Episode 89: The Three Buddy Problem

Handala wiper attacks, APT28 implant devs are back, Signal's verification problems

NOTE: The transcript is auto-generated by an AI tool, with obvious misspellings and errors. This doc is open for copy-editing via **Suggesting** mode. <3

SUBSCRIBE and WATCH on YouTube
https://www.youtube.com/channel/UCM-MHaHEKsQJOvHiV9k8jPg?sub_confirmation=1

LISTEN everywhere
https://pod.link/1414525622

Pre-Show Banter

JAGS (00:00)
...to at least take a little time. And like, I don't know, I'm scared to say this, but for once I'm enjoying San Francisco. Like it's nice — I went out, the weather's nice, the shopping's open again, and friends are around. I...

Costin Raiu (00:16)
You know what I remember about San Francisco when I was there? Like the last time I went to one of these — what was it, Walgreens or something like that — and I bought some duct tape and some

rope and things like that, simply because you couldn't find them in Romania. It was just crazy. Duct tape and rope and some matches, and I got like a big…

Ryan Naraine (00:32)

Jesus Christ, what were you doing?

JAGS (00:35)

Duct tape and rope? I was there with you the last time you were in SF. What the fuck were you planning? We could have…

Costin Raiu (00:44)

...package of all these things which looked like I was going to do something unthinkable, and I need all the rope and the duct tape. But I took it home and my wife looked at — I said, "Look, let me show you what I brought from San Francisco." And she was probably expecting, I don't know, Apple Watches and iPads — which I did have as well. But I showed her the rope and the duct tape and she was like, "What the fuck is this?"

Costin Raiu (01:13)

Well, it's good, because we don't have it in Romania, so it's valuable stuff from San Francisco. Duct tape and rope.

JAGS (01:20)

We take so much for granted.

Show Open & Sponsor

Ryan Naraine (01:20)

Good morning, everyone. It is the Three Buddy Problem, Episode 89. Juanito is out in San Francisco for the Nebula Fog hackathon, the AI hackathon tomorrow. I'm heading out there later — as soon as this podcast ends, I'm heading to the airport, which means this may be a very short episode. This episode is sponsored by TLP Black. Costin, I'll put you on the spot. What does TLP Black do? Let me see if you've got your marketing message down, pal.

Costin Raiu (01:46)

Everything you need, my friend. You need threat intelligence related — we'll do it for you. Anything — from passive DNS, IOCs, YARA rules, incident response, trainings — everything you need. And we do it with quality and class.

Ryan Naraine (01:54)

There you go. You see, the man has got his marketing message down. Quality and class. Bespoke services. Shout out to Curtis Baumgartner and the TLP Black team. Thank you for sponsoring the pod. Let's get right to the show quickly since I have to run.

Nebula Fog AI Hackathon

JAGS (02:19)

You're not going to ask us how we're doing? You're not going to ask us about RE//verse Orlando, nothing? Costin, yeah.

Costin Raiu (02:19)

Before you run — what is Nebula Fog? You keep saying it. You're going to Nebula Fog — what is Nebula Fog?

Ryan Naraine (02:35)

Nebula Fog is — it's not a cybersecurity event. It's an AI hackathon bringing in people to just build and ideate and just shoot the breeze around what is working, what's not working, what's connecting, and people actually building actual projects. I went to the first one last year. It's organized by Rob Regan at Bishop Fox. I went to the first one last year and helped him behind the scenes — kind of helped to host and moderate some of the things there. It was fascinating to see how these kids in San Francisco, a year ago, were thinking about things and use cases for AI. There was mostly a cybersecurity leaning, but it was beyond that. Juanito, you're excited to see why.

JAGS (03:15)

I mean, I'm kind of hoping there will be more of a cybersecurity leaning either way, but to be honest, look, you said this is an AI event, it's not a cybersecurity event. What I'll say is that there's no such thing as cybersecurity without AI anymore. Which is to say that this may not be a cybersecurity-focused event, but every cybersecurity event has to be an AI event in some way now, just for the sake of capacity. I mean, it has to — we have to figure out ways to make it apply.

We are an industry that's not like — it's not like we should be scared of AI for possibly taking our jobs. It's that we should be hopeful that AI will maybe make our jobs doable and tractable and potentially successful. And that's up to us, right? The AI people — yeah, they may have an interest in security, quote unquote, but they're not going to figure out how to properly leverage their stuff towards genuine cybersecurity problems. Like actual foundational important capabilities that we need to build and improve and eventually monetize. But we have to figure that out ourselves. So if cybersecurity is not successful in the next couple of years because of AI, quote unquote, it won't be because of AI. It'll be because we didn't figure out how to leverage it effectively. So yeah, I'm excited for Nebula Fog.

Ryan Naraine (04:44)

Sounds like the early makings of a keynote. Costin, that answer your question on what Nebula Fog is?

Costin Raiu (04:54)

It was more of a rhetorical question, just in case somebody who listens to our podcast doesn't know. I Googled it before, to be honest. I definitely thought, look, this is an interesting event and perhaps it doesn't get enough traction on X or on social media. It would be good to speak a bit about it. And I was also curious, by the way, if you guys have any talks yourselves there.

Ryan Naraine (05:22)

No, it's not a conference. There are no talks. I'm bringing my podcast gear. I'll be doing some individual episodes there based on some of the things I see. Juanito, just to piggyback on what you just mentioned there — and we'll go off the rails here for a little bit. How are you thinking about

hiring in this AI-plus-cybersecurity realm? Are you more interested in training your people, or are you going out to things like a Nebula Fog and looking for AI talent?

I'm just curious how companies are thinking about implementing AI internally — AI experts versus "we'll get our cybersecurity people ramped up to figure out the AI piece."

JAGS (06:03)

I mean, there's several pieces in what you just mentioned and I want to be candid, but I'm trying to be kind of careful in how I put this. I think I tend to be very quick and glib to make jokes and then — like things like this, right? Like employment and stuff that people are so touchy about right now that I just don't want them to get the wrong impression.

I don't see a lot of companies that are all that proactive about hiring in general right now. I think the fear is, you know, the corporate stances seem to be overall that everyone is overhired, everyone has too many people. So, you know, how am I thinking about hiring? I'm honestly not, most of the time. I am very interested in ramping up my team and our resources where we have a wealth of domain expertise for cybersecurity in our ranks. And we actually have a pretty good bench of AI folks, but we don't — you need to cross-pollinate that.

So a lot of what SentinelLabs has been doing — and if you look at the SentinelLabs publications from the past couple of months — all of it is security people engaging with AI things. And some of it, I think, has done a pretty good job of just kind of documenting some of that. Saying, "Hey, I tried this, I tried that, this doesn't work, here's all these terms everybody's using, I don't know what the fuck they meant, so let me kind of explain it to you." And now it's turning into experiments of, "We tried to do X, Y, and Z in threat intel, and this is how it worked, and this is how it didn't work," and so on.

But that's my way of saying we're spending a lot more energy trying to ramp up domain experts in security to be AI-native and AI-literate, and just giving them the resources and pushing in that direction. At the same time, I'm not saying there's no room for new talent. And if you have people who are very good at AI or AI development, I think anybody will hire them. Everybody wants them, whether they know how to use them or not.

But what I think is worth saying is that I don't know how much of an attractive proposition any of these companies can actually make for very good AI people these days. And it's not because you can go get a $2 million offer from Meta or a $10 million offer or whatever. Actually, everybody might be a little scared in the market. But the thing that seems to be working very well is: two or three very smart people go start a company to do one really cool hard tech thing with AI, and within a heartbeat you're getting acqui-hired into some big company for much more money than they would have offered you individually, with better stock, and so on. And you go on to kind of have done this cool thing that you wanted to do and grow it at a big company, or choose to rinse and repeat — drop out, start a new company, go again.

If anybody's looking for a model for success, I think helping ramp up these two-to-five-person companies to do a specific thing and actually build something legit — at a time when most big companies can't seem to build anything — is the move. Because you can't, right? You go to a Cisco, go to like, hell, even the mid-caps, and you go in like, "Hey, why don't you guys develop this thing?" And they're like, "Well, if you want to go and develop it, we'll buy you for $20 million." Six months — but it's better to do that than A, to try to hire, or B, to try to build in-house.

So that's a really roundabout way of not answering your question, but I think hopefully providing better feedback. Which is: I think something like a hackathon like Nebula Fog is interesting

precisely in that, A, you see how people are building — you see what it is that folks are doing, like how they go about their work, which is fascinating in itself. And B, whether they have the best idea or not, seeing a couple of folks that clearly have good synergy, they're clearly entrepreneurial, they know how to do the AI thing, they know how to build, they know to a percent how to talk about their ideas — that's the kind of talent I'm interested in right now. Not necessarily to bring into any team of mine in particular, but in order to motivate them to go do their own thing.

Ryan Naraine (11:18)

Yeah, no, it's fascinating because that's where I see a lot of the AI VC investments going as well — into building those little small tiny tools that you're talking about and getting very specialized AI tools and prepping them for acquisition. So I think you're onto something there. All right, let's…

JAGS (11:35)

I hope so. I mean, look, think about it. Sorry, last thing — think about it. Since you brought up the VCs, think about it from the VC standpoint. Getting into a Series A these days or beyond is an insanely expensive proposition. It's super costly.

