#181 - Inside the 2024 Verizon Data Breach Investigations Report

[00:00:00]

[00:00:00] **G Mark Hardy:** Hello, and welcome to another episode of CISO Tradecraft, the podcast that provides you with the information, knowledge, and wisdom to be a more effective cybersecurity leader. My name is G Mark Hardy. I'm your host for today, and we're going to be talking about the latest Verizon data breach investigations report.

But before we dive into that, a couple announcements. Congratulations to CISO Tradecraft. We have passed 30, 000 followers on LinkedIn. That's huge. We started out as two guys and a dog almost four years ago, and now we have had hopefully a big impact on the lives and the careers of a large number of people in our cybersecurity community.

If you're not one of our followers, please go ahead and do so at LinkedIn and look for CISO Tradecraft. We have more than podcasts. We'll put [00:01:00] out a high signal, low noise. Daily or every other day type of posting of information that we think you're going to find valuable in your career journey. Also, if you find that your cybersecurity needs expand to need some expertise, we can help build training programs for your team and for your organization.

We create and moderate tabletop exercises and we can provide career advice and counsel. Go ahead and contact us at LinkedIn or go ahead and send us an email and you can reach me at gmhardy@cisotradecraft.com. No spam, please. And we'll look forward to hearing from you.

Today's episode is about the most recent Verizon Data Breach Investigations Report, or the DBIR as it's often called, and this year they passed an inauspicious milestone with more than 10, 000 breaches analyzed in a single edition. A bit of history, back in 2008, Verizon's Business Risk team released their [00:02:00] first DBIR, looking at a four year period of over 500 forensic engagements by their investigative response team.

This modest, by today's standard, 29 page report offered a summary quote that I'll share with you, a little flashback. Most breaches resulted from a combination of events rather than a single action. Some form of error, often directly or indirectly, contributed to a compromise. In terms of deliberate action against

information systems, hacking and mail code proved to be the attack method of choice among cybercriminals.

Intrusion attempts targeted the application layer more than the operating system, and less than a quarter of the attacks exploited vulnerabilities. 90 percent of known vulnerabilities exploited by these attacks had patches available for at least six months prior to the breach. Wow, okay, 17 years back.

Now, a couple other facts and figures from that particular edition. 71 percent of breached systems had patches [00:03:00] that were unapplied for over one year. 70 percent of breach notifications were by a third party. 55 percent of attacks required low or no skill to complete, essentially the work of script kiddies.

External breaches averaged 30, 000 records each. 15 percent of attacks were targeted, 4 percent of incidents were detected through event monitoring, and only 3 percent of data compromises involved medical or patient data. Now think about how things were 17 years ago and what you have seen change.

Now let's move forward to today and look at the Verizon 2024 Data Breach Investigations Report that was released earlier this month and see the top ways that bad actors are exploiting companies according to this study. You can find a copy for yourself at verizon. com slant dbir. Now, once you've downloaded and have access to that report, [00:04:00] take a look at figure seven on page 12 out of 100 in the Verizon report entitled Select Ways in Variety and Vector Enumerations in Non Error, Non Misuse Breaches.

Holy mackerel, what kind of title is that? Anyway, seven different ways that attacks occur. And we thought we'd go ahead and take a look at them starting from the bottom up. Number seven is bad actors exploiting vulnerabilities in virtual private networks, VPNs. For simplicity, either of the engineer who designed it, or for ease of communications, some companies have a flat network.

So if you can get an internal IP address, then you can see almost every other internal device, making it a breeze to attack. VPNs receive patches from time to time, often to address security vulnerabilities if you haven't patched a VPN. That's an issue. So bad actors look for VPN concentrators that have known vulnerabilities.

Examples of that include multiple threat actors actively exploiting Ivanti Connect, secure [00:05:00] and Ivanti policy Secure gateways. This became so common that CISA put out a cybersecurity advisory on Ivanti and the 29th of

February this year to help companies understand how many actors were using this attack across multiple CVEs.

