way, but a four-level scale seems to be in common use, not just in our thinking but more broadly, where the bottom is, take whatever is given and there's nothing to be said about the validation.

The next level up is a syntactic level. If it's an e-mail address, does it look like an e-mail address, for example? And then the next level is, is it operational? If you send a message to that address, do we get an answer back? And the level above that is, and on top of whether you got an answer back, do you have reason to believe that that actually is an e-mail address that's assigned to that contact, for example? So, that's a core piece of the internal technology.

Another piece of the overall design is that different registrants are typically treated differently. So, if it's a business, it may not get the same level of privacy protection as an individual. That distinction is perhaps the most common, but there are other kinds of distinctions that can be made. For example, even within the class of individuals, there may be a distinction between an ordinary level of protection for them versus an extraordinary level of protection.

Think of celebrities or politicians or others that might be subjected to more attention than they want. And even in businesses, you have a distinction. One of the examples that we found quite compelling, a business that's actually a nongovernmental organization operating in a contested area might want very high level of protection about their physical address. And so even

though they're a business and there's no personally identifiable information included, they still may want and deserve a high level of protection.

And then we've encountered more recently the idea that a certain class of registrants may want to have the fact that the data that they provide subjected to a high level of validation so that it becomes more trusted when it's received by requesters.

So, we have one of the tools that's part of what we've built is a way of capturing all of this in a very precise way. What you're looking at is a snapshot of the registrant portion of one of the policies that Taiwan uses. The one you're looking at is for individuals that require extra privacy. And in contrast, here's one for ordinary businesses. And I'll just flip back and forth. And if you look on the right, sorry, in the middle.

KATHY SCHNITT Steve, I'm sorry, two-minute warning.

STEVE CROCKER Thank you. All right. I'm almost there. That's good. All right. So, here's an individual with extra privacy. And the privacy values are -- I've lost track of things. I think it's on the right. So, here's -- I'm just flipping back and forth so you can see the contrast there. Big differences in there. And then here, the third in this series is this is a so-called green. It's how they talk about their businesses that have all of their data extra validated.

And the middle column is where the validation levels. And you can see sharp differences in that. So, that's the basic system. Here's an example of the differences in what you would get back for somebody who has a low level of access versus a high level of access for the exact same data. And on the left, annotated in red, are the redacted portions.

And over on the right in green are what is the actual data underneath all of that. So, that's the basic system that is designed and is in the process of being implemented and deployed.

We are actively looking for others who may be interested either on the requester side of this equation or on the data holder side of this equation. And feedback, commentary, criticisms, suggestions, et cetera, are all quite welcome. And on this last slide, which is in the deck, are the ways that you can communicate with us, reach out, and so forth. Questions?

MARK ELKINS Thank you, Steve. Questions?

STEVE CROCKER Yes.

MARK ELKINS The gentleman next to you.

MACIEK PIASECKI

I was wondering if those levels you've mentioned would be the same globally or maybe country specific or able to modify because we've got countries like India who are more on the efficient policing side. We've got countries like Germany which are more interpolated policy oriented on privacy. I was trying to get some answers from GAC here, but they weren't able to answer.

STEVE CROCKER

It's a very good question, and I want to be careful and quite straightforward about it. The question you ask is, are these levels applicable across different jurisdictions and so forth? We think so. It's a sort of major bet in the design process. What we have not done and have absolutely no intention to do is to say what levels each organization or each jurisdiction should choose.

So, for example, taking the headlines that we see, India might say, we insist that a certain amount of data is treated as public and you're not allowed to code it as higher than that, whereas Germany might say, no, we want all of that data coded as much higher level of privacy. So, the scale is the same scale, but the application of it depends upon local policies.

MARK ELKINS

Thank you very much. We don't really have time for more questions. So, the last presentation this morning is from João and Quantum Threat. Straight over to you, please, João.

JOÃO MORENO Hello, Mark. Hello, can you hear me?

MARK ELKINS Yes, we can hear you. We are a bit strained for time, that's all.

JOÃO MORENO Benoit will say a few words before me.

