Insider Threat in the Post-Pandemic Workplace

featuring Matt Butkovic and Dan Costa as Interviewed by Matthew Butkovic

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.

Matthew Butkovic: Hello, welcome to the SEI Podcast Series. I am Matthew Butkovic. I am the director of Cyber Risk and Resilience here at the Software Engineering Institute. Today, I am joined by two of my colleagues, Dan Costa and Randy Trzeciak. We are here to discuss insider risk and the mitigation of insider risk. Welcome, gentlemen. Thanks for joining us.

Dan Costa: Thanks, Matt.

Randy Trzeciak: Thank you.

Matthew: I was hoping that you could start by just maybe explaining your roles here and maybe your path to the SEI. Let's start with Dan.

Dan: I am the technical manager for Enterprise Threat and Vulnerability Management, leading two bodies of work here: our Insider Risk team as well as our Applied Network Defense group. I have been at the SEI for 12 years

now, and I came here by way of a short career as a software engineer. I actually came here to do some contract software development, helping to write assessment data collection and analysis capabilities for an insider threat assessment that the team was currently developing. Since [I arrived] here, I have had a chance to wear a lot of different hats, within this Insider Risk team. As an individual contributor, then team lead, now, technical manager, I have had the great pleasure of going a lot of different places, working with a lot of different people on some really interesting problems, whether that was on afloat platforms, or at the Pentagon, or within the private sector. We have really jammed a lot of life into the 12 years I have had here at the SEI, and it is one of the reasons why I am still here.

Matthew: Excellent. Randy?

Randy: Thanks, Matt. So I am the deputy director of Cyber Risk and Resilience here within the CERT Division of the SEI. My transition to an educational environment started when I was working at the Naval Research Laboratory in Washington, DC. My background was historically in database development, design, application building, software development. Really one of the foundational components of database is database security, which was a great transition to focus on cybersecurity when it comes to protecting of data and critical assets within organizations. I was fortunate to have worked in a number of places across Carnegie Mellon [University] including different parts of the SEI, but really the foundational work of insider threat. I was very fortunate to have been part of a core group of three that really started on the early days of research into insider threat. That has evolved into a number of opportunities that provided ways to transition that knowledge out to different sectors within the Department of Defense, the federal government, law enforcement, industry, and academia. It provides a great opportunity to do foundational research into a specific area with a mandate to transition out to broad communities. I really appreciated that opportunity throughout the number of years I have been at the SEI. That has totaled now 21 years at the SEI.

Matthew: We have a wealth of experience here, and you have been working in this domain for quite some time. It sounds like, Randy, maybe from the inception of this work here at the SEI, essentially.

Randy: Absolutely. The early work in insider threat started back in 2001, looking at foundational incidents of insiders causing harm to organizations. That has evolved continuously into a corpus of insider data, of which we use as the foundation empirical data to build the models and patterns of insider

activity.

Matthew: Certainly identifying and mitigating insider risk is part of our mandate here as one of those organizations protecting the nation and one of the things that makes us unique as an FFRDC. So thank you for joining me to discuss this. Clearly, I have two prominent experts. My first question: you were doing insider risk before the pandemic. Then we suddenly decentralized the workforce in ways we would never anticipated with a technology stack that was not really intended to monitor folks in a way that may detect some of the things as I understand that you look for. A question rather about how the remote work, pandemic, the post-pandemic world is maybe changing expectations about employee behavior monitoring, insider risk detection. How has it changed the domain?

Randy: Let me start by what hasn't changed. Certainly, the need to address insider risk in an organization, whether that is in a physical building, or that is a remote workforce, working from home or some other location, there still is a need to address insider risk within organizations. That is the foundation that's continued over the course of the pandemic in the hybrid or completely remote workforce. Some of the things that we found is that from the very early and rapid transition from in a workplace to outside the workplace, the tools were typically configured to do basic anomaly detection. When I look for things like when do people traditionally work, where do they work from, the location, the tools were configured to have a pretty standard baseline of what was expected or normal or presumed good activity and looking for those specific anomalies within the context of the security operations. When the workforce goes completely remote from a Friday to a Monday, obviously some of the tools will automatically start generating alerts just from the proximity of where your workforce is connecting to and from. Tools, such as user activity monitoring or user behavioral analytics, those types of capacities would automatically start generating alerts just from the distance from which the employees are connecting into your systems. The way they are connecting and what they are doing has changed also as well.

