Getting the Most Out of Your Insider Risk Data with IIDES

featuring Austin Whisnant as Interviewed by Dan Costa

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.

Dan Costa: Hello and welcome to the SEI Podcast Series. My name is Dan Costa. I'm a technical manager for the Enterprise Threat Vulnerability Management team here in the CERT Division of the Carnegie Mellon Software Engineering Institute, and yes, I have a very long business card. Today, I am joined by my colleague, Austin Whisnant, one of our senior researchers in the Enterprise Threat and Vulnerability Management team. Austin and I today are going to discuss a new insider incident data expression standard that she has been leading in support for our team. Austin, welcome aboard. Good to have you.

Austin Whisnant: Hey, Dan. Thanks for having me.

Dan: Let's get things kicked off here. If you could, just give us a little bit of background about yourself, the work that you do here at the SEI, and maybe a little bit about one of the favorite things you do here at the SEI.

Austin: Yes. So I have been at the SEI for about 14 years at this point. I have

been on lots of different projects, which is something I really like about working here. I have done everything from sort of network analysis, cybersecurity operations kinds-of analysis and research, to workforce development, making training exercises, doing simulations to help cyber warriors train. I have been with the insider threat team for coming up on five years, I believe, now. So doing lots of different projects there, working with our big database that we use called MERIT, doing some cool AI stuff, looking at how AI can fit in and help with some of the issues and challenges that we have with insider threats. I don't know if I can pick a particular favorite, but I mean, one of the things that I really like about working here is that I get to do different problems, work on different things, and take on different challenges. That is probably my favorite part.

Dan: Fantastic. It's a refrain you hear a lot from folks that have been at the SEI for as long as you and I have, which is one of the reasons you stay at a place like this is because without having to move around companies or parts of an org. chart even, you can get a lot of different kind of experiences working a bunch of different problems in a variety of different organizational contexts. Certainly, I think you and I share that in terms of one of the reasons why we have stayed here and gotten a chance to do so many neat and different things during our time.

Austin: For sure.

Dan: Enough preamble. Let's get into the fun stuff, Austin. We have been working really hard in the past year plus on the development of IIDES, the Insider Incident Data Expression Standard. Recently, you released a blog post helping to introduce this data expression standard to the research and practitioner community. We are going to link the blog post, which also has links to how you actually get access to the standard, its documentation, all the supporting things that we put out there around it which we'll talk about. We'll link that in the details of this podcast recording. But, Austin, could you maybe just get us started with an understanding of what IIDES is, why we thought we needed to make it, and some of the problems that we are trying to address with it?

Austin: Sure. It is sort of a long story, and I think you probably understand some of the longer-term back-end kind of stuff better than I do, but I will do my best, and if there is anything that you want to add as well, of course, please. We have been building our own repository of insider threat information for years, more than a decade. We used that to do various research analytics, different kinds of things to help the community and our

mission partners deal with insider threats. We have had that database. We have been collecting information. We have learned the pros and cons of doing it different ways. We have sort of built out this schema over time for ourselves internally to use. It has gotten to a point where we think it is really helpful. It has sort of solidified. We wanted to transition that out to other people in the insider risk space who might be interested, who might find it useful. We have essentially codified our own schema into this standard called IIDES. We have produced the standard itself, information about how to use it, and then a JSON schema to go along with that for folks who might want to implement that in their own way in their own databases to use. That is the reasoning behind it, and it is now available on GitHub. People have asked us here and there over the years, *How do you collect your data? Where do you put it? What fields are in your database?* We are trying to answer those questions for people out publicly on the internet for folks to use. Is there anything that you want to add to the backstory of that?

