ICANN85 Mumbai | CF – DNSSEC and Security Workshop (3 of 3)
Wednesday, March 11, 2026 – 16:30 to 17:30 IST

KATHY SCHNITT

Hello, and welcome to the DNSSEC and Security Workshop. My name is Kathy, and I am joined by 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 keep this discussion moving smoothly, here is how you get involved today. 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.

You may also submit any questions for the speakers today through the Zoom Q&A pod. We will read those aloud as time permits and when directed by the moderator. While the chat is available for attendee interaction, please note that 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. With that, I am happy to hand the floor back over to Russ Mundy.

RUSS MUNDY

Thank you very much, Kathy, and it is my privilege to share this final part of our three-part series today with the DNSSEC and Security Workshop. At the end of the last Part 2 session, we heard about a number of the challenges and problems, especially things that are related to communications about what is going on. Our first pair of speakers are going to tell us a success story from Ghana. I would like to, with that, turn it over to Godsway, and he will also introduce his collaborator on that. Please go ahead. I will give you some time warnings as we get close to the time that we have allocated. Thank you.

GODSWAY KUBI

Thank you so much, Russ. This is Godsway Kubi, an ICANN Fellow. Today I am very pleased to be here, and thank you all for the opportunity. We are going to look at DNS security and also look at some success stories that Ghana has been having throughout this validation project. This is actually a collaborative effort that significantly has improved our scores from 24% to 99%, which makes us the first African country to have these highest scores in terms of these highest validation stats in terms of samples that are used.

We will describe our journey, challenges, and also look at lessons learned. In order to sustain these gains, we have also adopted some industry best practices, which ICANN has outlined in the Kind DNS framework. My colleague Isaac will be taking us through this

	presentation. Isaac, please, if you can hear me, kindly take the lead and take it up. Thank you.
ISAAC SAKYI	Yeah. Thank you very much, Godsway.
GODSWAY KUBI	You can share your screen from your end for your presentation.
ISAAC SAKYI	Hi, everyone. My name is Isaac Sakyi. I am a network engineer, and I will be running us through this presentation. Let me know if you can see my screen.
GODSWAY KUBI	It is now loading.
KATHY SCHNITT	Two seconds.
RUSS MUNDY	There it is.
GODSWAY KUBI	I think it is up. Yeah.

ISAAC SAKYI

Good. Like Godsway said, we will be sharing the Ghana DNSSEC validation success story and also highlighting regional best practices that we have adopted. In this presentation, I am going to give you a brief on the ecosystem in Ghana, the problems that we faced, the collaboration, and what led to the resolution. We will also look at the adoption of ICANN's Kind DNS framework. Then we would also share common myths that are associated with enabling DNSSEC validation, and then we will share our learnings and give our key takeaways from our project.

Ghana is a country in West Africa. It has a population in excess of 35 million, out of which 26 million have access to the internet. We say online penetration is about 75%. Currently, there are three major mobile network operators and over 70 ISPs. We have five submarine cables landing here and two internet exchange points. In this country, the Ministry of Communications, Digital Technology and Innovation has the oversight here, but the regulator in this case is the National Communication Authority.

What was the problem before? As you might be aware, the .gh domain itself is not signed, but apart from that, we also had very low DNSSEC validation scores. At the time we embarked on this journey, our score was very low, less than 24%. Today, as we speak, we are at 99%, and currently the highest in Africa in terms of samples used for validation. Having low DNSSEC validation and not having your domain signed comes with a lot of trust issues online. It also has a poor cybersecurity outlook for the country and for companies that are hosting online. These challenges informed our

decision to start with DNSSEC validation because it is the easier of the two.

How did we get around all these that I have mentioned? The government, through the National Communication Authority and the Ghana Domain Name Registry, have been very instrumental. They created policies and frameworks that were very supportive. We had a lot of key stakeholder engagement and consultation on the problems that I have described earlier. They also facilitated awareness creation trainings and capacity building. In 2023, on the back of the African Peering and Interconnection Forum, they organized an ICANN workshop for DNSSEC training. This was very good because a lot of the industry players were around and took the opportunity to take part in the training. After these trainings, we have groups where individuals can provide feedback for continuous improvement.

