BARCELONA – Tech Day (2 of 4)
Monday, October 22, 2018 – 10:30 to 12:00 CEST
ICANN63 | Barcelona, Spain

MACIEJ KORCZYNSKI: I have [poisoned] the local DNS recursive server and it will be serving the malicious information only to the users of Local Area Network. What about the attacks against authoritative end of the path, authoritative DNS servers?

So, in recent years, we observed the domain shadowing, for example, which is the process of using legitimate users, domain registration credentials, their logins and passwords, to create malicious sub-domains. So how does it work? The attack is very simple. The attackers first try to phish for credentials of legitimate users. Then they log in and they start adding sub-domains.

So this type of attack is very efficient for two reasons. First reason is that the users normally, the registrants normally, they do not check very often their accounts, and second thing is very often, the registrars, they have not only one but a few or several domains.

So in practice, it gives an almost endless supply of the sub-domains that can be used, for example, in phishing attacks or in [drive-by] download to the attackers.

A more ambitious attack victor is hacking the registrars directly. And here, we have an example from 2013 where Melbourne IT registrar was compromised and the attackers replaced a number of Twitter domains and also the domain website of New York Times.

So in this work, we explored an attack also against the authoritative end of the path, and that is the zone file of the authoritative name server using non-secure DNS dynamic update protocol extension.

So very briefly, I will give you a few technical details here. So the first RFC, RFC2136, edited by Paul Vixie, introduced a new type of DNS query which is DNS updates. It complies with the standard DNS message and the user can add or delete almost any type of resource record, [inaudible] and as you can name it.

The information, of course, propagates between this life and master server and once the server receives the query, the request for an update, it will verify if the prerequisites set by the requester are met. So for example, if the resource record exists and if the restrictions are met, if there are any.

So [while there are] security considerations, so we have three sentences there and basically, the RFC suggests in 1997 to use DNSSEC based on public key authentication between the client and server, and, which was, of course, resource-heavy. And then two years later, we had a new RFC suggesting using [SEEK], Secret Key Transaction Authentication for DNS based on shared secret between the client and the server, and that is one of now standards for securing DNS dynamic updates.

So as you can see, there was kind of a technological gap between the specification and security considerations.

So I will spend also one slide on implementations because this is also quite interesting. So in [BITE], DNS dynamic updates are disabled by default. The problem there is that we can put in the zone, the keyword "[hello] update" and then we can put, for example, IP addresses that are allowed to make updates. You can imagine that those can be easily guessed, and we can, for example, use IP spoofing.

What's more, instead of finding a set of IP addresses, we can use keyword "Annie" and then each single user in the Internet will be able to make changes in authoritative name server. So that's very interesting.

Even more interesting is implementation of Microsoft DNS, so by default, updates are only available via extended [sic]. On the other hand, non-secure updates are also allowed. And now, in Microsoft DNS, you can have two types of configurations, one with active directory and second with standard primary zones. And what is extremely interesting here, that when you set up your Microsoft DNS server with standard primary zones, then secure others are not available at all, only non-secure DNS updates are available. Please keep that in mind. I will come back to this later.

So what about zone poisoning attack itself? So what are the requirements? Non-secure updates are allowed and the attacker knows the name of the zone and it's NS server, which is pretty trivial.

Specifics of the attack? The great majority of implementations are based on UDP protocol so you need to send a single packet, no need to get response, and it's extremely difficult to detect because administrators do not check their zone files each day.

So what are the threats? There are a lot, but I will just briefly show you three of them. First, the attacker can simply replace the A record or MX record and point the domain to the IP address of the host under its control and he can then, or she can run a website or mail server.

We can also address reputations, similarly, to domain shadowing and add, for example, a sub-domain, Paypal.user.example.com. Even more interesting attack victor would be to make a delegation from the authoritative name server of example.com to our own authoritative name server for sub-domain. Let's say user.example.com and then we can add endless number of sub-domains, for example Paypal or BankofAmerica.user.example.com without interacting anymore with the authoritative name server itself.

So how to do that? We can simply use a standard unix comment, NS update, so here we go. So now maybe we could show the demo of the attack itself. So what we did here, so on the right hand, you see my website, internalmeasurements.online. It's just measureing the connectivity of IP Version 6. As you can imagine, I am the administrator of the authoritative name server and I misconfigured it. Then we just started checking on the left side, we're checking the IP address and NS server of internalmeasurements.com and now we are going to replace the A record.

So we will use the standard NS update. We define the server and NS1.internalmeasurements.online, the zone file name which is simply the internalmeasurements.online, and now we're going to make an update, delete internal measurements. Send the request. There are no errors, so let's check if it's resolving. As expected, it's not resolving.

So let's now use the same procedure and add our own IP address in association between internal measurements and the new IP address, so the same procedure again. Update@InternalMeasurements, and our new IP address of the host that me, as a malicious user, I'm maintaining.

Let's send the update. There is no error so it's perfect as expected. Now let's try to dig it, resolve it. We have a new IP address, which is great, and now let's take a look at the website and now it redirected to my new website.

So here we go. Can we go back to the presentation? So as you can see, the attack itself is trivial.

Now to make global measurements, we had some, actually, we had a lot of ethical discussions. Why? Because to verify the authoritative name server is vulnerable to a DNS zone poisoning attack, we need to explore the vulnerability. So as I mentioned, we had a lot of discussion with security experts, with legal, ethical partners and here is what we came up with. We were sending a single packet to authoritative name servers. We were not making modifications to existing records. We restored the previous state on all servers and what we were simply doing was we were adding a sub-domain to the existing domain and it was like research.example.com and we were redirecting it to an IP address of our host and on our website, we were explaining the study, all the problems, also how to fix it.

There was also opt-out mechanism if someone didn't want to [inaudible] in his servers. And finally, we send several notifications within two years to affected parties.

So we used, ALAC set up 1 million [as a] data set, also DNSDB from Farsight security. Thank you very much, Farsight. Project zone data, available zone files, and we only know we collected the 5 billion pairs between the domains and IP addresses of authoritative name servers. First, we ran the campaign to verify if the attack works and then we – five minutes only, so I need to be really fast – but all in all, we discovered 7,000 name servers and almost 420,000 vulnerable domains.

So as you can see, so here we examined the content of the website of ALAC, so 2 million. We've seen a lot of critical services, including nine financial services, and here, including the website of one of five biggest banks in the world.

