Threat Modeling and the Internet of Things

featuring Allen Householder and Art Manion as Interviewed by Suzanne Miller

---

Suzanne Miller: 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 and operated by Carnegie Mellon University. A transcript of today's podcast is posted on the SEI website at sei.cmu.edu/podcasts.

My name is Suzanne Miller. I am a principal researcher here at the SEI. Today I am very pleased to introduce to you to Allen Householder and Art Manion, both who are researchers in CERT's vulnerability analysis team. Today, we are here to talk about their work on threat modeling. Before we begin, let me tell you little bit about our guests.

Allen Householder, who is new to our show, is a senior vulnerability and incident researcher at the SEI's CERT Division. He has been involved in internet security since his first professional job in 1995, where a few weeks after starting at a Fortune 500 company he was told, *You are the IP and DNS guy*. Think about what that meant in 1995, and shortly thereafter, [he] was given the responsibility for the entire corporate firewall.

His recent work includes being the technical lead developer for the CERT Basic Fuzzing Framework, or BFF, and Failure Observation Engine, also called FOE, and research into the security, or insecurity, as you think about it, of the Internet of Things. His research interests include applications of machine learning and software and system security, fuzzing, and modeling of information sharing and trust among Computer Security Incident Response Teams, which we call CSIRTs.

Art Manion is a senior member of the Vulnerability Analysis team in the SEI's CERT Division. He has studied vulnerabilities and coordinated responsible disclosure efforts since joining CERT in 2001, where he gained mild notoriety for saying, *Don't use Internet Explorer* in a conference presentation. Manion currently focuses on projects including software component relationships, vulnerability management, and standards of development.

Prior to joining the SEI, Manion was the director of network infrastructure at Juniata College. Welcome Art and Allen. Thank you for joining us.

Art Manion: Thank you.

Allen Householder: Thanks for having us.

Suzanne: Let us start by having you talk to us about what threat modeling is, and how does it figure into today's complex cybersecurity landscape?

Allen: Threat modeling is a process that helps to reason about a system, a system that you care about its security. It has been popularized by Microsoft over the last 10 or 11 years. They actually published a book called *Threat Modeling* in 2004, and that went through a few editions. There is a new book by Adam Shostack called *Threat Modeling: Designing for Security*, which came out in 2014. That is probably the current definitive resource for learning about threat modeling, getting started with it, and understanding the landscape.

Suzanne: Threat modeling is not really what you think it is. Threat modeling is not about modeling the threats that could happen. It is really about modeling aspects of the system and how it responds. So, say a little bit more about that.

Allen: Right. It is really a way of thinking about a system and understanding the various attack surfaces that system may have. There might be multiple layers of attack surface. In the same way you can think of say a prison or a fort might have a fence at the outside. They might have locks on the doors at the perimeter. They may also have rooms inside that are locked and then a safe inside that. All of those things could be different attack surfaces, and you can address vulnerabilities at each of those.

Suzanne: Each of them have different vulnerabilities, so part of threat modeling is understanding the different character of vulnerabilities that could happen in different layers.

Allen: Right, and also understanding the assets that you are trying to protect and where they are and which possibilities you have for defending them.

Suzanne: OK. You use this concept of threat modeling in your research. How do you do that?

Allen: We have actually used it a few times in research we have done on Internet of Things devices as well as on automotive systems, connected automotive systems, in part because we often are asked to analyze systems that we weren't the developers on, we haven't necessarily been involved in, but somebody wants to know, *What vulnerabilities should I be concerned about in this kind of system*? So, we will go out and understand that system, and part of the way we do that is by using threat modeling.

Suzanne: So you are looking for patterns that can be seen after the system is already in operation because you don't have understanding necessarily of the design parameters, the data parameters, the things that the developers would know. You have to look at the system as it stands...

Allen: As it is. Right. Right.

Suzanne: So, that is a little different than trying to understand the vulnerabilities you are introducing as you design a system. That is another aspect of this, too, right? The data assurance kind of work is a design aspect, and how does threat modeling play into that.

Art: You talked about sort of the finished system, and you can certainly do threat modeling there, but it is probably worth mentioning you could use threat modeling techniques at different points in the process, and that could include not having a finished product but just having…. If you did have design specs, you could still reason about the design specs as to have a network interface, so that might be something you consider an exposure point, which could be a path for a threat. Even without a finished system, you can potentially do some threat modeling about the system you are trying to work on.