Ryan Naraine (11:47)

Costly. Everybody wants to get in at seed, and that's why I think there's going to be a bunch of VC guys coming out to Nebula Fog as well, just to peek around to see what people are building, where the opportunities are for these very early-stage seed investments. So it'll be fascinating to see. San Francisco is a good place to be if you want to connect to the AI hype for a little bit. You know, I've been the anti-hype for a while, so I'm going into the machine, into the fog. So maybe next week's episode will be…

JAGS (12:17)

Ryan's going to become a VC soon and that's when we're going to see the true monster be born here.

Stryker Cyber Attack / Handala Group

Ryan Naraine (12:22)

Let's switch to the news before we run out of time. The big story this week — Costin, Stryker. Stryker is a medical device manufacturer, a US medical device manufacturer, pretty big company, that sustained a cyber attack. This Handala group, which is a pro-Palestinian Iranian group — am I getting it right? — took credit for it. It's kind of a hacktivist group. The attackers posted a message saying they compromised over 200,000 servers, systems, and employee devices, wiping many of them, exfiltrating about 50 terabytes of data. Stryker confirmed a severe global disruption affecting all laptops and systems connected to their network.

The Iranians taking credit for this obviously spiked this to the top of the front page of the New York Times and some other publications earlier this week. Costin, did you get a chance to look at any of this? Did we get IOCs? Did we get a sense of anything at all beyond just these headlines?

Costin Raiu (13:23)

I thought, by the way, last week — remember we were talking about what does it mean, this conflict with Iran? Will it turn into cyber? Do we need to fear something? How capable are the

Iranian groups? Are they going to wipe or, I don't know, open the floodgates or break the dams or whatever? So I think this is a good example of what they can do.

And I think it's good that this happened, in the sense that people will become a bit more aware about what is possible, what is on the table, and what are the limits for this. Handala — they've been around for I think at least two years, maybe more than two years now. And it's one of those groups which I think are wearing several hats, in the sense that they are kind of an entity associated with nation-state-sponsored groups — MOIS in particular, so Iranian — and at the same time they've been doing these ransomware-style attacks. They have Telegram channels which they keep getting shut down, so they open new ones, and they constantly target someone or go after someone with wipers, with destructive activities. Remember, the BiBi wiper was one of maybe the most famous things. The wiper wasn't very capable.

JAGS (14:43)

Yeah, the video was better than the fucking wiper.

Costin Raiu (15:05)

What I think is, by the way, in this story with Stryker — which is an interesting case in particular — there's this military angle that I think multiple people pointed out. Maybe including Kim Zetter, who wrote this extended story on it on her website, which is zetter-zeroday.com. There is actually Stryker military equipment, which raises a question: is this what Handala were going after? And they somehow mistakenly hit this medical device manufacturer, Stryker, which just happens to have the same name.

Nevertheless, even if this was not the purpose and it's just an accident, I think it was quite disruptive. We know this has affected the employees all across the US, Australia, India, and so on. And in particular, I think that this attack is a really sneaky one in the way it was actually executed — through a mobile device management solution. You know that a lot of people are allowed to bring their mobile device to work and then the company installs some kind of an MDM solution, which essentially allows them to, let's say, wipe the device in case the device gets lost so the company data is still safe. So they implement Intune, right?

JAGS (16:18)

Is that the Microsoft Intune thing?

Costin Raiu (16:23)

Correct. But this also opens an interesting possibility and door for abuse, because the moment the attackers get control of Intune, then they can remotely wipe your personal device that you were using for work. Which is exactly what happened here. And I remember we were also discussing, back in our previous job, the IT security people really wanted to put some kind of software on our devices, and we were like — staying away from that thing like it's the devil. Like, we don't need that shit on our devices. On laptops too. For example, my wife has this MDM solution on her work phone, and now they're talking about potentially allowing employees to bring their own home devices to work and then integrate it in MDM. Which again, it opens very interesting doors for abuse.

And how Handala destroyed employee devices — how they wiped employee devices — in this case, I think it's an interesting lesson of what is possible and how these cyber attacks can actually happen in the future.

Ryan Naraine (17:38)

If you're an Intune administrator, isn't there a place where you can figure out roles, permissions, figure out who's in there, who's capable of wiping, lock that box down? I don't know how it works. I have no idea. But what is the mitigation for this? If it's meant to be a wiper, it's there to serve as a wiper.

Costin Raiu (17:58)

So I think it all breaks down to the way you implement security overall in your infrastructure. And it's the same issue — you remember that many, many times we've seen very interesting attacks in which the ransomware groups were actually leveraging the EDR, the security center, the interface to which the EDR talks, to distribute the malware. So they were using the security solution to distribute the wipers, to distribute the ransomware to all the endpoints where the EDR solution was installed. And this was possible because, I don't know, there was no password or no enforcement or no signature enforcement.

In the end, if you have powerful tools and powerful solutions like EDR, like Security Center, which can push scripts, can push packages and execute them in your entire network — allegedly for the purpose of remediation or to implement some kind of a playbook — it can also be abused. So what is the solution here? I guess you can still try to air-gap networks, to segment them, to limit the number of hosts from which these commands can be sent. But again, all of this increases complexity. It makes it harder to manage. And suddenly it means you cannot do that work from home. They either have to be in the office. The backups need to be secured. And it's always this balance between convenience and how easy it is to manage the network and even remediate attacks. And the moment somebody actually manages to breach you and they take over all that infrastructure, it's pretty much game over.

JAGS (19:53)

To be honest with you, this is where — I don't think — yes, there are ways that very engaged administrators could do better and we encourage them to do better, right? That's almost always the mechanism that we have available to us and we do well to encourage it. At the same time, I would place a lot more of the responsibility, or at least the call to action, on the people that are developing and maintaining these solutions. For domain management and all this sort of tooling — because frankly, the heuristics around this being abused are far more visible and immediate when you're developing that software than when you're administrating it.

The administrators are just trying to enable themselves to do what they need to do. And as Costin said, you want to be able to do things from home. God forbid you're on vacation and the whole network goes down and you need to help people make it work. And chances are that you've got five people that are just desperately trying to keep 200,000 systems online. So when I look at that, I think of the MDM solutions and the GPO policy capabilities in AD and these other places where you go, "Okay, what should these admins be able to do without any friction?"

Maybe you do need to wipe one device or two. Maybe you do need to push a package of this kind to this many systems. But how often do you need to be able to wipe every device in the company? How often should you be able to push an executable package to every system and have it executing with no friction whatsoever? I'm not saying you don't need the capability. I'm sure that there are good reasons. But no friction, no 2FA, no "Hey, I need you to" — fuck, for how often this might happen, you could literally have somebody call them and give them a code. There are ways that you can do this.

And I really do think this is an area where — I'm not trying to fault the manufacturers all the time or give Microsoft shit in particular — any EDR that has remote management capabilities, including our own, you need to stop and think about what is the actual workflow around these things. And some of these are nuclear options. And if you have a nuclear option in there, no, you don't necessarily take it away, but there should be two people pressing, unlocking the keys at the same time. There has to be something.

Just to kind of expand on this example — I don't know if Handala figured this out on their own, but this is basically the mechanism that has spread most of the wipers.

Ryan Naraine (23:13)

It's the ransomware chains everywhere. But even the ransomware groups everywhere — they're abusing these RMM tools everywhere, like ScreenConnect and these things.

JAGS (23:21)

But you're thinking RMM tools. What I'm thinking is GPO policies. So the Russians would log in — in the Ukraine conflict, right? They already had access to these garbage fucking appliances and they were just sitting on the appliances without any — nobody was going to notice that they were there, because there's no way to secure these stupid appliances. And they already had admin passwords, because they'd been sniffing the network, they'd just been sitting there watching things. So then you use the admin creds to hit the domain controller. And then you use the group policy overrides from the domain controller. And you can literally, as an administrator, push the wiper onto every computer and have it execute.

That's why they were able to just hit one network after another after another. Because if you were already in the perimeter, you already had the creds, and you knew where the AD was, then it wasn't going to take anything for you to have it execute a package across the entire fleet. So those are moments when I don't know that I can look at the admins and say, "Hey, you guys need to do better." You can and you should, but some of these problems are just far more ingrained in what the technology is allowed to do that it shouldn't be.

Costin Raiu (24:37)

You also asked about IOCs and I forgot to say that — at least I have not seen any IOCs from this case.

Ryan Naraine (24:48)

One of the things — do we know how Handala got in? I saw mention of identity-based attacks where they do exactly what Juanito just described: sniff some credentials, get a hold of some credentials, log in, figure out the Intune admin access, and do this mass wipe event. In that case — when the primary attack vector is identity-based — how do you go about thinking about your perimeter and thinking about how your network is set up when it's a password, a login, a pivot, and a boom?