Today, when we look at the map of validation for Africa, it looks like this. Ghana is green. Ghana is here, and it is currently showing the highest validation rate for a country. Having achieved this success, many of our operators, specifically the largest mobile network operator MTN Ghana, decided to voluntarily embark on the ICANN Kind DNS framework. In fact, in Africa, MTN Ghana is the only listed telecom company in the shared resolver category. It is very large in Ghana with over 20 million subscribers. Having them take the lead in this aspect creates the path for others to follow.

What is this Kind DNS framework all about? It is a list of recommended best practices for DNS operators. For MTN Ghana, we are looking at a shared DNS operator. In this framework, ICANN stipulates that every operator should enable DNSSEC validation and employ ACLs to limit queries that come to specific ranges of IPs. There should also be a separation of function. You do not host the authoritative DNS and the recursive resolvers on the same server. You have them separately, and if possible, the setup should be such that they are situated in two different geolocations so that in the event of an outage, you will not be completely taken out. There are other policies that can be implemented to enhance privacy, and they encourage you to implement those. Last but not least is to monitor this critical service continuously.

We have gone through the process, and we are glad to share that MTN Ghana has actually gone through the validation. We can debunk some myths that are associated with it. For instance, before we embarked on this journey, there was a lot of noise about DNSSEC validation being very difficult to configure. We also heard that it was a technology only reserved for the tech-savvy. Apart from that, we heard that it is very heavy on resources and that utilizations for memory, CPU, and links would considerably go up. Others also said once you enable DNSSEC validation, domains and URLs will fail to open.

From our experience, all these are not true. We recommend, and our learning along this journey has been, that before you embark on a DNSSEC validation project, make sure that you have properly

taken stock of your infrastructure. Ensure your hardware is not end-of-life and the software you are running for the DNS is up to date. We also saw that configuring DNSSEC validation in most software today is straightforward and simple. You just need to remove a comment or enable a particular line item and restart the software. If you are running open-source software, the community has a lot of expertise. There are many people who have embarked on this journey and have lessons you can leverage to do a similar thing in your environment.

All it takes is to ask. Most of the things that we experienced, we got support from the community. For instance, when we started, we realized there was a coincidence where the latency for browsing generally went up slightly. A lot of people attributed it to DNSSEC being enabled. We engaged the community, and Yazid from ICANN was very instrumental. He connected us with some experts in the field who helped us go through troubleshooting scenarios, only to discover that it was just a coincidence. At the time of enabling DNSSEC validation, there was an upstream challenge that resulted in rerouting of traffic and hence the higher latency, which had nothing to do with the DNSSEC validation itself. If you do not ask, you would not know. The community is enriched with people with a lot of experience. You can leverage that to give you smooth sailing in your implementation process.

RUSS MUNDY We have four minutes, please.

ISAAC SAKYI

Okay. We also realized that awareness creation and community engagement is very key. You cannot do without it because if the people going to do the implementations are not aligned and do not have knowledge of what will happen afterwards, it will be difficult for them to undertake such a step. We also encourage capacity training. The trick here is to leverage local technical expertise. Once someone has gone through it locally and is there to encourage his colleagues, it becomes easier to adapt. That is all for our side, and I am open now for questions. Thank you very much for your audience.

RUSS MUNDY

Yeah, we do have time for a couple of questions here. Kathy or Dan, could you run the question queue, please?

KATHY SCHNITT

You got it, Russ. So do we have any questions in the room, chat, or in Zoom? I am not seeing any in the queue or in the Q&A.

RUSS MUNDY

Well, I think this was wonderful, guys. We really appreciate hearing your success in getting to 99% of validated responses. That is very, very impressive. Thank you so much for sharing that with us and the lessons learned. And so now at this point, since we do not have questions, we can move on to Peter Thomassen, who for those that were in the last session already know Peter well. He is from deSEC

and going to give us an update on what is going on in the IETF arena relative to DNSSEC. Peter, over to you.

