Rethinking Risk Management
Transcript

Part 1: Why Traditional Approaches Fall Short

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 back Chris Alberts, a senior researcher in risk management, also at Carnegie Mellon Software Engineering Institute. I think you'll find it pretty interesting because today Chris and I will be talking about effective ways to assure success and manage risk in complex, distributed, system-of-systems programs, which many of us are facing today. We'll also be discussing some of the implications of risk management in those types of environments for security and software assurance. So welcome back Chris. I'm really glad to have you with us today.

Chris Alberts: Thanks Julia, great to be here.

Julia Allen: So in your most recent work, which we'll be talking about, you say that we need to rethink risk management. So what do you mean by that and why do you think it's necessary, particularly now?

Chris Alberts: Well there are three characteristics that we see when we look at most approaches – the traditional paradigm and traditional approaches for managing risk. The three characteristics are: (1) they rely solely on historical data; (2) they employ a tactical analysis of risk; and (3) they're point solutions. And I'll address each of these specifically beginning with the reliance on historical data.

I'm focusing my answers in the area of developing and operating software-intensive systems and systems of systems.

So these traditional approaches typically are based off of what we call "taxonomies of risk sources." And these just basically list the things that have happened in the past and things that we might want to look at and determine if we're concerned about them as we move into the future.

In addition, they look at probabilities based on statistical data or occurrences in the past as well. They'll very often look at frequency of occurrence of these different events. So basically what they're doing is, based on what's happened in the past, they're trying to look at and predict what might happen in the future.

But as we see in today's environments, especially in the business and the technological environments today, we see things are rapidly changing. And because of that the past is not necessarily always a good predictor of the future. And security is a really good example of this. If you think about the number of vulnerabilities that are identified over the course of a year and the new exploits that come with them, well these lead to new risks. And these new risks won't necessarily be in your taxonomies and they also won't necessarily – you won't necessarily have any probability data on them because they're new and novel types of risks.

Julia Allen: Right. I mean it seems like in the security arena we no sooner get a class of vulnerability or a class of exploit handled before – the attacker community is very innovative, very creative in coming up with the next big thing, right?

Chris Alberts: Exactly. And so what you can't necessarily do is rest on what's happened in the past will equal what's going to happen in the future. And so what we're looking at is moving from sole reliance on historical data. You don't want to forget about the historical data. But you want to augment what you're doing with more structured analysis based on models of system characteristics so that you can start looking at ways of identifying when the system or – system, I'm using it in the broadest context whether it be a process or a program or an actual IT type of system – when you're starting to see unusual characteristics from a performance point of view. And then start using that to make predictions about what might be happening.

Julia Allen: Okay. So what about your second kind of traditional approach? You talked about employing tactical analysis?

Chris Alberts: Right. And in tactical analysis what we tend to do there is we define risk as a simple cause and effect pair where the cause is an event, a potential event that might or might not occur (what I was talking about earlier in terms of the exploits and the). And you're looking and concerned, looking at analyzing the effect or consequence of that particular event.

So for example, let's say that I'm running an incident management help desk where I'm – people call in security events and incidents and my team is the first line of defense against – in mounting a response to those events and incidents. I'll be concerned about making sure that I have it staffed 24/7, and in terms of staffing I might be concerned about – on certain shifts where I don't have a lot of backup. So I might be concerned about losing someone, a key person on a permanent basis or it could be a temporary basis.

So from a tactical point of view, the cause of the risk that I'm worried about, or the event, is the possibility of losing a key team member. And the consequence that I'm worried about is how this might affect the quality of our response as we try to manage these events and incidents. So this is again where it looks at a cause and effect, a pair from a tactical perspective.

And the thing – when I do a risk assessment for a customer – and in the past I've done a lot of these from the tactical point of view – we can actually literally identify hundreds of risks for a single program or a single process or system. And when you think about how we now interconnect programs, processes and systems, the number grows accordingly. And so what happens is we get a huge number of these risk statements. And what we normally do is we evaluate them for probability of occurrence, the impact on the organization or the program or whatever if they occur. And then we derive what we call risk exposure from these two values. Risk exposure is essentially the product of impact and probability and gives us a measure of what the risk to us is in this given situation.