If you thought that Amandi was the only VPN provider in this category, you'd be mistaken. CISA has made similar announcements for Palo Alto, FortiGuard, Pulse Secure, and others. And this isn't anything new. Some of these announcements dig back to 2019.

Number 6 is bad actors exploiting vulnerabilities on desktop sharing software.

If you have Remote Desktop Protocol, RDP, or Virtual Network Computing, VNC, on your endpoints, then the bad guys just might use that tool to try to log in. Many employees travel, and when on travel, it's common to use airport and hotel Wi Fi hotspots. And this provides a potential means for bad actors to scan your laptops.

If there's something that allows remote access, bing, it's a target.

If these tools are not kept [00:06:00] up to date, known vulnerabilities could be exploited to create remote access when it shouldn't be there.

Number five is that bad actors steal VPN credentials. Passwords that are stolen in a public breach from Facebook or Okta or some other large IT provider often end up on the dark web.

If your employees reuse those same credentials and your systems don't have multi factor authentication, MFA, it's going to be trivial for a bad actor to reuse those same credentials in a password spray attack to break into your company's VPN while looking like a legitimate user. MFA has been considered best practice.

Change your thinking. MFA should be the only practice. Don't settle for 1990s era ID and password credentials when so many resources are available today to enforce MFA. Number four is bad actors stealing credentials to desktop sharing software. Have we noticed a theme here? If bad actors can do it on a VPN, they'll also try it on [00:07:00] desktop sharing software.

If your help desk uses this type of software to remotely troubleshoot your employees, please use MFA. If you don't know how, then Google RDP and MFA. There's a lot of easy fixes here. In addition, Bad actors may be skilled at social engineering. They'll call your employees, pretending to be the help desk, telling them they've identified malware and they need to fix it.

Let's open up a quick assist session, and please click give control when asked. See, non technical employees, there's a lot of them, will not recognize that during this help session, the attackers are inserting malware, a reverse shell, or some other tools while telling a narrative that sounds plausible. If there's a way that someone can remotely access your machines whether through a VPN or desktop sharing software. These tools need to be locked down and your employees need training to recognize when that help call is from someone other than your IT team.

Number three is exploiting vulnerabilities in web [00:08:00] applications. Most companies have a public website to allow customers to see their offerings, contact them, purchase things, view their account activities, or make payments. And each of these sites needs to be free of vulnerabilities and misconfigurations. However, it's nearly impossible to know when you're 100 percent secure, and just because your SAST and DAST tools didn't find any vulnerabilities, doesn't mean that they're not there. Additionally, if your team isn't patching in a timely manner, they're not secure.

It's another big issue. In December of 2023, the Qualys Threat Research Unit prepared a comprehensive blog series to review the threat landscape in 2023. One of their key takeaways is 25 percent of security vulnerabilities were immediately targeted for exploitation on the same day as the vulnerability itself was disclosed publicly.

And 75 percent of vulnerabilities were exploited within 19 days of publication. Now, if your patching goal is 30 days, you're holding the door open for an awful lot of attacks. [00:09:00] Now, what kind of compensating controls could you apply? Run multiple different tools in front of your web applications. Add a web app firewall or a runtime application self protection RASP tool to give you two layers of defense that allow you additional time to patch your systems and the ability to block exploits of zero days on your websites.

Number two on the top attacks bad actors are using to break into companies according to the Verizon Data Breach Report is phishing. This is no surprise since every company uses communications tools like email, Microsoft Teams, SMS, Slack, LinkedIn to communicate. That's where people are clicking links as part of their default behavior.

You can expect bad actors to insert bad links into those services. Now this year, ChatGPT has really gone mainstream. This means that non native English speakers are no longer at a disadvantage for composing phishing emails that look like this. not fishy. Anyone can go to ChatGPT and write a convincing

[00:10:00] phishing email in proper English or whatever language your target audience happens to be.

