ICANN73 | Virtual Community Forum – Tech Day (3 of 3)
Monday, March 7, 2022 – 12:30 to 14:00 AST

KIMBERLY CARLSON: Hi, everyone. And welcome back to Part 3 of Tech Day. As a reminder, this session is recorded and follows the ICANN Expected Standards of Behavior. Participants are welcome to post their questions or comments using the Q&A pod or via Zoom chat.

Again, thanks for joining. And I'll hand the call back over to Dr. Lisse.

EBERHARD LISSE: Thank you very much. Welcome to our session. Now at the first, we have a presentation from Kathleen Moriarty and Paul Vixie about the impact of more pervasive and corruption see about the impact a more pervasive encryption on correspondence communications can have.

Kathleen and Paul, you have the floor. We also have enough time, so we don't need to rush anything.

PAUL VIXIE: That’s good. All right. So, good morning to those in my time zone. Good evening to those around the world. Suzanne Woolf told me many years ago, and I’ve never forgotten, that it's always 5:00 in the morning somewhere. And I know that's difficult for some of you.

Today we're going to talk a little bit, ultimately, about DoH and so forth. But here we really want to talk about the impact on network operators and other on-path actors who would otherwise be trusted who will, nevertheless, be potentially disintermediated by some protocols that are coming out of the ATF, more or less in light of the Snowden disclosures of 2013.

So my co-speaker today is Kathleen Moriarty. And as you see on your screen, she is the CTO for Center for Internet Security. But she also, in the past, has served as the area director for security in the IETF. So I will come back when called for, but Kathleen is going to drive this morning.

KATHLEEN MORIARTY: Thank you, Paul. And yes, while I was an area director, it was just after Snowden. I started in 2014 through 2018, which were really key times in this transition.

And so—if you can advance to the next slide, please—for quite a while I was talking about encryption being a trend. It's gone beyond a trend now, and I would say it's reached the point of being called an inflection point. Something that will have really great impacts on industry, on protocols, and how we perform management—security, network, and other types of management.

So Paul made some great points on just encrypted traffic increasing. Snowden itself, for HTTP, right after that event, jumped us up to about 30% encryption in 2013. And that was the encryption people could deal with. So that was the sophisticated operators that could manage

installing by hand a certificate and associated key pair to be able to host an encrypted web server.

That big jump to 80% that we've seen ... And it's been about 80% for a few years, depending on what region you're in. And you can go look at the Let's Encrypt statistics. And that shows you the statistics for Mozilla browsers. That bigger jump is because of free certificates and keys through Let's Encrypt, but also because of the ACME (Automated Certificate Management Environment) protocol.

And in terms of shifts that we've seen in the IETF and elsewhere because of the Snowden revelation, the ability to publish a standard such as ACME was a large part of that shift. The ability to automate certificate management. We tried it many times in the past, but the purists— meaning those who would not accept any level of degradation in security in favor of some security— prevented such protocols from going forward.

The realization that some security was better than none and doing something to automate it would make a big difference is what allowed a protocol like ACME to get published and now is heavily utilized. We've seen that shift for a number of years now. So for HTTP traffic on the Internet, this is something we've all adapted to.

We expect, on the big Internet, that web traffic is largely encrypted. We don't expect to be able to intercept it. And any management is done on the layers above, on the IP stack, the IPv6. We've even heard, one talk today gave some statistics on their particular project in IPv6 adoption.

For service provider networks, it's much higher. And it makes sense because it has things like extension headers and IP flow labels. So there's the ability to do some management where you really lack that ability with IPv4.

And this adoption is starting to shift into internal networks, but I'll get a little bit deeper into that because the inflection point for internal networks is really starting to have an impact. And we'll get into that as well.

Increased use of QUIC is also of concern on the network. And this because signaling information is hidden—all but one bit, the spin bit. But again, on the larger Internet, service providers rely on IP and TCP headers. And this should not pose as much of a problem. And I'll dive into that more and explain why in a subsequent slide.

The other big change we have a zero trust architecture. Now this goes along with that inflection point to the encryption because one of the key tenants of zero trust architecture is pervasive use of encryption. So not just in transit, but also at rest.

The key idea behind a zero trust architecture is that you want to have detection as early in the kill chain as possible. So when an attacker begins their attack, you want to catch them as early as possible, whether that be reconnaissance or before an exploit is executed on your network. Or if they've gotten that far, before it can spread with lateral movement or privilege escalation and cause actual damage to your systems.

And so this use of encryption is very important because if your system is infiltrated and data is discovered and exfiltrated, if it's encrypted there's less chance of IP loss, intellectual property loss, or use of the data for disruption campaigns and such.

Other aspects of zero trust that are important to understand in this really large shift we're about to realize in the next, probably five years. It's been a bit more accelerated than I had expected. I would have projected 5 to 10 years, but we're seeing quite a bit of motivations to go a bit faster.

And so dynamic authentication. If you think of this on an application level, you might be a user on an HR system. And to access an employee record, you might be prompted again for authentication. Zero trust goes beyond that because it's also concerned with authentication of components. So one module would authenticate to another module, and it would require re-authentication.

So let's say an attacker infiltrate some part of an application. This might prevent them from surviving because they would fail the re-authentication attempt if it's a strong enough of a protocol in use for authentication.

Similar with verification. So verification is, now what are all the bits on my system? What are the configuration parameters? What are my policies? What are the measurements I expect to see? Does it all verify to be as expected? And if it doesn't, what do I do?

In the case of hardware when you boot from a root of trust, if the firmware doesn't match up to an expected value, there's a reset on that process. And a system could lock up until it boots as expected.

Now if we bring this up the stack which we're seeing with workload assurance, then you have the ability to manage your workloads, your hardware, and potentially your software to expected values. And then if something doesn't meet the expectation, there's an immediate action taken.

Additionally, zero trust calls for locking everything. And this also aids with your detection capabilities. So this is a large shift from that perimeter Defense in Depth that we've had for the last 25 years, that quite frankly hasn't really been working, to more pervasive detection and prevention. So you're getting much more granular within systems and it represents just an enormous shift, but also an enormous opportunity. If you could switch to the next slide, please.

In 2018, I believe, is when we published RFC 8404. And that was myself and Al Morton. And we thought it was going to be largely just simple documents. That we'd go out, we'd talked to operators—all different types—and we'd figure out how will they be impacted as encryption use increases.

For Internet service provider level, there wasn't much of an impact. And there's a few key reasons for that. They relied primarily on TCP and IP headers, and many of them had already shifted to IPv6. So if they needed some capabilities, they had them.

Then the other part is that any performance metrics were using dedicated packets and dedicated protocols. So they have measurement protocols that they use specific to the service provider that has their own infrastructure, doesn't utilize customer connections, and they're able to test for availability performance without impacting customer's data or intercepting any traffic.

And so that was very good to learn in the process and to be able to share that knowledge out for anyone that was concerned if there was data being used in that way. This also represented a big shift for content delivery networks to end-to-end only. But that was desired by many of the content delivery networks. They wanted to own their content. And by having end-to-end encryption, it allowed them to own their content.

What was the impact? Well, caching servers couldn't relay data out for them and speed up the time. Now this has a much larger impact in a mobile network where you have devices and at the edge, there's quite a bit of caching happening.

So the interesting part is that we as an industry have developed lots of solutions to handle these changes, and we're shifting things left—if you've heard that term—where the vendor is taking more responsibility for security, or we're shifting them more further up in stream so that we require a much smaller number of experts to perform a much larger task.

So this morning's opening talk gave an excellent example where the PIR has some funded research to look at how we take care of preventing malicious domains. And the research pushed this back to the registrar

and the registries. And the efficacy of that, if I got the statistic get down correctly, was 90%. That's a pretty impressive shift, where many of us are still relying on blocklists for DNS to make sure that customers and less capable entities are not going to sites that have malware on them.

So at Center for Internet security, we support the U.S. state, local, tribal, and territorial networks. And we do have such a service because we are supporting organizations, many of them of which don't have resources. So they need these types of capabilities. But if we can shift that left and see a real impact—and I'll give some other examples of how we can do things better at scale as we make all of these security changes—we can look at this as an opportunity.

But not all problems have been solved, so there's lots of opportunity for each of you to think about the problems that you can help to solve as we move forward because this is inevitable with executive orders requiring zero trust built-in security shifting left, and the EU NIS directives with a similar theme.

The organizational level. There's a much bigger impact to increased use of encryption. And this is because there's been heavy reliance upon the ability to see into traffic. Network operators ... And I've seen presentations on this. They would find it very disruptive because they've used their packet sniffers to detect login problems with applications.

Now if you really break that down, why didn't the application owner discover that problem? Why did it go to the network operator? And so the zero trust tenant of logging, and increasing logging so that we have

some uniform abilities to perform that early detection for troubleshooting, could make a large difference.

Intrusion detection and prevention systems. These in the past were largely based upon signatures—signatures of what we know to be bad. And that doesn't scale because what we know to be bad grows exponentially. And if you're looking at something like a hash of a file, it's really easy to make a very small change in a file and have that malware look completely different to a hash on a system that's comparing.

Behavioral detection is helping quite a bit. But what's making the most impact and what vendors are really moving towards, at least with the endpoint tools and even some on-network, is allow-lists—what traffic do I expect to see—and then looking at that to determine from the unexpected—is this a problem or not a problem?

So some of the changes we're seeing is research to do detection on encrypted traffic. We're actually working with a supercomputing center for our data that we host for SLTTs to see what's the efficacy of that is? Is that a viable solution? And at the same time, we also help with endpoint type solutions which have come a long way in the last few years. So they're really shifting to this allow-list type model.

But at the same time, vendors are working to build in security to the point where we might see some of those tools and the capabilities growing today diminish, and some of those tools diminish in need as well because of the shift to the endpoint and built-in security.

