#195 - Pentesting for Readiness not Compliance (with Snehal Antani)

[00:00:00]

Begin

[00:00:12] **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 going to be your host for today, and we're going to be talking about how offense informs the defense, and more importantly, some of the technical aspects of that and how to implement that in your environment.

And to check it. And with us, we have a special guest today, Snehal Antani. Welcome very much to the show.

[00:00:39] **Snehal Antani:** Welcome. Thank you for letting me be here.

[00:00:41] **G Mark Hardy:** So let me do a quick plug for CISO Tradecraft. This is your first time to the show. Don't forget to subscribe to us. We've got a lot more than podcasts. Please join us on our LinkedIn channel. If you're watching us on YouTube, give us a thumbs up or or on the podcast channel.

Let everybody else know how much you enjoy it, because it's very helpful to us in getting the message out to others to help them in their [00:01:00] cybersecurity career. But back onto our show here, we want to talk about the concept of offense informs the defense. And when I was teaching at SANS for a decade, we used to say that a lot.

And the idea was, is that if you're a defender, you learn more about how to protect your enterprise. By observing your attacker's behavior and adapting accordingly, rather than, for example, reading a manual or reading a book or just following the configuration checklist, because quite honestly, the attackers don't care if you've met your compliance requirements, they don't care if you're reading your manual, and they don't even care what your score is.

If they get in. So what are your thoughts in terms of that? first of all, how about a quick introduction and tell a little bit about yourself and how you got to where you are now with this really cool company that you're running?

[00:01:43] **Snehal Antani:** Yeah. once again, thank you for the opportunity to speak here. So my background, I'm a software engineer by education and trade, did my undergrad at Purdue. My freshman year, I saw Drew Brees play Tom Brady, which is always a fun experience for a college kid, and then started my career at IBM in the web sphere and middleware space and started on the [00:02:00] mainframe.

learned quite a bit about enterprise computing in that world. And then I was a CIO within GE Capital, which was an incredible experience to learn how to actually utilize and adopt enterprise technology to solve business problems. and from there became the CTO of Splunk to help grow and scale that company.

And in 2017, I chose to take a break from industry and served as the first Chief Technology Officer within US Special Operations. and for those that, for those that have served, thank you. For those that have not served, I think I grew more personally and professionally in that role than in anything else I've ever done.

and it's just an amazing experience all around. And that's where I met my co founder, Tony. And when he retired from the air force, he and I started Horizon 3.

Quote

[00:02:44] **Snehal Antani:** And the problem it was solved is in all my jobs, I had no idea I was secure until the bad guys showed up. Are we fixing the right vulnerabilities?

Are we logging the right data in Splunk? Does my team know how to respond to a breach? And the answer is, I don't know, either wait to be hacked [00:03:00] or hire consultants to show up once a year, poke me in the eye, tell me how bad we are, and then disappear, leaving me holding the bag.

[00:03:08] **G Mark Hardy:** Yeah. And if we think about that, typically when the consultants come in, and I've got to admit, I'm guilty of that because I run a consulting business, is that as a pen tester, you come in and say, okay, we're going to roll up our sleeves and let's see what we can do. The reality check is, and it's not criticizing pen testing, but a couple observations on it.

Number one is, that traditional pen testing is going to be a single point in time. And how well are you doing against that pen tester with the tools in his or her toolbox? And basically, how well are they doing that day? They might have an off day. It might not do everything. So to a certain extent, it's not that complete.

The other thing is if we look at different frameworks, such as the CIS critical controls, where you have now 18, we used to have 20, but they're prioritized. What's dead last? Pen testing. Now, it doesn't mean that pen testing is not important, but the construct of those who are writing it saying, You set up all your chess pieces [00:04:00] on the board, you get everything ready to go, and then at that point we'll call in the pen tester to see, can we find a problem?

But the thing is that works great for an absolutely static environment. But what happens if things change? And things change on a regular basis. You add users, you add applications, you change systems, you reconnect. New vulnerabilities are discovered. it seems to me that the idea of pen testing, some people have kind of said, yeah, it's kind of out of date.

I will disagree with that, but I think the way we do things, which is a single point in time, we line up like a military inspection, you suck in your gut, you look straight forward, sergeant walks past, and then you let out your breath, okay, fine, I survived this round, but there's got to be something more to it than that.

[00:04:43] **Snehal Antani:** Yeah. So during my time within special operations, I was fortunate to have this incredibly forward thinking commander, JSOC. And he and I talked about three things. The first thing he said was, don't tell me we're secure, show me, and then show me again tomorrow and then show [00:05:00] me again next week. Because our environment is constantly changing.

And the adversary always has a vote. And that was, an amazing philosophy to have from the leader of the organization and then driving or permeating that throughout, all aspects of what we were doing. The second key thing was, we started talking about this distinction of being secure versus being defensible.

Being secure is a point in time state. Being defensible is about rapidly adapting to the actions of the adversary. And what we cared about were. How defensible we were as an organization. You have to assume the enemy is going to get in. You have to assume they're going to assume breach and gain initial access.