Suzanne: I am assuming that the more modeled, as it were, your system is, the easier it is to do that. When you are looking at documents, you have to create mental models, but we have some modeling languages that are starting to become used that actually give you more explicit models about that. Is that something that you are starting to work with?

Art: I am not personally, but to your point: yes, if the system is already well modeled, it is probably an easier lift from there to doing some threat modeling. A lot of the threat modeling techniques we have looked at, they assume you do not have that in place already. So, they have techniques. In fact, there is actually a card game to get you to think about the exposures and surfaces. You do just enough modeling to do some of the basic threat analysis.

Suzanne: So, it is possible to do reasonable threat modeling without having very complicated kinds of models that are already in place?

Art: Sure. There are relatively lighter-weight threat modeling approaches that are still useful.

Allen: It can be as simple as drawing diagrams of the system and talking through those diagrams with experts.

Suzanne: Classic white board approach.

Allen: Right. Art mentioned the card game, which is an interesting way of facilitating that conversation. The card game has various prompts. The card game is called Elevation of Privilege. It came out from Microsoft. It has various conversation prompts that suggest ways that

you might have different problems in a system that you should explore. Given the system that you are analyzing, *Does this threat apply? Is it relevant?*

Suzanne: Is this relevant?

Allen: And you can play it that way. Some of the research I have been involved in was with the Architecture Analysis and Design Language, AADL.

Suzanne: Yes, AADL. I was thinking about that when I was talking about modeling systems.

Allen: The reason that became interesting is because one of the threat modeling techniques is attack trees. An attack tree is essentially just a tree diagram. The root of the tree is a bad thing that the attacker can do or an event you don't want to occur. Then each of the branches of those trees are different preconditions that lead up to that bad event. So, someone stole your car while you left the car unlocked and you left your keys in the car, and those are the two events. The way you would mitigate that is take your keys with you and lock your car.

Attack trees are also very closely related to fault trees. AADL has been designed.…We have done some extensions at the SEI, on applying fault tree analysis to AADL models.

Suzanne: Mostly for safety, but I can see where this would translate very easily to the threat modeling.

Allen: This past year we actually had a project where we looked at applying or doing getting threat models out of the AADL models. One of the things we found is that, as security analysts, learning AADL is pretty tricky. If you already know AADL, then there's a potential there for...

Suzanne: For collaboration.

Allen: For collaboration. That is actually going to be some ongoing research that we have got going this year to continue that work in AADL crossing over into threat modeling land.

Suzanne: Cool. Are there specific modeling languages and things that are being developed to make the threat modeling itself easier to do? Is that a direction that we are engaged in, or others that we are collaborating with are engaged in?

Allen: Not that I'm aware of.

Suzanne: So this has not yet made that sort of leap into, *We need to have our own toolset for this.*

Allen: Right. One of the things that I noticed, actually, when we were doing the Internet of Things modeling project was, in investigating attack trees and recognizing that connection there

to fault trees, I went and looked at a lot of research on fault tree analysis. Fault tree analysis has made it to where NASA is using fault tree analysis as math. They are treating it as math and modeling, and you can do calculations and proofs and those sorts of things. Whereas, attack trees have kind of gone the way of drawing diagrams and pictures in Visio. But there's no...

Suzanne: There's not a syntax or semantics to it.

Allen: There is no deep analysis and things going on there. That is actually a potential area for future work, is bringing some of the fault tree analysis techniques back into the attack-tree land and…

Suzanne: …looking for synergies between the language approaches and fault trees in the domain of threat modeling and attack trees in particular. Is that where your research is going? Where is your work going in relation to threats and how you use them in research?

Allen: Like I said there is the AADL threat modeling project that is going on. Our work has been more on vulnerability discovery. We are asked to *Take a look at this system and let us know what sort of vulnerabilities it has or might have*. Organizations that want to direct their testing on a system that they are not familiar with. We have done that for Internet of Things devices.

We have looked at an internet-connected light bulb system, which was kind of interesting because the system itself has a light bulb that has a wireless connection to a little device that sits on your LAN [local area network]. That device, in turn, can talk to a cloud-based service, which also has an Android or IOS for your phone. You can press some buttons on the app on your phone that, up to a cloud service, comes back to the little device on your network, which then tells the light bulbs to blink or change colors or anything like that. We built a threat model for that which let us see that there is a lot of different attack surfaces on there.

Suzanne: Even without your card game, I can think of several.

Allen: But it helps us direct our testing so that we were able to then use the CERT Tapioca tool to do man-in-the-middle attacks and do an analysis of the gateway device. We actually found some vulnerabilities in that, which turned out were part of the operating system that the little device runs.

