Zero-Trust Adoption: Benefits, Applications, Resources

Featuring Geoff Sanders as Interviewed by 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. A transcript of today's podcast is posted on the SEI website at sei.cmu.edu/podcasts.

Suzanne Miller: Welcome to the SEI Podcast Series. My name is Suzanne Miller, and I am a principal researcher in the SEI Software Solutions Division. My guest for today's podcast is Geoffrey Sanders, a senior network defense analyst in the SEI's CERT Division. Today we are going to talking about Geoffrey's work with something called zero trust adoption. Welcome, Geoff.

Geoff Sanders: Hey, thanks Suze. Thanks for having me.

Suzanne: I am excited about this because for some of the customers that I work with, this is a hot topic. So, before we get into the topic itself, tell our viewers a little bit about yourself and the work that you do here that led you into this zero trust adoption research.

Geoff: I have been blessed with a broad career in security operations, a large background in security architecture and engineering, intrusion analysis, penetration testing, vulnerability analysis, and even organizational leadership roles over the past 20-plus years. November will be my 10th year at the SEI. Through our work and our sponsorship, I have been able to leverage a lot of that practical experience for research. It is always a great adventure every day because we always do interesting things. Day to day is never a standard thing really, but most of my work focuses on security operations and situational awareness. We have had a large situational awareness research area for a long time at the CERT Division under SEI. This typically incorporates analysis of large datasets and the architecture and engineering for security operations that goes around that, especially like intrusion analysis and things like that, right?

There is more and more data coming in that you are trying to identify—adversary activity and things like that. Through this, I am privileged to participate in different projects, such as tech transition for the CERT Division's SiLK Tool analysis suite for network flow data collection and analysis. So, I have helped author some of our publicly released books on that and have trained other folks how to use that, even to the extent of implementing cyber threat frameworks for cyber intelligence, and that was done in the Cyber intelligence Tradecraft Report we did 2019, I believe. So, yes.

Suzanne: So, a lot of broad experience and applying it into this relatively new area in terms of security called zero trust. This is not a familiar term to a lot of our viewers. You did put a blog post out recently on this subject, and in that you wrote, *Zero trust adoption challenges many organizations. It isn't a specific technology to adopt, but a security initiative that an enterprise must understand, interpret, and implement*. So, this is a little bit outside of our normal SiLK tool technology stuff.

Geoff: Correct.

Suzanne: Why don't you explain what is different about zero trust and how it becomes this broader kind of topic?

Geoff: Sure, sure. There is a little bit to unpack there, so let me set some context before I get into the details. Fundamentally, zero trust is a security model, and those are sets of system design principles combined with a governance strategy. They determine how people, processes, and technology work together to implement security for an organization. It is usually technology in general. People refer to it as the network or the enterprise or something like that. But it's more than just technology. It is multiple things working together to provide that security solution.

Historically, cybersecurity has used a network security model, kind of like a castle/moat approach where the network functions as the drawbridge, moat, and wall to implement security for an organization, which would be termed the *castle*. A good example of this is VPNs, virtual private networks. These act as gateways into an organization's sensitive networking data, and they only allow authorized users access to that specific network or specific data that they can access. One of the larger problems with this model that we have encountered is the implied trust that goes with it. The model assumes that users authenticated through a VPN are authorized.

When a user is compromised, the VPN just sees that user coming through, and there is really no tracking or any additional things typically that go on to see what that behavior is. So, zero trust modernizes this approach by removing that implied trust, hence the name zero trust. *We are going to trust nothing implicitly*. With evolving technologies, such as mobile devices, Internet of Things, IoT, cloud computing, the network perimeter is no longer easily defined.

Suzanne: Yes.

Geoff: So, the previous approach of slapping a VPN on a gateway to protect your internal network and its data, doesn't scale well anymore, because you can't really define those boundaries as easily as you used to. Even organizations that are doing cloud understand this because they get in this hybrid model where they have their own data centers and external clouds. They are really trying to identify these boundaries and really can't anymore. Now you got to take the security from the network, and you have to move it to a different level. In this new model we are seeing with cloud and IoT, users can be humans. They can be applications. They can be devices that need to access those services. They can be anywhere and really belong to anybody. Just because you are using cloud doesn't mean it's necessarily your resource, right?

Suzanne: OK.