Dan: Well, a lot. We have learned over 20 years of collecting and analyzing incident data that there are a bunch of different discrete use cases for what you can and can't do with learning about a prior incident to help better support what it is operationally we are doing to manage the next threats or the next risks that those with authorized access or critical assets pose. First, it is how we refine and deploy the analytics that we use as the mechanisms for our technical detection capabilities. That is a really important part. There are value propositions to be had about the efficacy of our technical and administrative controls more broadly: *Hey, we spent $10 million on this widget that is supposed to stop all this bad stuff, but over the past year we have seen 25 instances of this bad stuff.* But that is a really powerful way to help folks understand what is working and what is not working, right? Shy of the bad thing happening, there is all the work that an analyst did to get to the decision of kind of what to do in the face of this threat-conducive behavior or activity. As our insider threat and insider risk management programs get more proactive at trying to address the root causes that contribute to somebody exhibiting concerning behaviors prior to actually sabotaging a system or stealing intellectual property. When we want to look at those proactive opportunities to manage those stressors that lend to the realization of concerning behaviors, we need to be able to have an understanding of all the precursor activities that led up to that and what we decided from an analytic perspective. Austin, maybe talk a little bit about how we had to factor that in and what we ended up developing with IIDE.

Austin: For sure. As far as the analytic portion and building that out and looking at trends across your organization, you are spot on with that. I would

say that having a standard, having something to reference, helps with figuring out which fields am I going to make that timeline out of, or what stressors am I going to track. One analyst might call them one thing; another analyst might call them another thing. So having it set in a standard fashion helps with that trending, helps with building those analytics.

Thinking through those use cases that different organizations might have, people call things differently. Maybe different sectors might have a different term for certain technologies or for certain stressors or things like that. Law enforcement might care very specifically about a certain set of things while, the financial sector, for example, cares about a different set of things. So that is something we had to incorporate as we were thinking through this. There were a couple of ways that we tackled that. We set ourselves some guidelines, like some guiding principles. Those principles were *simplicity*, *expertise*, *flexibility*, and *interoperability*. We didn't want to reinvent the wheel. There are some standards out there that are sort of related. Some of the cyber standards like MITRE ATT&CK or some of the other simulation standards from SISO, for example, that have the sort of technical components, the cyber components. We didn't want to reinvent those. So we kept an eye toward allowing those sort of standards in as far as people trying to pull from those. Flexibility. We didn't try to over-specify*, You should use this vocabulary. You have to use these particular components*, or things like that. So letting people kind of pick which pieces of IIDES they want to use that makes sense for their organization. And then expertise from us, like we just said, we have been doing this for quite a long time, and we have a lot of lessons learned from it. We tried to pick out what was most important in our own opinion and also from what we have seen operationally from different mission partners as well. That is the way we tackled it, just keeping in mind those guiding principles and allowing some flexibility in the system for people to use how they want it.

Dan: Fantastic stuff. Love the guiding principles. I want to poke a little bit at a couple of those. Interoperability is obviously a great thing for any kind of standard like this that we are using to kind of control vocabulary. For folks that may be not as steeped in insider threat and insider risk kind of program management/program building like we are, will you talk a little bit about why we felt like we needed a standalone controlled vocabulary for managing these types of incidents? What is different about an insider incident as opposed to what you might see in some existing cybersecurity standards that require us to fill that gap?

Austin: Sure. It might help to take a step back and describe a little bit what

the vocabularies are in IIDES versus some of the other components. To take that step back real quick, in IIDES we specify seven core components. Think about them as classes or major building blocks of the schema. That is things like *incidents information*, *insider information*, *organization information*, and *TTPs, and detection response*. Those are the core ones. We have others like the *targets*, *impact*, *accomplice*, stuff like that, that kind of fit in there as subcomponents if they are needed for the particular incident. We have these core kind of classes, if you will. Within those, we have the fields or the attributes that match those. For the insider component, it is going to be things like the insider's name, right, the job role, stuff like that. Then we also have vocabularies that we specify. For certain fields, for example, like organization sector, we have a specific vocabulary. In this case, we have pulled it from the North American Industry Classification System because why reinvent that wheel. They have already specified the different sectors that businesses can be in. We just repurposed that for that specific vocabulary for that field. That is what we are talking about when we have vocabularies. To your question about, why have insider threat-specific vocabularies, I think the best example of that is when we are talking about TTPs, the methods that insiders are using to conduct their actions. There are a certain subset or a certain set of vocabularies related to cyber TTPs specifically. They are typically more related to outsiders trying to attack. They will have a whole bunch of very specific technical methods for outsiders to attack. Sometimes, those might apply to insiders. Every now and then you get an insider using a logic bomb, for example, or something like that. There is also a lot of stuff that insiders do that isn't listed in those vocabularies. For example, using a colleague's account to do something, something like that. We wanted to specify that. When we did that specification, we looked at some other vocabularies. We used our own ontology from several years ago as well to try to bring in not just our expertise and experience but expertise and experience from other areas and other researchers that have done work in this area. Does that answer the question?

