

# SEI Podcasts

Conversations in Artificial Intelligence,  
Cybersecurity, and Software Engineering

## AI for the Warfighter: Acquisition Challenges and Guidance

*Featuring Carol Smith and Brigid O'Hearn as Interviewed by Eileen Wrubel*

*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](https://sei.cmu.edu/podcasts).*

**Eileen Wrubel:** The Department of War is looking to get software-driven systems into the hands of the warfighter more quickly than ever and eliminate all the barriers that stand in the way. On Friday, November 7, the department released [an aggressive acquisition transformation strategy](#) that doubles down on removing bureaucratic hurdles and streamlining acquisition processes to enable even more rapid adoption of technologies like AI [artificial intelligence]. AI has become a key player in helping the government acquire, develop, deploy, and maintain these systems at a cadence that keeps them relevant for our armed forces. But it is not a magic wand. Getting it right requires disciplined AI Engineering.

Welcome to the SEI Podcast Series. My name is [Eileen Wrubel](#), and I am a technical director here focused on transforming software acquisition policy and practice. Today I am pleased to introduce our podcast guests. My friend [Carol Smith](#) leads human-centered AI research in the [SEI's AI Division](#). And

my colleague [Brigid O'Hearn](#) is on my team and works with government and defense agencies on their software modernization policy and practices. We are here to discuss the [unique challenges associated with using AI in government and defense applications and how government teams can build effective acquisition processes for](#) differentiating AI systems, guidance on when to use AI, and matching AI tools to mission needs. Carol and Brigid, first, I want to thank you for joining me today on the podcast series. I appreciate you sitting down to talk with me about this important work.

**Carol Smith:** Thank you for having us.

**Eileen:** Before we dive into the business of putting AI in the hands of the warfighter, I would really like to start by introducing ourselves and talking about the cool experiences that led us to the SEI and what we do here. Carol, I would like to start with you.

**Carol:** Great. My undergrad was actually in photojournalism. I was following all kinds of interesting things for a brief while anyway in that career and really found that I needed something a little bit more engaging and looked for a master's program and ended up finding the [Human-Computer Interaction Program at DePaul University](#) when I lived in Chicago. Then I completed that degree and started consulting. I have been able to do all kinds of really interesting work observing people in a lot of different complex situations through the SEI with the military and then the manufacturing facilities, coal mines, hospitals, offices, in people's vehicles as well as in coal trucks. The work is really turning those observations into recommendations for computer applications. I have been working with AI systems since 2015, including with autonomous vehicles. In 2019, I joined the SEI AI division, and I have been able to collaborate with our teams here to make systems that are trustworthy, human-centered, and responsible. I also teach as an adjunct in the [CMU Human-Computer Interaction Institute](#).

**Eileen:** Oh, very cool. I didn't know so much about your journalism background, and I didn't realize that you were teaching now. That is awesome. Thanks so much for sharing that. Brigid, let's hear about you.

**Brigid O'Hearn:** Yes, so, I have a bachelor's in computer science, and then I have a master's in logistics management. My second master's is actually in national security resource strategy from [The Eisenhower School](#). If you are not familiar with it, it is one of the national war colleges that are part of the Department of War. It was a really neat degree where you get to learn about how national security impacts programs across the entire department. It is a

pretty neat experience. I got to do that 10 months in residence, and I focused on the space industry. You get to pick an industry that you want to follow and really had a neat opportunity to work with space industry, not just within the U.S. but across the globe. I got to see how space works across the planet, kind of neat. My background has always been embedded in software engineering, but I have worked in or around the Department of War for the last 30 years. That has been a pretty neat opportunity to work both for the Air Force, for the Navy, and help software acquisition programs, probably about 20 different software acquisition programs that I have worked with in my career. I have been at the SEI now for four years. It has been a really cool opportunity to be able to take all of that experience that I have gained over the 30 years and help build that in, help programs, and help build that back into policy that can be used to help even more programs. Pretty neat. Back to you, Eileen.