PETER THOMASSEN Thank you. Let me quickly bring up my slides. This will not be all too long. All right, here we go. I will give a quick update on what is going on in DNS security at the IETF. Let's dive right into it. Since the last time I gave this update in October in Dublin, three new RFCs have been published that relate to the maintenance of the DNSSEC security protocol. They essentially belong together. The way the list of signing and hashing algorithms that needed to be supported by signers and validators was maintained in the past was in an RFC table.

Now there is a more accessible and more easily maintainable IANA registry for that. That move was made in the first of these three RFCs, 9904. That did not actually change any of the content, just moved the table. Right after that, two changes were applied to phase out the SHA-1-based and GOST-based signing and hashing algorithms. That is now deprecated.

Second part of maintenance: DELEG is coming up, as people know, and DELEG will introduce a new record that lives in the parent zone as a new delegation type. Unlike the NS record, that DELEG record is expected to be signed by the parent so that the referral can be validated by the resolver. There already is another record that is signed by the parent, which is the DS record, but that record was specifically defined as special and signed by the parent. DELEG is

the second one of this sort, and instead of treating them all one by one, Roy Arends came up with the idea that maybe we should have a range of type numbers, and any type from that range would be signed by the parent.

That draft proposes that, and the DELEG record will likely be taken from that range once this has been adopted by implementations. If you do this, you need downgrade prevention because if you do not have that, you could just strip this new sort of record and the validator would not notice that it is missing because it would not expect a signature. You need to tell the validator to expect a signature for this type of record, and that is done by a DNSKEY flag bit that is set in the DNSKEY record of the parent zone. There are links here to these drafts.

Georgios is proposing what is called dry-run DNSSEC. Essentially the idea is, let's say you sign your zone and you have your keys, but you do not have a DS record yet, so the domain is still insecure. Before you actually set the DS record and turn it live, you might want to instead set a test DNS record that uses a signing algorithm that is not allocated, so it will effectively continue to treat the domain as insecure. The goal is that specialized validators that have testing mode support would be able to see that particular record and then report on whether everything is set up correctly. Once that is confirmed, you can remove the test bit from the DS record and your domain will be secure. I call that the soft fail mode here.

This has been recently adopted by the DNSOP Working Group, and the working group is refining that protocol to make this a generally available thing. There is also RFC 9364, which is essentially a reference list of all the RFCs that are relevant in the DNSSEC context. It is maybe two or three years old, and since then a few things have changed, including the stuff on the previous slide. An update is due on this RFC, and that was proposed by Paul Hoffman and is now being adopted by the working group as well. That should be a pretty quick update, just changing a few lists and it will probably be published quickly.

About automation, we talked earlier about the two documents I have been working on regarding consistency requirements for CDS and CDNSKEY records, and operational recommendations for DS automation. I will not go into details about that. Another proposal is from Johan Stenstam, who proposes a way to apply delegation updates using what is called dynamic DNS messages. A DNS packet has a field which determines whether the message represents a question, an answer, or in this case, an update. The DNS update protocol actually is not very new. It has been used in the past and is in use to maintain zones within organizations.

The new idea here is that maybe the child could send such an update message to the parent and then update the NS records or maybe DELEG records in the future or even DS records without any sort of other CDS records that would need to be published. The main question here is how does the parent believe that the update it gets is actually authentic? Some sort of trust has to be

established. That is done using the signal mechanism through the name server domain hostnames. The name server hostnames need to use DNSSEC. Then you can store under a subdomain of them the key that is used to authenticate the DDNS updates. That mechanism is inspired by RFC 9615, which uses the same mechanism for DNSSEC bootstrapping where you do not have a chain of trust to the domain yet.

Because of that property, the DDNS update mechanism for delegation maintenance is expected to work for both signed and unsigned children. You could, for example, update delegation NS records without the domain using DNSSEC itself; just the name server hostname needs that. That also has been adopted and will probably take a year or so to finish.