Dan: Absolutely. I think you hit the nail on the head. I think some of the differentiators, just for how we handle insider incidents as opposed to malicious cyber actors external to an organization, drives why we needed this in the first place. Detection and response happen very differently for insider threats and insider risk, particularly on the response side of the house. We have lots of opportunities to respond to the signs that somebody may be headed down a path to cause harm for an organization by retraining them, by getting them better recognition for the work that they are doing, by relieving whatever stressor personally, professionally, financially, is putting them in a position of exhibiting these concerning behaviors. You don't have

that chance with external attackers. You can't give an external malicious cyber threat actor a raise and make that problem go away. Maybe that is what ransomware is. That is a conversation for a different day. The point is these response options are different, and the work that we do leading up because these are people that we know a lot about, right? These are not unsubstantiated in many instances. We understand these are our people, and we are trying to help them protect their authorized access to our critical assets from being targeted by somebody external to the organization, right? We are here to help, and part of that is leveraging the data that we have about the individual in a way that we just don't have the chance to do with when we are responding to malicious, cyber threat actors. When you look from straight cyber, I think those are the pieces that we typically see missing. When we are developing things like controlled vocabularies, ontologies, taxonomies in the insider threat/insider risk management space, the information about the people, like you mentioned very specific TTPs, and then detection of response, and then what we do next in terms of the analysis and response processes that we almost always see kind of gaps that we are having to fill.

Austin: For sure. There are some other things in there, like, for example, job, right? There are lot of specific fields related to the person's job and the amount of access that they had and what their title was and that sort of thing. We also collect a lot of legal response court case-related information that we might not typically collect for external actors. That helps again, going back to the trends and the analytics. Is there a common job title, for example, or something like that that we are seeing in the data that we can look at? So, yes. Absolutely.

Dan: Fantastic. Okay. Put your practitioner hat on. Walk us through some examples of how and who might benefit from using kind of the work that we have done here with the development and release of the insider data expression standard. Who is our target audience here?

Austin: I think our target audience is a lot of different people. I mentioned flexibility as one of the guiding principles for the schema. Depending on who it is and how you want to use it, you might use it differently. That is okay. We have tried to account for that. Besides us, we have talked about the way we use it for trending and building analytics. We use it across a large amount of data related to mainly U.S. federal criminal court cases that have been prosecuted for insider threat. For someone who is using it more from like a operational or practitioner kind of standpoint within their own organization, I mentioned earlier, we have had people come to us and ask, *I've heard you*

guys have this database. Can you tell us what fields you are collecting, or can you essentially help us get started? What should we collect? What should we track? At a minimum this provides the fields that you might want to start with or even a subset of fields that you might want to start with. You can look through our documentation and pick out what makes sense to you. If you are doing this for your own organization, you might not include the organization component, for example, because it is always going to be the same for one organization. That is a simple example of someone using it. We also include a JSON version of the schema as part of this. If your team has developers or database administrators that want to turn IIDES into something very specific for your organization, a web app that you can log in and use and put in your information and do trending, that is available as well. There are some other use cases that we thought of. For example, there might be organizations out there that have to do collection across a number of other sub-organizations or maybe some sort of other information-sharing type of goal. In order to do that kind of information sharing, you have to have the same vocabulary, the same fields, that sort of thing, hopefully, the same format to make it easier on yourselves. That is another, I think, major use case for that. I will mention as well here, we just released PyIIDES, which is the Python implementation of this. It is linked through all of the stuff that will be linked in the show notes. That is our implementation that could be used as is or copied. But that has an anonymization function. If that is a particular concern during information sharing, you could anonymize things like the insider's name, for example, or the organization name, stuff like that, that might be a little too specific to share with other organizations. That is another major one is that cross-organization collection and sharing. Then on the research side I think is another major use case. For us, we are doing our own research in analytics, but also as a community, if we can use this to have some repeatability in the research to work off of each other as researchers across different organizations speaking the same language so that we can have repeatability in that research. Are there any others that you have thought of, Dan, for this?