I've seen examples where you can also ask generative AI to help write malware as well. We're now seeing attacks that were completely written. by AI, and bad actors will continue to innovate in the space. Here's another interesting figure. According to figure 9 on page 9 of this year's DBIR, 20 percent of users reported phishing and security awareness exercises, and 11 percent of users who clicked the email also reported.

okay. Sounds okay, but consider the average user clicked on a malicious link within 21 seconds of opening an email and starts entering data. with 28 more seconds. Think about that. The median for users to complete the sequence for an attack on a phishing email is less than 60 seconds. Continue to use email security gateway solutions.

Focus on decreasing the click rates of your employees and put in safeguards to minimize the harm of [00:11:00] employees being phished. For example, do you have conditional access policies that say that any user must come from a trusted device or on the VPN to access all of your applications. If you did. It doesn't matter if a bad actor steals login credentials from an employee, right?

They still lack the certificates that would be on that person's machine, and that will prevent them from logging in. And if you've instrumented your systems correctly, that attempt will fire an alert. Your SOC or MSSP can use to go act on that. Which brings us to the number one attack in the non error, non misuse breaches in the Verizon Data Breach Report.

It's credentials for web applications. What does that mean? Bad actors will try a lot of credential attacks to get into websites. Let's say users username and password log into their favorite email has been compromised and posted on the dark web.

Bad actors can purchase blocks of credentials there and try reusing those [00:12:00] credentials to log into bank accounts, public shopping sites, or even corporate systems. They might just go to those sites and select, I forgot my password. Only to have a new password or reset message emailed back to the now compromised account.

That was easy. And if this happens and they delete the password reset email before the user notices, that access could persist for quite some time before it's noticed, usually by an unexpected draining of funds or a surprise purchase being

charged to an online account. So ensure your users use unique passwords on each site by leveraging a password manager.

Insist on MFA for systems you control and strongly encourage your users to use MFA on sites that you don't control. Now, I use a Yubikey to log in as a global admin to my tenant. Plain old MFA isn't enough for that level of control in my humble opinion. Also, consider using pass keys for any login that supports it.

Check [00:13:00] out CISO Tradecraft episodes 74. Pass the passwords for a review of the eight levels in the consumer authentication strength maturity model where passwordless, that is a physical token, is considered the highest level of authentication. Even if a user utilizes unique passwords for every site, it doesn't mean everything is safe.

Bad actors will try other attacks. This year, we're really seeing a lot of man in the middle reverse proxy attacks, where a bad actor tries to intercept a user logging into a website, steal the session cookies, and then log in to the same site post authentication. Tools such as evilginx, Man in the Middle Framework, and Evil Proxy, Reverse Proxy can help perform this attack.

So perform Purple Team exercises and verify if you can protect or detect this type of attack. Direct all browser connections through a proxy service that's going to protect users from reaching known evil websites. Use passwordless authentication methods with [00:14:00] FIDO2 security keys, as I said, like a YubiKey, or apply conditional access policies with VPNs if you're struggling with these types of attacks.

Another form of credential attacks against web applications is credential stuffing attacks, brute forcing passwords until guessing a valid one. Remember that. Password 123 bang, with the uppercase P, while meeting most complexity rules, isn't a strong password. Even with mandatory password changes, the average user will increment the number at the end, or pick something current like their March Madness bracket winner, or their pet's name, and follow it by 123 and a piece of punctuation.

Smile if you've done that yourself. Bad actors are going to try these easily guessable passwords that could be brute forced or credential stuffed. The NIST Special Publication 800 63 BRAVO, which tells the recommendations for these things, has removed the recommendation for changing your passwords.

on a [00:15:00] regular basis. And also change the complexity requirements. Some government sites still say every 90 or 180 days, you have to do uppercase,

lowercase, number of special characters, something like that. you're not going to remember that. You got to write it down. You've got a password manager tool.