Regarding the future, this is largely the same as last time. The DELEG Working Group is doing its thing and moving forward. Some technical questions have been settled. For example, the syntax of the DELEG record is now clear. I expect that next week's IETF will have some more updates from the DELEG Working Group on that. There is also the DCONN Working Group, which is standardizing Domain Connect. Last time it was very new; now it is in the middle of its work and updates are coming. The Domain Connect protocol is expected to be standardized by the IETF, maybe next year.

For those who are not too familiar with it, the Domain Connect protocol is not actually a DNS protocol, which is why it is not in the DNSOP Working Group. Instead, it is a protocol where if you have a

domain and you would like to use a third-party mail provider, you usually have to set up an MX record for that and a few other things. You can do that manually, or you can use Domain Connect, which is a mechanism to authorize your email provider to make the changes in your zone. That requires your DNS provider to support that sort of interface and the corresponding authentication. That is called Domain Connect, and it is already available in the wild and has been deployed by a bunch of companies without having an IETF standard. Now that standard is coming.

A few times today, the post-quantum topic has surfaced. As usual, the IETF will have a side meeting next week on Tuesday, March 17th, at 08:00 AM Chinese time. There is a link here, sitemeetings.ietf.org, which has information about how to join that. If you want to learn more about the research results and the content presented there, there is a wiki link that you can follow to look at those results. That is actually all I have, and I will give it back to Russ unless there are questions.

RUSS MUNDY Do we have any questions for Peter?

DAN GLUCK Nothing in Zoom and nothing in the room.

RUSS MUNDY Well, that was great. I did have one question. The DCONN Working Group and their... Are they tied to DNS in that they are using DNS

records as part of their communication, or is that something that happens fully within DNS after some DCONN signaling goes back and forth?

PETER THOMASSEN: Largely, it is not a DNS protocol. It uses an HTTPS-based interface. The question is how to reach that. If I am not mistaken, the DNS provider of your domain that would support that would publish a text record at an underscore subdomain of your customer domain. Then the third-party mail service, for example, who would like to provision records for your domain, can look up the endpoint of the DNS provider in this underscore text record, and then the rest of it is an HTTPS API.

RUSS MUNDY: Okay, great. Thank you. Obviously, I haven't been following that very closely. That was real helpful. So any more questions?

ISAAC SAKYI: I wanted to ask, so with QUIC becoming a protocol that is being adopted in the DNS environment, how is the DNSSEC signing and then its application looking for it going to be affected by this protocol? Is it going to conform to the use of this, or are you just limiting it to what we have, the UDP and the TCP currently?

PETER THOMASSEN

Thank you for the question. It is a good question. The QUIC protocol is a mechanism to ensure encryption of the DNS responses and also of the questions. That is like HTTPS, which encrypts the website content and things you send, like a form input.

DNSSEC itself is inside that connection and contains signatures for the DNS responses that are transmitted inside the encrypted connection. Now, you could think, "If the connection is encrypted, why would you care about the signature in the content of it?" The reason for that is that it is not that your browser talks encrypted QUIC directly to a DNS name server that has the actual information that the browser needs. Instead, there are steps in between. There is a resolver and so on and so forth. At each step, the encryption channel is broken. The signature ensures that the recipient, which usually is not a browser, can still validate the signature.

Those are two independent security properties. DNSSEC ensures integrity so that the message is not manipulated, and QUIC ensures confidentiality, which ensures that the message is not observable or the content is private. On the technical side of it, QUIC is UDP-based and adds stuff on top that makes it more reliable and adds encryption, but it has a lot of UDP-like properties. It is not really a departure from UDP. It is equally fast.

RUSS MUNDY

Yes. The argument is QUIC is quick. Anyway, thank you, Peter. If there are no other questions, we will go on to our third presentation of the session. The name of our presenter is a very big challenge for

me, but I will do my best. Olabimpe Adeloro. Let me go ahead and pass it to you, and please feel free to correct my likely terrible pronunciation.