How well can you defend against that? How can you rapidly adapt to stifle actions and prevent the attacker from achieving their objectives? And then the third thing we often talked about was this distinction of between being compliant versus being ready. In the military, you always hear them talk about readiness.[00:06:00]

and readiness is all about your preparedness and ability to actually respond, to some sort of action. And once again, compliance was this point in time state and point in time mindset and readiness was this, continuous objective you strive to achieve. And I think the real crux or commonality between those three principles he and I would discuss was about continuous vs static to your point

[00:06:24] **G Mark Hardy:** Yeah, and when you think about that, the traditional approach to pen testing has been, in my observation, compliance requirements. Oh, I need to do that to meet this particular requirement, or I have a third party vendor who had said, Hey, before I'll do business with you, show me the results of your pen test.

And it becomes a check in the box exercise. Here we go. I've got other organizations that I have worked as a CISO for a while. And I say, you know what? A traditional pen test, in my opinion, is not that valuable. And it sounds like a bit of heresy coming from a professional security practitioner. But to a large extent, I think we behaved ourselves into a [00:07:00] corner that suggests that the traditional approach, compliance driven, is really not adding a whole lot of value.

You're right. It doesn't give us the readiness because for the most part, it's a point in time and it's not continuous. It also brings up an interesting question and something that we hear about a lot in the news of the shortage. of qualified security professionals, in particular, for example, good pen testers.

Now, for example, I'll do a kind of shout out to my son who is a pen tester and runs a pen testing team for a major auditing company. And, he's got his GPen, his GWAPT, his SANS certifications and a couple of certs in hand and some job experience under his belt. And he's gone out there and done amazing stuff.

And yet he's kind of chosen to be the red team guy. I'm the blue team guy. He said, I'll build the defenses. You try to break them down. But as we look then at what you're describing, which is a readiness based approach rather than a compliance based approach that seems to be superior. But then how do I engage my [00:08:00] traditional pen testing people or vendors or consultants, or do I not do that anymore?

Is this kind of the death knell of that particular line of business or just maturing into a new level and these organizations will adapt as well?

[00:08:17] **Snehal Antani:** there's a lot to unpack in there, so let's kind of start with forward thinking procurement teams, and I did this, and I kind of learned

and started to ask these questions when I was on the other side of the table. I actually don't care about your most recent pen test results. I care about the diff between your last two pen test results.

I want to know how many problems did you find in pen test result one? How quickly did you fix them? How often did they reoccur and why? And how effective were you at detecting and responding? And I should be able to run it. I should see the comparison of results between your previous two pen tests and understand those answers.

Because that's what's most important. How quickly are you fixing problems? [00:09:00] And how often are they reoccurring and why? Because if you fix them quickly, but you have high reoccurrence rates, you've got some fundamental shortcomings in your architecture and automation and talent. And that's super important.

That carries a greater risk than if you had a clean result, because you could have prepped for that clean result and it was a point in time state. But, your remediation time and reoccurrence rates between pen test results actually underpin your philosophy of security. And that you care about security by design and you're securing yourself.

I'll pause there. I'd love to get your feedback on that specific piece of it. And I'd love to talk about talent and then kind of the future, but does that make sense when I talk about the diff between your last two results?

[00:09:40] **G Mark Hardy:** It does. So let me say it back again. It's always the best way to make sure if you're in a two way conversation that you're hearing it, is that just simply taking a snapshot in time is insufficient, but comparing not only the differences between the last two pen tests, but the actions the organization took, both in terms of the completeness of remediation and the time to effect [00:10:00] remediation are important because, as you'd indicated, if it took Six months to get that last thing done before the next pen test is compared to, wow, everybody went turning and burning and they got it done in six days or probably not six hours, but it tells you the responsiveness of the organization.

And also if each pen test are identical, and that's one of the dangers I think of a traditional pen test is that if you bring the same tester with the same toolbox year after year you're measuring up to that exact same standard and you could score perfectly on that. But it's not necessarily a practical exam because the attackers are going to have a different set of tools.

They're going to have a different approach, and they're going to get past these essentially static defenses that you've built like a Maginot line to say, Hey, we know that our pen tester uses this in his or her toolbox. We're going to defend against this in his or her toolbox. And so I would think that having that diff is essential.

It's one of those things that's necessary, but maybe not even sufficient. You need more than [00:11:00] that.

[00:11:00] **Snehal Antani:** Well the next part of it is the concept of sample testing, which has kind of permeated through the pen test community and sample testing is in part, a function of labor shortage. So there are about 5, 000 OSCP certified ethical hackers within the United States, about 20 25, 000 globally. It takes 10 years to really become a senior penetration tester.

And so what we have is a massive shortage of pen testing talent that is, that has at least the base level of certification. Because there's a constrained amount of supply. And there's a significant increase in demand. PCI now expects you to be doing quarterly pen tests. DORA, GDPR, NIST 2, so on. All are starting to push quarterly pen tests or more.