Also, as we start looking at the research we have done—we focus primarily on malicious insider activity—but we also focused on the accidental, the non-malicious as well. So thinking about a remote workforce now with the pandemic, working from home. In the early days of the pandemic, everyone was working from home, all students were at home. So the employees with children and maybe elderly parents that care from home, the distractions in a work-from-home posture, at least, initially, were pretty significant. Some of those anomalies that were being detected could be because of the

distractions in the workplace that are now transitioned to the distractions at home. Looking for ways in which you can try to detect the malicious and non-malicious, the tools needed to have a longer time to start baselining what the normal or presumed good activity would be from a workforce which isn't physically in a facility owned and operated by the organization.

Matthew: Yes, that makes sense to me. What I heard is we need to re-establish what the profile of essentially standard behavior is. Things aren't anomalous, they are just different now. It is a really important insight for folks that are…Dan, do you have thoughts about this?

Dan: We went through really three stages of this. Everybody is normal, kind of, fundamentally shifted overnight. Then it took time to kind of re-baseline what normal looked like, and what we saw was that took organizations weeks and months. That was weeks and months where, kind of, the tools are spitting out alerts that says everyone within the workforce is potentially misusing their authorized access because they are behaving in a way that they weren't a week ago. By the time we got things caught up, it was just about time for everybody to start coming back to work, and normal changed again. Now, we are in this kind of hybrid environment where we have realized that, kind of, the nets we cast over developing these baselines, they have to be, kind of, really precise and specific almost to the individual level for us to have real trust in them. We also have to re-characterize what we do when we find a detection of something as being potentially anomalous. Pre-pandemic, we were able to put a lot more of our, kind of, analytic eggs in that basket of the Venn diagrams between anomalous and something of actual suspicious and malicious happening. We are seeing those kind of Venns separate a little bit. So it is kind of re-characterizing where anomaly detection fits in the analytical toolkit, and it is also put a lot of emphasis on solution providers to be more transparent about how they do that baselining and how much data we need before we can establish models that say, *OK, we're pretty confident that this is what normal looks like now*.

Matthew: Thanks. Maybe to put a capping sort of thought on this, it sounds like the fundamentals remain the same, and that makes sense to me. But the tools and the methods by which you detect malicious insiders need to evolve because the workplace has evolved. The nature of work is evolving, and we need to be vigilant in ensuring that our systems are tuned. I am sure there are folks in the audience that are operating insider risk, insider threat programs now. It sounds like this is one of the key considerations they should have.

Randy: They certainly would. Also, the tools have evolved over the past three years as well, where traditional tools would be an insider risk analyst telling a tool what to look for. Now, with the advent of machine learning, and artificial intelligence, the tools are getting better at looking at large data sets, looking for statistical anomalies in those data sets, and informing an analyst what potentially could be anomalous or malicious or suspicious activity as well. The evolution of technology is now being incorporated into these tools as well.

Dan: What constitutes a good control in insider risk management has also changed. This has been something that has been fun to watch, pre-pandemic, mid-pandemic, and post-pandemic as well. We published a study years ago about the critical role of positive incentives in managing insider risk, which was effectively groundbreaking in saying that better places to work, higher levels of engagement, connectedness to your job, and feeling supported by your organization were correlated with a decrease in the occurrence of insider incidents that organizations experience. Surprise, surprise, better places to work are less likely to have somebody want to maliciously cause harm to them.

Now, during the pandemic, what happened? Everyone started experiencing increased levels of personal and professional stressors, as Randy had mentioned earlier, financial stressors, just the change and what a normal day looked like. That exacerbated the importance of these administrative controls, better management practices in terms of better supporting the workforce and really driving home where those positive incentives fit when we are thinking about how to apply security dollars to help the organization drive down insider risk. Post-pandemic, now, we are back into that kind of hybrid thing. I would argue that we are still experiencing elevated levels of, kind of, personal and professional stressors relative to where we were pre-pandemic, but we're seeing kind of increased adoption within insider risk management programs for some of these more softer, kind of administrative, management, HR controls. I think moving forward, that it is always going to have to be a balance of the technical, detection, prevention, response, infrastructure, and increasing adoption of these management practices. Thinking of them as having causal security benefits.