**Eileen:** Thanks, Brigid. Just for a little bit about me, I often joke that I have been serving in the military or associated with it since the day I was born. I grew up an Air Force brat, lived all over Germany and various parts of the U.S. I attended Carnegie Mellon on an Air Force ROTC scholarship and spent about six years working in mobile communications and working in the private sector in the software field before Carnegie Mellon came calling once again. I have been at the SEI helping programs and agencies acquire software-enabled capability better for the last 22 years now. I am not good at leaving the military or leaving Pittsburgh it turns out. My area of practice here, the team that Brigid and I are on, focuses on integrating analytical insights and software engineering research and acquisition science to help the department scale workforce capability, reduce barriers, simplify acquisition processes, and accelerate the adoption of scientific and commercial innovation to serve warfighter mission needs. That is why you have the three of us here today to put all of this stuff together and talk about how does AI Engineering really affect what we are doing? We know that it is an essential discipline for enabling the acquisition, development, and deployment, and maintenance of AI-enabled systems. What I would like to do is start with talking about a workshop that Carol and Brigid and some of our colleagues convened back in June about acquiring AI-enabled systems. Can you talk to me a little bit about the why of the workshop that we put together? Why are our teams across the institute working together to team on AI and acquisition? What is the big problem we are trying to solve? And I will start with Brigid.

**Brigid:** All right. I have to say, Carol reached out to our acquisition team and said, *Hey, I'd really like to partner with you all.* Carol is an expert in AI. I'm not,

but I'm an expert in acquisition. So it was a perfect partnership when she said, *Hey, would you like to do this together?* Wow, I jumped at the chance. Carol is awesome to work with. I said, *Yes, I would love to.* But even better than that, to really be able to engage with the AI community and understand some of the unique challenges that they are seeing, to then, again, inform the policy work that we are doing, and to help build that into making software acquisition of AI even better. Carol, do you want to give your perspective?

**Carol:** Yes. I was thrilled that you said yes. It really worked out very nicely. Really the purpose of the workshop was to help people. They already had, in some cases anyway, some acquisition guidance. Certainly, the SEI provides a lot of that. That is a strong foundation. But it wasn't sufficient yet to identify and effectively assess dynamic AI systems and large language models. We wanted to present new ways of thinking about these systems and try to really help those teams who have concerns about falling behind to prepare for the future. We really were working to bring together these people so they could learn from each other but also share with them some of the research that we had found and help them with making the decisions that they need to make to move forward.

**Eileen:** Fantastic. Let's start with some of the big challenge areas. AI is exploding everywhere. My kids are using AI tools to help them execute school assignments or write invitations to parties or speed through the manuals to their cars. It is exploding everywhere. But there are unique challenges associated with using and adopting AI in national security settings. How are those settings different from what practitioners in more commercial settings might encounter, Carol?

**Carol:** Certainly, for all of the tasks that you mentioned, some of the newer systems, the generative AI systems are not only very helpful but fun in many cases to work with and can save a great deal of time. But with the Department [of War], we are looking at a very different risk profile. A lot of the projects, the work that is being done for mission is much higher risk. And those implications of training data or models or the abilities that we might give an AI system can lead to the systems learning the wrong thing, doing the wrong thing, or revealing the wrong thing. That could be of high consequence. In addition, it is very difficult to sometimes escape the hype. And so helping people just to understand these systems and the capabilities as well as the limitations. For example, generative AI systems, LLMs, they can create a really impressive facade of understanding and semblances of meaning. But it is mostly the spelling and grammar that is taken care of that

creates that perception. We want to make sure that people are able to identify errors in content, able to identify the provenance of the results, and able to have the time they need to double-check as necessary with those systems in particular.

**Eileen:** OK, great. Speaking of practitioners in national security settings, I know there was a lot of discussion at your workshop around [Project Linchpin](#), which is the Army's artificial intelligence accelerator. What lessons did the Project Linchpin team have to offer at the workshop? Brigid, can you start by taking us through some of that?