So I will skip server's implementation here, but as you can imagine, the most vulnerable were Microsoft Windows. And let's move to notifications. So after the first global scan, we send notifications to DNS server providers using DNS [certified] authority name records, generic e-mail addresses, and host name at the domain. Website owners, this option is not available for the moment anymore, and also network operators, which is important from the perspective of GDPR because now, probably one of the ways is to notify the owners of vulnerable resources through intermediarires.

So very fast, what are the lessons learned here? Obtaining WHOIS data at scale is a huge problem. Nothing new, but we are volunteering here and we have to pay for WHOIS information because we did not have resources to scan and to parse WHOIS data. Contact information is extremely unreliable. Forty percent of all e-mails sent to domain owners fail to be delivered, so that's registrant information. RFC standards are widely ignored. Seventy percent of e-mails sent to persons responsible for the name servers affected by zone poisoning failed to be delivered using SOA, our name.

Eighty-four percent of messages to host master and abuse.domain generated delivery failure. We were a bit more lucky with reaching network operators, and here, 8.6% only generated failure.

So skip maybe this.

And here is our also ongoing study, and we are not [inaudible] CERTS, which is, as I said, important from the perspective of GDPR. So we started notifying trusted introducer CERTS and also campaigns targeted to national CERTS and there is, as you can see in red and in blue, there is a stable decrease in vulnerable resources, which is great. There was one very successful campaign where the great majority of domains actually were fixed.

Here, also, please visit our website where we maintain the resource where we show and make aggregations per search, per AS operators, and also country overview. And here, this is maybe – no, I will discuss a little bit – so on the right, I am not going to discuss too much about survival analysis here, but each line corresponds to a region, like Africa, Americas, Asia, Europe and so on.

And we can see, of course, if it goes down within the time, it's better. We see that Africa is [cleaning] very effectively and why? Because there were not a lot of vulnerable resources and actually, the intermediaries did their job. So for Europe, it's much more difficult because there are a lot of vulnerable resources and a lot of parties involved. And it's not that one operator has tons of vulnerable name servers and domains that can fix at once. It's like really, one person, one vulnerable server, a few domains. So to see the global effect, it takes a lot of time and needs to involve a lot of people.

So also, in terms of our notifications campaigns, in two years, we managed to reduce the number of vulnerable domains by 97% and the remediation in terms of name servers, 45%, which are really, actually, nice numbers.

So notifications in the context of GDPR. So what are our available options now? So we can use, of course, generic e-mail addresses. But as we could see, deliverability is very low. We use the start of authority, our name field, but again, the same problem. Through intermediaries, it's more scalable, but there are variability rates, but requires a lot of effort actually.

Web form, which messages could be forwarded to registrant e-mail addresses, it will never scale when we take a look at large scale notifications.

And the final option is anonymized e-mail addresses forwarded to registrant e-mail address. If implemented correctly, that might be an option. So to conclude, so here, we studied the problem of non-secure dynamic updates that was overlooked by security community and it's still an existing problem since 1997. We observed relatively low percentage of affected [house], but multiple important services.

And now maybe the last comment here, help us to remedy the problem. So if you're an operator or a registry, just let us know and we can make a scan or provide the data that we have.

So I need one more thing.

UNIDENTIFIED MALE: I'm not. I'm just telling you.

MACIEJ KORCZYNSKI: Perfect, perfect. Acknowledgments, so many, many thanks to Paul Vixie because I went to discuss with him. I told him my research idea. He said, "Go for it," so he was really, really encouraging. So thank you, Paul, very much, and also to Farsight Secuirty for sharing DNSDB. Without this global research, wouldn't be possible. Also CERTS for their remediation efforts and a few colleagues from NSIDN for their comments. So thank you very much.

EBERHARD LISSE: Okay, thank you very much. I'm wondering whether I should be scared, but eh ran over time, so there will be, unfortunately, no question time. We can take this offline anyway over lunch. Next would be Greg Wallace from NetActuate. There he is.

And in the meantime, now that I know it's going over, understand that it's more [skiing] than [inaudible].

GREG WALLACE: Okay, good morning. Hi. Thank you. Good morning, everybody.

EBERHARD LISSE: First, [inaudible].

GREG WALLACE: Okay. And just this one advances. Oh, thank you.

EBERHARD LISSE: It works. If you run over time, your question time will shorten.

GREG WALLACE: That's fair enough.

Okay. Well, good morning, everybody. My name is Greg. I'm with NetActuate and we'll just go ahead and…

EBERHARD LISSE: Just point it.

GREG WALLACE: Oh, I see. Okay. Let's see. There we go. Great.

So this talk is going to be about Anycast, so very briefly, I'm just going to do a couple introductions, talk about some best practices and then give a couple of sinkhole examples.

So I just wanted to very, very briefly introduce myself. This is my first ICANN so thank you very much for inviting us to speak. I started my career, actually, in Washington D.C. in policy. I was with a place called the Office of Technology Assessment. It was part of Congress. Has anybody ever heard of that or remember the Office of Technology Assessment by any chance? Just curious. No, okay.

Next, I was working on a variety of technologies. I worked on core routers. That's actually a Nortel OPTera Packet Core and I just looked up before coming here that the largest core router that you can buy today, I think, is the Juniper T4000. Am I right about that? Which is three-something terabit. That was actually 9.2 terabits back in 2001.

And then most recently, I won't go through all the boring details, but most recently, I've been involved in Open Source. I was at the Lennox Foudnation prior to joining NetActuate where I worked on things like the Node Foundation and Hyperledger which is blockchain technologies.

NetActuate very, very, very high level. It's been around for quite some time. Been doing Anycast since 2008 and grown our network really organically in response to customer demand to the point where now we're about the seventh largest in terms of number of peers. The only other point I want to mention is we are very involved in Open Source. We host and support the varnish HTTP cache. We also host the FreeBSD Foundation.

So okay, quick show of hands if it's all right. Who here is at an organization that is running Anycast? Show of hands. Gotcha, so probably more than half. So that's great to hear.