Matthew: You said something really interesting, you said many things that are interesting, but one thing I would like to explore in a little more detail. So the World Health Organization says there is an epidemic of loneliness, and certainly during the pandemic, and now, after the pandemic, folks are feeling, some folks are feeling isolated, disenfranchised. You described engagement

as key to reducing the risk of inside, malicious threat by insiders. Thoughts about how this new work posture—maybe the research isn't there yet because it is mostly new, but how does this hybrid world where a portion of the workforce may be physically gathering and having interpersonal interactions, and then maybe you have folks that have never seen the team physically and know you only through the camera on their laptop and Zoom.

Dan: I think it changes the balance of kind of what we think of as insider risk management program (IRMP) operations, where, historically, we have been chasing these alerts from user activity monitoring and data loss prevention capabilities because some potential exfiltration event has occurred. Whereas as an organization shifts towards adopting these positive incentives, it is more kind of coaching with management and supervisors through the coordination of human resources to get really into, kind of, the needs of the individual employee, resolving colleague and coworker conflicts. And really accelerating and amplifying the capabilities of a really effective HR function enabled by the technology that we have available to us in 2023 within a modern or modernizing insider risk management program.

Matthew: Thanks, Dan. And it is certainly from the report that Dan mentioned as well, the positive incentive work, working with behavioral psychologists to identify when we can get people connected to their job, their job responsibilities, getting them connected to coworkers, as well as the organization support that's provided, that tends to reduce the risk of insider activity of someone maliciously causing harm against organizations as well. So we had to benefit over the course of the past couple of years of work of psychologists that really focus on the human element of insider risk and to try to combine that with the ability to use tools to detect both the stressors on the part of the individuals that are behaviorally based stressors, as well as the technical actions are taking that might be progressing down a path to cause harm to the organization.

Randy: One of the most interesting facets of what you do, in my opinion, is this joint between the technical and the physical in the sense that you are talking about the human element in a very overt way. So when folks think of CERT, they often think of things that are purely technical. I know in your bodies of work, you described the socio-technical, which is a really fascinating element of all this. You mentioned working with psychologists and people that maybe you'd think are far afield of cybersecurity, but for your work, they're front and center.

Matthew: Yes, they certainly are. When we talk about physical security

within the workplace, many organizations are far advanced in incident response when it comes to physical incidents within the workplace. But when an insider incident is a cyber component of that, how do organizations work together with physical security, physical response as well? Because really in the heart of cybersecurity is protecting the availability of a system, or service, or data, the confidentiality and integrity. The ways by which you can actually put those controls in place to prevent, detect, and respond does have a need of physical security protection, resiliency of the operation of the organization that includes things that are basic, such as power, and water, and AC that supports the data center. Because, certainly, if an insider has access to those facilities at a minimum, if they pulled the plug to a machine, that would disrupt the availability of that service as well. So the overlap between the physical and cyber is certainly a critical component of the work that we have done over the years.

Dan: I am going to brag on Randy and my predecessors for a while here becausethe first thing you saw from the technologists that were, kind of, analyzing this kind of cybercrime data in the early 2000s was a really prominent admission that this is a people problem. This is not a technology problem. This group has worked really closely with social and behavioral sciences for two decades to find ways to use technology to amplify the capabilities of those social and behavioral science practitioners. What you are seeing is that pairing really starting to catch on not only within the federal government here within the United States but even into the private sector as well. I am not saying that stumble on an insider risk management program, you are going to find a one-to-one between, kind of, psychologists and technologists. But you are seeing more government programs having folks within their programs that have that experience. And you are seeing industry follow suit and where the the happy middle ground is, is a much, more prominent inclusion of our HR partners in this type of decision making. Not only the reactive, *Something terrible has happened*, but in that proactive, *OK this person is kind of starting to display signs of needing some additional attention or support*. I think that is a really interesting progression. I am going to just say it. I think it is something that you all could have blazed the trail for with the way that you contextualize all this stuff early on.

Matthew: This progression from insider threat to insider risk is really interesting, which means being expansive in your approach and drawing on those stakeholders' best positions to create the conditions to prevent these things, and then not focus just on recovery but also on prevention. So, Dan, I have a question for you. Unfortunately, the world is full of examples of insider threat scenarios that have resulted in a realized risk, realized insider

risk. Now, I know we don't comment on kind of temporary cases, but all of this in some total led to a number of mandates within federal agencies and Department of Defense. I know Randy and Dan, the team has been sort of at the forefront of helping to articulate what organizations in government and the DoD should do and private industry. Dan, could you describe sort of what the key drivers are there and maybe specific instances where we have helped to determine what the appropriate outlines of a program look like for an agency or our stakeholders in the DoD?

