Barclay Damon Live Presents: The Cyber Sip Podcast

Episode 21: "All Things Ransomware, With Lizzie Cookson"

Speakers: Kevin Szczepanski, Barclay Damon and Lizzie Cookson, Director of Incident Response, Coveware

[Kevin Szczepanski]: Hey, everyone, this is a Barclay Damon Live broadcast of the *Cyber Sip*. Practical talk about cybersecurity. I'm your host, Kevin Szczepanski. Let's talk.

[Kevin]: Hey, everyone. We're here with Lizzie Cookson of Coveware. Lizzie, welcome to *Cyber Sip*.

[Lizzie Cookson]: Thank you.

[Kevin]: You are director of incident response for Coveware, which is one of the nation's leading ransomware remediation firms. But how did you get into the business of ransomware response?

[Lizzie]: That's a good question. I don't come from a tech background, actually. I was coming from a criminal justice and psychology background. I've always loved reconstructing incidents and putting pieces together, and I found out that you could actually do that in a way other than crime-scene forensics. So I went to school to study digital forensic investigation, got my start as a forensic investigator in D.C. And then a few years ago, ransomware just exploded as the premier cyber threat—not just against the US, but pretty much every enterprise globally. And so my attention became a lot more narrow on ransomware actors and profiling them so that we can predict their patterns and predict better outcomes for the victims impacted.

[Kevin]: So you and I were on a panel together about two years ago. I think it was right in the midst of COVID. But let's go back five years or so. Talk to us about how the ransomware threat has evolved in that relatively short period of time. But as we both know, not so short in the realm of cybersecurity.

[Lizzie]: Yeah, it's been a pretty phenomenal arc. When I started negotiating ransoms, the most we were paying was $500. Occasionally you'd see low thousands. I still remember the first day I saw my first $100,000 plus ransom, and I thought it was a typo. There's no way it would ever be that high. What really changed was a migration away from small-game hunting to big-game hunting. Ransomware used to be just focused on consistent, modest payouts, but as all criminals do, they became very greedy. They wanted to see if they could get more. So instead of just hitting small businesses and charging $10,000 a pop, some big game–hunting actor said, well, if we just do a little bit more research on our targets and figure out their financials, we can probably get a lot more. So that is where Ryuk ransomware really pioneered the big game hunting approach. And then everything really went sideways in 2019 with Maze ransomware. Maze was the first ransomware group to start stealing data and encrypting data. And people always wonder, you know, why did that happen? It must be because more victims were restoring from backups. People have been restoring from backups since the dawn of time. It's not like that

suddenly started happening at a higher rate. What I think happened is the economics of ransomware changed. Ransomware used to be a very low-cost venture for the actor; didn't cost them anything… these open source tools. But their partners that were selling them ingress methods like access brokers, started charging a higher premium. And the only way to make sure that they are making an actual profit and not going into the red is to force victims to pay, and the best way that Maze thought to do that was to dangle some data in front of companies and threaten reputational harm. And we can't overlook the most significant development in the ecosystem in the last six months, which was the invasion of Ukraine, which immediately segmented the Conti ransomware group and resulted in a pretty chaotic disbanding and re-alliance of a lot of the big-game groups. So now right now where we are in ransomware is a pretty volatile time. A lot of the actors that have cornered the market are new to the market and they don't really have that old guard, "this is how things work approach," that maybe are REvil used to do or Ryuk used to. So we're dealing with a lot of freshmen actors, I would say.

[Kevin]: So we're talking a little bit about the ransomware market. And oftentimes we get the question, when it comes to decryption: Why would a ransomware threat actor ever give a decryption tool to the victim of an attack? Can we talk for a little bit about the mindset of a ransomware threat actor? There's an old adage that there's honor among thieves…Why do they do what they do? And why is it that most of the threat actors do tend to try to work with the victims once they get the ransom to restore their data?