Costin Raiu (25:22)

Well, I would typically recommend a combination of canaries. We've had Thinkst Canary as a sponsor. Shout out to Haroon and his company — they've made some really interesting technology. So a combination of…

Ryan Naraine (25:38)

What's the potential use case of a Thinkst Canary that would help pinpoint this?

Costin Raiu (25:43)

Well, for sure — as they try to pivot inside the network and try to expand that access, they will for sure trigger one of these canaries. Maybe they manage to get their hands on credentials or things that they can use to manage the entire infrastructure. They will for sure trigger some of these canaries. Or even in the initial phase where they try to probe defenses, try to pen-test or check some of your websites, services, online portals — they can possibly, they can likely, trigger some kind of warnings if you deploy a canary in advance.

So yeah, combination of that together with multi-factor authentication. And you know my favorite part — it's always based on the idea that you cannot stop everything, you will not be able to stop a determined adversary. So then you need to focus on detection — try to catch them as early as possible. The moment you catch that someone else is using your mechanisms for MDM, or someone else is logging in with your account from a strange IP address — or even, remember what Amazon was doing when they spotted that key delay which indicated that whoever was on the keyboard was not actually in the United States but was likely somewhere else. Methods like those assume that you will be breached and try to set traps and try to catch the intrusion as early as possible and try to stop it before it turns destructive.

Ryan Naraine (27:15)

So: assume breach and set traps. And these identity-based attacks — these login attacks — aren't going to go away because these usernames and passwords are already out there, they've been shared out there.

JAGS (27:38)

"Identity-based attacks" is the sweetest fucking way of trying to tell people that your password's fucking out. We're trying to make this sound dignified and okay. It's fine, right? We don't need to make it shameful. But it's like, "Yo, someone stole your password," or "Somebody tricked somebody in your call center to give them admin access." So we say "identity," but really what we're saying is the classic problem of: we have a permission-based architecture that's based on you being you, and somebody has your social security number and your password, and we're fucked.

I'm glad — I want to point out that this episode is not sponsored by Thinkst Canary, but we just fucking love the tech and the fact that it makes sense. And it's funny, dude — why were we better about these things 20 years ago? Look at Linux back in the day: you had Tripwire, which was a really common thing for Linux distros to have. And the idea was, "Well, I can't make sure that nobody's going to get in here. So if somebody touches this file, if somebody tries to open this drive, if somebody tries to log into this account that nobody should use, then just pop smoke, just pop an alarm."

And that's the idea with these canary tokens, and it's just so fucking sound. It's the same way that when you do actual opsec trainings and stuff in proper intel, right? You're not trying to stop anybody from doing anything. You just want to know that it happened. Like Moscow Rules: don't look behind you, don't look around, you are being observed. Just assume that you're being observed and you just want to know where people have been and what they've done.

And that's really all we've got at this point with "identity-based attacks." That's the only level of defense that anybody's actually playing with in most of these enterprises — username, password, maybe Okta, maybe a 2FA code. And once that's done, how much segmentation is actually there?

How much proper segmentation of responsibilities by accounts, networks, et cetera? I don't know, guys — how much do you think is actually there? Because what I see is mostly nothing.

Ryan Naraine (30:07)

There are some technical controls to consider though, and I want to give a shout out to Palo Alto. Palo Alto put out a report warning about increased risk of these wiper attacks in this environment. One of the things they pointed out was abusing this Intune admin access to push remote wipe across organizations.

And the big recommendation is things like eliminating standing admin privileges and moving to just-in-time access. Why isn't that the default versus "We have to go figure that out"? And then the other thing on the mass wiping event — Palo Alto mentioned this need for automated alerts on mass wipe events. If more than five or ten devices start getting wiped in a short window, the system should auto-lock the admin. Those are like — I don't understand why that isn't the default.

JAGS (30:50)

But that's touching on what I was saying earlier about — I put this on the manufacturers, the actual maintainers of these codebases. And look, let's take that problem, let's put that responsibility a little bit closer to home. This is a problem of PMs — of product managers. There is this school of product management that thinks that user friction is a sin, no matter what. And it tends to constantly preach this idea of reducing user friction. Don't get in the way, you need to make everything easy. Product adoption and continued use is about how smooth and easy and "everything just works." And you forget that there are things that should not be easy and should not just work. You want them to work when you need them to work, but it should not be fucking easy to wipe your entire organization yourself.

You can make that as a mistake. Disgruntled admin. You fire an IT guy who still has creds. This shit happens all the time. People get prosecuted for this. But the technology didn't stop the admin from deciding to wipe the whole fucking network themselves. So there's an aspect here where it really is a product thing.

Ryan Naraine (32:12)

And I don't know — why can't these common sense ideas become the default thinking? You think it's because of that "security cannot get in the way, can't add friction, we need to make things as usable as possible"? That's the tradeoff we're fighting.

JAGS (32:23)

I think that's the part of the system that's consistently producing these results. But we shouldn't put all the negative responsibility on the product management people. We should also look in the mirror and say, clearly there's a failure of representation and proper advocacy from the security side.

You have a very powerful product management organization in most companies. They tend to shape how a product looks on the way out. They get to decide where a lot of priorities go. Yes, it's a mistake of theirs to apply the same measure they do for product features towards security things. At the same time, clearly it's an organizational failure that's a little too common that there is no one from the security organization that has a true seat at the table when it comes to product features, product deployment, product workflows, design.

And in that case, yeah, we can look at the PMs and be like, "You idiots, you should know better." And it's like, yeah, should they? No. They're not going to be security people. But the real problem

is: maybe you should have a real engagement with someone in the security organization who is a product manager involved in product management. Because I feel like any security person worth their salt would look at the ability to push an executable to every device in a fleet, or the ability to wipe any device remotely, and say, "Hold up." You don't have to be a genius. You don't have to be a 10x reverser to realize that that is a nuclear feature that you need to be careful of.

So yeah, I want to go a little — it's the systems thinker in me, it's the cybernetics obsessive in me that wants to take this a step further and say: look, this might be a problem we fix point-in-time. We go yell at Microsoft, we say, "Hey, you need to fix this thing." And we bitch enough and they go, "Okay, fine, we added a checkbox." And you go, "Okay, thank you." But the actual systematic problem here is: where are your honest-to-God involved security people whose job it is to look at the design of a workflow, a set of features of a security product, and who has the power to say, "No — we cannot ship this this way. We need to come up with something else. And if users don't like it, so be it, because we have some semblance of responsibility for how this feature is inevitably going to be misused."

Ryan Naraine (35:16)

You sound like Satya in a Secure Future Initiative when he said, "If you have to think about security versus something else, do security."

JAGS (35:23)

I wrote that speech for Satya. He paid me. They paid me under the table.

Ryan Naraine (35:29)

Costin, who should we turn to for IOCs? Does Microsoft have a role here? Who should we turn to? Who should be able to show what a TTP looks like as they're heading to this Intune admin interface? What should IOCs look like beyond samples? And who should we ask?

Costin Raiu (35:50)

I think the intrusion most likely has some other IOCs besides the remote wipe command. So to get to that point, for sure they used things like VPNs, things like tools, perhaps remote management tools. And if you ask me who we should turn to for IOCs, you know my answer — CISA. I think this is a moment that they need to step up. Wait — the moment they needed to step up was like a month ago.

Ryan Naraine (36:20)

Wasn't that during the COVID thing? No, it wasn't. Yeah, it was.

Costin Raiu (36:23)

Yeah, correct. Anyways, it would be nice if they can step up and provide information. At this point it's obviously turning into Handala or MOIS or Iranian hackers versus the US and Israel. And I would assume that the security organizations from these two countries — the responsible ones like CISA — need to step up and publish. Start publishing IOCs, start publishing advice, start publishing advisories. The good YARA rules — you know, the good ones that they've been publishing recently, not the previous ones, but the new good ones. IOCs — relevant IOCs. Not random VPN IP addresses, but things that matter, things that have context and things that can be used in practice.

And probably — I mean, we are laughing a lot at the Europeans, right? But every now and then they have some interesting ideas that could be reapplied. Some things such as NIS2, which is kind of a big catalog of critical infrastructure organizations and supply chain for those organizations — kind of gathering them around yourself, around your security agency, so that you can very quickly share information with them and test the readiness and see what the situation looks like whenever something bad or big happens. I think some of that could be applied in the States.

Ryan Naraine (38:03)

Let me ask Mr. JAGS over here. Is there any sharing happening in the background?

JAGS (38:08)

Yeah, yeah, yeah. Well, yeah — no, there was a CISA advisory on some of this IRGC stuff, and I actually got pulled into a White House briefing like two days ago or so. They were…

Ryan Naraine (38:24)

Woo! Look at TBP getting White House briefings!

JAGS (38:30)

Okay, guys. No, it was cool. It was nice to see sort of ONCD and these folks kind of spring back into action. There was, you know, FBI cyber lead spoke, CISA lead spoke. Did they provide anything ultra-actionable? Not really. It was just like, you know, "Hey, look, here's the past two years' worth of advisories that we've been publishing on the IRGC and Iranian-aligned groups." And, you know, "So you should be mindful."

