NETEVENTS GLOBAL IT SUMMIT

SAN JOSE, USA

OCTOBER 3 & 4, 2019

FINAL

Turning the Tide: Dragging Criminals Out of the Shadows
Keynote Speaker: Ted Ross, Authority on the Dark Web

Ted Ross

Yesterday on the panel, I promoted this concept of adopting a zero trust model. I wanted to ask you an important question. What does that mean to you? To me, it means that if I have a device in my hand, I can trust only the things that I am transmitting from that device, which means I don't trust the network that I'm attached to, even if I'm at work. I don't trust anything I receive, email or text. I don't trust anything from my spouse, especially my spouse. I don't trust anything from my board of directors, or co-workers. It seems like it's a very difficult position to operate in, but if you strive for a zero-trust model you will be safer, right?

What we do at SpyCloud is focus on interacting with criminals and bringing in data that they're stealing hours after breaching an organization. We take that data and we parse it, make it machine-friendly, and then we pass it down to our customers so that if they have a password that's exposed, or their employees have a password that's exposed, they can remediate that password, remove it from the rotation, change the password, before a criminal can take action.

Zero trust to us also means that at work, you don't trust your employees logging in. How do you know when somebody logs in that it's actually your employee? Furthermore, when a customer logs in to your website, how do you know it's actually your customer?

We've been using multi-factor authentication. We have 90-day password rotation policies in place, we have strong passwords in place. We've had those for years, and we're still having this problem with account takeover. So, what do you do?

(One other thing too, we have a new feature coming out that we've released to our customers where we can monitor their supply chain, so I should add: how can you trust

that a supply chain or a contractor is actually who they say they are when they're logging into your network?)

It's all about account takeover, this is what we focus on, and the reason we focus on account takeover, or ATO, is because it is the number one attack vector. Over malware, over social engineering, it's the number one attack vector. Even nation states, even really sophisticated actors, have now adopted this technique to start - whatever they're doing, they start here. Now, when you hear about low-hanging fruit, a credential that's in the wild is low-hanging fruit to a criminal. So, yesterday I met with a number of you, and I talked about our use cases. I'm not going to focus on the products that we sell to enterprises, the three that you see on the left. I'm going to focus today on how law enforcement is using our data to go after criminals, and the reason I want to do this is because ATO is how criminals are coming after us.

So, why is password reuse such a big issue? First off, it's a big issue because people underestimate the power of this simple technique. Often times, you think, okay, I changed my password, I've got this formula that I use for all my different sites, and there's four characters that are the same, then I'll change the date and I'll use two characters of the forum that I'm logging into or whatever, then I'll add a 23 or a 56 or whatever, and then bang, that keeps me safe. Somebody asked me yesterday, how long should my password be, and I told them 99 characters. You shouldn't know your passwords. Your passwords should look like encrypted strings. You should be using a password manager. Ultimately, that's the ideal approach: you don't know your passwords. You know one password, and that's to your password manager, and all the passwords that you're using to log in to all the different forums is only really known by the password manager because they're too complex to even remember.

But even if you're using a password manager, that complex password can be harvested from a third-party site, and that complex password can be used against you. It doesn't matter if it's 99 characters long. If a criminal has it, they can still log in with that 99-character password. So, that's a little bit more advanced, using a password manager. Most people don't even go there. Most people have five to eight passwords in their rotation, and it's human nature to do this. We've all done it. I've stopped doing it over the years, but the older we get, the more passwords we have in our rotation, which means the younger generations are really in trouble because they have fewer passwords to rotate, and they're way more active online, right? Does that make sense?

Then, there's this hidden attack vector. Often times security teams tell us about all the different things they have in place to protect their employees, like if they leave with a laptop and it has a firewall on it, and EDR, and all sorts of stuff on your laptop to make sure your employee's safe, and they go home and they sit on the couch next to their kids who are surfing game sites behind a residential firewall. Their children and their spouses are way more exposed than they are, and if I'm a criminal - I'm going to show you this later - it's easy to find out your personal email address and find out who your family members are, and go after them instead of your work credential. Alright?