Dan: Yes, absolutely. As Randy had mentioned, we collect and analyze all of this incident data. We develop these general models of incident progression for a bunch of different misuse cases, whether that's fraud or theft of intellectual property, IT system sabotage. Now that we have an understanding of kind of what these incidents look like and how they tend to evolve, we can then recommend some candidate controls for how organizations can identify that stuff happening within their organizations as well as some preventative measures that could have stopped some of those concerning behaviors and activity before the harmful act had actually occurred. We end up with these, the series of capabilities, both technical and administrative, that effectively forms the foundation for reference models that we use to help organizations lay out what they need to be doing at the enterprise level to manage insider risk effectively. We do that. Then, a couple of those very prominent incidents in the 2011-2012 timeframe, hit the papers. We see the federal government get serious about putting safeguards in place that address some critical capability gaps. The first thing that happens is Executive Order 13587, comes out that mandates the safeguards around protecting classified information on classified networks. It establishes a national insider threat policy and minimum standards, which we had the great pleasure of being able to kind of weigh in on and help shape. It also stood up the National Insider Threat Task Force, which was an organization that was effectively the action arm for ensuring that all federal departments and agencies implemented what was mandated, as well as was the clearinghouse in the concentration of subject matter experts to help organizations stand up these programs. The executive order, that pursuant national insider threat policy and minimum standards, [was] initially geared towards just the protection of classified networks and classified systems from US government employees. It is augmented in later years by a conforming change to NISPOM, the National Industrial Security Program Operators Manual. I nailed the acronym.

Matthew: Well done.

Dan: …which basically takes the national insider threat policy and minimum standards and brings it over into the defense industrial base [DIB] space and provides similar protections for cleared contractors that had the same level of access as our direct government employees. That regulatory landscape has really set the foundation for minimum standards for conducting user activity monitoring, for training and awareness, helping the general workforce understand what these misuse cases look like, and what their individual responsibility is in terms of protecting their authorized access to their organization's critical information, past even just classified. That is really what you see in terms of how we got to where we are from a governance perspective, particularly within the federal government and the DIB. Private industry has been following along, and some of the more heavily regulated sectors like the financial services and health care sectors. [They] have also kind of used those national insider threat policies and minimum standards as kind of proxies for their own set of, you know, requirements, guidelines, best practices. It has really followed that general pattern. Obviously, within the healthcare sector, we just replaced *classified* with *protected health care* information. Within the financial services sector, we replaced *classified* with *PII [personably identifiable information]*. But the frameworks are relatively similar.

Matthew: Thanks, Dan. I appreciate that. Let's do a thought experiment. I want you to forget that you are prominent experts in this field, and you are new to the subject completely. Let's say you have taken a new job somewhere, and one of the things in your job jar is create an insider threat detection mitigation program. Could you walk us through what the core components of that function are and maybe how folks should get started if they are given this responsibility?

Randy: So happy to start. And in this section, Dan, and I will kind of go back and forth what we describe as the critical elements of an effective insider risk or insider threat program. For more information, you can go to our Common Sense Guide to Mitigating Insider Threats, the 7th Edition, which was released in September of 2022. Best practice number two will walk you through how to build a formal insider threat program. I will start with critical component number one, which is developing a formality around the insider threat program or developing a formal insider threat program. Really at this stage, we want to get executive support, have senior leaders own the program, the ones that have the authority to mandate the participation enterprise-wide across the program, and provide information that will provide truly an ability to measure the insider risk within the context of an organization. Now, that comes with more formality, which we will outline in

these number of critical components, but it should start thinking about things such as, *Well, what is the scope of the program? Is it cyber or physical or overlap between the two? What is the authority of the program? What is in scope and what is out of scope of the program as well? When we operate the program, what are we looking for? How far do we go in terms of things like user activity monitoring?* All that should be clearly codified in the procedures, the policies, and the practices, which we will talk about the critical elements coming up here as well.

Matthew: Wait a minute, may I ask a question? I am sorry. I am thinking about, the first challenge there. Getting executive buy-in or senior leadership buy-in makes sense. There is a governance layer to this. You are going to be monitoring user behavior in a way that may be unusual based on past practices. Thoughts about how you strike that balance between privacy or expected privacy and operating an effective set of insider threat controls?