Ryan Naraine (38:58)

Shields Up — is this like the new version of Shields Up?

JAGS (39:00)

I mean, I don't know, but essentially in spirit it was equivalent, except that we didn't have a two-hour Q&A session of commiserating and everybody holding hands singing Kumbaya and hoping that means that people are not going to attack you.

Look, I actually don't know where to put the responsibility on some of these things. I can think of initiatives, right? Like what Costin mentioned about the Europeans. Frankly, with identity-based attacks — this nice euphemism — it's such low-hanging fruit. I think you have two things. There's an obvious initiative, which is: how about the government buys — pays for SpyCloud, pays for VirusTotal, pays for all these different really interesting places that are full of creds and have access to dumps very early on — and they monitor for what those fresh creds are, and they go, "Shit, here's a medical manufacturer. We're going to reach out to them and make sure that they have rotated these creds."

And then you don't have to hope that every fucking dumb two-bit 50-person company on earth is constantly rotating their usernames and passwords, and constantly not getting hit by phishing, and constantly somehow configuring their technology better than the manufacturer even intends them to configure it, and somehow deploying canary tokens, and they just so happen to have an IT person that got them all hardware 2FA. What we're asking for is impossible to actually meet under the current circumstances.

And I think the other issue — the reason that the targeting is so wide and basically intractable for us — is that journalists — this is where I've got to look at you guys. I love all of you, and especially

anybody we quoted with this story, like you're doing a great job and thank you. But everybody's been dying to see some kind of Iranian operation. Everyone's been dying for an example to be able to point and say, "Look, the Iranian cyber power is here to retaliate against us for our sins of having dared to go against the mighty Iranian Republic." And A, we haven't seen it. B, what we are seeing is fucking lame.

Ryan Naraine (41:34)

Well, if you're Stryker, it's not that lame, no? How lame is it that…

JAGS (41:48)

That is — I don't know, you read the fucking story that Brian Krebs put up? The demand thing from Handala, and it's like, "Stryker, a central Zionist company in the Epstein circle of blah, blah, blah." And you're like, "Bro, it's a random medical manufacturer." Good on you that you figured out how to fucking wipe this company. And all we're going to talk about is the fact that at some point maybe they had a government contract, at some point maybe they did whatever. What they are affecting is a bunch of fucking hospitals and people that actually need medical services. And it's not the drone company. It's a company nobody knew existed or mattered in any fucking way two weeks ago. And now we're going — it is fucking lame.

Ryan Naraine (42:34)

That doesn't mean it's lame though. How does that make it lame when there's an entire company with 200,000 servers compromised, 50 terabytes of data stolen, 50,000 employees affected — and let me tell you, recovery is going to be a bitch.

JAGS (42:44)

Yes, it's all going to suck. But guys, we just bombed the absolute fuck out of their whole government organization, their whole country. They killed the Supreme Leader, and your retaliation — your grand strategy, your amazing access into everything and critical infrastructure capabilities and "My god, they're going to melt us from the sky" cyber power — is to take out a 50,000-person company that has nothing to do with this.

And I think that needs to be said that way, because they're lame as fuck. They might have had other capabilities. Handala isn't it. Handala is just some two-bit fucking group that's been a nuisance ever since the Hamas-Israel war. And all they did was fuck around with a couple of random companies here and there. And half the time they were lying about what they were able to accomplish. And we don't mention any of that. We just go, "Oh my God, the great Iranian cyber power has just dinged us. They have brought our comeuppance. This is the find-out phase of our fucking around." And it's like — no, it's just some lame fucking group of assholes who just got creds to a company, which they probably — they didn't even probably know it was the wrong company, because it wasn't the drone manufacturer. It wasn't anybody involved with the war. And then they wiped it.

And if this was a ransomware group, we would be having all kinds of conversations about the fact that they targeted a medical organization and a medical support company. And instead, we're sitting here going, "Oh my God, the Iranians, they did us dirty, how the leopards have eaten our faces." It's like — fuck off, come on, guys.

Ryan Naraine (44:47)

What do you need to see that makes them not lame? See, this is what I'm trying to figure out from you.

Costin Raiu (44:55)

Iran in general, yeah. How is Iran in general going to look like not lame?

JAGS (44:59)

Look, we have spent 20 years cowering, thinking about how we should carefully negotiate with these nut jobs in the religious clerics in the Iranian regime, because God forbid we insult them a little too much and they may just end the world. And you go from that to — no, they wiped a company. If they had done ransomware instead of a wiper, we'd be like, "No, ransomware operation." And all I'm saying is: proportionality in this case entails they don't really have anything.

I'm saying the dams have not been compromised and opened. The oil has not stopped pumping. The gas has not stopped pumping to the entirety of the East Coast. The power plants have not all shut off overnight. Israel has not collapsed in its cyber infrastructure. And maybe that's because they can't do it. I'm not saying it's not doable. I'm saying that the Iranians don't seem to have that capability.

These idiots are going around shooting missiles at every one of their neighbors. Genius fucking strategy. Well done, sir. You're definitely going to get some support this way. And now they're just misfiring in idiotic ways online. And look, I'm not saying Stryker doesn't matter. Of course it matters. Of course we care about how this company gets back in order and how we try to support other companies. But I'm saying in the grand scheme of the war and Iranian cyber — this is when you start writing your stupid op-eds in the New York Times saying, "Holy shit, maybe the Iranians sucked at cyber too." And you're like, "Yep, yeah, I guess so."

Ryan Naraine (47:03)

Costin, do you share his thinking on capabilities?

Costin Raiu (47:05)

I wanted to ask — very fast, very quick, don't think about it — what is your advice for the sysadmins at Stryker?

JAGS (47:14)

I mean, I hope you have backups. Buy a bunch of Macs.

Ryan Naraine (47:20)

Well, let's talk about that. What's the path to recovery?

JAGS (47:25)

Well, it's not recovery at this point. What you're talking about is re-architecting. And I'll say this — I don't know Stryker as a company and I don't know anything about them because I didn't know they existed a week ago. And I think none of you did. Let's all be honest. You didn't know they existed. Just live with that. What was it — a Michigan-based company with a large presence in Cork?

Ryan Naraine (47:49)

50,000 employees. Fair enough.

JAGS (47:55)

Okay, well, I thank you all for your great work, but I'm just saying that nobody ever thought about this fucking company before. And that's okay.

Ryan Naraine (47:56)

Ireland. I know. The first time Costin shared a link about this story with me earlier this week, I was like, "Fucking Ireland. What did the Irish do?"

JAGS (48:07)

Leave Ireland alone! They were neutral this whole time. Look, my advice there frankly is: I'm assuming that these admins could not have gotten sizable budget to do anything three weeks ago, and now they're in a prime opportunity space to re-architect their entire infrastructure with probably a great deal of support and money that they wouldn't otherwise have. Hopefully.

So if that's the case, now is when you go: "What are these people? There's 200,000 people. How many of these people actually need full-fledged computers that can do everything? Half of them." Okay, well, the other half are getting Chromebooks. Welcome to the Chromebook life. All of you have Chromebooks now. We're just going to manage everything remotely through G Suite. The other half — what do you people do? What do you need this for? What kind of development, what kind of applications? Can that shit run on a Mac? Because if it can run on a Mac, let's get a bunch of Macs over here. And then this is going to be our new immutable backup solution, this is going to be our EDR that we deploy on the Macs, and how we're going to manage the stuff on the Chromebooks.

And you basically create the simplest version of how you're going to be able to manage the majority of the systems in the least complex way. And then, once you have that — that's like a fucking dream, right? Think about Google. Why is Google where it is now? Aurora did a lot for Google, because the day after Aurora got figured out, they literally took everyone's computers away and gave them all Linux boxes and re-architected how all of those systems worked. And all of a sudden, everyone was on the same kind of machine, everyone was using the same type of auth, everybody was using X, Y, and Z. So you can go and solve categories of problems.

I think that baked really well into the DNA of security at Google, to the point where when you look at Project Zero — which is not an internal security org — their thinking is also: we are going to go solve categories of problems in vulnerability and exploitation, because that's the only thing that matters. Otherwise, it's stunt hacking. Being able to come out and say, "I have an iOS chain that executes in this way and the other, and this is how you fix it on one device and no other device" — it's stunt hacking. You go, "Okay, cool, thank you for showing us this really cool thing." But security is: "Okay, how do I fix that problem for every fucking iPhone, or as many of the iPhones on earth?"

So this is the microcosmic version of that. Today might be the most painfully stressful day, but maybe the coolest day for the Stryker sysadmins. Because you could sit back and go, "Okay, this is the only time we're going to get these motherfuckers to pay for re-architecting this network. Whatever we buy today is probably going to be the same thing that's going to be here until the next time people remember that we're here, which is probably in 20 years." So now's the time — build the coolest, simplest — never waste a good crisis. In the words of DB: never waste a good crisis. A crisis is an opportunity.

Costin Raiu (51:39)

Small observation: deploying Chromebooks everywhere may be a case of the cure being worse than the disease, in case you used Chromebooks before.

Ryan Naraine (51:48)

Monoculture.