And to be honest, for every Patch Tuesday, you should be running a pen test Wednesday. That's the philosophy and the mindset that you should be operating with. Every time you onboard new employees, modify applications or modify your infrastructure, or make any sort of major change, [00:12:00] your attack surface has changed significantly.

And you need to assess the security posture as often as you're making major changes in your environment. So that's kind of the first thing. Your chain, your free, your pen testing, rate should be equal to or greater than your, rate of change in the environment. The next thing too, though, is sample testing.

If you're doing infrastructure and if you've got, 5, 000 hosts or 10, 000 hosts or 100, 000 hosts, You can't afford a pen test, do an infrastructure test against 100, 000 hosts. That'll take a year and millions of dollars. And so you end up sample testing a small little section. attackers don't sample test.

Once they get in, they take the time to find every avenue of attack across your entire environment. the lack of talent and the inability to test infrastructure at scale creates a false sense of security. And so what you want to be doing is, Testing yourself at scale, your entire environment as often as it [00:13:00] changes.

And what you end up running into is a human bottleneck problem because there aren't enough pen testers and it's too cost prohibitive.

[00:13:08] **G Mark Hardy:** And that then becomes essentially where a lot of arguments stop at that point. We just say, you've given me an impossible requirement, either financially or from an HR, it can't be done. So we're going to accept the risk of saying, I'm going to check 5%, 8%, 10%. And Oh, by the way, I can only do it on a periodic basis because I think our fundamental model of how we look at pen testing is broken.

And it's back based upon this whole compliance thing again. As you mentioned a number of compliance frameworks we're saying you need to do this. why are we doing it? Pen testing in and of itself does not involve somebody coming in there twisting the knobs and improving things. It's not like the pen testing fairy comes in on a Sunday night and corrects everything for you.

They basically come out and give you a report. And as a result, The idea is people don't like to be put on report. That's not [00:14:00] fun. There's a cost associated with it, which is driven by the scarcity of resources, i. e., skilled and qualified pen testers, as well as just the market value of that particular service and the external drivers, such as the compliance frameworks.

And so what we find then is that being able to come up with a solution using the traditional approach just doesn't work. And as a result, we have to then say, okay, if I'm a CISO, if I'm in charge of managing the risk or recommending to my senior leadership, what risks to accept, do I just keep on going and say this problem is too difficult to solve?

Or do I start to look at some of the directions we're seeing now? So we can say, Hey, I can automate things. And automation is one thing, but there's also autonomous, which is a little bit different. So if we say, Hey, could I automate a pen test or do autonomous pen testing? But how would I explore that?

How would I look into that and say, is this even feasible?

[00:14:58] **Snehal Antani:** So there's a great double click in [00:15:00] two dimensions. First and foremost, split pen testing up into application pen testing, which is looking for logic flaws in custom code, and then infrastructure pen testing. Humans are uniquely gifted at finding logic flaws in custom code. And it's very difficult for algorithms to do that.

So from a pen testing futures, I believe humans will be the scalpel focused on, Finding flaws in custom code and focused on very bespoke types of systems and environments. And I think that, what we've seen is algorithms are uniquely gifted at pen testing infrastructure at scale. Because under the covers, that's actually a graph analytics problem.

So when we run a pen test against 100, 000 hosts and we're able to scale up to pen test hundreds of thousands of posts in a single run. and this isn't like a cheap vuln scan. This is finding ways to harvest credentials, exploit misconfigurations, find misconfigured defensive tools, exploitable CVEs.

[00:16:00] And it's a combinatorics problem of how you combine these together in different ways. to achieve an objective like domain compromise or sensitive data exposure and so on. So algorithms are uniquely, capable of pen testing infrastructure at scale. Now, if you take that a step further, automated pen testing, in my opinion, is an impossible task.

And this is why this idea of autonomous, is the, avenue I pursue and the analogy is to think about when IBM played Gary Kasparov in chess a bunch of years ago, IBM did not hard code 40 million chess games. They looked at every piece on the board to determine the next best action. Similarly, trying to automate pen testing is attempting to hard code every possible permutation as runbooks.

And the environment is just changing too often and attack tactics are changing too often for that to even be feasible, which is why there's an entire graveyard of automated pen testing companies that are out there [00:17:00] formerly known or currently known as Breach Attack Simulation. Autonomous is about knowing nothing about the environment, conducting recon and enumeration to discover everything that's network reachable, and executing self directed actions to achieve goals like domain compromise, sensitive data exposure.

And that was the leap of faith that I took when we started Horizon 3, was could we build an autonomous system That knows nothing about the environment that can execute self directed actions to achieve a compromise. And sure enough, we're now four and a half years in and we're able to, pioneer this idea of autonomous testing.

[00:17:37] **G Mark Hardy:** And that sounds rather fascinating. Of course, it begs the question, as a lot of us are thinking, because if we look at, Artificial intelligence and generative AI, which in the last 18 months or so has really gone from being just a little backwater area where a few people knew about it to

pretty much everybody out there is playing around with something like a ChatGPT.