OLABIMPE ADELORO Thank you very much, Russ. Kindly confirm if you can hear me.

RUSS MUNDY Yes. Can hear you well. I see the image. Are you doing your slides or Kathy or Dan presenting slides for you?

OLABIMPE ADELORO Let me check with Kathy if she does not mind, because my internet seems to be slower this afternoon. I told her I would present it myself, but I have been off and on the room for quite a bit of time.

KATHY SCHNITT Dan is pulling up your slides now.

OLABIMPE ADELORO Thank you very much. Thank you so much, Kathy. I appreciate that. Good morning, good afternoon, and good evening, depending on where you are joining us from. My name is Olabimpe Adeloro. I work in cybersecurity and SOC investigations. Today, I will share insight on how DNS security gaps contribute to real-world attacks. Every day in security operation centers around the world, we investigate phishing campaigns, email abuse, and service

disruptions. In many of those incidents, DNS is not just infrastructure; it is the attack surface. Today, I am glad to share with you what we have observed from SOC investigations, and how gaps in DNS security and DNSSEC implementation contribute to real-world attacks.

Before looking at this attack pattern, however, it is important to understand why DNS plays such a central role in internet operations. As you can see on the screen, I will be talking about this slide at the moment. DNS is one of the most critical internet infrastructure components. Of course, experts in the room will agree with me. The reason being that every digital interaction starts with DNS. For example, if you were to log on to the ICANN website. When you type icann.org, which is the human language, the machine does not understand that. DNS translates it into an IP address that the machine understands. So you will have something like 192.0.43.7.

Without DNS, it is obvious that websites will not load, emails will not route, and cloud apps will fail us. That tells us that DNS is very important. It will not be surprising why attackers love it. For one thing, DNS is trusted automatically. Users rarely question DNS results. Attacks and exploits are carried out based on the trust that users have in DNS. If DNS resolution points a user to a malicious server, then the user might not even notice they are on a wrong site. This makes DNS a high-value attack surface.

Next slide, please. At this point, let's consider why DNS is an attractive target. Next slide, please. Thank you. For several reasons, the DNS is an attractive target to attackers. I will mention a few of them as highlighted on the screen. For one thing, the DNS is globally trusted. DNS is designed to be reliable and widely accessible. Systems assume DNS responses are correct. Another reason is registrations of domains are quite easy. Attackers can register domains within minutes, making it an easy prey. We have often poorly monitored DNS, and that makes it also a target. Organizations monitor endpoints, they monitor firewalls, and servers, but unfortunately, many do not monitor DNS logs. Attackers exploit this visibility gap.

We have frequently misconfigured DNS. DNS configurations can be very complex. Experts in this room will agree with me. Therefore, we have some common mistakes associated with DNS configurations. Mistakes like missing security records, weak policies, and incorrect validations. Apart from these problems, we have the problem of limited DNSSEC adoption, which is very big to this audience. DNSSEC adoption globally is still very much uneven. Many domains are not signed, and for those domains that are signed, some resolvers do not validate signatures. This weakens the protection of DNS.

Next slide, please. Thank you. Now let me walk you through the real-world attack life cycle. There is an African adage that if you want to catch a thief, you must have the knowledge of the tactics of a thief. For us to protect DNS or prevent DNS abuse, there is a

need for us to know how these attackers operate. What are the steps involved in DNS attacks? As you have it right there on the screen, we have six stages. First stage is to register a malicious domain. You could have a domain like Microsoft-login.com. If an attacker wants to take advantage of that, they can use Microsoft, but the 'o' character there can be replaced with '0'. An end user might not notice that that character has changed. When it is clicked on, it becomes a problem because what a user wants is not what they will eventually get. It is possible for this attacker to impersonate Microsoft.