Geoff: Zero trust security moves that model from the network to each session between a user, which is called a *subject*, and a service, which is called a *resource*. So, instead of trying to define your security at a computer level, which is like an IP address or something like that, we are actually saying, *This user is trying to access this particular service, and they have these attributes about them, and we have to start doing dynamic threat assessment*. They call it dynamic policy. This new approach removes that implied static trust at the network level, implements dynamic policy by merging multiple elements together, such as the person's identity, the device they are using and the risk that device presents, the subject location, meaning that user, where is that user located? Are they accessing things at a time they are not supposed to? Are they coming from a location they typically aren't supposed to be in? Those types of things. The requested resource they are trying to access, is a service, right? And, all kinds of other elements at the service level, right? Because this is a security model, and it is not one specific technology, organizations really must understand the model and how it relates to everything in the organization, including the organization itself. They have to have a real deep understanding of the people, the process, and the technology, along with the as-is and to-be architecture that it will take to get to that successful journey to implement all these things. From a simplistic level, it is a security model, but it touches everything in your enterprise across people, process, and technology, along with requiring that you really understand all those elements, and it ties all of that together.

Most organizations are really used to buying a particular technology and implementing it, even with cloud. *We are going to buy cloud, and we are going to implement it in the cloud*. But, when you start trying to implement security policy on all those things, you really have to have a good understanding of how all those things relate.

Suzanne: I am thinking of examples where this would impact people differently than happens now. In my case, outside of COVID time, I am on the road to different customer sites two-to-three weeks a month. If my IT security policy doesn't account for the fact that I am in multiple locations, then this zero trust position could put me in a position where I can't access the resource because, *You are not in Pittsburgh, and that is where all of our resources are expected to access our cloud services,* for example. So, that could happen if I don't have that enterprise understanding of, *But here's how we want the policies to work. There are certain people that travel. We have to allow for multiple locations for them. There are other people that don't.* Zero trust would extend to, *You are not where I expect you to be*. Is that a fair way to characterize one the things that might happen in this space?

Geoff: Sure, sure. I think one of the easier ways to relate it would be when you look at zero trust, it's raising the bar, so rather than just a standard performing organization, you really need to be a high performing organization. You have to really leverage different things like automation, distributed services, distributed compute, and be able to tie all of those things together. So typically, because that is a difficult thing, that is why you have the current network approach of saying, *They come in on the VPN. They have authorized. They can access these areas of the network.* Whereas zero trust raises it up to say you have a person authenticating into a service, which is not necessarily a VPN anymore. Because of all these remote users, a lot of people are putting things outside of their VPN and into the cloud, right?

Now you are having to authenticate that person into the service, and now you are wanting to look at, *Is their system running malware? Does it have any malware on it? Is it an asset that is controlled by the organization or is it a personal, bring-your-own-device type situation and all of those dynamic things come into play with zero trust now?* So, you are having to try to tie all of that together and make sure you have a good understanding of all those components working at the same level.

Suzanne: You talked about raising the bar and going to a more high-performing organization that essentially has a lot of deep and broad knowledge about what is expected in terms of how its services are used and in terms of the profiles of the users that are going to be using it. What are some of the other prerequisites, either organizationally or technologically, that need to be in play if someone wants to go this route of zero trust?

Geoff: Well, we term it broadly as like situational awareness. These environments are having to tie some dynamic policy together. Typically the way this is currently managed when we talk security operations, is you are relying on devices to generate alerts, and then have a human analyze those alerts to make a decision if it was good or bad. But the activities, in many cases, is still allowed to occur because there has just has been an event identified.

The goal of zero trust is to say, *All these things come together. This policy is applied, and for some reason, this activity looks strange, so we are not going to allow access to that resource.* That is kind of occurring on the fly, so you are trying to remove a lot of the latency in that human middle man making those types of decisions, and you are wanting to automate all of that. The implication would be that you are automating a lot of this understanding. You have a very deep understanding of your network and your environment to build these policies. You have a very good asset management capability, and it is continually updated. It isn't a static process where, *OK, here is the asset. It has been deployed. It has been handed over, and it never gets touched until there is a tech refresh*? There are dynamic things going on to understand the profile of that asset as it's working on the service, as it's accessing the service to say, *Oh, we see a different risk profile now, so now we are either going to deny, or we are going to put it into a different sandbox zone to do some remediation or something on there*. That is really where you are raising the bar is you are automating these things. You are defining what these things look like, and you are building the infrastructure and the enterprise architecture to handle it.

Suzanne: That level of automation, I think it really assumes that you are, for lack of a better word, working on modern processors. Your hardware-based technologies are modern enough that the latency that would be implied with some of this automation isn't going to affect you as much as it would have done 20 years ago with the processors that were available then. Is that a correct assumption?