I don't think I have a good place where I call that out, so I'm going to talk a little bit about it here in terms of what we're seeing. There are some requirements for things like the software bill of materials. And for the federal government SBOM software bill of materials), they're going to require this, and any vendor providing software to the government will have to have an SBOM on all of their software.

So what does this mean? It means that there's going to be a manifest. A manifest that describes everything that's in that software—what libraries are embedded in it. So if there's a log 4J again, can you go through and figure out from all of your SBOMs palms each place where log 4J exists. And that way you can remediate the potential vulnerabilities.

Going beyond that, it also provides software digital signatures. So you're getting a digital signature on top of your software through the use of an SBOM to provide you some level of assurance that the package is as expected from the vendor. And that doesn't preclude something like SolarWinds where there was an engineer embedded within the development team and malware in a backdoor went out through that expected software.

However, NIS put something out last week where SSDF (Software Security Development Framework) will be required, in addition to these SBOMs. So organizations will have to go through this process to better ensure the security of their software before it gets signed.

And so we're going to see some really big shifts in terms of knowing what we expect on an endpoint and then being able to see from that—

since we have all of this knowledge—what are all of the granular configurations, what are the expected software on my system to be able to predict from that and within expected ranges what might not be okay. And it might be that behavioral analytics look at just that gap of what looks unusual from what's expected. Because attackers will advance in that way.

So with that, use of things like indicators of compromise—I'm not sure if this audience would be as familiar with indicators of compromise, but it includes things like known bad IP addresses, malware, hashes, other information that indicates that an attacker's on your network—will become have less value. And that will be because we'll be shifting to allow-lists.

So we have some pretty important changes. And I for one am very optimistic after researching this and working on this for several years. If you could please advance to the next slide.

So with this, we're undergoing a very big transformation and we're going to see a lot happen within the next year. And so embracing those trends, it will be disruptive. It absolutely will be disruptive. We're changing how we've done things, but we should wind up in a better place—really, probably where we should have started in the beginning with built-in security expected from the vendor and, also, shifting left the management requirements of the security for our products.

And the important part of that is that we've been running with a 3 million or more person deficit for information security professionals. I

don't think that's a gap ... I personally don't think that's a gap we will ever fill. And I believe the numbers ...

We have somewhere around 4 million security professionals. It's just too large of a gap. I mean, we can train all the people we'd like, but a better approach really would be to do things at scale. Right? Can we shift left? Can we move things back and adopt architectural patterns that scale?

So this is something I've thought a lot about. I put out a book, I guess in 2020, on Transforming Information Security that helps you think about these patterns. And that's why I wanted to pull up the example from this morning, which was really great, where there was just this pattern of shifting back the detection of malware and domains in malicious domains to the registries and the registrar.

The power of that is pretty big because then you don't have to maintain all these different block lists and all of these different services in time if it were to be closer to the 100% or that distributed workload becomes less.

So it was really interesting example, but there are many more—the use of SBOMs, the use of knowing what is on your system, like with firmware when you boot up a system. The vendor is taking care of that. They're making sure your system is meeting those expected values. So this is a really big shift that will have a large impact. Next slide, please.

So here I'm going to go through some of the changes and pull Paul back in. I think I talked enough about network monitoring. You have IPv6 and

you have overlay protocols within a data center. That might be Geneva. That might be service function headers. But you can encapsulate packets. You can direct them as you want to. And so those capabilities exists. We have measurement protocols.

I think we're pretty well in hand for that. Just, we can't intercept traffic anymore and we can't look at the data anymore. And that data is going to be more protected going forward. And this will move into internal networks. The U.S. federal government, in a recent White House memo called for a lot of these actions to happen by the end of 2024, their fiscal year.

So we should see some pretty big advances in terms of capabilities from vendors to support that, which means it will disperse into other sectors and other nations because these capabilities will be available more so in products and our ability to shift to the use of allow-lists for detection instead of all of these add-on packages that were deny-list based.

Internet core services. So there's a lot to figure out here, still. XMPP. We've already had MLS in the IETF as a replacement protocol which provides end-to-end encryption for instant messaging. XMPP is still in use. It has many functions, but it operates as a relay protocol similar to SMTP (Simple Mail Transfer Protocol) in DNS because there are interception points.

There are points where TLS is not in use because there is an operation happening on a box for DNS. Your resolver is working and the TLS is terminated before those operations occur. Right? So you have points of exposure. MLS will see some reduction in that. But at the same time,

instant messaging is still siloed and vendors plan to keep it that way. So it's not one I would really focus on too much.

SMTP, I'll talk about a little bit, and then a little on DNS before opening it back up to Paul.

So for SMTP, we have things like MX record-based screening. And that's been an important tool to screen for malware, to look at URLs which rely also on DNS lookups and DNS blacklisting, services for detection of malware and preventing that from getting to an end host.

And what's been replacing that for many organizations is an API-based service. That API-based service lets you go into the mail server from a remote point and do those types of inspection services, whereas the MX record is a screening on the mail before it gets delivered.

So there are some core differences in that. And I can get in deep or have a separate call if anyone has questions on that.

But back to DNS. We are seeing some pretty big changes. DNS over TLS, really not a concern. Nobody should be looking at your traffic on the wire. So this is preserving what we've had, but adding encryption on the wire. Anyone that's doing work on a DNS server in terms of, let's say, providing a blacklisting capability or any type of measurement, they're still able to do that because the data is exposed on the resolver.

Some of the concerns you get into is the differences between DNS over to TLS and DoH. And one of those major differences is how the client configures itself. For an organization, for DNS over to TLS or just playing

DNS, it's at the host level and it can be done for the entire network with an update with DHCP and pushed out to every single host.

The big difference with DoH is it's on a per-application level and it can be hard coded. So if your organization has a Service Level Agreement with a service provider, it's not receiving the protection on your DNS and meeting the expectations because it's avoiding that particular server.

And so that presents a big challenge. If these services, the DoH servers, aren't aligning to your organization's policies or maybe perhaps don't have a privacy policy published like the ones of the early days with Mozilla and Google and Cloudflare ... They had clear policies published and they were really pushed to do that.

But if you get an application and it's hard coded to a DoH server and you have no say in the policy, that becomes quite complicated and your organization has no way to detect the use of that. There are certain complications there.

And so I don't have any good answers there, but I'm sure this is an area Paul can dive much deeper into.

PAUL VIXIE: Thanks. Yes, I have been thinking about the impact of DoH for a couple of years now, and I think ... Well, Kathleen mentioned an important difference which is that DoH is an application layer protocol, essentially. There is another one. And to understand this, I just want to remind everybody DoT came first. The idea that ...

We've had DNSSEC more or less for some years now, but all DNSSEC does is inform you that the data you're holding is authentic. We have other protocols like TSIG in DNS to just verify that you are talking to someone who you share a key with. So those were well-solve problems.

But the difficulty of being data mined, either for commercial or state actor purposes or criminal purposes, was big. And so DoT came along with the idea that, "We're going to encrypt this, and we're going to use whatever TLS happens to exist." It was 1.2 at the time. It's 1.3 now. But it was still on its own port number.

In other words if you looked at the IP and the TCP or UDP headers, you could tell this is DoT traffic. Without being able to dip into the payload at all, which you really can't with TLS, you could still say, "You know, that's against my policy as a network operator and maybe as just a host owner are dealing with applications that I've imported for various reasons."

If one of those things besides that it either doesn't know my policy or doesn't like my policy and they're using DoT, well, you can enforce it. Say, "This is my policy. It shall not pass." And I know that's a deny-list and I know that's out of favor, but you just have to remember there's a very big market of applications being pushed out into the edge and there's a long tail on changes.

If you want to make a difference to the security of the network, you've got to have a plan for what you do with the things that won't be updated for 5-10 years.

But anyway, DoH looked at that and said, “Well, we see what you did with DoT but the fact that it could be blocked is a problem. And so what we want to do is shroud it in an additional set, an additional layer, which is HTTPS, so that the traffic will be in indiscernible from desirable traffic at the layer of a network firewall. In other words, there'll be nothing in the IP or TCP or UDP headers that identifies this as a potentially policy-violating DNS transaction.

Now the reasons for that go back to Edward Snowden’s disclosures, but what this does is to kind of throw out the baby with the bathwater because, as I said at the outset, there are legitimate reasons why and on-path actor might want to be able to tell what's going on and differentiate traffic and make policy decisions. And that may be because you're concerned about poison supply chains, malicious intruders, malicious people on the outside trying to groom whatever—your employees, your children.

There are a lot of things that the app cannot be expected to know or do, and never will be. Same for the host operating system. So DoH, by trying to bypass one thing, ends up by bypassing quite a bit of currently very necessary, very desirable activities. And try to imagine that you're in a malware lab and you're torturing some piece of malware that was found out on the network. And so you're running it in a simulator and you're trying to figure out what it's doing.

In the DoH world, it's going to be very difficult to be able to tell what that thing is doing because we're used to looking at the DNS signal patterns that come out of malware. And now those will be completely

shrouded. And of some concern to me is that when I've spoken about this, people have said, "Paul, the firewall that you have is 1990s technology and you need to get up to speed. You need to understand how the world is working now."

So this is a gulf. This is a cultural gap where those of us who have been trying to secure networks are watching the last couple of behavioral signaling patterns that we had sort of disappear. And it's our friends that are doing it because they just insist that the world has to move forward. And that forward movement will not be without pain.

So as you look at QUIC, the UDP-based replacement for TCP, at least as concerns the web. And there will be DNS over QUIC as well. It has, as Kathleen said, deliberately no ability to manage it.

And there's going to be a lot of people who have been defended by behavioral edge security technologies—people, applications, data, etc.—that are about to go dark. They're about to become blind. Their defenders are about to become blind. And all of that endpoint stuff—the users, the applications, the operating systems—are not necessarily ready to defend themselves. And so that occasions this talk.

I'll pass it back to Kathleen.

