ICANN72 | Virtual Annual General Meeting – DNSSEC and Security Workshop (3 of 3)
Wednesday, October 27, 2021 – 16:30 to 17:30 PDT

KATHY SCHNITT: Hello and welcome to the DNSSEC and Security Workshop Part Three of Three. My name is Kathy and I am the remote participation manager for this session. Please note that this session is being recorded and follows the ICANN Expected Standards of Behavior. We will take questions and comments as we did for parts one and two and I will post those instructions in the chat pod. And with that, I want to hand the floor over to our moderator for this session, Russ Mundy.

RUSS MUNDY: Thank you very much, Kathy. And thank you, everyone who is here in attendance. We have a good collection of folks for our final session here. I think it is quite an interesting couple of presentations that both relate to what happens when the unexpected happens and how does one go about deal with things of that nature?

So our first presentation is from Kim Davies, who is the person who is very much on the pointy end of the spear when it comes to the DNSSEC for the root, and KSK management, and all of those type of activities. He has been through quite an exciting time in the last couple years. His presentation today is going to give us some good insight as to what occurs and what happens when one of the shared keyholders is no longer available and other related things. So, Kim, over to you.

KIM DAVIES: Thanks, Russ. Hi, everyone. My name is Kim Davies. I head up the IANA team and part of the responsibilities our team has is for the operation and the security of the root zone KSK. Today I wanted to use the opportunity to give you a little bit of insight into the recovery key shareholder system that we have—what it is, how it works, and so forth. And then I'm going to pivot to talking about what is arguably our most high-profile recovery key shareholder, Dan Kaminsky. Next slide, please.

So just by way of a little primer before we get into the details, everyone is familiar with this layout. There's seven people who control the Internet, right? Obviously, I'm joking here but this is the level of sophistication most people have when they think about the root zone KSK, if they think about it at all, is that there's seven people. Each have some fractional share of the system, and you need a number of them to come together to activate the KSK—in this case, three of the seven. That's kind of correct. At a high level, that suffices to get the concept across. But the reality is a little more complex than that. Next slide, please.

This diagram is a bit more reflective of reality. In truth, there are two sets of seven keyholders. We have four HSMs, the hardware security modules that store the KSK. Two of those HSMs are in one facility on one side of the United States. The other two HSMs are in another facility on the opposite side of the US.

Each of the set of seven is assigned to a facility. So they have the ability to unlock the HSM that is in their assigned facility. They can't fulfill the role in the opposite facility but they can fulfill the role in the same facility, on either of the HSMs. So we have some redundancy there. If one of the HSMs fails in one facility, we have an alternate HSM. And if both of them fail or are unavailable for whatever reason, we have duplicates on the other side of the country.

So that's roughly the model that we have in place today, with those 14 trusted community representatives, referred to as cryptographic officers, fulfilling that role roughly every quarter, coming together to sign the ZSKs for the root. Next slide, please.

So if we drill down one extra more layer, this is actually much closer to reality than the previous slide This is the full complement of active trusted community representatives that we have. There is actually 21 trusted community representatives. There's the 14 I just mentioned. But the additional seven is what we call the recovery key shareholders, which is the top of today's presentation.

So in each facility, in addition to the two HSMs depicted, we also have an encrypted backup. It's an export of the contents of the HSMs. It's then wrapped in a key. And the decryption key for that backup is what is shared between the seven recovery key shareholders shown on the righthand side of the diagram. Importantly, they don't have key or the backup, I should say. They have the decryption key for the backup. Next slide, please.

So what is the recovery key shareholders and why do we need them? Essentially, it's part of our disaster recovery planning. They don't play an active role in the day-to-day key ceremonies that we conduct.

They're really there to be activated in the event there's some kind of significant disaster that affects operation of our KSK equipment.