JAGS (51:51)

I mean, I'm not a fan of Chromebooks, but my point is: how many people actually need full-fledged machines?

Ryan Naraine (51:58)

Yeah, thin clients. We need to go back to the old thin clients that the banks in New York were using in the 2000s. There was one server and everyone just got a display.

JAGS (52:05)

Think about it this way. With the Chromebooks — there's a single point of failure in that you have this remote management for all the Chromebooks and you could probably wipe them or lock them that way. However, because they're thin clients, all that information is being stored in whatever way you see fit remotely. So then yes, they can wipe the Chromebooks and who gives a fuck, because we have everyone's meaningful data in some G Suite managed storage mechanism that we can half offload to Google, the safest organization on earth, to say, "Yo, you better keep copies of that fragmented across the biggest computer on the planet, and I'll let you know when I need that shit." Yes, you're going to pay for it, but frankly, it's cheaper than whatever the fuck's going on right now.

Ryan Naraine (52:54)

Costin, why is that solution worse than the problem?

Costin Raiu (52:59)

Assuming, of course, that the only things that Stryker's 50,000 employees do is Google Docs and Google Sheets and nothing else.

JAGS (53:07)

And that was my point, right? I don't think that's ever the full solution, and I said it. Half of them, maybe, right? If half of them are your Salesforce and forward-deployed, customer support people, your management, your HR folks — Chromebooks. You can all get Chromebooks, enjoy. And then the other half, you have to start segmenting. You basically break up the problem into smaller and smaller chunks that you can manage.

But the truth is, all of these companies that are just — everyone has a Lenovo ThinkPad. 200,000 people have Lenovo ThinkPads, whether you are a person answering the phones at the front desk or you're a critical engineer in the back room. I don't care what account management you're doing. You are giving everyone way more computing power, capability, and access than they need for their jobs. And it makes your job much, much harder to manage.

Costin Raiu (54:14)

Still, let's not forget one thing. Looking at the Iran state-sponsored cyber threat advisories from CISA, the most recent one is from June 30, 2025, when they say, "Iranian cyber actors may target vulnerable US networks." I think there's justification for another advisory. 2026: "Iranian cyber actors are actively targeting."

Ryan Naraine (54:48)

Update the date and the word.

JAGS (54:52)

I hate this word "may." And in any intelligence — I think we've talked a lot about intelligence assessments and language on this podcast. Personally, I think the "may" coupled with the passive voice is the worst sickness of all of the writing that anybody does in our threat intel space. Because note what you're not saying: who thinks this is going to happen? The passive voice there — you could say CISA analysts believe, assess, consider, suspect that Iranian forces may do X, Y, and Z. And you go, "Well, may they do this more than anything else they may do? They may not." And how likely are they? "May" talks a lot about "there's a possibility." And you go, "Sure, there is a possibility, and there's an interest from these people." And you go, "Yeah, I'm sure that yes, if they think about it, they would be interested in doing so." And you go, "Well, can they?" And they go, "I mean, in the grand scheme of things, they can, the same way that we can — we could perhaps."

Costin Raiu (56:04)

One hour on the Handala story. One hour already.

JAGS (56:10)

It's more than they deserve for any of their other fucking bullshit attacks. Fuck that.

Dutch Intelligence Advisory: Signal & WhatsApp Targeting

Ryan Naraine (56:11)

Which is a good way to pivot off to another report I wanted to flag, which was our Dutch intelligence friends — MIVD, AIVD — publishing this advisory: Russian state actors running large-scale global campaigns to compromise Signal and WhatsApp accounts of government officials, military personnel, diplomats, likely journalists, and so on. No technical vulnerabilities here. This is all social engineering.

JAGS (56:36)

They are massively…

Ryan Naraine (56:40)

They're actively compromising using two social engineering tricks. One is posing as Signal security support — like a chatbot — to trick people into handing over verification codes and PIN numbers. And then device linking, which we had seen previously from the Google warning in 2025 — I want to say shout out to Dan Black — for that earlier thing where they were linking accounts and you had no way of knowing that a third party was seeing things. Sending malicious QR codes with these group chat invitations and so on.

This stands out. Signal actually published a long response thread on Bluesky. I'll link to it in the show description. But this is essentially the same class of attacks we saw being targeted, apparently going to mass scale here. Costin, have you seen any of this stuff — any of the phishing messages, anything that you can add to what we're getting out of AIVD/MIVD?

Costin Raiu (57:34)

I have not seen any targeting myself, but I guess all of that is very precise and going after the really important people. It's not just anyone like myself or whoever, but the really important people would be seeing them. And I think that every time this happens — every time there's some kind of targeting or discussion around Signal — it's very interesting to note that there's a lot of people

immediately pointing out that Signal is crap, that the encryption of Signal has been broken, and pointing to blogs where they talk about the fact that the SQLite message database that Signal uses is not encrypted, the encryption key is hardcoded, and things like that.

There's for sure a problem of perception that Signal has been trying to address over and over again. I think even Elon Musk was saying that Signal is not secure, don't use Signal, use X messaging. Maybe this is why they got upset and they posted it on Bluesky.

Ryan Naraine (58:49)

There's a long thread on Bluesky basically just kind of…

Costin Raiu (58:54)

Trying to explain, I guess — yeah, nobody broke the encryption, everything is fine. And again, it's a technology which, to be honest, and we know this from our groups, there are a lot of confusing things about Signal, such as nicknames. Sometimes you have in a group people with the same nickname and you don't know who is who. Or sometimes people change their nicknames in a very confusing manner — they change their nickname to somebody else's. And then you start wondering who is who. I think people should admit this is a real problem.

Ryan Naraine (59:33)

This is a real problem. When people change their names to just "H" and now there are four H's in this group chat.

Costin Raiu (59:39)

Shout out to our friend H.

JAGS (59:50)

Bro, all of our friends are aliases. But okay, let's talk about this. This is where I disagree with Signal, and we've been complaining about this for a while, and I haven't heard any kind of meaningful engagement from anybody at Signal. So forgive me if they're engaging somewhere I haven't seen.

Look, we all stan Signal. We all use Signal a lot. The move with the usernames and decoupling things from phone numbers — and decoupling them from usernames where the usernames themselves are private — was asking for these quote-unquote "identity-based problems" to become a thing. This was not an issue with Signal two years ago.

Ryan Naraine (01:00:30)

Phone number was a problem though. Relying on people's phone number was an inherent design problem.

JAGS (01:00:33)

Phone number is an inherent problem when it comes to a certain amount of privacy and management of that — being able to attack somebody because you have the number. Yeah. But that's not the only possible problem. And what they ultimately did — it solved a theoretical problem and replaced it with a real problem. Which is that we are decoupled completely from any kind of identifiers. Even if we're friends, I look at Costin's thing and I can't see his username. So I cannot in any meaningful way verify who I'm talking to.

We have group chats where there's four people with a one-letter name and they can change that name and they can change their picture. And the last time I bitched about this on this podcast, I

changed my name to “Ryan Naraine” and put Ryan’s picture and nobody knew the difference between either of us. That’s a real and simple problem that’s affecting everyone in your user base. And you introduced it to fix a problem that exists but was not subject to, I think, a similar level of abuse.

Ryan Naraine (01:02:07)

What is the ask? You have Meredith and you say, “Meredith, I would love for you to fix this problem.” What is the ask?

JAGS (01:02:13)

They need to ground Signal identities on a visible marker. If I have your phone number, then that account is connected to your phone number. If I have your username, then that account is connected to that username and I can see that username. The idea of having these completely ephemeral instances that are just distinguishable by display — and I don’t give a fuck what you tell me about there being a verification number and a QR code. That’s a theoretical solution for the minimum 0.1% of nerds to get together and scan each other’s barcodes.

But the reality is most people are afflicted by this right now. You cannot tell who’s who. Which is why all of a sudden you start to see all these attacks that are based around linked devices or being able to just have a vaguely named account that sits in a room with 30 other people. That shit needs to be solved. It is a genuine weakness in Signal, and it’s the first time that we’ve looked at Signal as an unfavorable mechanism for the InfoSec space. They can solve it, they should solve it. It also sucks to not be able to share contacts. How the fuck do I connect you with somebody?

Ryan Naraine (01:03:29)

Costin, how often are you scanning QR codes for Signal contacts? You go to a conference — it’s rarely, right?

Costin Raiu (01:03:35)

Very rarely. We used to do that a lot with Threema, by the way, if you remember. Because in Threema, unless you do that, you are not verified, therefore the contact doesn’t have the full three green dots, which is super annoying.

Ryan Naraine (01:03:55)

Does anyone still use Threema?

Costin Raiu (01:04:05)

I know people who still use Threema. But one of the big — two problems in my opinion with Threema, which by the way is a great application, I’m happy they’re still around. I think they just sold the company to someone else, if I’m not mistaken.

JAGS (01:04:15)

VPN provider out of the UAE that nobody’s ever heard of before? No, I don’t know. I’m just guessing, but fact-check me, people. That’s usually how this fucking goes, right?

