ML-Driven Decision Making in Realistic Cyber Exercises

featuring Dustin Updyke Interviewed by Tom Podnar

Welcome to the SEI Podcast Series, a production of the Carnegie Mellon University Software Engineering Institute. The SEI is a federally funded research and development center sponsored by the U.S. Department of Defense. A transcript of today's podcast is posted on the SEI website at sei.cmu.edu/podcasts.

Tom Podnar: Welcome to the SEI Podcast Series. My name is Tom Podnar, and I am a senior cybersecurity engineer in the SEI CERT Division. Today I am joined by my colleague Dustin Updyke, who is also a senior cybersecurity engineer at CERT. Today, we are here to talk about our work using machine learning to make non-player characters more realistic.

Dustin Updyke: Hey, Tom, thanks for the intro.

Tom: Let's start by telling our audience about ourselves and what brought us to the SEI and the work that we do here and the coolest part of our jobs. Dustin, you go first.

Dustin: I had 20-plus years of experience in industry, and a large part of that is writing websites and web applications, mobile applications. The promise that I would be able to work on totally different things was one of the big factors in me deciding to come here. As a back story, the first thing I did work on was a web application. But it wasn't long after I got here that we as a team were talking about simulating users within an exercise or training event. Some of the things that we wanted to do weren't currently supported by any software that was out there. So that got me thinking in a space that I didn't have a lot of experience, but I was really excited to get to work in that area and think about how such a system would be designed, how the different agents would communicate with one another, and what kinds of activities and decisions they would ultimately make.

Tom: Thanks Dustin. My background is similar from having a long history of doing network infrastructure, network applications, and systems architecture. I really wanted to get involved in cybersecurity more than I was already in my previous lives. So I came to CERT to do more cybersecurity. I was very fortunate, and you were fortunate to join our team as well, whereas we have a good work plan where we work with cybersecurity exercise and training with the U.S. Army. It has become really interesting work. There are great synergies here between your talents and our team's talents in terms of developing applications that are useful for them from a cybersecurity exercise-and-training perspective.

Dustin: Yes, I think it is interesting how our team blends all of the different members' experiences and the background that they bring. We play to all of their skill sets. I didn't have a traditional cybersecurity background, but a great deal of cybersecurity deals with software and building and maintaining those systems. I was really excited to be able to contribute in that respect.

Tom: I think that is probably one of the most interesting parts of our jobs, is that we have a lot of freedom to do what we would like to do and put our input in and design our own parts of the process that, so far, we have made our Army customers pretty happy with what we have provided to them.

Dustin: Well, I would say, one of the things that has kept us engaged and differentiates us from other teams is our continual focus on realism. The Army trains as they fight, and we have taken that same mantra and looked at how these training and exercise events are managed, how to present it to the customer, what systems look like, how they are configured, and ultimately what that entire—we will call it an enterprise-class network—looks like when they get on there and they start to train. I think that you brought a lot of valuable experience in that area, in working with larger enterprise-level networks and the systems that live on those networks.

Tom: Dustin, you raise an interesting point, which is the concept of realism. Early on when I first started working with the Army over five years ago, it wasn't something that was really strongly identified as something that should be included in the cyber exercises we did with the Army. It really became clear to us after a couple of years of work that realism within that exercise environment is just so critical for multiple reasons. It's critical from the idea of *train as you fight*. We have heard that a lot, but it is also very critical from a participant perspective because participants are very busy people. They are based all around the world. They are facing cyber challenges on a daily basis, 24 by 7, and they are very busy people. So when they are trying to gain expertise in our cyber ranges, it is very critical that they are learning skills that are applicable to their daily lives or their daily work efforts, but as well as just buy-in from a participant perspective in that they are more likely to feel engaged and want to participate when they feel like they are immersed in an environment that is very realistic.

Dustin: It wasn't long before I got here that you and a couple other members of the team had published a framework called R-EACTR that outlined the importance of realism across a spectrum of concerns. That really set the tone for thinking about some of the tools that I would ultimately come to work on and build here, including GHOSTS [General HOSTS], which I am sure we'll talk about.

