Why are you telling me this?

Keynote for VB2022, by Hakan Tanrıverdi

This main purpose of this document is to help readers follow along the presentation. As I've mentioned during the talk: While I am the one standing on stage, our research is always teamwork (our team is called BR Recherche/BR Data).

These longform investigations take somewhere between four to twelve months of intense reporting. We write code, web-scrapers and check various databases to see if hackers made mistakes and if we can find out who they are (working for). If you're interested in the type of work our team is doing, check out these articles:

- Turla and their ties to the FSB
- Winnti and how they infiltrated some of Germany's most valuable companies
- OceanLotus and their targeting of dissidents
- How to identify a ransomware millionaire

Feel free to reach out anytime: I'm @hatr on Twitter, DMs are open.

Introduction

Slides 5 to 17

I'm aware that Virusbulletin is a very technical conference. I have relied on the outstanding research presented at this very conference time and time again. So, to be asked to give a keynote is humbling. And, to be honest, my first reaction was: What can I even talk about? I don't have a background in information science, I don't do incident response and I can't reverse engineer malware[1]. What else is there?

Then I remembered some of the discussions I followed online after articles on cyberespionage were published. The community[2] was vividly discussing how reporters who wrote the articles were able to know what they wrote about. It sort of seemed like watching them reverse engineering the news. And since I watched these discussions happen after we published some of our stories as well, I thought it might be interesting to learn, step-by-step, how the news gets made. The short answer is: Because you are telling me this. But why are you doing that? Why are you telling me this?

It's a question I have asked myself quite often, and, I'm happy to report, I came up with some reasons. They're not the Top 5, but the most interesting 5. Ideally you will learn from this presentation how reporters are listening to what you are saying – and what you're deliberately not saying.

[1] I can run strings on the command line, it's true!
[2] Pretty broad umbrella, I know.

Chapter One

You’re not telling me this, Slides 18 to 37

You’re not telling me this. As a useful example, let’s have a look at this document, published by Clearsky Cybersecurity. The title is: “Raw Threat Intelligence”. It has 239 pages of observations and technical analysis. As the title suggests, this doesn’t seem to be a finished product. Who do you think they wrote this for? I think it’s safe to say they’ve written it for a rather technical audience, for their community. They're not telling me this, a journalist.

But, obviously, nowadays reporters know that if they want to find information, it isn’t enough to just read blogs by companies or look at press material only. Discussions are happening on Discord servers, on Telegram, on Twitter and in documents like the “Raw Threat Intelligence”-document. So I want to walk you through how I have interacted with this document specifically. It might get you a better understanding of how reporters go about doing their job.

In 2020, Clearsky added a new entry about a group called “MuddyWater”, apparently this group is operating at the behest of the Iranian Government. The researchers write they were able to access a server and said server contained a txt.file, sort of a log with communications received by the C2-server. Elsewhere the researchers write about “potential” targets and victims that are located in Pakistan, the U.S., Turkey and Germany. For me, it is always interesting to be able to find out who might have been targeted, so I kept reading. They didn’t publish names of any victims.

Clearsky released the hash-file of the document (218fe2e33dcf16c3254ec05c395da7ed). It is up on Virustotal. Years ago, this is where my investigation would have stopped. I don’t have an account on Virustotal. How am I going to get thils file? Today I know, if you ask researchers, they will pull that file for you. Because they are nice people and want to help.

Looking at the file (Slide 28), you can see it’s just text, consisting of Base64-encoded data. I wrote a script to decode it. Two IP addresses jumped out at me: 195.47.249[.]18 and 212.64.228[.]100. Now, the interesting thing to note is: These IP addresses weren’t included in the Raw Threat Intel by Clearsky. To be clear: I’m not implying anything here, there certainly are valid reasons to not publish information. I just want you to understand how a reporter will both look at what it is you are saying and what you are not saying. Me noticing these two IP addresses had the effect of getting curious: What’s up with those IPs?