KATHLEEN MORIARTY: Thank you. That was a great ending summary. We do have a gap between now and the next few years when all of those allow-listing capabilities—you know, if they all come together—and the capabilities

that come with those. So, yeah, it's an uncertain time, but I also hold some promise for it. But we have challenges.

And so hopefully this community helps with some of those challenges and thinks about these problems sets, especially those Paul just enumerated so that we can ultimately get to a better place. Thank you.

EBERHARD LISSE: Thank you very much. It was a bit deep for me. But then you know, in my day job I'm going to gynecologist. So I touch on the surface of this technological stuff sometimes. But it's a good thing to get deeper into the issues sometimes. Thank you very much.

There are no questions that I can see in the Q&A pod. There are no hands of the panelists raised. There was a hand raised in between on the attendees, but we only take questions through the pod. There is none. So thank you very much, both for you, for this presentation. And I hope we'll meet in person again.

PAUL VIXIE: Yes. Me, too. Thank you very much.

EBERHARD LISSE: We'll have another Mexican dinner somewhere because that's what we did last time.

Anyway, now. There is one more question. Ayesha from Pakistan, "How can I join Center for Internet Security?"

KATHLEEN MORIARTY: So why don't we take that offline? I'm happy to help answer that question. We have memberships and many free offerings like our security best practice documentation if you're not reselling it. If you're reselling, it's different. So yeah, I'll take that offline. Thank you.

EBERHARD LISSE: Yeah. She can just click on your link in the agenda and get in touch with you.

KATHLEEN MORIARTY: Excellent, thank you.

EBERHARD LISSE: Thank you very much. Okay. And now we are going to have our DNS roundtable organized by Dan Owen and Graeme Bunton. I don't want to hear who was doing the more heavy lifting or not—I appreciate the work of all of you—but for the ones that I communicated with to get this going, and then the ones who are participating in the table.

I think Dan Owen can say a few words as an introduction, and then the roundtable can basically proceed among themselves.

DAN OWEN: Thank you. I think where we've got most of our panelists joining. I think we've still got one more that hopefully going to join soon.

Just an introduction to this. This roundtable is going to be discussing the respective efforts of five collaborating organizations to reduce the domain abuse problem and the challenges of addressing it to include a recent case study using new technology.

So with that, I'll hand it over to Leslie Daigle, who is the CTO for Global Cyber Alliance. She'll be the moderator and the first speaker. Thank you.

LESLIE DAIGLE: Great. Thank you very much, Dan. And thanks, everyone, for joining today.

By way of format for this, what I would like to do is, we'll have some opening remarks from the panelists talking about the subject and the technologies and efforts that they bring to bear on it. And then we'll have a bit of discussion, possibly between the panelists. And certainly, we'll be happy to open it up to questions from the floor after that.

On the whole, the more that we can get this, after opening remarks, to a discussion level, the better. So don't be shy when it comes time for asking for questions. Please do jump in.

Okay, so first up I would like to talk a little bit about the Global Cyber Alliance's own work in the space, and that is the Domain Trust project. Domain Trust is a project that the Global Cyber Alliance developed, aimed at helping reduce the amount and impact of domains registered for criminal purposes.

We're not a registry. We're not a registrar. So we can't make any of the changes ourselves. Our efforts are focused on bringing together those data about abusive domains and sharing that information from diverse sources in order to get a sense of the scope of the problem, bring more parties to the table to collaborate with each other, identify what else is necessary in order to provide those who can act—registries, registrars, protected DNS operators— with what they need in order to take action.

The kinds of entities that we're bringing to the table include registries, registrars, ISPs, public and private cyber responders, financial institutions, and CERTs. We believe that by sharing information and working together to reduce and, where possible, stop abuse, we can help build trust in the domain name system as a whole.

To that end, our initial offering is an information-sharing platform that provides actionable data against several forms of domain abuse used by global cybercrime, such as phishing, malware distribution, and command and control activities.

So let's back up and talk about where that all fits in the Global Cyber Alliance. We are a not-for-profit organization aimed at enabling a secure and trustworthy Internet. We build programs, partnerships, and tools that make connecting worlds safer and more secure for all. And we've covered a lot of ground in our six and a half years of existence, including a lot of effort to promote the uptake of DMARC and a little project that led to Quad9.

I'm the CTO and Director of the Internet Integrity Program at the Global Cyber Alliance. And in that program, our focus is on helping identify and

solve cyber security challenges within and because of the Internet's infrastructure, especially in terms of problems that can't be solved by any one infrastructure operator on their own. Domain abuse is such a problem.

So where we're going from here with the project is fostering as much dialogue as possible to get a shared understanding of the importance and scope of the problems, as well as partnering with organizations that have and can contribute lists of identified domains used for criminal activities and/or that will take the collected domains and take action.

Domain Trust is a project that is now over a year old, and we are transitioning to a higher technical value platform by recruiting more truly global partners. We're focusing on CERTs, as I mentioned earlier. Our hope is that better data on input will equate with faster action on the results. And I'll say more about that in a minute.

Data will be more current, more valid, more diverse geographically. That is our aim. So apart from the actual platform, we are building a Domain Trust community that leverages the platform and taking action by convening groups of Internet infrastructure owners and other domain abuse organizations— similar today's roundtable—leading collaboration within the community, encouraging communication for improved coordination and data sharing, facilitating community action to reduce domain abuse at all stages.

And we don't want to just talk. It's not what GCA is about. And I'd like to help contribute to the data-driven decision-making discussion. So I'll leave you with this observation from our data.

Our time-to-takedown indicator measures the length of time between when a domain is submitted to Domain Trust and when that same domain is taken down. Yeah, we know that's correlation and not causation because we're sitting somewhat outside of the system. But still.

We took one recent week as input data and observed that the average takedown time was 180 days. Six months. Of course, this is a just distribution with long tail. Some domains were taken down just four days after being reported in the Domain Trust. But some took more than a year. And of course, there are always reasons.

But I think through platforms like this, through some of the projects that you're going to hear about next, through ongoing dialogue, it would be great if we could actually drop that average takedown time.

So that's an introduction to our platform. Next up, I would like to ask Graeme Bunton, the Executive Director of DNS Abuse Institute, to say a few words about his.

GRAEME BUNTON: Thank you, Leslie. Hi, everybody. I'm Graeme. I'm the Executive Director of the DNS Abuse Institute which is an organization created by .org. Apologies to the people who are here at the session or the brief introduction I gave earlier today.

.org recognized—I think much like GCA—that DNS abuse is a rising problem and we need to do something about it, and quite rightly that cybercrime and DNS abuse is a complicated global problem and trying to resolve it at the level of individual registries and registrars isn't going to work and we need to do a bit more work to coordinate and centralize and address these problems in a single place. And so we have something like the DNSAI.

And briefly off the top, there's no sense of like competition here between GCA and the DNS Abuse Institute. I think our current initiatives are nicely complimentary, and the way we're able to work together and communicate seems to be a very viable model. And there's lots of room because this problem space is so big. And so I'm very pleased to communicate and work with GCA where we can.

So the Institute operates under three key pillars. We've got collaboration, education, and innovation. The education is relatively straightforward. We're looking at producing and have produced a number of best practices for registries and registrars to mitigate abuse or understand abuse, as well as for people like end users to keep their WordPress site, say, for example. So really trying to make sure that we're providing resources across the entire ecosystem.

Collaboration is really going to be bringing people together to share best practices, intelligence, useful things to mitigate abuse. And then the innovation is the fun bit. And I'll talk a little bit more about what we're working on for the community to address DNS abuse issues

because I think this is where the really interesting stuff is happening, especially over the course of the next year.

So we've got two key initiatives coming up, hopefully between now and July. The first one ... For lack of a better name, we're just calling it the DNS Abuse Institute's Intelligence Project which is basically to try and build a robust statistical understanding of where these types of domain abuses are occurring.

We're doing this for a couple reasons. One is for the Institute's own credibility. If we want to be an expert on DNS abuse, we need to understand where it's happening, who it's happening to, why it's happening, things like that.

The other bit is that we need to be able to push conversations along, especially within the ICANN space where we talk too much in broad strokes or vague ideas about bad actors. And we need to get past that. We need to be able to speak with specificity.

And so in order to do that, we're looking at partnering with academic institutions to build a robust reliable DNS abuse intelligence system unencumbered by commercial sensitivities or interests and/or community sensitivities and interests that I think undermine some of the existing efforts.

We're also really looking at making sure that data is evidenced, so it's not just a name on a list; that if we're going to call something abuse, it should actually be something that a registry or a registrar could meaningfully act on. And that, unfortunately, is not always the case.

We're looking at distinguishing between malicious versus compromised registrations because that is a distinction, with an increasing importance. And as an aside, there is a plenary on that on Wednesday, I think.

And then we're really looking at being able to publish reports at the registrar and registry level to demonstrate not just where abuse is, but to be able to celebrate the people who are really good at addressing it or have really low abuse on their platforms, as well as to highlight the places where they over-index and where there might be problems.

And then to do all of that not in a vacuum, but be able to share that information directly with registries and registrars and say, “Here are the problems that we're seeing. Here's how we think we can help you to make that better.”

The last thing I’ll try and touch on briefly without talking too long—and apologies if I’m going a little bit deeper here—is that we're building a Centralized Abuse Reporting Tool, something that the ICANN community has been largely concerned with. SSR2, CCTRT, and SAC115 all sort of touch on a similar idea of having a single place that solves a couple of problems.

One is that reporting abuse to registries and registrars is very onerous and difficult and you need to be able to identify the appropriate registrar. You need to find their abuse reporting page. You need to navigate forms. None of which is standardized. The levels of information are different. The languages might be different. So there are lots of impediments to reporting abuse.

On the registrar side—mostly registrar; registry, a little less so—the abuse reports that are coming in the door are terrible. They're unevidenced. The domains often don't belong to them. They're duplicative. And there are thousands. So registrars are spending substantial amounts of hours triaging basically useless tickets for very little value in making the Internet safer.

