An Alternative to Risk Management for Information and Software Security Transcript

Part 1: Why Risk Management Is a Poor Foundation for Security

Julia Allen: Welcome to CERT's Podcast Series: Security for Business Leaders. The CERT program is part of the Software Engineering Institute, a federally funded research and development center at Carnegie Mellon University in Pittsburgh, Pennsylvania. You can find out more about us at cert.org.

Show notes for today's conversation are available at the podcast website.

My name is Julia Allen. I'm a senior researcher at CERT, working on security governance and software assurance. Today I'm pleased to welcome Brian Chess, Founder and Chief Scientist at Fortify Software. Today Brian and I will be discussing why the concept of risk is a poor foundation for tackling software security and assurance, a bit of a contrarian view, and an alternative approach for what has worked well that Brian has seen in other disciplines. So welcome Brian, glad to have you with us here today.

Brian Chess: Thank you Julia, glad to be here.

Julia Allen: So effective risk management — I mean we talk about it all the time — is viewed as kind of one of the cornerstones for making the business case for information and software security and for helping us all figure out where to spend money. So why do you think that's emerged as an attractive approach and is being used so broadly?

Brian Chess: Well I'll tell you, I think risk has a lot of appeal to it. So before we try and tear it down let me try and build it up a little bit. So first of all, I think that there's a magic to probability that is just really, really beautiful. The whole idea that we can tame something that is essentially untamable. I mean I don't know what the next coin flip is going to bring but if I look at a long enough sequence of coin flips I can actually have a lot of control over that kind of system. So I think that the whole idea that we can tame something that essentially we can't control is really pretty cool.

And we start applying a concept like that to security where of course things are naturally unpredictable because we can't control those who seek to do us harm. Then it's a really great idea that we actually might be able to have some measure of control over what's going to happen. And also it's really useful for describing security because in security we know there are no absolutes. We can never tell you something is absolutely a hundred percent secure and the notion of risk goes a long way towards capturing that.

Julia Allen: So when you have this space, as you said, in security where there's always more to do than we could possibly do, it's really hard to nail down in any kind of quantitative terms. Risk tends to be kind of a natural fit don't you think?

Brian Chess: It's not just a natural fit. It solves a really natural problem that security people fall into and that is trying to explain why they're relevant to the rest of the world. So talking about risk really allows security people to talk to business people and security people like that a lot. They think if they can map security problems into risk management problems then they can easily map that into money and money is the language of business. Talking about risk is good for security because it makes us think we're relevant.

Julia Allen: Hmm, that's an interesting point and I think a very valid one. So heading down this alternative view, when I last heard you present, you said that risk management is failing us and serves as a poor foundation. So I'd love for you to describe how you've come to that conclusion or what your observations are about that.

Brian Chess: Well, I'll tell you it really struck me this last June. I was sitting in an ISE banquet, Information Security Executive banquet, and listening to James Nelms, who, at least at the time, was the chief information security officer for the World Bank. And he was talking about how he did risk management. And he was really very, very quantitative about it, exactly the way that you might read about it in a textbook. He was talking about the assets he had to protect, he was talking about the threats against those assets, and he was talking about the vulnerabilities that he had to compensate for in his system. And he had slide after slide with a lot of decimal points talking about defense-in-depth and how many attacks were being blocked by the firewall, and how many more attacks were going to be blocked by the anti-spam system, and so on and so on. And it occurred to me that he actually had quite a few more decimal places of precision up on his slides than I thought he could possibly have. And his outcome was, "And I think this is why the World Bank is adequately secure."

And what occurred to me, sitting there last June, was that he had an incredible amount of sensitivity to the assumptions that he was making — that actually the business that he was trying to protect had assets that were fluctuating every day. And that the threats that he was facing fluctuated every day too, as people develop new ways to attack systems, as new types of vulnerabilities are discovered. And that keeping up with all of that was actually going to be quite a bit of work and probably the end result was that he wasn't quite as safe as he thought he was.

And so that got me thinking about really all the problems that you might encounter with risk. And I'll tell you for the mathematically minded, I think the biggest problem that people encounter is the notion of composition — that for the most part we aren't trying to address one risk. We're trying to address a lot of risks and then take countermeasures to prevent the loss from becoming real. And the problem is that a lot of those risks are interdependent.

