Carnegie Mellon University
Software Engineering Institute

SEI Podcasts

Conversations in Artificial Intelligence,
Cybersecurity, and Software Engineering

Cybersecurity Metrics: Protecting Data and Understanding Threats

Featuring William R. Nichols 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 M Miller: Welcome to the SEI Podcast Series. My name is Suzanne Miller. I am a principal researcher in the SEI Software Solutions Division. Today, I am joined by my friend and colleague, Bill Nichols, who leads our work in Software Engineering Measurement and Analysis. Bill has been a frequent guest on our podcast series, and we've always enjoyed having him. We will link to his digital library author page in our transcript because he has lots of things that he talks about, and you can find all of his publications there. In this podcast today, we are going to talk about Bill's latest work in cybersecurity, measurement. Welcome back, Bill.

Bill Nichols: Well, thank you. Glad to be here.

Suzanne: For those members of our audience who don't know you already, can you please tell us a little bit about yourself and the work that you do here at the SEI. What do you think is the coolest thing about your job here?

Bill: Well, my background is in physics. That is where I got my PhD in experimental physics, which is where I kind of learned about measurement. I didn't learn a whole lot about software engineering, except that we did a whole lot of it, not always great. I did nuclear engineering for the Naval Reactors program for about 15 years. That is where I really learned the art of software engineering and measurement. We had the SEI come in, including Watts Humphrey, so I learned from one of the best there.

After a few years of that I came over to the SEI and started working with our Team Software Process team teaching software measurement and using those measurements to control your software projects and programs. I have been here now for another 15 or more years. My work now is leading the measurement program. I think the coolest thing about it is some of the people I really get to work with, really an interesting cast of characters, really smart. It is really great to be able to come in every day and get to work with these kind of people.

Suzanne: I love the people in in your group as well. I know who you are talking about. And cast of characters. Yes, we are. They are a cast of characters. Well, most of us at the SEI would fit into that as being a cast of characters. It is part of what makes it interesting.

To begin our discussion today, you have talked about measurement in some different contexts, but not yet about the importance of cybersecurity measurements. What is the important thing about cybersecurity measurement? What kinds of measurements do we use in cybersecurity, and what can they tell us about a system that that we are trying to analyze?

Bill: Well, you have a couple of different questions there. The first thing is the importance of cybersecurity metrics. That gets me down to the most important thing in measurement is not starting with the measurements [but] starting with the goals. What do you need to know? What do you need to do. What kind of decisions are you going to inform? In the world today it is hard to imagine anything being more important than cybersecurity because it is in everything. Everything we own is potentially vulnerable to cybersecurity, our identities, everything. It is about protecting ourselves. It is really critical to a modern digital, interconnected world. What was the second part of that question? What are they?

Suzanne: What are they? What can cybersecurity measures tell us about a system?

Bill: There are a number of different things. That gets down to the question of what aspect of the system are you looking for. The broadest categorization is everything to do with protecting information data and access to your system. The metrics are typically going to be involved in understanding what are the threats? What are the potential exposures in your system, whether it is in the environment, or whether it is something you could measure within the system? It would even extend to things like what is the value of what you are protecting. What are the risks if it is damaged or exposed? There are all sorts of things you could measure. What aspect you are looking at any given time will depend…that will drive what you actually do want to measure.

Suzanne: So that goes back to the goal-based measurement.

Bill: It goes back to the goals absolutely.

Suzanne: Alright and obviously at different points in the lifecycle we are going to have some different goals. We always want to protect access to data, access to even things like being able to manipulate the code. All of those things are going to have impact at different points in our lifecycle. Some of them sort of are similar I would think to some of our software engineering metrics in terms of how many of these things have happened, and how many of those things have happened. But what are some of the key differences between cybersecurity metrics and traditional software engineering metrics?

Bill: That is great question. I think the short answer is software engineering metrics by and large deal with the economics of building software engineering products. You are going to be talking about things like, *What is my productivity*? whether it is in lines of code, or whether you are talking about stories to try to meet deadlines. You are going to be talking about issues of quality. You will have measures of things like, *What is my defect rate? What are my find rates? How much are escaping?* Those are sort of pushing into security-related, but they aren't really there. When you get into security, you are really looking at things that potentially are going to expose data or expose your system to a compromise. There are some categorizations. For example, Mitre puts out the Common Weakness Enumeration of a bunch of code-based or environment-based paths that are potential weaknesses or the vulnerability list, which represents actual exploits. These are the subset of those qualities that you would be worried about in software engineering. You are normally going to be worried about things like extensibility, maintainability, in addition to the functionality. Cybersecurity focuses on those that specifically relate to securing your system and your data.

