Heather (00:12):

Welcome to the Hurricane Labs Podcast. I'm Heather, and today we're going to be talking about the sixth annual Collegiate Penetration Testing Competition or CPTC. Now, if you haven't heard about CPTC, it uses competition to address the cybersecurity skills gap by educating college students in risk technology and communication. Here with me today, we have Meredith, Tom, and Josh of the Hurricane Labs staff. They recently helped to coordinate the competition, and they're here with me today to tell us a little bit about that experience and how they use Splunk within the competition. Hi everyone. Thanks for joining me today.

Tom (00:51):

Thanks for having us.

Josh (00:53):

Thanks for having us on.

Heather (00:53):

Anytime! So go ahead and introduce yourselves and tell us a little bit about what you did with CPTC.

Tom (01:01):

My name is Tom I'm kind of the person who does all the crazy stuff that Hurricane Labs and in some of my copious spare time that doesn't actually exist I get to participate in help run the CPTC event. And I've been working with that event since it actually started back, I think 2015 and have kind of seen it built from just a regional security, offensive security competition to what is basically one of the world's premier offensive security competitions.

Meredith (01:32):

Hello, I'm Meredith. I work at Hurricane Labs on the Vulnerability Management and Pentest teams. I did spend a fair amount of time in the SOC, which we'll get to later. And for CPC, I first joined the competition in 2018 as a competitor and had so much fun that beginning in 2019, I joined the volunteer side and I am now the volunteer wrangler.

Josh (01:57):

I'm Josh, I worked on the Monitoring team, mostly handling audit alerting in Splunk. I also work at Hurricane Labs in the SOC mostly handling alerts that come in and fixing broken things, which is also what I ended up doing for CPTC as well, just troubleshooting random things in the environment.

Meredith (02:17):

And solving my problems when I screech them at you.

Tom (02:20):

All of our problems when we screech them at you.

Josh (02:23):

Finding out who broke the dam,

Tom (02:27):

Yes, there'll be a lot of them jokes this year, which is part of the thing we try to do. We have a good time when we're building this event.

Heather (02:34):

So tell me a little bit about the structure of the competition this year.

Tom (02:37):

Sure. So this year, obviously it was different because of many of the challenges that we faced with everyone shifting to remote. When this event started, the first couple of years it was a single event that took place at RIT Rochester Institute of Technology in New York and maybe ten teams competed. I think that was maybe 2015-2016, where that was the case. And we have been growing ever since. So, and then this year was the first time we branched out to where we had basically the top eight teams, one from each regional, made it, and then the top seven next teams at large across the world and ended up making it as well. And yeah, man, October and November, that was a bit of a marathon because normally we would actually do all the regionals all at once. And it was same day. Every event happened at essentially the same time based on time zones. But we decided not to do that this year and do things remotely because of the pandemic and basically created a situation where every weekend in October and half of November, we were doing a CPTC event that was either one or two regionals at the same time. And yeah, Josh, you and Meredith were both involved with all of that, even the ones that have crazy time zones.

Josh (04:07):

It really made things interesting. It actually helped a bit because we were able to do things in revisions, like in Splunk with our audit searches, like where we were able to like slowly iron out details and issues and just expand things, just using like the previous regionals data and just like, this just happened. Like, what did we learn? How can we improve our searches?

Tom (04:26):

Yeah, that is true. Because like how the systems were built, we couldn't change that between events. Cause we had to give all the teams the same experience, but using the data, we could keep on refining our process. And we only had one shot to do that in the past with one regional leading up to the finals. So yeah, that's a really good point. I know Meredith you don't get any sleep ever–now maybe a little bit more than you did in the past, but you really helpful with those overnights where time zones in Dubai and Europe just don't work too well for us here in the East coast.

Meredith (05:01):

I will say it, it was kind of fun to take a nap on a Friday evening and then hop in around, you know, 10 or 11:00 PM and get ready to start some fun. And then, you know, I'd be done by 8 or 9:00 AM and okay, perfect nap time. It's a Saturday I don't have work to do. And then wake up later in the evening and do some grading and stuff. So that was a lot of fun.

Tom (05:23):

Yeah. Your definition of fun is maybe a little bit different than some of us, but it still worked out.