I have detailed slides on a lot of this, not all of it. The first is pretty obvious, but surprisingly, not everybody's doing it. You want to avoid a single point of failure. Anycast is a distributed architecture, obviously. That's the whole idea, and so global monitoring needs to be distributed as well so we'll talk about that. DDoS mitigation, it's always important, but with an infrastructure as critical as DNS, you need to really think about making sure that that's hardened and automated as much as possible, so some best practices that we've learned doing Anycast since 2008. Something that's important to keep in mind is you have to announce with even AS paths. So this is a key as well.

Make sure of BGP communities. This is a way you can get around the sinkhole problem, so we're going to talk about sinkholes, how they come about and that'll be clear how you can use BGP communities to address those issues. And then lastly, you want to have consistent transit providers because if you don't and you're announcing those prefixes to a transit provider in one location and not in another, that can be one of the ways that you can have a sinkhole emerge.

So avoiding network or vendor dependencies, I was listening to Packet Pushers. I don't know if anyone listens to that. It's a pretty good podcast and they had some folks on from ThousandEyes and they were talking about this global DNS performance report. You have the link here. You do have to register. It's behind a form. But it's a good report, and as I was going through it, I noticed this, which I thought was quite surprising.

Anybody else find this surprising, that 68% of the Global Fortune 50 still only has a single DNS provider, 72% of the FTSE 100? So that's probably a little more of a single point of failure than we might indicate as a best practice.

So how might you deal with that?

This is just one way of doing it, so we've had customers come to talk to us and want to have a highly resilient and redundant network, so you can do that by setting up these groups. The key is end users are going to be the same distance from any of these pops regardless of which group they're being served by and you have full redundancy so the data centers in each group are different facilities and different upstream providers as well, so you're doing everything you can to kind of isolate that risk. But it's not the only way to do it. But it's a way to do it.

As far as DDoS mitigation goes, I won't spend a ton of time on this slide. The thing is there's a lot of different tools out there. There are plenty of Open Source tools that you can use and that we make use of. There's commercial tools as well, and then as far as your DDoS mitigation plan goes, you really want to make it as automated as possible so that a DDoS attack obviously never happens at a good time. We were joking as we were talking about this presentation that wouldn't it be really bad if one happened now with everybody here and not tending to their networks? So you want to run drills. You don't want to figure out when a DDoS happens that your automated re-routing programs aren't programmed the way that you intended them or aren't working the way you thought that they were going to.

For monitoring as well, so I'll talk quite a bit about monitoring. Again, the top-half of the slide is just some tools that we're familiar with. We do make heavy use of Open Source. We use Icinga, so I've got a slide that'll show you kind of how that visualizes things. But there's other ways to do it and we have the link in here as well to an article that does a really nice job explaining how to use RIPE Atlas probes.

But as I mentioned in the beginning, because Anycast is a distributed architecture, the whole idea is to make sure that end users around the world are going to have an equivalent experience as well as have some extra security and redundancy built in. But really, it's about a consistent high performance. Your monitoring has to be distributed as well, right?

So the other thing, and we've actually run into this with one customer who kind of was [chinced] out a little bit on their monitoring network. They put it on a gaming network and it was in the right location and it was proximate to their Anycast nodes, but it was a far inferior infrastructure than what they had a lot of false alerts. So you don't want to do that. It's important. If one of the things that you want to do is make sure that you're staying ahead of any performance degradation that your customers might see, then you've got to have monitoring that's distributed and on an equivalent network.

So this is just making the point that you don't want to have an Anycast network that's widely distributed, but only have your monitoring in GO because that's only going to tell you what your users or what your customers' users in that GO are seeing. It's not going to tell you anything about the kind of performance that they're seeing elsewhere. Instead, it really does need to be a network that looks a little bit more like this.

And you can do this. You don't have to necessarily spend a ton of money for a proprietary solution or a commercial solution. You can use Icinga which is a tool that we use, and you can see. We've got the set-up and it's showing nice performance everywhere except Sydney, so the left is just a graphical depiction of the table on the right, and we've established a performance threshold so that it's throwing an alert for Sydney.

So this is a distributed monitoring system so you have a centralized server that's going to visualize your data and then you've got satellites that are distributed around.

Okay, so now we're going to get into a question of a sinkhole. So who's seen a sinkhole as described here? Just show of hands. Anybody familiar with this? Yeah, we've got a few folks.

So basically, this is a situation where you end up with a sub-optimal routing path and it happens unintentionally and it's usually the case that it's when you're deploying Anycast across multiple regions. We've typically seen it when deploying NIXs and the thing that's kind of interesting is the traditional orthodoxy is more puring is better and that's been kind of the thinking. That's not always the case though, with Anycast. So sometimes if you're not careful about it, you can deploy Anycast into an IX and end up with something that you didn't anticipate or want, which is this thing called a sinkhole.

So we actually encountered one just last week and our customer DNS filter was kind enough to allow us to talk about it. So this is an actual example. And kind of the high level, and we'll get into a lot more detail of how we sort went through troubleshooting it. Looks like we're having some technical issues here. Anyway, okay. Maybe one's happening right now. I don't know.

EBERHARD LISSE: Carry on talking.

GREG WALLACE: Okay, sure.

EBERHARD LISSE: Kim will work on it.

GREG WALLACE: Yeah, no problem. So I'll just go off of what I've got here. But basically, what happened was DNS filters customers in Belgium were going online and instead of being served by either the Frankfurt or the Amsterdam pops, the traffic was going to Johannesburg, so that was ending up with a pretty bad end user experience so there was quite a bit of latency.

So what we were able to determine was that DNS filter had recently deployed to Johannesburg and they had done so because they wanted to provide lower latency for their customers in South Africa.

They announced prefixes to the Internet exchange which is called NAPAfrica in Johannesburg and then they were looking at their client traffic coming into that server and they saw that there were some IPs from out of South Africa. So we worked with them and we did some testing and we were able to determine that, in fact, there were IPs and users from Belgium that were going to Johannesburg. So clearly, this is not what you set out to see when you're deploying Anycast. So this is an example of a sinkhole.

The fundamental, and I'll get into the specifics of how we sort of identified exactly what was going on, but the fundamental reason that this kind of thing can happen is because, as I'm sure most of you know, maybe all of you know, BGP doesn't know anything about latency. It has no clue. It doesn't know anything about geography. And so when we tested this traffic going from Belgium to Johannesburg, as far as BGP was concerned, it was a perfectly valid route. It was actually a really good route. It was only four AS hops.