Dan: Before we open up the part of my brain that is going to talk about the art of the possible, I wanted to go back and emphasize a few of the ones that you mentioned. I think that that hub-and-spoke model of decentralized insider threat/ insider risk management programs that we see for really large organizations that are just too big to centralize all of this data collection and analysis. That is a really, really important piece of this. When information has to either flow up by way of, *Here is what we are required to report to you at headquarters with regards to kind of what we are seeing, what is happening, how things are going*. Or, how we want to kind of promulgate things down to those decentralized elements in terms of, *Are you seeing things that look like this? Get*

your stuff in this format so that we can do those training and analytics for you. I think that is going to be a massive, massive value add. We have worked with so many organizations, both public and private sector, that struggle with just having a mechanism they can use to facilitate that information sharing. I think we, you and the team, Austin, have done an amazing job in giving folks a solution to a problem that we see a lot. So I wanted to hit that one.

In terms of other places, I am the big thinker on this stuff, right? I like to shoot for the moon and land amongst the stars and have Austin pull me back to reality. When we look at case management systems that are being used by insider threat/insider risk manager programs, and we talk about interoperability, we are hoping that the folks that already have existing solutions in place can find mechanisms through what is commercially available to start adopting standards like these so that we can avoid vendor lock-in, so that we can reason about our potential risk indicators and the incidents that are associated with them in some consistent form or fashion, hop from tool to tool. Operating under the understanding that, for a problem space this diverse, we are really never going to see a single pane of glass. That is the only thing we stick in front of our analysts. From an analytics perspective to a case management perspective to a reporting perspective, we are going to end up with these different, views or application stacks or suites that we are putting in front of our analysts to make sure that we are happy with the consistency across those different parts of our jobs. Using an information expression standard like this could help us stay honest and not lose things in translation when we go from our analytic environment to our reporting environment to the work that we are doing in support of some data call or operation or prototyping new detection capabilities, trying to get some control efficacy measures or just understanding what happened in an actual incident. So operating with an understanding that these are very discrete functions that happen in the overarching business of an insider threat/insider risk management program. Then you have this one thing that works as the controlled vocabulary across those three contexts. I think that is going to be a really important gap for folks as well.

Austin: For sure. I agree with that. I think I would also add on there that part of the hope with this is that some of the vendors who play in the space or play in related spaces will pick up on at least some of it, so we can make this build a little more consistent in those tools as well that organizations use.

Dan: We have seen that happen with things like our insider threat indicator ontology in the past, too. We are hopeful that, if we can demonstrate a robust user base around this or you see market forces looking at folks to say,

Hey, this is the way I want to be able to at least get my stuff out of your case management system. Again, not because we are going to go use a different tool because we want to use the same data expressed in a consistent way in another kind of distinct, portion of our kind of insider threat/insider risk management work. I think we are going to see the demand signal for that in the not-so-distant future for sure. Okay. Talk to us, Austin, a little bit about the process of developing this thing. What was easy about it? What was hard about it? What challenges came up? What is it like just working on something like this? For challenges that arose, what did we learn about how to get around those?

Austin: Sure. I mentioned at the very beginning we have our own database, that has had several evolutions. I maybe can't speak as well as you can to the original evolutions of the database, but I know at some point there was an access version. Now it is up on 2.0, maybe 3.0, depending on how you want to define it. We had this database, which has its own associated schema, and we were using our own set of fields already. Part of that was the insider threat ontology. We developed this over time. Which fields make sense? Which fields can we collect? What do we want to collect? What can we collect, which are two different things. So that is one of the challenges there. Part of what we targeted with the schema is what would we want to collect. What would be beneficial? Even if we don't have access to it, maybe someone else does, or maybe that is something we can think about as a community that we want to be better about collecting certain information. We started with that as our schema. We also have a couple of decades' worth of research about what is useful to collect and that sort of thing. Part of the challenge was turning that into something more standardized that is not so specific to our organization. We have talked already about different use cases, different organizations. Some organizations might need, for example, the legal response data, the court case data. Some organizations might not care, right? That sort of thing. Different use cases were another challenge that we had. Different terminology, we have talked about that as well. So financial sector might have a certain set of terminology. Law enforcement might have a different set of terminology. We as a research and development institution might have our own. That was another challenge. And then just different goals of using it. So, you know, are you using this to collect information to do a, you know, legal prosecution, right? That's a certain standard of things that you want to get to. Are you using it just to kind of track what's going on in your organization and maybe do some trending, maybe, you know, presents to the C-suite, or something like that, what's going on? So different set of goals. And then we've also mentioned the related standards as well. So how do we incorporate those without reinventing the wheel, without stepping on