And so we're building, essentially, an intermediary to try and clean that problem up, which is that it will accept abuse reports from anybody, either via form API or embedded forms, so that they can be placed elsewhere. We're going to standardize those abuse complaints. We're going to enrich them with data from API-based sources from around the Internet. And that's moving the investigatory burden from frontline compliance people into the tool, and it's going to incentivize registrar adoption, I hope.

And then we're going to distribute those appropriately via API or e-mail to where registries and registrars would like those abuse reports.

And my hope is that this tool gets wide adoption and really helps move the reactive abuse reporting processes that right now are wildly disparate across the domain registration industry and will really help tighten that up.

And so maybe I'll stop there because that's a lot. Thanks.

LESLIE DAIGLE: That is a lot. Thank you very much, Graeme. And if we just scoot along right now, next up is Drew Bagley who, in this context, is the Director of Operations of the Secure Domain Foundation. Drew.

DREW BAGLEY: Thank you very much. Yes, I help lead the Secure Domain Foundation which has operated for nearly a decade and has been focused on empowering infrastructure providers to really take charge of abuse through proactive anti-abuse rather than merely waiting until those abuse reports that Graeme was mentioning a moment ago come in.

And I wear a couple different hats. My day job is with CrowdStrike where I'm the Vice President of Privacy and Cyber Policy. And then I'm very involved in this community, too, through several advisory groups—with Europol, with CISA, and with the DNS Abuse Institute as well. And I have previous experience working for the FBI.

And so my perspective is really informed by seeing what happens on the tail end when you don't do enough about abuse on the front end. And you can see some of these catastrophic data breaches, attacks on critical infrastructure, and of course attacks that are catastrophic for individuals on a one-on-one basis merely through the form of traditional, cybercriminal activity.

And so with the Secure Domain Foundation, one of the things that we've really advocated for, for years, with best practices is for infrastructure providers to recognize the position they might be in depending on the type of infrastructure provider we're talking about to

really be best suited to tackle abuse from a technical standpoint, either from ... You know, when we look at the dichotomy that Graeme presented where you have domain names registered for malicious purposes versus those that are compromised.

So either through adding friction to the process where you have a repeat abuser where an account with a registrar is being used repeatedly to register domains that invariably get taken down to add some friction there rather than just allowing for that repeat abuse; or for really being able to be engaged and quickly act to help customers who are the victims of compromised domain names that are then being used to perpetuate cybercrime on others,

And so a lot of the best practices that we've highlighted over the years have actually, really been repeated in multiple forums including the EU DNS Abuse Study that just came out recently where you're really seeing an emphasis on providers using simple things like fuzzy hash to determine if a domain name is matching something that is likely going to be abused; or, again, using that account information that they have access to without the need to disclose it to anyone externally, but to see if that's associated with known abuse internally.

And so the other way in which we've really been involved in advocating for being proactive with anti-abuse is to think about ways in which best practices that have been perpetuated by the CCT Review Team, which I was on the leadership of, are now being encouraged in the ccTLD environment as well, such as with this recent EU DNS Abuse Study.

And really, the notion is that when you have entities like GCA and now the DNS Abuse Institute putting out best practices, we think it's really important for the community to really adopt those best practices and look for those best practices.

One of the really interesting findings from a survey that the Secure Domain Foundation did several years ago was that many infrastructure providers that were well suited to help with abuse reported that, okay, they could engage in some of the proactive anti-abuse, paying more attention to fuzzy hashes that matched brand names or to DGAs being used. But that when they were dealing with reaction to abuse reports, the abuse reports were all over the place, like Graeme was describing.

And so one of the asks all those years ago was for a standardized system of reporting, a standardized protocol for that DNS abuse reporting. And that's why, with the Secure Domain Foundation, we're particularly hopeful with what the DNS Abuse Institute is coming out with, with the Centralized Abuse Reporting Tool, whereby members of the community can really join in in getting a standardized list of fields being reported on for abusive domains that will hopefully be very actionable.

Because I think, from our perspective, what we see is that those best suited to take action will take action if they have good information to take action. But that absent that, it really becomes difficult for everybody and nobody's happy. The victims aren't happy. Those in the cybersecurity community that want infrastructure providers to be able to do more aren't happy. And the infrastructure providers aren't happy because the information is not good.

So my hope is that now, after years of best practices being developed and now some actionable things like protocol shaped around standardized reporting, I think we can really see a lot of these things come together.

And similarly, what the findings have shown, whether you're looking at the DNS root abuse report that was commissioned as part of the CCT Review Team or you're looking at the recent EU one, is that DNS abuse is not coincidental. It tends to happen by repeat offenders. It tends to happen oftentimes in the same zones.

And so I think that there's a lot of low-hanging fruit that the community can really go after by utilizing these best practices, utilizing these tools, and really utilizing the entities that are speaking here today on the roundtable.

And so with that said, the other thing I would just say is that, unlike in years past, I think there's a lot more free tools available, too, for infrastructure operators to even check to see whether or not a reported domain name is associated with any known abuse rather than just looking even at that report that they're getting inbound itself.

So for example, CrowdStrike has a tool, Hybrid Analysis. And it's a free, community-based website where you can check to see if a particular domain name is associated with any maliciousness by running it in a sandbox. And there are countless other free tools out there like that, and so I would just encourage those involved in infrastructure to be on the lookout for those tools as well.

So with that said, I will pass the baton back to Leslie.

LESLIE DAIGLE: Great. Thank you very much for that, Drew. And now we'll turn to Danielle Deibler, who is the director of Threat intel and my favorite cat herder at Quad9.

DANIELLE DEIBLER: Hi. Thank you so much for having me. I appreciate it. Good morning and good evening to people across the world.

So a little bit about Quad9. We're a non-profit that provides an Anycast open recursive resolver. So we might be a little tiny bit different than some of the other presenters. We integrate a threat intelligence feed that's an aggregate of around 20 or 25 providers at this point. And we essentially filter domains based on fully-qualified domain names as the filtering targets.

So our goal is to provide a free, base level of defense [against] cybersecurity threats to the global Internet community. We don't moderate taste or access to information, so we don't do content filtering. It's more like helping your grandma avoid a phishing scheme that's targeting her banking credentials or keeping municipalities from downloading malware and ransomware. That's more of our edict.

We really don't, unless we absolutely have to, block content. Mostly it's malware, command and control botnets, phishing, and stalkerware domains that we really focus on. Our recent focus has pulled in a lot of

lists that are more focused on advanced persistent threats, so APT actors which are typically state-sponsored actors.

We're a little bit different. We're not trying to do ... I guess, like from some of the other panelists, we're not trying to do a massive initiative. We're actually doing the blocking right now for people who are using our service.

We also, in addition to a blocklist, we also have what we call an explicit permit list. So at the DNS level, it's not typically appropriate to block a domain like maybe drive.google.com. There might be some malware on it, and that is a definite issue. But we don't we can't get to the URL level. It's at the domain level.

So we participate at a level where the domain, we feel, is malicious; or, more precisely, a threat intelligence provider that we partner with believes that the domain is actually malicious. But if it's a few URLs on a particular domain, that's not our area of expertise and it's certainly not an area that we can impact in terms of filtering. So we do also have an explicit permit list in addition to the blocklists that we carry from our different threat providers.

And then we had an operational [incident] that is ... The one that we're going to discuss is, I think, probably why Luigi and I are on this panel together. So I'm going to give some data about the incident that we had, and then he's going to expand on that a little bit and talk about Bfore.AI's approach.

So we integrated the day Bfore.AI feed probably around the mid-December time frame. And within a couple of days, around December 21st, we saw this one domain really skyrocket in terms of the number of hits were getting. And the hits in this case were the number of blocks we were getting across our global network.

And went in ... And typically this, for us, would actually trigger a false positive. We'd look at it and say, "Oh wow. Yeah, this must be like EventBrite or some other kind of thing that we're blocking that's a really highly popular domain on the Internet."

And what we found is that it really didn't have the characteristics of that type of a domain for us. It was registered within the last 30 days. It pointed to a free, authoritative DNS provider. It had no Apex A record or www record for an A record. We were getting a lot of timeouts on the back end server. So were getting 522s. When the page loaded, it was a blank page.

It was an anonymous registration that didn't have any DNSSEC associated with it, just some basic characteristics of the domain. It just didn't feel like it was like some highly-popular domain that had been launched out there. And so we looked at it a little bit more closely and realized that we didn't know enough to essentially reported it as a false positive. So we worked with Bfore.AI.

Our query rates for this domain went to up to 94,000 blocks per second. So for a couple of hours and then over the next 24 hours, we saw about 365 million hits for this blocked domain. So clearly associated—we felt like, based on some of the other characteristics—with a malware or a

phishing campaign, something that was probably a malicious domain on the Internet.

And so we ended up keeping it blocked. We work with Bfore.AI. to get to the root cause. And Luigi can talk about that a little bit more.

And then one last point I want to make before I hand it off to Luigi is that we do frequently look at the reputation of a domain. And that means you look at how long it's been registered, how many incidents a domain might have had over years or decades. And the one thing that I would say that's pretty recent is that Russia has made this recent announcement to kind of nationalize or ... I'm going to say nationalized. They're not cutting themselves off from the internet, but they're trying to kind of nationalized their infrastructure around March 11th.

So I'd recommend that if you are a TI provider or registry or a registrar, you might want to keep an eye on previously-benign domain names because they could switch that benign behavior to a malignant state very quickly. And their domain reputation is actually pretty high. Right? Some of these domains could have been around for 10 or 15 years, relatively benign—not distributing malware, not doing phishing, not running command-and-control botnets, not downloading little things to you via JavaScript—but that could switch very, very quickly.

And that is it. I'm going to hand off to the Luigi. Thank you very much.

LESLIE DAIGLE: Great. Thank you very much. And next up is Luigi Lenguito, cofounder and CEO of Bfore.AI. Luigi.