Suzanne: You have really started to talk about what are some of the challenges of cybersecurity measurement. One of which is, *Where is my data? What is it that I need to protect?* I heard you say that earlier, *What is it that we* *need to protect?* What are some of the other challenges of measuring cybersecurity?

Bill: Well, one of the big problems in cybersecurity metrics is just scoping things down and getting clear definitions and standardization for what it is you are trying to measure. That is what are the underlying concepts that you are actually concerned about. It is very hard to get a consensus on what kind of metrics you should be measuring, or what are the specific things you should be measuring in a specific contest for purpose. There are lots of nuances also depending on the industry. You have got a very different scenario if you are securing, say, the software systems in your automobile versus an online bank.

All sorts of issues there. Then, you have the issue of data availability. Here I don't mean the data you are protecting. I am talking about the data related to security. Things like code bases tend to be held very closely. That is why a lot of the research is using open source because that is available. But the proprietary is very closely held. Very often you can't get good measures on the code, certainly not on the incidences that have actually happened. A lot of things are happening kind of behind the curtain that we aren't fully aware of. We are trying to measure this system, but we are only seeing parts of the system that we are allowed to see. There is a lot more out there, and that is a whole separate problem that leads to things like data silos. Then you have the unknowns. Because you have an active adversary who is trying to figure out there is like a cat and mouse game. You take certain actions, they react, they try a different path. It is not a stable system. You have these unknowns and this constant evolution as you are playing this game. It is not really a stable field.

Suzanne: I heard two things. One is the code bases that you would want to get into to analyze how certain things happen are essentially black boxes to external parties outside of whoever owns the code. Then the second thing is this dynamism of the system, the fact that as soon as we see an exploit, we try and recover from it. We change things, and, if we are lucky, we don't introduce any other exploits. But if we are not, we can introduce sorry exploits, vulnerabilities. If we are not, then we could still have problems. Just because we fix something doesn't mean that everything is fixed. Is that correct?

Bill: Correct, and to be honest, we don't even know about all the exploits. It is hard to be certain that certain things haven't been out there for years and are just being quietly exploited. These things can lay hidden. That is another challenge in the metrics. You just don't know what you don't know.

Suzanne: Okay? Well, sounds like you are going to be busy for a while.

Bill: Or someone will.

Suzanne: We have been talking kind of abstractly for the last few minutes. Can you walk us through a real-world cybersecurity incident and help us identify potential measurement possibilities and how that information could be leveraged to avoid something similar or even something different in the future.

Well, let's see. I think one of the nastiest that I can think of that is public…Part of the problem is for these incidents a lot of times they aren't really public. They are going to be investigated quietly. But working at the SEI, you are probably familiar with the vulnerability or the attack, and I think it was 2015 on OPM, the office of personnel management.

Suzanne: Oh, yeah. My data is in there too just like yours.

Bill: We don't know officially who was behind the attack, but they have their suspicions. The cost of that attack is hard to underestimate. Literally all of the QNSPs that we file. Those are the forms that those of us with security clearances fill out with all sorts of personal data were compromised. So all of that information is in the hands, presumably of some bad actor, probably a nation state. All right. That is about as serious as you can get.

What are the opportunities there for measurement? Well, one of the things that they were able to do is they did have to my understanding some good records and logs. So having fine-grained data available in the system logs can help you do things like track down where was the exploit specifically, how was the exploit executed. If memory serves, I think in that situation it was an access, some ability to get access to the system and then elevate privileges. With the data, you can try to determine where was the vulnerability exposed, so that you can potentially remove that in this system and harden other systems. One of the lessons from that one is probably the importance of doing things like having 2-factor [authentication].

Those are the kinds of information you can get. I think having the granular data, having robust access management in place and being able to trace. Another thing that I would think of here is from an operational standpoint, one of the things you can do is you can be constantly monitoring the network so that you can understand what is the normal profile of network activity where the IP is coming from. And, by monitoring this sort of information, you can identify outliers. Often you can identify IPs from suspicious sources.