After registration of a malicious domain, we have the configuration of DNS records. Attackers configure A records, MX records, and in some cases, TXT records. Thereafter, they launch a phishing campaign. Attackers send emails like, 'Your account has been locked. Click here,' just to lure users. When a user clicks on such an email, it goes to a phishing site. Then comes the next stage: connection of victims with this site. Because the DNS resolves the domain, the browser connects normally, so there is trust on the part of the victim. Next is credential stealing. When this user gets to that site, they likely put in their username and password. In case they are doing a banking transaction, they put in their banking information, and attackers will harvest these credentials. If they are successful with their malicious intents, here comes the final stage, stage six. Either the domain is abandoned if there is detection, or they will reuse the domain in some cases. We can see that there is a need for care.

Next slide, please. Thank you. At this point, I will be talking about case patterns. We have several of them, but in this delivery, I will be mentioning just three. We have phishing domain abuse. As you can see on the screen, attackers use the technique that we call domain weaponization. An attacker can replace a character with another character that a user will not readily notice. The implication of it is it increases phishing success due to users relying purely on visual trust. They feel they have always been on this site and that it is correct.

Next slide, please. Another case pattern is bypassing email security via DNS misconfiguration. The techniques that attackers use sometimes could be by exploiting common operational oversights. Oversights like missing or weak SPF records, missing DKIM signatures, and weak DMARC policies. When this happens, email gets spoofed very easily and sometimes even seamlessly. Malicious phishing emails are inherently trusted by a secure gateway, paving the way for attack success.

Next slide, please. The third case pattern I will be talking about is leveraging DNS for infrastructure evasion. Attackers take advantage of some things just to leverage DNS to evade infrastructure. They use fast flux DNS infrastructure. There could be that challenge with short TTL values. TTL, for the sake of simplification, refers to Time to Live values. When Time to Live values are very short, the resolver consistently fetches new IP addresses. This allows attackers to change infrastructure very quickly. We have abuse of domains also in some cases. We could

have some domains that mimic real domains, and users will think they are real. This makes it extremely difficult for detection on the part of SOC analysts or SOC teams, and this increases the persistence of malicious infrastructure.

Next slide, please. Thank you. As we can see from above, DNSSEC will be very important. We cannot just rely on DNS and think we are safe. There is a need for DNSSEC. DNSSEC has cryptographic validation to DNS responses. Normally, a DNS response could be spoofed. However, with DNSSEC, responses are digitally signed. We got that from Peter's presentation. Thank you, Peter, for that. Responses will be digitally signed. Resolvers will verify signatures using a chain of trust. DNSSEC works through a chain of trust that starts at the DNS root and extends through TLDs down to individual domains. This allows resolvers to cryptographically validate that DNS responses have not been altered.

It is important to note, however, that DNSSEC protects the integrity of data. It does not prevent attackers from registering malicious domains or launching phishing campaigns. In brief, what DNSSEC does is it protects against DNS spoofing, cache poisoning, and response tampering. Please note that DNSSEC does not stop phishing domains, it does not stop malicious domain registration, and neither does it stop social engineering. The focus here is more on integrity than intent.

Next slide, please. Thank you. Let's talk briefly about the DNS reality gap. DNS exists technically, but operationally, adoption is

inconsistent. The reason could be due to the complexity of setting up DNSSEC, operational risk, lack of training, or legacy systems. Some organizations fear a misconfiguration outage, so they delay deployment. As SOC teams, what we have learned so far is that DNS logs are extremely valuable in investigations.

Next slide, please.

RUSS MUNDY: You have five minutes, please.

OLABIMPE ADELORO: All right. Thank you. DNS visibility is the ultimate operational feedback. What we have learned so far is that DNS visibility is very critical. DNS logs provide the most valuable forensic evidence during an active investigation. Early DNS monitoring is the only way to reduce eventual damage. Acknowledgment that DNS security maturity varies significantly across organizations is something we have learned.

Next slide, please. Since DNS telemetry often provides early indicators of compromise and helps organizations monitor DNS, organizations that monitor DNS detect attacks faster. The recommendation is thus far basically for three categories: for organizations, SOC teams, and the ecosystem. Organizations, we encourage you to deploy DNSSEC, monitor DNS logs, and fix email authentication. For SOC teams, please integrate DNS monitoring tools, track newly registered domains, and investigate suspicious

DNS patterns. For the ecosystem at large, we encourage better awareness and easier DNSSEC deployment. For registrars, we encourage registrar abuse monitoring.

