#164 - The 7 Lies in Cyber

[00:00:00] Outro

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. Today, our podcast is about seven lies that we hear in cybersecurity. We want you to know what they are, why they might be lies, and how you can carry on despite knowing they exist.

Now, these are things that we have come up with and produced and They're not the official endorsement of any organization, so don't go looking for these in some other publication. As always, if you're following us on LinkedIn, great, and if not, you're missing out on an awful lot because we put a lot more information out there than just podcasts.

If you're following us on YouTube or other Platforms like that. Great. Let other people know where you can find your information. So let's get rolling. The first slide we see comes from the documentation that comes [00:01:00] out of the Center for Internet Security, or CIS. Now, CIS has done some awesome work, and we've got the critical controls, which are control sets that we look at.

which have kind of evolved quite a bit over the years. Well, according to these, you have to kind of do them in order. And CIS says we have an accurate inventory as a requirement before we can move to get to basic cybersecurity controls. And we see this in the guidance that comes out. Control one and control two say you need an inventory and control of all the hardware and software assets.

And it also says you should focus on basic categories like accurate inventories, vulnerability management, control of admin privileges, secure configuration of hardware, maintenance, monitoring, analysis of logs, before you really focus on organizational and foundational control. So, so let's get this right.

Before you apply any email or web browser protection, which is control number seven, I need to get an accurate inventory. Here's the rub, you'll never have a complete and accurate inventory to be able to move past controls 1 and 2. Now, again, I'm [00:02:00] interpreting this a bit too literally, because everybody's going to say, yeah, not really, but I've had more than one person come up and say, you've got to do them in order.

And so that becomes part of, well, the deadly lie. Here's an example to illustrate. Let's say your organization wants an accurate inventory of every application used by the company. The plan is to record each application into a centralized configuration management database, or CMDB, to provide an accurate listing of all the enterprise applications.

And then this CMDB will require developers to provide important metadata fields to denote which applications contain PII data, who the owners of the applications are, and then check boxes to determine if the application is in scope for additional laws and regulations, such as PCI or HIPAA. It seems doable, right?

The first issue is most organizations can't even define a shared understanding of what an application is. Let's look at the definition of an application from NIST. According to NIST, an application is a hardware and software system implemented to satisfy a particular set of requirements. Now in [00:03:00] this context, an application incorporates a system used to satisfy a subset of requirements related to the verification or identification of an end user's identity.

So that the end user identifier can be used to facilitate the end user's interaction with the system. Okay, great. Now we know what we can do for identity management. The definition is kind of vague. It doesn't really leave you a lot of technical guidance for developers to write a simple search query to find every application.

Is code the application? Does the application need to be on the website? What if there is no user identity captured? Does that mean I don't have an application? Uh, there's a lot of unanswered questions. So consider the example of a PowerShell script. Can they access PII data on your desktops? You bet. Can they identify user identities?

Yeah, you know that. Are developers in most organizations creating application entries for each PowerShell script and putting them into a CMDB today? Probably not. Now, what about web apps? How do you define a web [00:04:00] application? Is every file in an AWS S3 bucket, a different web application? What about every folder in an S3 bucket?

What about every Microsoft SharePoint site or a Jira Kanban board or a wiki page? See every Microsoft or Google form page may be used to collect input. You see, it's hard to draw the line and define what an application is and track every application. So please remember, well, the goal of having an accurate inventory is noble.

It's never going to be complete. Now another reason why having an accurate inventories will never be met is meeting business objectives always trumps having well, accurate inventories. Let's say the business needs to send its auditors, regulators, sensitive, Company data to pass an inquiry. They might send the data over email.

They might share the data via a Dropbox account because the files are too big to email. They might go to the auditor's SharePoint site and upload the data directly. Do employees in your organization understand which of these examples fall under the definition of an Application. Do you think they'll spend the [00:05:00] time to fill out your CMDB?