But it then sort of asks the question, [00:18:00] We've heard of attackers trying to go ahead and knock the guardrails off of some of these tools to say, Hey, can I use generative AI to create offensive attacks? And is that fundamentally what a autonomous pen testing is, just taking a generative AI, take off the guardrails and say, Hey, have at it bot, or is it a little bit more sophisticated than that?

[00:18:21] **Snehal Antani:** Yeah, it's actually a very different class of AI problem, to be honest with you. the best analogy for autonomous pen testing I think is full self driving and what we see with Tesla's. So first and foremost, there is no training data set available, for offensive attacks. Even the NSA doesn't, for as far as I know, they don't have an off a training data set specific to training offensive algorithms.

and there is no private data set unless you're building it yourself. And even for attackers, there is no massive data set that they could go off and pull unless they're collecting the data themselves, just like when Tesla [00:19:00] started to push self driving cars. There was no training data set available for them to build full self driving algorithms.

So what did they do? They put cameras in every car and they started to collect as much full motion video as possible, whether you paid for FSD or not, irrelevant. And they built this massive corpus of training data That allowed them to come up with Autopilot V2, and then FSD, and now SFD 12, you know, 12.

5, or whatever the latest version is. Because they've built the training data, and they're able to go off and build the algorithms to execute driving maneuvers. In offensive cyber, because there was no training data set available, the first thing we had to imagine Horizon 3 as being was a data company.

We're not a pen testing company. Pen testing is just a sensor. It is collecting a ridiculous amount of telemetry in every single pen test that's run. Anonymized scrub telemetry that allows us to build our corpus of training data so we can go off and improve our algorithms. And I think we're now at a point where we are collecting [00:20:00] a billion unique attack events a month that is proprietary training data for offensive cyber operations.

That data allows us to understand what defensive controls are most effective, least effective. What TTPs are most effective and least effective in achieving

our objectives, so on and so forth. So the first thing you've got to do is you think about autonomous pen testing. is accept that you're a data company first and architect your core technology accordingly.

And then you get into more advanced capabilities like shadowing. So how do I have, just like in Tesla, you've got two versions of autopilot running at the same time. The current version that actually has control over braking and swerving and so on. And the other with the new version algorithm, that's kind of a parallel testing to see how it would respond to those same conditions.

And so you need a whole bunch of infrastructure like that. So I don't think Gen AI is a good AI parallel to pen testing. autonomous driving is actually a much closer piece of it. and then if [00:21:00] I'll pause there and I'll talk about kind of Gen AI in general for pen testing and the, and a big fear I'd have as a CIO, but does that analogy make sense?

[00:21:08] **G Mark Hardy:** It does. And I think the idea that building that data company. I never really thought about that. It makes absolute sense. Of course, the question that I would have is that if you're going to be collecting and seeing all of these different actions per month, and you build them in this massive database, it almost sounds like you're not just there as a pen tester, but you're there as a massive sensor array to try to collect all this data, which then allows you to do a couple of things.

One is, we know what bad guys are up to, and maybe we're capturing threat Intel early on, which we could then federate. as long as it's anonymized across a number of customers. And then be able to then say, hey, an attack started over here, but we could kind of help you defend over here, but that still sounds like a defense role, rather than we often think of a pen testing role, which is a red team role.

So I'm getting the red and blue here [00:22:00] confused a little bit, and maybe some of our listeners are too. Could you help kind of align that for me?

[00:22:04] **Snehal Antani:** that's why we talk about how offense drives defense, right? The point of a pen test is not to check a compliance box. The point of the pen test is to assess and improve your readiness. And that's, that's at least, you know, my philosophy and to ensure that you are able to defend yourself. And so it's about this red and blue working together, to create a purple culture of proactively securing your environment and proactively improving your readiness.

I think the next part there that's super important is the more pen tests you run. So our customers end up shifting from running one or two pen tests a year to 40 or 50 pen tests a month. Constantly finding, fixing, and verifying. And here is actually the really insightful part. The bulk of our users are not security people.

They're IT admins and network engineers burdened with fixing problems that were found. [00:23:00] So the entire user experience is designed to enable those IT admins and network engineers to quickly fix and retest Issues that were found to be exploitable. And now think of all of that data that's being collected under the covers.

How many problems were found? How quickly did you fix them? How often did they reoccur and why? How effective were your security controls? And back to us being a data company. We have the more pen tests you run, the greater our resolution of understanding the change in results over time. And now you can start to use that to identify your self-identified findings reports, if you're a CIO or your compliance reports or your board reports, or to adjudicate, a lawsuit claiming you're not taking cybersecurity seriously.

here is your results over time. It's irrefutable. It's done by a third party. They had no ability to massage the sources of truth. And now you can show a clear trend line of reducing your exploitable attack surface, on a daily, weekly, monthly basis, which becomes difficult [00:24:00] to, claim that you're not taking security seriously.