Tom: That shows you from the value of the team. Normally, when people would say they want to do cybersecurity, they would say, *Oh you need to have cybersecurity experience.* Or someone would say, *Oh they are not a good candidate to work with our team because they don't have cybersecurity experience*. But, in the case that having all kinds of different people—software developers, engineers, planners—different skill sets are very important to our team.

Dustin: Yes, so obviously the cliche is *train as you fight*, right? But just like the military, the thing that they want to accomplish in that is a sort of muscle memory that has value in real-life situations that might occur. So, we focus on that, and part of that is to identify "game-isms" within training or exercise. Whether it's a system that they don't use in their real operations, or it's a configuration that they would never see that allows things to happen in the environment that they would typically not have to deal with. I think that's a big part of it as well.

Tom: It kind of brings us to the first efforts, one of the early efforts from our team was the development of GHOSTS and the need to have what the GHOSTS product provides. You want to talk about GHOSTS briefly?

Dustin: Yes, so at that time, I think CERT had a long history of effort in simulating users prior to even my arriving here. But some of the more complex things that we wanted to do weren't available out of the box. There were lots of competing products in that space, both open source and commercial products, that weren't quite what we needed. And what we needed specifically is that teams today are so good at what they do. Say they see something on the network that they want to investigate, they would track that back to host-based logs to correlate, *Oh this event happened, and therefore that is why we are seeing this traffic*. But even to go back and look for artifacts on the network that would, taking those three things to account, triangulate what exactly occurred. So none of the tools that we had seen before offered that kind of realism. The initial challenge to me was, *Hey, if we wanted to replicate someone doing something on the computer, we want it to look like it would like if you looked over my shoulder, and I was doing it myself.* So no sort of spoofing activity that looks like a user fired up Outlook and opened some email message and ran an attachment they shouldn't have. We want you to actually do that. Like I said, that was very different than anything I had worked on before. But that seemed like the most realistic way to simulate that sort of behavior, and the challenge of that just kept eating at me. I just had to start tinkering around with, how would you build such a system to do that?

Tom: So, Dustin, the delivery of the GHOSTS product adds a lot of flexibility to the cyber range. In the context of that, we talk about train-as-you-fight realism in a cyber range. There is a lot of different levels to that. Take, for example, that if a particular Army unit was used to using Cisco routers or products from McAfee or specific vendors, we can provide those services within our cyber range so they can be there. But the problem exists that you have to have real user activity, realistic network activity, for it to all make sense. Because otherwise if you just apply a bunch of Cisco routers or Microsoft servers, you have everything sitting there doing nothing. You have to bring that to life. You have to energize that somehow. That is where GHOSTS really fits in, is where we can deploy hundreds of instances of GHOSTS, when we say *simulate*, it's really not simulation. GHOSTS is actually playing the role on that NPC [non-player character] as that user creating network traffic, web-browsing traffic, file-system access, active-directory logins, file-share access, and things like that. It is bringing those different parts of realism together, which are really critical from the participant's perspective to make them feel comfortable and believe that things are very realistic.

Dustin: The team and probably you in particular are really good at pushing on that. If the baseline was to simulate activity that someone might be doing on a computer, that quickly became a series of activities. So the user does this and that, that spawns them doing something else in the future. Like you said, being able to link activity from maybe what they are doing on their workstation to things that they do on other machines quickly became the ask.

I think that is critical because that is what you see the people that we want to simulate doing on the actual ranges. I think you brought up a good point about *simulate*, because we pretty quickly moved away from that term as well because these are characters in a scenario that we are trying to play out. They actually have a role that they are trying to fulfill, activities that they need to go and do in order to promote that realism within that exercise. That actually bought into this whole idea of when we structure a scenario, it is almost like setting up a complex playbook for a series that is going to run multiple years on a TV network. Because our exercise season can be very long, if we portray someone as insider threat early on in the exercise, we want all that stuff to stay the same throughout that season, so we can refer back to that particular person and the other non-player characters around them that they may have affected in some way, shape, or form.