So if you remember back to your college probability class, when you start having interdependent probabilities you can't sum across them. And that makes the math a lot more complicated. And the end result is that you end up having to make a whole lot of estimations and those estimations might well be wrong.

Julia Allen: So in the case where you were sitting in this presentation in June then would it be fair to say that the precision, or the attempt to measure in the way that it was being described, really falls down or runs counter to the nature of the threat, the nature of the risk, you just talked about interdependencies. In other words, you're really trying to nail down a space that by its very nature doesn't lend itself to that kind of measurement and examination. Is that what you're saying?

Brian Chess: Essentially yes. The math works very well when you can be very precise about things like the value of the assets you're trying to protect or the probability that an attack is going to succeed. But in practice you really can't get that precise. And I'll tell you, the presentation that I listened to was, I mean, there are a lot of people out there who'll tell you that risk management is the way to go. But there are fewer people who will actually start putting numbers up about "Here's what I think we're looking at when it comes to risk management numbers."

Julia Allen: Typically we see the old high, medium, low, or stop-light, or some relative ranking as opposed to a hard number.

Brian Chess: Exactly. Most people get away from decimal points and they start talking about things like high, medium and low. And then they create a whole other set of problems for themselves. Because does a high trump four lows? What about four lows plus a medium? What if I've got two highs, how much is that worth? And all of a sudden you can't do math on this stuff anymore and that means it's much, much harder to convert into money. And now one of our perceived strengths, which is the ability to convert security problems into business problems, all of a sudden gets much, much weaker.

Part 2: Learning from Other Disciplines: Standards, Compliance, and Process

Julia Allen: So let's kind of keep running with this a bit and see where you take us. If identifying and managing risks can't deliver the results we need or help inform our actions and decision-making, what approach do you recommend?

Brian Chess: Well, you know, I felt a little lost after I walked out of that presentation in June because all of a sudden I was kind of disillusioned with the whole notion of risk and I needed something to fill that void. And I'll tell you, in my work, I help people build systems. And I try and help them build systems by understanding what could go wrong with those systems. So if risk wasn't going to be the answer, I sure started thinking about, "Well, how do other people build systems and how do they account for this stuff that they can't know for sure?"

And the first thing I thought about was construction. I thought about people who build buildings or who build bridges and how do they prepare for events that they know are very improbable but will eventually occur? How do you account for the hundred year storm, for example? And so I started thinking about building codes and from building codes I jumped to thinking about fire codes. How does the fire marshal know whether or not your building is fire safe or not? And certainly the fire marshal

can't ever guarantee to you that your building isn't going to burn down and can't guarantee to you that no one will be hurt by fire. But what the fire marshal does is go around with a checklist of things that we have learned through a lot of painful and sad events lead to worse fire problems and makes sure that a building is as safe as it can be.

And that led me to begin thinking about standards and compliance as really being the correct replacement for this notion of risk – that we're never going to be able to know what our asset values are or what the threats we face are to the extent that we can really rely on risk. So instead we need to rely on standards and checking compliance to those standards.

Julia Allen: I think analogies are always extremely helpful and looking at bodies of knowledge that are much more established than what we find in the security field. But when you talk about construction and you talk about fire codes and other types of disciplines, those have physical attributes. And systems and software and security other than looking at the servers, there's really nothing that we can get our hands around. So how do you counter the argument that, "Well, that may be fine and well for a physically-based discipline where you're trying to lay a standard around a building code." But how do you address the argument that that really doesn't apply to something as intangible as software?

Brian Chess: Well I think you raise a really good point and I'm going to answer you with yet another analogy. And I'll tell you, I really like the building analogy because it's easy to think about because you can picture it in your mind's eye. But exactly as you say, a lot of the things that we do that are related to software security you can't hold in your mind's eye nearly as well. And in fact, we can't measure nearly as well. So with a building you might think "Well, I can measure how strong that building is." But how do you measure how strong a piece of software is? Despite decades of work from a whole lot of researchers, we're just really not very good at measuring software or measuring digital systems. So I kept looking for more analogies, analogies where the systems that we needed to measure for safety in the physical world we inhabit are really hard to measure.

And that took me to food safety actually. It took me to thinking about how do you know whether or not it's okay to eat in a particular restaurant. We've got a real measurement problem there because you can't measure all of the food that comes out of a kitchen, it would be completely impractical. And yet it's really, really important to know that food is being prepared in a safe way because if you don't prepare food safely then people can be badly hurt. And so I started thinking about what is the essence of our public health system when it comes to food safety in particular?