BENOIT AMPEAU Yes, due to the time constraint, I will do a very short introduction to João's main presentation today. So, hello, everyone in the room and people attending online. Just to introduce the presentation, just to explain why AFNIC supports the International Security and Safety IGF Coalition and this ongoing work to address post quantum cryptography transition. So, I am Benoit Ampeau, Director of Partnerships and Innovation slash Labs at AFNIC, the French Internet Registry.

The problem, we don't know when relevant quantum computers will be operational. They could break RSA and ECC encryption algorithms, which are predicting DNSSEC, TLS and PKI today.

Anyone with sufficient resources can already harvest encrypted data now to decrypt them later, once quantum computers are available. Some early operational tests on the .FR and from also other registries or DNS operators show some concrete impacts. In our case, we have seen that PQC signatures are like five to ten times slower than ECC, for instance.

Zone files are significantly larger and its resource demands substantial higher resources. Adware HSM, for instance, also must be clearly upgraded. So, these are not theoretical, these are real outcomes from our initial tests. So, why did we find this first discussion on this topic? The transition affects the entire DNS ecosystem, registries, registrars, DNS hosters, resolvers, ISPs on every component in the DNS chain that is currently integrated cryptography.

So, as you know, IGF brings stakeholders together from technical operators to policymakers and industry to discuss standards and timelines. So, we started last year by founding an IoT PQC report, which we presented at IGF in Oslo. And we are founding currently a new phase this year, focusing on organizing expert discussion groups about DNS and routing security.

So, just to end, like IPv6 and DNS deployments, this will take time, over a decade? Question. This complexity and the complexity of the transition at scale is demanding. So, that's why we think we must plan strategically now for keeping resilient and secure Internet infrastructure. So, please welcome João and it's over to you. Thank you.

JOÃO MORENO Okay, thank you, Benoit. And good morning and afternoon, everyone. I'm here to make one main argument that the real risk is waiting too long to prepare for a transition that will take many years. And to do that, in the next few minutes, I want to do four

things. First, explain why the post-quantum transition matters for DNSSEC and RPKI.

Secondly, I want to show that the challenge is more complex than just cryptography. Third, share what our ISTC ethnic research found about how governments are organizing for this transition. And finally, outline what we are doing next.

So, we still don't know when a cryptographically relevant quantum computer will become operational. And on top of that, Internet infrastructure transitions take a long time. So, we know from experience, IPv6 and DNSSEC have both taken years of standards work, deployment testing, incentives, coordination and operational learning to start flying.

So, if the transition to post-quantum cryptography for DNSSEC and RPKI follows the same pattern, waiting for certainty would be a huge mistake.

The problem is simple. We do not know exactly when the threat becomes credible, but we already know that the transition will be slow, complex and distributed. So, this is why this is today's challenge. And as you can see here, we have less than half of the DNSSEC systems operational. And well, okay, but let me make this concrete. What is quantum unsafe?

The main quantum risk is today's public key cryptography. And when moving this to DNSSEC, this means that RSA and elite curve cryptography is unsafe to a quantum computer. And these are

used for zone signing in RRC signatures, DNS key public keys and the trust chain linked through DS records. So, when a resolver validates DNSSEC, it may be relying on RSA or elite curve signatures that are a quantum unsafe.

In RPKI, the picture is even more concentrated. The standard profile is central in RSA with SHA-256, including RSA-PKCS signatures. These are used in RPKI certificates, URLs, rules, manifest and other CMS signed objects. So, when relying parties validate routing attestations, they are also relying on classical public key cryptography, which is not post-quantum unsafe. So, this directly affect, on long term resilience of two trust layers, the Internet name and route. So, this is where the workshop context really matters.

The first step is to choose a new algorithm from Nestle, but the hard part is the operational transition. In DNSSEC and RPKI, changing signing algorithms have consequences. Larger keys and signatures can lead to larger responses, as Benoit said, affecting packet sizes, transport behavior, memory use, CPU use and validation performance.