LUIGI LENGUITO: Thanks, Leslie. Thanks, Danielle. Thanks for the invite today.

So how did we know that that domain was going to be malicious? Because we told Quad9 ... We actually had it in our threat feed about a month earlier.

What we do is a bit like a weather forecast. We observe the whole Internet on a daily basis and we use some smart predictive analytics and some other behavioral analysis to come up with these predictions.

And we call them predictions because 99% of the time when we are sharing them, there is no content, there is no traffic to those domains. And we are able to already know that they will be malicious, indeed, in the future.

Today about 99% of our indicators are unique. 80% of them cover an attack that others identify later. And in 61% of the attacks, according to one of our customers, we are the only one that can provide [an answer] even months later.

So technically speaking, this a massive AI model that is looking at network features of those domains—as they get registered, as they get changed, as they get modified—and predict, based on their behavior, how they will be used. Now we're unable to say the specific type of maliciousness that ...

Is it going to be phishing? Is it going to be malware? Is going to be a botnet? We don't know that. That we cannot predict. But we can predict

with a very high reliability at the moment. Our false positive rate is 0.05% that they will be malicious or benign. We also do the allow-lists and provide it to other security features like Quad9 as a DNS resolve or anti-phishing filters or firewalls and the like, this information.

The objective for us is that we have realized that nowadays the cost of an intrusion and the speed of detection makes it unbearable for commercial users to, you know, the situation. So allowing people in is not acceptable. Of course, teams have to have strong detection and response capabilities, but if we can avoid that intrusion at all, I think everyone would be happier.

And so I'm very proud to work with Quad9. It's a great use case for our technology. But I'm also very interested in this conversation, as we have a second service that uses the prediction. That is brand protection where we have our customers block abuse on their brands that may cause impersonation [inaudible] partners of large brands may be scammed of hundreds of thousands of euros through business e-mail compromise or for banks like Volksbank, one of our customers. They want to prevent their users to be phished for credentials and then lose tens of thousands in wire transfers that shouldn't be allowed.

And so it's very interesting to hear all of the effort in the industry to reduce domain abuse. Today one of the challenges—and I definitely join what Graeme was saying—is the diversity, let's put it this way, of reporting. So we work with a lot of registrars to raise these domain abuses and try to get it down.

And certain are more responsive than others. Let's put it this way. But definitely, the degree of information that they require and the speed with which they act on it is too varied. So more standardization would help.

Now to that extent, I have to comment on one of our partners. This is Namecheap, that in the last year has definitely set the standard. Again, yesterday night we completed the DNS takedown—sorry—a malicious domain takedown in like six minutes from reports, from our prediction to their takedown that I think is quite an astounding value. And I think, obviously, may not be always so good, but think we should try and aim for that.

And with that, let's get back to the panel. I think there is a good conversation ahead.

LESLIE DAIGLE: Great. Thank you very much for that, Luigi. And indeed, let's get to some discussion. And I appreciate that there have been a couple of questions that have come in, in the Q&A, and I'd like to get to them.

But before we quite get there, Danielle brought up a point that March 11th there's an expectation that the world may shift significantly in the Internet. And [now] the only context in which we think we're headed one direction. All of a sudden, we're headed in another. We're still living in the wonders of the pandemic.

So I guess one question is, what does that kind of tectonic shift do in the world of predictive analysis?

LUIGI LENGUITO: So I'm going to take it maybe from my end first. So first of all, I'm going to give you a couple of numbers. So every day we observe about 6 million changes in subdomains and domains that gets called by our systems. That releases about 90,000 malicious FQDNs in our feed. I can tell you that about 60% of them activate more than six months later. That means that for six months, they sleep.

So for whoever has rules in the firewall that blocks newly-registered domains or less than 30-day domains, I'm sorry, it's not very valuable. The malicious domains are much, much more like wine. They tend to age. And everything that is very easy to detect, maybe script kiddies can be blocked or young people trying—security researchers.

But unfortunately, the criminal gangs are very well organized. Nowadays there are very good DGA engines and other scripts, and our system sometimes actually tends to, say, reverse engineer some of these algorithms and give us the behavior. And so it's quite interesting to see how complicated the matter is.

But to the point, I don't think the March 11th thing is going to be a big change. If they disconnect Russia, if anything Russia will be disconnected from the world as well. So there will be maybe a slightly different perspective of who is going to be the criminal profile. But most of the attacks don't come from Russia ASNs or .ru domains. So we should be a little bit more sophisticated.

Most of the attacks are coming from cloud operators where the virtual private servers being loaned to criminals. And those may be in China. Those may be in Thailand. They may be everywhere. So I'm not personally buying that this would be a big change for us. Unfortunately, most of the criminality is commercially driven.

LESLIE DAIGLE: Great. Thank you. Anybody else want to take a bite of that Apple? Graeme, I don't know if your hand is up because you want to say something about that or if you have a new questions to introduce.

GRAEME BUNTON: It's probably tangential to that.

LESLIE DAIGLE: Go ahead.

GRAEME BUNTON: Boy, the observation that domains are like wine would be very interesting. I'd love to see some more data on that, and we should take that offline because most of my understanding is that a lot of malware and phishing is very quick. And so that you're seeing lots of aging would be a substantially new data point. So, cool. Let's hear more about that.

I think I want to make, though, and observation and maybe hear from other panelists about something that I've heard a bunch because we're talking about lists and predictions. And it feels to me like there are three

things we're trying to do here. And the level of evidence or information required for each one is escalating. And I don't think we differentiate enough between these things.

So I think network protection and network blocking, be it at the corporate level or at the resolver level, where you can do that with a broad list with minimal harm done by over blocking in that circumstance. And then we're talking about a next step which maybe like friction in the registration process where you're trying to identify potentially malicious domain names.

And I was talking about this, this morning, where someone's trying to register paypal-login.qrs or something and you're like, okay, preventing this from registering is likely to prevent a harm and isn't going to have particularly dire consequences because it doesn't exist yet. And it's not like we're bringing a domain name that exists offline and potentially harming someone.

And then the last one is abuse mitigation at the registry or registrar level where something already exists. It may or may not be compromised. It may or may not be malicious. The domain exists and so we need not just a name on a list, but some corroborating evidence for it.

And I think it's important that we don't conflate those three tiers of activities and what's required for each of them because I think we put ourselves in a pretty difficult spot if we don't recognize those subtleties. And I don't know if that resonates with everybody else, but maybe I'll leave it as food for thought. Thanks.

LESLIE DAIGLE: Yeah, I think that resonates pretty strongly for me, at any rate, because ... Well, one of the things I appreciate significantly about Quad9's efforts in the productive DNS base is the emphasis on trying not to have false positives.

Because when you're doing things like blocking the actions that come from protecting your networks—I think was the characterization you used—too much blocking and can leave you in a space where you effectively Swiss cheese the Internet. You know, there are big holes in it that weren't meant to be there.

So that's where, earlier, Kathleen Moriarty talked about moving things to the left and trying to understand better how to identify and stop domains from being registered for malicious purposes—I think is a key element of that.

And maybe this is a good point to tackle one of the questions that's come up in the Q&A, which is when we're talking about abusive domain, without getting too philosophical, what do we mean by "abusive domains"? In my remarks I tried to stay focused on the "registered for the purposes of criminal activity."

I'll leave it there. I don't know if other panelists would like to share their perspectives.

DREW BAGLEY: Sure. Yeah, I'm happy to jump in there. So there are many definitions of DNS abuse and many that are accepted by various members of the community. So the EU DNS Abuse Study actually did a great job, I think, providing a list of all of the community definitions of DNS abuse.

So for gTLD providers, some of those definitions of DNS abuse come from the contractual language. In the CCT Review Team's report, when we issue that report one of the things we highlighted was the definition of DNS abuse that ICANN had previously defined and identified before we kicked off our work as a review team. It was part of what informed the scope of our work.

There is a definition within some members of the community from the DNS abuse framework that I believe Graeme worked on in his prior life. So there are several out there. And essentially you can see a Venn diagram where the core types of DNS abuse that really can be seen as cybersecurity threats seem to have nearly unanimous commonality being identified as abuse.

And then there are many other areas where there is not uniform agreement at all and where some groups see certain activity as abuse and others say, "No, that's outside of the scope." But the bottom line is that for DNS infrastructure providers, what is going to definitely be abuse is whatever is in the policies of those providers as being defined as abuse and not in line with those terms.

And so to the extent those include some of the threats we've been talking about in this session that are cyberthreats, that's where, similarly, there's really a lot of commonality with the approaches you

can take such as the PayPal example Graeme just took. When you're talking about that and when you're using just basic fuzzy hashing and you're realizing, “Okay, this brand-new account is trying to register a new domain name that's associated with a brand, and it does not appear that this account is that brand,” is that likely going to be used for something good or not?

There are things like that that are probably a lot easier to decide to add friction to than some others which are going to be much more ambiguous and then what's going to be compromised.

But that's where I think, too, providers now are in a unique place where, to help protect customers that are against being compromised, that's where they can promote best practices such as registry locks, or registrar locks—to prevent domain names from being transferred to other registrars without the owner truly wanting that to happen—two-factor authentication being pushed if a provider is providing both hosting as well as the domain name registration.

And that's where other things, such as promoting just good, basic cyber hygiene like endpoint protection and things like that can be really helpful and useful as well. So I think a lot of this stuff depends on the provider and what services are being offered and whether they're taking a holistic approach to what's being offered. And if they are, it's really good then to tap into some of the best practices already out there for cybersecurity and be promoting those as well.

I’m sure others have opinions on this as well.

LESLIE DAIGLE: Others?

LUIGI LENGUITO: Maybe I'll just close the loop as well on what Graeme was saying. I think he was very right. There is no one-size-fits-all, one-solution for-everything. Right? So there is network locking. There is brand protection. There is takedowns. There are different challenges behind each one of those.