Randy: One of the things we will talk about is an enterprise-wide approach to building the formality around the insider threat program. We will talk about the critical elements, but one of the foremost subject matter experts we want in the discussion is your general counsel or the legal aspect of your organization, the ability to deploy tools and to monitor with the approval of the general counsel in your organization. Human resources and physical security is part of that as well. So striking a balance truly of what we are doing to protect the critical assets of the organization. Those critical assets do include your people in the organizations and the facilities in which they work, but also the information and technology. At the heart of it, we are trying to assure the availability, the integrity, and confidentiality of those assets to protect the assets. Now, that can include, with the approval of general counsel, an ability to monitor the employee activity within the context of your organization.

Dan: Yes, so part of formalizing the program is getting an understanding of what already exists within organizations and really unpacking why we need this standalone carved out thing. One of the reasons why we end up needing that is because, historically, traditionally stovepiped parts of organizations like IT, HR, and physical security don't share information freely or proactively without somebody pretty high up the org chart, not only requiring them to do so, but actively facilitating that flow of information. Part of the formalization is, *OK, this is what we are after in terms of these types of threats to these critical assets. We understand all of these different components… Information security has this part to play and HR has this part to play, but to coordinate the proactive kind of response to the, just the bad stuff, but the harm*

or the concerning behaviors and activity that precede the bad stuff. *We need you all kind of talking and doing these things in a way that isn't necessarily happening organically.*

Matthew: A question then occurs to me. Imagine, again, we are new to this job, and we are presented with the incident response plan for the organization, and they call it *all hazard*. But truth be told, it focuses on physical disruptions and then purely technical challenges like malware. Do we need a dedicated insider threat response planned? Should we just simply augment the overall plan with these elements? Thoughts about how best to position insider threat in the realm of incident response?

Dan: Sure. Yes. I mean, let's contextualize it, right? An external threat actor doing bad thing. We need to make sure everybody knows exactly what is going on at all times. It is a colleague, a coworker, someone whose office is two doors down, or someone who you are on Zoom with a couple of days a week. You have to handle the communication of the investigation of that potential incident a little bit differently. It calls for additional confidentiality. It calls for a little bit more precision in terms of how you progress from step-to-step within an investigation and analysis process. Whether you are doing that within its own standalone policy or playbook from an incident response perspective, or you have got specific and special call-outs within some generalized plan, you have got to consider what has to change when the threat actor is one of your own. Because the response options change. It is one of the good things about insider risk management. Our response options aren't just disable ports or protocols or throw people in jail. It is get them training, get them help, get them better connected. It is make the organization a better place to work. Because the range of response options varies so significantly, potentially with insider risks as opposed to what you worry about from an external threat actor perspective, that really drives what has to change in your incident response.

Matthew: What you are describing is sounding familiar. Prior to the SEI, I worked in private industry, working in IT security and audit. We had very specific procedures around any financial crime that we encountered. If we thought there was fraud or some sort of other nefarious activities, there was there was a policy and procedures related to the preservation of evidence, and privileges communications, and basically ensuring that counsel was involved early on. It sounds like, not to sort of be reductive, but it sounds like the steps you take for an insider risk program are very similar.

Randy: Yes, very similar. Again, it is the protection of the organization's

critical assets. As Dan mentioned, we have models of insiders who intentionally sabotage an IT system. The way you would prevent, detect, and to recover and respond to an IT sabotage incident will be different from, as Dan described, someone stealing intellectual property. The goal of this is the remediation, the recovery, and then changing the control stack that you have in place to hopefully reduce the likelihood a similar incident will happen in the future. The same would be true with fraud and recognizing that an insider who commits the activity with malicious intent may be part of the incident response for that type of incident as well. So building that defense in depth in your incident response is critical when it comes to insider incidents as well.

Dan: That is a really interesting one to riff on. There was an incident a couple of years back where an organization that was experiencing a ransomware incident. One of the responders to that ransomware incident had changed the Bitcoin address that was associated with the ransom payment to a wallet that that individual had control over. I challenge you to tell me that your incident response plans have that particular vector well covered, unless it is something that you have either explicitly had happen to you, or you have really taken the time to cover your bases from an insider risk management perspective.