So, I thought I'd share this story with you. This actually came from one of our customers. We have a large financial customer that manages 401ks, and every once in a while,

they'll get a phone call, and their customer will say, where the hell did my 401k go? It's missing. So, they have to do some research and figure out what happened, and unfortunately what they're finding is that three months prior to the 401k going missing, they found that this email address that they're using at the bank was part of a breach. In this case I'm going to pick on fantasy football. Let's say you logged in and you played last year, you forgot you even had an account and a password out there, and fantasy football gets breached. Their credentials get harvested, and now they're in the hands of a criminal.

Well, in this case, the criminal tried to log into the bank with that password and it didn't work, because the bank's password requirements were different than fantasy football's. But the criminal was able to log into their Gmail account. Once they logged into their Gmail account, they did a password reset at the bank, and then they were able to access the bank. The criminal is fairly sophisticated too. They know that if they were to drain the 401k immediately that would have raised some flag at the bank and they would have stopped the transaction, so what they did was, they changed one thing each day for seven days. Day one, they changed a phone number. Day two, they changed an address. Day three, they changed something else, and then seven days later they drained the funds. True story, happens more than you would know, and what I didn't realise at the time when our customer shared this with us, they use our data to go back to the customer to show they weren't culpable, number one, but also to hopefully retain that customer when they explain what happened. I didn't know that 401ks are not insured like checking and saving accounts. It's mind-boggling.

So, we talked yesterday as well about BEC fraud, business email compromise fraud. This is the number one form of fraud, period. I mentioned that 65 percent of all fraud being reported to the FBI is BEC fraud. So, what is this? It's CEO fraud, account takeovers where somebody accesses your inbox and they send emails on your behalf, and it's things like romance schemes. This is a number that came from the FBI two years ago, so it's a bit old now; I'm sure the number it a lot is higher now, but two years ago, the scary part was that the amount of money that has been recovered from these schemes is only three percent. Tells a grim picture of the reality that we live in right now. The best way to combat this is to adopt a zero trust model.

Alright. So, let's talk about how the criminal operate. I like to use the account takeover timeline to walk through what they do and how they operate, and the different levels of sophistication across the timeline. So, day zero, the organisation is breached, or their credentials are spilled. A criminal will harvest credentials from a site. When they harvest credentials, they're not just taking email addresses and passwords. They're taking everything that they can get their hands on. If the site logged your IP addresses every time you logged in, they're going to take that. If your site remembers your old passwords when you change passwords so that you can't use a password twice, they'll get the old passwords as well as the new passwords.

So, another story really quick. I was presenting in DC to the Federal Identity Audience, and after my presentation somebody came to me with his phone and said, hey look, while you were presenting, I got this email, and it says that they have all my information, and they included a password and they're asking for $1200 in bitcoin or they're going

to blackmail me. When I first saw it, I said, you know, we get those all the time. There's billions of credentials that are out there. Your old passwords are out there, they've been out there for years, so they're just trying to use your old password to scare you out of your money. He said, yeah, but that's my password that I use at a particular bank, and I only use it in one place. So, that was an oh shit moment.

So, we took his password. We put it into our system to find out which breaches in the past had that password associated with it, and there were two. When he looked at the forms, he realised that this password that he uses at that particular bank - first off, that bank was not breached. This wasn't a problem with that bank, but the password he used at the bank, he uses every once in a while when some system asks you to change your password, and you don't want to come up with a new one and change it, you use the password that you know temporarily, and you go back a couple of days or a week later and you change that password, because it's your special password and you don't want it out there. Well, during that short period of time, a forum could be compromised, and your special password could be revealed, but what's more likely is the forum that you put that password on remembered the old password, and when it was eventually breached your old password was leaked, and that's what happened in his case. So, he realised very quickly that, oh, great, this is a bigger problem than I thought.

