VERODIN

WEAPONIZING THE ADVERSARIAL MINDSET

12/18/2018

Brian Contos: Welcome to the Cybersecurity Effectiveness Podcast, sponsored by Verodin. The Verodin Security Instrumentation Platform is the only business platform for security that helps you manage, measure, improve, and communicate security effectiveness. I'm your host, Brian Contos, and we've got a really special guest today. Joining me is Mark Bagley. Welcome to Podcast, Mark.

Mark Bagley: Thanks a lot for having me, Brian.

Brian Contos: So, Mark, let's give the audience a little bit of background on you.

Mark Bagley: So, hi, everybody. I'm Mark Bagley. I'm Verodin's Vice President of Product, and that means I'm responsible for a few different things at Verodin. I'm responsible for the Verodin engineering organization. I'm responsible for our product management function as well as the Verodin Behavioral Research Team, which I think is what you want to talk to me about today.

Brian Contos: Sure, why not? Well, let's go ahead and expand on that then. So, what is the Behavioral Research Team?

Mark Bagley: So, The Behavioral Research Team is Verodin's own security research team. And that's a little different than, you know, what most people think of when they think, "Hey, like what's a security research team?"

Brian Contos: So, Mark, the Behavioral Research Team, or BRT... how is that different than the security research teams that we see in virtually every other security company?

Mark Bagley: Well, Brian, I think it's different in the way that you actually sort of teed up in your asking me of the question, which is that Verodin's team is focused on adversary behavior. And that's a very, very big difference when you think about the types of security research that almost every other organization is doing.

Brian Contos: Okay, so it's based on behavior. You know, I look at other security companies and they're chasing the zero day and there's the endless stream of CVEs and there's tons of threat intelligence players that are pumping this information in and lagging indicators of compromise. That to me seems to be the traditional approach. Why is the way we've done it maybe not the best way to continue doing it and why we're doing it differently?

Mark Bagley: So, Verodin's team is focused on adversary behavior. And quite frankly, the adversary behavior is the most durable thing that we could be researching and codifying. Especially if we're focused on then evaluating the performance of defenses. And, you know, the example that I'll use here is SQL injection. So, the specific behavior for performing SQL injection is roughly the same regardless of the specific vulnerability that an adversary would be attempting to exploit.

Mark Bagley: Yes, there are many different downstream systems that could be vulnerable. And in every single one of those cases, we record a CVE. And that's a fine thing to do, because we do need to actually enumerate all of the different systems that are vulnerable. But chasing every single one of those CVEs inside an environment doesn't necessarily make sense because it's not done in a way that takes into context the defenses that exist and also the rest of the environment that is being defended.

Mark Bagley: So, if I spend my time focused on producing a behavior inside the Verodin Security Instrumentation Platform and then executing that behavior, I'm inherently providing a more durable piece of content to my customer than, you know, stamping out content that addresses CVE after CVE after CVE, in specific.

Brian Contos: You know, I think you brought up a great term there which is durability. And I'll stay with your SQL injection example, because SQL injection has been around for decades. And there's -- I can't count how many CVEs there are for a SQL injection -- but you don't need to run 10,000 disparate SQL injection attacks to determine if your WAF, for example -- there could be other tools, but let's say your WAF -- is successfully blocking. Or if it's not blocking, at least detecting. If it's detecting, is it actually alerting? It's just a question of scalability, right? And why [you should] evaluate your security controls predicated on things that are actually measurable and unique.

Brian Contos: And what it brings to mind is WannaCry. When WannaCry came out a few months ago and all its variants and there was a number of other ransomware flavors that came out at that time, the interesting thing was they were primarily based on a very finite set of behaviors. And to me, that was one of the greatest examples of the scalability that behavioral-based analysis achieves juxtaposed to, again, your point of chasing the latest 0-day or CVE. So, it really just comes down tactically to “How do I scale in the face of an almost limitless barrage of different attacks?”

Mark Bagley: Extremely well said, Brian. And that durability and that sort of scale, being able to provide that to a defender that we know to be overloaded, overworked, and quite frankly, in need of some force multipliers, was the genesis for this concept.

Brian Contos: So, let's switch gears a little bit to frameworks. And it seems like these days everyone's got one. There's the MITRE ATT&CK Framework. Lockheed Martin Kill Chain. NIST [Cybersecurity] Framework. And, you know, and on and on. How does taking a behavioral focus, which, you know, I believe is primarily how MITRE ATT&CK is working their approach as well for endpoint activities post-compromise. But, how does a behavioral focus sort of help when you're trying to address this multitude of frameworks?