Matthew: That is really interesting. That is another lesson. Let's build on the past mistakes of others. This is the way it works in aviation, for instance, where every accident is seen as a way then to prevent future accidents. I would say then ensure that those responsible for the crypto wallet don't have some other role where there is a temptation to do this. But I understand there is a limit to checking the checkers. It is really interesting. If you are hearing this for the first time, if you are given the job of standing up an insider risk program, it may seem really daunting and a bit confusing at first. A question, Dan, what barriers do organizations typically face initially when they are trying to stand up this capability?

Dan: Sure. Well, this is a good one to start with, right? Within the federal government space…

Matthew: Meaning money.

Dan: Within the federal government space, this was an unfunded mandate. I walked you through kind of the executive order and the National Insider Threat Policy and minimum standards. They said, *Do all this stuff*, and no checks got cut. It was finding a way to amplify the resources that you already

had and be that grease, that coordinating component, to help those things share information more effectively and efficiently. Let's get past just, new programs are expensive and can take time. We also hit on one in the beginning of this, too, which is in the formalization of the program, insider risk is such a broad, a set of challenges that we house within this one specific area within organizations or this one specific domain. We see a lot of organizations just struggle with wrapping their head around what that means to them or what specific part of the insider risk problem they are carving off to try to apply some resources to buy down risk in a particular area. We see organizations struggle with just articulating the scope of the program. A lot of that also has to do with to clearly articulate the scope of the program, you have got to have a relatively in-depth inventory of what is already in place. This is meant to be additive, not necessarily repetitive. Now, you mentioned a good example in your time in the financial services. If an insider threat program got stood up within that organization you were working at, and they didn't do an amazing job checking with the folks that were handling the fraud part, duplicative capability we bought the same tool twice. And now bosses are being called into boardrooms and having to fight about territory, who is responsible for what.

Matthew: Sorry, is that how insider threat programs wither, is that they have that duplication of function? I am curious, not only is it the barriers you have to constructing it, but barriers to the operation and enduring way of insider threat programs.

Dan: Oh, yes. Sure. Yes. I mean, we have written a whole paper about this. *Effective Insider Threat Programs: Understanding and Avoiding Potential Pitfalls*.

Matthew: That is available on our website.

Dan: It is, yes. A couple of themes you will see in there. You hit the nail on the head, which is you have got to have the right metrics and measurements of effectiveness in place. Because it is not just, *We stopped these many pieces of company confidential information from flying out the door*. Eventually, as programs mature, that decreases. You can't rely on that metric without being able to explain that to a board that says, *Well, it is down, way down, but we are applying the same dollar amount to it as we were*. So finding ways to incorporate all of those different response options that we talked about as success metrics for the program ends up being really important. Another way to really easily erode trust within the program is to do something outside of what is policy, procedure, legal.

It is easy with the tooling and telemetry that is available to kind of get outside of your swim lane. You are a couple of knobs or switches away from doing something that is going to put you in hot water from a compliance perspective. We have seen organizations struggle with that, not even necessarily on purpose. They have got a stakeholder within the organization that says, *Hey, real quick, I know you can see this. So run this for me, and let me know if anything comes back*. It can get easy to erode trust in what ends up becoming a critical asset in and of itself that the organization must protect, which is this aggregated collection of information that insider risk management programs build.

Matthew: Let's talk about that temptation to do more with the tools you have and do more in a way that maybe is not appropriate to the mission that the tools were purchased for. We talked about the decentralized workforce, folks working from home. There is a temptation, and organizations often now are contemplating monitoring their employees for productivity. Thoughts about the division you should have between monitoring for productivity and monitoring for insider risk.

Dan: It is a matter of what you do with that analytic and whose hands you put that into. Has the right set of conversations happened within the organization that says, *Yes, we are okay with this being kind of an outcome of this type of analysis*? Whether that is measuring for productivity, or connectedness, or engagement, same question there, too, right? Putting that into the hands of an HR practitioner, they say, *Well, how do we support this person?* A different stakeholder in the organization might say, *Well, they are disengaged*. *They are not next up, and we can start to build the dream team for this particular project*. When not very clearly articulated what we use these data sources and analytic capabilities for, it is really easy for those things to kind of go sideways, which is why I could have governance of these programs, and having it at the right levels of the organization is so important.

Matthew: Thanks. Randy, shifting gears slightly. As a federally funded research development center, our focus is on national defense and ensuring national prosperity. One of the key goals of an organization like this is transition. Would you describe to me sort of how we have transitioned the things we have learned about insider threat to our partners in the DoD and elsewhere? Are there specific resources you would point our audience to if they are contemplating these things for the first time? Or, if they are long-term and frequent visitors to the website, are there things that are new that you would highlight as important?

