Assurance Cases

featuring Chuck Weinstock interviewed by Suzanne Miller

Suzanne Miller: Welcome to the SEI podcast series, a production of the Carnegie Mellon University (CMU) Software Engineering Institute (SEI). The SEI is a federally funded research and development center headquartered on CMU's campus in Pittsburgh, Pennsylvania. A transcript of today's podcast is posted on the SEI website at sei.cmu.edu/podcasts.

My name is Suzanne Miller. I'm a principal researcher here at the SEI. Today, I am very pleased to introduce you to Charles Weinstock (we call him "Chuck") a principal researcher here at the SEI whose work focuses on dependable computing with a recent emphasis on assurance cases and measuring confidence, one of the topics that I am very fond of. In today's podcast, we are going to be talking about assurance cases and their role in verifying safety-critical systems. Welcome, Chuck.

Chuck Weinstock: Hey. Thanks for having me, Suzanne.

Suzanne: Chuck, let's start off by having you explain exactly what an assurance case is and where they came from? How have they evolved since their inception?

Chuck: So, before talking about assurance cases, I need to talk a little bit about safety cases because assurance cases are a generalization of the idea of a safety case. The safety case, of course, is a means of justifying that a system will meet its safety properties. It was realized along the way that the same ideas could be applied to properties like security or reliability or usability or any of the *-ilities* or any attribute you would care to apply it to. So, we renamed it assurance cases. It's not just us; it's the community-at-large that's doing work in the area of assurance cases. The assurance case is similar to a legal case in structure. The term *case* in the assurance case or safety case is meant to evoke a case in the same sense as legal case. Just as a lawyer makes an argument based on evidence that his client is innocent, or a prosecutor makes an argument presumably based on the same evidence that the client is guilty, an engineer would make an argument based on some evidence—for instance, testing results, static analysis, formal proofs or what have you—but the system that he has developed is safe or secure or reliable or whatever claim they are after.

There's a long history of development of safety cases in Europe, particularly in the United Kingdom. The practice is required for systems in diverse areas such as avionics, as exemplified by the United Kingdom Ministry of Defence's regulation 00-56, railway signaling systems, nuclear power plants, and other areas.

Suzanne: So, these safety cases or assurance cases, I've looked at some of these, and one of the things that is really striking is—even if you don't really know the system very well—what's really apparent is where you've got lots of evidence for some claims and not very much evidence for others. It's a really wonderful visualization of what is it that you have available to make a claim stick in terms of your reliability or your security and that sort of thing. Have you found that aspect to be useful in communicating the systems with…

Chuck: You certainly can look at it and see [that] there's a lot of evidence here, and there's a little evidence there, but quantity doesn't necessary make for quality.

Suzanne That's also true.

Chuck: So, you may have a ton of testing results that are superficial and a very deep formal proof over in another part of the case that nails it basically. So, the visualization of the case—which is the way they're usually presented, by the way, via notations such as Goal Structuring Notation, developed in the United Kingdom by Tim Kelly, or as text—doesn't necessarily give you a clue, but it's something.

Suzanne: It's just the beginning of a conversation.

Chuck: Yes it is. Exactly.

Suzanne: Without that visualization, you really don't have an understanding of what kinds of evidence you have for the claims that you're making.

Chuck: Right. I guess the point is given a pile of evidence and a claim that the system is safe, somebody looking at this pile has to decide, *Is it safe*? Without the argument that links them together, it's very hard for especially somebody who's not deep in the guts of the system to understand how this test result actually shows that this system is safe because it does this and it does that and that does this, and so on.

Suzanne: So, the art part of this right now is creating those arguments in a way that is compelling, one, and also correct. So, that you can understand is this a set of trivial test cases or a set of test cases…

Chuck: Right. And, again, even a trivial test case might show something important.

Suzanne: Sure. True.

Chuck: The idea is to start with the top-level claim, or start with the evidence and head towards the top-level claim, and make baby steps along the way that lead the reader inexorably towards the conclusion. Yes, this evidence supports this claim, and the system is safe.

Suzanne: Or secure, or reliable, or dependable.

Chuck: I'm just going to stick with safety just so I don't have to keep saying *or, or, or*.

Suzanne: Fair enough. So as I understand it, we've got some federal agencies that are starting to embrace this idea of assurance cases. Could you tell us a little bit about some of the initiatives you're seeing in our own federal agencies here in the United States?