[00:24:02] **G Mark Hardy:** That's a good point. We've brought up a term there that I want to dive into a little bit, and that is the attack surface. And traditionally, the attack surface is something we look at to say. Local to us if we were in a local environment, but now I have to extend the attack surface to the cloud because we have to say, that could come under attack.

But does it also make sense to try to extend the attack surface to suppliers and distributors? And if so, because it seems reasonable because of all the supply chain attacks that we've heard of recently, both in terms of software as well as just upstream from somebody who is going to go ahead and provide a part or a component for an assembly requirement.

How do we extend our model? to that larger ATT& CK service. if we don't have control over those entities.

[00:24:45] **Snehal Antani:** Yeah. so the bane of my existence as a CIO was having some of these third party risk management companies come in, and do

some cursory scan of my perimeter, make some [00:25:00] claim that my security isn't any good. Okay. And then, hold me hostage and pay them to improve my score. Drove me nuts. And it drove all of our suppliers nuts, you know, at Capital and elsewhere as well.

And the struggle with the way third party risk management is done today. And I'm not talking about source code third party. I'm talking about actual, like the 50 person ball bearings company that's supplying parts to Boeing or whatever else there might be there is those third party risk management scores are not based on any sources of truth.

They're mostly based on form fills that, the suppliers and the users are filling out. And so back to, I understand the change in results over time because we're actually attacking you. what we've seen our customers do is as they secure themselves, they make the capability available for their critical suppliers first.

So really think about the long tail of suppliers, that 50 person ball bearings company, that Or that hundred person, you know, [00:26:00] data company, whatever else, or design company, whatever else there might be, because that's who the attackers are going to target. When you look at, how SpaceX got the Falcon designs stolen a few years ago, it's because the attackers targeted Visser Aeronautics, a small design firm in Colorado, because they couldn't fend for themselves in the same way that SpaceX could.

And I think that there's a massive shift that's going to occur, or starting to occur. In how third party risk management is executed. we're leading the way we're actually working with a, a high profile government agency in the DoD, to help proactively secure the entire defense industrial base.

And what they focused on was the long tail of suppliers first, not the top 10. You know, the Raytheons, the Boeing, and so on, they're expected to secure themselves. But it was that 100 person, 50 person company, that sole source provider, critical to just in time logistics and lean manufacturing. Those are the people that are being pulled into this third party risk management program [00:27:00] where essentially procurement buys a pool of my license.

They make it available for the long tail to enroll into. We'll automatically enroll them, train them, enable them, onboard them, and get them to use continuous pen testing. And then those suppliers opt in to share their results with procurement. And procurement can now see at an aggregate level. the security posture and how it's changed over time across all suppliers, and on an individual

level, answer those questions of how quickly are you fixing problems and how often are they reoccurring?

And sources of truth are paramount to actually improving third party risk management in my opinion.

[00:27:38] **G Mark Hardy:** And that I think is rather profound because as we, some, many of us on this call may be aware of, or listening in, there's an issue called CMMC, the Cyber Security Maturity Model Certification, which is now in about its, Third year, it's been going for longer than that, but I've been in the program for three years and we've gone through some tectonic changes and things, but the fundamental [00:28:00] assertion is that the defense industrial base or the DIB is at risk from third party nation states and other actors who are saying, Hey, you're right. Let's not go after the Boeing secure enclave, or let's not try to steal the secret formula from Coca Cola by breaking into the safe in Atlanta. But let's go ahead to one of these second or third party suppliers, these smaller organizations that, as you said, can't afford things.

Now, one of the changes that took place in CMMC from the first version to today, version 2. 0, was the allowance now of self attestation for level one. And that was considered the initial problem in the first place, is that someone would say, you have to self attest that you are secure before you can get this contract.

You say, is anybody gonna double check? no. Okay, sure, I'll self attest that I ran a marathon last week. If no one's gonna double check, all right, you're in. And as a result, it created This environment where you really didn't have any fidelity. Now, as we move forward, [00:29:00] and they've gone back to the self attestation, because a manpower constraints, again, there's so many people that are available, all these tests that have to be done, the costs of them were prohibitive.

They said, okay, fine. maybe this is only be a few thousand bucks to do, but Oh, by the way, now the government is wielding their big stick and their big stick is to say, the false claim act says that if you make a self attestation and it's not true, we can come ahead and. and zap you in court for three times the amount of money that you got on the contract.

Ooh. but what you're saying here, if I can say it back is, that if I am a major supplier to the defense department and I have a number of smaller companies that are out there that are providing components or I'm sharing design data with or something like that, that I don't have to go ahead and start beating them with a stick to say, you need to get your act together.

It's almost sounds like I could extend this umbrella of protection that I've got to reach into these to say, Hey, You guys are not going to have to sweat the CMMC because we can help you find stuff pretty soon. [00:30:00] Why? Because we care about it because we're the ones at risk as well as you. Am I getting that correctly or did I

[00:30:05] **Snehal Antani:** That's exactly right. And it's been really awesome to watch this start with the Cyber Collaboration Center as an example, and then push down and then these tier two suppliers, they're getting excited because it's a win They don't have to spend that money out of their pocket because the government's paying is subsidizing or actually covering the cost of, quarterly pen tests for them, they're able to go off and redirect that money otherwise spent on consultants towards other things.