The idea is, basically, should the contents of the HSM be corrupted, or unavailable, or unrecoverable in some way, this provides a mechanism that we can take the backup, basically restore the contents either to an existing HSM or newly-procured HSM, reconstitute its contents and then restore service. And we do that with the encrypted backups that the RKSHs or the recovery key shareholders can unlock. Because we are unable to unencrypt the backups alone, we need their participation. That's one of the fundamental protections around the backups that we've generated.

However, as I mentioned, they don't attend ceremonies regularly. This means that we don't see them very often. We don't get to put into play some of the things we do regularly with normal participants. So when we have our key ceremonies every three months, we try to exercise all the smart cards that are held by those cryptographic officers. We don't have that luxury with the recovery key shareholders so we don't regularly exercise the equipment that they hold.

Because they're not used regularly and are only there to be retrieved in the event of an emergency, they're typically expected to be stored in a more secure location—for example, safe deposit box in a bank. And essentially, our ongoing operating procedure around these is through an annual attestation. We effectively reach out to each of those seven on an annual basis. We ask them to confirm the safety of those credentials and to make an annual attestation to that effect for our records. That's really what happens with the recovery key shareholders. Next slide, please.

So Dan Kaminsky. He was one of the seven recovery key shareholders for the root zone. I think Dan needs on introduction here. Obviously, he was pivotal in raising awareness in DNSSEC and really helped make the business case for signing the root zone. If it wasn't for his discoveries, then I think we would have had a greater challenge getting to where we are today.

So Dan, in his role as a recovery key shareholder, much like the other six that hold that role, he attended our first key ceremony back in 2010, in Culpepper, Virginia. That was the ceremony where we first created the initial KSK. But in this role, he's never actually attended a key ceremony since.

Essentially, the reason that he hasn't returned, or any of the other recovery key shareholders, is we've never had an emergency of the kind that we would need to make use of a recovery key shareholder. HSMs have functioned as expected. We've never needed to dip into that pool of volunteers to make use of those exported backups.

In his role, apart from attending ceremony one, as I mentioned, he faithfully made his annual attestations but we never actually needed to use his key share. As we all learned earlier this year, he passed away in April after a long-term illness. Obviously, this was a surprise and somewhat devastating at the time. And there was certainly an outflowing of recognition for all the work that Dan did when that news came to pass.

For us, though, in addition to that, it also presented a challenge. It was exercising a contingency that we hadn't necessarily considered in great depth, which is what to do when a recovery key shareholder becomes permanently unavailable. So that's what I'm going to turn to next in this presentation. Next slide, please.

This is the story of recovering the recovery key share. And I'm just going to walk you through some of what happened, just for background, for interest, insight. And then we'll talk about lessons learned and next steps after this. In the beginning, obviously we heard the news, much like everyone else—mailing list and so forth, and then in the media.

The first step was to make contact with people that might be familiar with the situation and our first thought was with his family. We didn't reach out right away. We waited a few weeks. But once that time had passed, we didn't have any leads, initially. But we were able to make contact with his extended family. Thankfully, interviews were being given about Dan's role. We came to know some of his extended family through that. We reached out to them.

They, in turn, redirected us to Dan's attorney. Initially, we had to explain what we're about, why we were reaching out. I have to imagine that they were receiving a lot of contacts and part of the attorney's role is sorting out legitimate and illegitimate attempts to reach the family.

But eventually, we were able to make contact with the family and establish the situation, explain to them what was going on and so forth. So in the end, we were directly speaking to Dan's parents. And they were actually quite familiar with what ICANN did. They were very familiar with his role. Apparently, Dan had spoken to them about it, and explained how it works, and so forth. So they weren't completely naïve or unfamiliar with what he did. And I think, most importantly, they were really keen to honor his legacy and were very supportive to us in what we needed to reobtain his key share in our dialog with them.

So that was a good first step. But the next step was probably the trickiest of them all, which is ascertaining the location of the key share. Fundamentally, Dan, in his role, being a recovery key shareholder, was to keep his key share secret, and private, and keep it basically in an undisclosed location. And he did that. We did not know where he stored it. His family did not know where he stored it. And that was the next step in this process.