Costin Raiu (01:04:24)

No, no, no. I think either they got some external investment or they sold the company. But again, you can fact-check me. However, I think they still have two problems. One: they don’t have these disappearing messages that even WhatsApp and Telegram have nowadays. They still haven’t

implemented disappearing messages. It's 2026 — I mean, how is it possible they still don't have this feature?

And two: the quality of the calls is worse than the free Signal. And keep in mind that Threema is a paid application — you pay two dollars, two euros, three euros, something like this — and it's a paid application but the quality is worse than Signal, which is free. Which is why most people use Signal. By the way, there's also the possibility — and we were joking about it — that someone is messing with the Threema calls just to get people to use something else and stop using Threema. Because you call someone and the call dies, or the quality is bad, or it doesn't connect, or you try to dial them five times before you get through. And Signal — it just works. The first moment you call someone, it just works immediately. With Threema, there are all these kinds of issues. Which is a pity, in my opinion. And disappearing messages as well.

Ryan Naraine (01:05:54)

Is there a Signal best-case use-case document? "Never send classified information," regularly — is there a way to regularly check your device for this linked device thing? And we talk about Signal here, but there's also a lot of WhatsApp targeting going on along the same lines. And the linked devices on WhatsApp is also the same problem — it's kind of having the same WhatsApp experience on the browser, which is web.whatsapp.com. You can link the devices there and have both. What is the — how do people…

Costin Raiu (01:06:25)

I think Signal has really grown a lot. They've grown to the point where they probably need to start thinking about a threat intelligence team, some more sophisticated internal security team. Maybe they should start thinking about trying to identify rogue connected linked devices. Trying to identify that maybe both of your devices are in the US and then there's a third one that just appeared — it's connected by a US VPN, and it's not the only one, but sometimes it's connected by a VPN and sometimes it's just beaming out of Russia or UAE or whatever.

I think they've grown large enough to the point where they need to approach the issue of security in a more serious manner. You can't deal with it entirely from a cryptographical point of view and say, "This is our application, it's cryptographically sound, everything is secure. If you link a device or if you fall victim to phishing, it's not a vulnerability, it's not a problem, it's your problem." They've grown enough — it's a mature app now. The US government was using it — at least that version, the sabotaged version with logging that was made by an Israeli company. I think they are a mature solution now and they need to pay attention a bit more to things like this and deal with them, instead of just denying that these problems exist.

Ryan Naraine (01:08:06)

It calls for re-architecting the way identities are managed, the way identities are displayed, the way linked accounts are done — something new.

Costin Raiu (01:08:14)

Possibly. The other thing here is — I know Signal kind of survives out of donations. I donated myself a number of times and you get that tiny rocket logo. So if you want that cool logo on your Signal profile, donate and you'll get it. It looks really cool in my opinion. And probably they may have reached the point where they can no longer survive just with donations. So maybe they need to consider some other alternative schemes — like commercial schemes — that could support the whole infrastructure, the whole ecosystem for the next 20 years.

APT-28 / Fancy Bear: ESET Research on Sofacy Return

Ryan Naraine (01:08:57)

Staying in Russia, we got a new report from ESET with the return of APT-28 — Fancy Bear, a.k.a. Sofacy. ESET documented the return of this group's advanced implant development team, which had gone quiet since around 2019 — a long time ago — when the group shifted to mostly phishing operations. They're saying here the primary targets are Ukrainian military personnel in long-term surveillance operations. Did you get a chance to look at this custom Fancy Bear? APT-28 is one of your favorites.

Costin Raiu (01:09:29)

I did. This is not only my favorite — it's actually, I think, the first APT that I ever studied in detail. I remember back in the days, as we were starting to look into the issue of advanced persistent threats, sophisticated adversaries, there were three blog posts from a French researcher about Sofacy, about their implant that was written in assembler. Back in the days, it was really tiny — in the range of six to eight kilobytes of code, maybe a 10–12 kilobyte binary. Nowadays we see megabytes, 10-megabyte binaries. But back in the day, Sofacy had this tiny implant written in assembler that was just a few kilobytes in size.

And yeah, that was fascinating. It was pretty much the first thing that I looked into in detail, and then I started tracking — like the dog with a bone, you know — for the next decade, more than a decade, 15 years. And they're absolutely right. I thought that this research from ESET is fascinating because they saw the same phenomenon, and I think a lot of researchers have seen it as well. Everyone was wondering what happened with Sofacy. Remember, back in 2014–2015, they had zero-days — five zero-days in one year. They just burned through five zero-days in one year. Silverlight, Windows EOPs, all sorts of things. And then suddenly they just collapsed, or went behind the curtain, or just stopped all these amazing exploitation and implants and switched exclusively to phishing.

This is what I saw as well, around — I don't know if it was 2016, 2017, or 2018, but somewhere around there — their operations wound down and they switched to phishing. And actually, their phishing was insanely successful. This is what we need to say here, and this may also be one of the explanations why they switched entirely to phishing — because with phishing, it was working very, very well. They were just harvesting credentials, and then they had a pretty sophisticated system that monitored thousands of email accounts, copied all those emails nicely, got them indexed, got them scanned for keywords. All that system, in my opinion, was incredibly powerful, incredibly successful.

Just for historical reasons, back in 2008 or maybe even before 2008, Sofacy was a particularly interesting entity because one of their main goals was related to nuclear weapons. The group, the military unit that Sofacy was part of back then — their goal was mostly to understand when to make the decision: do we launch our ICBMs with thermonuclear warheads, or do we not? Getting that understanding, getting that awareness, was the primary goal. And of course that may have changed over the years, but this is pretty much where they started.

Over the years they evolved — kind of the counterpart to Turla — and evolved in a different direction, in my opinion. And I think that had a lot to do with their military budget being massively boosted somewhere around 2013. I think they got over $160 million for cyber operations, and I

think that may have been visible in all the zero-days they were deploying. And then maybe they just ran out of budget, switched to phishing — insanely successful, the phishing part.

And there was another very interesting Sofacy side of things that people didn't discuss much: their compromising of routers, notably Cisco routers, which they've been doing for a very long time, more than 15 years, perhaps close to 20 years now. Compromising Cisco routers and backdooring them in RAM, through all these mechanisms that are built into the router such as redirecting the traffic, building VPN connections, tunnels in particular — tunnels that redirect all the traffic to their servers. I've even seen cases where they were reflashing routers with malicious implants or malicious firmware that had hardcoded backdoors.

And it's interesting that they did stop all that crazy stuff for a while. And it's also very interesting to see this research from ESET. I was so happy when I saw it because they do point out what happened — where did these developers go? Because they had the developers writing X-Agent, X-Tunnel — which are the two most famous tools from Sofacy — what did they do for all these years? Were they just phishing, or were they on vacation?

Ryan Naraine (01:15:21)

Or did you guys suck at finding them? Is it possible that they just got more careful about exposure and just never got caught?

Costin Raiu (01:15:27)

I think so — no, it just felt that they switched to phishing. Like a philosophical change of mind: "We don't do implants, maybe now all this implant business is the other guys" — you know which other guys I'm talking about — "and we're going to refocus our efforts for phishing." And I guess at some point they realized this is not enough, so let's do a bit of implanting again, and rebuild some of our tools, improve them, improve X-Agent, improve X-Tunnel, and redeploy them in Ukraine.

Ryan Naraine (01:16:04)

ESET shipped IOCs — is there anything useful to help people go hunting?

Costin Raiu (01:16:07)

ESET shipped a bunch of IOCs. Yeah, they do have two hashes in the blog and they have a lot more on their GitHub. But if I remember correctly, most of the hashes on their GitHub are for the old stuff. If we're talking exclusively about the new stuff in the ESET blog — they do discuss some of these new tools, and I'm not entirely sure, maybe I failed to grasp what they're saying there, but it felt that they do have these two hashes, but they're also discussing some newer tools. It's not clear to me if they share the hashes of those newer tools that look similar to these older tools. So it may very well be that they're sharing some hashes, but at the same time they have a lot more for their private intelligence customers or their private reporting service.

Ryan Naraine (01:17:06)

Juanito, did you get a chance to look at it at all?

JAGS (01:17:08)

Yeah, I did actually. I'm glad I did, because frankly it was a good report, but it was a good report in an old-school way that I've missed a lot. They figured out the code similarity, they're moving backwards and forwards in a way that I really love. It's like, "We see this implant from 2024, and when we look at that, we realized that there were samples going all the way back to 20-whatever

that we hadn't realized, and that makes us find this other thing that's connected to it." And then you can connect that to this name that many of us are familiar with — remotekeylogger.dll from back in the day. And then when we look at this other component, it's clear that it's a similar codebase, same source code, but they include this calculation and this other thing. And that connects it to this thing.

That's the sort of — it's not even paleontology, because you're still in the same decade — and you can just see how much consistent development there is. I think it also shows something really interesting that a lot of us have been looking out for with APT-28 and kind of demurred on. It was like APT-28 suddenly threw everything out the window and then started to make entirely new components using new technologies, a different approach. And it became really hard — A, really hard to track APT-28, but B, really hard to define APT-28. It was just kind of like there's some different sub-units and maybe this part and that part that work with each other and don't.