anyone's toes, trying to make it easy for interoperability and those sorts of things? How we address those, it really comes back to those guiding principles that I mentioned, it really helped us stay focused, especially the idea of simplicity and flexibility. Don't over-constrain the schema. Provide more than what people might need because we don't necessarily know every use case, so let's specify what we can so folks can use it. Then the other thing I would say is we are hoping to get feedback about this. It is supposed to be sort of a living standard that we can update as a community. We can get feedback and sort of iterate and make it more helpful for folks.

Dan: Let's talk a little bit more about that part of this because we are still kind of in the initial release. We are just super excited that we have gotten it out the door phase. We certainly encourage everybody who has taken the time to stay with us this far in our conversation, to check it out and to tell us what you think about it. Austin, can you help our listeners, you know, better understand kind of what specifically we're interested in hearing about with regards to, you know, feedback we're looking for for kind of subsequent iterations of the data expression standard?

Ausin: Sure. So we are always open to feedback about just how awesome and useful it is. That would be much appreciated. But besides that, you know, we're looking for really any kind of feedback. So anything from really specific corrections like, *Hey, I found this typo in the field description*," or something like that. Vocabulary additions, *You guys have specified 15 different vocab words, but we use this one. Can this be added to that particular vocabulary?* New use cases, just sort of general feedback is really helpful as well. Like, *Hey, you know, the schema kind of, sort of covers our use case. But, you know, if you made this particular change, added this class, changed this relationship, whatever it is, you know, this would help our use case a lot more.* That would be super helpful as well. Just any kind of requests for clarification would be good. You know, if there's something we can change in the documentation, for example, we are absolutely happy to do that. And as far as, you know, how you get us that feedback, this is all up on GitHub. There is a repository there. That is open for pull requests. If you happen to have developers that want to just get in there and change stuff, that is also available. There is also an issues tab and a discussions tab on that GitHub repository as well. That is probably the most efficient way to get those discussions going. There is also just the general SEI feedback address that we can hopefully link in the show notes for this.

Dan: Fantastic. Maybe, Austin, can you give our audience some tips and tricks on kind of the easiest way to get started with IIDES?

Austin: Sure. I guess it depends on where you're coming from and why you're using it. If you're just trying to figure out what in the world you should track for your organization, the easiest way might be to just go to the GitHub repo and look at what classes and fields are in there. There is lots of different documentation in there about that. There is easy-to-read kind of text formatted stuff. There are pictures about how all this stuff is related. You might just go in there and say, *This is the subset of fields I want, and I'm going to put that in a CSV, and that's where I'm going to start*, and that's perfectly fine. I also mentioned the JSON schema. There are ERD, Entity Relationship Diagrams. If you are looking to make your own database out of this, that's available to you as well to develop your own. I mentioned PyIIDES, Python version of it. So it's sort of intended as a reference implementation, but it is available for use on PyPy. If you are a Python shop, you can use that. We are also hoping to put at least a version of our database available online, the schema itself, and the web application that we use at some point here in the next few months.

Dan: Fantastic. Helping organizations find ways to kind of, A, get the most out of the data that they have in this space, but then B, understand where are the opportunities for the collection of additional data to again address those kind of unique aspects of insider risk management. In my mind, that is the highest level kind of description of the value proposition using this? And, you know, near and dear, you know, to my heart is getting everybody to call the same thing about the same thing, or at least have a translation mechanism to get from when you're saying it like this in that organization or context, and this other one that's, *We mean this*. I am really proud of this work. I think you've done a great job kind of getting the team to the spot where we have this capability out there that I think will really help address some kind of operational gaps and challenges and kind of enable additional kind of research and refinement and capability development. I am thrilled that we have this out there. Really proud of the work that you and the team have done on this, and can't wait to see what we get in terms of feedback from the community on this one.