First order of business, his personal effects. Didn't see any evidence of it there—his home and so forth. But after speaking with his family and talking it out, we came to the realization, based on a variety of factors, that it was quite likely in a safety deposit box in another city. We didn't know which bank it would be at. We didn't know which branch it would be at. We were pretty confident it would be in New York, in fact. His parents said he was living in San Francisco at that time.

So this process, this phase, played out for many, many months. Probate, where his affairs were settled, took many months. Eventually, we came to learn of the bank or the branch where it probably was located. But attaining access to that branch by his estate was a laborious process. It wasn't immediately possible. But eventually, access was available.

His family was able to go to the bank. They actually flew to New York specifically for this purpose and this purpose alone. They went to the bank. His safe deposit box was drilled. They obtained the contents. And sure enough, it was exactly as we expected.

They flew back to San Francisco with the package. I met them the following day. And then, the following day after that, I traveled back to Los Angeles, where ICANN is headquartered. And myself, along with our team, we stored his key share in a secure facility within our office.

So that brings us to where we are today. We now have his key share in our possession. It’s still wrapped in the tamper-evident bag that was issued to him in 2010. We have full confidence that it hasn’t been tampered with in any way. And it is now sitting in a safe, waiting for us to take the next step.

Obviously, our plans are a little complicated due to COVID. This is not quite how we would desire to do it. But our travel policy and so forth and restrictions related to COVID, our current thinking is we expect to identify and induct a new replacement for Dan as a recovery key shareholder at our next key ceremony, which is in the middle of Q1, 2022, so probably roughly around the February timeframe. Next slide, please.

So that’s what’s led us up today. Here’s some of the food for thought that has gone into this. Basically, what have we learned and what are we thinking about as a result of what happened?

Firstly, the most immediate gap that comes to mind is that in those initial days and initial weeks, we didn't have a point of contact to reach out to for Dan. In a sense, we were lucky that there was so much media coverage and we came to learn family members that we commenting to the media and were able to reach out to them. But I have to think that for all the TCRs, we wouldn't necessarily have that visibility and that ability to quickly identify a point of contact.

So I think one of the first orders of business is establishing emergency contacts or next of kind for each of our trusted community representatives. This will help obtain contact with them but it isn't necessarily the complete solution to the problem.

There is a question about even if we can reach the family, if the TCR's job is to keep their credentials safe, and secure, and private, the family won't necessarily know where it is stored. This might involve setting an expectation that the TCR leaves enough clues or some pointer that might help them recover it in the case of an emergency.

Or it could be codifying an assumption that they just aren't recoverable—that we should assume that this kind of event, unless it's critically ascertainable where they're located, that we just assume that they're unrecoverable and make adjustments accordingly.

I think the next thing to think about is retesting the recovery key shareholders. I mentioned before that we saw them in 2010. We haven't seen them since. The smart cards don't last forever. It's probably useful to have them reconvene every, let's say, 5 to 10 years, retest their credentials, issue new credentials.

This is, in fact, something that was already on our roadmap. We had actually planned to do this in 2020–2021 but COVID had really put a pause on that activity. So as soon as the travel scenario allows for it, I think we would like to resume that. So we would like to see this and then have some kind of regular cadence where recovery key shareholders come together to regenerate their credentials.

I think the more fundamental question is whether this model is even useful. Is it a useful paradigm? Bear in mind, the recovery key shareholders hold the decryption key for a backup but the backup is stored in the same facilities as the HSMs that it backs up. And there's relatively few disaster recovery scenarios you can think of that would simultaneously make our entire fleet of HSMs inoperable. But at the same time, the backups would be operable and we could use them to restore HSM function.

So is it tailored correctly to the kinds of disaster recoveries we can conceive of? Is there a better way to do it? I think these are the questions we will want to explore a little bit. If the recovery key shareholder mechanism is just not a viable scenario for any kind of realistic threat, then maybe it's just not a useful structure to have at all. So is there a way to have disaster recovery that doesn't have this kind of fate sharing as a fundamental component of it.