And you get into the situation where APT-28 had become a little bit meaningless, other than when somebody from the government or wherever came out with an advisory and said, "APT-28 is using this, that, and the other to do X, Y, and Z hacking." And the only reason we think it's APT-28 is because they call it APT-28, but we have no other way to connect it. So to now see — no, look, the codebase from the original developed stuff for APT-28 is still active, still available, and being maintained. We just hadn't noticed this part, we hadn't seen this other thing.

It tells us that there's more to the mystery that looks the way we would expect it to, and it's just right there — we just hadn't picked up a thread. So this is, I think, really important work by ESET. And I think it kind of invites everybody else to go backwards and go, "What the fuck did we miss? Did APT-28 just look a little bit different and we just didn't fucking catch it? And they've been around same as they always were since 2018, or what?"

Ryan Naraine (01:19:54)

Shout out to ESET. These guys are always pushing out good work.

Costin Raiu (01:19:58)

I think they're doing an amazing job at publishing about these Russian APTs. I remember there used to be a time when there was a constant race between ESET and ourselves — who's going to publish more reports about Russian APTs? And I think around 2020-something, they just started publishing so much. And I think that was also partly because of their visibility in Ukraine, because it was widely being deployed in Ukraine, so they were getting all these goodies from there. And yeah, the amount of Russia-APT research they were publishing just went through the roof, beating almost everyone else out there.

Poland Nuclear Research Center Cyber Attack

Ryan Naraine (01:20:47)

What do you guys make of Poland? Polish ministers confirming that Poland identified an attempted cyber attack on the servers of the National Center for Nuclear Research. They said the attack was thwarted — in a statement, all safety systems operated according to procedure. We think this is real, or we think this is just someone scanning something? Because there's no IOCs, there's no details, just kind of — and by the way, we suspect it might be Iranian. What do you guys make of that story?

JAGS (01:21:23)

Well, look, there's no details, but also there is an attempt at nuance, right? Like even the story itself is kind of like, "They say Iranian, it could be Iranian, but it could not be Iranian." And the fact that we don't have details — to me, this could very well be a DDoS thing or whatever. At the same time, I refer you to our previous conversation about Handala, where Iranian groups are loath to do really stupid shit left, right, and center. And it's not beyond them to try to do some attention-grabby bullshit at a completely random target in a totally unrelated nation.

To bring in maybe the strongly worded concerns of a bunch of people that have just sidelined this conflict completely. But it could very well just be a fucking DDoS from people that Handala have used before — you know, if Anonymous Sudan were still around or whatever, or somebody else just kind of pretending. So I appreciate the heads up. I also appreciate the nuance of it. And I think it's important to keep it as tamped down as possible, until told otherwise.

Because as far as Poland goes, I'm still hung up on the part where the Russians very really actually tried to shut down the power in late December. And I don't super care if the Iranians kind of sort of touched a nuclear research center the wrong way. That is a bullshit thing. That's a maybe and has no real effects, versus our very real thing with a very clear attribution and very clear ill intent at a time of not-war against a civilian population that we all seem to just kind of be like, "Well, okay, no problem, keep going."

Costin Raiu (01:23:28)

And by the way, I think there's an interesting angle here that maybe people missed. However, us in Europe have seen as a big thing. It was just at the beginning of March that Poland announced that they are willing and interested to take the steps to develop their own nuclear weapons, which is something that hasn't happened in a very long time. The amount of nuclear powers who have nuclear weapons has been pretty much constant since the '90s. Maybe the last one was Pakistan, I think, probably.

And so the fact that Poland wants to acquire or develop their own nuclear weapons — I thought that was a huge story. And the Prime Minister, Donald Tusk, also said that he was in discussions with France to get nuclear-capable jet fighters to be stationed on Polish territory. So I wonder if this is one of the reasons for this attack. I don't know if it's proper to call it an attack, because the language is more along the lines that there was something that was stopped before it had any chance to do anything. So maybe that's not necessarily an attack, depending on how you define it.

Ryan Naraine (01:24:52)

If I send a phishing message to you and your IT guy flags a phishing message that would have led to some sort of identity-based pivoting, is that an attempted cyber attack? And I'm curious about it because I see a lot of companies will ship a threat intel report on some targeting they're observing, and it's based on two emails they saw being sent to very specific people, which are targeting.

Costin Raiu (01:25:05)

I fully understand the perspective and the optics. It's like, let's say, Iran is firing a missile towards some country but it falls into the sea. Is it a missile attack? Or just because it fell into the sea, it's not a real attack? I guess it depends on the impact, how serious that attack is.

In this particular case, the Minister for Digital Affairs — and I like how you said in our document, "the Polish digital minister." And I was wondering, is that a real person or is it some kind of an AI, a digital AI which is a minister as well?

JAGS (01:26:19)

You're being a little speciesist. Don't be so speciesist, Ryan. Just because it's a synthetic being does not mean that it's not a real being, okay? It's a little holophobic is all I'm saying.

Costin Raiu (01:26:33)

Of course, of course. The Doctor in Star Trek Voyager is a real person, of course. Photons Be Free.

JAGS (01:26:39)

According to Anthropic, it has consciousness. Look, you had the same thing in Westworld, right? When you look at the board meeting of Delos and there's an AI in charge of the decision-making. Look, all I'll say is this — that's all sci-fi whatever, you know, as Ryan would say, hooey-phooey. But at this point, I will say: to what extent would we actually have better judgment in government and different policy positions if these people were actually consulting some decently well-built, actually purpose-driven, high-end reasoning model? I think that's meaningful.

Not just "Can they go talk to Copilot?" No, we're better off if they don't talk to Copilot. But if you could actually do research briefs and summaries and things, they might actually read some of the bills that they're signing. That's all I'm saying.

Costin Raiu (01:27:40)

One thing we shouldn't forget to mention: remember we had a very similar discussion about this cyber attack against the Polish power grid in December, and we didn't know if it was real, how serious it was, all these politicians — and then the Polish CERT came with this bomb of a report. That amazing report which showed that it was really serious. So I think the Polish CERT has the potential to publish something really cool about this nuclear cyber attack. And you see — it works! Did you see that? It works.

Apple Patches for Coruna Exploit Kit & Google Samples

Ryan Naraine (01:28:14)

Costin is always asking for goodies. He knows how to ask for goodies. And I was just going to say — you know how to ask for goodies, and it actually works. Because last week, we were all complaining about the absence of documentation from Apple and the absence of samples on Coruna, this exploit kit. Lo and behold, Apple released two patches this week: iOS 15.8.7 and iOS 16.7.15. One patches a kernel bug. The other one patches four bugs — three WebKit bugs and a kernel bug, if I remember correctly.

And Apple's documentation says this update brings that fix to devices that cannot update to the latest iOS version. And they actually document that these were WebKit vulnerabilities associated with the Coruna exploit. So Apple calls out Coruna, ships some patches for older versions of the operating system — iOS that cannot update to the latest version. That part I'm a little lost on, if you can explain that bit to me.

But in addition to that, we got Clément from Google actually sharing samples. So, Costin, not only did you get word from Apple confirming and patching and updating things, but we got samples. Did you get to play with the samples? What do you make of your requests coming to pass?

Costin Raiu (01:29:34)

Of course. First of all, thank you to Clément and to Google and to all the people who approved and actually made it happen. And we did get some really interesting samples. By the way, one of the things that attracted my attention with Coruna was that the YARA rule — Google actually in their blog published two YARA rules, one for the Plasma loader for the implant itself, and another one for these HTML pages which had this obfuscated JavaScript. So it's more like a signature for the obfuscated JavaScript used in the exploits.

And from the other rule they published for the implant itself, actually only three strings were hitting on the samples that were available from the Chinese — all the samples that are now everywhere, all over the internet — only three of those 20-ish strings in the other rule were hitting. And I was wondering, why aren't the other strings hitting? What is going on? And it does appear that somehow Google has a set of different implants — or different modules would be maybe a proper word here — than the ones that are being delivered from this website. So I think that they probably have been watching this for much longer than everyone else and they got their hands on some really interesting ones, together with samples that have the equivalent of PDB paths — the macOS or iOS Mach-O equivalent of PDB paths — that allow you to see some of the internal file names.

So when they compile this thing — and it looks to me that the people who are deploying the kit do have access to the source code and they have access to the whole framework — when they compile it, sometimes you can see where it was sitting and you see the internal names of these modules. And one of them, for instance, "Hellion" — it didn't appear anywhere else. It only appears in these new samples that Google posted. So it's a wonderful capture of samples.

Perhaps the only thing which is not entirely clear to me is if all these samples are coming from the Chinese, or some of those samples that Google posted are actually related to the Russian deployment of Coruna, which took place in the summer.

Ryan Naraine (01:32:07)

Why does that make a difference?

Costin Raiu (01:32:09)

Well, it would make a difference in my opinion because it would indicate perhaps that the Chinese didn't get it from the Russians. And then you start wondering: how did the Chinese get their hands on this thing, if not from the Russians? And if the Russians were the first to acquire it, where did the Chinese get it from? And these subtle differences in the samples that are now in the wild and those samples that Google posted makes you wonder — where did the Chinese get this stuff from?