To BGP, it was like, "Well, that's fine, no big deal." What you would be able to see if the slide was working, and he slides are going to be available so people will be able to see this and we have a lot of details about running the trace route in RIPE Atlas, but basically, the traffic, everything was fine. But between – let's see here – it was AS HOT 7 and 8. There was 150 millisecond delay that was introduced and so that ended up giving that end user an RTT of 171 milliseconds, which clearly is way, way, way past any [inaudible].

Is that okay? All right, cool.

So how did this happen? Right? I mentioned, obviously, BGP is geography and latency sort of ignorant. So what turned out to be the case was that… Oh, okay. I think I'm on Slide 18.

EBERHARD LISSE: Carry on. Carry on.

GREG WALLACE: Cool. So basically, what was happening was one network in the EU was peering with the out of region. It's easier to read it that way. They were peering at the Johannesburg IX, but they were not peering at either Amsterdam or Frankfurt or London for that matter.

So we were able to just kind of work backwards using Traceroute, Trace MON and a few other tools, establish peering with the carrier in those pops and within one day of contacting them and working this out, they were great. They were very responsive and after we established peering, this is now what you can see is the sort of after Traceroute. And again, this is using RIPE Atlas. So now we're down to 15 milliseconds. So that's obviously the kind of performance that you want when you're using Anycast, not 171, and the whole idea is to make sure that users are served by the pop that's closest to them.

But because you're leveraging BGP to do this, you can encounter situations where if a carrier, as in this case, isn't peering in those local IXs, and it only knows about the announcement in Johannesburg, BGP can say, "Well, sure, that's fine. That's only four hops, so we're going to go ahead and do that," and then that ends up kind of defeating the purpose of Anycast, or one of the key purposes.

So let's see here.

EBERHARD LISSE: Can you advance the slide?

GREG WALLACE: All right, great. Thanks.

EBERHARD LISSE: We're reaching that time.

GREG WALLACE: Okay. All right, I'll move quickly. So what we did with working with DNS filter was we performed pings, we identified where that high latency was coming from. There are ways to work around, if you don't get responses to ping, you can match up some of those questionable IPs or ones that have high latency with a database like Maximind.

Go to the next slide. So this is a second example, and I'll speed through this one.

EBERHARD LISSE: It's time to close.

GREG WALLACE; Oh. Okay, all right. Well actually, I'll let you all refer to the sinkhole example too, but this is where some AS prefixes were being announced to a carrier in the United States but not in Europe and so that can produce a sinkhole as well.

So there you go. That's kind of my presentation, but happy to take any questions that you might have and I, and a couple of my colleagues, will be here all week, so we'd love to talk to you about anything.

EBERHARD LISSE: Thank you very much. I found this quite interesting as a ccTLD manager where we use a lot of Anycast. We don't really want to have a clue about any of this, but quite interesting. I'll take one question from the floor if there is one. Otherwise, I would like to thank him very much. Give him a big hand.

And Jacques Latour is next.

JACQUES LATOUR: Forty minutes, all right. Good afternoon. No, good morning. My name is Jacques. I am CTO with CIRA and today were going to talk about the Secure Home Gateway Project.

In the ICANN meeting in Johannesburg, that's when I first presented about this. I had a bunch of slides. You can look them up, but it was like, "Would it be nice if we built a home gateway project or a home gateway device to attack? ccTLDs? Back then, all it was, was a bunch of slides and a story and not really anything concrete. We didn't even know if we could do anything around that.

But since then, I think it's been a year almost. We're actually working on a project, a prototype, and this is what I want to share with you today, or a status update.

Yes. So that's kind of a repeat of the old slide, but the idea is today, the way the home network are designed, they're not secure. They lack security. When you put a home gateway, they're wide open, anything can connect out of your home. There's no control in place and we need home networks that have active monitoring and that can mitigate attacks coming from your house to the Internet, and that's something that doesn't exist today.

So the home network, obviously, they have to be safe, private, and I think the key criteria for this project is easy to use. So Grandma has to be able to configure her home network to make sure that it's secure and that's one of the objectives that we have. We can't build a prototype that only the kids can use. We need to see if we can build this technology that is super easy for end users to support it. That's the key criteria here to do this.

The other thing, also, is that the home network has to be reachable from the Internet because now your home network extends to your IoT device that you bring with you and they need the access inside your home network. And today, there is a lot of cloud computing platform solutions that have an IoT device connect to the cloud and you connect to the cloud with your app, your mobile phone or whatever, and then there's a lot of information that you don't want to come close to your house that goes in the cloud because that's the architecture that works today. Since he can't connect him, you've got to expose out and make your information on the cloud.

So we need to come up with a solution that whatever is in your house, you can reach it and make it secure. IPV6 is going to make things worse once it's more out. You'll be able to connect your IoT device anywhere and that's something we need to make sure that won't happen.

So being able to reach out. So that's what I basically said, but IPV6 is going to make things worse in the future and this project is about, it's not a lot of people so while today, this doesn't work, this is about the future. In a couple of years, when people download Open Source software, they've got to support this framework that supports… Yeah, that's it.

So the key thing is for this project also is that typing passwords is a thing of the past. So we have to be able to rethink the way we provision the gateway. We have to rethink we provision the Wi-Fi connection. Having password, post your home Wi-Fi password on the board because that's how everybody can connect to your home network, we need to get rid of typing password, and part of this project is to leverage some crypto. Potentially, we're looking at getting DNSSEC on the home gateway to create security keys to allow things to connect in a trusted manner without typing passwords. So that's we need to get away with.

So those are the main highlights of the home gateway.

So how do we get here today? So when we look at the post-Mirai attack, that's when we determined that there was a lot of IoT device. They lack the security mechanism on the IoT device themselves. They were sharing passwords, so that's how they got compromised. So some people are addressing IoT security, the device themselves.

This project is focused on the home gateway, and having security measures in the home gateway to prevent IoT device from attacking ccTLDs or any Internet resource. So you don't want all your home camera, in the future, if you have thousands of IoT devices in your house, you don't want them to attack Internet resource. You want to control very strongly what has access to what.

So the goal is we need to create this, an Open source framework. Part of this is “I want to have all of this software, all of this code, open source and available, let’s say in the open WRT project, so that when people in the future, download the code, it comes with a security framework built in that protects the Internet from home attacks.