So that leads to questions about a more fundamental rethink of trust dispersion. Let's set aside COVID impacts, in which has really hampered our ability to travel internationally and has had a significant impact on the way we operate the KSK.

Is the recovery key shareholder model the right one? Would it be better to store our backups in alternate locations, not in our key management facilities? Or would it be better to invest our time and effort into, for example, a third KMF rather than into backups? That would provide an additional level of redundancy but it also increases the attack surface as well. So everything here is a tradeoff of different complexities and it's something that we will want to explore and talk about in the coming years.

So that's a bit of whirlwind tour through some of the thoughts that we've had. And very welcome to hear the thoughts that the community has on where we should focus—perhaps additional ideas the perhaps we haven't thought of yet. Not really a call to action here. Just a high-level brainstorming of some of the things that have come to mind going through this process. So as always, we're very welcome to get input into the model. The model has been constantly evolving over the last 11 years and we expect it to evolve forward as well.

So with that, that's the end of the presentation. I'm not sure if we're holding questions or comments until the end or now. But yeah. Thanks for listening.

RUSS MUNDY: Thanks very much, Kim. I think since the presentations are about two different areas, it might be best to go ahead and take our questions now. There's two in the Q&A pod and Steve has his hand up. But the Q&A pod questions came first. So can you look at the Q&A pod, Kim, and see if you can address those, please?

KIM DAVIES: I certainly can. The first question was, “Why was it necessary to recover the key share? How many of the seven key shares are required for disaster recovery? What would it take to roll the backup key and hand out new key shares to the other recovery keyholders?”

That’s a very good question. That was actually, originally, in my slide deck but I stripped it out in terms of brevity. It wasn’t absolutely necessary to recover the key share in terms of numerically satisfying the quorum for the key shares. We do only need four of the seven for disaster recovery. So in effect, with Dan Kaminsky’s key share unavailable, we now had a quorum of four of six.

I think the bigger challenge was really related to COVID. Should we need to pull the trigger on this, we know COVID was limiting international travel. That key share in particular was based in the US so that’s one less thing to worry about. And also regenerating the keyset does require all seven of the recovery key shareholders to return to generate a new set. Again, because of COVID, it’s not really viable—at least it wasn’t at that time—to have all seven come together to regenerate a new keyset.

So it wasn’t a fundamental problem that we were down one of the seven for an extended period of time. It reduced our margins a little bit but the other options were not as good as the ones that we had available. So our priority was definitely on recovering his key share if we could.

I think if it wasn't for COVID, maybe we would have made a different decision. Maybe we would have decided, after a month or two of lack of success in recovering his key share, that we should recall all seven, or the six and induct a new one, to generate a new set. But that wasn't something we thought was practical at the time.

The next question, "My apologies for a bit off-topic question but may I ask you to tell what the procedure and what exactly a TLD registry should do, if it's possible at all, for making a second backup DS record in the root zone, which would have an appropriate DNSKEY in a TLD zone."

I'm just going to answer this briefly. In terms of IANA's procedures for the root zone, if you wish to list an additional DS record in the root zone, we do ask that the matching DNSKEY be in the apex of the TLD zone. Happy to go into more detail on that in the chat or what have you but that's the basic requirement that we have.

I see Steve's hand's up. Steve?

STEVE CROCKER: Thank you very much. I like Peter's question because it had occurred to me too and I was pleased to see that he asked it. Let me take Peter's question and enlarge it, leveraging some of the things you said. What you're finding—what you've found over time in both this experience and some others—is that the nature of the problems that you encounter don't necessarily match what you might have thought when the whole system was put together around 2010, before, when it was being planned.

So clearly, one general question to ask is, is it time to rethink along the lines that you said, about fate sharing and so forth, but from a more comprehensive view about the whole design, not just the key recovery portion of the system, about where the problems arise and then what the mitigations or responses are to that? So that's part one of what I want to say.