Restaurants actually have more similarities to the kinds of business that we encounter in the digital world than you might think. Restaurants are a tough business to be in. There's a lot of competitive pressure in them and there's a lot of opportunities to make mistakes. I'll tell you the thing that people don't leap to immediately when I give them this analogy is they don't think about the fact that restaurants are actually

constantly under attack. And they're under attack by a global threat. And those global threats are diseases and they're global because our food supply is global. And those diseases are adaptive. They might not adapt quite as quickly as new viruses pop up on the Internet but they do change. And so I started thinking about "Well how do restaurants make sure they're safe?" And they make sure they're safe by really checking the processes they use to prepare food. As opposed to trying to measure the end result, they measure the process by which they achieve that result. And I think that's the same thing we need to do for software. We need to be looking at the processes used and make sure that best practices are used in the preparation of that software.

Julia Allen: So this is great because you've given us a couple of ways to think about that. So you talk about standards and compliance as mapping into — or a useful analogy being the construction industry and other types of fire safety. And then you talk about food safety. And here we're focused on process and outcome. In other words, if I can measure all the different ways that products come to me, how I actually go about preparing and delivering and taking into account all the health code and safety standards there, that process, I take it then, is an additional element of where you think we can get some real traction?

Brian Chess: Exactly right.

Julia Allen: Well, it's interesting because as you and I were dialoging and preparing for our conversation, I was reminded that Donn Parker (who's a noted security expert) in his work, has strongly stated that intangible risk reduction, something that we can't get our hands on, is really a weak justification for security. And he also argues about that the real tenets of strong security is — are due diligence, compliance, and enabling the business. And so I wonder — is there some resonance between what you're observing and formulating here and what Donn Parker talks about?

Brian Chess: I think Donn Parker's writing in this area is absolutely tremendous. And I'll tell you, I think we came from different places. I came from thinking about constructing systems and I think he came from decades of experience watching risk management achieve somewhere between so-so and abysmal results. And we got to the same place, which is the answer has got to be standards and compliance. So I hope that Donn takes heart when he sees other people reaching the same conclusion that he's been preaching for quite awhile now.

But let me tell you, it's not just me and Donn out there. If I might recommend a book, the name of the book is *The Black Swan*. It's by [Nassim Nicholas] Taleb and Taleb comes from a very different background. Taleb comes from – well he's an economist by training and he's a hedge fund manager for the last ten or fifteen years or so. And in his book *The Black Swan*, he talks about how people didn't know that black swans existed. Europeans didn't know black swans existed anyway until they got to Australia. And all of a sudden they found out that the world is quite a bit different than they thought it was. In other words, this assumption they'd been

making for, well, hundreds and hundreds of years turned out to not hold. And it's things like that that can really turn our world upside down.

For instance, if a year ago you were operating under the assumption that housing prices could never decline on a national or international scale then you might all of a sudden find that your credit default swaps weren't quite worth what you thought they were going to be worth — that this whole notion of risk management that all the quants on Wall Street were basing their business upon wasn't quite as firm a foundation as we thought it was.

Julia Allen: So when you talk about the black swan case and what's happened with housing prices, what you're really talking about from a risk management point of view, is kind of the low probability catastrophic events — and how we prepare for those or the things that are big disruptors that risk management doesn't really help us there. Is that what you're saying?

Brian Chess: If you're doing your risk management by watching the attacks that arrive on your firewall yesterday then your risk management will be centered on what you saw yesterday. And that one big catastrophic event is likely to be the thing that you overlook. What you've got to do if you want to prepare for that catastrophic event is you've got to start pooling all of the knowledge we've got about all of the attacks that have gone on in all sorts of relevant systems. And hopefully that knowledge is coming together in terms of a standard that you can actually measure your compliance with.

So I think that Taleb does a great job of arguing the case against risk, or against the misuse of risk perhaps I should say, but standards and compliance I think that's something that's for us in the security domain.

Part 3: A Case in Point; Examine Your Environment

Julia Allen: Well, it's been a big frustration of mine. I think in some respects we as security practitioners have really done our community a disservice with the proliferation of standards, compliance guidelines, process definitions. I mean you could make a long list of all the things that we all recommend that we should be doing in our organizations and of course there's a great deal of commonality amongst those. But when you talk about standards and compliance and process for information and software security, where do you point? Or what sources do you find to be most useful and most reputable?