They're getting, security posture assessments that are very detailed and actionable, back to find, fix, verify, and quickly fixing and retesting. they're able to prove their posture up to reduce that liability risk. And then those folks want to push it down to the tier three and tier four suppliers that they're dependent on.

And so it only works if you can execute pen testing at scale. And so back to this idea of [00:31:00] autonomy. we, by, by declaring ourselves a data company on day zero, we had to design a multi tenant SaaS architecture, but with very strict data isolation controls and other things in place in order to do what we do.

And we kind of came from that world. And so if you're a supplier, if you're a procurement, you can either be the multi tenant owner and all of your suppliers are sub clients and you're gonna be able to run assessments, Or the sub clients are there are customers directly and they can opt in to share your data up.

So you've got all sorts of flexibility in here. What we're actually seeing with the DIB is for the first year or two, they're taking advantage of the government providing this capability to them. And then they get addicted to running pen tests as often as possible. And they graduate out of the program into using continuous pen testing themselves and then just opting in to share the results.

Thanks. up and out. And so it's a really interesting network effect that has [00:32:00] demonstrably improved the security posture of UAV companies, armaments companies, operating on sensitive data, food service companies, all sorts of vectors that allow the bad guys to get into, to DoD entities.

[00:32:14] **G Mark Hardy:** Wow. As I say, that's just really, a well thought out. And I'm wondering if your experience at SOC had anything to do with thinking

about that, because in that role, you ended up dealing with some of the folks that are really physically charged with defending the nation.

[00:32:31] **Snehal Antani:** I think that my big epiphany. At that time, and even my time at Capital and elsewhere, was the attack surface has changed so dramatically over 10 years. So when I was a CIO in 2012 to 2015, the primary concern was insider threat. Like we were mostly an on prem shop with some very specific cloud workloads.

And the fear was insider threat targeting Crown Jewels data. If you fast forward 10 years later, if I was back in the CIO seat, my attack surface [00:33:00] is my on prem environment, My SaaS applications, all of my multiple cloud environments that are hosting critical workflows and processes. Insider threat, credential reuse from personal systems like Netflix to your corporate email.

and the attack surface or the risk posed by my, critical suppliers. And so your attack surface has grown significantly over that time, yet your resources haven't increased and you need some scalable way. to assess where you are exploitable, and this is the key term that we started focusing on.

Exploitable is more important than vulnerable. Just because you're vulnerable doesn't mean you have to take immediate action. Are you exploitable right now? And if so, what is that ease of exploitation? What is the impact of exploitation and how do you fix that as quickly as possible?

Quote #2

[00:33:50] **G Mark Hardy:** if we look at that then, let me offer this as a thesis and you can go ahead and agree or disagree and explain why. Using CVEs as a defense [00:34:00] priority for patching is a way to solve this problem.

[00:34:04] **Snehal Antani:** I would vehemently disagree and flip a virtual desk. Because, you know, I gave this talk at Black Hat where I said I, I showed the, bulk of the ways that we get in and break into systems. And there's a lot of data that makes clear that CVEs are maybe 2% or less of, of the attacks that you see out there.

It's misconfigurations dangerous product, default misconfigured, defensive tools, harvested credentials. That's how attackers get in. They don't hack in using zero days like the movies. They log in with credentials that they've harvested through a variety of techniques, guessing usernames off of LinkedIn

searches, finding that one employee that reused their compromised Netflix password as their corporate email.

That's how they do it. And, you know, those spending millions of dollars on EDRs, forgetting to check the advanced configuration that prevents OS credential dumping from occurring, you don't need CVEs in zero days. And I think [00:35:00] that we have created a false sense of security across CIOs. And CISOs that your vulnerability management program of driving CVEs to zero is good enough.

It's not, it's the exact opposite. It's giving you a false sense of security.

[00:35:15] **G Mark Hardy:** Yeah. I mean, one of the best thing you do, at least for admin, I would tell everybody, give your admins one of these things, little YubiKey and they're a little bit pain to get you started, but once you, yep, there we go. We're kind of a member of the club here, but absolutely. It's a stolen credential. A, however it got done is not going to complete the launch sequence for someone to come in and become a global admin in my tenants.

[00:35:38] **Snehal Antani:** And on top of that, it's how you combine these things together. So you could take, a misconfigured LLM and R service, which in certain tools is informational. In other tools, it's a one severity score. Combine that with SMB signing not required, which in certain products is informational. In other products, it's a one CVS score.

And if you [00:36:00] combine them together, you become domain admin, which is a 10 out of 10. And so we see this over and over again of how attackers, including us, are able to combine a variety of lower severity issues or informational issues into a sequence of steps that lead to full domain compromise or full, Sensitive data exposure, and that's just how attackers operate too.