From that knowledge we were able to then go look for those vulnerabilities in other things that ran the same operating system, eventually leading us to a lot of home routers and problems that are occurring in home routers as well. Exact same problem, just a different domain. In one case it is the Internet of Things, the other is home routers. That is one way the threat modeling has been useful to us.

The only place that we have used it is in the AADL work that I mentioned. The system that we were modeling was an internet-connected car. And we did not have a car.

Suzanne: Lots of attack surfaces in that domain.

Allen: Right. Actually, the interesting thing about that is, we did this modeling in 2014. One of the things that came out of it was, *Well, internet-connected radio can talk to the devices in the car, which includes potentially the adaptive cruise control that controls the throttle and the brakes on the car.*

Many of the attacks that were demonstrated this year at Black Hat and DEFCON were things that came up in our threat modeling game that we played almost a year ago. We did not have a car. We were not analyzing the car directly. We were just drawing pictures on a white board, talking through the process, and we came up with most of vulnerabilities that, in turn, were validated by other people's research that came out at this year's security conferences.

Suzanne: Really, in terms of organizations who would want to take advantage of this research, one of the takeaways is you do not have to be completely sophisticated to do this. You can just use thought experiments, some heuristics in terms of how to look at these problems and get some reasonable vulnerabilities that you can then use to inform your testing, to inform your supply chain choices, other things like that. That is actually very important for organizations to be able to have that understanding, that they do not have to have the most sophisticated tools in the world to be able to address this problem.

Art: Right. I would like to really reinforce that point. There is another aspect of our work that is much more operational. We receive reports of security bugs and vulnerabilities and things. We process them and try to see them through to a coordinated disclosure process with fix software and announcements and minimize harm to everyone involved.

This happens, and we are seeing…. Allen has mentioned light bulbs so far, cars. He said *Internet of Thing*s a couple of times. There are lots and lots of things that are connected. The manufacturers making these things might have been a business for 50 or 60 years. They are great at making cars or refrigerators or light bulbs. They have now, in some cases, literally bolted on a small embedded computer with a number of network connections.

My impression, and I cannot prove this in any way, is that there is not a lot of threat modeling going on. They know what the refrigerator does. When you stick the embedded network connection to it, there was not a threat modeling process to say, *Now, what is changed about the model? Why do you care about security of your refrigerator*? It keeps the food cold. It is electrically safe. I do not know what else it has to be. Now it has to have, potentially updates, modern software.

Suzanne: And if it is connected to my electrical system, if I am trying to do adaptive power management, now it is in my electrical system, and now my electrical system has a vulnerability. Yes.

Art: It is just an observation based on the reports we are seeing. Home routers are another example. There are these device-like, thing-like connected things, devices, IOT, that they are new to the Internet connectivity. They haven't done some basic threat modeling with that change to the system. We are seeing vulnerabilities, debug ports left turned on, default passwords, very basic types of vulnerabilities that threat modeling probably would have caught even very lightweight threat modeling.

Suzanne: One of the things we might see, I don't know if this is part of your standards work, is, I can envision places like the UL labs adding standards that require this as part of their certification. Because, as you say, the people that have been developing these kinds of things have done it without, really, knowledge or understanding of the implications of security. As soon as they add any kind of computing power onto their light bulb or anything else.

Art: Another new area here, as you mentioned UL, light bulbs might not be… well, UL might cover light bulbs, but vehicles, avionics, these are regulated industries already. They already have rules for safety in place. Now you have to consider core network and computer security has to be part of the safety.

Suzanne: It is part of safety.

Art: In these cases. We are talking with those regulators, in fact, because that is a new area for them as well.

Suzanne: There is standards work. There is tools work. I think you guys are going to be busy dealing with this stuff for a little while.

Allen: We already are. Yes.

Suzanne: I do want to thank you both for joining us. I think you are probably going to scare a few people that need to be thinking about things a little differently, and sometimes that is part of our job.

I do want to tell our listeners that if you are interested in learning more about this work, you want to visit the CERT/CC blog, that is the Coordination Center blog, and that is at insights.sei.cmu.edu.

I want to thank you for joining us today. Remember that today's podcast is going to be housed at sei.cmu.edu/podcasts It is also going to be available on the Carnegie Mellon University's iTunes

U site. As always, if you have any questions, please don't hesitate to email us at info@sei.cmu.edu. Thank you for listening. Thank you for watching.