The other part is there's a couple of numbers deeply buried in there, deeply enmeshed in the system. Seven is a magic number. Everybody likes the number seven. And you said you have four out of seven to do the key management recovery.

My understanding, from a common-sense point of view, is that the selection of those numbers try to balance two basic factors. One basic factor is prevention of joint bad action—a cooperative undermining of the system. So that's why you don't have a single person able to operate the system. You want to have a coalition of people so you don't get a conspiracy of some sort.

The other big factor is prevention of denial of service. You could improve the resistance to the cooperative undermining by increasing the number of people who have to participate. You could say all seven have to show up but then you would leave yourself open to a denial of service if any one person is unavailable.

So just from those basic principles, the question could be asked, “So what are the right numbers? Is seven the right number? Is four the right number? And how could you figure out what the right numbers out to be. So you ought to have a model based upon what your estimate of the probabilities are of conspiratorial collusion—that’s the word I was trying to go for—on the one hand, versus absentee for one reason or another.

Better yet, we’ve gone down the road quite a ways. We now have more than a decade of actual experience. I don’t think we have any data that gives us a way of estimating the probability of collusion but we certainly have some data on absenteeism and unavailability. And it would be interesting to see what the relationship is between the protection levels that you think you are trying to … Again, you want the probability of collusion to be what? Less than one in a million? Less than one in a billion? You want the probability of denial of service because there’s too many people absent to be less than 1 in 1,000? 1 in 10,000?

I did a lightweight exercise some years ago, trying to run those numbers, not for a system as complicated as you have but just for a single level, just if you had k out of n people and none of the other layers of this. And I’ll just leave it with a little tease. The numbers didn’t match at all. If you wanted the level of protection that you might guess at, the numbers were going to have to be enormously bigger.

So just a reality check on what are the design goals of the system and then what is the actual behavior of it—has it been—in addition to reacting to each individual event and saying, "Oh. We didn't think of that. Therefore, we should tweak the system some."

And that kind of analysis may be best done by people who are not connected directly to operational activities on a day-to-day basis but are more theoretical in their approach or living in other quarters and so forth. So that's my long-winded comment.

KIM DAVIES: Thanks, Steve. You touch on some really great points. And certainly, many of them, we've talked about and discussed. I think one of the great features of the current model is we do have these trusted community representatives. Maybe it's an unofficial role that they have but they act as a sounding board for the evolution of the system. And we're in constant dialog with them about adjustments and changes we want to make to the way we operate. And they act as, like I said, a sounding board for the community on what makes sense and what doesn't.

Part of the things we've proposed over the years, one of them involved adjusting the way that they're provisioned in a way that there was more redundancy but that got shut down at the time by that group so we didn't pursue it any further.

And indeed, if we took a clean sheet approach and rearchitected the model for scratch based on what we know now and based on the current environment, particularly in light of the fact that international travel is just not as vital as it was in 2010—at least for the moment—then maybe we'd come up with a relatively different design.

I think that that exercise is useful to explore alternatives. But I would also be mindful that we don't convey a false impression that the current model is fundamentally broken and doesn't work, therefore we need to redesign it from scratch. I think there are various accommodations we've made. But regardless, it's actually still working quite well in its current form.

STEVE CROCKER: Let me offer one small comment about not knowing where Dan had stashed his key. I could imagine making a requirement that everybody write down and share, in a cryptographically-secure way, where they have put that. By cryptographically-secure, have it divided up so that it requires several people to decrypt that piece of information and an agreement that you're going to do that so that you have a recovery process that has the political aspect of needing a certain number of people to break the secrecy, just of where that stuff is stored. That would have saved you quite a bit of time.

And that's all just software-distributed, doesn't require travel. And then you'd wind up with that piece of information and then speed up your recovery process.

RUSS MUNDY: So, Kim, do you have any final thoughts? We hope you're able to stick around if more people have questions at the end after Stefan's presentation. But did you have anything you want to close with?