Suzanne: Gotcha. So there is a combination of developmental things like monitoring in terms of the code base. Looking at access points, where are the access points that are available in the code. Can we make sure that those are hardened, as you said, to having actions that we don't want to happen.

Bill: Right. You have a combination of being able to look at the source code. You can look at designs. You can look at the operational characteristics like the network. And then forensically, you can try to go through some of the logs and try to understand how exactly was this exploit executed. What part of the system did they attack? What was the exposure or the attack surface. But it helps you get better for the next time.

Suzanne: One of the things that I heard you say is that that we don't have a lot of standards in this arena for measurement. When I think of standards in this case, I would think of having standards for example, how we comment code that is meant that deals with say PII [personally identifiable information] private data. Or, how we comment code that is going to make the data available for access. Are those the kinds of things you are talking about where there is really no standard for that, even though those are things that I, as a software engineer, would think about doing to help make it easier for people to find where there was a vulnerability if something does happen. Am I on the right path?

Bill: You are on the right path, but I think it is even more meta than that. Even something like the Common Weakness Enumeration has a lot of issues with duplications and redundancies. It is certainly not perfect, but it is probably the best thing we have right now for categorizing the data. Then you have these other questions about well, what exactly should you be measuring on the network? There are certainly things that people use, and there are things that have been found useful, but the research doesn't have a gold standard. If you take a measurement, you know exactly how that measurement was taken at different locations. What does mean time to recovery really mean, just as an example? Well, that depends a lot on how you count. So that presents all sorts of problems.

Suzanne: Let's just take that incident right? The meantime to recovery has multiple components. One of them is when did the incident occur.

Bill: Which you may not know.

Suzanne: And what does recovery mean? Does recovery mean back to a normal state, or does recovery mean back to an acceptable state? Those may not be the same thing.

Bill: Exactly. And those are the kinds of semantic definitional problems you run into if you try to aggregate data and look at trends across the industry, a lot of data cleaning.

Suzanne: I was just thinking about that. The whole world of data cleansing has just really exploded between AI and cybersecurity. You know the people that enjoy doing data cleansing should have jobs for a long, long time. Are there solutions or approaches that you have run into that you would recommend for cyber practitioners in the field who want to capture metrics of the code that they own and the systems that they are operating?

Bill: The best advice that I would give right now is a little on the generic side, because it is hard to be precise. Everything depends on your context, but always start with the objectives of what you are trying to measure. That is always what drives what you should measure. So try to understand what is the concept? What is it you are actually trying to get to, and then look for the mechanisms that you can get indicators. Rather than just collecting a bunch of data, I always recommend, try to do some sort of validation. Make sure that the data you are collecting is actually consistent. To the extent possible you would like to automate the data collection. You would like to see does the data have some variability, some trends. You want to understand what is the normal behavior? How would you recognize outliers? Anomaly detection is often one of the clues that there is an exploit or vulnerability out there.

As far as frameworks, there are a couple. CERT has the Capability Maturity Model for Cyber, which is worth looking at. NIST has some standards out there. They have some strengths and weaknesses, but that is certainly one of the things that is approachable. I think the most important thing is when you are measuring know what you are trying to get at, and make sure that you define the measures that you are actually using to get at that underlying concept very precisely so you can reproduce measures consistently.

Suzanne: What you are talking about is actually a framework that we have been using at the SEI for decades now is called GQIM, Goal Question Indicator Metric.

Bill: Exactly. Yes.

Suzanne: That is one the things we will reference in the transcript, because it is not only applicable to cybersecurity, it is applicable to measuring anything that is of interest. Let's face it, measurement costs money. One of the reasons for GQIM and for narrowing down not just what you can measure, but what you should measure, is so that you don't spend money on items that aren't going to give you answers that are going to be useful to your context. That is a whole economic aspect.

Bill: That is another one of the big challenges in cybersecurity metrics. One of the things you are trying to do with the cybersecurity metrics is secure your system. How much are you willing to spend to secure it? Am I going to have a little master padlock, or am I going to have a bank vault, a big difference in cost? The cost of cybersecurity can go way beyond a bank vault. Unless you have unless you have some consistent ways of measuring, even getting that trade-off is going to be a real challenge.