Tom: I think that we had a lot of success with GHOSTS being able to provide those NPCs in a range. What we did find by watching our players and our participants in our cyber range, was that some of the better players would find ways to take advantage for, say, game-isms of the NPC generators and NPC traffic generators in the GHOSTS system to try to put them at an advantage, where they could see that certain browser user agents were used by the NPCs. So if they found anything different from that, then that could lead them to maybe detecting malware quicker on their network. Some of the Army teams would even keep master lists of all the sites

being accessed. As soon as some malware site which was not part of the normal NPC browsing sites was accessed, they knew they could track on that pretty quickly to put them at advantage within that exercise. So, we were always looking for ways to improve GHOSTS, and adding the ML [machine learning] piece to GHOSTS was one of the first steps to improving that realism, such that the user…It was even harder for that participant to use those game-isms to their advantage and make the environment even more realistic. So, Dustin, you want to talk about that aspect of GHOSTS?

Dustin: I think that is a funny story because we were using GHOSTS to create all of these NPCs that blue teams would defend. We pretty quickly asked the question, *Could we use GHOSTS to do different sorts of activities from the red-team side?* We do all sorts of different kinds of training within our directorate, and some of those are on demand. Those are usually targeting a very specific skill or tool that you want to get better at. I thought it would be really interesting to try to build GHOSTS as a thinking adversary for more advanced, on-demand exercises. I pretty quickly came to the realization, that is a really hard nut to crack. I was sort of disappointed because there was no amount of resources that our team was going to have to build that piece of software that would actually be realistic.

The problem with that is because, based on our experience in red teaming, you have to be very adaptable to whatever the blue team responds. We just don't have the resources to build all those possible responses to responses. But one day you said to me, *If you could just clean up the browsing patterns so that it wasn't so predictable, that would be a huge win*, because of all the problems that you just stated. It is really easy to spot random computer-generated behavior because that is not really what teams would see in their real-world operations. The user-agent string is a great example as well. Again, it was a challenge that got me to thinking about how we might do that. Of course, technologies like machine learning and buzzwords like AI [artificial intelligence] come up pretty often. We, as a team, talked about how could we do some good in this space in a way that other teams could look at it and say, here's a very simple…it's not a simple problem, I shouldn't say that. *Here is a simple solution to a pretty hard problem that you can do yourself on your own hardware.* You can certainly do it in the cloud, and we can talk about that as well. But I thought it was a really interesting way that that project came to light and how we decided to approach it.

Tom: I think that we were a little bit naive in terms of our ability to do a "thinking adversary" easily or how to use ML to create the idea of a persona within NPCs. We thought, well, there are lots of efforts out there from a cloud perspective, using different vendors, different toolkits. We really felt that we could probably accomplish that pretty quickly, but I think we hit a number of roadblocks with that. I think ultimately, I would like to have you talk more about this, because a lot of your effort, was some of the initial paths, some initial concepts, of how we identified NPCs

in terms of having personas, having identities, some of the different things we looked at that actually failed at first. Then, ultimately, the work that you did to actually implement it in GHOSTS.

Dustin: Yes, within our exercise network there are lots of different enclaves: there are operations people, there are MPs, there is the med [medical] unit, there are logistics. All of those different characters do different things by virtue of the role that they play. So we would always have a different list of activities that those users would do in comparison to one another. The talk around, *Well, how can we personalize browsing patterns? We have to create more of a persona for that user*. That was a formalization of those groups that we had that we would manage before. Like I said, logistics versus operations.

So we formalized that and we said, *Hey, the persona for this particular NPC is, they work in operations, and so their concern is operational sorts of activities.* If you think about, *Well, an agent is going to sit down and perform some list of activities.* They would have a preference based on their persona to do the things that align with their job, right? We simply came up with a key-value list of activities, which turns out to be preferences. We assigned those to given personas. A persona is sort of a master bucket of all these preferences. People in operations have preferences to use these particular file shares versus some other, these particular printers versus some other, these particular websites and applications. Typically we use web applications. But they use these web applications to manage their operations day-to-day functions or versus a different list that logistics might use, right? You can imagine how that plays out in large organizations. Different people are doing different things, and this sort of approach just aims to replicate that.