Chuck: The U.S. federal government has gotten slowly involved in assurance cases. We've been pushing this ourselves since about 2003, so for a decade. We are making some progress, and that's pretty good. In around 2007, the National Research Council of the National Academies of Sciences issued a report called *Software for Dependable Systems: Sufficient Evidence?* In it, [the report] studied what it takes to make dependable systems, and issued some recommendations. A key one that it issued was that for testing to be a credible component for a case for dependability, the relationship between testing and the properties claimed will need to be explicitly justified.

Suzanne: Sounds like an assurance case to me. How much did you pay them, Chuck? [Laughter]

Chuck: We had nobody, unfortunately, on that committee, but we were very happy with results.

Suzanne: That actually makes the recommendation stronger.

Chuck: Right, exactly. The point of the assurance case and the thing that the government is understanding is that it creates an artifact that allows them or their independent assessor or whatever to evaluate that the evidence shows the claim's been satisfied as we've been saying. Because the argument has to be understandable to the independent assessor, it means it's got to be much more detailed than an engineer might do internally.

Within the government itself specifically adopting assurance cases, the major U.S. initiative in this has been led by the U.S. Food and Drug Administration. The FDA is responsible for assuring the safety and effectiveness of medical devices, among other things. They found that they were approving devices called infusion pumps that were causing fatalities down the road, and there were more of these than they were comfortable with. They are not comfortable with any, but there were a fair number of fatalities coming from this.

An infusion pump is a device that injects fluids into a patient under treatment. That might be in the hospital; it might be at home. An insulin pump is an example of an infusion pump. If there is an internal error in the pump that causes it to shut down….

Suzanne: Which could be a software…

Chuck: It could be a software…

Suzanne: These pumps do have software that helps to monitor, provide data, all kinds of things as well as the hardware that actually administers…

Chuck: Exactly. The pumps are more and more software. Thank you for mentioning that. There are software enabled libraries of, drug libraries in these pumps that…

Suzanne: Give you dosage information…

Chuck: Exactly.

Suzanne: Per kilogram of weight and that kind of thing …

Chuck: And in particular settings. So in the emergency room it's a different setting. You might allow more drug-per-second in an emergency room because it's critical as opposed to in a hospital room.

Suzanne: So, making sure that these devices do what they say they're going to do and do it in a way that maximizes safety and prevents as many errors as possible is a huge goal for the FDA.

Chuck: Yes. Exactly.

Suzanne: Where do assurance cases play into that?

Chuck: OK. So, the typical submission—I don't give numbers so you're just going to have to take me, my word on this—big piles of evidence, and other materials like hazard analyses and things like that. An FDA assessor is allowed, by statute, something like 90 days from receipt to make an approval or rejection. They can't just reject because they don't have time to look at it.

Suzanne: They have to have a rationale.

Chuck: They have to have a rationale. They're given this huge pile of data, and they may or may not understand what they're seeing. There's no way they can look at it all anyway. By having an assurance case, you've got a link from the big pile of data and evidence to the claims. The FDA examiner would be able to look at the assurance case and cherry pick evidence based on the argument and understand how it all fits together, and make a good assessment. That's what they are trying to do because it makes their life better, and it makes the assessments of…

Suzanne: Higher quality

Chuck: Higher quality, right, exactly. They have issued draft guidance on this. I think it is on the way to being accepted as full guidance, but in the meantime the industry is starting to use

assurance cases because they see the writing on the wall. It's going to happen whether it's exactly this way or not. That's a big one.

The DoD has used it in several situations. We did a report on supply-chain security that looked at how you deal with the fact that you don't always know where pieces of software that you're integrating into your system come from, and how do you deal with that in a way that assures that. For instance, there's no back doors in that software that you don't know about and things like that. We used an assurance case technique to lay out how one would go about doing that.

We've used assurance cases with the Department of Defense to show that testing a component of an aircraft in a research lab was equivalent to testing it on the aircraft but much cheaper. That is, showing what needed to be done in the research lab to make the environment suitable for doing…

Suzanne: Sufficiently relevant.

Chuck: Of sufficient fidelity.

Suzanne: OK.

Suzanne: Relevance and fidelity.

Chuck: Yes, relevance and fidelity, so that the test results would carry over. Now, if you can do the majority of testing of your software component in the laboratory as opposed to on the aircraft, well, testing on the aircraft costs lots of dollars, and testing in the lab costs some dollars, but not lots of dollars in a relative…

Suzanne: So, you have talked about aircraft. You have talked about medical devices. I know that there's also some activity in this area in the automotive industry. We've got one statistic that says that 85 percent of the functionality of cars today is in software. So, how are [assurance cases] used in some of those other commercial areas?