Meredith (05:29):

It's okay. I'm weird. We all know it.

Tom (05:30):

Lucas, Morris, and I are probably the, some of the leads in the technical side of things and making this event happen. And we basically say this event was created by crazies, but we say that in the most affectionate and kind version of crazy that we can because we all do this because we're passionate about it and want to create the best event we can, but we are all very much insane to make this all happen.

Heather (05:59):

So, from what I understand, you all really go all out with designing this competition. What all goes into pulling this off?

Tom (06:06):

One of the themes of CPTC is to base the event off of a real life business and basically create the most realistic and immersive experience that we can in order to give the students that are participating a real life experience of what it's like to work with an actual company. And that doesn't just mean, you know, we build up a couple of systems to mock up that environment. We take that a lot more seriously and it varies obviously by year, but in the past we've been a number of different organizations. So that could have been a hospital one year, an elections provider, we were a ride sharing company one year, last year we were a bank and Meredith can tell you all kinds of fun stories about the stuff we did for that. And then this past year, we were a utility company called the Next Generation Power Electric and Water or NGP. Some of the systems that we build for that environment basically are mock-ups or simulations of what you might find in that type of organization. So for the utility, we had actual power plant like components that we built. For the hospital, we had a system to simulate that type of environment. We had voting machines that we created for the election service. And then for the banking system, we had all kinds of banking services and actually a truckload of actual ATM's that we took from a bunch of waffle houses, legally, of course from somewhere in Ohio and took them to New York. And that was an adventure. But we definitely take the getting into the theme of what we're trying to do very seriously. The history of all of the dam jokes that we have is because one of the components in this environment had a dam, and that is D-A-M for those of you that are listening, which is everyone, because this is a podcast. And basically this was a very, very sensitive environment that depending on what traffic was sent to it, this dam controller would crash. And it would result in the water level of the reservoir behind the dam. We ended up seeing like, depending on, you know, what teams were doing this water level would go up and down quite rapidly whether the point where it would go over the top of the dam and obviously cause a lot of flooding or below the like minimum level and the nuclear power plant would just like explode. What was it that we saw teams do that actually caused that to happen?

Josh (08:48):

Yeah, I can talk about that. Mostly we saw teams just running Nmap over the entire environment. Just regular service scans would do it, just sending any kind of data to the ports that the dam controller was listening on. And it was a common port for this kind of device too. And it's left out of the default scans of Nmap generally. So what we saw was OpenVas doesn't exclude it. So when teams would run OpenVas over the subnet range, it would start hitting the controller and sending random data to it, trying to see

what service it is and would start bringing down the dam and they would just have these running automated all the time. Like it would just be down fluctuating for an hour or so.

Tom (09:35):

So like, one of the things we were trying to teach the students in this year's event is you have to be careful when you're working with industrial control systems. And that is something that if you were assessing these types of systems, you would have to be careful about. So yeah, like, Nmap does exclude that. And I think you're right, Josh, OpenVas doesn't and, you know, we'd inevitably come to a situation where we had a Splunk dashboard going on that was showing all of this water level going up really fast and down really fast. And then someone in character would get on the phone as a rather frustrated customer dealing with a national emergency, or at least a regional emergency at that point asking what the heck's going on. And the team has been like, no, we're not scanning anything. So I think–how many times did we deal with the teams that weren't scanning anything that they definitely were based on the logs that don't lie?

Meredith (10:32):

The teams that actually were not scanning anything or teams that were in fact scanning something,

Tom (10:37):

Aren't they both the same in this case?

Meredith (10:39):

Uh very close. We did have the one inject where we said, "Hey, are you scanning this thing?" And some of them were able to say no and be truthful. And some said no and logs determined that that was a lie.

Tom (10:52):

That is true for regionals. We wanted to make sure all the teams had that experience. So we were the typical customer that blames the consultant for anything that happens and how we responded actually varied based on what our data was at that point. But yeah, you're right. There were teams that actually hadn't done anything at that point that kind of gave them a little bit of an exercise in how do you explain that to a client?

Meredith (11:17):

It's a thing that happens and it's best if you own up to it right away, which I think is a really good lesson that the students got to learn this year.