And even if you could get them to register their application, have you told them to ensure each application meets every general IT control required by your organization's security policies? Okay, now they know that they have all this overhead, are they still going to put it in? And how can you verify you didn't miss one app?

So having a full complete inventory of your hardware or software before you do anything else under critical controls, kind of a lie. Second lie we tell ourselves. comes from ISO and FAIR. They both claim that we in cybersecurity can perform an accurate risk assessment and perform cyber risk quantification.

Now, remember, one of the major focuses of ISO 27005 is to perform risk management. That is, if we can correctly perform risk identification, risk analysis, risk evaluation, and risk treatment, we're doing risk management. Well, we believe risk assessments are helpful. But we think they're also flawed. You see, it's highly unlikely you'll find [00:06:00] every threat or have very accurate likelihoods and impacts for every cybersecurity risk assessment.

Take, for example, third party risk assessment process. Sales teams just inform you that they want to use the Acme company as their CRM tool to host customer data and send messages via a DocuSign like functionality. System will contain PII data from customer contracts, payment card information, PCI, bank account data, as well as a lot of other sensitive data.

And they've asked you to perform a third party risk assessment to identify all the cyber risks before they sign a contract with Acme. Now, the cyber team uses a DAST tool to find if there are any cross site scripting or SQL injection vulnerabilities on the site. Luckily, we find zero. You look at the security policies from Acme.

They're good as well. The company had a pen test from a company that you never heard of, but this pen test came back clean and had zero findings. Now, however, Acme has yet to perform a SOC 2 Type 2 or an ISO 27001 [00:07:00] certification. Tell me, what is the likelihood that Acme may be breached in the next year?

Can we truthfully say we've got a 70 percent chance of it breach happening and it will cost 273, 512? But we just got a SOC2 Type 2 report that the likelihood goes down by 40%? See, we don't have any real world evidence that is accurate to support these very specific claims. Anyone who tells you it's a fair analysis to ballpark is kind of selling you snake oil.

Also, the value of a pen test is entirely dependent upon who performs it. Your results of a pen test are how well did you do against a particular team on a particular day with their particular tools? It's a point in time. If it were the same all the time, it would be like sports games. Why bother to go ahead and have all these games in the season?

Let all the teams play each other once. That's it. We figure it out. Nothing changes. They'll always win. They'll always lose. Season's over. No, things are different than every day. And so what we want to find [00:08:00] out is a pen test might be more of a check the box exercise via a true validation of risk, especially if it's a company you haven't heard of before.

Risk assessments can identify things they've found, but they can never identify the things that we missed. You know, this unknown cyber risk means that while risk assessments are helpful, they're never going to be perfect. And finally, if the vendor doesn't fill out your third party risk assessment, you have a lot of unknowns.

And even if they did, if a third party audit was never performed, such as the case of no SOC 2 or an ISO 27001 CERT, then you really don't even know if the vendor was lying to you on their questionnaire. Lie 2. Lie three. Third lie, we tell ourselves, comes from DevSecOps. Man, I'm picking on everybody today, aren't I?

We just need to focus all of our security efforts on shifting left. Well, let me tell you something. If Barry Sanders and Bo Jackson only shifted left in football, they wouldn't have been as impressive. [00:09:00] Although you got to figure out that it does work with Josh Allen from the Buffalo Bills. He goes left almost every single time.

I think 23 quarterback sneaks and 21 of them worked. And they all go left! But Barry Sanders, Bo Jackson, they shift everywhere. We should do the same. Here's an example to illustrate a little bit further. Does everything in your production environment 100 percent mirror all the code that's actually in your GitHub repository?

Can you tell me that once you're in a large organization with Thousands of developers that no one's logged into a server and made changes that are not reflected in the code base. Can you tell us developers haven't modified an AWS setting without changing the Terraform file every single time? Even the interns?