KIM DAVIES: Not particularly. Hopefully this was insightful, gave you a little bit of a taste of some of the operational stuff we deal with. Perhaps this is the most interesting caper we've had for the last 11 years but our team is dealing with all sorts of things on fairly regular basis that perhaps you don't get to see and are part and parcel of operations, I suppose. But thanks, everyone, for listening. If you have additional ideas, very welcome to hear them. Thank you.

RUSS MUNDY: Thanks very much, Kim. And if people do have additional suggestions, what would be the best thing to do? Send you e-mail or something of that nature?

KIM DAVIES: Yeah, certainly. E-mail is fine, kim.davies@iana.org. Thank you.

RUSS MUNDY: Okay. Great. Thanks very much. Now let's go on to Stefan and we'll hear a different slant on disaster recovery kinds of activities. So, Stefan, please.

STEFAN UBBINK: Thank you. Hello. My name is Stefan Ubbink and I am a DNS and system engineer at SIDN. SIDN is the registry for .NL zone, which is the ccTLD for the Netherlands. SIDN also has some other TLDs under its guard and I will be talking about disaster recovery with DNSSEC. This is an overview of the things I'm going to talk about in this presentation.

Why this talk? Earlier this year, we replaced our HSMs to get a better understanding of what we had to do in case of an emergency. We created some what-if scenarios. One of the scenarios was what do we do when we lose all our keys? Since we make a backup of our keys, we could recover it from that incident by restoring the keys from a backup. Or not because we noticed that we didn't make backups of every ZSK that was produced by [our setup]. So there was a possible issue.

To be able to show you how our setup is set up, it's this slide to get an understanding of how it works. We have two datacenters, one on the right and one on the left. And we have one active signer, which signs the zone, and standby signer, which can be used in case of emergency.

We have two HSMs in each datacenter but they are on the same network and they are HA-connected. So if someone would, say, destroy this key to the HSM, then those keys would be lost. That's an issue. So if someone wouldn't do that, you would have an impact because you no longer have those keys.

And if you lose those keys, you will start with the loss of service because no updates for the zone can be done without the keys. And you cannot do much until the keys are back. And if you don't take any action, the domains will start to fill because the RRSIGs will expire. The last RRSIG that will expire will be the one of the server records.

If you cannot restore the ZSK, you have to introduce a new ZSK. This will mean that all DNSSEC-signed zones will be unavailable, or at last for .NL, it will be unavailable. And as you can see from the graph from APNIC, about 60% of the resolvers in the Netherlands are validating resolvers, which means there is a lot of people who noticed this issue. And this would make a very bad headline in the news for us.

What can you do to prevent this? A sort of prevention is removing DS from the parent. That still has impact because you no longer have the possibility to use DANE services. And the A records will no longer work and SSHFP resource records will no longer work. So that’s not really an option. If you do it the wrong way, you will have a similar issue that Slack.com had a few weeks ago.

To be able to restore from a backup, we would have to make sure every time a ZSK is created, a new backup has to be created. For .NL, this is every 90 days. Because we also have other TLDs, it would mean a lot of backups. And since this is a manual process at our registry, it would take a lot of time and effort to do this all the time. Besides this, you will have to have the right procedure to do everything—to have a consistent backup of all the keys you need.

After this after this what-if scenario, we came with improvements. Since we are using OpenDNSSEC for our signing process, we implemented the required backup setting in OpenDNSSEC with another setting which is needed, the AutomaticKeyGenerationPeriod. We set it to one year so we have keys for one year in our HSMs, and we can create a backup and don't have to do a backup for every ZSK rollover when it happens because we have a backup for one year. When we do backups, we have all those keys available.

We scheduled tickets in our ticketing system to ensure that we will start a new backup when it's needed. And we implemented some checks to alert us when we forgot or didn't know that we had to do this. So we should be sure that we will do this backup.