**Brigid:** Yes, Project Linchpin, for those who aren't familiar, they are the Army's artificial intelligence accelerator. They actually started out as a Software Acquisition Pathway program. That is how we were first introduced to them. But their strategy has evolved over time, so now they are more of this enterprise function. But they had several years of building up their capabilities including things like creating this very modular and open architecture. That is one of the lessons that they shared with us. They talked about having this flexible environment such that you could bring in different vendors. They were easily able to connect into the architecture. Using [APIs](#), [MOSA](#), things like that are one of the lessons that they learned over the time that they evolved. Another one is planning for legacy integration. So, of course, I think we all know that the department has thousands of legacy systems. They are not going to change overnight, so having the ability to integrate with those systems, knowing that they are built on older architecture, infrastructure, even older languages—the ability to have that flexibility to be integrated in. They also talked about making the outputs explainable. In order for the operators to really embrace and adopt and trust something, they need the ability to understand how those outputs came about. How did they start? Where did they get their information from? How did they make those decisions? Making sure those operators are comfortable is another lesson they shared. Focusing on the field performance. So really focusing on how is it going to work in operations, not how is it going to work in a lab, but how is it actually going to work in terms of the operational environment in which it is going to live? That is a big key because we know things will work differently on a battlefield than they will in a hardware-in-the-loop lab. Really making sure they focus there.

Lastly, managing data bias carefully, so making sure that the models are trained correctly, that they really understand how things are introduced, and that they continue to look at how things evolve over time. I will say as one of the neat exercises we did in the workshop. We actually had groups break out

and look at data that they might be provided for a system that they were acquiring. Look at how the data was put together, what the images told us, what might be missing. It is a really neat way to connect that. And [Project Linchpin](#), by the way, gave us these lessons learned. We already had this as a plan for the workshop. It really connected very nicely to be able to pull that together. But really Project Linchpin, the lessons that they shared with us really emphasize testing and transparency and responsibility in AI. And they are a great example of how that has evolved over time as they have continued to build their expertise.

**Eileen:** OK. I want to roll back just a little to you talked about focusing on field performance and understanding how a model or how an AI-enabled system is going to perform in live missions. Can we talk in a little more detail about the importance of aligning these tools with real mission needs? What insights can we offer to project and program teams about making sure that they have the right team or the right tool for the mission that they need to accomplish? Brigid or Carol, I don't know which one of you would like to bounce that ball.

**Carol:** Yes, I can start off anyway. Part of it is recognizing that AI is just another tool. It is technology that is not appropriate for every situation, and figuring out what the right tool is, understanding what those needs are, what the problem you are trying to solve is, and if it is an AI-appropriate problem to begin with. Often, we find people are trying to solve an area of consideration that is just too broad, and AI systems typically are not good at broader problems, or they have limited access to data, or their data itself is limited, or their data includes some harmful imbalances. It is just low quality. And so that can limit the ability to make a system that's really going to be helpful.

Then in some cases, humans are faster or more accurate, more flexible. There are many things that we are still better at. In some cases, it doesn't make sense for a variety of reasons to automate an activity, particularly if there are limited resources, so helping them understand what resources are needed. AI systems can, of course, do many things that humans are not as good at. But we still need people to support them, maintain them, make sure that the data is what is necessary, and that it is not undergoing drift. There are a lot of activities that need to be done that are new in addition to software management-type activities.

**Brigid:** Oh, thank you. I was going to say I wanted to add on to what Carol said. One of the things that I thought was a nice connection is one of the

pieces that we talked about in the workshop was making sure you get the feedback, right? You get the feedback from the users. That is something that, for those of you that are familiar with the Software Acquisition Pathway, that is a really big integral piece to the software pathway is making sure you have the users involved to get that feedback in real time as quickly as possible. You build in what they need, and you give them the capabilities that they can use as quickly as possible. That was one big piece that I really appreciated from the workshop, was that connection there.

**Carol:** Yes, and just to build on that, too, prototyping, even without doing any coding, just having a prototype that people can start to use in their environment, so a clickable prototype that might need to be thrown away. It might need to be discarded. It may not be the right thing, but at least you find out quickly before spending the time and money and investment in building an AI system. Doing those operational experiments with people who are actually going to use the system in their environment is really important and can save a lot of time and frustration later.