But I think the effort to move in a more proactive and one-step-ahead type move is critical. I was on this panel earlier on today. I think the idea that there is kind of a hold in a while, the quality of the registrar—or registrant, sorry—is validated. It's very important. It covers a certain degree of abuse, not all the other abusers. The same thing as what we do with Quad9 will protect certain people but not everybody. Right?

So it's a bit like what ... And it will resonate to this. In cybersecurity we say Defense in Depth. This is a bit like an onion. You need to have many different layers of protection for different types of challenges. But to me, the shift left as to tie with a timeline aspect. Right?

Let's not wait for a criminal to operate and we collect proof because it's too late. It's just never going to work from a time perspective. We need to be more in advance and find a new paradigm that is no more detection response, much more prediction and prevention from our end, or at least prevention and preemption [if we] cannot do the predictions.

LESLIE DAIGLE: So riffing off some of the comments in the chat just now, maybe we can turn that question around a little bit, too, and say, you know, I expressed concerns with ... If you block too much, you end up breaking the Internet. But if you insert too much friction in the context of registering new domain names, you also risk making the Internet difficult to use particularly for smaller entities that don't have the squad of lawyers in their back office.

So I suspect that most of the panelists here believe that they understand that problem and this is not about breaking the Internet in that regard. But maybe people would like to share some thoughts on how getting a little bit closer to introducing friction at the registration level doesn't necessarily break the Internet for those entities as well.

DANIELLE DEIBLER: I have a comment on this one, which is, so, I figure if somebody can figure out how to get through Facebook and their ability to place political ads or do certain things, you can probably have a tiny bit of friction in the registration process that verifies a little bit more about who you. And that doesn't seem insurmountable to me, even looking at some ...

Recently login.gov for filing your taxes online added this new thing where you have to upload your passport. I'm not saying I want people to upload their passport to do a domain registration. Let me put that right out there. But they introduced some friction in the process for

being able to do something. And I think it is about education. Not just of the registrars, but of the public.

And maybe you make it a little tiny bit harder, but it could make a big difference if you're doing a little bit more validation. Especially with so little ... Especially with the privacy concerns around WHOIS information, you actually ...

Like, if you can make it a little bit more friction up front, I don't think you're going to impact the bottom line. A person wants a domain name and they want to launch something on that domain name. And if they're a legitimate business, that's the brand that they want. That's the website that they want to launch.

And it doesn't feel to me like it would be a huge problem for there to be a little tiny bit more upfront validation for that particular user or domain. And certainly it would make our mission a little bit easier because we would know that there was some upfront validation that had happened in a particular ...

I mean, we do take that into account in terms of certain registries. We do require more upfront validation. But a generic domain name that got registered? It would certainly make our lives a little bit easier. We'd know they went through something up front.

And it looks like Graeme has his hand up.

LESLIE DAIGLE: Yeah. Thank you for your comment. And yes, Graeme, what would you like to say?

GRAEME BUNTON: Thank you. So, with the caveat that I did a presentation this morning on, essentially, ways that we can look at introducing some friction into the registration process—more friction. Because I do think there is room for that. And there are existing tools that we can leverage for doing that. And that should be part of a responsible registrar's processes.

But—and it's sort of a big but—this is one of the things that the DNS Abuse Institute is really concerned about, understanding the ecosystem that registries and registrars operate in. And again, I acknowledge that this going to apply more to gTLDs and cc's.

But I was chair of the Registrar Stakeholder Group for four years. I worked for a registrar for a decade. And so whether you buy a .org from a registrar in the U.S. or a registrar in Japan or wherever, it's the same thing. And it is an extremely competitive marketplace. It is extremely price sensitive. And asking registrars to unwind 20 years of optimizing that process of domain registration is a big ask.

And it's easy for us to say, "Hey, just go do that. It's going to make the Internet better." But we need to be able to do that in a way that recognizes the global competitive marketplace that they're operating in and ensure that there is fairness in how we do that. Because I think if

we don't, we're going to have some problems. And the economics of the industry are just going to undermine the work.

So we need to find the places where our interests in mitigating and reducing DNS abuse and making the Internet safer align with an industry that suffers under pretty serious global competition. I'll stop there.

LESLIE DAIGLE: Yeah. I think that's a really fair point. And I would add that I think that introducing friction in cases where there is some reason to be a little suspicious seems more right-size than just introducing friction across the board.

I mean, I can explain to you why I wanted to register internetimpossible.org and did. And it's not because I'm a large business doing something. But neither is internetimpossible.org going to go and phish and try to claim to be somebody else either.

DREW BAGLEY: I would just say a lot of these ideas about adding friction or doing other things on the front end have really been around in the community for years. And in fact, in the CCT Review Team's final recommendations to the ICANN Board, one of the recommendations actually dealt with incentivizing the adoption of best practices and to consider whether or not financial incentives should be included as part of that. And that's something ...

And obviously we're talking about GTLDs in the one context with that. And then ccTLDs, similarly governments could take similar measures there. But essentially, if you're able to identify these best practices, encourage adoption of these best practices, and then find some sort of competitive means to incentivize it—such as through, perhaps, discounts—then that's something that really could mitigate the impact to the competitiveness on the one hand and then really enhance cybersecurity on the other hand.

And granted, the adoption would need to be in a mechanism where you're really looking at this like you would any other risk; where you're doing risk mitigation where the risk is the highest and not just painting a broad brush and using the same methodology no matter who the registrant is or what the characteristics are.

LESLIE DIAGLE: Thanks.

EBERHARD LISSE: Can I [inaudible]?

LESLIE DAIGLE: Yes, please. Eberhard.

EBERHARD LISSE: As the ultimate timekeeper, yeah. We are running at the end of the session. So if everybody wants to give a short, closing statement, that would suit me quite nicely.

LESLIE DAIGLE: Great, and thank you. Maybe we should go in reverse order. Danielle, would you like to start? Oh, sorry. That would be ... Yeah, you go ahead.

DANIELLE DEIBLER: I think it would be Luigi, actually. But I can make a short statement. So I think it's definitely some interesting discussion here. I want to thank everybody for participating. And I look forward to working with everybody going forward to try to make the Internet a safer place.

LESLIE DAIGLE: Thank you. Luigi.

DANIELLE DEIBLER: You're muted.

LUIGI LENGUITO: Sorry, a double muting system. Really, thank you. I hope that the conversation continues in a very interesting subject. I think this is only growing every month we see about a 5% increase in Internet registration and notification. So it's up to us to altogether work on putting a stopgap to this problem.

LESLIE DAIGLE: All right, thanks. Drew.

DREW BAGLEY: Sure, yeah. I would just encourage everyone to really check out some of these core recommendations that have been around for years from the CCT Review Team, SSR2, now the EU DNS Abuse Study; and then look at the fantastic opportunities that are now being made available by some of the tools that are out there that are being [actualized].

And I'm especially, again, excited about the Centralized Abuse Reporting Tool from the DNS Abuse Institute. I think, hopefully, that will really help with some of those issues that I know the Secure Domain Foundation has identified for years in the community. And I really hope that we're at a state where the incentives are right for a lot of these best practices to be adopted.

LESLIE DAIGLE: Great. Thank you. Graeme.

GRAEME BUNTON: Thank you. First, thanks to GCA and Dan for putting all of this together and having me on the panel. Briefly, because Drew mentioned it, I think incentives, especially from registries to registrars to reduce their abuse levels, appear to be extremely effective. And that's a model I would like to see more of. Full caveat, they pay my bills. But the QPI program from .org, I think, is a really good example of this.

And lastly, thank you, Drew for that shout out. Please stay tuned to what the DNS Abuse Institute is working on, the centralized Abuse Reporting Tool. We're going to be inviting registrars to participate—likely late-March/early-April—sort of a beta test. And then really hoping to have it out there and in the public and available for anyone to use by June. Thank you.

LESLIE DAIGLE: Great. Thank you. All right, well I'm a long-time proponent of building out a free and open Internet. I have to say that, to me, it's really important that we continue to do that.

That said, I think one of the largest threats to a free and open Internet at this point is all of the unwanted traffic that exists on the Internet. We see that at GCA in our IoT honey farms and, as well, in the context of the Domain Trust project, all of the domains registered for malicious purposes. So I think the right best-step forward is for all of us to work together collaboratively, because that's what makes the Internet go, to try to reduce the amount of that nonsense going on and continue to build this out.

So with that, I would like to thank all of the panelists. And also I would like to thank the organizers for this opportunity.

Before we close out, I would like to remind you to please fill out the survey in the chat because the link is in the chat. We would certainly appreciate feedback. And I think that many of the panelists, and

certainly myself and Dan Owen included, would be happy to continue the discussion offline.

Dan, would you like to have the final word?

DAN OWEN: One more. Thanks to Kim and Kathy for their help with the chat lines and making sure that we have the survey links up there. And just to reiterate what Leslie said, we very much appreciate any time you can talk to respond to the survey. It's only 10 very short questions. Thank you, Eberhard, for the opportunity to present.

EBERHARD LISSE: Thank you, everybody. It was really interesting. And I'd like to do this roundtable once in a while because it gets a lot of people together to give different opinions. And often people who didn't know find contacts, get people informed, and start thinking about things. Thank you, again.

Now we come to what I hope is going to be the highlight of today. Dimitry Kohmanyuk from .ua is going to show us what happened in the last two or three weeks in his area of work while the Ukraine is under this attack by Russia. And I had to be careful not to start saying something impolite.

Dimitry, you have the floor.

DMITRY KOHMANYUK: Thank you very much, Eberhard. And I hope I can enable my video, although it's not strictly necessary. Can I? I think it's not allowed to me.

EBERHARD LISSE: No, no, no. You are on the panel. You can just click it on. It works.

DMITRY KOHMANYUK: Okay, now it ... Okay. I wasn't trying ...

EBERHARD LISSE: There you go.