Tom: Yes, I think that is a great first step in terms of making GHOSTS and NPCs more realistic. What challenges did you face from an implementation perspective?

Dustin: We have these preferences, and we can use those to make different decisions. So, *I have a strong preference to do this, or a weak preference to not do that*. But it was hard to figure out which model would best implement that and make future predictions that were aligned with those preferences. I remember we were looking at different technologies available on different cloud providers. One of the really early suggestions that you made was, *Hey, the next site that I am going to browse is not unlike thinking about the list of sites that I can choose from. It's not that different from a site like Amazon recommending the next product for me. Based on what I have done in the past, you are more likely to like or enjoy this smaller list of things that align with that.*

I think we had done some initial testing with different systems that provide that recommendation. It turned out, this was our first step into this area. We were spending a lot of time refining that

model. Part of the problem is you don't want it to get too tightly aligned to a preference, because then all you are doing are the things that, it's more like an *if* statement. We didn't want it to be like that because human beings randomly go to websites here and there out of curiosity, or interest, or whatever, right? We wanted to be able to replicate that. So we were spending a lot of time, which ultimately in the cloud—you could talk more to this—that winds up being more money. We were looking very quickly for cheap ways to, or I should say more cost-effective ways, of doing those sorts of tests, because we were doing lots of tests and trying to figure out the best way to approach the problem.

Tom: I think we went in early on and thought that, well, the cloud providers, the ML services that they provided could really boost what we were trying to do. I think that we were very surprised by two things. Surprised at how expensive the cloud bill was for the first six months. But two, how much it really didn't help us that much, and how we really had to do a grassroots customized design internally that we felt was compatible with what we were trying to do.

Dustin: Yes, for sure. I think that we spent a lot of time getting data into and back out of a format that they expected to use as opposed to just building our model locally and playing with what we already had. That turned out to be much more effective. It turned out to be actually faster to run those tests. You can see some of those results in the paper that came out of that. But we ran a lot of different tests, and we were able to do that I think in part because we just did it on our own hardware.

Tom: So it starts to raise a lot of questions like, how do you validate that success of that persona, being able to preference and align? How did you start to create those personas in the first place? Like, how is someone, a sports persona mixed with a news persona? Things like that.

Dustin: I thought that those are just arbitrary choices for the paper's purpose, because those are the sorts of examples that I think most people can visualize. Obviously, it extends to any NPC that you want to replicate. I think probably the easiest way to think about those is, what role is that NPC playing and what activities do they do in that role for their daily operations? Again, I think any reasonably sized network would have different types of people doing different types of things, and those personas are an easy way to break that out. So you can move your operations people into a particular persona and then still randomize the preferences that they have for that persona. So you get that differentiation between person A and person B on the range. But that would be very different from a person, say, in logistics that is doing different sorts of activities using different kinds of applications and browsing different kinds of websites.

Tom: Stepping back from it, I think it is interesting to also point out here the passion you have for this particular problem. When you first started at the SEI you came in as, say, *I am a developer, software developer, software engineer*. Now you are faced with very interesting

problems. I think that that is one of the attractions for…We see this in some of our data science folks that we brought on recently. In the private sector, they are busy doing tasks related to making money for a particular company and boosting the bottom line. Whereas, in your case, you and our team faced very interesting problems. Like, *OK, how do we make this more realistic from an NPC perspective?* It's not something that you are going to find out there as a typical problem that needs to be solved in those places. So, having an interest and a passion in trying to just incrementally make things better and sometimes making bigger leaps is one of the attractions of working at SEI.

Dustin: I totally agree, and some of those problems we find on our own because we identify, *Hey, here is an area that could be better for training an exercise*. It came up in an after-action report or something like that through the customer. But there are other ones where the customer comes directly to us and says, *Hey, here is a really hard problem. Do you or the SEI have any publications that can assist us in this particular area that we are struggling with, or we are trying to get better at?* We actually have some recent examples of that if you want to talk about them.