You have a limited time to react to the loss of keys. This is based on the refresh time and re-sign time in the OpenDNSSEC configuration. This is a picture of how it's calculated. If you lose your ZSK, you have the refresh period minus the re-sign period to act. And for us, we now have a reaction time of about seven days. So if we lose the keys, we got notified immediately because can no longer publish zones. And then we can start the process of restoring the keys from the backup.

I would like to thank. Berry van Halderen from NLnet labs for providing a very excellent explanation about the possible impacts and ways to act in this what-if scenario. I would also like to think SIDN colleagues for helping improving our DNSSEC. If there are any questions, please ask them in the Q&A pod. Thank you very much.

RUSS MUNDY: Thank you very much, Stefan. That was really an interesting look into something that dramatically affects every signed TLD. We have a question in the pod from Peter Thomassen. So could you open that and see if you can address that, please?

STEFAN UBBINK: I’m sorry but I don't see any question at the moment. Maybe you can read it out loud.

RUSS MUNDY: Sure. Peter Thomassen asks, “The backup strategies all were about the ZSK. Wouldn’t the KSK be more critical?”

STEFAN UBBINK: Peter, thank you for your question. The issue in our situation was that we always created a backup of the KSK. And the KSK was always available. But since the ZSK was rolled automatically, that would walk out of our backup, I would say. So that’s why my presentation is about the focus on the ZSK. I hope this answers your question.

RUSS MUNDY: And I suspect, from what you had in your presentation, that a part of the challenge of getting all of the ZSKs backed up is the quantity of them and the frequency of the rolling that, with an annual process, it was really hard to effect the backup there. Okay. Great. Thank you.

So do we have other questions from folks for Stefan or any more for Kim? Well, we are approaching the end of the workshop. Oh, Ulrich. Yes. You’re on the panelists side, I think. Please go ahead.

ULRICH WISSER: Yes. Hi. Yes, Stefan. I would like to ask you. You said that the signatures are valid for seven days. How did you come to choose seven days?

STEFAN UBBINK: I said that we had a recovery time of seven days. The seven days is the refresh time minus the resign time. We calculated that. We want to be able to have a holiday—Christmas holidays, and weekend, and etc. So seven days was a good fit for us to be able to restore everything, even when there are weekends and not everyone is available.

RUSS MUNDY: So as in all of these—and certainly, we heard it in Kim's presentation—there's lots of operational procedures involved, as well as has been pointed out by Mark here in the Q&A pod, there are written formal procedures that need to be followed, that all need to come together and work cohesively. So those factors have to be rolled in together. And coming up with what's the right balance is challenging. Even, as some of the questions earlier reflected, in just about everything that you would look at operationally, there are trade-offs involved. So this is good. Do we have more thoughts, questions by anyone here?

Okay. Well, at this point, I want to very much thank Kim and Stefan for the presentations. Very interesting to try to get more insight into how people are handling day-to-day challenges, both expected and unexpected. These are the type of topics that we want to be sure to cover in our DNSSEC Workshop and Security Workshops. So I thank both or two presenters here, our earlier panels.

And I think, if we have no more questions, then I want to also extend great thanks to our tech support team and our staff support team because I can truly say, as one of the longtime program committee people, these would not happen without the wonderful support that we get from Kathy and the other support staff.

Thank you very much. And please think about what you'd like to hear and talk about in the next workshop, for the next meeting, and be watching for our call for participation because we want to continue to present these varied and interesting topics. So thank you. I'll close this panel and turn it back to Kathy.

KATHY SCHNITT: Thank you very much, Russ. I too want to thank all our presenters and panelists for their very interesting presentations and also for joining us at some not-so-pleasant hours for some. I want to thank the Program Planning Committee. Russ, you and the team just put together another fantastic agenda. And it's always a pleasure to work with you and to come up to the point of these workshops.

I'll thank my colleagues, Kim and Andrew, for always supporting me, as well as our techs, [Sare] and Scott. Thank you so much for having this run so smoothly. We'll see you at ICANN 73, either virtual or face-to-face. We don't know yet. And with that, please stop the recording. Thank you, everyone.