So the first thing we do is we put them in a list and then we focus our management attention on the top 10 to 20% on the list. And so – but what we find in practice is very often managers get caught off guard by something in the other 80%. Over time something may become more important. Or what we're seeing a lot of is a lot of small, relatively benign risks in and of themselves combine to cause a greater failure. And so if you're not looking at that last 80%, you're not necessarily doing a very broad-based view of managing your risk.

Julia Allen: So in that kind of an approach, you collect these hundreds of risks. You have ways of grouping, prioritizing, aggregating, assigning various values to them – as you said, focusing on the top 10 to 20%. But it seems to me then that you can really – it's almost like what we talk about a false sense of security. You can have a false sense that you've identified the key risks and go merrily along your way, as you said, and be caught off guard. Right?

Chris Alberts: Right. And to counter that what we're suggesting is moving to more of a systemic solution. What we essentially do is aggregate these individual risks into groups and then focus management action at the group level. So in some sense we try to get about 10 to 20 of these aggregate areas – we call them drivers – and we focus our attention on them and then periodically reassess them over time. So hopefully when you do that, when some of the other 80% start having more of an impact, you start seeing that because you're paying attention to them on a continual basis.

Julia Allen: Okay. And then you also had a third characteristic of how we've traditionally managed risk, where you talk about point solutions, single events, single domain type of focus. Can you say a little bit more about that?

Chris Alberts: Sure. And that is the third characteristic. And really what we do is when we manage risk, we very often focus on what I'll call silos of risk, based on a life cycle phase and on type life cycle phase and entity. So when you think about it, we manage operational risk, or we'll manage security risk, or we'll manage architecture risk or program risk; and they're all done separately. But and so that's what I mean by type of risk.

Usually the type is based on the cause of the risk. And so what we find is that – and today we're in an interconnected world, and so what happens is that actually these risks affect each other. So, for example, security risks will affect business processes and programs.

So think about an information-related risk. If a competitor gains unauthorized access to proprietary information about a new product that you're developing, for example, that security breach will affect that development program and the organization as a whole. There is an example where a security risk affects a business program and potentially the financial prospects for the organization.

And so what we see is that security risks in this case do affect programs and a lot of these different types of risks are interconnected with each other. But program and business managers are not really always involved in managing these risks. That's really – security risks are in the domain of the IT or the security departments.

So the questions are, do the people in those departments, really are they implementing the appropriate controls based on how that information can impact the process? Do they know how sensitive the information is? And this is an example where a security risk can affect an entire program, and so what we need to start doing is breaking down these silos and really start assuming a more holistic, integrated view of risk.

Part 2: Managing Risks across the Life Cycle; Using the Mosaic Toolkit

Julia Allen: So let's build on that a little bit because I'd like to run with that thread, which is managing – how important it is to manage risk across the system life cycle. And also I would assume when you talk about business process, you've got partners. So you need to talk about managing risk across the supply chain. So how do you get to this more holistic integrated view and move towards it from this traditional picture that you've painted?

Chris Alberts: If you look at the traditional focus, again it focuses on what I'll call a single entity. It could be a process, it could be a system, it could be a program. But today things are interconnected. So what you're really doing is looking at collections of these things; in fact, it can be across types. A process is supported by systems, IT systems, as are programs and they all are related in some way, shape, or form.

And so the life cycle view of risk we're talking about is that you're developing and deploying systems, and then maintaining them and operating them in the operations space. Decisions that are made early in the life cycle can have a profound impact on operations. And this is an example of what we say – that decisions made early in the life cycle can impose risk on later activities in the life cycle and ultimately on operations. And those people later in the life cycle inherit risk based on decisions that were made earlier in the life cycle.

Julia Allen: Okay, that makes sense. And typically you find there's handoffs that occur at life cycle phases where you don't necessarily have that view of inherited risk and something that you need to consider downstream by something that was introduced upstream, right?