So far, we’ve built a prototype. I’ll do a quick demo of it. But there’s significant gaps in Open Source project for this. There’s a lot of people build a firewall feature to protect the Internet from breaking in your house, but there’s absolutely nothing that protects your house from attacking the Internet and this is what we’ve discovered and we’re trying to fill the gap.

So we started to build code and we have a functional prototype that actually works. But the goal is to protect the Internet from IoT attacks. That’s what we’re trying to do. And that’s the picture. So when we started to do this project, we did a bunch of research. I wanted to re-use a lot of Open Source stuff to create a gateway that would protect outbound connection and that did not exist. So we had to build this.

The second thing that we discovered halfway through the project is that there is a thing called MUD, Manufacturer User’s Description. It’s an Internet draft that just started pretty much at the same time that I started the Gateway project, and we decided to integrate this and this project. But the goal of the [inaudible] project for now is to build a prototype. So I wanted, instead dob having a bunch of slides, I wanted to have a real prototype of a solution that meets all of the requirements in the beginning. That means you don’t type passwords.

It's easy to use. I want the code to be, everything is Open Source. It's all on GitHub. There's a link at the end, how to get there.

We want to identify new standards that needs to be developed or updated to make this work, and luckily, MUD was one of the things that made this old thing easy to use or easier to use. And it meets the requirements, so we started to build a prototype. So this is what I talked about. The slide's out of order, but you can read it later. The goal is to protect against IoT attacks.

Same thing. That should be in the beginning.

So when we started the project, the first thing, a lot of people are saying we're going to connect our IoT device directly on the Internet. It's not part of this project and the prototype but this is just common sense. Somebody somewhere has to prevent people from connecting IoT device directly in the Internet because an IoT device is a dumb device, potentially doesn't have software support in the back. It can have vulnerabilities that can be exploited and the IoT device always needs to be placed behind the home gateway and enterprise firewall, a utility network that's controlled with security access.

So number one thing that we need to tell the world is don't connect IoT device directly on the Internet. So that's the number one rule.

But for this project, we wanted to mimic enterprise firewall framework into a home network framework. That means in the business, you do zoning. You have a database network. You have application server network. You've got [proxemic] work. You have a bunch of different networks. Each of the networks have different security access controls. And then you set all that up using your security policy and then you have a bunch of people that support your security footprint in your enterprise network.

So I tried to do all of that in the home Gateway where Grandma can manage this and I think we got most of it working. So we got to figure out what the IoT device is and what it needs to access. That's where the MUD profile comes in. So when you plug in a network on your home gateway, it doesn't have access to the full Internet. You need to tell it what it can access. If it has access to anything, like a water sensor is not have access to all of the Internet. They should have access maybe to GE which is the manufacturer of the water sensor. It should connect to GE to get the firmware update, but that's it. It shouldn't be able to go to Google or [inaudible] zone or whatever or any IP addresses on the Internet that you don't know what it's for.

So you need to have a per device access policy and the other thing you need to do is you need to monitor your network, so if there is a device on your home network that starts to scan the entire network, that's a bad behavior. So you want to put that device in quarantine.

So enterprise-type security control and the home gateway because people are saying we're going to have 100 to thousands of IUT devices in you home network. This is a large scale network, and you need to have the right security control in place.

So MUD started. You can read the spec there. But basically, it's adjacent [blog] or it's a descriptor of the device that the manufacturer makes available on the Internet.

So the example there is acme.corp. They make the water sensor and there's a URL there where you can grab the profile for the water sensor. And the MUD file basically says, "You know what? I made a water sensor. I need to upgrade my firmware at this domain name and then you can configure me here and alerts available here." So it finds the characteristics of the IoT device and this can be used to actually build the security policy associated to that device.

I think the biggest challenge for this project is to define what an IoT device is because people say, "Well, I have a smart TV that has a keyboard and a screen and I can have a web browser and I can connect anywhere I want on the planet with my TV. Well, it's not an IoT device anymore. It's a computer with a big screen. It's not an IoT device.

So we need IoT device. There's tons of definitions, but I haven't seen the right definition yet. Our definition is it's a small device that, it's an actuator or a sensor type of thing. It's not a computer. It's not a router. An IoT device is not a desktop or a mobile. Those are Internet devices. They're not IoT devices.

MUD profile are for devices that have fixed functionality. For home, for enterprise, for utility companies where you have 1,000. Electric meter, they should have the MUD profile. That's very specific, so that's where the idea is.

So part of our project that we have in Canada with ISOC is that we're trying to define a labeling convention for IoT devices. And part of this is MUD started halfway through this project, so we started to consider MUD as a label that can be associated to an IoT device. So in the demo that we do, we have a QR code that basically has a link to a URL.

But the MUD profile, it can also be electronic, so when the device boots with the DHCP, in the DHYCP request, it can have its MUD profile in there or it can be through LLDP. That means on the Wi=Fi connection, when it tries to connect to the Wi-Fi network, it passes the MUD profile information along with the Wi-Fi connection request.

And at that point, we can provision the device correctly. I think that's what is pretty good around the mud. So if, right at the beginning, the Wi-Fi, we can determine its water sensor and it meets these requests. Then we can provision it much easier and the user doesn't need to scan the device or do anything with it.

The other thing with this project so far, I think we agreed that the MUD profile is usable in connecting our IoT device for home networks. So that's one thing we wanted to know is can that be used as a viable solution? Because the new protocol, we built all the code around it and I think it is something that we should continue enhancing and using for provisioning IoT device.

So this is the architecture of the solution that we built. So on the secure home gateway, so we're using the [inaudible] from [season] [Nick], their home gateway because it's got a lot of CPU power to do this.

So I'll decode that we have actually can download and put it on that gateway for others. I guess it would still work.

So we built a MUD supervisor, a MUD controller and in the application, to support this, all that code is on get up. The idea of the MUD supervisor is when… See the challenge that we've discovered so far is if the vendor has a MUD profile, we discovered that we needed to enhance the MUD profile is Gateway-specific information. So we are very repository at CIRA for MUD profile and it's enhanced by the vendor's profile.

So the reason we have to create the MUD repository at CIRA is in case a vendor goes offline, they disappear or the device is not supported and the MUD profile disappears. Then it's not available on the Internet anymore, so we need to keep copies of the old MUD profile in case people provision the old IoT device, so that's what we're still trying to figure out but we need to have a repository of all the MUD profiles to ensure that devices can be supported.