Yeah, that's not practical. Let's take another example of a developer or system administrator misconfigured a proxy server setting from Being a reverse proxy to a forward proxy. Now, would anything in your DevSecOps tool chain catch that misconfiguration as a new server side request forgery. SSRF vulnerability [00:10:00] probably not 'cause.

Forward proxies and reverse proxies are both acceptable patterns when used in the right context. Now, since DevSecOps, security tools like SAST or DAST or SCA don't understand the context of your app to know if it's customer facing or employee facing, they don't know which configuration is appropriate, so they don't report the finding.

Now, if the developer never gets notified of the vulnerability, let's hope your pen test team finds it before an attacker does. Otherwise, you might face the same fate that Capital One did with their large breach. Now, one little tidbit that I've, uh, learned about by doing a little bit of research on the Capital One breach.

According to various reports, it looks like Capital One had a ModSec WAF proxy, which was, well, misconfigured from a reverse proxy to a Forward proxy, which then enabled a server side request forgery attack. This allowed the attacker the ability to forward requests to the [00:11:00] AWS EC2 metadata service. Thus, the attacker was able to gain access to the EC2 access key and the EC2 secret keys.

And once the attacker had these, the attacker was able to use the AWS command line interface to query files in the S3 buckets. Now, I don't understand why Amazon would ever allow a command line interface used by humans to impersonate a server or an EC2 instance. This seems like it kind of violates the concept of least privilege.

What's worse is when the attacker did this, the attacker also used the AWS command line interface from the newly created AWS Middle East region in Bahrain. Another insight learned is when Amazon created new regions, not all the AWS capabilities are stood up instantaneously at the same time. They're rolled out. So when AWS GuardDuty was not available in the new region at the time of the breach.

So Capital One might've had [00:12:00] DLP capabilities in AWS East or AWS West had no native AWS capabilities to detect attacks over AWS command line attacks in Bahrain. Now, when we look at this sophisticated attack, you would not have been able to stop it from a CICD test with a SAST. or Software Composition Analysis tools.

And additionally, if AWS had issues with least privilege, like what we see with them not limiting EC2 metadata access calls to only EC2 servers and not AWS command line interface, which we believe may still remain true today, you can go test that. You really need a more verbose cyber response strategy.

Now, note AWS has now added two capabilities. AWS added a second version of the EC2 metadata service and the capability to limit using new regions. Where not all features are accessible or even available. Just a side note, if you wanted to see the attack commands to do this, you can go to CloudGoat and look at the [00:13:00] AWS EC2 SSRF attack.

Take a look at our notes, show notes, we'll give you a reference. See, this is why it's helpful. To adopt breach and attack simulation tools, in addition to things that scan in your DevSecOps toolchain, you also need to build secure software training into required training programs for developers. Don't forget to add threat modeling into your architecture practices to identify contextual problems that are difficult to find.

This way, you can identify where fraud and abuse are likely to occur. Through multiple practices, you start shifting everywhere, not just shifting left. And also You should also trust tools that pull from your production environment more than what's actually in your code. All right. Fourth lie we tell ourselves comes from auditors.

That being attestations and certifications that mean an organization is secure. Now, I don't think any real cybersecurity executive would believe this, but there's a whole audit [00:14:00] industry that lies to us about this on a daily basis. Take, for example, a SOC 2 Type 2 report. SOC 2 assessment is a point in time assessment.

Documenting the ability for systems within an organization to meet security controls. Now, let's say an organization chooses to meet every control in NIST, special pub 800, TAC 53. They have a SOC 2, type 2 attestation, then surely they met every NIST control and have great security, right? Well, not so fast.

You see, the organization gets to choose Which 853 controls are tested by the auditor who may also pay to perform the audit. So crafty companies will typically choose the easiest controls to meet and say, auditors, Hey, look over here. While never having to provide evidence for the hard controls, which they may be unable to meet.