Key rollovers may also become more complex, especially in hybrid or transitional stages. It is about migrating an ecosystem built on existing certificate profiles, validation software, repository behavior, object formats, operational procedures and trusted assumptions. This is why this topic belongs in this room.

Now let me turn to the ISUC AFNIC research. One useful contribution of the report was to show that governments are really already organizing themselves, but not in the same way. The United States is moving in a more top-down, mandate-driven way with NIST standardization and a 2035 migration target for federal systems. And we also saw this in the G7 document released last month.

The European Union is moving in a more coordination-driven way, using existing governance structures and placing more emphasis on coordination across member states and on hybrid approaches during the transition. The report also identified concrete national examples like the French, German and Dutch approaches.

One key finding is that governments are not waiting for one perfect moment. They are already using roadmaps, guidance, inventories and a staged migration model. The report also makes practical recommendations for governments to do these kinds of roadmaps, public-private partnerships, funding, crypto-agility, incentives, procurement and leverage the international coordination.

For industry and service providers, we suggest to do the cryptographic inventories, readiness plans, dual-phase migration, risk-based prioritization and also invest in hybrid approaches and testing. And this is exactly where the research becomes relevant to DNSSEC and RPKI. It helps us move from policy awareness to infrastructure planning. And that brings me to the next step.

The follow-up phase on ISGC ethnic work is designing around two critical tracks, one for DNS and one for routing. The current phase includes outreach, pre-public webinars and the creation of two working groups engaging with the IPM community, RIRs, IETFs, TLDs, commercial ccTLDs and routing businesses. That matters because DNSSEC and RPKI are related but are not the same transition. They have different deployment realities, software ecosystems and operational bottlenecks.

So, they should be worked in parallel. And what we need now is structured readiness work. So, map the dependencies, identify blockers, test assumptions, improve interoperability and connect standards with operations. And to do that, we are also seeking partnerships. We are looking for partners who want to contribute expertise, join the working groups and help us shape this work.

So, we are also looking for opportunities to co-fund this effort so it can be sustained and bring together the right mix of operators, standard experts, researchers and policy stakeholders. Because if this transition is collective, the preparation also has to be collective.

So, let me end with this. The transition to post-quantum security for DNSSEC and RPKI is not just a cryptography issue. We need to start organizing seriously now and we welcome partners who want to help building that work with us. Thank you very much. I would be happy to take one or two questions. Maartin, you can leave this. Thank you.

MARK ELKINS We can take one question. Yes.

KATHY SCHNITT We have a question in the room. Peter, go ahead.

PETER THOMASSEN Hello, this is Peter Thomason. So, this is all quite an important effort. I just wanted to point out that there is an effort in the ITF. It's the PQC DNSSEC Informal Research Group. They meet at ITF meetings in the site meeting mode. Next time is next Tuesday. If you go to sitemeetings.ietf.org, you can find details. And that's where the protocol issues, at least for DNSSEC, are being considered.

So, I invite you to join that as well. And I'd like to make the point that I don't think it's realistic to think that all the DS records will be updated manually once PQC DNSSEC algorithm has been specified. So, I'd encourage you to include in your considerations the goal of also deploying DS update automation to make PQCD DNSSEC work in practice.

MARK ELKINS Thank you very much, Peter. We literally have 30 seconds.

JOÃO MORENO: Okay, yes. So, I invite everyone to join us. This is a collective effort. The ITF is doing an amazing work with the protocols. But we need a coordination effort, a huge coordination effort. Because, as you all know, the DNSSEC is a huge challenge for us. It's still today. And we have a path forward to work together.

MAARTEN BOTTERMAN: Yes, just to finalize on these 10 seconds, I realize, Maarten Botterman, iIt is really preempting and preparing for what may become a problem in the future. And therefore, this ethnic-supported and I3C initiative is really asking for you also to engage to make sure that the problem is coped well and the issues are coped well. We already had good discussions here in the room. And please step forward to João, to Benoit, or to me if you want to get involved. Thank you. Thank you very much.

MARK ELKINS: Thank you very much. Next session, half an hour.

KATHY SCHNITT: You may stop the recording. Thank you.

[END OF TRANSCRIPTION]