So, anyhow, day zero, passwords or credentials are spilled, and then they're quickly shared with a small team of individuals that focus on monetizing this information, and they do some pretty interesting things along this timeline. So, between day zero and day 500, usually a year and a half to two years, they'll do things like they'll fingerprint your organization. They'll try to log in with a credential using a botnet, and if you can detect a botnet, they'll try a different technique until their login goes undetected. Once their login goes undetected, they'll try your password. If it doesn't work, that's fine, because there are tools out there, like a tool called Purple Spray, which will take your password and try thousands of different variations of your password, and it will figure out the site's brute force threshold detection and it will stay under that. If it has to work for years, that's fine. It's all automated. They don't have to do anything. They just start the job and let it go. Eventually, they'll get a password that works, right?

This is why even if you have four or five characters as part of your rotation and that password leaks, you have to throw away all the passwords that have used those characters in your rotation.

So, they'll continue to do sophisticated things all along the way, and once they finish monetizing the information - by the way, I should also mention at this point in time, if your company has a 90-day password rotation policy in place, I'm just curious, how many of you are forced to change your password every 90 days? If you'd raise your hands. Okay. Go back to your security team and tell them that Ted Ross told you to stop doing that. The reason for that is, every time you change your password, if your password hasn't already been exposed, you now give the criminal another chance. You're playing into their hands. You only change your password if it has been exposed.

NIST no long recommends a change of password every 90 days. Large, sophisticated security companies have removed that policy altogether. So, eventually, day 500 - yes, sir?

Audience – Male

What do you mean by password has been exposed?

Ted Ross

That means that your password, you may have used it at a third-party site. Like, do you like cars? Do you log in to Porsche.com, or Audi? Well, maybe one of those sites will get breached, and the credentials that you used to log in to that site will be in the criminal's hand at that point, so that password has been exposed at that point. You didn't do anything. Your company didn't do anything wrong. Your company wasn't breached, you weren't breached, but a third party had a problem and your credentials leaked out into the underground.

So, day 500, they finish monetizing this information. It somehow always ends up leaking to the deep and dark web on a forum somewhere, and that term deep and dark web is something that we don't like at SpyCloud. It's very misleading, but because the industry has adopted it, we'll stick with it in this presentation. When the data ends up on a forum somewhere on the deep and dark web, hundreds of criminals will have access to it. Whereas before it was a small team of sophisticated actors, now there's a lot of people, a lot of criminals, fraudsters, who have this information. They're not sophisticated, but they're creative.

It also turns out that law enforcement and security companies are all over the deep and dark web. We're scanning it constantly. Then, eventually, the data that was leaked on day zero will end up in a combo list, and this is usually when we read about it in the press. Somebody will pick it up and realize there's 1.4 billion credentials that are out there, it's one of the largest exposures ever, and all it is is hundreds of prior breaches put together in one database, and now it's being marketed in the underground as a new database, but they're credentials that have been used already.

So, we like to think of credentials in this space in two different categories. When it gets to the deep and dark web, it's too late. It's too late to remediate. It's already been used by a sophisticated actor, most of the time. So, credentials before that are very highly valued to criminals, and this is an area that as an industry we need to do better at: discovering credentials in this part of the timeline so that we can remediate the problem before a criminal can take action. So, this number came from a customer of ours, a large financial institution that focuses on cryptocurrency. They have a way of measuring the attacks that come into their company, and they have categorized attacks into two different buckets.

There's targeted, meaning that somebody has your credential but they're also trying to get a new token to bypass your multi-factor authentication. What they've found is that the targeted attacks are 10 percent of all the attacks that they receive. Then, once the credentials are a commodity, this represents 80 percent of the attacks that they receive.

As an industry, we have been focused completely on stopping that 80 percent for the last decade. So, we have tools like Akamai. Akamai does a fantastic job at detecting botnets and malicious IP addresses, so if you're using Akamai, if your servers are behind Akamai, they're probably taking care of the 80 percent for you. Think of it as an account takeover firewall. That's what you need to block the 10 percent on a massive scale.