Tom: Yes, one of the key aspects here that is concerning SEI as an FFRDC [federally funded research and development center], we are here for technology transfer out to others. We are here to incrementally make things better, improve things, and publish that and make that research available so others can take it and use it and make it better for their own uses.

Dustin: Yes, I love that we were able to open source GHOSTS and that entire ecosystem. There are a couple of other, what do you call them, plugins, or applications that assist in creating the NPCs using publicly available information. We try to create NPCs randomly but that are realistic for different kinds of purposes that you might be interested in. Like when we generated a new range for this season of exercises, we have a tool called ANIMATOR that does that. We have this tool that we are specifically talking about called SPECTRE that increases the realism of the browsing patterns for NPCs that you already have deployed to a range. There will probably be other things that come up in the GHOSTS ecosystem that we will be able to open source. I know we have some very specific tools that help us do different red-team activities. You have a tool that helps diversify IP addresses for that traffic coming into a blue-team network that they are defending..

I always get really excited about that because the range of people that are then able to take advantage of those tools for their own training and exercise scenarios just grows exponentially. We have been contacted by people all over the world, both in the United States for the Guard and Reserve units that want to train and exercise on their own terms; and from the Five Eyes. We are sort of big in Belgium. But I just think that is a great opportunity, and I am glad that the SEI and our customers in the Army enable us to do that.

Tom: I would agree. It is always interesting to see who is interested in some of our, essentially ecosystem, of products that surround GHOSTS and supplement those realism activities within our cyber ranges. It is always very rewarding too to see other people using what we have created or the documents that we have published. Like you really discuss the R-EACTR framework, that's a realism framework we published a number of years ago. There was actually a follow-up to the GHOSTS white paper, a technical publication from SEI. There is a publication for actually implementing realistic cyber ranges in very much more detail in this recent paper from you about NPCs and the ML enhancements for that.

So I can say that in addition though, GHOSTS and NPC and increased NPC browser activity for realism, actually we find out the uses for it as well within our ranges. For example, we are doing some research now on detecting network beacons. We are interested in having GHOSTS generate very, very realistic traffic such that we can produce datasets that can be used to validate our beacon-detection models. Our data-science folks can say, *Hey, I would like to use these methodologies to detect beacons*, new things that they are working on, but they really need to have some datasets that confirm their different technologies or different approaches for detecting beacons. Using GHOSTS to generate traffic, essentially hide the beacons that would exist within the network traffic that the beacon detectors are trying to pick up, is just so critical.

Dustin: I can't believe I didn't mention that. But yes, GHOSTS has been used by other researchers at the SEI, Shing-hon [Lau], and some others to generate that realistic network traffic, and it is awesome to support them. It just occurred to me, do you want to talk a little bit more about what beacons are just so people are clear on what you mean by *beacon detection*?

Tom: Essentially some of the teams that we have worked with previously, some of our customers have come to us and said, *Hey, we have a problem essentially detecting malware command-and-control servers that we call sneaky in the sense that they don't communicate very often*. So you may have some sort of advanced persistent threat, which may be able to gain a foothold within your network and only communicates back to its command-and-control server, say, once or twice a day or once a week. Those sorts of network activity are really, really hard to detect, especially on larger networks that have millions and millions of network-traffic events. We are trying to develop fast methodologies to be able to detect that adversary within your network. The concept of a beacon is that call back to the command-and-control server by the adversary.

Dustin: Can you talk about some of the techniques that we've been looking at for doing that in a new and novel way?

Tom: There have been a lot of different efforts from CERT and different organizations to do beacon detection. Starting last fall, the summer of 2021, fall 2021, we decided to take a more

original, more unique approach to beacon detection where we looked at some machine-learning techniques related to unsupervised machine learning, related to clustering. We applied some different clustering algorithms to actually detecting network beacons. We found actually some pretty good success with that early on and persisted in that research over the past year. But you have been heavily involved in that as well, do you want to comment on it from a clustering perspective in terms of those algorithms?