Chuck: Right. But you talked about U.S., though, and I think, to my knowledge, there is interest in the U.S. industry, but I don't know that it's actively being used yet. I don't have visibility into that.

Suzanne: You have to connect with Peter Feiler and the Society for Automotive Engineers and get on one of those committees so that we can get going on that one.

Chuck: As I say, there is interest. The question is: is it actually something that's impacting what you step on in your U.S.-made car. The answer is I don't know. If somebody listening to this wants to put me in touch with the right people, I'd be really pleased to find them.

Suzanne: I would be happier if the automotive industry was using assurance cases.

Chuck: Again, that's not to say that they're not doing good things.

Suzanne: The thing that I keep coming back to with this is it's just striking how much more accessible assurance cases make the knowledge about what's going on in your evidence gathering, and how you see that connecting to claims like safety, in particular. I just think it's marvelous.

Chuck: There's always pushback, right? *We've been doing it this way for years. We have the hazard analyses. We have the fault-tree analyses. Why do we need to do this? It's just going to add to our work.*

Our belief is, yes, they've been doing this for years. A lot of the materials they've gathered, these analyses, are, in fact, suitable evidence for an assurance case. The assurance cases—in the situation where they've been doing a good job in the first place—the assurance case is a way of organizing that evidence, so that somebody else can tell that they've been doing a good job in the first place.

Suzanne: I don't know if you've been working in this area yet, but I know our legacy systems, especially in the Department of Defense, systems are in play and operations for 40, 50 years, not even 10, 20 anymore. So, having something like this as a way of understanding, *When we change a system, what evidence has to change?* How does the argument change when you change that? That's got to be an amazing…

Chuck: Of course, the problem is you've got the legacy system, but you don't have the case. So going forward, this is important. In fact, one of the areas that we're beginning to work on, we're just beginning to work on, is the problem of incremental certification. How you might structure a system—both the architecture and the assurance case—so that if you have to make a change, perhaps add a feature, perhaps fix a problem, you don't have to go about recertifying the whole system.

Suzanne: That's very important in our Department of Defense systems or aircraft worthiness. I know that the aircraft worthiness testing that has to occur after you've made even minor changes sometimes can take as much as a year. That kind of time—if you can reduce that by helping people to understand, *No, you really don't have to touch this part of the testing again because of these arguments*—that is what saves the taxpayers money.

Chuck: This work is happening under our value-driven incremental development project. The whole overall goal there is to get things out to the field a lot faster: incremental development and get it out there in an assured way.

Suzanne: There are a lot of different ways and places that I think assurance cases can be used. I'm very happy to see that they're starting to make inroads into some of those places. I do invite

those of you in the audience that think that assurance cases may be an interesting thing for you to pursue to get in touch with Chuck and his team. I know you're looking for interesting pilots for doing different things with assurance cases.

Chuck: Right. I just want to add that the area we've been concentrating on for the last couple of years is confidence, how you achieve confidence in an assurance case. We didn't cover that, but it's a very interesting area that has brought together ideas from law and philosophy and artificial intelligence, and we have been looking at something we call eliminative argumentation, which is showing a lot of promise to make it easier to achieve confidence in assurance cases.

Eliminative argumentation makes use of eliminative induction, Baconian probabilities, and defeasible reasoning, all of this would be a great topic for a future podcast. It's also beginning to be discussed on the SEI Blog. We have several reports and white papers that you can find on the SEI website about this.

Suzanne: I definitely want you to come back and talk to us about that because I have no idea what eliminative reasoning is, and I think that sounds like a really interesting thing to talk about for a little while. That's wonderful that you are not only pursuing what we can do with assurance cases but how we make the assurance cases themselves improved and increasing the confidence in them.

I think that's going to go a long way towards making them accepted. It sounds like you've got a lot on your plate, and eliminative reasoning looks like your immediate future. So anything else that you want to tell us about that you're working on in this area before we close?

Chuck: Well, I think we've pretty much covered everything, at least at a superficial level. I would be happy to come back sometime. So thanks very much, Suzanne.

Suzanne: I look forward to you coming back and talking some more about this. Thank you so much for joining us today, Chuck. For more information on the research that Chuck is doing on assurance cases, please visit sei.cmu.edu/library. At the bottom of the page, under SEI links, search on Chuck Weinstock in the author index.

Chuck: Or Chuck.

Suzanne: Or Chuck Weinstock because, yeah, we have to be more formal on the papers. This podcast is able on the SEI website at sei.cmu.edu/podcasts and on Carnegie Mellon University's iTunes U site. As always, if you have any questions, please don't hesitate to mail, e-mail us at info@sei.cmu.edu. Thank you.