[Lizzie]: Yeah, that's a great question, probably one I hear the most often from our clients. The rate of paying and receiving a tool in return, in our experience is 99%. It's very high. The likelihood that they're going to take money and walk away is actually lower than you might imagine. And there's two variables that are driving that consistency. One is reputation. It gets around very quickly in the cybercrime community, not just for victims. That's going to impact their future revenue if people find out that they're taking money and not giving a tool. But there's also a fear of retribution from their own peers. A lot of these actors have accounts on exploit.in and XSS and forums where they have resources, they get tools, they get advice. If any of those forums find out that there is a group out there withholding the key to decrypt data, they will immediately get banned. So it's somewhat of, you know, they do want to give the victim a means to restore— they don't like causing extended business interruption. But unfortunately, I think the bigger driver is just fear of retaliation from their own criminal ecosystem.

[Kevin]: Right. All right. So let's switch gears and talk about incident response. An organization suffers a ransomware attack. The breach coach comes in, retains the forensic investigation unit. Tell us about how you get involved and how you assist the response team in responding to the attack.

[Lizzie]: Sure. So best way to engage us is…we're probably going to get the call from the breach coach because they've worked with us before or they know of us and they know that we're a vetted negotiating incident response firm. We don't have any shady, under-the-table practices—which do exist on the Internet. You have to be really careful when you're searching for ransomware recovery services. We're usually brought in at the same time as containment and forensics just to get a lay of the land and assess the urgency—and know no two cases are alike. So on day one, we're brought in to profile who we're dealing with. Is this a known actor? Is it a sanctioned actor? Is it an actor that has a problematic decryption tool? Getting those questions answered on Day One pretty much sets the stage for the rest of the incident resolution.

[Kevin]: If you know it's a sanctioned actor, I know that has significance. Can you tell us why?

[Lizzie]: Sure. So if the actor is either sanctioned by OFAC or just internally restricted by Coveware (because there's two different kinds of restrictions, there are some groups that we do not pay, even though

they're not technically sanctioned by the government), it means that they well, it means a few things. So for Evil Corp., for instance, they're kind of the most obvious example. They're the named group that is sanctioned by the government. They have been tied to non-financially motivated attacks. They've been linked to sanctioned nexus. It means that we have identified there is a likelihood that a payment to that actor or that group will flow into a sanctioned nexus or directly fund terrorism. There's also restrictions that can come into play because it just...it violates what our attestation letter says. We say before we make any sort of payment, that all evidence points to this being a financially motivated crime. When the Conti group made their very ill advised, very loud announcement on their website back in February that they would be launching attacks against the West on Russia's behalf, they suddenly became a non-financially motivated group. So they were restricted overnight and we have to use everything in our arsenal to analyze every case, not just how the actor talks and behaves, but what tools are they using? Are they using tools that are written in Persian Farsi? All of that goes into the calculus.

[Kevin]: So let's suppose that you find out early on that the threat actor is a sanctioned group. Can you still assist in that situation? And if so, what types of assistance can you give? How can it be helpful even if, in the end. you're not going to be negotiating a ransom?

[Lizzie]: Sure. Even if payment is not something that we can do, we can still assist in the intelligence-gathering aspect of it. The testing of the decryption utility. So we, you know, sending them sample files, making sure they even have a key. A lot of companies out of the gate are resilient. They have good backups. They're not intimidated by the threat of a data leak, but they want to know what was taken. They want to be able to notify custodians in a timely and accurate manner. So we are often brought in just to collect details on what the actor has taken so that they can make those informed decisions.

[Kevin]: So let's suppose it is a situation in which a ransom can be negotiated. I know that you can't get into the means and methods, but can you share with us a little bit of an overview as to the steps you take along that negotiation process?