Mark Bagley: So, a behavioral focus helps with embracing frameworks because it sort of goes at the dimensions that we discussed previously. Both, you know, a durability and a scaling function. We know that defenders are often overworked. Teams are not as large as they need to be to defend the environments that they're dealing with. And so finding those force multipliers is important. Now, the frameworks that you described earlier are sort of yet another great way to once again start codifying how defenses perform and to put a language around the type of behavior that could be executed at a given time. You know, sort of breaking that all down into the TTP, or Tactics, Techniques, and Procedures, that are commonly observed in the face of the adversary.

Mark Bagley: So, being able to then quickly identify the behaviors that map to the TTP means that I'm not moving throughout a list of specific attacks that may be relevant against a specific vulnerability. I'm working with roughly equivalent first order concepts, and that puts me further down the path to demonstrating my defensive readiness. Did that make sense?

Brian Contos: It did. It did. And it really takes the whole idea of these threats... And again, it brings in durability, it brings in scalability. But it also allows you to operate more strategically. But with that, I always like to take a couple steps back and, you know, we'll let our listeners know, yes, it's great to have everything based on behavior, but hey, if you want to run a very specific attack to do a very specific bit of analysis that might overlap with, you know, a multitude of behaviors that already exist, that's absolutely fine, too. Because Verodin has an open content platform within the Security Instrumentation Platform.

Brian Contos: Which means when the boss calls and says, "Hey, I just heard about this thing called Apache Struts. Are we vulnerable to this? Or if it happens, what are we gonna do? How do we know how to react?" You're like, "Well, I'd like to very specifically be able to get back to them and say, 'I've validated our security controls specifically against Apache Struts across these various zones within my environment. And the answer is yes we are, or no we're not. And if we're not, Verodin SIP has said these are the steps we can take to mitigate that risk.'"

Brian Contos: So, you really get the best of both worlds in that you have this very strategic, scalable approach predicated on behavior. Again, with the ability to do these very, very sort of laser-focused tests against a specific threat. And I think that's really the power that an instrumentation platform brings. But it also shows you the flexibility that it doesn't just have to be this very broad scale behavior focus. It can actually be extremely tactical.

Brian Contos: With that, Mark, are there any final words you'd like to share about the BRT before we wrap up?

Mark Bagley: Sure. And I'll backtrack briefly in regard to what you described there. You know, Struts is a perfect example of a case where a latent vulnerability in the system and the behavior used to exploit that vulnerability should a similar vulnerability exist in the future, because we've identified and captured the behavior itself and not just something that exploits that specific vulnerability. You can be assured that as you test against that behavior, not only have you validated your defenses against that specific VULN today, but also any example in the future of a related behavior, you now also understand your ability to defend against the next variant therein.

Brian Contos: Very cool. And I've had the great honor to spend a lot of time with the folks that make up our Behavioral Research Team, and it's important I think in any type of group like this to have a very diverse group of researchers. And, you know, we've been pushing out tremendous content across network, endpoint, email, cloud security controls, at a pretty regular tick. Additionally, you know, when there's an emergency, right, Mark? I mean, generally speaking, we have content ready for our customers that the BRT has put together and many times in less than 24 hours. Isn't that right?

Mark Bagley: That's absolutely true, Brian. I mean, we have our monthly cadence of releases that are, you know, continually extending the classes of behavior that our customers can execute out of the Verodin Security Instrumentation Platform. But we also understand that when there's a threat in the news, being able to help executives speak to what's going on and whether or not the defenses are ready to respond to the thing that is the headline of the day is important. And that's why in many cases, we have content available for our customers within hours of a new threat entering the news.

Brian Contos: Fantastic. So, Mark, final question before we conclude this podcast and arguably the most important question of the entire podcast. Who's your favorite superhero or super villain and why?

Mark Bagley: Hmmm... I think I'm gonna take it back to The Thundercats and I have to say that while Lion-O wielding the Sword of Omens... He was sort of inspirational in, you know, how he had to sort of wield all of his capabilities against the tests. So definitely a better sort of superhero analogy. I do have to say that he wasn't my favorite. And I would have to say that my favorite was perhaps maybe Cheetara.

Because she was definitely, like, bringing together those elements of being brave and also being empathetic.

Brian Contos: Very, very nice. And you are absolutely our first reference to The Thundercats. So, very exciting to have that in today's podcast.

Mark Bagley: Also, [she was] a swimmer.

Brian Contos: And a swimmer, there you go. Well, thanks, Mark. And thanks to all our listeners for joining. And be sure to check out other Cybersecurity Effectiveness podcasts, sponsored by Verodin.