Turns out, they belong to Bayer AG. One of the most valuable companies in Germany. At this point, I’m hooked, because in 2019, our had published an investigation into Bayer getting hacked by Winnti. So at this moment I’m thinking: Maybe Bayer was not only breached by Winnti, but MuddyWater was successful as well. I reached out to researchers, talked with them for quite some time and wanted to get additional information. I never was able to get a proof. All I had was the information provided by Clearsky. Interesting, but not a proof.

As a journalist, you can't publish allegations without at least reaching out for comment to the affected entity. Because reporting might influence the stock price, because you might be wrong and there might be a different explanation than the one you are thinking of. So we reached out, asking: What are you doing connecting to these servers, apparently operated by hackers? Bayer AG told us (Slide 36) that there was no breach. The IP addresses I highlighted were part of their Cyber Defense Center, used for checking malware. So there was no story there. Which is OK and pretty common actually. Not every investigation results in a publication. Still, there are other lessons I could take away from this. Mainly: A simple connection to a C2 doesn't have to mean anything.

All of this was jump-started by a few lines in a document released by a security company. Which I'm mentioning because, I think, for you it can be helpful to realize that you can't choose your audience. Once it's online, even if it is pretty technical, people you do not have in mind might come across your publication and draw conclusions.

Chapter Two:

You don't know you are telling me this, Slides 38 to 57

Much in the same way your industry has rules to follow when passing along information or not doing so – TLP – journalism has a set of rules when having conversations with people. There are three rules, described here[3] by the Associated Press.

1.) Conversation on-record: the information can be used with no caveats
2.) Conversation on background: The information can be published but only under conditions negotiated with the source. (Are they going to be quoted, and if yes, anonymously or is there going to be a descriptin of their vantage point?)
3.) Conversation off-record: the information cannot be used for publication.

I almost always keep my conversations off-record. Unless I'm pursuing a specific story, our conversation never happened. This will mean that I might miss out on a bigger story. Because people are talking to me off-record and I cannot use that. This puts me at a disadvantage. So why do it at all? Because I want to have a conversation in which people feel safe. If they do not have to worry that I'm potentially using everything they're saying, they can open up. They can tell me things. And if these things are interesting, I can always tell them that I didn't know about a specific fact and that I would like to follow up on that, i.e. asking other people about it.[4] And since they know what I am interested in, they can then make a judgment as to whether this is okay to do or not. Most often they just say yes.

I have one example for you. It starts with gossiping and ends up in a story. I'm meeting with a researcher, we talk about this and that and then they tell me about an alert a government

[3] AP actually has four rules, but I'm counting on background and deep backgound as one, since it's all a variation of coming to an agreement how to proceed.

[4] Technically, there is no need to do this, because off-record just states the information cannot be used for publication. It doesn't stop reporters from getting the same information from other sources, if it is widely known.

agency had recently sent out. That alert was about a hacking group, called OceanLotus, targeting the automotive industry in Germany. The researcher told me that included in the alert was a pretty specific IP address. It sounded interesting to me and I asked, hey, can I ask around? The researcher said yes.

I was able to get that alert and had a look at that IP address. Specifically I was interested in PassiveDNS data. Which domains hosted that IP address, and when? As you can see on Slide 54 and 55, among the domains there were two connected to BMW. Now, I had that alert, suggesting the campaign was ongoing, I had the IP address and I now had the name of at least one company being targeted. I did some further reporting and was able to find out that BMW's networks in fact had been breached and the attackers used Cobalt Strike to move within the network. That's the benefit of talking off-record.

What researchers consider to be interesting might not necessarily overlap with what journalists find interesting. Journalists are interested in "newsworthiness". While researchers talk about the newest techniques, journalists are also interested in finding out how these were used and against whom? So there might be a gap and to bridge it, it is helpful, for me at least, to be able to talk as freely as possible.