Geoff: Well, technology is part of it. One of the innovators of this that I've seen in the past is Google and that is mainly because they have published their work for what they call the Beyond Corp model. They have done multiple papers on this to show you how they went about engineering the solution and things like that. It isn't always that you have high performing hardware. It is that you are engineering a distributed solution that can leverage that technology, of course, but understands that there are static states of existence. It's not going to be dynamic all the time. That threat profile we talked about where you are assessing that endpoint to understand it's now changed its risk profile. That would be its own specific service that is doing that type of thing, and that would have to access the asset management. It would have to access the identity management, maybe the HR database to understand who that user is. That itself would mesh all that information just to come up with that kind of dynamic risk assessment over time.

Suzanne: OK. So, you mentioned engineering, and there is a subdiscipline in security called cybersecurity engineering. We have had some podcasts on that topic. What role does that play? It sounds like that really is an essential element of this that the engineering of the entire enterprise, not just the network and not just the devices.

Geoff: Actually, cybersecurity engineering is fundamental to zero trust adoption. We call it cybersecurity engineering. Some organizations may just call it engineering and architecture, but

the fundamental piece is that you are combining the engineering rigor with operational security and building that into all aspects of the system lifecycle. Specifically, when you are trying to determine risk, it takes the three core areas: the people, process, and technology, and assesses them from multiple perspectives.

One of the bigger things is people look at zero trust, and one of the commercial solutions in implementing some of these ideas is called software defined printer. People will look to buy that solution and implement it. In some cases, they are following some type of road map or something, but they are not looking at it from the business-mission perspective. In addition to the security, engineering, and architecture-risk perspective, and then the acquisition perspective. So, there are multiple viewpoints to look at this process of implementing this journey from as-is to be. And, in many cases, it is always just a technology focus. The cybersecurity engineering that we do takes all of these aspects into play including the fact that you have multiple systems, a system of systems architecture that is putting all of this together. Our research has had unique opportunities to be able to look at those multiple perspectives, all of those systems together, and really pull out risks and things that people are missing on this journey.

Suzanne: When I think about some of our customers that have very…Their products themselves are at a very high security posture. There is a need in my mind that the architecture of that product is reflective of that zero trust adoption cybersecurity-engineering perspective, so talk a little bit about sort of architectural implications for products that are living in and are being leveraged in this very zero trust adoption kind of environment.

Geoff: Well, so one of the larger things a lot of folks are looking at right now is DevOps, or DevSecOps because now you are trying to incorporate security into your DevOps environment, and that is a very Agile process, right? Because it is Agile that drives it. So some of the challenges, especially when you talk from a DoD perspective are things like authority to operate.

Suzanne: Yes.

Geoff: Now you are trying to tie these two things together, this Agile environment and some automated way of putting all that together. So, I think the larger challenge is looking at how that technology works, defining how it works, looking at the areas of threat, so this implied trust, right? If I am a developer, and I am writing code, then I am an entry point for an attack.

One of the larger things that the NIST Zero Trust Architecture, document 800-207, I believe it is, they specify what are called non-person entities as one of the higher risk areas in these architectures. That is because they are nonhuman. When we implement zero trust, you are having an understanding of, *This is a human or this is a nonhuman*, because with a human I can deploy things like two-factor authentication and strong authentication. I have different ways of doing

that. But when it is an automated program or device or something, I can only do things like certificates. When you are building these complex architectures and pipelines, you have to start really understanding where your risk is, not only from a technology perspective, but the people and the process too. So, that coder with the IDE who is writing code and pushing that up to a code repository, there may be people and process risks in there that you are not taking into account as part of the pipeline. Those are some of the architectural implications of looking at this from a much broader perspective.

Suzanne: There are architectural patterns that have become much more common in the DevSecOps kind of world, the microservices architecture is the one pattern that comes to mind.

Geoff: Correct.

Suzanne: Are there specific advantages or disadvantages to architectural patterns like that within organizations that are trying to use zero trust?

Geoff: Well, the advantages are the main reason everyone uses those types of patterns; the automation, the ability to do some repeatability and to apply policy in a particular way, repetitively and consistently. The challenge there is how do you understand that your policy is the way it is supposed to be, and that it is actually being deployed in the manner you assume it is. Those are some of the people and process risks you're looking at, right? I'm automating something with Kubernetes in a container. Well, how am I validating that the configuration is what I assume it to be? How do I understand what the aggregate policy is that I am deploying on those containers? And, how am I analyzing the behavior of those environments to make sure it is acting in accordance with the way I expect it to be?