Dustin: I think the first thing we looked at was *k*-means. Obviously, if your beacon is checking in every five minutes consistently, *k*-means will catch that. But if there's jitter, which I think most systems we have been looking at try to do, then *k*-means will give you a lot of clusters that necessarily don't match up. So we started looking at ways to pick out the outliers. But we found this DBSCAN algorithm was an improved *k*-means where it would throw those outliers out by default.

So we were able to identify different clusters. There are still some challenges with that. Part of it is how far to zoom in or out on a timescale. Like you said, you can have beacons that check in every five minutes, every five days, and so you have to adjust the lens for how you are looking at that data to match that. We are still doing some research in that area.

One of the things that a recently joined team member is making a big difference in, is making that algorithm faster. If you are adjusting the lens, obviously you have to rerun those algorithms on that different set of data that the time span that you are covering. He's done a lot of work, Sean Huff has done a lot of work in speeding that up and making it as close to real-time as possible. Is it real-time? No, not quite yet. But our thought is that, wow, what a tool it would be to be able to provide teams with something that's watching for those beacons continuously. When it finds something with a high likelihood of, *This is a high likelihood of beaconing activity*, it could throw an alert and they could go investigate further.

There are a lot of tools out there that have chipped away at this problem. But one of the things I think you always want to be mindful of is that, if you throw too many alerts, then they sort of become useless because SOC [security operations center] teams don't have time to go and investigate every single one that you throw if you are throwing hundreds of them a day. We were trying to stick to very high-likelihood examples of beaconing that other tools wouldn't catch. Like you were very quick to say, *You should probably ignore anything that is talking a lot because that would just come up in a top-talker report and eventually get investigated and white listed or black listed*. But those other ones that are longer but have a high likelihood, that's something that security teams would want to take a look at.

Tom: The one concept that we always tell the data-science folks helping us with this effort is, we are trying to find the most interesting things to look at. Because you may give a security analyst

at the SOC a hundred things to look at, but if we can help them look at the most interesting things and really narrow that down to just a few items, a few IP addresses, a few websites to take a look at, that really helps a lot from finding advanced persistent threats and finding adversaries in your network.

But, going back to the original, the concept of this podcast which is, we have found yet another reason to use GHOSTS. The realistic and the ML pieces of GHOSTS are really beneficial to helping us generate these datasets from a data-science perspective. If you were to say, *Oh, I want to test my new beacon algorithm*, well, good luck going out there on the internet somewhere and finding the dataset that has source IP diversity from hundreds of different source IPs going to thousands of destination websites all doing realistic browsing as well as mixing in active-directory traffic and file-server traffic and other internal SharePoint traffic, things like that. And finding that data so you can test your beacon algorithm is really hard to do. It is not really publicly available out there. Whereas, we can generate those datasets alone provided by the GHOSTS activities, it is just super valuable to us.

Dustin: It's all just layers of realism, right? That wouldn't be possible if we didn't have a best-in-class range that allows us to run GHOSTS on it and have that sort of diversity. But, yes, I agree. That is actually one of the things we never talk about, having that ability to run those sorts of experiments on a range like that, with that kind of diversity of IPs and such.

Tom: Yes, and the benefit to the data-science folks is that, if they want a change to the dataset, or they want more beacons, or less beacons, in our range, we can generate more traffic with the aid of GHOSTS and generate that on demand. If they want a week's worth of data, a day's worth of data, a month's worth of data, we can refine that and get that to them to continue their research.

Dustin: So what do we do next? I mean, not to put you on the spot, but what should we build next? What is interesting?

Tom: That is always the big question. I think that we are just going to continue to refine these, the idea of a more realistic GHOSTS and having GHOSTS do more things. Can it be a chat, can it be chatting to a team server? Is it accessing SharePoint? Is it uploading, downloading content? Some of that, those features are already available, just continual refinement of that work I think in terms of GHOSTS. Just find other aspects on the range to make it more and more realistic. What are your perspectives on next steps for some of our efforts?