[Lizzie]: Sure. I mean, for obvious reasons, we can't really talk about how we get a price down or what our strategy is. Threat actors are always very curious about that. That's why they routinely break into emails to try and eavesdrop on Coveware-related correspondence. But we get feedback from the client first on...what is your goal? Is your goal to pay quickly because you are experiencing critical business interruption? Is your goal to delay because you're not experiencing interruption, but you don't know yet if you want to pay? So it depends on the objective of the case, and we're there to help them figure out what those objectives are. First and foremost, we just need to know who we're dealing with. Who is the actor group? Because they're all really different. They all have different personalities and triggers and behavioral tics. So, you know, when I get on the phone with a client who's been attacked by Black Hat, one of the first things I tell them is this group will launch DDos attacks against you. That is not a bluff. That is a credible threat. Do you have DDos protection? If not, get it. Whereas with other groups that's not always a very credible threat.

[Kevin]: So let's say you're at a stage where you're getting questions and little pressure from the client wanting to make that decision as to whether to pay the ransom. What factors go into that decision? What are the pros and the cons of paying?

[Lizzie]: Sure. So we lead with paying a ransom is a last resort option. It's not something to be considered in lieu of restoring from backups. It's not something that's considered as well, if we can, you know, improve the chances the data won't be leaked, we should probably pay something. No. Paying the ransom is only really justified when you have business or customer critical data encrypted that you cannot get back through any other means. That makes that the judgment of whether to pay or not

pay pretty simple. And very few of our cases really have that element. A lot more of them fall in that gray area where they have backups, but it turns out their cloud provider lost a month of it. So they're going to lose a month of data if they don't pay. Can they reconstruct that? How much time is it going to take to reconstruct that? On the data theft side, and those are the most difficult conversations is... are we willing to pay for the threat actor to not publish data, even though we know that it's not really preventing them from doing anything, they keep the data. There's ample proof. Lots of articles have been published showing that the threat actors do not delete it. They keep it in perpetuity on their servers, online. So that's the first question is just, do you need to pay to get your data back? And then the second is, and you're often a part of these conversations, what is the actual reputational harm if this dataset gets leaked? And I mean, to be fair, there's a huge difference. Sometimes people are talking about, well, I don't want some of my W-2s getting out there versus a company saying, well, I don't want all of our patients' STD results getting out there. I mean, there's a very wide range of potential harm, but people have kind of a misguided understanding of how that data actually gets misused. People think that, okay, it's published on the dark web and then all of these other people are going to go and download it and they're going to commit identity theft and wire fraud. The actual anecdotal reports we've had of that happening are nil. People who commit, you know, identity fraud and bank fraud en masse, they do that in a really targeted way. They purchase PII from PII brokers. They don't go to ransomware leak sites and download terabytes of data and hope to find a couple of Social Security numbers.

[Kevin]: Right. So I feel like the trend has changed in the last five years. When I first started talking to forensic firms, it was iffy as to whether or not a ransom would be paid. There wasn't as much of a stigma. But over time, and particularly after the Colonial Pipeline attacks and some of the more recent high-profile events, not only have industry experts started to weigh in, but the US government has started to weigh in right now and it does seem to be some heavy pressure against paying ransom. What do you think about that, and do you foresee any change in that trend?

[Lizzie]: We are also against paying ransoms. We will support our clients to the point that they need us to do that and we give them all the information they need to make that decision. But to actually curb the threat of ransomware; trying to make crypto illegal or paying ransoms illegal, that's not what's going to get us there. Just like with banning anything. It'll just continue. It'll happen underground. People will stop reporting to law enforcement and the visibility to track and capture these actors will cease to exist. What I think will and maybe already has started to curb ransom payments is enforcing reporting of the incident, you know, in a quick fashion, in an accurate fashion. A lot of companies, especially ones that are only concerned about the data leak part, they don't need data back. They definitely think there's this veil of secrecy that can be preserved and, oh, if we just pay, we can make this go away. If there was some reporting requirements that just had they had to pierce that veil, then it sort of...everyone breathes out a little bit. They said, okay, well, now, now people know. Now we're not going to pay, you know, $6 million, $10 million, $20 million because it's not a secret. And that will dramatically reduce the number of these data suppression payments.