So on the gateway itself, so the way it works is we have a D app, the secure home gateway app. So the green is the first step, so use the app to scan the QR code on the IoT device and based on that, the app sends their request to the home gateway and the gateway goes to the MUD repository and says, "Do you have this MUD profile?" If it doesn't, it goes to ACME.CORP, it grabs the MUD profile and then it comes back to the gateway. The MUD supervisor decodes the MUD profile and then he sends it to the MUD controller and the controller actually, based on the Mac address of the device that was being added to your home network, it builds the IP table for that mic address. So every device that gets added by default is deny, deny. It doesn't have access to anything on the Internet except the gateway itself, and then when it gets a MUD profile for he device, it adds the allow roles, so it allows the device to connect to ACME Corp to send alerts internally and to be managed by the app if it needs to.

So that's where we configure the access control to allow that sensor to have access only to a few things. The challenge with this is that the security lives in domain names and the security that we have on the gateway works with IP address and domain name and, you know they can change.

So we need a way of updating the domain name and the IP address on a regular basis on the gateway itself to do the mapping in case ACME.Corp changes IP address every five minutes, then we need to make sure the MUD controller keeps updating that.

So that's one of the things is we need a better… So I assume there was an Open Source project somewhere that would manage domain name to IP address and all of that automatically and it doesn't exist so I think we'll make that available. It is available, so we built something.

But we're in 2018 and stuff like that should exist. I shouldn't be writing this. Well, I'm not writing because I can't code but I get people to do this so far. But it's beyond me. This is basic stuff. This should have existed in 1992. We shouldn't be building this today.

So I have a demo. Kim?

So that's my IoT device. Is there some? No.

IoT device like this toaster would have thousand next generation toaster. It's got Internet access. Typically, a home Internet IoT device like this toaster would have access to all of the Internet when it's connected so it would have access to, let's say, [inaudible] website. Also, it would have access to all of the malware websites on the Internet.

So obviously, this is a behavior that we want to stop is having this IoP device, accessing malware site to download malware and compromise your home Internet network.

So most of the projects that we have around CIRA Lab [inaudible] project. So we got the MUD gateway, we have the slides, we have aroud 20 different projects to make this. So it's all open source, all of it local.

So that's the [inaudible] home gateway, so when you boot, it loads the MUD controller, the MUD supervisor that sets up the secure connection with the app, and that's a part where we forgot to add some voice to the homemade video.

[1:00:00]: All right, so in this device, we have the CIRA app already preconfigured with certificate, can talk with the router. And this, of course, we could have brought it in your cell phone or any other device that we want to have the ability to do this configuration. So I'm just going to launch our CIRA app here and then we have the option to just view devices.

And then this right is already talking to the router and we have a list of the devices that are already online. I'm just going to select our toaster here and then have some basic information from the device. We can see the IP that it's using, the Mac address, the version of the primer, some of the information here, and then we, of course, have the option to apply the money file. That's where you're going to apply all the security rules and the gateway.

So you're going to click "Add MUD file" and what that gets us is essentially the camera for the QR code, I'm going to scan it, I'm going to hit it again because I was too slow. All right, so it's telling us that this is a valid QR code and the information that it has is essentially a URL for the MUD file so it's going to download the MUD file and I'm going to hit "@device". It says "Apply MUD profile."

So it just finished what it did. Here's that, essentially downloaded the MUD file. It applied all the rules in our file and now the device is secure and it has only traffic that it can actually access and everything else should be blocked and secure.

JACQUES LATOUR: So the [inaudible] has now been set up on the home network and it should have access to the CIRA network. That should work? That works. Now all the different malware sites, that this could have access before so if we go back to this malware site, the access should be blocked.

No Internet access, so the access lists are preventing, only allowing the access to [inaudible] and block the rest. And that's the profile that we have for this next generation toaster 2000.

It's our Toastmaster 2000. so we actually had a couple of Google, Amazon, IoT device and we put boards and had everything working for the toaster, and in the end, the best way to just put an iPad and provision the iPad as an IoT console for the toaster, so it worked.

So the key thing with the MUD profile, the architecture, all of that is that it's complex to provision your IoT device, but we need to make it simple. That's the key objective, just move the… All right.

So the end user interface is the other critical part of this project. It doesn't matter what we build on the Gateway. If the general population can't use it, then it's not worth doing this project. So Grandma has to be able to provision this for her home network and the guideline that I gave the team is four directions. They are actual direction.

I said, "The only thing you can do with the app is swipe up, swipe down, swipe left, or swipe right." Diagonal, not good. Up, down, left, right. And that's the goal of the project, so all the security and the back end, the automatic password change roll over, the key exchange between the app and the Gateway, the Wi-Fi password, the Wi-Fi credential changing all of that has to be automatic. It's got to be user-friendly because if somebody has to write the password, has to type the password down, they've got to write it somewhere and then that's all they become vulnerable because typically they have weak passwords and make the whole thing less secure. So there's tons of reasons to try to build this with a simple framework on the front end.

So the other part that we haven't worked on but we need to do, so far, the only thing we did is we put a [inaudible] on the home gateway and you can detect just standard stuff, bad behavior. So if you have bad lettuce in the fridge and it's starting to scan your entire home network, you should be able to detect that and put the device in quarantine.

The thing we don't know yet with the quarantine is that what is Grandma going to do if she has a bad fridge? That's where we got to figure out. What's the simple process to quarantine, to tell a user, "You know what? This device is compromised. We know it's compromised. We've isolated. You can't use the fridge anymore. You can't connect to the camera in the fridge." So that's what we need to figure out because then, there's a process. Who do I call? Do I call the fridge vendor? All of that, so we need to figure out that part, but I think it's important to acknowledge that that's something we need to look at, and address, and build a solution so that this works.

So right now, people have compromised device at home and they don't care and it can impact ccTLD or any TLD operator or any Internet business to be compromised or being DDoS, so when we quarantine something, we need to figure out a process to manage that.