Chris Alberts: Right, exactly. And so you're not really tracking what the residual risk is as you hand it off down the, from one life cycle activity to the other. And usually you really don't get a good feel for that until you actually deploy the system and then you see all the problems that are occurring.

Julia Allen: And then how does that show up in supply chains?

Chris Alberts: Well in supply chains it's a similar principle. If you think about it, if you're dependent on products and services provided by someone else in a supply chain, their decisions are going to affect the quality of the products and services and the timeliness of the products and services that are supplied to you. So in that sense they're imposing risk on you. And likewise, your decisions can impose risk on other people in the supply chain.

So what we see is this idea of inherited and imposed risk really calls for more holistic solutions that really link to mission and objectives, rather than just point solutions where they're based on specific entities like programs, processes, and systems.

You can think of it as a chain, when you look at a collection of systems or processes. And if you focus your activities on just one link in the chain, that's what we call local optimization. And which you can then – when you locally optimize based on that, your weakest link defines your overall risk based on mission. And so what we're looking to do is try to take a look at, based at what you're trying to achieve as a group, how can we manage risk best throughout the chain and not just focus on specific links?

Julia Allen: Okay. Well then let's start to turn our attention to how we get to kind of some tangible methods and approaches for a more integrated, holistic view. So your new risk management methodology is called Mosaic. And why don't you tell us a little bit about the focus of Mosaic. And you've certainly introduced some of the motivators, but a little bit about its development history.

Chris Alberts: Sure. What we noticed when we were – and this goes back about four or five years – we started noticing that people were beginning to really struggle when they were trying to apply some of the more traditional approaches for managing risk, especially in these multi-enterprise, multi-system management environments. So our goal was to look for better ways to manage risk in these types of environments. And in the process what we developed is SEI Mosaic, which is basically a suite of methods that can be applied across the life cycle and supply chain.

And so what we've done is we've developed a variety of methods from those that are very basic, some of which can be self-applied, some very highly advanced methods that really require considerable expertise to apply. And what we found is that different methods are needed in different situations. And the best analogy that I can draw is the medical paradigm.

And when you think about it – at the very front lines when you're looking for treatment for a given set of symptoms that you might have from a health perspective, there are things that you can do over the counter. You can do some selfdiagnosis and use some over-the-counter cold medications and things of that sort for instance. So that's the first line of defense.

And if that doesn't work, you might go and see a general practitioner and get just a general health check where they run some simple tests like blood pressure and look at your heartbeat and things of that nature, and they can diagnose there. Or they might call for more specialized tests, let's say like an MRI or something of that nature where they get more data to do a more sophisticated analysis of what might be going wrong. And in some cases they may send you to a specialist for some very specialized care in certain Areas. It could be a cancer specialist, like an oncologist, or a heart specialist and so forth.

So when we designed Mosaic, we had this paradigm in mind. And so much like – we have some self-applied assessments that are just like looking at very basic health checks that people can do for themselves and then they can make corrective actions as they see fit. We also have some diagnostics that serve as health checks that can be applied by risk management experts. They're a little bit more sophisticated but still are fairly basic in nature.

Then we have some more advanced analyses that provide a more in-depth view of a process or a program or a system and this is analogous to the MRIs. So you're getting some much more depth of information collection and you're really taking a much more analytical look at what really the root causes are of the symptoms that you're seeing.

And then finally we, in some cases, we might recommend a full-on, specialized risk assessment such as a security assessment. And fact those very often are likely then outside of the Mosaic suite of methods but you would go and look at specific security issues using various security assessments. And so that would be equivalent to the specialist in the medical paradigm.

So what we've developed with Mosaic is what we call a risk toolkit that includes various types of methods of varying degrees of depth.

Julia Allen: You peaked a thought here, which is when you talked about a security assessment as being outside, potentially outside, the scope of Mosaic, but being something that Mosaic would recommend. So would it be fair to extend that idea and say if an organization already has – you talked about these kind of stove-piped or domain-specific risk assessments. Might there be potential to use something like Mosaic to give the overarching, holistic integrated view and then maybe point to already an existing risk assessment method to drill down, as you say, like the medical specialist? In other words, could they take advantage of the risk assessment methods that they already have?