[Kevin]: Right. So there are so many considerations that go into it. One of the things...we had a client recently that...whose first reaction was to pay because the data was frozen; client needed data to do business. There were some significant reputational harm threats. But one of the things we said is, well, first of all, we have no idea yet what the threat actor has. So there's no guarantee you're going to get back the...you're going to get access to the data that you need. The second issue that we often hear about, I want to ask you about it now, is we are hearing and seeing increased statistics that victims who pay for ransomware attacks are increasingly victims of second attacks and even third attacks. Is that your experience, Lizzie? How common is it for someone to pay a ransom and then suffer an attack from the very same threat actor within a certain period of time?

[Lizzie]: So I'm going to give a nuanced answer. The percentage of our clients who are deliberately reattacked

by the same group twice is really, really small. Most actors have sort of an internal blacklist that says, we already attack these guys, leave them alone. It's also just not economical for them to go after a company that just reinforced their security posture. It's easier to go after someone unsuspecting. Conti Ransomware...they were a group that would deliberately attack victims that paid and victims that didn't pay. That was not an insulating factor. Now, a lot of our clients report being reattacked by different threat groups than the original threat group. And that has to do with how access brokers work. Access brokers are extremely skilled in identifying everybody with X vulnerability or everybody with a public-facing IP address, and they're not going to stop at selling it to one buyer. They are going to sell it and resell it as long as there's someone willing to buy it. So it's not uncommon for one access broker to sell the same victim multiple times, and the extortionists are probably not even aware of each other.

[Kevin]: So if you pay, are you increasing the likelihood of a future attack from a different threat actor?

[Lizzie]: Not in our experience. The only thing that pretty much ensures your risk of reattack is going up is failing to contain the original incident, failing to patch the vulnerability that led to the original attack. That's really the only factor that makes that risk higher. But paying does not elevate or diminish that risk.

[Kevin]: I think that's helpful to know. And I think that might come as a surprise to some people. I think we see a lot of statistics and stories out there that aren't data-tested. Speaking of which, can you talk a little bit about how Coveware gathers its data? Because you have some great reporting on your website and you and I have talked, you keep... Coveware keeps statistics. How do you go about doing that and how does that affect your ability to help companies prepare for and respond to ransomware attacks?

[Lizzie]: Yeah, that that's one of the most rewarding things about working here is that our data set is the most robust, I think, on ransomware in the world because from the get-go, from the inception of the company, we have just meticulously tracked data points on every single case. And what emerges very quickly, as I'm sure you've noticed, is that these actors are repetitive. They behave the same way on multiple attacks. They charge similar pricing. They use similar tools and tactics. So when we are consulting with a client about, should we pay, should we not pay, what are the risks, what can we expect? It's not just well, we feel like it might happen this way or I'm guessing it'll happen this way. We have specific data that will forecast the likelihood they will get 100% of their data back, the likelihood they will get deletion proof if data theft was part of the negotiation. And so by the time they make that decision, they're doing it based completely on an informed data discussion, not just how they're feeling or what they think about ransomware. Because ransomware is a lot of things, but intuitive is not one of them. And a lot of what is written about kind of in the media space is based on conclusions that sound good in a vacuum. And are logical but don't actually reflect how the ecosystem actually works.

[Kevin]: Yeah, I think you've got to have the data and you've got to know the numbers, and what really happens from attack to attack and I know you guys do. You mentioned "deletion proof." That sounds like a really good thing. What is it and how reliable is it?