[00:36:23] **G Mark Hardy:** Yeah, and also backing up that comment that I made on the CVEs, I was just reading this morning about the fact there's tens of thousands of CVEs waiting to be evaluated and classified. They're all turned in, I think it's now 400 different entities that are out there and we just can't keep up.

so what that is telling us though, is that something is happening and it's the rate of change is accelerating. And so it's a second order derivative, if you will, instead of speed, you know, acceleration. And by the way, the next derivative past acceleration is a jerk of the first order. And it's not meant to describe

anybody that you might've worked for, but is the rate of change [00:37:00] exceeding our rate of learning?

And if so, how do we bridge that gap?

[00:37:05] **Snehal Antani:** I would, I would take that a step further and use CISA KEVs as a great example. So a key exploitable vulnerability, isn't as a annotation or a, classification from CISA, and it just indicates that it has now been observed that this particular vulnerability is massively exploited in the wild.

So what does that mean? one great. That. CISA has illuminated cybersecurity and elevated it to the point where people are reading about it in the boardroom that are not security practitioners and they've done a great job raising awareness. What's the second thing it does? The moment they mark something at KEV, it's a spotlight and you've got all of these other fast follower attackers that now know what to go off and use because it's massively exploited.

So it creates a fast follower dynamic with the bad guys. But what does it do to the poor CIO who now has [00:38:00] to treat it like an IT outage? Stop everything they're doing and rapidly surge to find and patch every occurrence. And the recent Verizon DBIR report said that 85 percent of CISA KEVs are still not remediated 30 days after it's been announced.

In part, the guess is that those shops don't know that it's there, so shadow IT or bad inventory, or they don't have the capacity to fix it. Or they're not operating with urgency because they're not aware or they don't care. You know, I would, assume it's not the third one. It's probably more the first two.

And then imagine yourself. One of the hardest things I had to do as a CIO was look at my IT admins in the face and say, I need you to skip your kid's basketball game, stay late and patch a bunch of servers because my bone scanner says these things are critical. Even though we all know that it's nonsense.

You know, and that was before we was talking about KEVs. Now I've got to go look at them and say, you need to stop [00:39:00] what you're doing and massively test and patch and so on. And so I think what we're seeing now is this huge impact that KEVs are having on CISOs their ability to do their jobs with any sort of predictability.

And so rapid response is a really important area we focused on where we can actually run analytics on your historical pen test results. Thanks. And notify you proactively that we're quite confident you're exploitable to this F5 CISA KEV on these eight hosts. Click here to run a pen test just on those hosts for just this problem and verify that you're not exploited.

And if you are, you can fix it and run a retest and verify that you're good to go. So a lot of our CISOs now take these before after screenshots. of pen tests, where in the before, there's a bunch of things that are, that require remediation or mitigation. After, they show it's all green and that's evidence that they can send up to the board.

[00:39:58] **G Mark Hardy:** And I love that [00:40:00] because one of the things that I face, I am doing a board presentation for a client a couple of weeks and I just finished my deck, for the briefing as a CISO. And of course, it's a snapshot in time, and I sent my draft in a couple days ago, and one of the things we look at sometimes is Microsoft's security score.

for better or worse, by the way, if you want to max that out, you've got to buy a lot more things from Microsoft that you might not already have. But I thought it would get approaching my theoretical maximum, sent it in, and then over the weekend, I did a couple more tweaks, and I realized I nudged it up just a little bit more.

In a way, it's a bit of a false sense of security because although I'm trying to make more things look green, it's really going up against a metric that is not necessarily reflecting my current threats. And so one of the things I want to share, and again, we're getting to the last few minutes of the show, it's gone quickly, but I'm loving this thing, is a quote from Thomas K. Adams. from an essay he wrote called Future Warfare and the Decline of Human Decision Making. So I'll give you three brief quotes out of this paper, which I really loved. [00:41:00] It's published by the U. S. Army War College, of which I'm a graduate. And it says, warfare has begun to leave human space. We're faced with a prospect of equipment that not only does not require soldiers to operate it, but may be defeated if humans do not attempt to exert control in any direct way.

It's easy to see a steadily decreasing role for humans in direct combat as the 21st century progresses. Victory does not always go to the commander with the best observation. It goes to the one that can best process observation into data into information. Information to orders and then orders into action.

This process is continuous. The results of action are observed, starting the process all over again. The individual functions involved have been enshrined in military jargon as the OODA loop. By the way, He wrote that in 2001.

[00:41:53] **Snehal Antani:** Yeah? Yeah.

[00:41:55] **G Mark Hardy:** if you go back and read that at rest of that essay, and it's up, I'll put the link in the show notes.

It's [00:42:00] wow, this guy was really onto something. A lot of things that have happened. So one of the things we're looking at then is going ahead for the future is we really have shifted. from compliance to readiness, or we need to. And if we haven't, we need to start doing that because in the military jargon, we say we are going to train the way you fight. And if we look at the Admiral Hyman Rickover quote that he had made years ago, the father of the nuclear Navy, who had said, any man who stops to think in a crisis shows a severe lack of training. All right. Being just men at the time they were in the submarine force, but in a way we don't have the luxury to stop and think anymore.