Tom (11:25):

Yeah, that's true. Being honest with the client and talking about what they actually did is really never the wrong solution. And that ties into what we try to do with CPTC, it's making it not just a security competition and a event where you showed your technical skills, but also practice those professional skills that you don't really get to see until you're working. And for me, that's one of the parts that I really put a lot of effort into trying to make happen. And I think it's one of the things that really sets this event apart.

Josh (11:58):

Yeah. It really makes things different from something like Hack the Box, where you can pretty much throw whatever you want added or any other normal CTF where you can run any kind of intensive scans you want, and if you bring it down and you just reboot the box yourself. When you're in a competition like this, where you're dealing with a virtual customer, it becomes a lot different where you need to be careful and you need to keep track of exactly who's running what on your team. So you can say with confidence that if you didn't do something, you didn't do it. And if you did do something, you can tell what it is and stop it.

Tom (12:31):

That awareness of what your team is doing is something that is really important. And I think it's an area a lot of teams actually struggle with because when we talked to the teams, as the people that are the customers and all of this is done, like in character as if you were actually a customer. So I've played this role before and we actually get industry professionals who know what they're doing, as opposed to me do that. And they get to have those conversations with the students and we get to hear how they react to it. And it's definitely something they're not necessarily used to hearing, but it's a good exposure for sure, because I know I never actually got to compete in CPTC, but both Josh and Meredith have at least a year or two of that under our belts. And hopefully it was valuable. And that's what I hear from the students.

Meredith (13:21):

I'll throw in a plus one for valuable. Out of all of the competitions I've been in that one was the closest to real life for what I actually get to do on a daily basis. Now, granted, of course, you know, we do take some liberties, but it is the closest to real life that you will ever experience.

Tom (13:40):

Yeah. That, that's really what we try to do for this sort of thing too, is how do we make it realistic and how do we take our professional experiences and put them into the event and give students exposure to things that we see in industry.

Heather (13:57):

Why don't you go ahead and tell me a little bit about how Splunk gets wielded in the competition.

Tom (14:03):

I think I'll talk a little bit about infrastructure and data, and then Josh, you can talk about the enterprise security and SOC components of it.

Josh (14:13):

Okay.

Tom (14:14):

So basically the approach I take for the Splunk side of things is I know what's in the environment and the type of data I want to collect is yes. Basically anything that we can get data for, we will turn that on and try to get as much detail from the systems in the environment as we can. I would say there's things that we do that are targeted towards certain types of activity that we know we're going to use. And then there's also things that we're not really sure what they're going to be used for, but they could be

potentially beneficial for integrity monitoring or research purposes. So good examples of things that are things we know we're going to do: the dam this year is a perfect example. That actually was a custom written service that operated, I believe it was like a Docker container or something like that, that presented a lot of these virtual, like industrial control type systems. And what that actually did was monitored–it had its own ability to self-regulate like a system would be that has that logic built in and desired state and every couple seconds or so it would actually send the values of those sensors out into Splunk. So it was a little bit, maybe you wouldn't see that exact type of logging mechanism in a real industrial control system, but they're going to be using some kind of syslog or other mechanism to send that data depending on what platform you're using. So we model that they behave similarly. And that's a good example of something like we know this component of the game was going to be something we want to pay attention to and something that we want to have dashboards to show what's happening and also collect that data. Other things we do are more standard data sources that you would see in really any type of environment. One of the good examples that we do is like log all the bash history and we've configured the system so that it does that dumps it out to Splunk in basically real time, so we can see commands that are coming in as they're entered, as opposed to waiting for our session to end. So it gives us, you know, within, you know, a minute or so of releasing all the teams to the environment, we can already see what the teams are doing to set up their systems and get ready to knock things out. Some of the other things we do are monitoring the target systems in the environment to see what types of things students are doing to them. Authentication logs and security logs, and windows are a great example for that. So we can see if they're creating situations where like accounts are trying to be brute forced. We're also grabbing a lot of metadata on the networking side of things, using the Splunk App for Stream. And what that's doing is monitoring TCP, UDP connections from all the systems in the environment. So we can see what talks to what and how. So it's, it's really just, you know, a highly monitored corporate environment with maybe a little bit more detail than you would normally collect. Because it's like the intervals we're collecting stuff is a lot more compressed than you might see at a large organization. So that's kind of the the Splunk side of things and all of that soup gets dumped over to Josh to do I guess it's not even soup at that point, it's just like a bunch of vegetables and other stuff going to deal with. Right? So I know we've wanted to use ES, which is Splunk Enterprise Security for some time, but you're the one who made the magic happen this year. So how'd you approach that?