So the other part is VPN access to your home network so that if you want to see the camera in the fridge, you don't have to connect to a cloud provider to proxy the camera out there. You want to be able to have a secure connection in your home network to view the camera or

	any device on your home network securely without having to type passwords and stuff.
	So right now, we're provisioning the home gateway with a third level, .CA domain. I think it's going to be for free and then people, if they want to have their own domain for the home gateway, they can buy it as part of the provision process. But the idea is that it's signing with DNSSEC. There's a private key in the home gateway and then we can leverage a private key to set up all the keys for the VPN and the exchange that securely with the mobile app so that you can VPN securely to your home gateway that has a domain name and the key exchange is based on the DNSSEC.
	So this is something we're building now. WE haven't put code out for that yet and whatever standard do we need to write ITF drafts or whatever, we'll make that public.
UNIDENTIFIED FEMALE:	[inaudible]
JACQUES LAFOOT:	Eventually. Right now, all we are doing is provisioning the up, down, left, right on the VPI. But the MUD profile doesn't disappoint us yet, but there are things that the IoT device can do that I think we can provision using MUD, so knowing there's status, URL and the IoT device or status control or alerts or is there a camera, media? So different types of service IoT devices have and then you can provision

that on your home gateway to understand you can get the alerts from your water sensor and then what to do with those.

So I think there's more opportunity to build more functionality for a home IoT device using a standard way instead of being every device having their own proprietary way of communicating information in your home.

So the gateway is based on the [inaudible] gateway so far. All the code is on the CIRA Lab Github and these are the drafts that we've touched, looked at. Some, we're writing a few. We're writing. You can have a MUD profile and you can also have a signature for a MUD that you can prove that true signatures that this MUD profile is signed by GE and that's the real spec you need to because the issue is a device could be saying, "I may [inaudible] CERTing, and I need full access to your network." So you need to prove that Amazon actually signed that MUD profile, so we're trying. I think we have implemented that in the prototype so far.

That's it. Any questions?

EBERHARD LISSE: Yes, I have two questions. But can it cook coffee?

JACQUES LATOUR: No.

EBERHARD LISSE: This regards to the video that it was homemade. I would propose you ask your daughter to do it on your cell phone. She will tell you how to video it better.

JACQUES LATOUR: It was six minute video and it took us seven minutes to do it.

EBERHARD LISSE: Yes, and if your daughter is under the age of 14, she can do it much better.

JACQUES LATOUR: That's a problem.

EBERHARD LISSE: Any questions? Please identify yourself for the record, please, for the remote participants.

HOWARD BENN: Is this one working? It's Howard Benn from Samsung Electronics. I think I raised this issue last time around, There are other organizations working on similar projects. OCF is the largest. That's got about 300 of the IoT manufacturers in there today and they're taking a slightly different approach in that they're trying to ensure that the devices themselves are secured and the way that the protocols and the signing works is that devices could talk to other devices int eh house without having to go out to the internet.

But it's quite interesting seeing the main problem you'll find to solve which is limiting the access those devices have. So I was wondering whether you've actually interacted at all with OCF.

JACQUES LATOUR: I think I'll look at that quickly. The focus is like you said. It's on the home gateway to protect the Internet because we allow people who are working on the IoT device themselves, but if their vendor stops supporting the device and it becomes vulnerable and you can't update it, then you need that second layer of security that's there by default.

Any other questions?

EBERHARD LISSE: There is one more.

[CHRIS]: I'm with a security company called [Highus] from Vancouver Island, Canada.

Jacques, I really, really like what you're doing. I think that's fantastic. So one of the things I was wondering about, when you're talking about dealing with a device that maybe it's compromised, you quarantined it. Would it be possible to work with the manufacturers maybe to maintain a repository of the latest versions of the firmware for the given IoT device just like you're doing a local repository of the MUD profiles and so if you do quarantine a device, because it's an IoT device, it's not going to be storing tons of interesting data that you're worried about wiping.

Maybe you just grab from a firmware repository and re-flash the firmware to correct the device. Just a thought. It's not really a question, just an opinion. But yeah, really cool. Thanks.

JACQUES LATOUR: No, that's a good idea. Yeah, I like that. All right.

EBERHARD LISSE: Okay, thank you very much. Give a good hand.

Now we have our sponsor giving us a presentation about next gen firewalls and the pitfalls of commercial software can get you in. Afterwards, Kim and I will hand out tickets at the door and if you want to engage with the sponsor, please do so over lunch.

JOHANNES LOXEN: Okay. Thanks, Eberhard. And now from the Home Protection Tool corporate Networks, first of all, I want to thank Eberhard and everybody here for all the Tech Days that I enjoyed in the last years, and so I talked with Eberhard that there is no lunch opportunity this event, so I jumped in and I think ti's fun to do.

So Eberhard offered me a talk and I had. We had some discussions what Eberhard does find less boring and it was now about Next Gen Firewalls and the threats that we have because Next-Gen firewalls are a threat by themselves and we need them and this is the talk, what I'm going to give because in the corporate environment, you need a firewall that is secure and compliant and you have to survive a security audit and yeah, let's jump in.

We're here. Next slide.

Oops, it moves around. Okay, thank you.

Okay, we are a company. We are active since '96. We do lots of open source stuff in [inaudible] Samba which is an S&P implementation in open source and very nice, which is in information security management software. We deal with GDPR and ISO $27,001. We do lots of firewalls and VPN for mid-size and large companies and yes. We take a lot of care of inter-operability and Samba and very nice two examples of open sourcing an already existing close source solution. This is what we think that sooner or later, you have to have an open source alternative to a proprietary software juts to mention, open source is commercial because you can sell it so we have the opposition of open source and commercial this morning a few times, and we think that open source is commercial too.

Oops, I have to press forward. At least my remote does not jump forward. Oops.

Okay, I don't want to bring sand to the beach, but the close update, old school firewalls are [inaudible] where you just have the IP number in port and you have all those gun ports like, let's say Telnet and POP3 and you have still alive old school ports like DNS and NTPs, and today, more or less, everything is TLS-encrypted so you have fancy ports where you talk in an encrypted way.

You have proxies that are application level stuff, which I don't want to go into detail just to have them named for future use in this talk. And all of this is available as open source, so this is the most important part, that you can build an old school firewall with open source components and some people try to build them up in one machine, they invented Unified Threat Management, UTM, which is a fancy word but it's an evil thing because you, on a shiny weapon to face, you hide a lot of complex stuff that is all combined on one open source and on one operating system kind of like the firewall, the ALG proxy, male anti-virus, all the stuff that you can imagine and it's internally complex and it's in a crowd of attackable software combined by some crypt on the shiny web interface and this will not survive an assessment regarding ISO 27,001 in most of the cases.