The attackers aren't giving that to us. The time from an exploit being discovered to being cataloged to being exploited is narrowing and narrowing to the point where we don't have that luxury anymore of putting a person in the loop. So how do we do this human machine thing? Teaming, if you will, in the future, and how is this type of approach with the pen testing in particular going to be one of the critical elements of success going [00:43:00] forward for CISOs in their enterprises?

[00:43:02] **Snehal Antani:** Yeah, it's a great question, so I'll give you some data. the definition of an infrastructure pen test in this example is at least 2, 000 hosts, so a decent sized environment, at least for a typical pen tester or human. And a, critical impact is an issue that spans at least three hosts and, involves at least two different weaknesses.

Harvesting credentials from one, misconfiguration from another. a pretty decent base test for achieving a critical impact. In 2023, we were able to go from initial access to domain compromise critical impact with those, you know, at least three different, three different hosts, at least two different issues in a scope of 2000 in seven minutes and 19 seconds. In early 2024, we're able to get, do something similar. In four minutes and 12 seconds, no humans

[00:43:55] **G Mark Hardy:** I was impressed with the first number, but the second one's, you know,

[00:43:58] **Snehal Antani:** And I suspect [00:44:00] if we optimize for Blitzkrieg in this example, we could do it in under 60 seconds. Right now we optimize for, breath first, which is going to go wide and then go deep. And, but if we'd optimize for depth first, we could probably get it done in 60 seconds or less, which is pwn you as quickly as possible.

So in four minutes and 12 seconds. Can a defender characterize the risk of alerts, get permission to actually take some evasive maneuvers and stifle the attack and be defensible, and then actually execute it before the attacker has locked them out of the system? And the answer is, you have to be pretty damn good to be able to do that.

And so imagine it's 60 seconds. So my hypothesis when we started Horizon 3 was that the future of cyber warfare will be algorithms fighting algorithms with humans by exception. And that's a pretty big leap. And it's a, it's what you would expect from an obnoxious Silicon Valley guy in skinny jeans.

Quote #3

[00:44:54] **Snehal Antani:** But what the data kind of supported, how quickly we were able to compromise you. And [00:45:00] so I don't believe we have a tools problem in cybersecurity today, we have an effectiveness problem because the tools weren't the bottleneck is how effective they were at tuning and your processes to stop them. But we are very quickly going to run into this bottleneck where every defensive tool today is designed with humans at the center.

It is humans in the loop. A few of them have humans on the loop. But if the future of cyber warfare is algorithmic, you must redesign all of your core security tools to be humans out of the loop. And I think you're going to see a massive reckoning over the next few years and a fundamental change in architecture and an entire new generation of security tools that are humans out of the loop that start to emerge.

And I think it's going to take a few high profile Blitzkrieg style cyber attacks to get to that point. where the attacker was in and full domain compromise in 60 seconds or less, locked everybody out, despite the Gucci tools that they already had in place, and it's going to be the time that gunpowder rendered [00:46:00] the city walls obsolete.

[00:46:02] **G Mark Hardy:** wow, it's kind of an interesting, prognosis. I'd love to talk more, but I see where they were out of time. So if anybody wants to learn

a little bit more about you or the organization or your products you build, how would they follow up?

[00:46:13] **Snehal Antani:** Yeah, so go to, horizon3. ai, the website, I joke it's ai because I couldn't afford com when I started the company back in 2019, but no, we actually do some pretty amazing stuff. if you go to the main website, by the way, at the very bottom of the main page is a talk I gave for the National Science Foundation that talks about AI for offensive cyber operations.

So you can get a much deeper dive technically of what we're doing and kind of our point of view of the world, but that's the best way to do it. I'm very active on LinkedIn. So find me on LinkedIn, follow or add, shoot me a message. I'm pretty good at replying. And, I look forward to the continued dialogue and engagement with the broader community.

[00:46:49] **G Mark Hardy:** this was an awesome talk. So Snehal, I want to say thank you very much for your time and your insights. You got me thinking along a different element for those who have been listening in or watching it on our show. [00:47:00] Snehal Antani, the founder and CEO of, or one of the co founders, I guess, of Horizon 3 AI, who's really given us some fascinating ideas about what to think about not just pen testing, but continuous pen testing and the role of being able to have that as part of our automated and autonomous response to the increasing threats that we're facing.

Hey, if you love this show, make sure you subscribe to CISO Tradecraft. We got a lot more for you than just shows on our LinkedIn page. We'll go ahead and we'll provide information for you on a pretty regular basis. High signal, low noise. We'll look forward to having you back for another listen in a week from now.

But if you have been enjoying our show, again, appreciate the thumbs up. In the meantime, this is your host G Mark Hardy. Thank you very much for listening or watching. And until next time, stay safe out there.