Unfortunately, the 10 percent that comes from sophisticated actors causes 80 percent of the loss for organisations, and that's a space that in the past has been fairly underserved. The only way you really can attack it is to use tradecraft, like human intelligence, and social engineer the criminals out of the data as early as possible in that timeline so that you can basically beat the criminal to the password. Change the password before a criminal can take advantage of it.

So, I thought I would share this with you. MyHeritage, MyFitnessPal, the Under Armour breach, these are fairly recent breaches. They were in the press. I think most people have read about these, and probably a lot of people in this room were part of these breaches. That's why I highlighted them in yellow, but there's a lot of other breaches on this slide. We've been measuring when the breach actually occurred, and when it leaked to the deep and dark web and everybody was able to gain access to it. All the security companies should have access too at that point. When it became a commodity. As you'll see, criminals have a long time with these credentials to do sophisticated, nefarious things with them before it becomes a commodity. The last column is what I'm pointing at here.

Alright. So, now I've set the stage for what criminals are doing, how they're using it, I'm going to walk to the back of the room so that your attention is not on me, it's on the screen, and I'm going to show you our investigation tool.

So, we have law enforcement using our data to go after criminals. We also have large enterprises that have sophisticated investigator teams going after criminals, and I'm going to flip to a tool called Maltego. So, let me run you through an example of an investigation. For those of you that are recording this, please do not record the screen right now. I'm going to show you things that we do not want in the public's eye. I'm sharing this with you as a special preview, but there are things that might be revealed on the screen that a criminal could use to understand how we're tracking them, and we don't want that to happen. So, feel free to use the words and the story, but please don't take pictures of the screen as we go through this.

So, are you all familiar with the Silk Road story? Yeah? Good. So, Ross Ulbricht is the guy who created Silk Road. He's now serving a life sentence in jail. Brilliant kid, went to Westlake High School in Austin. Silk Road operated between 2010 and 2012, and while it was operational it pushed about $1.2 billion worth of drugs through the site, so law enforcement was really, really focused on bringing this guy down, but they couldn't find him. They couldn't find the site, because it was on the TOR network, they couldn't find the creator of Silk Road. All they knew was, somebody who went under the moniker Dread Pirate Roberts was operating the site.

At one point in the investigation, an agent did a simple Google search, and he looked for the mention, the promoting of Silk Road at the same time the site came online. He found some forums that promoted it, and he subpoenaed those sites, and he got back an email address, rossulbricht@gmail.com. So, this is one of the first times in the investigation they actually have a name of an individual.

So, they profiled Ross Ulbricht. They looked for three things. Was he involved in bitcoin, was he a web developer, or was he a computer science major in college? Ross was none of those things, so he didn't meet the profile. The investigation continued for many more months. Had they had more data, they could have plugged in his email and done an exact match search on his email, and they would find very quickly that we have a lot of records on rossulbricht@gmail.com, so let's see what we can learn about all these records that are in the 13,000 breaches that we have in our database, the 80 billion assets that we have in our database. It's a database full of criminal activity, in addition to all of our activity.

So, let's ask some questions. Let's find out, first off, what alias is Ross Ulbricht using with these various identities and, more importantly, what forums is he logging into? What's he doing on the internet? Let's study this guy a little bit. We'll find out very quickly, on the screen you'll see this Rossman alias is logging into two different forums, PHP Freaks and DevShed. Both of those are developer forums, so now we check one box of the three that we're looking for. Over to the right, we'll see [BondiBlues] is logging into bitcoin talk forums. That's the second thing we're looking for, so we've checked two boxes out of three. This alone is enough to go pull Ross into a room and have a conversation with him. Now, he looks like he might be a person of interest. But because we crack passwords and we do other things, we have a whole bunch of other data on these individuals, let's bring up some additional information on the screen. Let's see what kind of passwords Ross was using.

Over here, we see that there was one password in particular that he was using a lot, so let's pivot on that password. Remember, password reuse is human nature. If we started out with a Gmail account, we're betting here that he reused this password across other identities, and sure enough, we started with his Gmail, and now we find his Yahoo account. He did reuse the password.