**Eileen:** I would like to spin back a little bit to something that you alluded to earlier, Carol, which is the role of human oversight. AI can't assess its own effectiveness, right, or understand the significance of its output. Can you talk to us a little more about the importance of human oversight of AI-enabled systems?

**Carol:** Yes, certainly. The systems can of course be monitored to some extent in an automated way. But often, the statistical information that you can learn about a system isn't necessarily a representation of how it is actually working. A system could potentially be 100 percent accurate but actually be providing unhelpful information to the end users or making recommendations that aren't applicable or just wrong. And so having those individuals who understand the subject matter, so subject matter experts, who are able to see the system in action and note when it is not working properly, as well as providing some of that information to the end users so that they can decide, *The system is just not performing as I expect*. And being able to choose to use a different tool or choose to solve the problem in a different way, if necessary, is really important.

**Eileen:** Brigid, can you talk a little bit about maybe any other guidelines that came out of the workshop that government teams should consider when evaluating AI solutions?

**Brigid:** Absolutely. One of the things we talked about is the Software

Acquisition Pathway has some interesting pieces that are geared toward either applications or defense business systems or embedded systems, but it doesn't yet have an AI pathway. That is something that is coming. That is actually a recommendation that came out of the workshop, is to really make sure that we build some of the things that we talked about within the workshop into the pathway. We have been, Carol and I and others, have actually been working with a policymaker in the department on that AI-acquisition subpath. That is a nice piece that allows us to be able to influence some of what is there. Carol, did you want to add some other things there?

**Carol:** Yes, and one of those aspects is the data. Just making sure that people understand how key the data becomes within an AI system. Whereas in the software system, the software system is displaying information, and humans are determining how to interact with that. Within the AI system, the data is actually the system. It becomes even more important to ensure that we really understand that information, that is the right information, and that we are aware of the inherent biases that the information holds so that we can determine if it is applicable to the particular situation. That is one of those aspects that we hope will be included in the new guidance.

**Eileen:** Fantastic. What are our next steps at the SEI on this front? What can we get the gang back together to talk about on a new podcast in the near future?

**Carol:** Well, in the AI Division, we have been conducting a national AI Engineering study. We anticipate sharing the report early next year. After our workshop, there was also an AI data quality workshop for data pipelines. And then just last week, a few of our SEI colleagues were presenting at the [AAAI Symposium on Engineering and Safety-Critical AI Systems](#). There is a lot of information coming that we will be able to hopefully talk about soon and looking forward to it.

**Eileen:** Great. And Brigid, what would you like to add there?

**Brigid:** Yes, some of our viewers might be interested in the [Software Acquisition Go Bag](#). Eileen and I did [a webcast last month on the Software Acquisition Go Bag](#) and the work that we are doing there to help acquisition programs, making things a little bit easier, providing them some of our expertise and lessons learned. Carol, you and I, I think, are going to collaborate on an asset for the Go Bag. We can actually take our AI and acquisition expertise and build something in that will help those that are out there acquiring new AI programs. So that is a pretty neat collaboration.

**Carol:** Very cool.

**Eileen:** It will be just like having Brigid and Carol sitting at your desk right with you.

**Brigid:** Wonderful.

**Carol:** It will be fun.

**Eileen:** That is the goal. Well, thanks, folks. I want to thank both of you first for taking the time to sit down with me today. I know everybody is super busy. And I want to thank all of our listeners for taking the time out of their day to click on this podcast or visit us on YouTube. We will include links in the transcript to all the resources that we mentioned in this podcast, including the [Software Acquisition Go Bag](#). The SEI Podcast Series is available in all the places that you regularly find your podcasts, [Apple Podcasts](#), [SoundCloud](#), [Spotify](#), and the [SEI's own YouTube channel](#). As always, if you have any questions for us, please don't hesitate to email us at [info@sei.cmu.edu](mailto:info@sei.cmu.edu). Thanks.