Randy: That's a great question. You mentioned earlier the resources that are available. There's certainly resources that are here within the government, within the work we have done here. *The Common Sense Guide to Mitigating Insider Risk, Version 7* is a great starting point if you are looking to build formality around the program. It identifies 21 best practices organizations can implement as we look to build the controls that can prevent, detect, and respond insider activity. But there is also a lot of great work that's being done across the federal government. Dan mentioned the National Insider Threat Task Force. They have a lot of great resources available on their website. The DoD Management and Analysis Center, the DITMAC, has done a lot of great work. So a lot of publicly available information is a great starting point. As we look specifically for other resources that we have done, we have published over the years over 125 reports specifically on the topic of insider threat and insider risk. If you are looking to build a program that is looking to reduce the risk of someone sabotaging your system, an insider doing that, we have reports on IT systems sabotaged. [If you are] looking specifically about insider fraud or theft of intellectual property or accidental disclosure of information, a lot of information on our website, sei.cmu.edu, a lot of great resources that are available.

That is a great place to start the program. But what if you have a program that's up and running? How do you start measuring the effectiveness of the program? There are a lot of controls that can be put in place that can be measured. We try to build that into assessment methodologies that we have built that we can train you on. We have insider threat vulnerability assessments that we can perform against organizations to see how well they would prevent, detect, and respond. We have insider threat program evaluations that we can do on your organizations. Those same technologies that we can do as an organization against your programs, we can train you how to do that as well. We have a wealth of information on training that is available as well, such as insider threat program manager training. If you are the program manager, the one that is responsible for the day-to-day operation, we can train you how to implement and operate an insider threat program. Or if you are an analyst in the insider threat program team, we have insider threat analyst training that can be very, very helpful as well. Again, there are a lot of great starting points.

I would also like to mention that we here at CERT and the SEI have a group which is called OSIT, O-S-I-T, the Open Source Insider Threat Information Sharing Working Group. What that is, over the years, over 300 organizations have got together to share information around building formality around insider threat programs. We get together on monthly phone calls where we

just share challenges, successes of building insider threat programs, and really allow the practitioners to talk in a trusted environment around ways by which they have learned from, they have succeeded, they are failed. It is just a great way to connect with other organizations that are challenged with building insider risk programs in their organization as well. For more information on the OSIT group or the public cases we have talked about, feel free to go to our website. A lot of more information is available on our website that can walk you through how to build, implement, operate, and measure the effectiveness of an insider threat program.

Matthew: There is a catalog of things that we offer here at CERT. And also, Randy, you piqued my interest with OSIT. How does one become a member of OSIT?

Randy: If you want to send an email to either Dan or I, we can get you connected as well. Or, we have an insider threat feedback alias that is available on our website. Send us a quick, brief email, that you would like to get connected, and we will get you an invitation out to join the group.

Matthew: Great. Excellent. Thank you. So Dan, a question about what comes next. So AI is the free space in bingo. I know your answer will include AI in some part and large language models. Tell us what is coming next? What is on the horizon for us in insider risk?

Dan: A couple of things. I will broaden the aperture, and then we will narrow it down to AI. We are continuing to assist organizations in the development evidence-based recommendations for the most effective combinations of data sources, analysis techniques, and response options for the various incident progression components for the incident types that we study. There are a million different ways to make some detections of a particular concerning behavior and activity. Our real deep engineering and research in this space is done to identify the best ways to find disgruntlement or data exfiltration, or all of the different components of these models that we study. As you mentioned, artificial intelligence is now providing real promises in this area to help reduce analyst burden, find patterns that would be really hard for humans to be able to identify, standardize and streamline portions of an analysis process, or even remove analysts from the loop entirely. Our research within this space is designed to help illuminate where the boundaries are for human-machine teams to include those that use artificial intelligence to wade through large collections of data, whether that be network sensing data all the way through employee performance review data to identify the insider risks that are present within specific

organizations. Taking the human out of the loop when deciding on who needs training, more coaching, support, or having their clearance revoked is probably something that is a bridge too far for the risk appetite of most organizations. Finding where that line is in terms of what we can and can't let computers for us as well as articulating and identifying the biases that are not only present in human analysis within this space, but also when we bring AI to the party, we are just switching out the different types of biases that we have to worry about. We are doing a lot of work right now, partnering with organizations that are trying to lean in to find the most effective combinations of AI algorithms and data sources that automate bits and pieces of the analysis and response process. Through experimental design and executing these experiments, whether with synthetic data or giving them the test plans and letting them do this in their own environments, giving them empirical bodies of knowledge they can build on to say, *I want to look for this. These are the most effective and efficient ways to do it*.