Brian Chess: Well, I'll tell you, I think there's one really big success story over the last few years in this area and that's the PCI Data Security Standard. And it's an interesting success story because the Data Security Standard hasn't been successful because it's the be all, end all standard. In fact, if you talk to the people who've created the Data Security Standard, they know that they've actually got a long way to go in terms of creating the best possible standard. But what they did is they carved out just a small portion of the universe of software. They said "We want to talk about software that handles credit card numbers." And that's enough for them to actually

get quite a bit of traction in terms of describing the properties they want around that system. And I've seen, in terms of talking to my customers at Fortify, a tremendous amount of rallying and actually changing the behavior of their organizations around the PCI Data Security Standard, so I think that's a tremendous success.

Julia Allen: Well, and obviously you've got the credit card providers pushing their merchants in the direction to comply. So you need to have a well-defined standard as well as a catalyzing force to make people stand up and take notice, right?

Brian Chess: That's a really good point. In the case of that standard, and one of the reasons why it's been so successful, is because security people haven't had to sell its business relevance. Its business relevance has been pretty obvious. And that is the people who handle the credit card payments say we have to do this. And so that's been very powerful, and I think it's the kind of thing that we're going to need to see more of if we want to see standards really take off.

Julia Allen: So even though we've kind of built it up and tore it down, is there value in your experience and in the customers and clients that you've worked with, in continuing to use risk management perhaps to complement some of the standards and compliance and process and outcome work? Do you see those being kind of complementary in any respect?

Brian Chess: Well, I'll tell you, I think that there are scenarios where you can talk about risk management, and should talk about risk management legitimately and with authority. Those places are largely places where the math works, where you actually can get around to talking about math. Other than that you're going to be stuck talking about sort of risk as a motherhood and apple pie subject.

But those particular areas are areas where you have well-defined and well-controlled scenarios. So if you actually know all of the possible outcomes and you can measure them then that's some place where you certainly can talk about risk management. For example, if you're running a casino, you won't be running a casino very long if you're not thinking about the risk set up by the rules you constrain your games by.

I would say there's another area too and that is if you actually have a very large number of honestly independent experiments, then you can talk about risk. So for instance, when Dan Geer talks about risk in the context of all of the hosts out there on the Internet, then I think he can reach some valid conclusions that way. But he is not suggesting risk management decisions there in terms of a single organization or a single entity.

Julia Allen: Well Brian this has been terrific, kind of a mind opener and an alternative point of view to some of the prevailing wisdom. Do you have, in addition to the sources that you've already called out, do you have some places where our listeners can learn more?

Brian Chess: Well, I guess I'd like to suggest two courses of investigation to people who are interested in this topic. And I'll say both of them are indirect and they're

indirect because as we've already said, this is really the contrarian view in terms of software security these days but let me give them to you anyway. The first is take them out of your environment. I'm sure you've been in an elevator recently where the elevator said "Inspection report available upon request?" Well, go request that elevator inspection report. It's actually pretty darn interesting to look at what elevator inspectors do, elevators turn out to be a really safe mode of transport but it's because they've got really good guidelines around how do you know if that elevator is safe or not. Similarly, I'll tell you the last time the fire marshal came and looked at the office at Fortify, I trailed him around. And I pestered him, and I said "Hey, could I have a copy of that checklist you're using? And do you have friends who use different checklists? How does this all work?" And I learned a lot from talking to my fire marshal. Similarly, you might end up getting a restaurant's health inspection report and there's a lot to be learned from that too. So that's my first suggestion is go and learn from your environment about how people are managing risk in the world around you all the time.

My second recommendation here is to support the idea that people are just not that good at the gut check form of risk management — is a book titled *Stumbling On Happiness* by Daniel Gilbert. He's a psychology professor at Harvard. And he makes the case that people aren't good at making themselves happy because they're not good at predicting their own futures. Human beings just aren't very good at predicting improbable events. And so I think that's a good, again, rather indirect way to look at the subject of risk.

Julia Allen: Well, very interesting, kind of outside the box thinking. And I think as practitioners we owe it to our community to try and bring the best ideas from other disciplines and where things have kind of been more examined over a longer period of time, more tried and true, and bring them to our discipline as well. So I'm really extremely appreciative of your time and your perspective and your experience today.

Brian Chess: Well thanks for having me Julia.