And if you need to be compliant, you need to do much more. You need to set up a compliant firewall which consists of an external firewall, an internal firewall, so at least two firewalls. You have separate machines. You have separate services like Firewall ALG proxies, web server, SQL server, client VPN, site to site VPN, and this is what some people call shielded network so that you distinguish between several DMs at several local area networks and stuff.

So this is the rough picture of a compliance set up and now we face some challenges. Okay, so we face some challenges in the Internet area. That is, first of all, we have platforms. We have social platforms like Facebook, Instagram, Snapchat. We have shopping platforms, Amazon and eBay. You have CDNs like cloud and [inaudible] and recently now, we have DNS platforms like Cloud [inaudible] 4x1, Google 4x8 and 4x9, and all these directly talk via ports that your old school packet filter cannot look into. And this is a problem. You can just block it or you can let it go and if there is some [inaudible] traffic and other reverse SSH traffic in your connection, you cannot really see them in your old school packet filter.

And the browser goes full screen in the most of the sessions that our customers have. You can do almost everything via the browser. You don't need to install software anymore because you can do your office packet, you can do your media streaming, everything via browser and so the browser is your new operating system and this is encapsulated services in the closed environment, and this is where DNS over HTTP comes now, that you have a platform which is a browser and that directly determines you to ask a pre-configured DNS server. This is a discussion that we will have here on this conference tool between DNS over HTTP or DNS with TLS.

And this means that you can be subject to more sophisticated attacks that are above network level, so above the TCP level that your packet [inaudible] control because you are attacked on application level, you are attacked on session level, and even on identity and user level with social engineering attacks and you have to have some service that can identify your encrypted traffic without interception, without doing a man in the middle attack, and this is what a next generation firewall promises to do. The Next-generation firewall promises to make Port

443 visible for you for so the different kinds of traffic that is encrypted running over your firewall is now identifiable so that you can make new firewall rules for shopping, banking, remote control, voiceover IP or gaming so that you now can say some people are allowed to do Amazon shopping, some people are allowed to use PayPal or team viewer or Skype or even World of Warcraft. But you cannot do it with an old school firewall.

And you have some companies that often exchange firewalls and the next four slides are about why this is just UTM reloaded and you have all these companies. There are some more like, say, [inaudible] and Elvis, but these are, I think the companies that offer the most mature Next-Gen Firewalls and they all have some problems and the main problem with them, they're all closed source. You cannot dig into them. You just have to rely on the [inaudible]'s promises.

What are the problems? First of all, now you have inside, one machine on one corner. You have the full awareness of packets like before, of applications, of sessions and of identities. They all promise to connect to your internal active directory inside your local LAN so you can have user regulations like this user is allowed to use Team Viewer on afternoons or whatever you want.

So these are really machines that allow you to do anything on the old packet level and also on the highest session and identity level. So they have GDPI, is used a lot because if you say, "My corporate compliance management is based into my active directory users," so implements what people are allowed to do or not and you give that to your firewall, which with the other leg standing in the Internet, you have a big problem.

And if you do SSA encryption, this is another thing that I don't want to go into detail because this is another thing. But even if you don't do SSA decryption, you have big GDPR issues and you have a closed source operating system so if you look to the CVE's course, ten in the last three years, both companies, [inaudible] Juniper [inaudible] and Cisco had scored 10 CVE bucks and these are not bucks like that you encounter from time to time with careful engineering.

These are bucks like fixed passwords for all the machines with root access and stuff like [Fortunet] had a few years ago. So you could dig into every [Fortunet] machine if you know the root password from remote and in the moment, there is no open source solution inside and if you like, they all offer cloud service. You can call a cloud server that you cannot control the back door. And so these firewalls have back doors and you cannot control them.

So what to do? And this is the last slide. If you use Next-Gen Firewalls because you have to, because you have to look into the 443 traffic, you have to build firewalls to enter just so build them multi-tiers, put one firewall in front of your Next-Gen Firewall. Put something behind. You have to protect and observe your next generation firewall very carefully. Look into the look files. You really have to stand next to them and really have to have a daily or weekly inspection.

You have to utilize your information security management system really carefully. So that means reflect and implement the controls and measures that ISO27,001 or other standards really wrinkle [inaudible] for firewalls, especially for a Next Generation Firewall. This is very important. You have to know your information security and data protection officers. You have to tell them what you're doing if you're running a Next Generation Firewall. You really have to make them aware that such a machine is in your corporate network.

You really have to carefully relegate cloud services for a while. Don't use them. Don't send your files to send box services to get them analyzed and get some results back, but to send your secrets by this way to that remote sandbox. This was from a provider that you, perhaps, do not Trust because cloud is just computers that are owned by other people, not you, and so this is very important that you see which machines are under your control, which machines do you send secrets to and which not, and if you have a chance to color bright on the open source solution, that we get this open source too, just take the chance and that's what I have to tell to you. Thank you.

EBERHARD LISSE: Okay, thank you very much. I must say I wasn't even aware that such clever firewalls that could, they're not really firewalls. It's just in a packet so traffic analysis on particular ports are running. Is it really necessary to use those? For my particular system running registry, it's not but the more competent this becomes, the more interesting this becomes.

But the point here is, we, Europeans, and I count myself as a European for this purpose are more aware of data protection than Americans are. Americans call this over-reach whereas when you go to the States, if you want to buy a house, you're supposed to have credit. You only get credit if you have loaned money. When you loan money, you have a credit score. A credit score is three different companies and you don't want to give them your information. They just have it.

When you drive around in New Jersey, you find police cars that have readers that read routinely ever number plate passing by and check whether they… That kind of thing, I, as a European, don't want. Now we have to use firewalls or we may use firewalls where we send every single thing that is outlawed. The data that I don't want to get around to a company outside of the Jewish diction to analyze it. I find that concept not acceptable, personally.

Any questions from the floor?

Okay, then the last question remains. It's where you get your meal ticket. Kim, thank you very much for this. The presenters, of course, have first dibs and all others can get a ticket from Kim or me at the door when we leave, so just wait for us to stand there.

[END OF TRANSCRIPTION]