And also if you don't like your audit results, just buy a new audit from a different company that better aligns to your organizational view on things. Now that wasn't bad enough. You should really investigate what the standards require. [00:15:00] Are your organizational policies based on ISO 27001 guidance? Did you know that standard doesn't require penetration testing?

Control Objective 812. 6 states, Quote, Information about technical vulnerabilities of information systems being used shall be obtained in a timely fashion. The organization's exposure to such vulnerabilities and appropriate measures taken to address the associate risk. Okay, so. If I run an NMAP scan every two or three weeks and fix the findings, I've satisfied the control.

I might even show a Qualys report, which tracks my vulnerability burndown each month, if I want to get fancy. There you go auditors, I've met the objective. Now, leave me alone, no need for a pen test, right? When really investigate what the controls and policies require, you're going to see gaps. And if you're smart, consider adding additional controls and requirements.

To fill those gaps. Multi-factor authentication, MFA. It's more of a nice to have than a true requirement in both SOC two [00:16:00] and ISO 27,001. I kind of insist on it in my enterprises for everything but an ISO 27,001 a 9 42 use of privilege utility programs encourages limiting access to privilege utilities by using authentication techniques like MFA. A 9 44 use of secret authentication information.

emphasizes the importance of ensuring that secret authentication is managed securely, which could include MFA. So, I could include it, but I'm kind of too busy, so I'm going to skip it and just get the certification anyway. Um, thanks, but now we see why audits and certifications may not be worth as much as you think.

Now, one final thought on this fourth risk. Do you think cyber standards from 2017 understood the threat of ransomware attacks in 2023 or 2024? Probably not, yet it's common to find policies in use today by large organizations that were written in 2017. You'd be surprised how many people use, [00:17:00] and companies use, ISO 27001 2013.

Might be better to align with the framework like the MITRE ATT& CK framework, which keeps up to date with real cyber attack methodologies. And finally, Many of our controls don't have metrics. For example, most standards or certifications will check if you have a vulnerability management program. Okay, that makes sense, but they don't hold a measurement to say you're patching in 15 days or 15 months.

Now, if we have a program, but we only patch internet facing vulnerabilities in 15 months, we might be certified, but we're not gonna be secure. Another example on this topic is tooling. You look at the most important safeguards we have, like email security gateways to stop phishing attacks, antivirus and EDR to prevent malware, data loss prevention tools to stop data exfil attacks, web app firewalls to stop web application attacks.

Now, look at the standards. and see if they tell you what percentage of email is protected by your email gateways. What percentage of your endpoints have the [00:18:00] latest EDR agent on it? What percentage of applications sit behind the WAF, etc.? There's zero guidance on these types of metrics in the standards, which means you don't get audited on it.

Thus, audits become checking the box exercises, which lie to our executives saying that we have good security, while masking real cybersecurity issues that never get publicized. The fifth lie we see in the cyber industry is one we tell the regulators that we'll report all cyber incidents in 72 hours. Or, if we're under the new guidance from SEC, four business days after we've declared that it's actually been an incident.

Here's the problem. Sometimes it takes way more than 72 hours to investigate a breach, like most times. Example, let's say one day your pen test team informs you there's a major vulnerability on your customer facing website that, if exercised, could allow customers to see other customers PII.

Unfortunately, after talking to the developers about it, you learn the logging system on that [00:19:00] application was terrible. So you're not able to look through the logs to determine Bad actors have already exploited this

vulnerability. Do you need to disclose there's been a data breach within 72 hours? Some lawyers might say yes.

Others might say we can't show any material impact or harm has occurred, so we have no requirement to disclose. All right, here's a different example. Let's say the help desk finds a ransomware file on a laptop and informs the incident response team. The IR team opens up a cyber investigation. They start looking around and successfully determine there was a bad actor on the network. Now. What if you could only see encrypted SSH sessions, but you couldn't determine what the bad actor did, how long should you look and how long should you take to identify what files were stolen? It's all encrypted. All your telemetry is encrypted. How long should you keep looking to identify where the attacker went?