Our head of investigations loves to take all of these new Yahoo records and take the investigation all the way down to where we find his criminal identity. I don't have time to do that, so I'm going to just shortcut it. We always recommend to our investigators, if you start an investigation by looking up an email address, you also want to do the same lookup as a username, because we pull from different records in our database. So, if I do an exact match on Ross Ulbricht as a username, I stumble across this email, dprjkt - remember, he ran Silk Road as Dread Pirate Roberts. So, now I'll tap into our fuzzy logic capabilities, and bring back a number of Ross' criminal identities, and there's a lot of them.

Remember, we track IP addresses, full names, passwords, so if you're law enforcement this is just a wealth of information, but the fact that he's involved in bitcoin, he's looking at development of websites, and now we've linked it to a moniker that has

Dread Pirate Roberts, this is the time when you just go and pull him in and put him in handcuffs. So, that's all fun, and you can understand it pretty well because it's a cyber-criminal. How would an enterprise customer make use of this information? Earlier, I told you that if you hired an employee, it is one of the weakest links in the attack surface, right, and that slide that I had where if you hired one employee, you had to think about their family members? Well, this is one way that criminals are using this type of data to identify your personal email addresses.

This is a real investigation. We started with a bank, and we plugged in to their domain, and then we redacted it. So, this came from a real investigation, but we hid the name of the organization. It was a medium-sized bank. Well, when you plug in a domain, it's pretty easy to find out what IP addresses a company is using. This is an open source pivot to find an IP address. Anybody can do it. Once you find an IP address, you can then start querying the database and bring in all the records from prior breach sources that have used that IP address before to surf the internet and go to forums, right? So, now we have the bank's employees' personal email addresses. They went to the bank, they logged in to fantasy football. Fantasy football eventually got hacked, and now we can tie this person's personal identity to the bank. If I'm a criminal, I'm now going to go after this personal identity, and not the work identity because I'm expecting that to be secured.

There are a number of other use cases we can demo here. I showed one yesterday a couple of times. I should probably show it again. This goes beyond hunting for cyber-criminals. There was a bomber in Austin two years ago. He was putting bombs in packages, and he was putting those packages on people's doorsteps, and when they would move the package they would blow up, and he was killing people. So, during the time he was doing this, the SpyCloud team did some research, found some forums that promoted this idea of bombing in Austin, and these are normal criminals who think they can hide behind an alias. Now, information on those criminals can be revealed, and I'm going to zoom out so I redact it, in a way, and you can start bringing in all sorts of information very quickly on these people that are doing pretty sophisticated terrorist activity. Scary terrorist activity.

I'm just showing you a quick glimpse of the type of data you can pull in. We have a geolocation here, we have an email address, we're going to have some full names. We have passwords of this person who claims to be the Austin bomber, and we have IP addresses where they're logging in, and here's another alias that we have. All these are assets that I can drill into and explore further. So, again, a great set of data to go after criminals in a number of different ways.

So, with that, I'm going to flip back to the presentation. Did all that make sense? I know I went through it really fast. Any questions on the investigation tool?

Audience – Male

What does it look like when it's a nation state rather than an individual?

Ted Ross

We typically don't focus on nation states. So, if we find that somebody's involved in nation state activity, we'll flip it over to somebody who actually can act on that, and we'll share all the information with them. We focus on cyber crime specifically. We actually don't want to tackle other domains, because we want to be experts in what we're doing. I'll tell you a quick story. This is an email that came in soon after we came out of stealth mode, went to the guy who runs sales, and it came from me, apparently. He asked, can you send me some money? It wasn't a very good attempt either. So, we sent it over to our CTO, who loves playing around with these criminals. He set up a whole storyline, and convinced the criminal that we were lovers and we wanted to run away and embezzle money, and we needed accounts to put the money in. He was going to help us out. So, the criminal gave us a whole bunch of accounts.