[Lizzie]: "Deletion proof" is when the threat actor provides some sort of artifact that supposedly proves they've deleted the data that they stole. It can come in a few forms, come in the form of screenshots showing them dragging the files to their recycle bin. It can come in the form of a log, so they generate what they call a shred log that shows all of the files being recursively deleted. And we have a few threat actors that provide actual videos. Some of them are 20 minutes long showing a screen capture of them erasing and deleting...all of that. The amount of credibility these artifacts have is zero. It is at best a security blanket. It's something that's nice to have. It's nice to put in our back pocket. It should

not be taken as credible proof that they deleted anything. Logs can be faked. Let's say they actually are deleting that in a video. All that's showing is them deleting one copy. They can and do keep multiple copies of these data sets. So that's why when we're talking about negotiation terms, the only thing we can reasonably expect to receive is a decryption tool. Everything else is very volatile.

[Kevin]: It is very volatile. And that leads me to another question, which is, this is an industry that is constantly changing, constantly evolving. What have you seen this year so far and what do you see coming in the next six to 12 months that may give us pause?

[Lizzie]: Something I've seen this year that has really kind of taken me aback is the level of social engineering involved in big game hunting attacks, the level of sophistication of the spear-phishing attacks. There's a new group, newish group, Black Basta that popped up around April. The first few cases we handled with them, the group targeted not the actual person they ended up encrypting. They targeted vendors and partners and customers of those companies...got into their email and then started emailing the target, looking very credible because the email chain looked familiar. It was coming from a trusted sender. That's a lot of effort to put into attacking a single victim, which doesn't really align with how ransomware's always been, which is path of least resistance, opportunistic. So that's a little bit surprising to me, the level of effort to go after a single victim. At the same time, that was sort of inevitable given how prominent EDR is now. EDR solutions used to be only used by the highest echelon of enterprises. It was a luxury. It's expensive. Now, it's pretty widespread. So threat actors have to be creative. And one of the only ways to circumvent that heavy monitoring is by going after employees...the weakest link; the end user, and hope that they can move laterally once they get in on that one end point.

[Kevin]: You mentioned employees and I was thinking the same thing. I wonder if we could talk a little bit about how the threat vectors have evolved. And I'm going to give a silly example, but I remember 10 years ago, we all got the emails from the Nigerian prince who wanted to deposit money in our account, and it just was so obviously sketchy. It didn't stop some from falling prey to it. But those were the sort of obvious attacks that we could train our employees to recognize. What we're seeing these days are increasingly sophisticated spear phishing attacks with emails and threats that look so much more legitimate than they ever did before. Would you talk to us a little bit about that? Because I think that our listeners and viewers would benefit from hearing your voice and appreciating that as time goes on, it's going to be harder and harder to spot these threats.

[Lizzie]: It's absolutely evolved to a point where the layperson is going to be tricked by one of these attacks. Generally, sophisticated phishing attacks don't even really use attachments anymore. We've all been trained not to click on attachments, not to enable macros. You know, that kind of thing links, hyperlinks that lead to a, you know, a crack site or it pulls down a payload in the background without you noticing. That's far more common. What's newer, though, is getting the emails that have no attachment, no link, but a directive to call a phone number, whether it's to finish setting up accounts or DocuSign. There's a big DocuSign campaign going out right now that direct you to call a phone number. And what happens is you call a phone number and the person on the other end convinces you to download remote access software like Zoho or Anydesk, which again is a very common technique to compromise individuals and commit really low level, low budget cybercrime. But it's being weaponized against enterprises, which is a pretty startling evolution. And those and we had a case recently where that was the attack vector. It was the phone call, person downloaded the software, luckily realized five minutes later what had happened and you know, frantically tried to uninstall it. But the thing is that...that is working. And the only way to stop that from blowing up into something worse is to have really good network segmentation, assume that the threat actor is going to get in somewhere, but make it impossible for them to go anywhere else once they're in.

[Kevin]: Right. And have those basic security controls as well, like multi-factor authentication, smart

passwords that aren't just password with a couple of numbers after it. We're finding as we all get more sophisticated and increase our ability to train our clients to improve their controls, it's the basic controls that are still lacking in so many places, Lizzie. Does that surprise you today that in 2022 you might still encounter folks who do not use MFA, whose passwords are not smart, and whose network segmentation is so flat that once you get the keys to the kingdom, you've got access to everything?