In these presentations, we have been able to talk about DNS and the fact that it is not a passive infrastructure; it is active security. It is an active security layer. If DNS security is weak, attackers gain an advantage. If DNS security is strong, many attacks will fail at an early stage. Next slide, please. Dear viewers, may I leave you with these final notes. Attackers do not break DNS. They exploit gaps in our implementation. DNS is not just infrastructure; it is a primary attack vector. Strengthening DNS security through visibility, proper configuration, and DNSSEC adoption is essential to protecting the global internet ecosystem. Thank you very much for having me.

RUSS MUNDY: Thank you so much for this presentation. We really appreciate this. I see we do have a question in the Zoom room. LJ, why don't you go ahead with your question?

LJ: Thanks, Russ. On your second to last slide on bridging the operational reality gap, you highlighted the risks and gave some recommendations. You highlighted in your presentation risks of self-inflicted outages from DNS misconfigurations as a major barrier to adoption. For organizations that are just beginning their DNSSEC journey, what would be a phased implementation

approach, or tools, or testing practices that you have seen work best in practice to enable signing and resolver validation safely without disrupting operations? That would be helpful if you can. If it's a long answer, we can take it offline.

OLABIMPE ADELORO Sorry, can I get the last part, please?

LJ The last part is: what would be some phased implementation approaches, tools, and testing practices that you have seen work best in practice to enable signing and resolver validation safely without disrupting operations for small organizations that are just beginning their DNSSEC journey?

OLABIMPE ADELORO Thank you very much. I appreciate the question. The best thing I would recommend is monitoring. It is good for you to know what you have in terms of taking stock of your assets. When you implement DNSSEC, it is not enough to put it in place. Please kindly monitor the logs. When logs are monitored, then in case there is any attack, it can be worked on, handled, and prevented further. Another thing I would recommend is awareness. When education and proper education is given, and we know why we need DNSSEC, then it will be helpful for the SOC team. That is why in my last slide, you will notice the recommendation because what we talk more about is awareness and easier deployment. DNSSEC deployment

could take a bit of effort. However, we should not shy away because it takes effort; we should do the right thing. That would be my recommendation.

RUSS MUNDY Thank you. Do we have any questions? Isaac, go ahead.

ISAAC SAKYI Thank you very much. I would like to know what constitutes misconfiguration. How do you classify misconfiguration? Also, you suggested that SOC teams investigate suspicious...

OLABIMPE ADELORO Can you please speak a bit louder?

ISAAC SAKYI I am asking what in your estimation constitutes DNS misconfigurations?

OLABIMPE ADELORO What constitutes DNS misconfiguration? I got you. Thank you. Generally speaking in layman's language, if configuration is not done the right way or is not properly adopted, then it is misconfiguration. However, if you want to go deeper, if you are talking about configuration that lacks SPF records at all or has overly permissive SPF, if DKIM is not enabled, or you ignore or miss DMARC, then you would be said to be having misconfiguration.

RUSS MUNDY Okay. I think hopefully that answers your question, Isaac. We are about out of time. If anyone has any last-minute urgent questions, we possibly could take a short one. I do not see any in the Zoom room. Dan, Kathy, is there anyone in the room?

DAN GLUCK Not that I can see.

RUSS MUNDY Thank you. Thank you to all of our presenters for all three of the sessions of the workshop. It was a really interesting set, and I know it is work to put it together and to do the presentations. On behalf of the program committee, I want to thank you very much and hope that you found it rewarding also to do so. In the future, consider to come back and do more from what you have learned and what you have presented. I think with that, I will turn it back to Kathy for closing the session.

KATHY SCHNITT Thank you, Russ. Thank you to all the presenters, and this concludes the session. Please stop the recording.

[END OF TRANSCRIPTION]