Dustin: I can't disagree with any of those. One of the hard things still is, if you wanted to script out very specific commands for a group to do, that can be time-consuming and sort of tedious. I don't know if the answer is to make the scripting language for doing that easier and faster to use,

or to build some tool that automates the creation of those user-activity scripts or whatever we want to call them. That has always been a challenge. I think we could probably make progress in that area. I agree other tools, being able to automate other tools. Some of those are difficult though, because by definition you don't want them to be automated because bad people on the internet use them for nefarious purposes. But I agree, like chatting is a big part of network activity that we should be able to support. It is probably some other things around that ecosystem like you said, downloading docs from SharePoint, or updating the wiki page, or something like that.

Tom: I think that ultimately that goes back to the concept of the original blog and this paper on NPCs and ML, which is, having that ML piece and the persona piece and the customization piece to make these NPCs even more realistic is really the big bundle of bringing all these things together.

Dustin: Any other application of ML to make that NPC make more realistic decisions, I would definitely be interested in that.

Tom: As usual we could probably talk for hours about these areas and you can see that, like I said before, the passion you have for pushing this forward, it's not something that you stop thinking about it at five o'clock. You are thinking about it, and you find it as an interesting problem to solve. It is good that it really aligns with problems that are very critical to the defense of the United States from a cybersecurity perspective going forward. Do you have any final closing comments?

Dustin: Yes, I think we have covered a lot of things, good stuff. I actually can't stop thinking of it, because I am still a graduate student at Carnegie Mellon's School of Philosophy. There are lots of interesting courses on decision theory, game theory. Some of the concepts that made it into GHOSTS came directly out of some of the studies I have done there.

Our team has been great in augmenting the things that I am able to do. You are more of a systems network person and I am a software person. Put us together, and that works really well. I think our team really does a good job of that diversification of skills. I suspect that cybersecurity needs more of that. Because the challenges that we see, a lot of the issues that we try to replicate in scenarios are not always technical problems. They are process problems or people problems. They are social issues. I think many of the things that organizations are struggling with today are not necessarily just tech problems. So maybe we need some other viewpoints on how to solve some of those issues. I hope that GHOSTS is someday able to replicate some of those problems, because you can coordinate activity between different agents and agents can tell each other different bits of information, and that could obviously influence the decisions that they make,

update their preferences, and all that. We will just keep evolving it until it starts doing things in exercises that we can't really account for, and we have to dial it back a little bit.

Tom: Yes, that would be pretty exciting to see that, a little scary I guess too. But I 100 percent agree that not all of these cyber issues can be solved with just straight cyber people. There are a lot of different types of people that come to SEI that have lots of different skill sets that are great fits for helping us solve these sorts of cyber challenges that are so critical. We will wrap up here. A quick note to our audience. A transcript will include links to all the resources mentioned in this podcast. We talked a lot about some white papers, some blogs, and actually some open source software that exists out there already because GHOSTS is open source. So those will all be included as part of the transcript.

Dustin, I really appreciate you talking with us today. It's been a very great conversation, I am sure we could talk for a lot longer. We will wrap up here. Finally, a reminder to our audience, our podcasts are available on SoundCloud, Stitcher, Apple Podcasts, and Google Podcasts, as well as SEI's YouTube channel. If you like what you see and hear today, give us a thumbs up. We would really appreciate it. Thanks again everyone for joining us, and we will see you soon.

Thanks for joining us. This episode is available where you download podcasts, including SoundCloud, Stitcher, TuneIn Radio, Google Podcasts, and Apple Podcasts. It is also available on the SEI website at sei.cmu.edu/podcasts and the SEI's YouTube channel. This copyrighted work is made available through the Software Engineering Institute, a federally funded research and development center sponsored by the U.S. Department of Defense. For more information about the SEI and this work, please visit www.sei.cmu.edu. As always, if you have any questions, please don't hesitate to email us at info@sei.cmu.edu. Thank you.