Chapter 3

You want me to do something, Slides, 58 to 71

This reason seems to be obvious, but it's still worthy to think it through. While doing the investigation into the BMW-hack, some people I spoke with asked me: Why are you so interested in industrial espionage when most of what this group is doing is targeting civil society both in Vietnam but also abroad? There are a lot of blog posts detailing that, among others Volexity in 2017 and TrendMicro in 2018. One researcher told me to have a closer look at the indicators of compromise. And in the TrendMirco blog there was a SHA256-hash included (2bb855dc5d845eb5f2466d7186f150c172da737bfd9c7f6bc1804e0b8d20f22a), pointing to a "delivery document". This document was both up on Virustotal and on HybridAnalysis. When HybridAnalysis runs malware in their sandbox, sceenshots are taken.

The document is in Vietnamese and mentions "Stuttgart", a city in southern Germany. It's an invitation to a bi-yearly conference. The researcher sent me this document because they wanted me to reach out to the organizers. As best I can tell, in this industry some researchers will tell people they've been targeted, but others won't. The reasons I'm hearing for not alerting them are: It can feel intimidating/threatening, remediation can be hard, notifying 100s of people might not be feasible. But it's literally my job description to reach out to people. Which I did. The researcher telling me about the IoCs clearly wanted me to find out more about the document and then tell someone in Stuttgart they might've been hacked. If they then talked with me, I could maybe publish an article and that in turn would help researchers to refine their understanding of this particular hacking operation. So the researcher gave me a tip to receive information (not directly, but by way of us publishing).

On Slide 66 you can see Vũ Ngọc Yên, who is very well known in Vietnam. He organized the conference. He didn't know that hackers sent this spearphishing document, but he wasn't too surprised that the government of Vietnam would target him, he told us. He had arrived in Germany in the 70s and published a magazine on culture. Dissidents and poets and policy people discussed pathways to democracy in Vietnam. The government doesn't like that, Vũ Ngọc Yên said.

Another person targeted was Thanh Hiếu, who is known by his handle "The Wind Trader". At the time, around 200 000 people subscribed to his page on Facebook. He received the invitation mail, some weeks before the conference took place. Crucially, around that time Bùi Thanh Hiếu and Vũ Ngọc Yên heard a rumor that the government of Vietnam paid somebody to attend the conference and then report back what they were talking about. In other words: To spy on them. This is nteresting as it showcases that cyber-espionage can be intertwined with regular espionage.

Bùi Thanh Hiếu chose not to attend the conference, but if you have a look at the phishing-mail (Slides 68-69), you can see how well it was timed. Receiving an email like this, supposedly from the organizers of a conference you want to participate in might ensure you will take the mail seriously. Bùi Thanh Hiếu clicked on the link. When we had his laptop analyzed, there were no signs of malware. When asked why he might be of interest to the government of Vietnam, Bùi Thanh Hiếu replied that he regularly publishes stories about goings-on in Vietnam and that the government is interested in who is providing him with sensitive information. If they have malware on his laptop, they can then see who he is interacting with, thus identifying the source and retaliating against them.

We published our findings and shortly after, Amnesty released an analysis of their own, shedding light on Vietnam and their government-sponsored hacking operations. (To be clear: When reached out for comment, Vietnam denied having anything to do with cyberespionage, as you can read in our story.)

Chapter Four

You want to use me, Slides 72 to 81

If we take a look at which APT groups are written about most often, the answer is going to be: the ones operating out of China, Russia, Iran and North Korea. There are many reasons for why that is the case. Just as an example: We don't know of a ransomware group operating out of Germany, and if there were[5], we'd think the government would try to prosecute them. Conveniently, these four countries also happen to be strategic adversaries to the U.S., Germany and others.

So if I talk about groups hacking in the name of Russia, the people I'm speaking with have a strong incentive to see these hacking operations fail. For that reason, they might need

[5] Do let me know!

publicity. A reporter writing a story can be helpful. This is a situation I can use to my advantage.