Chris Alberts: Mosaic is – we've done, actually applied it in two different ways. The first is the most general application, and I think it's worth – because we talk about this holistic view. When we perform any of our assessments, we include information and technology, since they're critical to any program or process these days, as areas to look at. So based on the findings, we might find some issues related to information from a security perspective – confidentiality, integrity and availability, things of that nature – or security issues related to technologies that we can actually diagnose and correct based on that basic assessment.

We try to incorporate some basic security issues into all of our assessments when we apply them to programs or processes. Now specifically we've also applied Mosaic in a specific security context which is cyber security
incident management.

So in this particular case what we did is we wanted to look at the effectiveness of the incident management process. So basically from the time that you see an indication of an event or an incident until it's actually responded to, there's a sequence of steps that people follow in doing that and different groups are involved. So again it gets to that multi-enterprise, multi-system kind of an environment. And so to kind of answer that question I'll start actually by giving a little bit of background on how we do a structured analysis using Mosaic.

So we start at the top with the identification of the key objectives for the program or system or process that we're looking at. So in the case of incident management it was the incident management capability that we were looking at.

And there were three basic types of objectives that we identified. One was the quality of the response to an event or an incident; the second was the timeliness of the response; and the third was customer satisfaction. So you're actually responding to an event based on, in many cases, a call to a help desk. And so are the people who are receiving these services that you're providing, are they satisfied with those services?

So these form the core objectives for incident management. From these objectives then what we do is identify a small set of what we call drivers. And a driver, as I use it here, is a factor that has a strong influence on whether or not the objectives will be achieved.

So for incident management we identified about ten of these drivers. An example of one was focused on the ability to execute tasks and activities by the incident management team. So we created a question around this that was phrased,
"Is task execution effective and efficient?" It's a yes/no question but we allow five possible responses: yes; likely yes; equally likely yes and no; likely no; and no. And the idea is that you're looking at yes and no, and the shades of probability in between.

So as we look to answer each of these ten or so questions, what we're doing is collecting information. In this case for task execution we're looking at things like the experience and the expertise of the people performing the task execution; the experience of management and the actions of management; the staffing levels and resources. Do they have the tools that they need to do their jobs? Did they have training and how effective is that training?

And then we also look at various types of events such as what is the possibility of losing key staff. If you lose a key, a tool or system that you rely on as part of this incident, do you have backups? Are you able to keep the process going? So we look at a lot of different aspects under each of these drivers.

Julia Allen: So you basically set an objective. You have a scope or a domain, in this case incident management. You set, you determine what the key objectives are for that process or function. And then you factor them into a series of drivers that you can get more tangible input on, right?

Chris Alberts: Right. And so it's a top-down analysis where we start with objectives, the drivers, and then we look at details underlying each of those drivers. And then you can keep going and drilling down to any depth of information that you need to really answer and evaluate each of those drivers effectively.

Part 3: Dealing with Preventable Failures

Julia Allen: So Chris this is a great example of how we might take something like Mosaic, apply it to a security function like incident management. But let me bump back up a little bit. And I think we've all observed in working with our security community, our customers, and our clients, that almost all organizations have some type – I mean it might be ad hoc and more implicit than explicit – but some type of risk management program or process. Yet preventable failures – even risks that have been identified – preventable failures continue to occur. And obviously this is also true for security.

So from your experience, from your observation, what can business leaders do to improve their risk management practices along some of the lines you've described?

Chris Alberts: Well several issues lead to preventable failures. And the first is an uneven or inconsistent application of current practice. In some cases what we've seen is some organizations have great risk management plans on paper at least, but they're really not executing them well. And so basically you need to be able to follow through with what your plans are, and gaps in execution are one aspect.

Second, they might not be implementing methods that are well suited to their environment. So this is the point that I've touched on throughout our conversation is that if you're looking at very interconnected environments, you need methods that are kind of designed for the kind of environment that you're in. And maybe the traditional paradigm is not best suited to what your risk management needs are.