Suzanne: One of the things that I have spoken with some of our other colleagues in CERT [about] is there has been a lot of work in the last few years on threat analysis. Is that one of the sources for the kinds of measures, the kinds of goals? If I have got a goal to keep my system secure, then one of the things I need to know is about what kinds of threats are there to my system. Threats to a banking system, and the threats to a healthcare system may have some similarities, but also probably a lot of differences. Is that an area where…. That is something that we don't…In software engineering I don't really have a threat space for software engineering. That seems like that is an area that is of difference, where cybersecurity is looking at those external kind of factors that I may not be as conscious of when I am thinking about software engineering,

Bill: Yes, exactly. Part of the big difference…Cybersecurity is a system. It is not just a property of your software. It really is a property of your system in use. One of the things I definitely would recommend is before you set up a measurement program you really have to start with things like doing a threat analysis, examining the attack surface because that gives you some real understanding of what are the sort of things you really need to focus on for your measurement goals.

Suzanne: Also, how much should I spend, right? If my attack surface is very limited, I have a lot of air gaps and things in my system that really limit, you know where there could be vulnerabilities. That is a different case than something that is out on the web and available to the whole world that presumably needs to be. That is going to change the cost. That attack surface is going to change the cost of collecting the measures, correct?

Bill: Oh, absolutely. You are from the same generation as I am. You have probably seen many of the same examples of these systems that were perfectly useful in their day because they were on this air-gapped PC. Before we had Wi-fi networks. Then these things somehow ended up on network systems and became vulnerable. That kind of problem just repeats over and over again, as we have become a more and more connected world.

Suzanne: Agreed.

Bill: It is not just where the system is today. Every time you move a system into a different environment. you have to reconsider the entire security posture. One of the nice things about the GQIM approach is that the environment is one of the explicit considerations you go through, and the steps through, for understanding the measurement that you need to do to answer your questions.

Suzanne: Right. And when that environment changes that could either directly change the measures or indirectly change them, possibly not at all. But chances are you are going to have some change.

Bill: But you always should look at it. As they say, context is king or queen, depending on your preference, but the context can change everything.

Suzanne: Agreed. A big part of our work as an FFRDC is transition. If you can highlight what are some of the things our audience can do to learn more about cybersecurity measurement, how to implement cybersecurity metrics? You mentioned NIST. You mention MITRE CWE. What are some things either from the SEI or other places you would also suggest people look at.

Bill: Well, one of the practical books I would recommend is look up Carol Woody and Nancy Mead's book, *Cybersecurity Engineering, A Practical Approach for Systems and Software Assurance*. There is a lot of great technical information in that one.

For a more general measurement approach you might look up Doug Hubbard's book on cybersecurity risk [*How to Measure Anything in Cybersecurity Risk*]. He goes into a lot of examples on how to do estimates without a ton of data. It is a very principled, economically practical approach.

Suzanne: Excellent. Okay. I don't know Doug's book. I may have to look that one up. What is next for you? What are you working on now? What are we going to be able to bring you back to talk about in the future?

Bill: Well, I have, the next stage of a research program we have had going for several years on extracting data from the DevSecOps pipeline and the surrounding information systems, like the ticketing systems, to get a better handle on software development and productivity and tracking. We call it ACE TOPS.

We are taking the automated continuous estimation, using the data we collect to estimate the completion dates, schedules, and uncertainty in delivery for software engineering.

We are working on transitioning it with the kinds of things that would include the data collection tools, the types of visualizations you would like to have available. It is trying to take a lot of the approaches that you are seeing in modern software development with gathering all of these metrics. But it is a much more focused approach on understanding how do we use this to control software projects and larger software programs and take action at a reasonable time?

Suzanne: Oh, excellent! Well, that sounds like you are you are going to be busy with that. And you are going to be busy with cybersecurity measurement. I can see that. I want to thank you for talking with us today. I look forward to talking with you in the future as well as some of these other things emerge. For our audience, we will include links in the transcript to the resources mentioned during this podcast.

Finally, a reminder that our podcasts are available pretty much any place. You can access podcasts, including SoundCloud, and Apple. And my favorite, the SEI's YouTube Channel. If you like what you have seen here today, please give us a like. If you have any questions, please don't hesitate to contact us at info@sei.cmu.edu. Thanks again for joining us.

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.*