That's great, but the average person doesn't. Furthermore, length is always better than complexity. And you've heard me talk about this before. You go back to the day of the PDP 8, when you're trying to cram as much entropy into eight bytes as you can, or eight characters. And as a result, that's where our complexity rules came from.

Use something really long. The little girl went to the store to buy some food for her cat. NSA is not going to crack something that long. Length beats complexity every day of the week. So crank open your password length. Did you know, and Microsoft will support at least 127 characters? Not that I think you need that many.

Although if you do have a break glass account, that is a global admin without MFA, just in case your phone network [00:16:00] goes down. I would consider approaching that limit just in case. Hey, ask your developers what happens if a user has 10 incorrect logins in a row. Does the app just keep allowing them to retry different passwords?

Or does it disable the account or suspend it until the user contacts a help desk to reactivate it? Remember though, if somebody steals the password file, even if it's encrypted, they get an infinite number of offline tries to brute force the password. Now, we at CISO Tradecraft recommend MFA for all users on all systems, and a look at recent breaches indicate that the lack of MFA is a major contributing factor to these breaches, and it's probably going to be a difficult courtroom defense if you had this level of authentication available.

But you didn't use it. Now a couple other notes worth mentioning in this year's report. Denial of Service remains the champion again in this year's DBIR, responsible for more than 50 percent of the incidents analyzed. Distributed Denial of Service, or DDoS, attacks are relatively cheap and they work [00:17:00] well against a number of targets.

Content Delivery Networks, or CDNs, took the brunt of the traffic. And although the median attack size dropped from 2. 2 gigabits per second to 1. 6, on the high end, the standard normal 97. 5 percentile or 1. 96 Sigma. Go figure out standard normal. You take 95%, throw the two and a half off on either end because those outliers could really screw with you.

Go look it up in the statistic book. Put that. Standard normal max increased to 170 gigs from 124 gigabits per second. The majority DDoS attacks can easily overload a gigabit switch, let alone a website and the connectivity you have. Now here's what could make that troublesome. Ah, I don't have to worry about that.

I'm in the cloud! Great! What's the impact on your cloud spending? Do you think about that? Your website might scale elastically to meet increased [00:18:00] customer demand. That makes sense. Bring on another virtual server. Keep that going. But what happens if you get a sustained 40 gigabits per second of traffic?

Not only do you have to worry about the availability of your website, but your AWS costs could go to the roof. Could cost you thousands of dollars. Do you have the budget to cover that? Do you have someone monitoring it 24 by 7 to turn it off if it does happen? Now, we recently shared one story on our LinkedIn feed.

Again, that's the reason I say go take a look at our LinkedIn feed. Where a developer reported he had an AWS S3 bucket. Okay, so what? Unfortunately for him, one popular open source tool, which he did not want to name, had a default configuration that stored their backups in S3 and as a placeholder for bucket name, They use the same name that he chose for an S3 bucket.

That means that every default configuration attempted to store its backups in his S3 bucket to the tune of a hundred million [00:19:00] S3 put requests in a single day. AWS charges for unauthorized requests, by the way, in a 1, 300 bill for 24 hours of this was, to say the least, a bit too high. Unexpected.

Apparently that's part of the contract. So read the fine print carefully. Use unique and unusual names for your S3 buckets. The name can be guessed easily. You might be in for an expensive surprise. Now note that AWS employee, Jeff Barr later mentioned in a public tweet that S3 engineers are now working to make unauthorized requests that customers did not initiate free of charge.

But in the meantime, don't change your phone number to 8 6 7 5 3 0 9, S3 bucket name is something that everybody's going to potentially try. Now, perhaps the only other thing that we didn't see much data on in this report is the use of generative AI enabled social engineering. What do you mean by that?

Say your CEO gave the presentation. It was recorded on YouTube or CNN. That's great content. Reflects well on the company. Positive Outlook perhaps.