Suzanne: I am guessing those are some of the areas of research that we are probably working in and trying to figure out how do we make some of those things more obvious, that we may have -- I'm just saying you have a case where people are trying to reuse, take advantage of reuse and that's a place where the Kubernetes container policy in this environment, then the policy in this environment may actually not be the same or may not need to be the same, and if that's not something that is highlighted, the automation may make it seem like, *We're fine*. And, really what you're saying is we have to still maintain that posture of zero trust. Don't assume that what was, you know, true over here is necessarily going to be true over there and finding ways to make that obvious seems like something that we would want to be helping those kinds of resources with. Is that one of the areas that we're working in right now?

Geoff: Well, kind of what we are doing is we are looking at how we do the cybersecurity engineering, these assessments as how to look at all these research areas we have done with like

mission risk diagnostic and SERA and to fold the zero trust specifics into those and to apply those to the systems-of-systems work that we do all the way from mission to security engineering to acquisition.

Some of the larger implications that you are talking about though is yes, you want to be able to define those things, but part of what the zero trust architecture includes is that you are monitoring for these types of things. This is where the situational awareness piece comes in. You are not waiting for necessarily an alert for one specific event. You are trying to tie these events together to paint a picture that hopefully your automation will mitigate early enough in the cycle, because when you keep trying to push events to a human, they are trying to go through a bunch of data themselves and make those interpretations. So, you want to be able to define your environments well and then monitor against that baseline definition to understand what the change is or the deviations or the anomalies are.

Suzanne: When you get into system of systems environments with multiple governance structures, I can see where would be challenging.

Geoff: Right.

Suzanne: Yes. So, OK. We have kind of explored a lot of the edges of that space, but we also like to talk about transition, and you have mentioned that you have been involved in other kinds of transition activities. For people that are new to this and that are interested in trying to apply zero trust adoption in their organizations, what resources are available, both SEI, non-SEI? Where should be people start in bringing that kind of thinking into their organization?

Geoff: Well, one of the first places I always point folks to is the NIST Special Publication 800-207, which is Zero Trust Architecture. It is a must-read. It gives you all the architectural patterns of what makes up zero trust. It is kind of foundational to zero trust architecture, and it is background information organizations should understand and incorporate into their journey. There is also a really great new publication that was released called Zero Trust Security. It is by Garvis and Chapman through Apress. These authors have actually helped implement zero trust in several organizations. There is a lot of not necessarily lessons learned in there but the actual implementation of particular patterns and how zero trust is implemented in this journey and the things you should be aware of and what is practical real-world experience advice what you might just come to mind when you start reading some of these higher level architecture documents

It just provides a lot of great detail from organizations that have implemented these patterns. I also like to reference folks to the Google Beyond Corp initiative, which is their native zero trust implementation. They were early adopters. A lot of their engineering solutions and why they did what they did and the choices they made leveraging common technology and with custom

development, you can actually look at how a large organization approached this. Keep in mind it is high performing. They are able to build these own solutions themselves.

Suzanne: OK.

Geoff: Keep that mind as well. Last but not least, our work with the cybersecurity engineering and how we have basically stepped [in] to pull these multiple views together and understand risks from multiple dimensions, not just what you would call tactical risk for one particular event on one system. We are looking at the aggregate events together in all of these systems put together.

Suzanne: And we tend to look at things from a mission viewpoint, which is different from some of the other risk taxonomies that are out there.

Geoff: Correct. The mission risk diagnostic is one of those specific ones that looks at mission itself.

Suzanne: I want to thank you for talking about this with us today, Geoff. I think this is news to some people. It is old hat to others, but it is definitely worth talking about how we need to be changing our posture as the world becomes more dynamic. Fundamentally, that is really what we are talking about, is the static view of security is no longer sufficient for us to really guarantee that we are minimizing threats to our organizations. I want to thank you for highlighting that and talking about that really clearly, because this is not stuff that is always easy to talk about.

Geoff: Oh, yes.

Suzanne: We have several things that we talked about that will be included as links in the transcript, the NIST publication and others, so our viewers can look there for in the transcript particularly for any of those resources. Thank you again, and I look forward to seeing how some of this research plays out over the next couple years.

Geoff: Oh, great. Thanks, Suze. It is really an exciting area to be working in right now.

Suzanne: It is. It is. Thank you to all of you who are viewing the publication today, the podcast today, and have a wonderful rest of your day.

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.