We published a podcast[6] about APT28, specifically about Dmitri Badin making a crucial mistake while he tried to exfiltrate the emails on a client in the parliamentary office of then-chancellor Angela Merkel[7]. While doing the research (and before that as well), I asked around and had the impression APT28 didn't seem to pose a big problem anymore. This had me curious, so I asked around.

I was sent a research paper by ESET (pdf). In it, the Xtunnel-tool is being described. (Slide 78). Basically, an infected machine needed to authenticate its connection to the server. The way APT28 chose to do it meant the server would use ultimately[8] a pre-chosen key and encrypt the string "ok". All possible keys could be derived from a Table included in the malware. The tbale stayed the same in all observed samples by ESET[9].

If you are an intelligence agency, this is something you can use. If you are monitoring the internet, you can just listen for that encrypted OK and once the alert goes off, you know who you are dealing with.

Chapter Five

It's already public, Slides 82 to 99

There's too much information out there, just look at Slides 85 to 88. This is interesting to know because I don't think that most reporters can keep track of it all. So if there is a viral tweet or blogpost going around, I wouldn't bet on a reporter having seen it.[10] It doesn't matter, though. At least not for the investigations I'm interested in doing.

Let's talk about Winnti. When we looked into their cyberespionage-operation, a lot had been written about "them"[11] already. We had received a tip that Winnti stored the name of the targeted company in the overlay of a PE-file. Knowing that, we read a blogpost by Kaspersky (Slide 93) in which this technique was described. We got the binary (again: Virustotal) and wrote a script to verify the tip. After doing this we knew that writing a YARA-rule would be useful. We got hundreds of binaries and teamed up with reverser Moritz Contag to analyze all of them (you can find the relevant scripts in this repository). Kaspersky had published their blog in 2015, we were doing all of this in 2019, at a time when the hackers were not using this technique that often anymore.

[6] In German. Tell a friend!

[7] He forgot to delete the PDB-path when re-compiling malware because the first version wasn't able to deal with german umlauts. The PDB-path included a nickname he was using ("Scaramouche"), which law enforcement were able to tie to him.

[8] I'm stepping over this and that. ESET's blog is great.

[9] Not mentioned in the report, but I reached out to one of the authors to make sure

[10] Just send it their way

[11] For the purpose of this article, it is not important whether Winnti is a group or an umbrella, ella, ella. e.

But for us, this isn’t a concern per se, since we are interested in getting the bigger picture. And in this sense, it can be useful to start investigating an operation at a later date, for then you might be able to find more relevant information (being stored in databases, written up in blogs, discussed on twitter etc.).

We also learned that in 2018 the CERT of ThyssenKrupp open-sourced a tool for scanning hosts for Winnti-infections (see our repository). We used their script to find infected servers at the time of our investigation. In the end, we were able to find out that six stock-listed companies had been targeted with Winnti: Siemens, Covestro, BASF, Bayer, Thyssenkrupp and TeamViewer. All of this with information that had been known for years.

Time having passed actually made it easier to speak with people because a lot of the techniques had already been described publicly. So instead of breaking an NDA, researchers could just send us a link and (in some cases) help us understand it.

Conclusion

The main reason why I think it is important for journalists to have a look at malware and hacking infrastructure is: It’s part of their jobs to look at the primary evidence. But also, reporting on this industry can bring with it the problem of having to rely too heavily on sources that can only be quoted anonymously. While I appreciate readers trusting me, I find it liberating to create a situation in which they don’t have to. Getting the information and verifying it for ourselves puts us in a comfortable position. The reader can, if they want to, just double-check some of our findings and make a judgment call whether they find it to be credible. If we then quote somebody anonymously it is always going to be on top of the available sources.

In summary, I wanted to walk you through some of the reasons why people reach out to me. It is equally important to mention that if researchers wouldn’t reach out to me and send me this or that information, there’s no way I could be doing this work. I want to use this very last sentence to say thank you; for reaching out and for trusting me. Thank you for telling me this.