Josh (17:56):

Yeah, I mean, we just started out like just brainstorming–me and Meredith and you–of what types of things do we not want the teams to do and then taking those logs and figuring out a way to detect them? You know, for example, one of the last things we did at the, at the competition was just account lockouts, like penalizing teams for locking out accounts from brute force activities. We were just able to take the, the win event logs and have a search that spits out all the accounts that a team has locked out in an hour. And like some of the earlier ones that we had were just ensuring competition integrity, having correlation searches, to make sure that people were logging in to Google Drive from the correct locations that I guess one thing I should mention is that we had the teams fill out a form for what IP addresses they were to be expected from logging into the competition environment. And we were able to take that asset data, import it into Splunk through some rather interesting means, and then write correlation searches comparing their log-on activity to the asset data that we received for them.

Tom (19:03):

Basically that that's really based on something that a lot of clients do in their SOC alerting with us. Right?

Josh (19:10):

Yeah. I know, we have geographically improbable searches, which is basically if you see two log-ins close to each other from geographically distant locations, we also at some kind of, some of our customers will have asset data to determine like, okay, is this their normal working location? If so we can, we can exclude it from these results.

Tom (19:30):

Got it. So instead of like just a list of employees at a customer, you took the student competitors and treated them like the same kind of thing and put that into the UC customer.

Josh (19:43):

Exactly. I mean, we had asked that data of like their IP associated with which competitor was coming from that IP.

Tom (19:50):

Yeah. I'd imagine that kind of the brainstorming process that we have is, is probably based on at least two factors. One of it would be you two are both previous CPTC competitors. So you have a pretty good idea of the mindset of someone participating in this event.

Josh (20:06):

Yeah. It was fun brainstorming, like, okay. So how would we try to cheat? Things like DNS tunneling and just data exfiltration to strange places.

Tom (20:19):

The other thing I would say is as SOC Analyst and working with clients, you have a good sense of the type of alerts that clients are looking for and what are relevant to an organization that might be getting pentested.

Josh (20:34):

Yeah. And one of the great things about CPTC was just the amount of logs that we had. Like, we don't always have process logs at our customers because they can get so large and so much data for every one of these systems we had process logs. So when we an alert come in of like a large outbound transfer, like we can actually go in and figure out like what file they were trying to transfer and where.

Tom (20:55):

That's really cool. And honestly, that's not something I necessarily expected when I turned on that monitoring for process monitoring. So the fact that you're doing something with that and benefiting from it is really awesome. Is there data types that you found in the CPTC environment that were really useful that you don't see at clients a lot of times that maybe might be good ones for them to consider and maybe new security use cases?

Josh (21:20):

The windows process logs. I forget what it's called, like the win host MAN logs, where it'll log traffic based on, on process and that's like, the Holy Grail to me is an analyst of, of logs is–Okay, so I see this traffic coming to the firewall, going to weird place. What process did it like, was this just from Chrome

browsing or was a piece of malware, like a VB script doing this. And that, that is amazing from my perspective much, much better than just having firewall logs is having a process associated with that.

Tom (21:56):

I imagine if you're just getting firewall logs or, you know, stream events, you might know what name was involved for DNS request and where it's going to from what hosts. But besides that, you're kind of almost stuck just to the host level, right?

Josh (22:13):

Yeah. Like once I can find it, I can figure out how to stop it and having something like that is great at customers. It's just the volume of data that it generates becomes an issue.

Tom (22:23):

CPTC is a good example of insane amount of volume of data from what we're collecting because of how verbose we are, but that's kind of by design.

Josh (22:32):

Yeah. I think it didn't. We have to expand the hard drives a couple of times.