Matthew: Dan, thank you. There is something really interesting, several interesting things. But when I think about as we give machines, essentially, *autonomy*, as the name would imply, or agency, you can imagine us evolving to a point where a source of insider risk that is a trusted entity. In this case, a synthetic entity, could do things that are malicious or unintendedly consequential. Do you think in the future, we'll be talking about…You talked about human-machine teaming, we'll be talking about the insider risk posed by AI itself?

Dan: We are talking about that internally. What we hope to do is follow a very similar pattern to what we use within this organization for getting started with the traditional research within this space, which is a massive body of misuse cases and then use that to build these generalized or generalizable models of incident progression. Now, we are going to be talking about applying some very different controls on the other side, but the basic methodology and the process we think is going to be the same. We will end up with these general models of incident progression where the threat actor isn't necessarily a person with authorized access to an organization's critical assets, but an artificial intelligence system.

Randy: If I could, Matt, if we could describe this as we apply the AI into the tools and technology, the ethical application of AI, the explainability of AI as well, because if you think about this being the detection of an insider threat, that is potentially identifying an individual. How can we explain why that trigger generated alert on this individual? Is that something we can justify to our human resources or to our general counsel as well?

Dan: Or to a court?

Randy: Or to a court?

Matthew: Which is an important test in all this, absolutely.

Randy: Very, very important. I think about insiders that potentially could influence the data that's being provided to the AI engines. Could that data be poisoned from an insider to then avoid detection of future insider incidents as well? As Dan mentioned, these are things we are thinking about. We would love to partner with organizations to continue that evolution of our research and love to work with organizations on the application and transition of that thought and that thought process as well.

Dan: There is another really key component to where I think we are going next and where I think this domain is going next. I mentioned the critical role of positive incentives here. We have got early research out there that establishes this kind of correlative relationships. There is more work to do to drive forward this causal understanding of reframing these management practices as security controls. I think another really key component to that is us doing a much better job a domain at being able to quantify these risks and being able to drive towards return on risk investment calculations that allow us to measure re-upping the license for the DLP tool versus bringing in some executive coach that is going to also slip a couple of security sweet nothings into the ears of our folks with authorized access. We are breaking ground in other parts of CERT in terms of how we are, tackling this problem of cyber risk quantification. We are partnering very closely with those folks as they crack some new code, find ways to adapt that to this. The positive return stuff, because it is management-practice-focused, that is going to require its own set of, kind of, dimensions that we need to factor in. So it is another real key area of, kind of, growth for us in the future.

Matthew: Thanks, Dan. So if you want to connect with the team, if our audience wants to connect with the team, and I don't mean Digital Dan or Robo Randy, we have an event coming up in September, do we not?

Dan: We do, yes.

Matthew: Could you speak to that, please?

Dan: Yes, September, National Insider Threat Awareness Month. This will be

year four of September being National Insider Threat Awareness Month. It will be year nine for our Insider Risk Management Symposium. This year we are excited to get that back in person. We will be in our Arlington offices, really justbringing practitioners, researchers across government, industry, and academia together to talk shop, to get a sense of kind of what is new, what organizations are struggling with, and what organizations are being successful with as the threat landscape changes and as these new challenges surface and emerge. The last two of these have been YouTube events, and we are excited to get back in person.

Matthew: Excellent. Well, Dan, Randy, thank you so much for the discussion today. I really enjoyed the conversation. I am sure our audience has found many things that can help them influence and inform the way they construct insider threat programs. There will be links in the transcript that will take you to the artifacts that Dan and Randy have named. A final reminder to our audience, these podcasts are available on SoundCloud, Stitcher, Apple Podcasts, Google Podcasts, as well as the SEI's YouTube channel. If you have liked what you saw today, please give us a thumbs up. Randy and Dan, thanks again for this conversation.

Dan and Randy: Thanks for having us.

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 do not hesitate to email us at info@sei.cmu.edu. Thank you.