And third, risk management tends often not to be well integrated with other management practices. So one of the things that we find – we see this quite often – people will do a large assessment, come up with findings and then stick them on the shelf and not really actively mitigate them. So knowing about the risk doesn't mean you're addressing the risk. So there's that.

But many organizations, people – at least some of the feedback that I've received – is that people talk about how time consuming and bureaucratic their risk management processes are. So it's really not surprising that people aren't fully embracing them because it becomes another activity that they have to do rather than something that's just a part of their normal day-to-day activities.

Julia Allen: Right, and something that actually provides them good and useful information for where to focus their attention, right, on a day-to-day basis.

Chris Alberts: Exactly. Exactly. And so what business leaders need to do is they need to really take a look at their current risk management practices and take appropriate steps to improve them. Are their practices effective? Are they getting the information they need when they need it? Can they improve their current methods? Do they need better methods? And also looking at what the gaps in their risk management practices are.

The first step for these organizations is to answer these types of questions and then really start charting a course for improvement.

And one other point that I want to make, that I think is really important in terms of risk management, is the difference between effectiveness and following a process. Because we've come up recently with a risk management framework. And we do actually two types of evaluations against the framework. One – and the framework basically defines what are the best risk management practices. What are the expected activities and outputs from when you manage risk?

And so we look at it from a process point of view, adherence to the process. So are you doing the things? Are you coming up with a risk management plan? Are you identifying risks? Do you have a way of expressing a risk, like through a risk statement or some other means?

So that tells you kind of are all the pieces in place but it doesn't really tell you whether they're working well together. And so you need to take a second look at effectiveness. Are you actually – if you have a risk management plan, is it a good plan? If you're looking at risks, identifying risks, are they well constructed risk statements, and in the end are you keeping your risk within tolerance? So these are the things that, from an effectiveness point of view, complements the process point of view. And so we look at it in both ways.

And so business leaders need to keep this in mind. It's not just about following a process. It's about actually making better decisions based on the risks that are confronting you.

Julia Allen: Yes, it's almost like they can do a sanity check and say, "Is the risk management process that I have in place informing my business decisions day to day? Is it helping me at critical juncture points; like if I'm looking at a merger and acquisition or if I'm looking at some major new system, product and service development? Is my risk management process helping inform those decisions?" Right?

Chris Alberts: Exactly. That's the bottom line. We see, and some organizations are very concerned about whether they're following some standard or some guideline. And they lose sight of why they're doing it in the first place, which is to make better decisions.

Julia Allen: Well this is great Chris. So as we come to our close, just a few final questions. So what's next for Mosaic? What are you looking at down the road and what are some of your development plans going forward?

Chris Alberts: Well we've just released a new set of courses and evaluation services based on the Mosaic approach. And so those are available to the community at large. I did allude to, in the last question, the risk management framework. In a few months, fall of 2009, we plan to be releasing that. And then we also have evaluations to support that framework as well.

We're also – we've come to a point where we have enough material together to begin thinking about our next book. And so we're hoping to publish that, either in late 2010 or early 2011, in this whole area of systemic risk management.

And then finally we're looking at refining, continuing to pilot, and then to codify some of our more advanced methods. And, in fact, we're starting to look at the possibility of and to applying – developing risk simulation models. So some very sophisticated simulation techniques based on system dynamics that are focused on evaluating risk in mission critical systems. And so we're real excited about the prospects for this work as well.

Julia Allen: Well that sounds like plenty to keep you busy. Right?

Chris Alberts: It certainly is keeping us busy.

Julia Allen: So where can our listeners learn more about this work including how to get in touch with you?

Chris Alberts: Well they can visit our webpages at www.sei.cmu.edu/risk. And we have a lot of information available there: various technical publications, overviews of the contrast between our systemic approaches and some of the more traditional approaches, and things of that nature.

Julia Allen: Well Chris this has been a great introduction and summary of some very exciting work in progress, and I think very relevant for the large complex system of systems that we're all having to contend with these days. So I thank you very much for your time.

Chris Alberts: Thanks Julia.