DMITRY KOHMANYUK: Right. So let me see myself. I'm just going to quickly ...

EBERHARD LISSE: Yes, it's fine.

DMITRY KOHMANYUK: Well, I am glad to be here. I hope to see you all one day. And let me go to the gist of it. You have my file, so let me just say we got, what, 15 minutes? And I don't know if I can get the screen ... I mean the ...

EBERHARD LISSE: Dmitry, the slide is on. Kim can advance the slides, on your instructions. And you have as much time as you want.

DMITRY KOHMANYUK: Great. I'll just say next. I still don't see that.

EBERHARD LISSE: Next slide.

KIMBERLY CARLSON: Dmitry, are you seeing your slides up?

EBERHARD LISSE: Dmitry, you're on mute. You are on mute.

DMITRY KOHMANYUK: I'm getting muted for some weird reason. I apologize. Okay, I got them. Thanks.

EBERHARD LISSE: Okay.

DMITRY KOHMANYUK: So, well, the original presentation was about a denial of service attack that we had experienced on the 15th. Next slide, please. Thank you.

And I'll try to highlight the impacts. And there was supposed to be a long and boring talk about how we were affected. So I'm trying to shorten that up.

So we had been attacked ... Well, a lot of servers in Ukraine were attacked on the day that included the TLD and also gov.ua domains, Anycast, and known Anycast machines.

What happened is that one of our Anycast machines went down, but it happened to be one of those so called hidden secondary or hidden proxy servers. So that was part of the infrastructure that was responsible for distributing zones from third parties and from ourselves to the Anycast. And that means that knocked down—domino effect—some other zones, meaning that ... Well, they were not knocked down, but they were not updating.

Also, some of our providers have disabled our machines because they say, "Look if the normal traffic is 10 megabits and we're seeing 50 gigabits, maybe something's wrong. We'll just turn you off for 24 hours. Right?"

So I guess the lesson learned with that is you must, okay, separate public [facing and private facing] things. But the more important lesson, which I'm talking about later, was to always have a direct contact with as high a level person in the company using [inaudible] as possible. Of course, that means if you're running Google servers, unless you know Sergey Brin or somebody else like Pichai, you're probably out of luck.

Our main channel was Signal. And in fact, I migrated my entire company to Signal in what followed. And I saw Cloudflare reporting about traffic. Interestingly, the Signal downloads and the Signal traffic increased tenfold in Ukraine. So the lessons and ... Oh, next slide, please.

The post-impact actions and lessons we have learned were that we had standby Anycast providers because before the whole thing started, we were in the process of adding the secondary or tertiary DNS Anycast providers. I just switched that on. It was about 2:00 AM So nothing worked.

There was a [inaudible] cut and paste error. Luckily, I had the CEO in my contacts, so I sent him a message at like 3:00 AM Surprisingly, he answered. Surprisingly, he fixed the stuff at about 4:00 AM or 5:00 AM

Well, like I said, lessons learned. Know your CEOs. I know, not everybody can do that. But you know, if you have been served by the small company, that's easy. Maybe you know your VP or Engineering, whoever.

We did a press release. We had a big post-mortem write-up. And in fact, nowadays I'm encouraging everybody to do that. Even the small attacks are something you can learn from, and therefore you must have that analysis.

We did create a spare transfer infrastructure, and we started to rethink our entire zone transfer ecosystem because it was a mess. It's still a mess and ... Well, we'll get to this later. Next slide, please.

So this is the part where I'm getting a bit ... Okay, I'm hoping to stay politically correct. But I'm not going to be very polite here. The date of that was the 24th of February, which was exactly ... Well, not exactly. It was, I think, 11 days ago. Next one, please.

So, here I'm just stating the facts. Kyiv was attacked before on the 22nd of June. And I guess Hitler use the 4:00 AM time, and Mr. Putin decided to use the same. I'm not making it up, folks. That's really the truth. I was accidentally awake at 6:00. Something happened. My brain clicked.

I just logged into the Messenger. I mean, I had no idea. I think ... Or maybe my girlfriend woke me up. Anyway, my reactions were, of course, panic and denial. I thought, "It's a mistake." I thought, "It's a bad, stupid, rude joke." I'd been warned by some intel from my friends in the United States, but I never thought it would come to that.

I needed to call ... I was abroad at this time. I'm still abroad. I was on a trip to the European Union that kind of saved myself from going crazy ... Anyway, I called everybody in my team. Within the next 72 hours, we managed to migrate most of the infrastructure abroad. It was about like a 70/30 or maybe 80/20 rule.

Should I speak up, anybody? I'm not hearing myself, so I'm not sure if that's clear enough.

EBERHARD LISSE: Quite good, go ahead.

DMITRY KOHMANYUK: Okay, good. Thanks, yeah. So we had created what they called the "to save" list, like in the military, I guess. Okay, I haven't really served, but they say you have to triage. Who can you save? Who's already dead? And who can wait to be saved?

So I collected the list of what I call the business processes chains. So you know, you need an X, and X needs B. And the B needs QZ and so on. So it was [inaudible]. And so we started to migrate that.

My rule of migration was as follows. As we had established IP access lists, all the process and stuff. Nothing was touched in our production. Instead, I created a mirror of every system that was important. Okay, for Anycast that's easy. Right? Because then we just create an extra node.

But for each primary server or for each database replica, you have to decide ... You have to split our database cluster, for example, in two and create a sub-cluster with its own master and then switch over the replicas to that.

So we had some hardware abroad, but we never focused it on non-Ukrainian services except the DNS. So for example, we had none of WHOIS servers abroad. We had none of the replicas ... Okay, I'm not getting into details.

Again, we used the Signal chat and I was having daily scrum meetings. I'm not a software developer, so I heard that's what they use nowadays. Next slide, please.

And then I have to put my priorities. Well, I put them here. Next slide, please.

And those are ... Without the people, nothing works. That means also your customers. That means also your users. Right? But all of that, all of

that Internet, it's not existing for the machines. Right? I never saw a Docker script being injured.

Data. That means that if you save the data off of your database, worst-case scenario, you can rebuild the whole thing from scratch. Get a new machine. Install the [inaudible] or whatever, your EPP server of choice. Install your DNS servers of choice. Get some software. Get some help. You'll be up and running. Maybe you will be up and running faster.

Well, services. Obviously, you want to keep the services. But again, you have to prioritize. Setting up a second-layer SQL master was less important than actually having a dump.

Well actually, luckily we had the data synched. But for some stuff, we have to turn off, copy over, set it up, set it back. That included, say, our financial database which had our customer data. So for a while, our registrars can just use this thing for free because we're not even tracking their spending.

Money was, of course, the least priority. And it's easier said. But I started to get free offers from day two. So we are paying our bills, but as you may have not known, the Ukrainian government immediately prohibited all transfers abroad. That mean that every customer—sorry—every supplier ...

Let's say you've got Cloudflare for free. Yeah, I'm thankful to the Cloudflare CEO. And I'm not going to mention all of the companies. But the Cloudflare service we got for free is probably in 10K dollar per month range. I'm not really sure. I don't know. Maybe he'll turn it off

later. But I think that if you're in real trouble and people know you're in trouble, you get things being done for you. And don't hesitate to ask. I was hesitating. Next slide, please.

Components. That means that what we have done, well you know, as they say, Russian components, American components all made in Taiwan. Well, every ISP and ... Sorry, very ISP is probably pretty much the same. Every ccTLD is pretty much the same. The main component organization ...

And I'm speaking as, let's say, [inaudible] CTO. I don't have the title, but that's pretty much what I'm doing now. My formal title was like Director of Strategy or something.

People. Those are unreplaceable items in your workflow—called people. They can do everything. They can review the [inaudible]. Nothing else matters if your people are in danger.

We moved our EPP service [inaudible]. We moved our DNSSEC signing and key management [partially] abroad. Zone generation script and stuff, still being processed because we have some domains [signed] locally.

DNS service itself. The Anycast and other stuff, WHOIS and RDAP. I think our RDAP box is still in Ukraine. We don't really care because it's just a mirror of the normal database.

Oure website for the government, gov.ua registrations for the normal registrar backend, and for the public, they've all been migrated. Actually, that was the most important stuff.

Paradoxically, [how do you do services?] It's more important to have a website when you can publish a simple press release than an actual backend website where customers can login. Yeah. Because the first one says you are alive or, "We are under attack." That's more important.

Again, it's kind of a paradox of being a sys admin. I'm just trying to think that in reverse. You know, we do the work and then we communicate. Wrong. Communicate first. Do it later. Communicate again when it's done. Maybe do a status update that includes your team and everybody you talk to.

E-mail, chat, and phone support. Well, phone went out the window. Everybody's in Messenger nowadays. E-mail. Tell you something. People don't read it anymore. People using Signal, Telegram, Viber, whatever. No, seriously. People are more likely to call than e-mail in a moment of crisis. We can e-mail them all you want. Don't expect them to write you back. Don't expect them to read their e-mail in the next 24 hours. Maybe they'll read it next week. Next slide, please.

Other components or parts of our infrastructure. Thank you, Kimberly. While data center, Internet, networking hardware, that stuff was ... Well, okay, we are going to talk about this.

My main component is DDoS protection, already mentioned who we're getting it from. But this is not a complete list because some of our undisclosed partners are now providing us with DDoS-protected servers for the DNS. And that's part of the offering.

Cloud computing. There is no cloud. There is somebody else's computer. I think the cloud computing is a hoax to make people to basically centralize their data within the "safe confines" of Amazon or ... Do they have cloud in Russia? Oh, we would never find out. It's gone now. I mean, it would be local right? Somewhere in Siberia.

Anyway, the business back office. And that was actually the worst part. That's a system you never usually want to talk to the cloud abroad. It's just somewhere.

And again, I repeat the word "people" again. People are the most important component you can have. Next slide, please.