[00:20:00] However, what bad actors are doing is taking this content and making deep fakes of the CEO. They'll take that YouTube video, use a site like Speechify, where they can just type the language that they want the CEO to say, it trains on the voice and out comes a really close voice.

There's also more sophisticated sites that will do just that. Video, as well as just audio. But if you're just doing audio, you call a procurement team selling in the CEO's voice. You need to change the banking account number on one of our suppliers to a new account. We're just talking to him today. CEO calls, tells you to do something, you do it, right?

this modern play on a business email compromise is going to be common. In fact, a multinational company who least recently experienced a multi million dollar attack where the bad actor created multiple deepfake video calls of the CFO and other colleagues. They're all fake! But they convinced an employee to send 25 million U. S. dollars to five different bank accounts across 15 transactions. that's not, as you would say, a [00:21:00] career enhancing move. Deepfakes are also being used to harm brand reputation. For example, celebrities have been faced with. Reputational damage associated with deepfake porn videos made with their likeness, but not a video isn't the only reputation threat What if attackers create and release a deepfake video of your company's CEO?

Making pre release statements on the economic outlook of the company and that panicked the stock market into a major sell off. The attackers knew that so of course they shorted your stock and they profited from that now the SEC is going to be involved in that Your company has to respond to that. Your investors need to respond to that.

The press is going to want to respond to that. And it doesn't get fun. How about a deepfake when your CEO makes racist or sexist statements they would never do normally? How do you investigate to prove it's fake or real? And would that even matter? Once the reputational damage has been done. These are the things a lot of CISOs are going to have to be thinking about in more detail in the next few months.

It's a little bit early in [00:22:00] the election season, but I've heard it said that 2024, maybe the last year that humans get to decide who wins an election. We're getting so good at deep fakes and so good at AI and manipulating people's perceptions that unless you have extremely well honed critical thinking skills, you're And you can spot the phony.

The average person is going to go ahead and get led along to be able to believe something that they should not. Now there's a ton of other fascinating facts and figures in this DBIR and I encourage you to download it at no cost. And at least peruse a number of the sections even if you don't have time for the 100 pages.

And note that you'll encounter a squeeze page where you are asked to provide contact information in exchange for document access. Thanks. I'm usually okay with that. I found out that if you, I put in as an individual in this case, and they just gave me the PDF, they didn't send me an email link. So maybe that's a feature if you're not a company.

As the past years, the authors [00:23:00] of the DBIR hid some puzzles that when solved first could result in a nice reward. Now, if you haven't looked for it yet in this year's report, let me suggest you listen to Warren Buffett's advice on reading annual reports. Start with the footnotes. Let me know if you have any flashbacks to the subtitles for Monty Python and the Holy Grail.

A moose once bit my sister. Okay. Anyway, we're going to have a shorter episode this week, because if you're at RSA, you're busy going to the parties. And if you're not at RSA, I didn't go this year either. Last time I went, I was speaking and there's a ton of things going on. Didn't necessarily get the return on investment I was hoping for.

And a lot of people go there for the tchotchkes that you get off the Floor for the exhibitors or for the parties and the networking presentations are of course good. Meanwhile, if you're not there, good luck with your CISO Tradecraft journey. And if you are, hope you had a wonderful time out there in San Francisco.

And we hope you've enjoyed listening to us talk about the 2024 [00:24:00] report from Verizon on data breach investigations. So thanks again for listening. And if you found the show interesting, do us a favor, invite someone you know to follow the CISO Tradecraft Podcast. If you're not watching us yet on YouTube, take a look and subscribe.

Getting that number up is going to help us get rid of those ads that show up in the middle that we don't control. But once we get to a certain level, so they said that we understand why people are kind of grade grubbing for likes and subscribe on YouTube. You get control of the channel. So until next time, this is your host, G Mark Hardy.

Appreciate you tuning in and get smarter about cyber. Keep your system safe and stay safe out there.