In your network. You see, 72 hour investigation may have very different results than a [00:20:00] three week investigation, or maybe one that takes a month or longer. You might find out that while the attacker got to a lot of places in the network, they were never really able to exfiltrate any sensitive data. And other times, it's just the opposite, and you learn the breach was a whole lot worse than you initially thought.

Since either outcome is likely, some organizations are going to use legal language loopholes to lie. A cyber incident can only be determined after a cyber investigation has concluded. So if the cyber investigation takes three weeks, then you don't really need to follow the three day rule or the four day rule for SEC.

Be very careful with this one. There's a lot of guidance that could quickly bite you in this lie, could make you as a CISO land you in the courtroom, or possibly even some jail time. The sixth lie we see in the cybersecurity space is one we tell the CIO and her or his direct reports. We tell them that our application security tools are accurate and developers Just need to fix the vulnerabilities.

See, the truth is our [00:21:00] SAST tools are just not that accurate. In 2015, OWASP created something called the OWASP Benchmark. And the OWASP Benchmark study took a software application with just over 21, 000 use cases to test software vulnerability identification by open source and commercial SAST and DAST tools.

And this assessment allowed OWASP to benchmark How accurate SAST and DAST tools were in finding real software vulnerabilities. The result of that

study found that commercial SAST tools could only find 54 percent of the true vulnerabilities in the applications. That's a lot of vulnerabilities to go.

Unnoticed. Now, what's worse is when SAST found vulnerabilities in code, they have an average false positive rate of 26%. This means developers go on wild goose chases a quarter of the time that they're investigating vulnerabilities. Gee whiz! Can we just be honest with the CIO and said These tools are far from perfect.

They got a high rate of false positives. Yeah, we know we're missing vulnerabilities and chasing rabbit holes is [00:22:00] terrible, but we promise it'll be worse if we don't go look for it. Just depend on it. All right. Hopefully not having a thing too depressing, but number seven, the last lie is we tell our own cyber employees.that being we don't want to be viewed as a cost center. Now some of you might be thinking, that's true. Well, where's the lie? The lie is we tell ourselves and everyone else that being a cost center is a bad thing, and we don't want to be viewed that way. The truth is being a cost center can be a good thing, as well as a bad thing.

We overemphasize the bad because we fear losing our cybersecurity budget, right? However, if we consider the positive, we might just see how being a cost center could increase help. Both cyber and also help the organization. For example, cost centers can monitor effectiveness more effectively than profit centers.

See, cyber departments that take a cost center approach and provide executive briefings which show how the effectiveness of the cyber program [00:23:00] improves will create an engaging conversation with executives and you'll know that in exchange for this amount of investment, you're getting this coming back.

Profit center, you're always trying to bring in more and more and more and we're not going to generate excess profit in a cybersecurity organization.. Also, cost centers try to update processes to be more effective and save money so that they can reduce expenses. People love that. Bean counters love that. Cost centers usually try to cover all their costs with offsetting revenue by reducing expenses and preventing losses.

What if cyber showed the value by showing the cost of each attack the company prevented last year and the ROI of expenditures that were used to stop the attack? For example, last year, the XYZ company spent 250, 000 on an email security solution, and this solution stopped 10, 000 phishing emails from entering employee inboxes.

On note, last year, the average cost of a successful breach of a business email compromise, valued by the FBI and Larry Poneman's Institute, It'd be about 6, 000, 000 and we [00:24:00] think that spending 250, 000 to stop a potential damage of 6, 000, 000 is a really good investment to protect our hard earned revenues.

We might also highlight that we taught our employees to report phishing attacks by clicking a button in Outlook or to flag the phishing emails to the Incident Response Team. This capability allowed our organization to stop over 5, 000 confirmed phishing attacks. We view this as one of our best safeguards to protect the company.