[Lizzie]: Yeah, unfortunately, I wish that surprised me. It does not. The reality is that operational efficiency is always going to be butting heads with cybersecurity, and you can have a phenomenal IT department and they can push to lock things down in a way that makes it impossible to launch a scalable attack. But they're going to get a lot of complaints very quickly from employees, from executives, saying it's impeding their workflow. And what happens is the reins get loosened and the controls get a little bit lax. Everyone's happy because now they don't have to go through all these steps to access their sensitive resources. But then once the attack happens, it also makes it easier for the adversary. The other part of that is the classic, "This will never happen to me. We're not vulnerable because we don't work...we're not a bank. We're not a huge public company. We don't have intellectual property."

[Kevin]: Right.

[Lizzie]: "Why would anyone bother with us?" That is when you're at the most vulnerable, because ransomware doesn't care what you do. They just know that you care about your data and that if you lose access to it, you will panic and you will do what you can to get it back. So that fallacy of thinking that you'd never be targeted is one of the biggest risk factors.

[Kevin]: No, it is a fallacy. So in the remaining minutes that we have left, Lizzie, if you had some advice for every organization that could improve their posture for responding to a ransomware attack right now, what would you say?

[Lizzie]: Stop focusing on prevention, focus on resiliency. You have to assume it is going to happen and know how you're going to recover from it. How are your backups tested? It is so sad when I get on a call with a client who has been paying for a cloud backup service for months or years and when they actually have to call to use it, it's oops. We've been backing up 2017 data this whole time. Sorry. Oops, data is gone. It was deleted. Sorry. Because no one's testing it. So resiliency should be the focus, not just preventing something bad from happening. And if, you know, if unfortunately, let's say something bad does happen, do not try to handle it yourself. Call in the big guns early, quickly. And let them take control early on to prevent some of these huge mishaps that can derail the instant resolution.

[Kevin]: Right. Have an incident response plan ready. Call your lawyer. Have cyber insurance coverage if you can get it. It's increasingly difficult to get it, but you still can get it and make those phone calls early. And if you need to negotiate with a threat actor, reach out to you. Coveware or someone who has specific experience right, Lizzie, because not every forensic vendor has the expertise to be able to negotiate with a threat actor.

[Lizzie]: So that's correct. And certainly laypeople, well-meaning laypeople, are not equipped to know what to do and what not to do. I can't tell you how many clients we've had who got phone calls from the threat actor and said, Hold on, I'm on the phone with cyber insurance or like it...wrote an email, "let me just check on my policies, see how much I can pay you." Just stuff that they're not thinking and that I mean, then it's just derailed from the beginning, right? So it's just better to not touch and just call in the help first.

[Kevin]: First, do no harm. The Hippocratic Oath applies to ransomware negotiations as well. Well, Lizzie, let's leave it there. Thank you for such an informative and wide-ranging discussion. I'm so grateful that you

could join us on *Cyber Sip* and I hope you'll come back some time soon.

[Lizzie]: Thank you. I would love to. I've really enjoyed our discussion today.

[Kevin]: Oh, me too. Thanks so much. And thanks to all of you for joining us on *Cyber Sip*. We're back soon with another episode.

[Kevin]: The *Cyber Sip* podcast is available on barclaydamon.com, YouTube, LinkedIn, Apple Podcasts, Spotify, and Google Podcasts. Like, follow, share, and continue to listen.

Disclaimers:

This material is for informational purposes only and does not constitute legal advice or legal opinion. No attorney-client relationship has been established or implied. Thanks for listening.

Barclay Damon *Live podcast transcripts and captions are automatically generated through artificial intelligence, and the texts may not have been thoroughly reviewed. The authoritative record of Barclay Damon Live programming is the audio file.*