Every time he'd give his account, we'd call the bank and we'd have it checked out. It was a lot of fun. The way we found this, the way we could verify it - and it's something you could do very easily - when you have an email that comes in your inbox, go to the raw data. I'm showing you a screenshot of an Apple Mail raw data check and somewhere in the raw data you're going to see the reply to email address.

Now, I use a different email for this and I see [Vik] over there laughing because this is [Jennifer], our CFO, who got a similar email later on and when we looked at the raw data it came from this other email address. So now we have a way of looking at who this person is, finding out a little bit more about them, right? I highly recommend you adopting this technique when you get an email that you think is strange. Look at the raw data, see the reply to, see if it actually matches the display email address. If it doesn't, you know it's fraudulent. But even if it does map to the email address you have to be careful of one thing.

I have a story, a title company, true story. Title company was breached. The criminal was watching somebody's email. They knew the day that somebody was supposed to put a down payment on a house of their life savings. So they sent an email to that person from the title company's inbox, with the title company letterhead all formatted from prior emails the title company sent, so it looks completely legit. They got it on the day they were expecting it and they asked them to wire transfer the money, they gave them the account number and they did. Hours later, or a day later, they get an email from the title company asking them to wire transfer the money. This time it actually came from the title company. So if you look at those emails, the reply to is actually the title company because they were compromised and the criminal is sending email from their company. So you still can't trust it, even if it comes from the company.

So I typically like to leave with recommendations. We've already talked about this enough, I'm not going to pound it into you. We know the password managers don't scale to tens of thousands of employees, but as an individual you should be using a password manager everywhere you go. Use a VPN, don't trust this network that you're on right now, don't trust coffee shops, don't trust your work network. With that, I'll take some questions real quick before I get kicked off the stage. Yes, sir?

Audience Q&A

Audience - Male

[Inaudible]. There's a variety of password managers out there, so my first question is are there particular ones you recommend? The second one is a topic I always bring up with friends, it's like where should you actually keep passwords other than your head? Because a lot of this advice is sometimes practically impossible when you set 12 passwords. I imagine I have over 50 passwords and I actually keep them in a secure location on a piece of paper. Is that a way to do it? Or if you keep a file on your laptop, how do you keep passwords?

Ted Ross

If you were from Texas, I would ask you, is it in your gun safe? My neighbours had their house broken into the other day and their gun safe was stolen. So I think it's pretty common for people to do that, quite frankly. If you're in cryptocurrency you probably have your own hardware-based wallet and if you ever lose that wallet, the way to recover it is to have a piece of paper with the answers to the questions so you can recover your wallet. So if you do that, just make sure that piece of paper is really, really secure.

But I would recommend using a password manager because if your password is compromised, password managers will have access to datasets like ours and they'll come back and let you know this password, even though it's really crazy, 99 character string, it's been compromised and they'll force you to change it. So I do recommend a password manager as a consumer; Dashlane, they do a really good job as an enterprise, Password Keeper. Both are great companies, I'd recommend either one of them for enterprise or consumer, but it just seems like consumers are going to gravitate towards Dashlane and enterprises gravitate towards Password Keeper.

Wayne Rash, eWEEK

Hi, Wayne Rash from eWEEK. You recommended using a VPN as one of the things you should do, but one of the things that we found in looking at VPNs is that I think the vast majority of consumer VPNs are run by companies you can't find anything out about. They terminate at servers in very odd countries and the question I have is…

Ted Ross

I think you're talking about NordVPN probably, which has thousands of termination points.

Wayne Rash, eWEEK

Yes, Nord is one of them, but there are several of them. How do you figure out a VPN you can actually trust?

Ted Ross

Great question and when we started to search for a VPN provider, we went and asked the same exact question and we ruled a number of them out because we simply didn't trust their termination points. You have to be able to trust your VPN provider more than you trust your corporate network, in a way. So there are companies that are headquartered in the US and they put their policy out there very clearly. They'll tell you how long they retain records, what they retain and what they would do if they were subpoenaed. They give you all the details so you know exactly what you're dealing with. If they don't do that, don't use them. Any other questions? All right, thank you very much.