Decisions. I used to make maybe one or two big decisions during the day, [normally]. Okay, what to have for lunch? Oh, no. I'm not serious. But, you know, should I check in now? Should I check into changes later? Should I move to that or that? Should I renumber IP blocks?

Sometimes you do things because you're bored. Really. You're like, renumber your Anycast address or you sort your zone files alphabetically. No, seriously. I was making tens of decisions per hour. Next slide, please.

So I'm not going to bore you with all of the details, but those are the most decisions I'm trying to put down. I made this presentation a couple hours before this talk, although I've been thinking about it.

What do you outsource and what are you doing yourself? Well, hardware and data center? Yes. When you're under attack, pretend you have nothing but your computer. Maybe nothing but your phone. I did

50% of my work from my phone, just talking to people, writing them Signal messages. Communication work, that is.

DNS secondary service? We got [multiple here] and we're getting more. Like I said, we were in the process of ... Okay, we short-listed several providers. You can privately contact me and I'll tell you what I recommend. I'm not going to share my partners because I don't want my partners to get attacked.

Russian government people can literally come and kill you with Novichok or whatever. And I'm not making this up. Ask somebody who lives in the UK. So don't tell people where you have your infrastructure. By the way, that wasn't there.

EPP and WHOIS servers. I recommend do not outsource that. You may, of course use open-source software. But don't put your data ... We have been approached by multiple parties who said, "Oh, why don't you move the UA domain to Company X?" I said, "No and no." We want to own our data, privacy relations. It's like, I don't know, putting your family jewels to the banking safe deposit box. Okay, maybe the bank would go bankrupt. Maybe they would be raided.

Business and financial operations. Also I think no because if you can run the DNS, running the company is easier, especially now with things being online.

Virtual servers. Well, you run BGP for everything. Like, we use almost none of the IP addresses provided by the ISPs except the point to point. I can imagine somebody ... I think Belgium did that. I know RIPE NCC,

for example, has moved some of their stuff to the cloud—the United States cloud. I know people have been really complaining about that. So, no.

They would be interesting things like, oh, you spin off their Amazon box—and it has no v6, because, Amazon. Right? So you found out they'd have to stop it, include the DHCPv6 and restart.

Things are weird in virtual servers. If you run your own virtualization, though, it's great. And that's what we've been doing for quite a long time. We use Proxmox. It works.

EBERHARD LISSE: Dmitry, we have a hard stop in nine minutes.

DMITRY KOHMANYUK: Oh, great.

EBERHARD LISSE: And I want one or two minutes just to wrap up. I don't want to [inaudible].

DMITRY KOHMANYUK: I'm hearing you, Eberhard. No problem. I was listening to my own timer myself.

Registry and DNSSEC signing. I said no. And I did outsource our e-mail to the not-so-evil company called Google. Well, costs. Next slide, please. And that is. Yes, thank you.

Things that we have been paying for and things we're not paying. Well, I have reached out to 42 ... Well, probably more. I haven't really counted, but that's my estimate. A lot of people had offered me help. My rule of working with new suppliers was that if they're slow to work with—slow means takes more than four hours—I just drop them. I'm still tracking the costs. We wouldn't be able to get things done without lots of free services. Some of the services we got aren't free, but they're being free to us for now. Some people have offered consulting help. That's been immeasurable.

And I'm repeating my values slide. People are more valuable than computers, any day. And if you can get the one consultant or the 10 boxes, choose one consultant.

Again, time more valuable than money. Things have to change quickly. You may have the daily goals and you have to reach them today because tomorrow you'll have another fire. Like our database administrators spending 24 hours driving the car through multiple checkpoints like 20 kilometers per hour instead of 60 kilometers per hour and being unsafe.

Smaller companies generally react faster. Let's say I'm using ... Okay, I don't want to compare now, but Company A and Company B. Company A has 100 people on payroll and Company B has 1000 people on payroll. Well, I guess Company A would be easier to get to the CEO. The

company with 10 people may be even easier, but that company may not have enough, let's say, scale to help you. Next slide, please.

I'd like to express my thanks. Next, please. Next one, please.

Am I heard? Oh, yeah. Thank you. The phone wasn't updating.

All of my fellow colleagues, my team, and my suppliers—my product suppliers—acting quickly. And all of the members, too, of the ccNSO and TLD community and other people, too. The IANA staff which went through emergency change, and Kim Davies did update [inaudible] checking that everything works.

I'm thankful to CENTR for terminating—or suspending or whatever—.ru membership. And I'm thankful to RIPE community for its reaction. I don't always appreciate all the things they say, but I like that they're not ignorant. And I see something's been now said in the ccNSO community, but we have the plenary for that.

And I'd like to thank specific people in DNS-OARC staff. You know who you are, and thank you very much. Not just for now, but for all of the other meetings. And I hope to see you all and have you all and have beer with you all. And I hope we'll prevail in this war.

Please, next slide. And I'm open for the queues.

EBERHARD LISSE: Okay [inaudible].

DMITRY KOHMANYUK: And let me remind everybody. The military attacks of Russia started in 2014, and we had a ceasefire agreement which was essentially violated. Well, it was not the only thing. There, please.

EBERHARD LISSE: Thank you very much. There is no doubt that this is an illegal war of aggression which is, under no circumstances, tolerable which is a crime against humanity.

DMITRY KOHMANYUK: Agreed.

EBERHARD LISSE: There is one question before, but I want to quickly abuse the prerogative of the chair. When my colleague died five years ago at the last ICANN meeting in Puerto Rico, we had started already with writing down a handbook. It has grown up to 1,000 pages now, separated in many little chapters so that update is easy. But that helps.

Having a book, having a plan that you can execute in case any form of disaster helps.

DMITRY KOHMANYUK: I do agree.

EBERHARD LISSE: Rubens Kuhl asked one question. I have to rush because we need to get out of herein four minutes. Rubens, I'm going to not ask this question. I want Régis Massé to able to wrap up. Régis, you have the floor.

RÉGIS MASSÉ: Yeah. Do you hear me?

EBERHARD LISSE: Yes

RÉGIS MASSÉ Okay. I just put on my camera. Do you see me?

EBERHARD LISSE: Yes.

RÉGIS MASSÉ Okay. I have the slide on the screen, so I don't see my face. But if you see me, that's okay. So I will do my best to be as quick as I can.

So good evening to all attendees from Paris. And thank you, Eberhard, for asking me for making this quick wrap-up of today's Tech Day session. It was a very interesting session, as usual.

So we had nine presentations today. First of all, Graeme talked about the work for the fight against DNS abuse and the way to stop various threats. And a particular focus was made on action to be taken by a

registrar to stop the creation of fraudulent domain names before they were sent to the registries.

After that, Edward from ICANN made an update on DNS Core Census, already presented during the last Tech Day session. And the goal is to collect public DNS data from ccTLDs, gTLDs, and so on to conduct analysis and better understand the DNS architectures [and its specificities].

The thing important here is dataset are now available for use as well as a set of tools. So feel free to use them and get your feedback to ICANN. They will appreciate them.

After the break, Andrew presented the website on API for tracking DNS activity on root servers. Five years of data allow us to have a good [evaluation of the usual use and some evolutions like creation of short attacks.]

A particular focus was made on the Chromium Browser during the presentation to illustrate the works done on that.

After that, Craig, whose company operates new gTLDs showed us how to protect e-mail exchanges with an evolution of DMARC. These developments are based on the work of IETF Working Group, especially on the RFC 9091. Again, it's a matter of increasing the level of security of the [chance] to fight against the abuse of the [electronic mails.]

Gustavo from ICANN talked about ICANN monitoring tools and the [API called] MoSAPI which allows the monitor registry service such as DNS, RDDS, and soon [inaudible]. This service has been offered for more than

five years. For example, at AFNIC we use it on a daily basis to monitor .FR and gTLD registry services from the outside, and also to declare maintenance periods. We have very few incidents, but it's always a pleasure to talk with the ICANN team when they call you in the middle of the night.

[inaudible] .FR will soon join the DAAR program for ccTLDs, so I encourage all of the ccTLDs and gTLDs to join this kind of program.

Pablo made the usual host presentation. Thank you, Pablo, for sharing what the new convention center looks like. I wish I could have entered in the meeting in San Juan like I did five years ago instead of sitting behind my computer in my office. But another time, I hope.

The .pr registry offers a new platform of Internet services for small- and medium-sized businesses to develop their online presence. A very nice incentive, I think. Bravo for that. It's very fun.

After the second break, Kathleen and Paul took us in the heart of the evolution of encryption of the Internet in recent years. And the change of challenges that impose on operators in terms of architectures and services.

The DNS Abuse Roundtable. So, there were four panelists moderated by Leslie, and they shared their view and works on the fight against DNS abuse. It was about collaborative works, tools, project, and detection methods—all the things we need to make the Internet safer.

And at last, Dmitry spoke about the last DDoS attack on UA zone and the consequences of the war. [inaudible] And thank you for your testimony. It was very interesting and very [appreciated].

So that was a very short resume of the day. I propose a virtual round of applause for all of the speakers who did an excellent job this afternoon. Thank you all for that.

EBERHARD LISSE: Thank you very much. I would have really liked to be able to do a Q&A with Dmitry, but I'm told that we have a hard stop and we are already a little bit over it already. So it makes me feel quite ambivalent about it. But we can't do it.

We can repeat this or look at this in two to three months when we have our next Virtual Tech Day. And we have a standing offer, Dmitry, if you want, when you want. You just give me the word. You have 20 minutes or even longer at your convenience.

DMITRY KOHMANYUK: Thank you very much.

EBERHARD LISSE: I admire your work and the work of your colleagues very much, and I hope that you and your families will all be safe.

Thank you very much. We got it two minutes after the hard stop, so I hope we don't get punished. And I hope we'll see each other again face to face in The Hague in three months. Good night.

KIMBERLY CARLSON: Thank you, all.

[END OF TRANSCRIPTION]