So the next time you hear that you don't want to be a cost center, think about the benefits. Okay, well, thanks again for tuning into today's show. We hope you've appreciated hearing about and maybe agreeing or disagreeing with the seven lies. And if you don't like them, Let's go talk, you know, come on, talk to us on CISO Tradecraft at LinkedIn or engage us.

But let's recap. Number one, the first lie is you need an accurate inventory before you can move on. The truth is you'll never have an accurate inventory. So do the best you can and then move on. Do your other controls. The second lie comes from places like FAIR. Where we can do accurate [00:25:00] cyber risk quantification, where it's true that you can quantify some things if you have the data.

Mostly we have a vague understanding of likelihood times, a vague understanding of the true impact to a system, which means a vague squared measure of risk. If you don't believe it. Look at the risk estimates from your third party risk program and see how often those end up being realistic. Third lie is we need to just shift left and fix all the vulnerabilities before they go out.

The truth is your production environment doesn't always look like your code. So you need to shift everywhere. Also, there's a lot of vulnerabilities you won't find in your tools. So focus on attack surface mapping, threat modeling on your production, because that's where you're most likely to get attacked.

Fourth item is certifications show security. Absolutely not. They show evidence you met a control, but that's only as good as what the control requirement requires. [00:26:00] See, we see examples where they don't really require MFA or patching an X number of days or percentage of endpoints to run EDR. But the audit industry has convinced us to pay for things that aren't as helpful to stopping the real world attacks that we need.

Understand this lie and be sure to meet legal compliance requirements. However, once you do that, Go focus on real world security mechanisms that stop the attackers. Our fifth lie is where you report cyber incidents in 72 hours or four working days for the under SEC. Remember, you don't know about an incident until after some investigation work has been done.

You have some suspicions, you have some rough stuff. But remember what George Patton had said, The first report is always wrong. I found over the years as a military officer and also in cyber, it's almost guaranteed to be true. The first report's going to be wrong. Someone's rushing to get you information, it's incomplete data, you jump to conclusions.

You need time. Okay? So understand this. Talk to your lawyers, get on the [00:27:00] same page, get stuff in writing, report as soon as you see something. So you don't end up with the hot water of the SEC, but make sure as you're finding out with some of the issues of these 8-Ks, do they contain the information? They're not going to be comprehensive.

They're not going to have all the data, but they do meet the reporting time deadline. So you've got to get something out fast. So do what you can as you could, but then follow up and trace it to ground. Our sixth lie is AppSec tools are accurate. We saw they often have false positives. So let's just be truthful and fix everything by also giving developers a chance to say, these are false positives.

So it should be removed from the cyber metrics that penalize them. Nothing is worse than getting dinged for something you didn't do wrong. False positives do that. Our last lie focused on. Cyber not being a cost center. Remember, cost centers get money too. There's some good benefits with cost centers, so don't forget to use them in your budgeting exercises when you show the value of each cost being [00:28:00] expensed.

Okay, how about that one? So this has been this week's episode of CISO Tradecraft. Thanks for stopping by and listening to our podcast. We're glad we could provide some education on the seven lies we see in cybersecurity. And if you enjoyed our show, give us some feedback. Give us hopefully some comments on a podcast platform that you subscribe to, or go ahead and follow us on the YouTube channel so that you can go ahead and see what's coming out.

If you share it on LinkedIn, that's wonderful. We put out a lot more than podcasts. We got a very high signal, Low noise stream that we put out on LinkedIn. Let's get that out to help out as many security leaders as we possibly

can And so we also have a newsletter on Substack so we can put out we're putting on newsletter now So if you'd like to get an email that highlights key lessons learned from our weekly podcast Subscribe and we get that to you.

Go to Substack, search for CISO Tradecraft to find it and you will. So that's it for now Until next time, I'm your host G Mark Hardy. [00:29:00] Thanks for listening. Stay safe out there.