Tom (22:36):

Well, they're not hard. They're, they're really fast SSDs, which helps. But but yeah, if I think we're up to 10 terabytes of fast SSD and the indexers that we have for that environment, which now from a Splunk architect perspective, some of the stuff that we have to do makes me cringe a little bit, but there are some practical limitationsIn the scaling that we have, we essentially have a Splunk environment that means to handle a massive amount of data, a couple of weekends every year, and then handle search requests outside of that, you know, normally during, you know, anything that isn't an event we're ingesting internal logs from the Splunk systems. So basically zero megabytes or gigabytes per day, and the, the weekends we're running these events and some of the times leading up to it, we can see, you know, one and a half to three terabytes of data going into this environment. Basically, we've just accounted for a lot of that by in the Splunk architect in me is cringing, but just having powerful indexers to handle it and kind of really just jamming as we can through those systems. It is way over what is recommended for our Splunk environment, but it does show that, you know, a burst can be handled by reasonably powerful systems, keeping it up and running at a cost-effective solution is hard. So RIT actually just had a cyber range that they're opening a would've opened in person this year, but they, they had done the virtual version of that. And we were using quite a nice size of some of the new VMware servers that they have that are designed to handle basically the entire CPTC competition.

Heather (24:19):

What's your favorite part about doing the competition?

Josh (24:21):

I know for me, it was looking at all the data coming in, just watching what the teams are doing and how they go about their pentest and like the similarities between things, you know, so the calls that are made out when someone breaks something and we need to call a team because their dam has overflowed and just hearing the responses and then going and looking at the data and be like, "Oh,

here's what they did". I see this process log here with this traffic going to the going to the dam control port. It was just very fun to look at. Kept kept things from getting boring over the day.

Tom (24:58):

I honestly, with all the lack of people contact, we had, it was just those Saturdays that we spent on, you know, all day random calls talking about this stuff while we were watching what the teams are doing that's like the closest to human interaction I had in all of 2020. And it sounds really weird, but it was, it was a good change from what was going on. Otherwise. How about you Meredith?

Meredith (25:23):

So as much as I do enjoy the fun and the shenanigans with the calls, my absolute favorite, I guess it's a concept more about the competition is our competition integrity, because I firmly believe that your competition is only, you know, a solid competition if you do in fact stick true to your guns and uphold the rules and uphold competition integrity. And, you know, we do utilize Splunk for competition integrity, and I love to be a part of that because I don't want to see, you know, I don't want to see anybody trying to get by our monitoring. And when they do, you know, that's a perfect time to build out something new or for us to, you know, think about, okay, is this something that would be considered a competition integrity violation, or is this just, you know, somebody flubbed something on the keyboard and did something bad? So I enjoyed that part.

Josh (26:16):

Yeah. It's interesting. Like the same logs used for researching with like just network traffic logs and firewalls and process logs and Windows and Linux. Just being able to take that and use it for auditing the environment, just making sure people are logging in where they're supposed to be logging in from and any other strange behaviors going on and building out searches for them. It's really good time. And I learned a lot about using some of the, some of Splunk that I haven't touched before.

Tom (26:50):

Now. I would say for me that the thing that is my favorite part is the students. Like we've basically given up on tracking how much time we put into this, but it is safe to say it's a lot, but the part that makes it all worth it is the feedback we get from students and the improvement you see year over year, but also talking to people after they graduate and having them come back, help us make this event even better. And when I, you know, I'm at a conference or an industry event, and someone comes up to me and says, "Hey, I competed in CPTC this year and is the best thing that I did in college". That is something that just makes it all worth it for me.

Heather (27:32):

Alright. Well, thank you very much for taking the time to talk to me today. I appreciate it.

Tom (27:38):

Yeah, Heather, thanks for letting us ramble and hopefully something coherent came out of that for the audience. But I would just say for those of you who are interested in learning more about the event you can go to the CPTC website, which I'm sure will be linked somewhere in a way that people can go to it. We use the Twitter account for CPTC to share information and I'm always willing to answer questions and we release a ton of information about the event and try to be transparent about what's going on,

but also make everyone be given the opportunity to get better with the information that we have available. So I'm always looking for ways to improve the whole competition and make it something more valuable for everyone who's involved.

Speaker 1 (28:22):

Absolutely. Well, again, thank you very much. And as Tom says, you can check out our links to learn more. Thanks for joining us, stay safe and we'll catch you next time.