Austin: I am excited about that as well. Looking forward to feedback. And I just want to pull on that thread with the team. It is two decades' worth of work. So there are people that worked on sort of the original version of the schema who I have never even met. And I know just getting IIDES across the finish line really required the whole team and some of our interns as well. So it was a lot of work, and we're very proud of it, and we're happy to show it off.

Dan: Fantastic. All right, Austin. Now IIDES is out in the wild, and while we

anxiously await kind of community feedback on it, what's going to keep you busy over the next couple of months? Where are we headed from that kind of insider risk management research perspective? What are some things that you are thinking about here in the context of the role we play in this space as an FFRDC?

Austin: I know at least on my side, I really am interested in how we use the data to make changes. There will be some almost foundational trending and pattern analysis and things like that coming up here in the next few months to few years as we learn the right ways to use this and we learn what the community needs for that, so sort of statistical analysis on various types of insider threat data. I know I personally enjoy the AI side of things. There is a big opportunity to look at how AI can fit into insider threats and addressing those challenges. Everything from the risk scoring and how we do that appropriately and ethically and efficiently and accurately to just using AI as a tool to help us address some of the really tedious aspects of data analysis and data collection and things. Where does AI fit in that space? We are also doing some work with simulations and data generation from simulations and things like that. Hopefully, helping folks address some of the simulation and testing needs that we see sometimes. So that's some of the stuff that I'm interested in and I know about. Dan, you probably have a few other things as well.

Dan: I think you have hit the nail on the head, Austin, with a lot of kind of where we are focused from a research perspective, understanding the bounds of human-machine teams in the context of insider risk analysis. Helping organizations come to the realization that insiders now aren't just people. We've got to kind of redo everything we learned about human behavior in the context of machine behavior. We are talking about autonomous systems that are artificially intelligent, that have been given the same level of access to our organization's critical assets that we used to give to people. Removing the human out of the loop didn't make that risk just go away. It has transferred the risk to something we know less about how it's motivated and what it does. And then we do people, which is another thing we are going to have to collectively wrap our heads around. understanding the tech, the insider risk implications of technology-driven changes to this landscape. These are the things that we are here at the SEI from an insider risk research perspective. These are the things that are continuously keeping us trying to find, innovative solutions to these semi-indirect challenges. Austin, one other thing to put on your calendar here in terms of what is going to keep you busy over the next couple of months is the Insider Incident Data Expression Standard is going to be prominently featured at this year's Insider

Risk Management Symposium. Save the date for that, June 12, 2025, Arlington, Virginia. We would love to have you there. You will get a chance to see from Austin and team how this wonderful capability was built. We will walk you through a little bit of the guts of the implementation of putting this to practice. We will have slides to make it a little bit easier to follow. If you like what you see on the blog, spend some time on the GitHub page. We are going to dive a little bit deeper into this during our annual symposium this year as well, so be on the lookout for that. Austin, I want to thank you for taking the time to sit down and talk about this work. I'm really excited to get some community feedback with regards to where we take things next. I want to thank our listeners for joining us today. We'll make sure that we include all the links in our transcript that we've talked about, how to get IIDES, the blog posts, information about our upcoming Insider Risk Management Symposium. So anything else that may be applicable to this conversation, we'll do our best to kind of get it up there. Austin, thanks so much for the time. Enjoyed the conversation.

Austin: Thanks, Dan. My pleasure.

For our audience, we will include links in the transcript to resources mentioned during our conversation. Finally, a reminder to our audience that all our podcasts are available on SoundCloud, Spotify, and Apple podcasts, and the SEI's YouTube channel. If you like what you see and hear today, please do give us a thumbs up. Thanks again for joining us and looking forward to your feedback.

Thanks for joining us, this episode is available where you download podcasts. Including SoundCloud, TuneIn radio, 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 e-mail us at info@sei.cmu.edu. Thank you.