Ryan Naraine (01:32:41)

And this is why people are always hesitant to post samples, because they are always wondering, "What am I giving up in the background that may not be obvious?"

Costin Raiu (01:32:48)

Oh yeah, and then people start poking at the samples and then they see all these things that you may have not noticed or spotted. That's true. But again, thank you for sharing those samples. I think they should be able to advance research into this topic a lot. People can write additional signatures, improve their existing signatures. I did that — my YARA rules, I improved them based on the samples from Google, with more things, more data.

And I'm still hoping — I'm always asking, I'm always pushing — we need a bit more research, if you ask me, into the Russian deployment of Coruna against Ukrainian websites, because there's almost

no information whatsoever on that. And while at it, if iVerify could also post their samples on VirusTotal — if it's not too much to ask — that would be nice, because their hashes are different from what Google posted. So it looks like they captured something else, something different, which is also different from whatever is in the wild at the moment from the Chinese sites.

Ryan Naraine (01:34:00)

In Apple's documentation, they said the fix is associated with the Coruna exploit and it was shipped in iOS 17 since September 18, 2023. "This update brings that fix to devices that cannot update to the latest iOS version." What makes a device unable to update to the latest iOS version? And what did they tweak here that allows it to?

Costin Raiu (01:34:19)

iPhone X. Yeah, devices like iPhone X. I still have one and I love it. It's one of the most wonderful iPhones Apple ever made. And still, I think that somehow — I don't understand how it's possible — but the camera or the software running on iPhone X produces images which have some kind of soul that the most modern iPhones, with much higher resolution, colors, processing — whatever — they just lack that kind of soul in the image that the iPhone X has. The photographer in me — I love that. I still have that iPhone X and I use it to take photos.

I can share some of them, just to get an idea. They look so different from whatever the latest iPhone produces. I don't know if you know about this trend which is known as Lomography? There are these super old cameras, film cameras called LOMO, with terrible optics, terrible quality. But because of that, they make beautiful images. There's soul in it.

Ryan Naraine (01:35:36)

It's like going into these new modern coffee shops and there's all these old bicycles and rusty nails hanging from the ceiling and you're just like, "What — are we just clinging to the past?"

JAGS (01:35:46)

This is cool. Sorry, I'm Googling the Lomography thing, and that's dope. You've definitely seen images — I just didn't know that was a thing, that was how they were produced.

Ryan Naraine (01:35:54)

But iPhone X cannot patch to new things?

Costin Raiu (01:35:54)

They're stuck. They're stuck on iOS 16, for instance, I think. Maybe iOS 17 beta may have been available at some point, but I think I broke it, so it went back to iOS 16. I don't use it for much except two things: to take photos, and to run Waze whenever I'm driving somewhere.

Ryan Naraine (01:36:21)

Is it a surprise to you that Apple didn't make it available earlier? These patches were not made available for iOS 16 — they're not supported any longer. But Costin, if you have documented evidence that it's being used in these kind of targeted attacks against extremely targeted people — that Apple always uses this language — they know that these people are on iOS 15 and 16. Don't you have a responsibility?

Costin Raiu (01:36:43)

I know people who still have iPhone 8, and I know people who even have iPhone 6, and they're super happy with them. Really happy. It's going to last forever.

JAGS (01:36:58)

Well, there is a difference with the iPhone X part. When you — I believe that's when you get whatever vault, right? There's a certain amount of generational change with the Secure Enclave. I believe that was iPhone 10. And I bring it up because I assumed you had an iPhone 10 because, my understanding is, that's sort of the gold standard for being able to do a certain amount of debugging or reversing of certain components in the iOS world. After that, you don't have the right hardware.

Costin Raiu (01:37:33)

It's the last one supported by the checkm8 exploit for jailbreaking. After that, you can't do it any longer.

Ryan Naraine (01:37:51)

All right, we're out of time.

Costin Raiu (01:37:51)

I mean, I never sell — I never throw away my old iPhones for two reasons. One, it's good to keep them in case you get some new IOCs. Maybe you find something cute that somebody gave you on one of those older iPhones. And it's always good to keep them around in case you want to jailbreak them or play with some old iOS versions or see how they differ from the new ones.

Ryan Naraine (01:38:22)

Are you surprised we haven't seen a Kaspersky report?

Costin Raiu (01:38:26)

I'm very surprised. I don't want to sound like I'm bragging or anything, but I do remember back in the days, whenever a story like this dropped, we would just jump on that story and dig the shit out of it. Search everything you can imagine in all available databases, caches, detections… I hope they are — yeah, I mean, Boris for sure, Georgy…

Ryan Naraine (01:38:49)

You have to imagine Anton and these guys are doing that though.

Costin Raiu (01:38:56)

Probably. I guess it takes longer. I'll tell you — when it takes longer, the only time it takes longer is when you find a lot of shit. So when you find a lot of shit, then you need time to go through it. That's when it takes longer. Because otherwise, you find something very fast, very quick, and then you jump and publish. The bigger the cannonball will be.

JAGS (01:39:12)

Suffering from success.

Ryan Naraine (01:39:17)

So the longer it takes Kaspersky to push something, the more they're finding?

JAGS (01:39:21)

The more they have. What's the meme of, you know, guy just sweating bullets?

Closing & Shout-Outs

Ryan Naraine (01:39:25)

All right. I wanted us to get to this cybersecurity strategy from the White House along with an executive order around fraud, but we're out of time. I promise we're out of time, but we'll get to it. So I've got to run to the airport. Literally, I'm closing the laptop, pushing it in my bag, and heading to the airport. Close the show with some shout-outs. Juanito, you first.

JAGS (01:39:48)

Okay, since we didn't talk about Reverse, which was the second iteration of a really great conference — very cool. Just a different environment, especially if you've gone to REcon, which is historically the reversing conference, but also a lot of drinking and some disorganization and stuff around it. Reverse is a much more nicely put-together con and also a little more wholesome. Everybody was just very nice. I didn't really see many people drinking too much. You could if you wanted to. Folks were just hanging out.

Shout out to Jordan and Peter and all the Vector 35 folks for putting on a really great conference. More importantly, dude, I had never seen — like I've watched Lori Wired's videos before, and she's clearly very smart and has great content. But I watched her talk at Reverse and — dude, well, first of all, obviously very engaging and she knows how to express herself and captivate the audience. That's just a talent that comes with being a producer on your own.

But don't let that fool you. Her talk was basically the best use of a one-hour slot I've seen in a very long time. Usually when I see a one-hour slot, most people — you're like, "Oof, that's a little too much for you." It lets people be lazy and not condense their points. And you have a tendency to misuse the audience's time because you think that you have a lot of it.

Instead, Lori basically walked us through — she wanted to talk about how we could use compilers and compiler abstractions and compiler mechanisms to make reversing easier. Which is a topic that I'm fascinated by. It's the same philosophy around — I tried to grasp at that very poorly when I was doing AlphaGoLang, which are the Go reversing tools for IDA, which are kind of outdated now. And Nicole Fishbein and I worked on something similar with Project Dock Salic, where we were trying to do Rust reversing tools. Most of what we were doing was coming at this same approach of trying to get the compiler abstractions to work in our favor.

But we didn't have — I wish I'd had Lori's talk as a guide years ago to approach that project. Because she went through the universe of everything that goes in the compiler: how the compiler works, how the compiler and the linker and everything fit together, what levels of abstraction each flag provides, how each thing would be defeated, then how we would work through that part. And she mapped all the concepts of how these things build on one another. Like, "We do this, we fix this, we can do this, now we can do that." It was just a master class in compilers — taking advantage of compilers to deal with more abstract languages and a lot around LLVM.

So I'm excited for that video to come out. But more than anything, it was just a phenomenal talk. If folks watch nothing else, go for that the minute the videos drop.

Ryan Naraine (01:43:11)

Shout out to Jordan, shout out to all the folks who presented at Reverse. PivotCon just published their agenda — that's on May 6th to 8th. I don't know if they even have tickets available, but the agenda is now on the pivotcon.org website. That's in Malaga, Spain. Juanito, you and I are going to be at Ekoparty in Miami in the middle of May, the 21st and 22nd. Shout out to those guys organizing that. You're keynoting the conference as well, right?

And we'll be doing a live recording of the podcast. I have an idea about a special guest that I want to try to attract there to do that. Shout out to Zack Allen, who runs the Detection Engineering newsletter. He's in our group. He is constantly promoting our podcast and sharing episodes with his audience. So shout out to Zack. Any closing thoughts from you, Costin? Close the show with some shout-outs.

Costin Raiu (01:43:58)

No, man. Just stay safe today. Two drones — Shahed drones, or Geran, or whatever, it's the same platform — they just crashed in Romania, or passed through Romania. Just be safe, everybody. Take care of yourselves. Travel safely. Get back home safely. And enjoy the conference.

Ryan Naraine (01:44:21)

Cheers, everyone. Catch you next week.

Costin Raiu (01:44:23)

Run to the airport. Bye.
